
Ansible Tower User Guide

Release Ansible Tower 3.0.2

Red Hat, Inc.

Jun 06, 2017

CONTENTS

1	Overview	2
1.1	Real-time Playbook Output and Exploration	2
1.2	“Push Button” Automation	2
1.3	Enhanced and Simplified Role-Based Access Control and Auditing	2
1.4	Cloud & Autoscaling Flexibility	3
1.5	The Ideal RESTful API	3
1.6	Backup and Restore	3
1.7	Ansible Galaxy Integration	3
1.8	Inventory Support for OpenStack	3
1.9	Remote Command Execution	3
1.10	System Tracking	4
1.11	Integrated Notifications	4
1.12	Satellite and CloudForms Integration	4
1.13	Run-time Job Customization	4
2	Tower Licensing, Updates, and Support	5
2.1	Support	5
2.2	Trial Licenses	5
2.3	License Types	5
2.4	Node Counting in Licenses	6
2.5	License Features	6
2.6	Tower Component Licenses	7
3	Logging In	8
4	Import a License	9
4.1	Adding a Tower License Manually	10
5	The Tower Dashboard and Interface	11
5.1	Tower Admin Menu	12
5.2	Settings Menu	13
5.3	My View – User Menu	14
5.4	Dashboard Views	16
5.5	Activity Streams	17
6	Organizations	19
6.1	Organizations - Permissions	22
6.2	Organizations - Notifications	24
6.3	Organizations - Users	25
6.4	Organization - Administrators	26

7	Users	29
7.1	User Types - Quick View	31
7.2	Users - Organizations	32
7.3	Users - Teams	32
7.4	Users - Granted Permissions	33
8	Teams	34
8.1	Teams - Users	35
8.2	Teams - Granted Permissions	37
9	Credentials	39
9.1	Understanding How Credentials Work	39
9.2	Getting Started with Credentials	39
9.3	Add a New Credential	42
9.4	Credential Types	43
10	Projects	56
10.1	Add a new project	57
10.2	Updating projects from source control	59
10.3	Add a new schedule	61
10.4	Ansible Galaxy Support	62
11	Inventories	63
11.1	Add a new inventory	63
11.2	Groups and Hosts	65
11.3	Running Ad Hoc Commands	84
11.4	System Tracking	86
12	Job Templates	93
12.1	Utilitizing Cloud Credentials	100
12.2	Surveys	102
12.3	Scan Job Templates	106
12.4	Provisioning Callbacks	115
12.5	Launching Jobs	116
12.6	Scheduling	118
13	Jobs	121
13.1	Job Results - Inventory Sync	121
13.2	Job Results - SCM	123
13.3	Job Results - Playbook Run	124
13.4	Job Concurrency	133
14	Notifications	135
14.1	Notifier Hierarchy	135
14.2	Workflow	135
14.3	Create a Notification Template	136
14.4	Notification Types	136
14.5	Configuring the <code>towerhost</code> hostname	141
15	Best Practices	143
15.1	Use Source Control	143
15.2	Ansible file and directory structure	143
15.3	Use Dynamic Inventory Sources	143
15.4	Variable Management for Inventory	144
15.5	Autoscaling	144

15.6	Larger Host Counts	144
15.7	Continuous integration / Continuous Deployment	144
16	Security	145
16.1	Playbook Access and Information Sharing	145
16.2	PRoot functionality and variables	146
16.3	Role-Based Access Controls	146
17	Index	152
18	Copyright © 2016 Red Hat, Inc.	153
	Index	154

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

The *Ansible Tower User Guide* discusses all of the functionality available in Ansible Tower and assumes moderate familiarity with Ansible, including concepts such as **Playbooks**, **Variables**, and **Tags**. For more information on these and other Ansible concepts, please see the Ansible documentation at <http://docs.ansible.com/>. This document has been updated to include information for the latest release of Ansible Tower 3.0.2.

Ansible Tower Version 3.0.2; Aug 31, 2016; <https://access.redhat.com/>

OVERVIEW

Thank you for your interest in Ansible Tower. Tower is a graphically-enabled framework accessible via a web interface and a REST API endpoint for Ansible, the open source IT orchestration engine. Whether sharing operations tasks with your team or integrating with Ansible through the Tower REST API, Tower provides many powerful tools to make your automation life easier.

1.1 Real-time Playbook Output and Exploration

Watch playbooks run in real time, seeing each host as they check in. Easily go back and explore the results for specific tasks and hosts in great detail. Search for specific plays or hosts and see just those results, or quickly zero in on errors that need to be corrected.

1.2 “Push Button” Automation

Access your favorite projects and re-trigger execution from the web interface with a minimum of clicking. Tower will ask for input variables, prompt for your credentials, kick off and monitor the job, and display results and host history over time.

1.3 Enhanced and Simplified Role-Based Access Control and Auditing

Ansible Tower allows for the granting of permissions to perform a specific task (such as to view, create, or modify a file) to different teams or explicit users through role-based access control (RBAC).

Keep some projects private, while allowing some users to edit inventory and others to run playbooks against only certain systems—either in check (dry run) or live mode. You can also allow certain users to use credentials without exposing the credentials to them. Regardless of what you do, Tower records the history of operations and who made them—including objects edited and jobs launched.

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.) Refer to *Role-Based Access Controls* for more information.

1.4 Cloud & Autoscaling Flexibility

Tower features a powerful provisioning callback feature that allows nodes to request configuration on demand. While optional, this is an ideal solution for a cloud auto-scaling scenario, integrating with provisioning servers like Cobbler, or when dealing with managed systems with unpredictable uptimes. Requiring no management software to be installed on remote nodes, the callback solution can be triggered via a simple call to ‘curl’ or ‘wget’, and is easily embeddable in init scripts, kickstarts, or preseeds. Access is controlled such that only machines in inventory can request configuration.

1.5 The Ideal RESTful API

The Tower REST API is the ideal RESTful API for a systems management application, with all resources fully discoverable, paginated, searchable, and well modeled. A styled API browser allows API exploration from the API root at `http://<Tower server name>/api/`, showing off every resource and relation. Everything that can be done in the user interface can be done in the API - and more.

1.6 Backup and Restore

The ability to backup and restore your system(s) has been integrated into the Tower setup playbook, making it easy for you to backup and replicate your Tower instance as needed.

1.7 Ansible Galaxy Integration

When it comes to describing your automation, everyone repeats the DRY mantra—“Don’t Repeat Yourself.” Using centralized copies of Ansible roles, such as in Ansible Galaxy, allows you to bring that philosophy to your playbooks. By including an Ansible Galaxy requirements.yml file in your project directory, Tower automatically fetches the roles your playbook needs from Galaxy, GitHub, or your local source control. Refer to [Ansible Galaxy Support](#) for more information.

1.8 Inventory Support for OpenStack

Ansible is committed to making OpenStack simple for everyone to use. As part of that, dynamic inventory support has been added for OpenStack. This allows you to easily target any of the virtual machines or images that you’re running in your OpenStack cloud.

1.9 Remote Command Execution

Often times, you just need to do a simple task on a few hosts, whether it’s add a single user, update a single security vulnerability, or restart a misbehaving service. Beginning with version 2.2.0, Tower includes remote command execution—any task that you can describe as a single Ansible play can be run on a host or group of hosts in your inventory, allowing you to get managing your systems quickly and easily. Plus, it is all backed by Tower’s RBAC engine and detailed audit logging, removing any questions regarding who has done what to what machines.

1.10 System Tracking

Introduced in version 2.2.0, Tower's System Tracking brings a new level of visibility to your infrastructure—you can see exactly what is happening on your systems, comparing it to both the prior state of the system and to other systems in your cluster, which helps you to ensure compliance. The rich and extensible store of data available in System Tracking is accessible via Tower's REST API, enabling you to feed it into other tools and systems.

1.11 Integrated Notifications

Starting with version 3.0, Ansible Tower allows you to easily keep track of the status of your automation. You can configure stackable notifications for job templates, projects, or entire organizations, and configure different notifications for job success and job failure. The following notification sources are supported: - Slack - E-mail - SMS (via Twilio) - HipChat - Pagerduty - IRC - Webhooks (post to an arbitrary webhook, for integration into other tools)

1.12 Satellite and CloudForms Integration

Ansible Tower 3.0 also adds dynamic inventory sources for Red Hat Satellite 6 and Red Hat CloudForms.

1.13 Run-time Job Customization

Bringing the flexibility of the command line to Tower, you can now prompt for any of the following:

- inventory
- credential
- job tags
- limits

TOWER LICENSING, UPDATES, AND SUPPORT

Tower is a proprietary software product and is licensed on an annual subscription basis.

Ansible is an open source software project and is licensed under the GNU General Public License version 3, as detailed in the Ansible source code: <https://github.com/ansible/ansible/blob/devel/COPYING>

2.1 Support

Ansible offers support for paid Enterprise customers seeking help with the Tower product. If you or your company has paid for a license of Ansible Tower, you can contact Ansible via the Red Hat Customer portal at <https://access.redhat.com/>. To better understand the levels of support which match your Tower license, refer to *License Types*.

If you are using Ansible core and are having issues, you should reach out to the “ansible-devel” mailing list or file an issue on the Github project page at <https://github.com/ansible/ansible/issues/>.

All of Ansible’s community and OSS info can be found here: <https://docs.ansible.com/ansible/community.html>

2.2 Trial Licenses

While a license is required for Tower to run, there is no fee for managing up to 10 hosts. Additionally, trial licenses are available for exploring Tower with a larger number of hosts.

Trial licenses for Tower are available at: <http://ansible.com/license>

To acquire a license for additional servers, visit: <http://www.ansible.com/pricing/>

2.3 License Types

Tower is licensed at various levels as an annual subscription. Whether you have a small business or a mission-critical environment, Ansible is ready to simplify your IT work-flow.

- **Self-Support**
 - Manage smaller environments (up to 250 nodes)
 - Maintenance and upgrades included
- **Enterprise: Standard**
 - Manage any size environment
 - Enterprise 8x5 support and SLA

- Maintenance and upgrades included
- Review the SLA at: <https://access.redhat.com/support/offerings/production/sla>
- **Enterprise: Premium**
 - Manage any size environment, including mission-critical environments
 - Premium 24x7 support and SLA
 - Maintenance and upgrades included
 - Review the SLA at: <https://access.redhat.com/support/offerings/production/sla>

All subscriptions include regular updates and releases of both Ansible Tower and Ansible core.

For more information, contact Ansible via the Red Hat Customer portal at <https://access.redhat.com/> or at <http://www.ansible.com/pricing/>.

2.4 Node Counting in Licenses

The Tower license defines the number of nodes that can be managed by Tower. A typical license will say ‘Enterprise Tower Up To 250 Nodes’, which sets the maximum number of nodes that can be managed at 250.

Tower counts nodes by the number of hosts in inventory. If more nodes are in the Tower inventory than are supported by the license, you will be unable to start any Jobs in Tower. If a dynamic inventory sync causes Tower to exceed the node count specified in the license, the dynamic inventory sync will fail.

If you have multiple hosts in inventory that have the same name, such as `webserver1`, they will be counted for licensing purposes as a single node. Note that this differs from the ‘Hosts’ count in Tower’s dashboard, which counts hosts in separate inventories separately.

2.5 License Features

Note: Ansible Tower version 2.2 introduced a separation of features for Basic (now called Self Support) versus Enterprise (now Enterprise: Standard) or Premium (now Enterprise: Premium) licenses.

The following list of features are available for all new Enterprise or Premium license users:

- Custom rebranding for login (*added in Ansible Tower 2.4.0*)
- SAML and RADIUS Authentication Support (*added in Ansible Tower 2.4.0*)
- Multi-Organization Support
- Activity Streams
- Surveys
- LDAP Support
- Active/Passive Redundancy
- System Tracking (*added in Ansible Tower 2.2.0*)

Enterprise license users with versions of Ansible Tower prior to 2.2 must import a new license file to enable System Tracking.

2.6 Tower Component Licenses

Ansible Tower includes some open source components. Ansible, Inc. supports Tower's use of and interactions with these components for both development and production purposes, subject to applicable terms and conditions. Unless otherwise agreed to in writing, the use of Ansible Tower is subject to the Ansible Software Subscription and Services Agreement located at <http://www.ansible.com/subscription-agreement>. Ansible Tower is a proprietary product offered by Ansible, Inc. and its use is not intended to prohibit the rights under any open source license.

To view the license information for the components included within Ansible Tower, refer to `/usr/share/doc/ansible-tower-<version>/README` where `<version>` refers to the version of Ansible Tower you have installed.

To view a specific license, refer to `/usr/share/doc/ansible-tower-<version>/*.txt`, where `*` is replaced by the license file name to which you are referring.

LOGGING IN

To log in to Tower, browse to the Tower interface at: `http://<Tower server name>/`



 **ANSIBLE
TOWER**
by Red Hat®

Welcome to Ansible Tower! Please sign in.

USERNAME

admin


PASSWORD

.....

SIGN IN

Log in using a valid Tower username and password.

The default username and password set during installation are *admin* and *password*, but the Tower administrator may have changed these settings during installation. If the default settings have not been changed, you can do so by

accessing the Users link from the Settings () Menu.

IMPORT A LICENSE

Tower requires a valid license to run. If you did not receive a license from Ansible directly or via email, or have issues with the license you received, refer to <http://www.ansible.com/license> for free and paid license options (including free trial licenses) or contact Ansible via the Red Hat Customer portal at <https://access.redhat.com/>.

To add your license, you must browse to the location where you saved your license (or save the license contents to a text file locally, if needed) and upload it. The uploaded license may be a plain text file or a JSON file, and must include properly formatted JSON code. Once uploaded, agree to the End User License Agreement and click **Submit**.

TOWER

TOWER LICENSE

Welcome to Ansible Tower! Please complete the steps below to acquire a license.

1

Please click the button below to visit Ansible's website to get a Tower license key.

REQUEST LICENSE

2

Choose your license file, agree to the End User License Agreement, and click submit.

* LICENSE FILE

BROWSE

No file selected.

* END USER LICENSE AGREEMENT

TOWER SOFTWARE END USER LICENSE AGREEMENT

Unless otherwise agreed to, and executed in a definitive agreement, between Ansible, Inc. ("Ansible") and the individual or entity ("Customer") signing or electronically accepting these terms of use for the Tower Software ("EULA"), all Tower Software, including any and all versions released or made available by Ansible, shall be subject to the Ansible Software Subscription and Services Agreement found at www.ansible.com/subscription-agreement ("Agreement"). Ansible is not responsible for any additional obligations, conditions or warranties agreed to between Customer and an authorized

☐ I agree to the End User License Agreement

SUBMIT

Copyright © 2016 Red Hat, Inc.

Once your license has been accepted, Tower navigates you to the main Ansible interface for the Dashboard (which you can access by clicking on the Ansible Tower logo at the top left of the screen as well).

For later reference, you can view this license from the Settings () Menu's 'VIEW YOUR LICENSE' link.

TOWER

PROJECTS

INVENTORIES

JOB TEMPLATES

JOB

admin

SETTINGS / LICENSE

DETAILS

LICENSE

Valid License

VERSION

3.0.0

LICENSE TYPE

Enterprise

SUBSCRIPTION

Enterprise Tower Up To 10000 Nodes

LICENSE KEY

E6242476b1e52223d599936fd50994e642618ad77478186f558e80234ac19f0a

EXPIRES ON

07/19/2016

TIME REMAINING

34 Days

HOSTS AVAILABLE

10000

HOSTS USED

1

HOSTS REMAINING

9999

If you are ready to upgrade, please contact us by clicking the button below

UPGRADE

LICENSE MANAGEMENT

Choose your license file, agree to the End User License Agreement, and click submit.

LICENSE FILE

BROWSE

No file selected.

END USER LICENSE AGREEMENT

TOWER SOFTWARE END USER LICENSE AGREEMENT

Unless otherwise agreed to, and executed in a definitive agreement, between Ansible, Inc. ("Ansible") and the individual or entity ("Customer") signing or electronically accepting these terms of use for the Tower Software ("EULA"), all Tower Software, including any and all versions released or made available by Ansible, shall be subject to the Ansible Software Subscription and Services Agreement found at www.ansible.com/subscription-agreement ("Agreement"). Ansible is not responsible for any additional obligations, conditions or warranties agreed to between Customer and an authorized distributor, or reseller, of the

☐ I agree to the End User License Agreement

SUBMIT

Copyright © 2016 Red Hat, Inc.

4.1 Adding a Tower License Manually

If you are in a situation where uploading a file is not allowed due to a locked down environment, you can add the Ansible Tower license by hand using Tower's API.


In Tower's REST API, at the `/api/v1/config/` endpoint, scroll down to the POST text entry box.

Add your valid license, the one you received directly from Ansible, to the POST box using the following as an example:

```
{
  "eula_accepted": "true",
  "subscription_name": "Enterprise Tower up to 100000 Nodes",
  "features": {},
  "instance_count": 100000,
  "trial": false,
  "contact_email": "maddux@hotmail.com",
  "company_name": "Dr. Maddux Golden",
  "license_type": "enterprise",
  "contact_name": "Dr. Maddux Golden",
  "license_date": 0000000000,
  "license_key": "xxx111xx111xxx1x1x1x1x1x1x1xxx1x1xx1x1xx1x1x1xxx111xx1x1xx1x"
}
```

Once done, hit the "POST" button and review your license.

THE TOWER DASHBOARD AND INTERFACE

Note: Ansible Tower 3.0 provides a streamlined interface, with the Settings () button offering access to administrative configuration options. Users of older versions of Ansible Tower (2.4.5 or older) can access most of these through the top-level navigational menu or from their “Setup” menu button.

The Tower Dashboard offers a friendly graphical framework for your IT orchestration needs. Across the top-left side of the Tower Dashboard, administrators can quickly navigate to their **Projects**, **Inventories**, **Job Templates**, and **Jobs**.


Across the top-right side of this interface, administrators can access the tools they need to configure organizations, users, groups, and permissions as well as view related documentation, access portal mode, and log out.



At the top of the Dashboard is a summary of your hosts, inventories, and projects. Each of these is linked to the corresponding object in Tower, for easy access.




On the main Tower Dashboard screen, a summary appears listing your current **Job Status**. Also available for review are summaries of **Recently Used Job Templates** and **Recently Run Jobs**.




Note: Clicking on the Ansible Tower logo at any time returns you to the Dashboard.

5.1 Tower Admin Menu

The Tower user menu is accessed by clicking  **admin**.

From here, you can:

- View/Edit the properties of the Tower admin account
- View the activity stream for that user (by clicking on the Activity Stream  button)
- View the **Organizations** which have been setup for the Tower user
- View the **Teams** to which the Tower user has been added
- View the **Granted Permissions** for this Tower admin account

TOWER

PROJECTS

INVENTORIES

JOB TEMPLATES

JOB

admin

SETTINGS / USERS / ADMIN

ADMIN

ADMIN

DETAILS

ORGANIZATIONS

TEAMS

GRANTED PERMISSIONS

* FIRST NAME

* LAST NAME

* EMAIL

admin@example.com

* USERNAME

admin

PASSWORD

SHOW

CONFIRM PASSWORD

SHOW

USER TYPE

System Administrator

CANCEL

SAVE

USERS

4

+ ADD

USERNAME


SEARCH

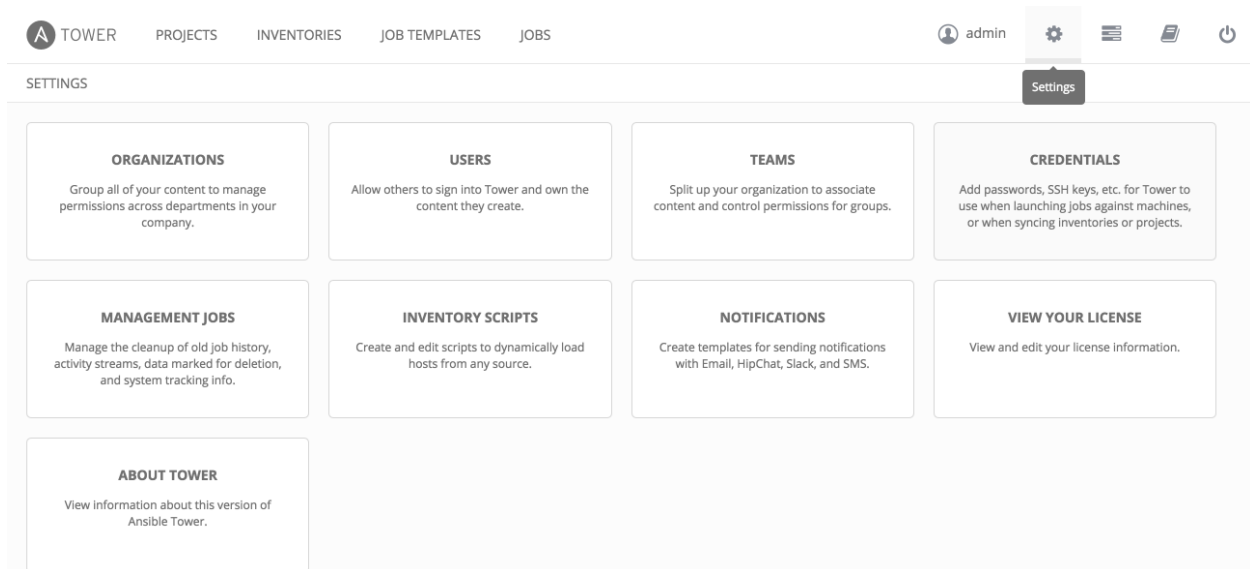
USERNAME	FIRST NAME	LAST NAME	ACTIONS
admin			
cdoge	Carter	Doge	
gdoge	Ginger	Doge	
jdoge	Josie	Doge	

ITEMS 1-4 OF 4

Copyright © 2016 Red Hat, Inc.


5.2 Settings Menu

To enter the Settings Menu screen for Ansible Tower, click the  button. This screen allows you to create your organizations, add credentials, add users and teams, schedule management jobs, and more. You can also view your license from the Settings Menu's 'View Your License' link.




Copyright © 2016 Red Hat, Inc.

5.3 My View – User Menu

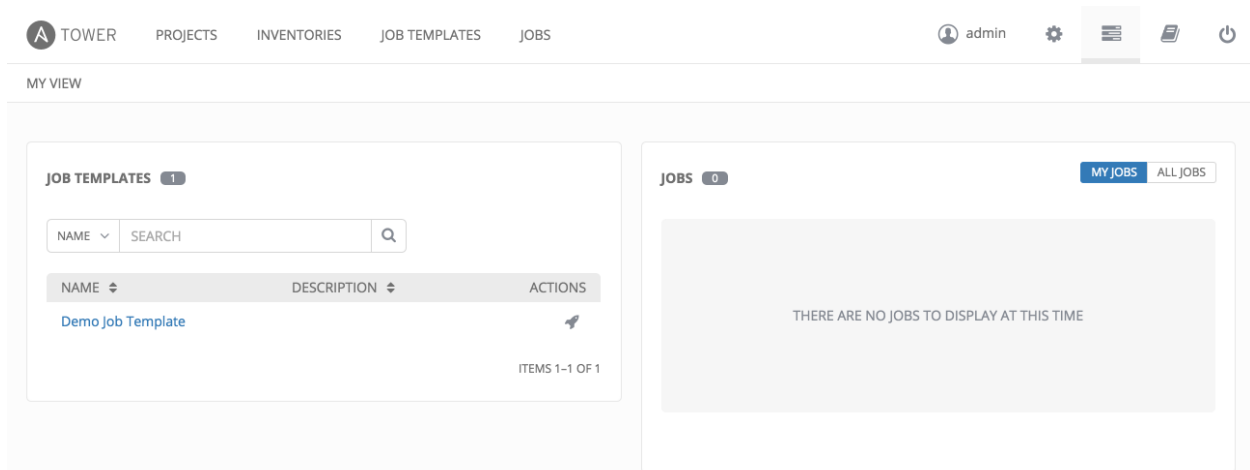
My View, a user's single-page view of jobs and job templates, can be accessed by clicking the My View () button.

My View is a simplified interface for users that need to run Ansible jobs, but that do not need an advanced knowledge of Ansible or Tower. My View could be used by, for instance, development teams, or even departmental users in non-technical fields.

My View offers Tower users a simplified, clean interface to the jobs that they are able to run, and the results of jobs that they have run in the past.


Pressing the  button beside a job in My View launches it, potentially asking some survey questions if the job is configured to do so.

Other portions of the interface are hidden from view until My View is exited.




Copyright © 2016 Red Hat, Inc.

My View can be accessed in two ways:

- via the My View () button at the top-right of the Tower interface
- by navigating to `https://<Tower server name>/portal`

My View displays two main sections—*Job Templates* and *Jobs*.

5.3.1 Job Templates

This shows the job templates that are available for the user to run. This list can be searched by **Name** or **Description**, and can be sorted by those keys as well. To launch a job template, click the  button. This launches the job, which can be viewed in **My Jobs**.

Note: Unlike Tower’s main interface, you are not automatically redirected to the Job view for the launched job. This view is still accessible via the **View Details** button for this job run in the **My Jobs** panel. This is useful for instances when a job fails and a non-technical user needs an Ansible expert look at what might have gone wrong.

5.3.2 Jobs

This shows the list of jobs that have run in the past.

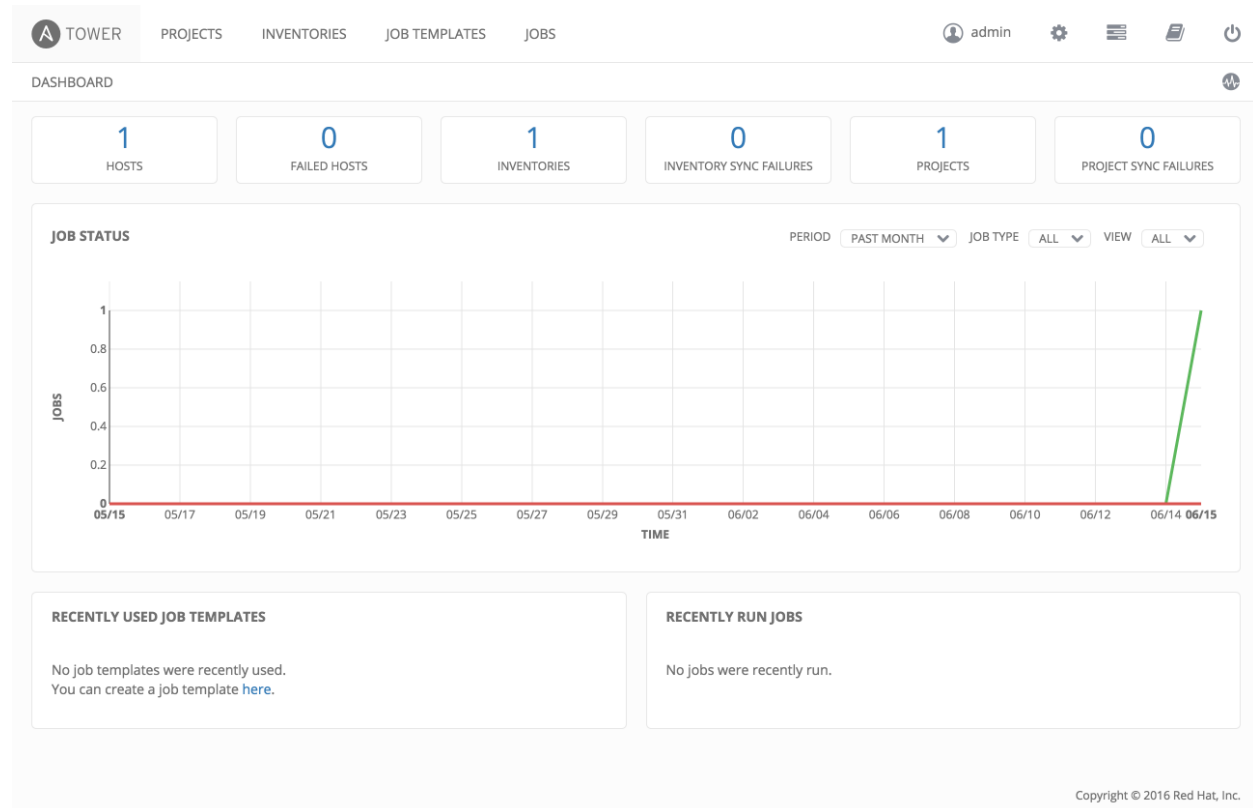
Sort for jobs specific to you by clicking on the **My Jobs** button or review all jobs you have access to view by clicking on the **All Jobs** button, next to the search bar.

- **My Jobs:** View jobs that you (as the user) ran .
- **All Jobs:** View your team members’ completed jobs, viewable based on your RBAC permissions.

For each job, you can view the Job **ID**, the **Status** of the job (*Running*, *Pending*, *Successful*, or *Failed*), its start time, and the job **Name**. The job list can be sorted by any of these fields. Clicking on the *Details* button opens a new window with the **Job Details** for that job (refer to *Jobs* for more information).

5.4 Dashboard Views

The central interface to Tower is the Dashboard. You will use the dashboard to quickly view job statuses, recently run jobs, and recently used job templates.



5.4.1 Job Status

The Job Status graph displays the number of successful and failed jobs over a specified time period. You can choose to limit the job types that are viewed, and to change the time horizon of the graph.


5.4.2 Recently Used Job Templates

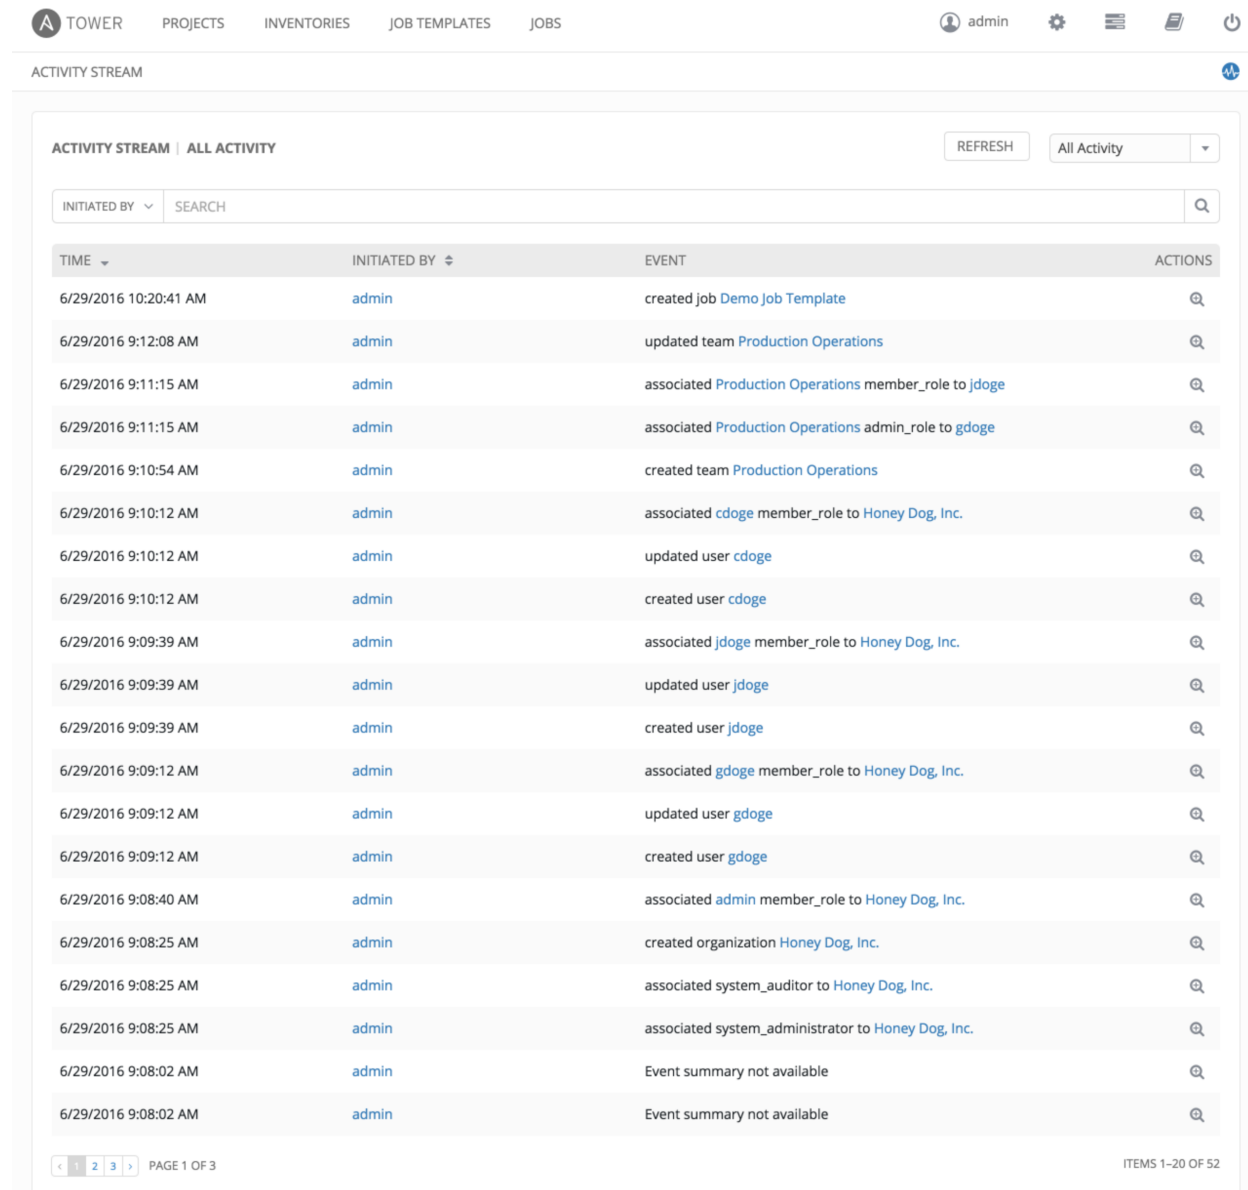
The Jobs section of this display shows a summary of the most recently used jobs. You can also access this summary by clicking on the **Jobs** entry in the main navigation menu.

5.4.3 Recently Run Jobs

The Recently Run Jobs section displays which jobs were most recently run, their status, and notes when they were run as well.

5.5 Activity Streams





















Most screens in Tower have an Activity Stream () button. Clicking this brings up the **Activity Stream** for this object.



ACTIVITY STREAM | ALL ACTIVITY


REFRESH All Activity

INITIATED BY SEARCH

TIME	INITIATED BY	EVENT	ACTIONS
6/29/2016 10:20:41 AM	admin	created job Demo Job Template	
6/29/2016 9:12:08 AM	admin	updated team Production Operations	
6/29/2016 9:11:15 AM	admin	associated Production Operations member_role to jdoge	
6/29/2016 9:11:15 AM	admin	associated Production Operations admin_role to gdoge	
6/29/2016 9:10:54 AM	admin	created team Production Operations	
6/29/2016 9:10:12 AM	admin	associated cdoge member_role to Honey Dog, Inc.	
6/29/2016 9:10:12 AM	admin	updated user cdoge	
6/29/2016 9:10:12 AM	admin	created user cdoge	
6/29/2016 9:09:39 AM	admin	associated jdoge member_role to Honey Dog, Inc.	
6/29/2016 9:09:39 AM	admin	updated user jdoge	
6/29/2016 9:09:39 AM	admin	created user jdoge	
6/29/2016 9:09:12 AM	admin	associated gdoge member_role to Honey Dog, Inc.	
6/29/2016 9:09:12 AM	admin	updated user gdoge	
6/29/2016 9:09:12 AM	admin	created user gdoge	
6/29/2016 9:08:40 AM	admin	associated admin member_role to Honey Dog, Inc.	
6/29/2016 9:08:25 AM	admin	created organization Honey Dog, Inc.	
6/29/2016 9:08:25 AM	admin	associated system_auditor to Honey Dog, Inc.	
6/29/2016 9:08:25 AM	admin	associated system_administrator to Honey Dog, Inc.	
6/29/2016 9:08:02 AM	admin	Event summary not available	
6/29/2016 9:08:02 AM	admin	Event summary not available	

PAGE 1 OF 3 ITEMS 1-20 OF 52

Copyright © 2016 Red Hat, Inc.

An Activity Stream shows all changes for a particular object. For each change, the Activity Stream shows the time of the event, the user that initiated the event, and the action. Clicking on the Examine () button shows the event log for the change.

The screenshot displays the Ansible Tower web interface. At the top, there's a navigation bar with 'TOWER' and tabs for 'PROJECTS', 'INVENTORIES', 'JOB TEMPLATES', and 'JOBS'. The user 'admin' is logged in. The main section is titled 'ACTIVITY STREAM'. On the left, there's a sidebar with 'ACTIVITY STREAM | ALL ACTIVITY' and filters for 'INITIATED BY' and 'TIME'. The main area shows a list of activity events. A modal window titled 'EVENT 45' is open, showing details for a specific event.

EVENT 45

INITIATED BY: admin on 6/15/2016 10:28:29 AM

ACTION: updated tower_settings

CHANGES:

```
{
  "value": [
    {"subscription_name": "Enterprise Tower up to 500000000 Nodes"},
    {"subscription_name": "Enterprise Tower up to 10000 Nodes"}
  ]
}
```

OK

The background activity stream table shows the following entries:


TIME	INITIATED BY	ACTION
6/15/2016 10:28:29 AM	admin	updated tower_settings
6/15/2016 10:23:22 AM		
6/15/2016 10:15:13 AM		
6/15/2016 10:15:13 AM		
6/15/2016 10:04:09 AM	system	updated user admin
6/15/2016 9:58:42 AM	system	created job_template Demo Job Template
6/15/2016 9:58:42 AM	system	
6/15/2016 9:58:42 AM		


The Activity Stream can be filtered by the initiating user (or the system, if it was system initiated), and by any related Tower object, such as a particular credential, job template, or schedule.

The Activity Stream on the main Dashboard shows the Activity Stream for the entire Tower instance. Most pages in Tower allow viewing an activity stream filtered for that specific object.

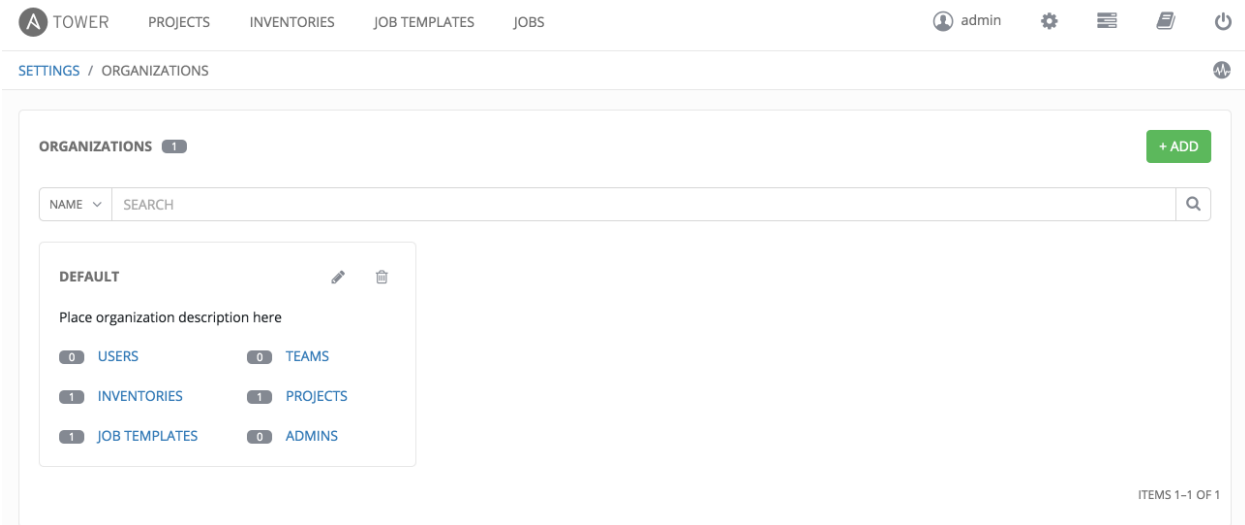
ORGANIZATIONS


An **Organization** is a logical collection of **Users**, **Teams**, **Projects**, and **Inventories**, and is the highest level in the Tower object hierarchy.



The **Organizations** link from the Settings () menu displays all of the existing organizations for your installation of Tower. Organizations can be searched by **Name** or **Description**. Modify and remove organizations using the **Edit** and **Delete** buttons.

Note: Tower creates a default organization automatically. Users of Tower with a Self-Support level license (formerly called Basic) only have the default organization available and should **not** delete it. Users of older versions of Tower (prior to 2.2) will not see this default organization.



Enterprise: Standard and Enterprise: Premium Tower licenses allow you to create a new Organization by selecting the  button.

Note: If you are using Ansible Tower with a Self-Support level license (formerly called Basic), you must use the default Organization. Do not delete it and try to add a new Organization, or you will break your Tower setup. Only Enterprise: Standard or Enterprise: Premium Tower licenses have the ability to add new Organizations beyond the default.

1. Enter the **Name** for your Organization.
2. Optionally, enter a **Description** for the Organization.

Click **Save** to finish creating the Organization.

TOWER

PROJECTS

INVENTORIES

JOB TEMPLATES

JOB

admin

SETTINGS / ORGANIZATIONS / CREATE ORGANIZATION

NEW ORGANIZATION

DETAILS

PERMISSIONS

NOTIFICATIONS

* NAME

DESCRIPTION

CANCEL

SAVE

ORGANIZATIONS 1

+ ADD

NAME

SEARCH

DEFAULT

Place organization description here

0

USERS

0

TEAMS

1

INVENTORIES

1

PROJECTS

1

JOB TEMPLATES

0

ADMINS

ITEMS 1-1 OF 1

Copyright © 2016 Red Hat, Inc.

Once created, Tower displays the Organization details, and allows for the managing of users and administrators for the Organization.

TOWER PROJECTS INVENTORIES JOB TEMPLATES JOBS admin

SETTINGS / ORGANIZATIONS / HONEY DOG, INC.

HONEY DOG, INC.

DETAILS PERMISSIONS NOTIFICATIONS

* NAME DESCRIPTION

Honey Dog, Inc. A capable company making capable things

CANCEL SAVE

ORGANIZATIONS 2 + ADD

NAME SEARCH

DEFAULT

Place organization description here

0 USERS 0 TEAMS

1 INVENTORIES 1 PROJECTS

1 JOB TEMPLATES 0 ADMINS

HONEY DOG, INC.

A capable company making capable things

0 USERS 0 TEAMS

0 INVENTORIES 0 PROJECTS

0 JOB TEMPLATES 0 ADMINS

ITEMS 1-2 OF 2

Copyright © 2016 Red Hat, Inc.

6.1 Organizations - Permissions

Clicking on **Permissions** (beside **Details** when viewing your organization), allows you to review, grant, edit, and remove associated permissions for users as well as team members.

SETTINGS / ORGANIZATIONS / HONEY DOG, INC.


HONEY DOG, INC.

DETAILS PERMISSIONS NOTIFICATIONS

USER SEARCH + ADD

USER	ROLE	TEAM ROLES
admin	MEMBER SYSTEM ADMINISTRATOR	
cdoge	SYSTEM ADMINISTRATOR	
jdoge	SYSTEM AUDITOR	

ITEMS 1-3 OF 3

Click the  button to create new permissions for this organization.

HONEY DOG, INC. | ADD PERMISSIONS

1 Please select Users / Teams from the lists below.

USERS

TEAMS

FIRST NAME

SEARCH

FIRST NAME	LAST NAME	USERNAME
<input checked="" type="checkbox"/> Ginger	Doge	gdoge
<input checked="" type="checkbox"/> Josie	Doge	jdoge

ITEMS 1-2 OF 2

2 Please assign roles to the selected users/teams

KEY

Ginger Doge USER

Admin

Member

Josie Doge USER

Auditor

Member

SAVE

CANCEL

In this example, two users have been selected and each have been granted permissions for this organization.

You can also assign permissions to teams. Note that you do not have to choose between teams or users, and that you can assign permissions to both at the same time.

HONEY DOG, INC. | ADD PERMISSIONS

1 Please select Users / Teams from the lists below.

USERS

TEAMS

NAME

SEARCH

Q

NAME ^

☒ Production Operations

ITEMS 1-1 OF 1

2 Please assign roles to the selected users/teams

KEY

Ginger Doge USER

Admin

Member

X

Josie Doge USER

Auditor

X

Production Operations TEAM

Admin

Member

X

SAVE

CANCEL

6.2 Organizations - Notifications

Clicking on **Notifications** (beside **Details** when viewing your organization), allows you to easily manage notifications for this organization.

Click on the  button to create a notification.

Supported notification sources include Slack, Email, SMS (via Twilio), HipChat, and more. Refer to [Notifications](#) for more information.

SETTINGS / NOTIFICATION TEMPLATES / CREATE NOTIFICATION TEMPLATE

NEW NOTIFICATION TEMPLATE

* NAME

DESCRIPTION

* ORGANIZATION

* TYPE

Choose a type

CANCEL

SAVE

6.3 Organizations - Users

The **Users** submenu of an Organization displays all the Users associated with this Organization. A User is someone with access to Tower with associated roles and Credentials. Expand the Users menu by selecting **Users** in the Organization Summary.

TOWER

PROJECTS

INVENTORIES

JOB TEMPLATES

JOB

admin

SETTINGS / ORGANIZATIONS / HONEY DOG, INC. / USERS

HONEY DOG, INC. | USERS 0

+ ADD USER

PLEASE ADD ITEMS TO THIS LIST

ORGANIZATIONS 2

+ ADD

NAME SEARCH

DEFAULT

Place organization description here

0 USERS

0 TEAMS

1 INVENTORIES

1 PROJECTS

1 JOB TEMPLATES

0 ADMINS

HONEY DOG, INC.

A capable company making capable things

0 USERS

0 TEAMS

0 INVENTORIES

0 PROJECTS

0 JOB TEMPLATES

0 ADMINS

ITEMS 1-2 OF 2

Copyright © 2016 Red Hat, Inc.

This menu allows you to manage the user membership for this Organization. (User membership may also be managed

6.3. Organizations - Users

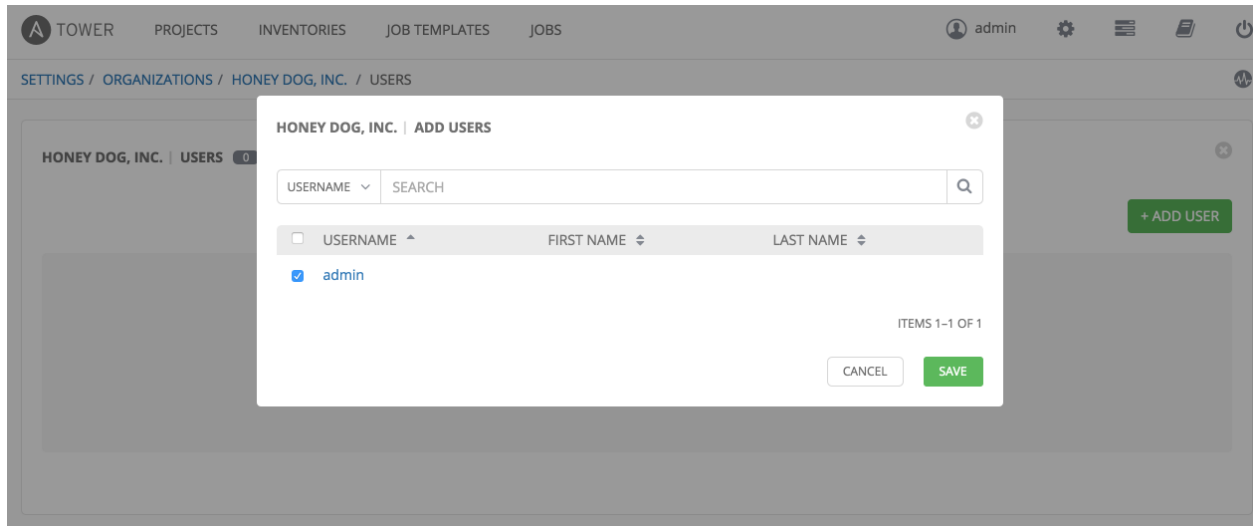
25



on a per-user basis via the **Users** link available from the Settings menu.) The user list may be sorted and searched by **Username**, **First Name**, or **Last Name**. Existing users may also be modified and removed using the **Edit** and **Delete** buttons. Clicking on a user brings up that user's details, which can then be edited. For more information, refer to *Users*.

+ ADD USER

To add existing users to the Organization, click the **+ ADD USER** button. Then, select one or more users from the list of available users by clicking the **Select** checkbox or clicking anywhere on the user row. Click the **Select** button when done.



6.4 Organization - Administrators

An Organization Administrator is a type of user that has the rights to create, modify, or delete objects in the Organization, including Projects, Teams, and Users in that Organization. Expand the **Admins** submenu by selecting **Admins** in the Organization Summary.

TOWER

PROJECTS

INVENTORIES

JOB TEMPLATES

JOBS

admin

SETTINGS / ORGANIZATIONS / HONEY DOG, INC. / ADMINS

HONEY DOG, INC. | ADMINS 0

+ ADD ADMINISTRATOR

PLEASE ADD ITEMS TO THIS LIST

ORGANIZATIONS 2

+ ADD

NAME SEARCH

DEFAULT

Place organization description here

0 USERS

0 TEAMS

1 INVENTORIES

1 PROJECTS

1 JOB TEMPLATES

0 ADMINS

HONEY DOG, INC.

A capable company making capable things

0 USERS

0 TEAMS

0 INVENTORIES

0 PROJECTS

0 JOB TEMPLATES

0 ADMINS

ITEMS 1-2 OF 2

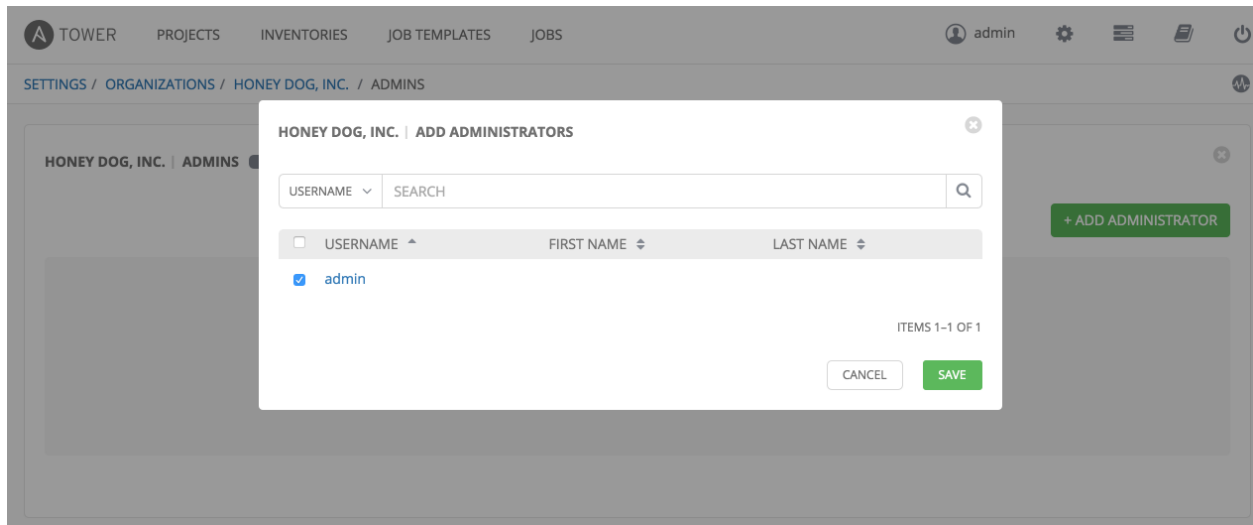
Copyright © 2016 Red Hat, Inc.

This menu displays a list of the users that are currently setup as an administrator of the Organization. The administrator list may be sorted and searched by **Username**, **First Name**, or **Last Name**.

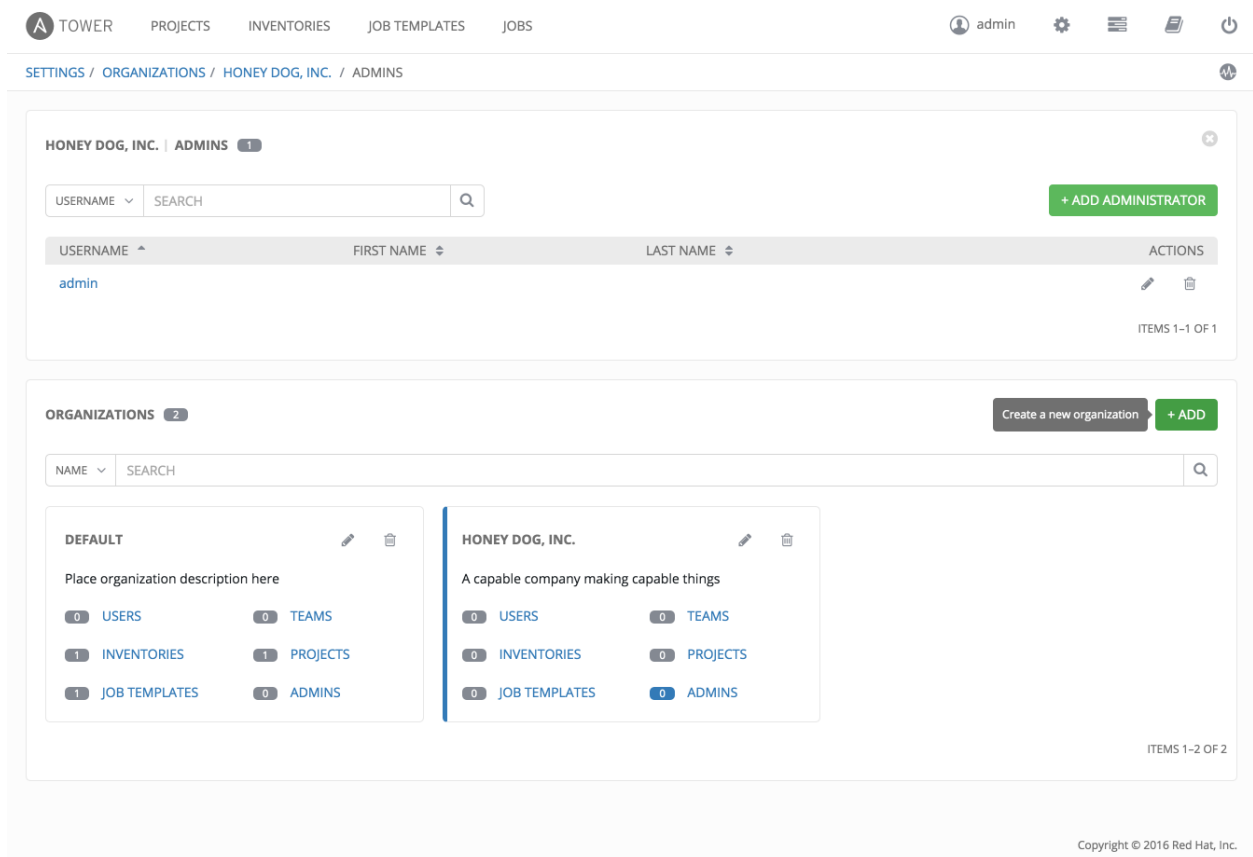
Note: Any user marked as a ‘Superuser’ is implicitly an administrator of all Organizations, and is not displayed here.

+ ADD ADMINISTRATOR


To add an administrator to the Organization, click the button. Select one or more users from the list of available users by clicking the **Select** checkbox or clicking anywhere on the user row. Click the **Select** button when done.




Note: A user must first be added to the Organization before it can be added to the list of administrators for that Organization.



USERS

A **User** is someone who has access to Tower with associated permissions and credentials. The Users link (found by clicking on the Settings () menu and selecting **Users**) allows you to manage the all Tower users. The User list may be sorted and searched by **Username**, **First Name**, or **Last Name** (click on the **USERNAME** button beside the search field to toggle your sorting preference).



There are three types of Tower Users:

1. **Normal User:** Normal Users have read and write access limited to the inventory and projects for which that user has been granted the appropriate roles and privileges.
2. **System Auditor:** Auditors implicitly inherit the read-only capability for all objects within the Tower environment.
3. **System Administrator:** a Tower System Administrator (also the Superuser) has admin, read, and write privileges over the entire Tower installation. A System Administrator is typically responsible for managing all aspects of Tower and delegating responsibilities for day-to-day work to various Users and System Auditors.



Note: The initial user (usually “admin”) created by the Tower installation process is a Superuser. One Superuser must always exist. To delete the “admin” user account, you must first create another Superuser account.

To create a new user click the  button, which opens the Create User dialog.

Enter the appropriate details into the following fields:

- First Name
- Last Name
- Email
- Username
- Organization (Choose from an existing organization—this is the default organization if you are using a Self-Supported level license.)
- Password
- Confirmation Password
- User Type (The System Administrator, superuser, has full system administration privileges for Tower. Assign with caution!)

Select **Save** when finished.

Once the user is successfully created, the **User** dialog opens for that newly created User. This is the same menu that is opened if the Edit () button beside a User is clicked from the **Users** link within Tower's Settings (). Here, the User's **Organizations**, **Teams** and **Granted Permissions**, as well as other user membership details, may be reviewed and modified.

TOWER PROJECTS INVENTORIES JOB TEMPLATES JOBS
 admin

SETTINGS / USERS / JDOGE

JDOGE AUDITOR

DETAILS ORGANIZATIONS TEAMS GRANTED PERMISSIONS

* FIRST NAME: Josie * LAST NAME: Doge * EMAIL: josie@honeydog.com

* USERNAME: jdoge PASSWORD: SHOW CONFIRM PASSWORD: SHOW

USER TYPE: System Auditor

CANCEL SAVE

USERS 3 + ADD

USERNAME SEARCH

USERNAME	FIRST NAME	LAST NAME	ACTIONS
admin			
gdoge	Ginger	Doge	
jdoge	Josie	Doge	

ITEMS 1-3 OF 3

Copyright © 2016 Red Hat, Inc.

7.1 User Types - Quick View

Once a user has been created, you can easily view permissions and user type information by looking beside their user name in the User overview screen.

SETTINGS / USERS / JDOGE

JDOGE AUDITOR

DETAILS ORGANIZATIONS TEAMS GRANTED PERMISSIONS

* FIRST NAME: Josie * LAST NAME: Doge * EMAIL: jdoge@honeydog.com

View user labels here for Auditor, Admin, LDAP, etc.

If the user account is associated with an enterprise-level authentication method (such as SAML, RADIUS, or LDAP), the user type may look like:

SETTINGS / USERS / JDOGE



JDOGE
RADIUS

DETAILS
ORGANIZATIONS
TEAMS
GRANTED PERMISSIONS

* FIRST NAME
* LAST NAME
* EMAIL

Josie
Doge
jdoge@honeydog.com

If the user account is associated with a social authentication method, the user type will look like:

SETTINGS / USERS / JDOGE



JDOGE
SOCIAL

DETAILS
ORGANIZATIONS
TEAMS
GRANTED PERMISSIONS

* FIRST NAME
* LAST NAME
* EMAIL

Josie
Doge
jdoge@honeydog.com

7.2 Users - Organizations

This displays the list of organizations of which that user is a member. This list may be searched by Organization Name or Description. Organization membership cannot be modified from this display panel.

A TOWER
PROJECTS
INVENTORIES
JOB TEMPLATES
JOBS

admin

SETTINGS / USERS / JDOGE

JDOGE
AUDITOR

DETAILS
ORGANIZATIONS
TEAMS
GRANTED PERMISSIONS

NAME SEARCH

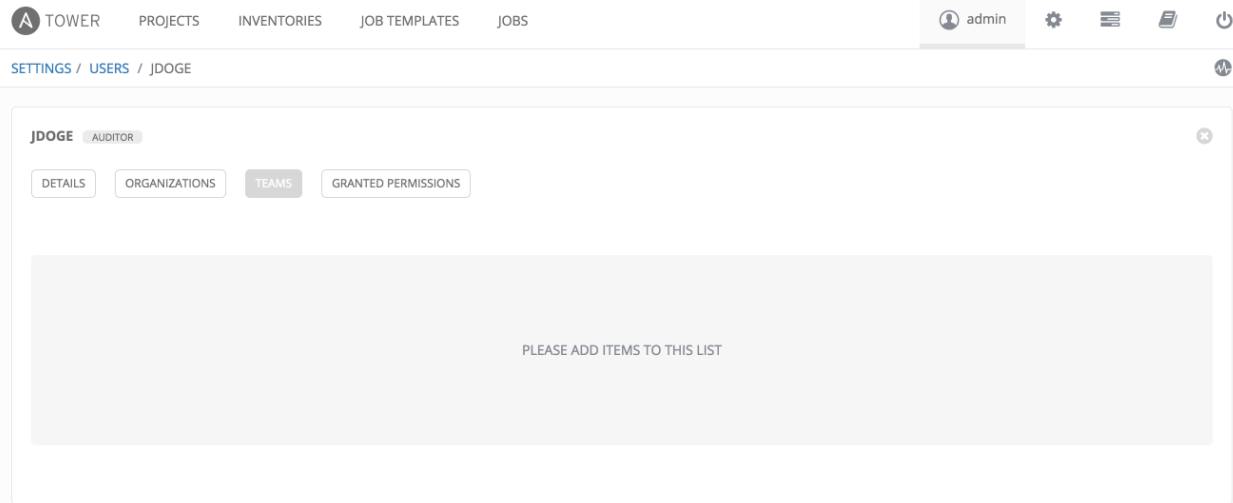
NAME	DESCRIPTION
Honey Dog, Inc.	A capable company making capable things

ITEMS 1-1 OF 1

7.3 Users - Teams

This displays the list of teams of which that user is a member. This list may be searched by **Team Name** or **Description**. Team membership cannot be modified from this display panel. For more information, refer to [Teams](#).

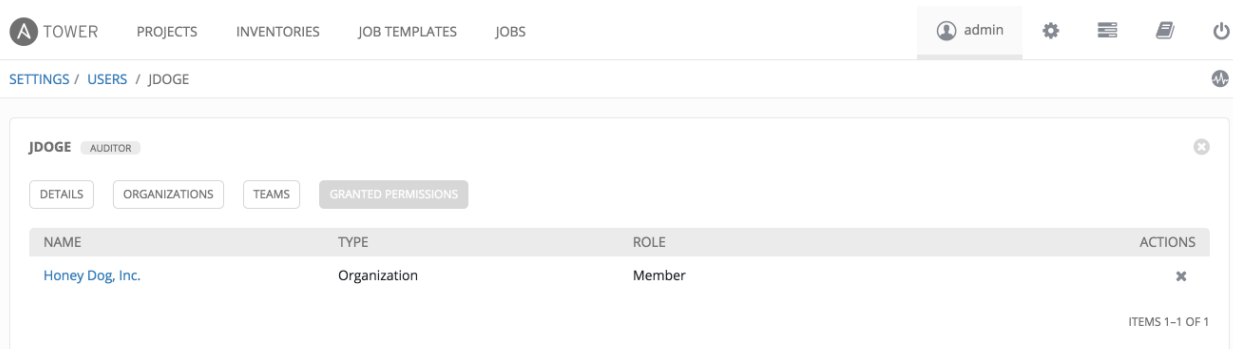
Until a Team has been created and the user has been assigned to that team, the assigned Teams Details for the User appears blank.




7.4 Users - Granted Permissions

The set of Granted Permissions assigned to this user (role-based access controls) that provide the ability to read, modify, and administer projects, inventories, job templates, and other Tower elements are Privileges.

This menu displays a list of the privileges that are currently available for a selected User. The privileges list may be sorted and searched by **Name**, **Type**, or **Role**.



To remove Granted Permissions for a particular User, click the Disassociate () button under **Actions**. This launches a **Remove Role** dialog, asking you to confirm the disassociation.

TEAMS

A **Team** is a subdivision of an organization with associated users, projects, credentials, and permissions. Teams provide a means to implement role-based access control schemes and delegate responsibilities across organizations. For instance, permissions may be granted to a whole Team rather than each user on the Team.

You can create as many Teams of users as make sense for your Organization. Each Team can be assigned permissions, just as with Users.

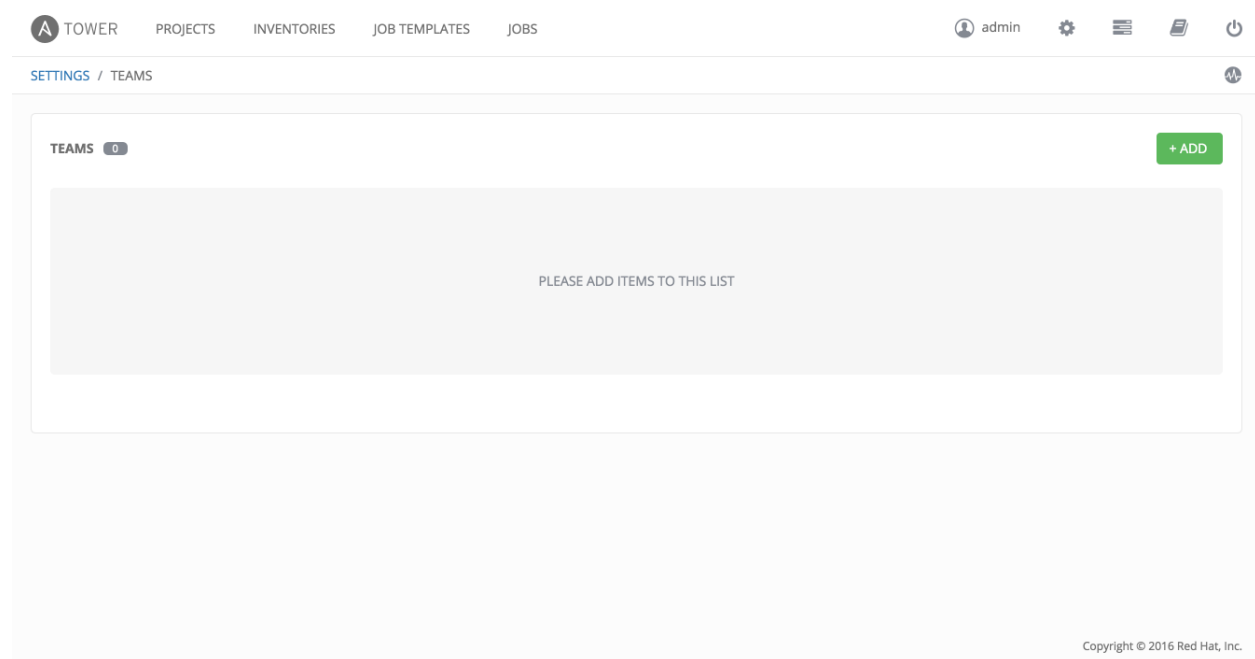
Teams can also scalably assign ownership for Credentials, preventing multiple Tower interface click-throughs to assign the same Credentials to the same user.




The Teams link, accessible by clicking on the Settings () button and then selecting **Teams**, allows you to manage the teams for Tower. The team list may be sorted and searched by **Name**, **Description**, or **Organization**.

Buttons located in the upper right corner of the **Team** tab provide the following actions:

- Create a new team
- View Activity Stream



To create a new Team, click the  button.

NEW TEAM

DETAILS USERS GRANTED PERMISSIONS


* NAME DESCRIPTION * ORGANIZATION

CANCEL SAVE

Enter the appropriate details into the following fields:

- Name
- Description
- Organization (Choose from an existing organization)

Select **Save**.

Once the Team is successfully created, Tower opens the **Details** dialog, which also allows you to review and edit your Team information. This is the same menu that is opened if the Edit () button is clicked from the **Teams** link. You can also review **Users** and **Granted Permissions** associated with this Team.

PRODUCTION OPERATIONS

DETAILS USERS GRANTED PERMISSIONS

* NAME DESCRIPTION * ORGANIZATION



Production Operations ProdOps team Honey Dog, Inc.

CANCEL SAVE

TEAMS 1

+ ADD

NAME SEARCH

NAME	DESCRIPTION	ORGANIZATION	ACTIONS
Production Operations	ProdOps team	Honey Dog, Inc.	 


ITEMS 1-1 OF 1

Copyright © 2016 Red Hat, Inc.

8.1 Teams - Users

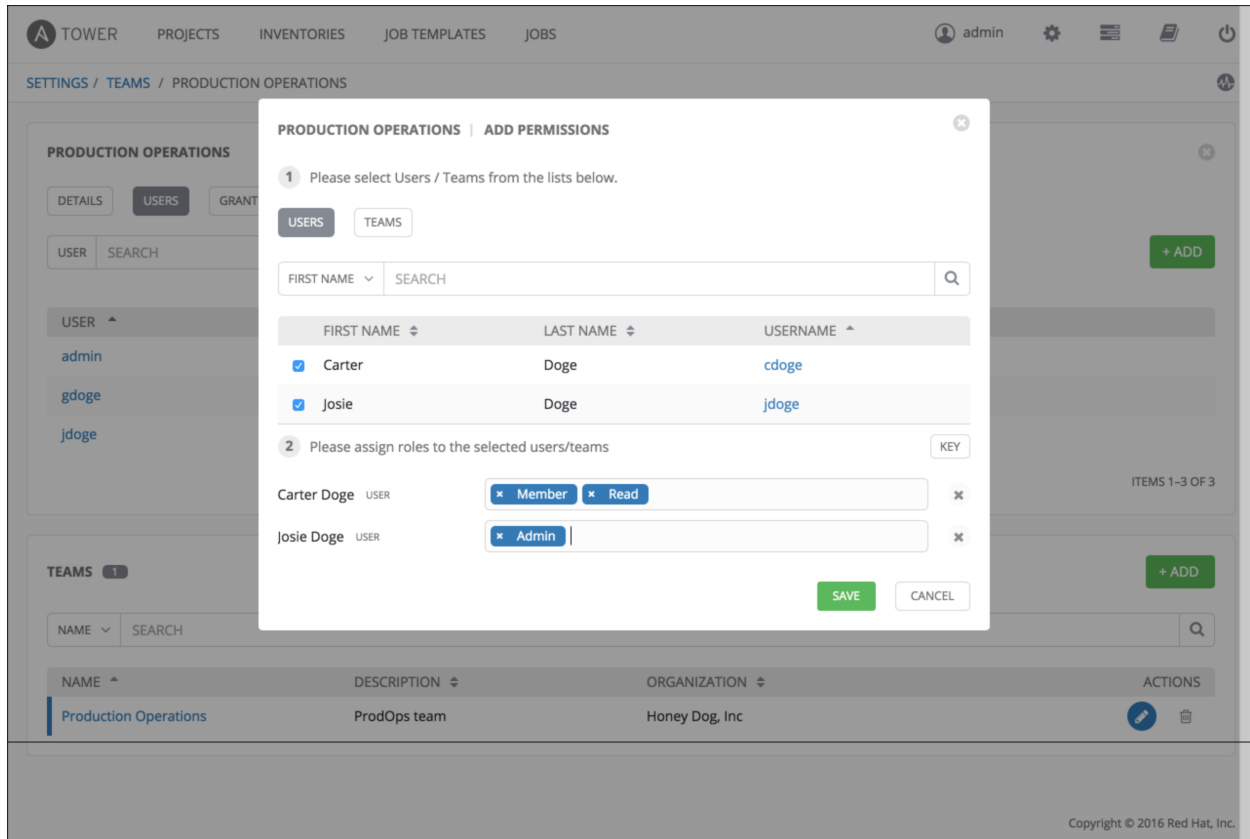
This menu displays the list of Users that are members of this Team. This list may be searched by **Username**, **First Name**, or **Last Name**. For more information, refer to *Users*.

The screenshot displays the Ansible Tower interface for managing teams. At the top, navigation links include TOWER, PROJECTS, INVENTORIES, JOB TEMPLATES, and JOBS. The user 'admin' is logged in. The breadcrumb trail shows 'SETTINGS / TEAMS / PRODUCTION OPERATIONS'. The main content area is titled 'PRODUCTION OPERATIONS' and contains three tabs: 'DETAILS', 'USERS', and 'GRANTED PERMISSIONS'. The 'USERS' tab is selected, showing a search bar and a table of users. The table has columns for 'USER', 'ROLE', and 'TEAM ROLES'. Three users are listed: 'admin' (SYSTEM ADMINISTRATOR), 'gdodge' (SYSTEM ADMINISTRATOR), and 'jdodge' (SYSTEM AUDITOR). A '+ ADD' button is located in the top right corner of the users section. Below the users section, there is a 'TEAMS' section with a '+ ADD' button and a table. The table has columns for 'NAME', 'DESCRIPTION', 'ORGANIZATION', and 'ACTIONS'. One team is listed: 'Production Operations' with description 'ProdOps team' and organization 'Honey Dog, Inc'. The 'ACTIONS' column contains edit and delete icons. The footer of the interface shows 'Copyright © 2016 Red Hat, Inc.'

To add users to the Team, click the  button. Then, select one or more Users from the list of available users by clicking the **Select** checkbox or clicking anywhere on the user row.

Next, grant permissions for the User(s) selected for this team. Choose one or more of the following permission settings:

- **Admin:** This User should have privileges to manage all aspects of the team
- **Member:** This User should be a member of the team
- **Read:** This User may view settings for the team



8.2 Teams - Granted Permissions

Selecting the **Granted Permissions** view displays a list of the permissions that are currently available for this Team. The permissions list may be sorted and searched by **Name**, **Inventory**, **Project** or **Permission** type.

TOWER

PROJECTS

INVENTORIES

JOB TEMPLATES

JOBS

admin

SETTINGS / TEAMS / PRODUCTION OPERATIONS

PRODUCTION OPERATIONS

DETAILS

USERS

GRANTED PERMISSIONS

NAME	TYPE	ROLE	ACTIONS
Production Operations	Team	Read	✕

ITEMS 1-1 OF 1

TEAMS 1

+ ADD

NAME

SEARCH

NAME	DESCRIPTION	ORGANIZATION	ACTIONS
Production Operations	ProdOps team	Honey Dog, Inc	✎ ✕

Copyright © 2016 Red Hat, Inc.

The set of privileges assigned to Teams that provide the ability to read, modify, and administer projects, inventories, and other Tower elements are permissions.

By default, the Team is given the “read” permission (also called a role).

Permissions must be set explicitly via an Inventory, Project, Job Template, or within the Organization view.

CREDENTIALS

Credentials are utilized by Tower for authentication when launching Jobs against machines, synchronizing with inventory sources, and importing project content from a version control system.

Tower credentials are imported and stored encrypted in Tower, and are not retrievable in plain text on the command line by any user. Once a password or key has been entered into the Tower interface, it is encrypted and inserted into the Tower database, and cannot be retrieved from Tower. You can grant users and teams the ability to use these credentials, without actually exposing the credential to the user. If you have a user move to a different team or leave the organization, you don't have to re-key all of your systems just because that credential was available in Tower.

Note: Tower encrypts passwords and key information in the Tower database and never makes secret information visible via the API.

9.1 Understanding How Credentials Work


The encryption/decryption algorithm uses Electronic Code Book (ECB) as the mode of operation with AES-128 as the block cipher. The 128-bit AES key is derived from the `SECRET_KEY` (found in the `awx` settings). Specific, sensitive, Model fields in Tower are encrypted and include:

```
Credential: password, ssh_key_data, ssh_key_unlock, become_password, vault_password
UnifiedJob: start_args
```

Data is encrypted before it is saved to the database and is decrypted as is needed in Tower. The encryption/decryption process derives the AES-128 bit encryption key from `<SECRET_KEY, field_name, primary_key>` where `field_name` is the name of the Model field and `primary_key` is the database assigned auto-incremented record ID. Thus, if any attribute used in the key generation process changes, Tower fails to correctly decrypt the secret.

Note: The rules of encryption and decryption for Ansible Tower also apply to one field outside of credentials, the Unified Job `start_args` field, which is used through the `job`, `ad_hoc_command`, and `system_job` data types.

9.2 Getting Started with Credentials

The **Credentials** link, accessible from the Setting () button, displays a list of all available Credentials. It can be sorted and searched by **Name**, **Description**, **Type**, or **Owners**.

TOWER

PROJECTS

INVENTORIES

JOB TEMPLATES

JOB

admin

SETTINGS / CREDENTIALS

CREDENTIALS 1

+ ADD

NAME

SEARCH

NAME	DESCRIPTION	TYPE	OWNERS	ACTIONS
Demo Credential		Machine	admin	<div></div> <div></div>

Credentials added to a Team are made available to all members of the Team, whereas credentials added to a User are only available to that specific User by default.

To help you get started, a Demo Credential has been created for your use.

Clicking on the link for the **Demo Credential** takes you to the **Details** view of this Credential.

TOWER

PROJECTS

INVENTORIES

JOB TEMPLATES

JOB

admin

SETTINGS / CREDENTIALS / DEMO CREDENTIAL

DEMO CREDENTIAL

DETAILS

PERMISSIONS

* NAME

Demo Credential

DESCRIPTION

ORGANIZATION

* TYPE

Machine

TYPE DETAILS

USERNAME

admin

PASSWORD

SHOW

☐ Ask at runtime?

PRIVATE KEY PASSPHRASE

SHOW

☐ Ask at runtime?

PRIVILEGE ESCALATION

Choose a privilege escalation

VAULT PASSWORD

SHOW

☐ Ask at runtime?

PRIVATE KEY

CANCEL

SAVE

CREDENTIALS 1

+ ADD

NAME

SEARCH

NAME	DESCRIPTION	TYPE	OWNERS	ACTIONS
Demo Credential		Machine	admin	<div></div> <div></div>

Copyright © 2016 Red Hat, Inc.

Clicking on **Permissions** shows you users and teams associated with this Credential and their granted roles (owner, admin, auditor, etc.)

TOWER

PROJECTS

INVENTORIES

JOB TEMPLATES

JOBS

admin

SETTINGS / CREDENTIALS / DEMO CREDENTIAL

DEMO CREDENTIAL

DETAILS

PERMISSIONS

USER

SEARCH

+ ADD

USER	ROLE	TEAM ROLES
admin	<div>OWNER</div> <div>SYSTEM ADMINISTRATOR</div>	
gdoge	<div>SYSTEM ADMINISTRATOR</div>	
jdoge	<div>SYSTEM AUDITOR</div>	

ITEMS 1-3 OF 3

CREDENTIALS

1


+ ADD

NAME


SEARCH

NAME	DESCRIPTION	TYPE	OWNERS	ACTIONS
Demo Credential		Machine	admin	<div></div> <div></div>

Copyright © 2016 Red Hat, Inc.

You can click the  button to assign this **Demo Credential** to additional Users or Teams.

9.3 Add a New Credential

The  button located in the upper right corner of the **Credentials** screen allows you to create a new credential.

CREATE CREDENTIAL

DETAILS **PERMISSIONS**

* NAME DESCRIPTION ORGANIZATION

* TYPE
Choose a type

CANCEL SAVE

CREDENTIALS **1** **+ ADD**

NAME SEARCH

NAME	DESCRIPTION	TYPE	OWNERS	ACTIONS
Demo Credential		Machine	admin	

Copyright © 2016 Red Hat, Inc.

Enter the appropriate details depending on the type of credential and select **Save**.

9.4 Credential Types

Topics:

- *Machine*
- *Network*
- *Source Control*
- *Amazon Web Services*
- *Rackspace*
- *VMware vCenter*
- *Red Hat Satellite 6*
- *Red Hat CloudForms*
- *Google Compute Engine*
- *Microsoft Azure Classic (Deprecated)*
- *Microsoft Azure Resource Manager*
- *OpenStack*

9.4.1 Machine

Machine credentials enable Tower to invoke Ansible on hosts under your management. Just like using Ansible on the command line, you can specify the SSH username, optionally provide a password, an SSH key, a key password, or even have Tower prompt the user for their password at deployment time. They define ssh and user-level privilege escalation access for playbooks, and are used when submitting jobs to run playbooks on a remote host.

The screenshot shows the 'CREATE CREDENTIAL' form in Ansible Tower. The 'DETAILS' tab is active. The form includes the following fields and options:

- *NAME**: Text input field.
- DESCRIPTION**: Text input field.
- ORGANIZATION**: Text input field with a search icon.
- *TYPE**: Dropdown menu set to 'Machine'.
- TYPE DETAILS**: Section containing:
 - USERNAME**: Text input field.
 - PASSWORD**: Text input field with a 'SHOW' button and an 'Ask at runtime?' checkbox.
 - PRIVATE KEY PASSPHRASE**: Text input field with a 'SHOW' button and an 'Ask at runtime?' checkbox.
 - PRIVILEGE ESCALATION**: Dropdown menu set to 'Choose a privilege escalation'.
 - VAULT PASSWORD**: Text input field with a 'SHOW' button and an 'Ask at runtime?' checkbox.
- PRIVATE KEY**: Large text area for pasting the private key.

At the bottom right, there are 'CANCEL' and 'SAVE' buttons.

Machine credentials have several attributes that may be configured:

- **Username:** The username to be used for SSH authentication.
- **Password:** The actual password to be used for SSH authentication. This password can be stored encrypted in the Tower database, if entered. Alternatively, you can configure Tower to ask the user for the password when necessary by selecting “Ask at runtime?”. In these cases, a dialog opens when the job is launched, promoting the user to enter the password and password confirmation.
- **Private Key:** The actual SSH Private Key to be used to authenticate the user via SSH. This key is stored encrypted in the Tower database.
- **Private Key Passphrase:** If the SSH Private Key used is protected by a password, you can configure a Key Password for the private key. This password may be stored encrypted in the Tower database, if entered. Alternatively, you can configure Tower to ask the user for the password as necessary by selecting “Ask at runtime?”. In these cases, a dialog opens when the job is launched, prompting the user to enter the password and password confirmation.
- **Privilege Escalation:** Specifies the type of escalation privilege to assign to specific users. This is equivalent to specifying the `--become-method=BECOME_METHOD` parameter, where `BECOME_METHOD` could be `sudo` | `su` | `pbrun` | `pfexec`.
 - **none:** Assigns no privilege escalation to this credential.

- **sudo**: Performs single commands with super user (root user) privileges
- **su**: Switches to the super user (root user) account (or to other user accounts)
- **pbrun**: Requests that an application or command be run in a controlled account and provides for advanced root privilege delegation and keylogging.
- **pfexec**: Executes commands with predefined process attributes, such as specific user or group IDs.

PRIVILEGE ESCALATION ⓘ	PRIVILEGE ESCALATION USERNAME	PRIVILEGE ESCALATION PASSWORD
<input type="text" value="Sudo"/>	<input type="text"/>	<input type="text" value="SHOW"/>
		<input type="checkbox"/> Ask at runtime?

- **Privilege Escalation Username**: The username to use with escalation privileges on the remote system.
- **Privilege Escalation Password**: The actual password to be used to authenticate the user via the selected privilege escalation type on the remote system. This password may be stored encrypted in the Tower database, if entered. Alternatively, you may configure Tower to ask the user for the password when necessary by selecting “Ask at runtime?”. In these cases, a dialog opens when the job is launched, promoting the user to enter the password and password confirmation.

Note: Sudo Password must be used in combination with SSH passwords or SSH Private Keys, since Tower must first establish an authenticated SSH connection with the host prior to invoking sudo to change to the sudo user.

- **Vault Password**: If your playbook uses Ansible Vault, add the Vault password to your credentials here. Alternatively, you may configure Tower to ask the user for the vault password when necessary by selecting “Ask at runtime?”. In these cases, a dialog opens when the job is launched, promoting the user to enter the password and password confirmation.

For more information about Ansible Vault, refer to: http://docs.ansible.com/playbooks_vault.html

Warning: Credentials which are used in *Scheduled Jobs* must not be configured as “Ask at runtime?”.

9.4.2 Network

Network credentials are used by Ansible networking modules to connect to and manage networking devices.

CREATE CREDENTIAL

DETAILS PERMISSIONS

*NAME DESCRIPTION ORGANIZATION ?

*TYPE ?

Network

TYPE DETAILS

*USERNAME *PASSWORD

SHOW

☐ AUTHORIZE

SSH KEY ?

Not signed in

CANCEL SAVE

Network credentials have several attributes that may be configured:

- **Username:** The username to use in conjunction with the network device.
- **Password:** The password to use in conjunction with the network device.
- **SSH Private Key:** The actual SSH Private Key to be used to authenticate the user to the network via SSH.
- **Authorize:** Select this to add an **Authorize Password** which signs the RSA key with a password (the **Authorize Password** field is only seen if this option is selected).

9.4.3 Source Control

SCM (source control) credentials are used with Projects to clone and update local source code repositories from a remote revision control system such as Git, Subversion, or Mercurial.

CREATE CREDENTIAL

DETAILS | PERMISSIONS

*NAME DESCRIPTION ORGANIZATION

*TYPE

TYPE DETAILS

USERNAME PASSWORD PRIVATE KEY PASSPHRASE

SCM PRIVATE KEY

CANCEL SAVE

Source Control credentials have several attributes that may be configured:

- **Username:** The username to use in conjunction with the source control system.
- **Password:** The password to use in conjunction with the source control system.
- **SCM Private Key:** The actual SSH Private Key to be used to authenticate the user to the source control system via SSH.
- **Private Key Passphrase:** If the SSH Private Key used is protected by a passphrase, you may configure a Key Passphrase for the private key.

Note: Source Control credentials cannot be configured as “Ask at runtime?”.

9.4.4 Amazon Web Services

Selecting this credential type enables synchronization of cloud inventory with Amazon Web Services.

Traditional Amazon Web Services credentials consist of the AWS **Access Key** and **Secret Key**.

Ansible Tower version 2.4.0 introduced support for EC2 STS tokens (sometimes referred to as IAM STS credentials). Security Token Service (STS) is a web service that enables you to request temporary, limited-privilege credentials for AWS Identity and Access Management (IAM) users. To learn more about the IAM/EC2 STS Token, refer to: http://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_temp.html

AWS credentials consist of:

```
AWS_ACCESS_KEY
AWS_SECRET_KEY
AWS_SECURITY_TOKEN
```

Note: If the value of your tags in EC2 contain booleans (yes/no/true/false), you must remember to quote them.

Warning: To use implicit IAM role credentials, do not attach AWS cloud credentials in Tower when relying on IAM roles to access the AWS API. While it may seem to make sense to attach your AWS cloud credential to your job template, doing so will force the use of your AWS credentials and will not “fall through” to use your IAM role credentials (this is due to the use of the boto library.)

9.4.5 Rackspace

Selecting this credential type enables synchronization of cloud inventory with Rackspace.

CREATE CREDENTIAL

DETAILS | PERMISSIONS

*NAME DESCRIPTION ORGANIZATION

*TYPE

TYPE DETAILS

*USERNAME *API KEY

CANCEL SAVE

Rackspace credentials consist of the Rackspace **Username** and **API Key**.

9.4.6 VMware vCenter

Selecting this credential type enables synchronization of inventory with VMware vCenter.

CREATE CREDENTIAL

DETAILS | PERMISSIONS

*NAME DESCRIPTION ORGANIZATION

*TYPE

TYPE DETAILS

*VCENTER HOST *USERNAME *PASSWORD

CANCEL SAVE

VMware credentials have several attributes that may be configured:

- **vCenter Host:** The vCenter hostname or IP address to connect to.
- **Username:** The username to use to connect to vCenter.
- **Password:** The password to use to connect to vCenter.

Note: If the VMware guest tools are not running on the instance, VMware inventory sync may not return an IP address for that instance.

9.4.7 Red Hat Satellite 6

Selecting this credential type enables synchronization of cloud inventory with Red Hat Satellite 6.

Satellite credentials have several attributes that may be configured:

- **Satellite 6 Host:** The Satellite 6 hostname or IP address to connect to.
- **Username:** The username to use to connect to Satellite 6.
- **Password:** The password to use to connect to Satellite 6.

9.4.8 Red Hat CloudForms

Selecting this credential type enables synchronization of cloud inventory with Red Hat CloudForms.

CloudForms have several attributes that may be configured:

- **CloudForms Host:** The CloudForms hostname or IP address to connect to.
- **Username:** The username to use to connect to CloudForms.
- **Password:** The password to use to connect to CloudForms.

Please note that during the testing of Ansible Tower 3.0, it was discovered that the CloudForms inventory plugin has a known failure with the current release of the CloudForms Management Engine (CFME), version 5.6. This failure is expected to be resolved with the release of CFME 5.6.1.

Additional Resources:

Red Hat is writing a blog post series on Ansible Tower Integration in Red Hat CloudForm 4.1 which can be accessed at <http://cloudformsblog.redhat.com/2016/07/22/ansible-tower-in-cloudforms/>.

9.4.9 Google Compute Engine

Selecting this credential type enables synchronization of cloud inventory with Google Compute Engine.

The screenshot displays the Ansible Tower web interface. At the top, there's a navigation bar with 'TOWER', 'PROJECTS', 'INVENTORIES', 'JOB TEMPLATES', and 'JOBS'. A user 'admin' is logged in. The breadcrumb trail shows 'SETTINGS / CREDENTIALS / CREATE CREDENTIAL'.

The main form is titled 'CREATE CREDENTIAL'. It has two tabs: 'DETAILS' (active) and 'PERMISSIONS'. The form fields are as follows:

- * NAME**: A text input field.
- DESCRIPTION**: A text input field.
- ORGANIZATION**: A text input field with a search icon.
- * TYPE**: A dropdown menu with 'Google Compute Engine' selected.
- TYPE DETAILS**: A section containing:
 - * SERVICE ACCOUNT EMAIL ADDRESS**: A text input field.
 - PROJECT**: A text input field.
 - * RSA PRIVATE KEY**: A large text area for pasting the private key.

At the bottom right of the form are 'CANCEL' and 'SAVE' buttons.

Below the form is a section titled 'CREDENTIALS' with a '+ ADD' button. It contains a search bar and a table of existing credentials:

NAME	DESCRIPTION	TYPE	OWNERS	ACTIONS
Demo Credential		Machine	admin	

Copyright © 2016 Red Hat, Inc.

Google Compute Engine credentials have several attributes that may be configured:

- **Service Account Email Address:** The email address assigned to the Google Compute Engine **service account**.
- **RSA Private Key:** The PEM file associated with the service account email.
- **Project:** The GCE assigned identification. It is constructed as two words followed by a three digit number, such as: squeamish-ossifrage-123.

9.4.10 Microsoft Azure Classic (Deprecated)

Selecting this credential type enables synchronization of cloud inventory with Windows Azure Classic.

CREATE CREDENTIAL

DETAILS | PERMISSIONS

* NAME DESCRIPTION ORGANIZATION

* TYPE

TYPE DETAILS

* SUBSCRIPTION ID

* MANAGEMENT CERTIFICATE

CREDENTIALS 1

NAME SEARCH

NAME	DESCRIPTION	TYPE	OWNERS	ACTIONS
Demo Credential		Machine	admin	<input type="button" value="edit"/> <input type="button" value="delete"/>

Copyright © 2016 Red Hat, Inc.

Microsoft Azure credentials have several attributes to configure:

- **Subscription ID:** The Subscription UUID for the Microsoft Azure Classic account.
- **Management Certificate:** The PEM file that corresponds to the certificate you uploaded in the Microsoft Azure Classic console.

9.4.11 Microsoft Azure Resource Manager

Selecting this credential type enables synchronization of cloud inventory with Microsoft Azure Resource Manager.

CREATE CREDENTIAL

DETAILS | PERMISSIONS

* NAME: DESCRIPTION: ORGANIZATION:

* TYPE:

TYPE DETAILS

* SUBSCRIPTION ID: USERNAME: PASSWORD:

CLIENT ID: CLIENT SECRET: TENANT ID:

CREDENTIALS 1

NAME SEARCH

NAME	DESCRIPTION	TYPE	OWNERS	ACTIONS
Demo Credential		Machine	admin	<input type="button" value="edit"/> <input type="button" value="delete"/>

Copyright © 2016 Red Hat, Inc.

Microsoft Azure Resource Manager credentials have several attributes to configure:

- **Subscription ID:** The Subscription UUID for the Microsoft Azure account.
- **Username:** The username to use to connect to the Microsoft Azure account.
- **Password:** The password to use to connect to the Microsoft Azure account.
- **Client ID:** The Client ID for the Microsoft Azure account.
- **Client Secret:** The Client Secret for the Microsoft Azure account.
- **Tenant ID:** The Tenant ID for the Microsoft Azure account.

To pass service principal credentials, define the following variables:

```
AZURE_CLIENT_ID
AZURE_SECRET
AZURE_SUBSCRIPTION_ID
AZURE_TENANT
```

To pass an Active Directory username/password pair, define the following variables:

```
AZURE_AD_USER
AZURE_PASSWORD
AZURE_SUBSCRIPTION_ID
```

You can also pass credentials as parameters to a task within a playbook. The order of precedence is parameters, then environment variables, and finally a file found in your home directory.

To pass credentials as parameters to a task, use the following parameters for service principal credentials:

```
client_id
secret
subscription_id
tenant
```

Or, pass the following parameters for Active Directory username/password:

```
ad_user
password
subscription_id
```

9.4.12 OpenStack

Selecting this credential type enables synchronization of cloud inventory with OpenStack.

The screenshot displays the 'CREATE CREDENTIAL' form in the Ansible Tower web interface. The form is titled 'CREATE CREDENTIAL' and has two tabs: 'DETAILS' (selected) and 'PERMISSIONS'. The form fields are as follows:

- *NAME**: A text input field.
- DESCRIPTION**: A text input field.
- ORGANIZATION**: A text input field with a search icon.
- *TYPE**: A dropdown menu with 'OpenStack' selected.
- TYPE DETAILS**: A section containing four fields:
 - *HOST (AUTHENTICATION URL)**: A text input field.
 - *USERNAME**: A text input field.
 - *PASSWORD (API KEY)**: A text input field with a 'SHOW' button.
 - *PROJECT (TENANT NAME)**: A text input field.
 - DOMAIN NAME**: A text input field.

At the bottom right of the form are 'CANCEL' and 'SAVE' buttons. Below the form is a section titled 'CREDENTIALS' with a '+ ADD' button. It contains a search bar and a table of existing credentials:

NAME	DESCRIPTION	TYPE	OWNERS	ACTIONS
Demo Credential		Machine	admin	

Copyright © 2016 Red Hat, Inc.

OpenStack credentials have several attributes that may be configured:

- **Host (Authentication URL)**: The host to be used for authentication.
- **Username**: The username to use to connect to OpenStack.

- **Password (API Key):** The password or API key to use to connect to OpenStack.
- **Project (Tenet Name/ID):** The Tenant name or Tenant ID used for OpenStack. This value is usually the same as the username.
- **Domain name:** The FQDN to be used to connect to OpenStack.

If you are interested in using OpenStack Cloud Credentials, refer to [Utilizing Cloud Credentials](#) in this guide for more information, including a sample playbook.

PROJECTS

A **Project** is a logical collection of Ansible playbooks, represented in Tower.

You can manage playbooks and playbook directories by either placing them manually under the Project Base Path on your Tower server, or by placing your playbooks into a source code management (SCM) system supported by Tower, including Git, Subversion, and Mercurial.

Note: By default, the Project Base Path is `/var/lib/awx/projects`, but this may have been modified by the Tower administrator. It is configured in `/etc/tower/settings.py`. Use caution when editing this file, as incorrect settings can disable your installation.

This menu displays a list of the projects that are currently available. The list of projects may be sorted and searched by **Status**, **Name**, or **Type**. For each project listed, you can edit project properties and delete the project, using the edit and delete icons.

The screenshot shows the Tower web interface. The top navigation bar includes 'TOWER' and 'PROJECTS' (highlighted), with other options like 'INVENTORIES', 'JOB TEMPLATES', and 'JOBS'. On the right side of the navigation bar, there is a user profile 'admin', a settings gear, a list icon, a document icon, and a power icon. Below the navigation bar, the 'PROJECTS' section is active. It displays a table with one project: 'Demo Project'. The table has columns for 'NAME', 'TYPE', 'LAST UPDATED', and 'ACTIONS'. The 'ACTIONS' column contains icons for edit and delete. A '+ ADD' button is located in the top right corner of the projects list.

Status indicates the state of the project and may be one of the following (note that you can also filter your view by specific status types):

- **New:** The source control update has been created, but not queued or started yet. (To be deprecated.)
- **Pending:** The source control update has been queued, but has not run yet. (To be deprecated.)
- **Waiting:** The source control update is waiting on an update/dependency.
- **Running:** The source control update is currently in progress.
- **Successful:** The last source control update for this project succeeded.
- **Failed:** The last source control update for this project failed.
- **Error:** The last source control update job failed to run at all. (To be deprecated.)
- **Canceled:** The last source control update for the project was canceled.
- **Never updated:** The project is configured for source control, but has never been updated.

- **OK:** The project is not configured for source control, and is correctly in place. (To be deprecated.)
- **Missing:** Projects are absent from the project base path of `/var/lib/awx/projects` (applicable for manual or source control managed projects).

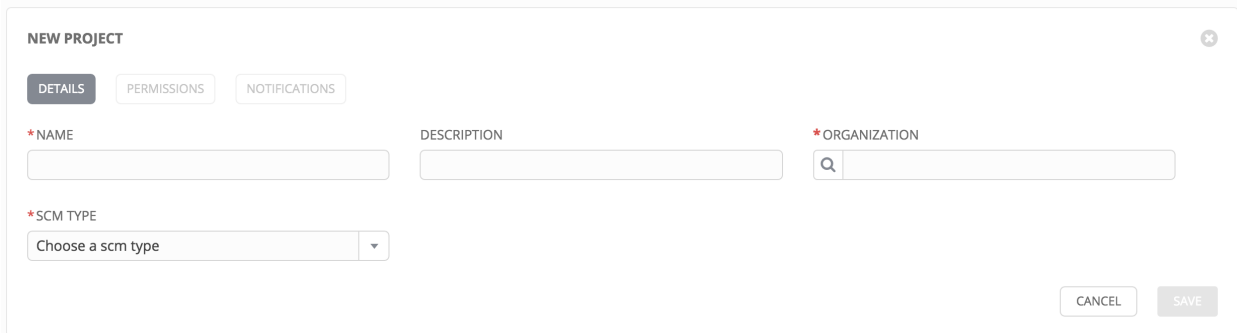
Under **Actions**, the following actions are available:

- **Update:** Invoke an immediate update from source control, if configured for this project
- **Schedule:** Schedule an update from source control, if configured for this project
- **Edit:** Edit the project
- **Delete:** Delete the project

Note: Projects of credential type Manual cannot update or schedule source control-based actions without being reconfigured as an SCM type credential.

10.1 Add a new project

To create a new project, click the  button, which launches the **Create Project** dialog.



Enter the appropriate details into the following fields:

- **Name**
- **Description**
- **Organization** (A project must have at least one organization. Pick one organization now to create the project, and then after the project is created you can add additional organizations.)
- **SCM Type** (Select one of Manual, Git, Subversion, or Mercurial. Refer to [Manage playbooks manually](#) and [Manage playbooks using Source Control](#) in this guide for more detail.)

All fields are required.

Note: Each project path can only be assigned to one project. If you receive the following message, ensure that you have not already assigned the project path to an existing project:

“All of the project paths have been assigned to existing projects, or there are no directories found in the base path. You will need to add a project path before creating a new project.”

10.1.1 Manage playbooks manually

- Create one or more directories to store playbooks under the Project Base Path (for example, /var/lib/awx/projects/)
- Create or copy playbook files into the playbook directory.
- Ensure that the playbook directory and files are owned by the same UNIX user and group that the Tower service runs as.
- Ensure that the permissions are appropriate for the playbook directories and files.

If you have trouble adding a project path, check the permissions and SELinux context settings for the project directory and files.

Warning: If you have not added any Ansible playbook directories to the base project path, you will receive the following message from Tower:

NEW PROJECT

DETAILS PERMISSIONS NOTIFICATIONS

* NAME: Example DESCRIPTION: Ansible example playbook * ORGANIZATION: Honey Dog, Inc

* SCM TYPE: Manual

WARNING: There are no available playbook directories in /var/lib/awx/projects. Either that directory is empty, or all of the contents are already assigned to other projects. Create a new directory there and make sure the playbook files can be read by the "awx" system user, or have Tower directly retrieve your playbooks from source control using the SCM Type option above.

PROJECT BASE PATH: /var/lib/awx/projects

CANCEL SAVE

Correct this issue by creating the appropriate playbook directories and checking out playbooks from your SCM or otherwise copying playbooks into the appropriate playbook directories.

10.1.2 Manage playbooks using Source Control

Select the appropriate **SCM Type**. Then, enter the appropriate details into the following fields:

- SCM URL
- SCM Branch (Optionally enter the SCM branch for Git or Mercurial.)
- Revision # (Optionally enter the Revision # for Subversion.)
- SCM Credential (If authentication is required, select the appropriate SCM credential.)
- SCM Update Options
 - Clean (Remove any local modifications prior to performing an update.)
 - Delete on Update (Delete the local repository in its entirety prior to performing an update. Depending on the size of the repository this may significantly increase the amount of time required to complete an update.)

- Update on Launch (Each time a job runs using this project, perform an update to the local repository prior to starting the job. To avoid job overflows if jobs are spawned faster than the project can sync, selecting this allows you to configure a Cache Timeout to cache prior project syncs for a certain number of seconds.)

NEW PROJECT

DETAILS
PERMISSIONS
NOTIFICATIONS

*NAME
DESCRIPTION
*ORGANIZATION

Example
Ansible example playbook
Honey Dog, Inc

*SCM TYPE

Git

SOURCE DETAILS

*SCM URL
SCM BRANCH
SCM CREDENTIAL

http://github.com/ansible/ansible-examples.git

SCM UPDATE OPTIONS

☐ Clean
☐ Delete on Update
☐ Update on Launch


CANCEL
SAVE

Click **Save** to save your project.

Tip: Using a Github link offers an easy way to use a playbook. To help get you started, use the `helloworld.yml` file available at: <https://github.com/ansible/tower-example.git>

This link offers a very similar playbook to the one created manually in the instructions found in the [Ansible Tower Quick Start Guide](#). Using it will not alter or harm your system in anyway.











10.2 Updating projects from source control

Update an existing SCM-based project by clicking the  button.

Note: Please note that immediately after adding a project setup to use source control, a “Sync” starts that fetches the project details from the configured source control.

PROJECTS 2
+ ADD

STATUS FILTER


NAME	TYPE	LAST UPDATED	ACTIONS
 Demo Project	Git	6/22/2016 11:22:38 AM	   
 Example	Git	6/22/2016 11:33:36 AM	   

Click on the dot under **Status** (far left, beside the name of the Project) to get further details about the update process.

The screenshot shows the Ansible Tower web interface. At the top, there are navigation tabs: TOWER, PROJECTS, INVENTORIES, JOB TEMPLATES, and JOBS. The 'JOBS' tab is selected, and the breadcrumb 'JOBS / EXAMPLE' is visible. The main content area is divided into two panels. The left panel, titled 'RESULTS', shows job details: NAME (Example), STATUS (Successful), STARTED (6/22/2016 11:33:26 AM), FINISHED (6/22/2016 11:33:36 AM), ELAPSED (9.772 seconds), LAUNCH TYPE (Manual), and PROJECT (Example). The right panel, titled 'STANDARD OUT', displays the command output. It starts with 'Using /etc/ansible/ansible.cfg as config file' and 'PLAY [all]'. It then shows several tasks, all of which were skipped because of a 'Conditional check failed'. The tasks are: 'delete project directory before update', 'update project using git and accept hostkey', 'update project using git', 'update project using hg', 'update project using svn', and 'update project using svn with auth'.

Copyright © 2016 Red Hat, Inc.



To set a schedule for updating the project from SCM, click the  button. This will navigate to the **Schedules** screen.

The screenshot shows the 'SCHEDULES' screen for a project named 'EXAMPLE'. At the top, there is a header 'EXAMPLE | SCHEDULES' with a count of '0' in a grey box. On the right side, there is a green '+ ADD' button. The main area of the screen is a large, empty light grey rectangle with the text 'PLEASE ADD ITEMS TO THIS LIST' centered in the middle.

This screen displays a list of the schedules that are currently available for the selected **Project**. The schedule list may be sorted and searched by **Name**.

The list of schedules includes:

- **Name:** Clicking the schedule name opens the **Edit Schedule** dialog
- **First Run:** The first scheduled run of this task
- **Next Run:** The next scheduled run of this task
- **Final Run:** If the task has an end date, this is the last scheduled run of the task.

Buttons located in the upper right corner of the **Schedules** screen provide the following actions:

- Create a new schedule
- Refresh this view
- View Activity Stream

10.3 Add a new schedule

To create a new schedule click the  button, which opens the **Add Schedule** dialog.

PROJECTS / EXAMPLE / SCHEDULES / CREATE SCHEDULE

* NAME

* START DATE (MM/DD/YYYY)

* START TIME (HH24:MM:SS)

 : :

* LOCAL TIME ZONE

* REPEAT FREQUENCY

SCHEDULE DESCRIPTION

every day for 1 time
OCCURRENCES (Limited to first 10) DATE FORMAT ☒ LOCAL TIME ☐ UTC
06/22/2016 12:00:00 EDT

EXTRA VARIABLES ☒ YAML ☐ JSON

1

CANCEL

SAVE

Enter the appropriate details into the following fields:

- Name (required)
- Start Date (required)
- Start Time (required)
- Local Time Zone (the entered Start Time should be in this timezone)
- UTC Start Time (calculated from Start Time + Local Time Zone)
- Repeat Frequency (appropriate scheduling options are displayed depending on the frequency you select)

The **SCHEDULE DESCRIPTION** allows you to review the set schedule and a list of the scheduled occurrences in the selected Local Time Zone.

Caution: Jobs are scheduled in UTC. Repeating jobs that run at a specific time of day may move relative to a local timezone when Daylight Savings Time shifts occur. Essentially, Tower resolves the local time zone based time to UTC when the schedule is saved. To ensure your schedules are correctly set, you should set your schedules in UTC time.

Once done, click **Save**.

You can use the **ON/OFF** toggle button to stop an active schedule or activate a stopped schedule.

The schedules overview screen for the project also shows you when the first, next, and final runs are scheduled.

There are several actions available for schedules, under the **Actions** column:

- Edit Schedule
- Delete schedule

EXAMPLE | SCHEDULES 1

+ ADD

NAME SEARCH

Q

NAME ^	FIRST RUN ↕	NEXT RUN ↕	FINAL RUN ↕	ACTIONS
ON Test 1	6/22/2016 12:02:00 PM	6/22/2016 12:02:00 PM	6/22/2016 12:02:00 PM	

10.4 Ansible Galaxy Support

At the end of a Project update, Tower searches for a file called `roles/requirements.yml` in the top level of the Project directory. If this file is found, the following command automatically runs:

```
ansible-galaxy install -r requirements.yml -p ./roles/ --force
```

This file allows you to reference Galaxy roles or roles within other repositories which can be checked out in conjunction with your own project. The addition of this Ansible Galaxy support eliminates the need to create git submodules for achieving this result.


For more information and examples on the syntax of the `requirements.yml` file, refer to [Advanced Control Over Role Requirements](#) in the Ansible documentation.

INVENTORIES

An **Inventory** is a collection of hosts against which jobs may be launched, the same as an Ansible inventory file. Inventories are divided into groups and these groups contain the actual hosts. Groups may be sourced manually, by entering host names into Tower, or from one of Ansible Tower’s supported cloud providers.

Note: If you have a custom dynamic inventory script, or a cloud provider that is not yet supported natively in Tower, you can also import that into Tower. Refer to the [Tower Administration Guide](#).

This tab displays a list of the inventories that are currently available. The inventory list may be sorted and searched by **Name** or **Organization**. This list may also be filtered by selecting **Cloud Sourced**, **Failed Hosts**, or **Sync Failures**.



The list of Inventory details includes:

- **Inventory Sync:** If configured, you can “Sync” to fetch the project details from the configured source control.
- **Status Dot:** This shows the status of recent jobs for this inventory.
- **Name:** The inventory name. Clicking the Inventory name navigates to the properties screen for the selected inventory, which shows the inventory’s groups and hosts. (This view is also accessible from the **Action** menu.)
- **Organization:** The organization to which the inventory belongs.
- **Actions:** The following actions are available for the selected inventory:
 - **Edit:** Edit the properties for the selected inventory
 - **Delete:** Delete the selected inventory. *This operation cannot be reversed!*

11.1 Add a new inventory

To create a new inventory click the  button, which opens the **Create Inventory** window.

Enter the appropriate details into the following fields and select **Save**:

- **Name:** Enter a name appropriate for this inventory.
- **Description:** Enter an arbitrary description as appropriate.
- **Organization:** Choose among the available organizations.
- **Variables:** Variable definitions and values to be applied to all hosts in this inventory. Enter variables using either JSON or YAML syntax. Use the radio button to toggle between the two.

After clicking save, the Groups and Hosts Management screen appears.

[INVENTORIES](#) / DATABASE SERVERS

11.2 Groups and Hosts

Topics:

- *Groups*
 - *Add a new group*
 - *Credential Sources*
 - * *Rackspace Cloud Servers*
 - * *Amazon Web Services EC2*
 - * *Google Compute Engine*
 - * *Microsoft Azure Classic (deprecated)*
 - * *Microsoft Azure Resource Manager*
 - * *VMware vCenter*
 - * *Red Hat Satellite 6*
 - * *Red Hat CloudForms*
 - * *OpenStack*
 - * *Custom Script*
 - *Scheduling*
 - * *Add a new schedule*
- *Hosts*
 - *Add a new host*

Inventories are divided into groups, which may contain hosts and other groups, and hosts. There are several actions available for inventories.

- Create a new Group
- Create a new Host
- Run a command on the selected Inventory
- Edit Inventory properties
- View activity streams for Groups and Hosts
- Obtain help building your Inventory

11.2.1 Groups

Under Groups, you can view which groups belong to this inventory, easily filtered or searched by group name.

Additional actions may be performed on the group by selecting the buttons to the right of the group name:

- **Sync status:** Show the status of inventory synchronization for groups configured with cloud sources. If synchronization is configured, clicking this button shows the synchronization log for the selected group.
- **Host status:** Show the status of successful and failed jobs for the selected group. Clicking this button shows the list of hosts that are members of the selected group.

- **Start sync process:** Initiate a synchronization of the group with the configured cloud source. (A synchronization process that is in progress may be canceled by clicking the cancel button that appears here during synchronization.)
- **Edit Group:** Edit the properties for the selected group
- **Copy Group:** Groups can be nested. This allows you to copy or move the group to a different group.
- **Delete:** Delete the selected group. *This operation cannot be reversed!*

Add a new group

+ ADD GROUP

Create a new group for an inventory by clicking the **+ ADD GROUP** button, which opens the **Create Group** window.

Enter the appropriate details into the required and optional fields, and click **Save**.

- **Name:** Required
- **Description:** Enter an arbitrary description as appropriate
- **Source:** Choose a source which matches the credential type against which a host can be entered.
- **Variables:** Enter definitions and values to be applied to all hosts in this group. Enter variables using either JSON or YAML syntax. Use the radio button to toggle between the two.

In prior versions of Ansible Tower, the **Source** default was manual, meaning that the hosts must be entered into Tower manually. Beginning with Ansible Tower 3.0, host can be added via multiple methods/credential sources. (Refer to [Add a new host](#) for more information on managing hosts individually.)

To synchronize the inventory group from a cloud source, choose the appropriate source from the **Source** menu.

Note: Starting with version 2.2, Ansible Tower supports Amazon Web Services EC2, Rackspace Cloud Servers, Google Compute Engine, VMware vCenter, Microsoft Azure, OpenStack, and custom scripts added by the administrator. With Ansible Tower version 3.0, Microsoft Azure Classic (deprecated) and Microsoft Azure Resource Manager

were added to expand upon the support offered for Microsoft Azure, and support for Red Hat Satellite 6 and RH Cloudforms credentials were also added.

All cloud inventory sources have the following update options:

- **Overwrite:** When checked all child groups and hosts not found on the remote source are deleted from the local inventory. When not checked any local child hosts and groups not found on the external source remains untouched by the inventory update process.
- **Overwrite Variables:** If checked, all variables for child groups and hosts will be removed and replaced by those found on the external source. When not checked a merge is performed, combining local variables with those found on the external source.
- **Update on Launch:** Each time a job runs using this inventory, refresh the inventory from the selected source before executing job tasks. To avoid job overflows if jobs are spawned faster than the inventory can sync, selecting this allows you to configure a Cache Timeout to cache prior inventory syncs for a certain number of seconds.

The “Update on Launch” setting refers to a dependency system for projects and inventory, and it will not specifically exclude two jobs from running at the same time. If a cache timeout is specified, then the dependencies for the second job is created and it uses the project and inventory update that the first job spawned. Both jobs then wait for that project and/or inventory update to finish before proceeding. If they are different job templates, they can then both start and run at the same time, if the system has the capacity to do so.

Note: If you intend to use Tower’s provisioning callback feature with a dynamic inventory source, “Update on Launch” should be set for the inventory group.

When done, click **Save**.

Credential Sources

Topics:

- *Rackspace Cloud Servers*
- *Amazon Web Services EC2*
- *Google Compute Engine*
- *Microsoft Azure Classic (deprecated)*
- *Microsoft Azure Resource Manager*

- *VMware vCenter*
- *Red Hat Satellite 6*
- *Red Hat CloudForms*
- *OpenStack*
- *Custom Script*

Choose a source which matches the credential type against which a host can be entered.

Rackspace Cloud Servers

To configure a group for Rackspace, select **Rackspace Cloud Servers** and enter the following details:

- **Cloud Credential:** Choose from an existing Credential. For more information, refer to [Credentials](#).
- **Regions:** Click on the regions field to see a list of regions for your cloud provider. You can select multiple regions, or choose “All” to include all regions. Tower will only be updated with Hosts associated with the selected regions.
- **Variables:** Enter definitions and values to be applied to all hosts in this group. Enter variables using either JSON or YAML syntax. Use the radio button to toggle between the two.

You can also configure **Update Options**.

- **Overwrite:** If checked, all child groups and hosts not found on the external source are deleted from the local inventory. When not checked, local child hosts and groups not found on the external source remain untouched by the inventory update process.
- **Overwrite Variables:** If checked, all variables for child groups and hosts are removed and replaced by those found on the external source. When not checked, a merge is performed, combining local variables with those found on the external source.
- **Update on Launch:** Each time a job runs using this inventory, refresh the inventory from the selected source before executing job tasks.

CREATE GROUP

DETAILS

NOTIFICATIONS

* NAME

DESCRIPTION

SOURCE

Rackspace Cloud Servers

* CLOUD CREDENTIAL

REGIONS ?

UPDATE OPTIONS

☐ Overwrite ?
☐ Overwrite Variables ?
☐ Update on Launch ?

VARIABLES ?

☒ YAML
 ☐ JSON

1

CANCEL

SAVE

Amazon Web Services EC2

To configure a group for AWS, select **Amazon EC2** and enter the following details:

CREATE GROUP

DETAILS

NOTIFICATIONS

* NAME

CMS Web Group - West

DESCRIPTION

CMS Web West Coast

SOURCE

Amazon EC2

* CLOUD CREDENTIAL

REGIONS ?

INSTANCE FILTERS ?

ONLY GROUP BY ?

UPDATE OPTIONS

☐ Overwrite ?
☐ Overwrite Variables ?
☐ Update on Launch ?

VARIABLES ?

☒ YAML
 ☐ JSON

1

SOURCE VARIABLES ?

☒ YAML
 ☐ JSON

1

CANCEL

SAVE

- **Cloud Credential:** Choose from an existing credential (for more information, refer to [Credentials](#)).

If Tower is running on an EC2 instance with an assigned IAM Role, the credential may be omitted, and the security credentials from the instance metadata will be used instead. For more information on using IAM Roles, refer to the [IAM_Roles_for_Amazon_EC2_documentation_at_Amazon](#).

- **Regions:** Click on the regions field to see a list of regions for your cloud provider. You can select multiple regions, or choose “All” to include all regions. Tower will only be updated with Hosts associated with the selected regions.
- **Instance Filters:** Rather than importing your entire Amazon EC2 inventory, filter the instances returned by the inventory script based on a variety of metadata. Hosts are imported if they match any of the filters entered here.

Examples:

- To limit to hosts having the tag TowerManaged: Enter `tag-key=TowerManaged`
- To limit to hosts using either the key-name staging or production: Enter `key-name=staging, key-name=production`
- To limit to hosts where the Name tag begins with test: Enter `tag:Name=test*`

For more information on the filters that can be used here, refer to the [Describe Instances](#) documentation at Amazon.

- **Only Group By:** By default, Tower creates groups based on the following Amazon EC2 parameters:

- Availability Zones
- Image ID
- Instance Type
- Key Name
- Region
- Security Group
- Tags (by name)
- VPC ID

If you do not want all these groups created, select from the dropdown the list of groups that you would like created by default. You can also select `Instance ID` to create groups based on the Instance ID of your instances.

- **Variables:** Enter definitions and values to be applied to all hosts in this group. Enter variables using either JSON or YAML syntax. Use the radio button to toggle between the two.
- **Source Variables:** Override variables found in `ec2.ini` and used by the inventory update script. Enter variables using either JSON or YAML syntax. Use the radio button to toggle between the two. For a detailed description of these variables [view ec2.ini in the Ansible GitHub repo](#).

You can also configure **Update Options**.

- **Overwrite:** If checked, all child groups and hosts not found on the external source are deleted from the local inventory. When not checked, local child hosts and groups not found on the external source remain untouched by the inventory update process.
- **Overwrite Variables:** If checked, all variables for child groups and hosts are removed and replaced by those found on the external source. When not checked, a merge is performed, combining local variables with those found on the external source.
- **Update on Launch:** Each time a job runs using this inventory, refresh the inventory from the selected source before executing job tasks.

Google Compute Engine

To configure a group for Google Compute Engine, select **Google Compute Engine** and enter the following details:

- **Cloud Credential:** Choose from an existing Credential. For more information, refer to [Credentials](#).
- **Regions:** Click on the regions field to see a list of regions for your cloud provider. You can select multiple regions, or choose “All” to include all regions. Tower will only be updated with Hosts associated with the selected regions.
- **Variables:** Enter definitions and values to be applied to all hosts in this group. Enter variables using either JSON or YAML syntax. Use the radio button to toggle between the two.

You can also configure **Update Options**.

- **Overwrite:** If checked, all child groups and hosts not found on the external source are deleted from the local inventory. When not checked, local child hosts and groups not found on the external source remain untouched by the inventory update process.
- **Overwrite Variables:** If checked, all variables for child groups and hosts are removed and replaced by those found on the external source. When not checked, a merge is performed, combining local variables with those found on the external source.
- **Update on Launch:** Each time a job runs using this inventory, refresh the inventory from the selected source before executing job tasks.

The screenshot shows the 'CREATE GROUP' form in Ansible Tower. The form has two tabs: 'DETAILS' (selected) and 'NOTIFICATIONS'. The 'DETAILS' tab contains the following fields:

- * NAME:** A text input field containing 'CMS Web Group - West'.
- DESCRIPTION:** A text input field containing 'CMS Web West Coast'.
- SOURCE:** A dropdown menu with 'Google Compute Engine' selected.
- * CLOUD CREDENTIAL:** A search input field with a magnifying glass icon.
- REGIONS:** A text input field with a question mark icon.
- UPDATE OPTIONS:** Three checkboxes: 'Overwrite' (unchecked), 'Overwrite Variables' (unchecked), and 'Update on Launch' (unchecked).
- VARIABLES:** A section with a question mark icon, radio buttons for 'YAML' (selected) and 'JSON', and a large text area for entering variables.

At the bottom right of the form are 'CANCEL' and 'SAVE' buttons.

Microsoft Azure Classic (deprecated)

To configure a group for Microsoft Azure Classic, select **Microsoft Azure Classic (deprecated)** and enter the following details:

- **Cloud Credential:** Choose from an existing Credential. For more information, refer to [Credentials](#).
- **Regions:** Click on the regions field to see a list of regions for your cloud provider. You can select multiple regions, or choose “All” to include all regions. Tower will only be updated with Hosts associated with the selected regions.
- **Variables:** Enter definitions and values to be applied to all hosts in this group. Enter variables using either JSON or YAML syntax. Use the radio button to toggle between the two.

You can also configure **Update Options**.

- **Overwrite:** If checked, all child groups and hosts not found on the external source are deleted from the local inventory. When not checked, local child hosts and groups not found on the external source remain untouched by the inventory update process.
- **Overwrite Variables:** If checked, all variables for child groups and hosts are removed and replaced by those found on the external source. When not checked, a merge is performed, combining local variables with those found on the external source.
- **Update on Launch:** Each time a job runs using this inventory, refresh the inventory from the selected source before executing job tasks.

Microsoft Azure Resource Manager

To configure a group for Microsoft Azure Resource Manager, select **Microsoft Azure Resource Manager** and enter the following details:

- **Cloud Credential:** Choose from an existing Credential. For more information, refer to [Credentials](#).
- **Regions:** Click on the regions field to see a list of regions for your cloud provider. You can select multiple regions, or choose “All” to include all regions. Tower will only be updated with Hosts associated with the selected regions.
- **Variables:** Enter definitions and values to be applied to all hosts in this group. Enter variables using either JSON or YAML syntax. Use the radio button to toggle between the two.

You can also configure **Update Options**.

- **Overwrite:** If checked, all child groups and hosts not found on the external source are deleted from the local inventory. When not checked, local child hosts and groups not found on the external source remain untouched by the inventory update process.
- **Overwrite Variables:** If checked, all variables for child groups and hosts are removed and replaced by those found on the external source. When not checked, a merge is performed, combining local variables with those found on the external source.
- **Update on Launch:** Each time a job runs using this inventory, refresh the inventory from the selected source before executing job tasks.

CREATE GROUP

DETAILS | NOTIFICATIONS

* NAME: CMS Web Group - West

DESCRIPTION: CMS Web West Coast

SOURCE: Microsoft Azure Resource Manager

* CLOUD CREDENTIAL: [Search]

REGIONS: [Search]

UPDATE OPTIONS:

- ☐ Overwrite
- ☐ Overwrite Variables
- ☐ Update on Launch

VARIABLES: ☒ YAML ☐ JSON

1 [Text Area]

CANCEL SAVE

VMware vCenter

To configure a group for VMware vCenter, select **VMware** and enter the following details:

- **Cloud Credential:** Choose from an existing Credential. For more information, refer to [Credentials](#).
- **Variables:** Enter definitions and values to be applied to all hosts in this group. Enter variables using either JSON or YAML syntax. Use the radio button to toggle between the two.
- **Source Variables:** Override variables found in `vmware.ini` and used by the inventory update script. For a detailed description of these variables [view vmware.ini in the Ansible GitHub repo](#). Enter variables using either JSON or YAML syntax. Use the radio button to toggle between the two.

You can also configure **Update Options**.

- **Overwrite:** If checked, all child groups and hosts not found on the external source are deleted from the local inventory. When not checked, local child hosts and groups not found on the external source remain untouched by the inventory update process.
- **Overwrite Variables:** If checked, all variables for child groups and hosts are removed and replaced by those found on the external source. When not checked, a merge is performed, combining local variables with those found on the external source.
- **Update on Launch:** Each time a job runs using this inventory, refresh the inventory from the selected source before executing job tasks.

CREATE GROUP

DETAILS | NOTIFICATIONS

*NAME: CMS Web Group - West

DESCRIPTION: CMS Web West Coast

SOURCE: VMware vCenter

*CLOUD CREDENTIAL: [Search Icon]

UPDATE OPTIONS:

- ☐ Overwrite ?
- ☐ Overwrite Variables ?
- ☐ Update on Launch ?

VARIABLES ? ☒ YAML ☐ JSON

SOURCE VARIABLES ? ☒ YAML ☐ JSON

CANCEL SAVE

Red Hat Satellite 6

To configure a group for Red Hat Satellite 6, select **Red Hat Satellite 6** and enter the following details:

- **Cloud Credential:** Choose from an existing Credential. For more information, refer to [Credentials](#).
- **Variables:** Enter definitions and values to be applied to all hosts in this group. Enter variables using either JSON or YAML syntax. Use the radio button to toggle between the two.

You can also configure **Update Options**.

- **Overwrite:** If checked, all child groups and hosts not found on the external source are deleted from the local inventory. When not checked, local child hosts and groups not found on the external source remain untouched by the inventory update process.
- **Overwrite Variables:** If checked, all variables for child groups and hosts are removed and replaced by those found on the external source. When not checked, a merge is performed, combining local variables with those found on the external source.
- **Update on Launch:** Each time a job runs using this inventory, refresh the inventory from the selected source before executing job tasks.

Red Hat CloudForms

To configure a group for Red Hat CloudForms, select **Red Hat CloudForms** and enter the following details:

- **Cloud Credential:** Choose from an existing Credential. For more information, refer to [Credentials](#).
- **Variables:** Enter definitions and values to be applied to all hosts in this group. Enter variables using either JSON or YAML syntax. Use the radio button to toggle between the two.

You can also configure **Update Options**.

- **Overwrite:** If checked, all child groups and hosts not found on the external source are deleted from the local inventory. When not checked, local child hosts and groups not found on the external source remain untouched by the inventory update process.
- **Overwrite Variables:** If checked, all variables for child groups and hosts are removed and replaced by those found on the external source. When not checked, a merge is performed, combining local variables with those found on the external source.
- **Update on Launch:** Each time a job runs using this inventory, refresh the inventory from the selected source before executing job tasks.

CREATE GROUP

DETAILS | NOTIFICATIONS

* NAME: CMS Web Group - West

DESCRIPTION: CMS Web West Coast

SOURCE: Red Hat CloudForms

* CLOUD CREDENTIAL:

UPDATE OPTIONS:

- ☐ Overwrite
- ☐ Overwrite Variables
- ☐ Update on Launch

VARIABLES ☒ ☐ YAML ☐ JSON

CANCEL SAVE

OpenStack

To configure a group for OpenStack, select **OpenStack** and enter the following details:

- **Cloud Credential:** Choose from an existing Credential. For more information, refer to [Credentials](#).
- **Variables:** Enter definitions and values to be applied to all hosts in this group. Enter variables using either JSON or YAML syntax. Use the radio button to toggle between the two.
- **Source Variables:** Override variables found in `openstack.yml` and used by the inventory update script. For a detailed description of these variables [view openstack.yml in the Ansible GitHub repo](#). Enter variables using either JSON or YAML syntax. Use the radio button to toggle between the two.

You can also configure **Update Options**.

- **Overwrite:** If checked, all child groups and hosts not found on the external source are deleted from the local inventory. When not checked, local child hosts and groups not found on the external source remain untouched by the inventory update process.
- **Overwrite Variables:** If checked, all variables for child groups and hosts are removed and replaced by those found on the external source. When not checked, a merge is performed, combining local variables with those found on the external source.
- **Update on Launch:** Each time a job runs using this inventory, refresh the inventory from the selected source before executing job tasks.

CREATE GROUP

DETAILS

NOTIFICATIONS

* NAME

CMS Web Group - West

DESCRIPTION

CMS Web West Coast

SOURCE

OpenStack

* CLOUD CREDENTIAL

Q

UPDATE OPTIONS

☐ Overwrite
 ☐ Overwrite Variables
 ☐ Update on Launch

VARIABLES

YAML

JSON

1

SOURCE VARIABLES

YAML

JSON

1

CANCEL

SAVE

Custom Script

Tower allows you to use a custom dynamic inventory script, if your administrator has added one.

To configure a group to use a Custom Inventory Script, select **Custom Script** and enter the following details:

- **Custom Inventory Script:** Choose from an existing Inventory Script.
- **Environment Variables:** Set variables in the environment to be used by the inventory update script. The variables would be specific to the script that you have written.

Enter variables using either JSON or YAML syntax. Use the radio button to toggle between the two.

You can also configure **Update Options**.

- **Overwrite:** If checked, all child groups and hosts not found on the external source are deleted from the local inventory. When not checked, local child hosts and groups not found on the external source remain untouched by the inventory update process.
- **Overwrite Variables:** If checked, all variables for child groups and hosts are removed and replaced by those found on the external source. When not checked, a merge is performed, combining local variables with those found on the external source.
- **Update on Launch:** Each time a job runs using this inventory, refresh the inventory from the selected source before executing job tasks.

SETTINGS / INVENTORY SCRIPTS / CREATE INVENTORY SCRIPT



NEW CUSTOM INVENTORY

*NAME: Host-a-nator

DESCRIPTION: Host creation custom script

*ORGANIZATION: Honey Dog, Inc.

*CUSTOM SCRIPT ⓘ

```
#!/usr/bin/env python

# Python
import json
import optparse
import os

nhosts = int(os.environ.get('NHOSTS', 100))

inv_list = {
```


CANCEL SAVE

For more information on syncing or using custom inventory scripts, refer to [Custom Inventory Scripts](#) in the *Ansible Tower Administration Guide*.

Scheduling

For groups sourced from a cloud service, the inventory update process may be scheduled via the **Schedule** view. To



access the **Schedule** view, click the Schedule () button beside the Inventory Group name to open the **Edit Group** dialog.

INVENTORIES / WEB SERVERS

GROUPS 1 + ADD GROUP

NAME SEARCH

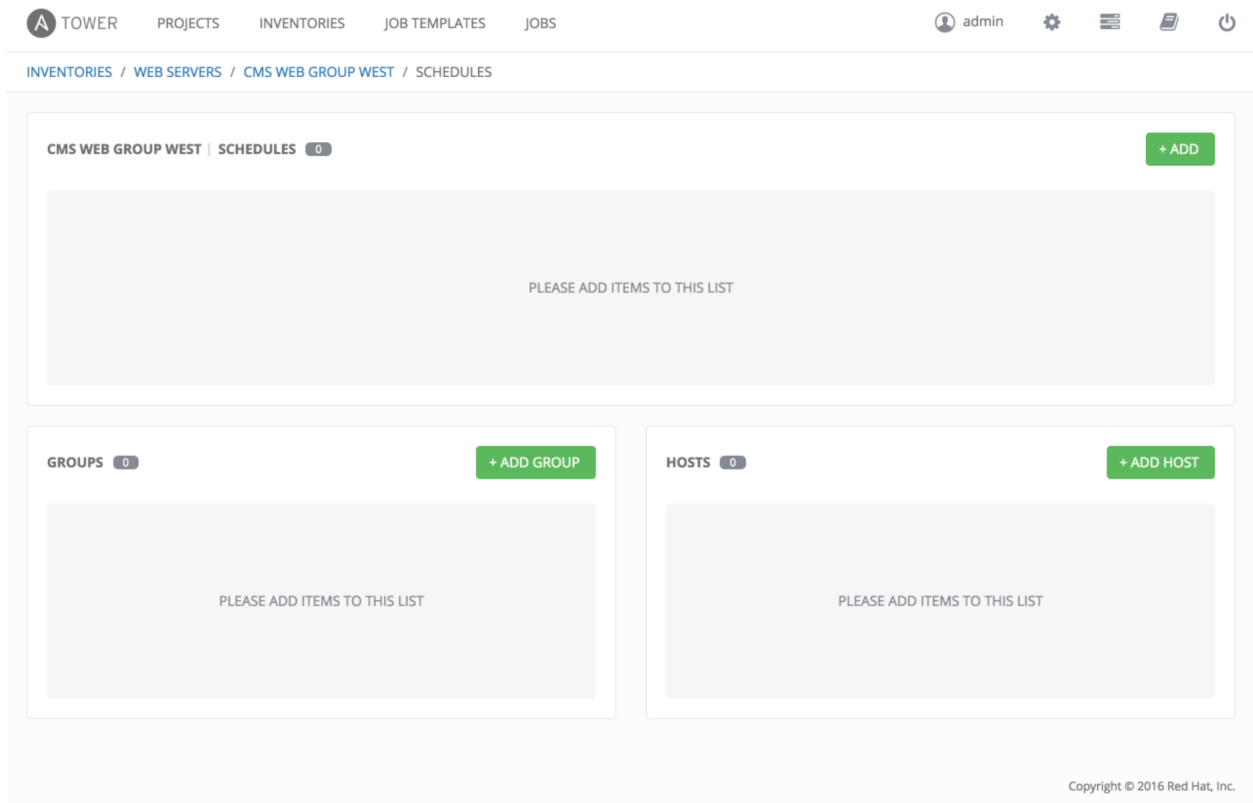
GROUPS	ACTIONS
<input type="checkbox"/> CMS Web Gro up West	

ITEMS 1-1 OF 1

HOSTS 0 + ADD HOST

PLEASE ADD ITEMS TO THIS LIST

This screen displays a list of the schedules that are currently available for the selected **Group**. The schedule list may be sorted and searched by **Name**.




The list of schedules includes:

- Name (Clicking the schedule name opens the **Edit Schedule** dialog)
- First Run
- Next Run

Buttons located in the upper right corner of the **Schedules** screen provide the following actions:

- Create a new schedule
- Refresh this view

Add a new schedule

To create a new schedule click the  button.

WEEKLY SYNC

* NAME

Weekly Sync

* START DATE (MM/DD/YYYY)

06/27/2016

* START TIME (HH24:MM:SS)

23:30:00

* LOCAL TIME ZONE

America/New_York

* REPEAT FREQUENCY

Week

FREQUENCY DETAILS

* EVERY

2

WEEKS

* ON DAYS

SUN MON TUE WED THU FRI SAT

* END

After

* OCCURRENCE(S)

6

SCHEDULE DESCRIPTION

every 2 weeks on Thursday, Sunday for 6 times
OCCURRENCES (Limited to first 10) DATE FORMAT ☒ LOCAL TIME ☐ UTC
06/30/2016 23:30:00 EDT
07/03/2016 23:30:00 EDT
07/14/2016 23:30:00 EDT
07/17/2016 23:30:00 EDT
07/28/2016 23:30:00 EDT
07/31/2016 23:30:00 EDT

EXTRA VARIABLES ☒ YAML ☐ JSON

1

CANCEL

SAVE

Enter the appropriate details into the following fields and select Save:

- **Name** (required)
- **Start Date** (required)
- **Start Time** (required)
- **Local Time Zone** (the entered Start Time should be in this timezone)
- **Repeat Frequency** (the appropriate options are displayed as the update frequency is modified)

- None

QUICK SYNC

* NAME

Quick Sync

* START DATE (MM/DD/YYYY)

06/30/2016

* START TIME (HH24:MM:SS)

00:00:00

* LOCAL TIME ZONE

America/New_York

* REPEAT FREQUENCY

None (run once)

SCHEDULE DESCRIPTION

every day for 1 time
OCCURRENCES (Limited to first 10) DATE FORMAT ☒ LOCAL TIME ☐ UTC
06/30/2016 00:00:00 EDT

- Minute, Hour, or Day

DAILY SYNC

* NAME: Daily Sync

* START DATE (MM/DD/YYYY): 06/30/2016

* START TIME (HH24:MM:SS): 00:00:00

* LOCAL TIME ZONE: America/New_York

* REPEAT FREQUENCY: Day

FREQUENCY DETAILS

* EVERY: 1 DAYS

* END: After

* OCCURRENCE(S): 5

SCHEDULE DESCRIPTION

every day for 5 times

OCCURRENCES (Limited to first 10) DATE FORMAT ☒ LOCAL TIME ☐ UTC

06/30/2016 00:00:00 EDT
 07/01/2016 00:00:00 EDT
 07/02/2016 00:00:00 EDT
 07/03/2016 00:00:00 EDT
 07/04/2016 00:00:00 EDT

• Week

BI-WEEKLY SYNC

* NAME: Bi-Weekly Sync

* START DATE (MM/DD/YYYY): 06/30/2016

* START TIME (HH24:MM:SS): 00:00:00

* LOCAL TIME ZONE: America/New_York

* REPEAT FREQUENCY: Week

FREQUENCY DETAILS

* EVERY: 2 WEEKS

* ON DAYS: SUN MON TUE WED THU FRI SAT

* END: On Date

* END DATE (MM/DD/YYYY): 08/26/2016

* END TIME (HH24:MM:SS): 00:00:00

SCHEDULE DESCRIPTION

every 2 weeks on Monday, Wednesday, Thursday, Sunday until August 25, 2016

OCCURRENCES (Limited to first 10) DATE FORMAT ☒ LOCAL TIME ☐ UTC

06/30/2016 00:00:00 EDT
 07/03/2016 00:00:00 EDT
 07/11/2016 00:00:00 EDT
 07/13/2016 00:00:00 EDT
 07/14/2016 00:00:00 EDT
 07/17/2016 00:00:00 EDT
 07/25/2016 00:00:00 EDT
 07/27/2016 00:00:00 EDT
 07/28/2016 00:00:00 EDT
 07/31/2016 00:00:00 EDT

• Month

MONTHLY SYNC

* NAME: Monthly Sync

* START DATE (MM/DD/YYYY): 06/30/2016

* START TIME (HH24:MM:SS): 00:00:00

* LOCAL TIME ZONE: America/New_York

* REPEAT FREQUENCY: Month

FREQUENCY DETAILS

* EVERY: 1 MONTHS

* ON DAY: 1

* ON THE: first Sunday

* END: After

* OCCURRENCE(S): 3

SCHEDULE DESCRIPTION

every month on Sunday for 3 times (~ approximate)

OCCURRENCES (Limited to first 10) DATE FORMAT ☒ LOCAL TIME ☐ UTC

07/03/2016 00:00:00 EDT
08/07/2016 00:00:00 EDT
09/04/2016 00:00:00 EDT

• Year

YEARLY SYNC

* NAME: Yearly Sync

* START DATE (MM/DD/YYYY): 06/30/2016

* START TIME (HH24:MM:SS): 00:00:00

* LOCAL TIME ZONE: America/New_York

* REPEAT FREQUENCY: Year

FREQUENCY DETAILS

* EVERY: 1 YEARS

* ON: January 1

* ON THE: first Sunday January

* END: After

* OCCURRENCE(S): 4

SCHEDULE DESCRIPTION

every January on the 1st for 4 times

OCCURRENCES (Limited to first 10) DATE FORMAT ☒ LOCAL TIME ☐ UTC

01/01/2017 00:00:00 EST
01/01/2018 00:00:00 EST
01/01/2019 00:00:00 EST
01/01/2020 00:00:00 EST

The **Schedule Description** panel displays an detailed overview of the schedule and a list of the scheduled occurrences in the selected Local Time Zone or in UTC (be sure to select **Local Time** or **UTC** as based on your needs).

Note: When using UTC time settings, repeating jobs that runs at a specific time of day may move relative to a local timezone when Daylight Saving Time shifts occur.



Once you have saved the schedule, it can be viewed by clicking on the Schedule () tab beside the group name.

The screenshot shows the Ansible Tower interface. At the top, there's a navigation bar with 'TOWER', 'PROJECTS', 'INVENTORIES', 'JOB TEMPLATES', and 'JOBS'. On the right, there's a user profile 'admin' and several icons. Below the navigation bar, the breadcrumb 'INVENTORIES / WEB SERVERS / SCHEDULES' is visible.

The main content area is titled 'CMS WEB GROUP WEST | SCHEDULES 1'. It features a search bar and a table of schedules. The table has columns: NAME, FIRST RUN, NEXT RUN, FINAL RUN, and ACTIONS. One schedule is listed: 'Weekly Sync' with a status 'ON', first run '6/29/2016 11:30:00 PM', next run '6/29/2016 11:30:00 PM', and final run '7/30/2016 11:30:00 PM'. The ACTIONS column shows edit and delete icons.

Below the table, there are two sections: 'GROUPS 1' and 'HOSTS 0'. The 'GROUPS' section has a search bar and a table with one group: 'CMS Web Gro up West'. The 'HOSTS' section is empty and contains the text 'PLEASE ADD ITEMS TO THIS LIST'.

Copyright © 2016 Red Hat, Inc.

There are server actions available for schedules:

- **ON/OFF** – Stop an active schedule or activate a stopped schedule by using the toggle button.
- Edit schedule
- Delete schedule

11.2.2 Hosts

Hosts are listed on the right side of the Inventory display screen.

The screenshot shows the Ansible Tower interface. At the top, there's a navigation bar with 'TOWER', 'PROJECTS', 'INVENTORIES', 'JOB TEMPLATES', and 'JOBS'. On the right, there's a user profile 'admin' and several icons. Below the navigation bar, the breadcrumb 'INVENTORIES / DATABASE SERVERS' is visible.

The main content area is titled 'DATABASE SERVERS | HOSTS 0'. It features a search bar and a table of hosts. The table has columns: NAME, FIRST RUN, NEXT RUN, FINAL RUN, and ACTIONS. No hosts are listed.

Below the table, there are two sections: 'GROUPS 1' and 'HOSTS 0'. The 'GROUPS' section has a search bar and a table with one group: 'CMS Web Gro up West'. The 'HOSTS' section is empty and contains the text 'PLEASE ADD ITEMS TO THIS LIST'.

The host list may be sorted and searched by **Name** or **Groups**, and filtered by hosts that are disabled, by hosts with failed jobs, and by hosts synchronized with an external source.

This list displays information about each host and provides for several actions:

- **ON/OFF** Indicates whether a host is available and should be included in running jobs. For hosts that are part of an external inventory, this flag is set by the inventory sync process and cannot be manually changed.

- **Name:** Opens the **Host Properties** dialog
- **Available:** A toggle indicating whether the host is enabled to receive jobs from Tower. Click to toggle this setting.
- **Jobs:** Shows the most recent Jobs run against this Host. Clicking this button displays a window showing the most recent jobs and their status.
- **Edit host:** Opens the **Host Properties** dialog
- **Copy host:** Copies or moves the host to a different group
- **Delete:** Removes the host from Tower. *This operation is not reversible!*

Add a new host

Hosts can be added manually, by IP address, or hostname. Tower can also sync inventory directly from AWS EC2, Google Compute Engine, MS Azure, VMware, Rackspace Open Cloud, or OpenStack.

+ ADD HOST

To create a new host and add it to an existing group, click the button.

This opens the **Create New Host** dialog.

[INVENTORIES](#) / [WEB SERVERS](#) / [CREATE HOST](#)

Enter the appropriate details into the following fields and click **Save**:

- **Host Name:** The hostname or IP address of the host
- **Description:** Enter an arbitrary description as appropriate
- **Enabled?:** Indicates if a host is available and should be included in running jobs. For hosts that are part of an external inventory, this flag cannot be changed. It is set by the inventory sync process.
- **Variables:** Variable definitions and values to be applied to the selected host. Enter variables using either JSON or YAML syntax, using the radio button to toggle between JSON or YAML.

11.3 Running Ad Hoc Commands

To run an ad hoc command, first select an inventory source. The inventory source can be a single group or host, a selection of multiple hosts, or a selection of multiple groups.

INVENTORIES / WEB SERVERS

The screenshot shows the Ansible Tower interface with two main sections: 'GROUPS' and 'HOSTS'. In the 'GROUPS' section, there is a 'RUN COMMANDS' button highlighted with a red box. The 'HOSTS' section shows a host named '127.0.0.1'.

Then, click the **RUN COMMANDS** button.

The screenshot shows the 'EXECUTE COMMAND' form in Ansible Tower. The form includes the following fields and buttons:

- *MODULE**: A dropdown menu with 'Choose a module' selected.
- ARGUMENTS**: A text input field.
- HOST PATTERN**: A text input field with 'CMS Web Group' entered.
- *MACHINE CREDENTIAL**: A dropdown menu with 'Demo Credential' selected.
- ENABLE PRIVILEGE ESCALATION**: A checkbox that is currently unchecked.
- *VERBOSITY**: A dropdown menu with '0 (Normal)' selected.
- FORKS**: A dropdown menu with '0' selected.
- RESET** and **LAUNCH** buttons at the bottom right.

Enter the details for the following fields:

- **Module:** Select one of the modules that Tower supports running commands against.

command	apt_repository	mount	win_service
shell	apt_rpm	ping	win_updates
yum	service	selinux	win_group
apt	group	setup	win_user
apt_key	user	win_ping	

- **Arguments:** Provide arguments to be used with the module you selected.
- **Host Pattern:** Enter the pattern used to target hosts in the inventory. To target all hosts in the inventory enter `all` or `*`, or leave the field blank. This is automatically populated with whatever was selected in the previous view prior to clicking the launch button.
- **Machine Credential:** Select the credential to use when accessing the remote hosts to run the command. Choose the credential containing the username and SSH key or password that Ansible needs to log into the remote hosts.
- **Enable Privilege Escalation:** If enabled, the playbook is run with administrator privileges. This is the equivalent of passing the `--become` option to the `ansible` command.
- **Verbosity:** Select a verbosity level for the standard output.
- **Forks:** If needed, select the number of parallel or simultaneous processes to use while executing the command.

Click the **LAUNCH** button to run this ad hoc command.

EXECUTE COMMAND

* MODULE ⓘ

ping

ARGUMENTS ⓘ

HOST PATTERN ⓘ

CMS Web Group

* MACHINE CREDENTIAL ⓘ

Q Demo Credential

☐

ENABLE PRIVILEGE ESCALATION ⓘ

* VERBOSITY ⓘ

0 (Normal)

FORKS ⓘ

0

RESET

LAUNCH

GROUPS 1

RUN COMMANDS

+ ADD GROUP

NAME SEARCH

Q

GROUPS ^

☒

CMS Web Gro up

↺

ITEMS 1-1 OF 1

HOSTS 1

SYSTEM TRACKING

+ ADD HOST

NAME SEARCH

Q

HOSTS ^

☐

127.0.0.1

ITEMS 1-1 OF 1

11.4 System Tracking

Note: System Tracking, introduced as a new feature in Ansible Tower 2.2, is only available to those with Enterprise-level licenses.

System Tracking offers the ability to compare the results of two scan runs from different dates on one host or the same date on two hosts.

Data is grouped by fact modules:

- Packages
- Services
- Files
- Ansible
- Custom

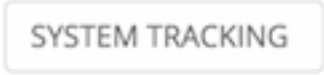
Tower is designed to make every attempt to find your data. If you select a date without any scan runs, Tower gathers the previous year's worth of scan runs to verify possible data to include. Successful comparisons display results from the available dates instead of the specified dates. Unsuccessful comparisons display a message indicating why they did not work.

Note: Service scan jobs should not run against an inventory with hosts that point to the same physical machine.

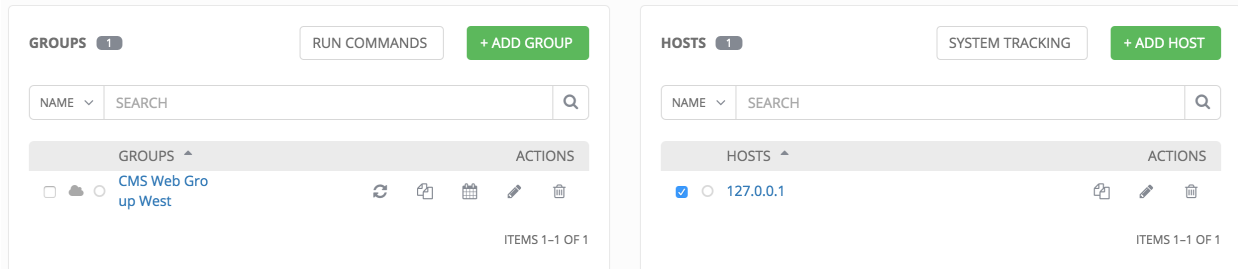
11.4. System Tracking

86

11.4.1 Single Host Workflow

Select a single host in an inventory to compare against two dates and click the  button.

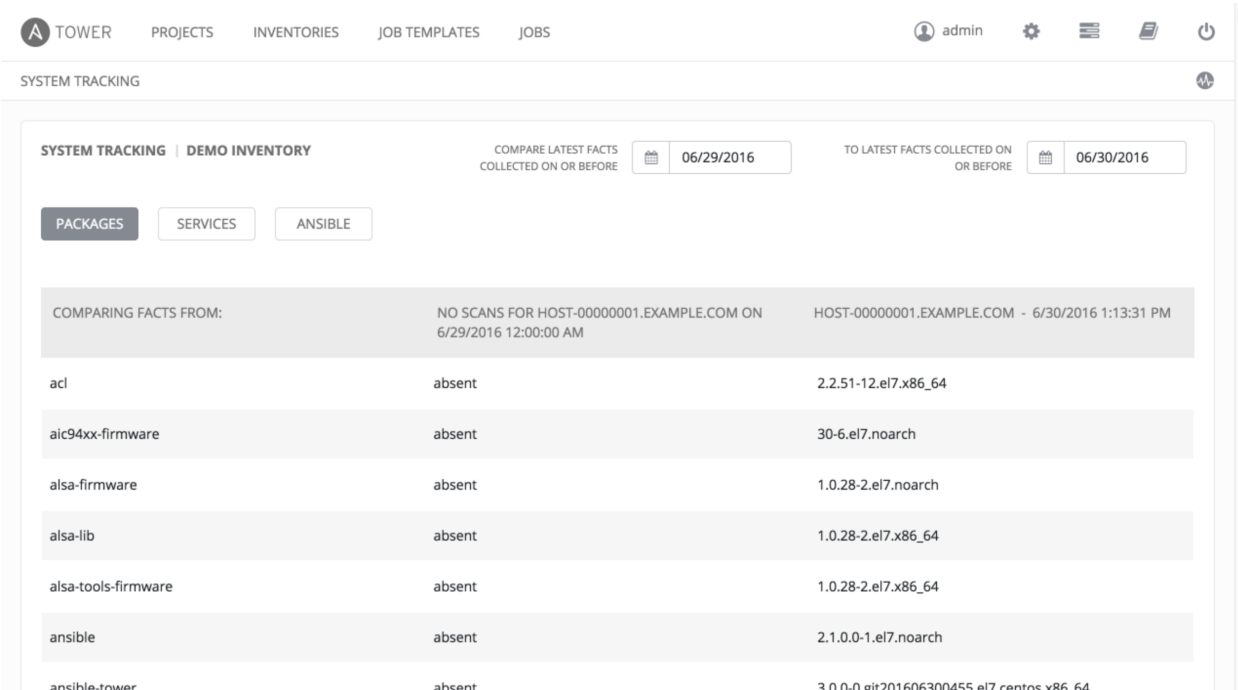
[INVENTORIES](#) / [WEB SERVERS](#)



The screenshot shows two panels. The left panel, titled 'GROUPS 1', has a search bar and a list of groups including 'CMS Web Group West'. The right panel, titled 'HOSTS 1', has a search bar and a list of hosts including '127.0.0.1'. Both panels have a 'SYSTEM TRACKING' button and a '+ ADD' button. The 'HOSTS' panel also has a 'RUN COMMANDS' button.

Note: If you have not already created a job template set to scan, you will not be able to proceed until the correct job template has been created.

Select two dates on which you have scan data for the host, with the earliest date to compare on the left and the latest date to compare on the right.



The screenshot shows the 'SYSTEM TRACKING' page for 'DEMO INVENTORY'. It has a navigation bar with 'TOWER', 'PROJECTS', 'INVENTORIES', 'JOB TEMPLATES', and 'JOBS'. The 'SYSTEM TRACKING' section is active. It shows a comparison of facts for a host (127.0.0.1) between two dates: 06/29/2016 and 06/30/2016. The 'PACKAGES' module is selected, and the comparison shows that no scans were found for the host on the earlier date.

COMPARING FACTS FROM:	NO SCANS FOR HOST-00000001.EXAMPLE.COM ON 6/29/2016 12:00:00 AM	HOST-00000001.EXAMPLE.COM - 6/30/2016 1:13:31 PM
acl	absent	2.2.51-12.el7.x86_64
aic94xx-firmware	absent	30-6.el7.noarch
alsa-firmware	absent	1.0.28-2.el7.noarch
alsa-lib	absent	1.0.28-2.el7.x86_64
alsa-tools-firmware	absent	1.0.28-2.el7.x86_64
ansible	absent	2.1.0.0-1.el7.noarch
ansible-tower	absent	3.0.0-0.git201606300455.el7.centos.x86_64

Select the module (**Packages**, **Services**, or **Ansible**) for which you want to compare differences. To change modules, click on the module button with the button navigation to filter by different types of facts. Note that differences among the “ansible” and “files” modules changes are highlighted in red, while only changes for “packages” and “services” are shown.

You may also choose the same date in both date selectors if you want to compare multiple scan runs against a single date. If two or more scan jobs runs are discovered on a particular day, Tower compares the most recent and the second-most recent. If there is only one run for the selected date, Tower may display a message saying it could not find any

scan job runs in one of the columns. (Also noted in [Known Issues](#) in the *Ansible Tower Release Notes* manual.)

SYSTEM TRACKING | DEMO INVENTORY

COMPARE LATEST FACTS
COLLECTED ON OR BEFORE

06/30/2016

TO LATEST FACTS COLLECTED ON
OR BEFORE

06/30/2016

PACKAGES

SERVICES

ANSIBLE

COMPARING FACTS FROM:			HOST-00000001.EXAMPLE.COM - 6/30/2016 1:13:31 PM	HOST-00000001.EXAMPLE.COM - 6/30/2016 1:20:11 PM
ansible_product_serial	NA		NA	
ansible_form_factor	Other		Other	
ansible_product_version	1.2		1.2	
ansible_fips	false		false	
ansible_service_mgr	systemd		systemd	
ansible_user_id	awx		awx	
ansible_user_dir	/var/lib/awx		/var/lib/awx	
ansible_userspace_bits	64		64	
ansible_ssh_host_key_rsa_public	AAAAB3NzaC1yc2EAAAADAQABAAQDU0nk47tYOcejwWgW6/LaRA3Xwfo2j6U+h5G49o6kUwQIKAcTGhAEukYmhIRfGFwaEzpEI5Q3lynFop1VskyGHdTq9I0BSDvbbar2xM6Me/p0MNXKaYaloiH6X1t6DI4aQeo3RYHncPh7dm/W2KHs79hRhqaqucYZ4YmgjhQ32TeureeHq6Pqj8DIOFHH+r6UgwZS2Z2+x5mGja427yMLqnD82cz3aFRu3pcErdfxWj181q+Dg2KE1K7wjpuHEFeioaSreJxVDJ4cIUQyb6MMDQkvNwORFci7K48LejOSrcldlMrGisPF4cnn1iK74Bip1yMeCoWe1JQhAOPGPfZT		AAAAB3NzaC1yc2EAAAADAQABAAQDU0nk47tYOcejwWgW6/LaRA3Xwfo2j6U+h5G49o6kUwQIKAcTGhAEukYmhIRfGFwaEzpEI5Q3lynFop1VskyGHdTq9I0BSDvbbar2xM6Me/p0MNXKaYaloiH6X1t6DI4aQeo3RYHncPh7dm/W2KHs79hRhqaqucYZ4YmgjhQ32TeureeHq6Pqj8DIOFHH+r6UgwZS2Z2+x5mGja427yMLqnD82cz3aFRu3pcErdfxWj181q+Dg2KE1K7wjpuHEFeioaSreJxVDJ4cIUQyb6MMDQkvNwORFci7K48LejOSrcldlMrGisPF4cnn1iK74Bip1yMeCoWe1JQhAOPGPfZT	
ansible_ssh_host_key_ecdsa_public	AAAAE2VjZHNhLXNoYTItbmlzdHAYNTYAAAAIbmlzdHAYNTYAAABBBcsd4s1dE22hUtoxQmFt9N+ZQFcxllhRQIYSAfTC2Nlc+Q8Iok/oKgphfE9coUyUD3lWF9M9TZar7JfI87pGcj4=		AAAAE2VjZHNhLXNoYTItbmlzdHAYNTYAAAAIbmlzdHAYNTYAAABBBcsd4s1dE22hUtoxQmFt9N+ZQFcxllhRQIYSAfTC2Nlc+Q8Iok/oKgphfE9coUyUD3lWF9M9TZar7JfI87pGcj4=	
ansible_distribution_version	7.2.1511		7.2.1511	
ansible_domain	localdomain		localdomain	
ansible_ssh_host_key_ed25519_public	AAAAC3NzaC1lZDI1NTE5AAAAILSAqff3xMyHiFxsRZUN680fFOMcmYYRD3h1FXN1joY		AAAAC3NzaC1lZDI1NTE5AAAAILSAqff3xMyHiFxsRZUN680fFOMcmYYRD3h1FXN1joY	
ansible_processor_cores	2		2	
ansible_virtualization_role	guest		guest	
ansible_processor_vcpus	2		2	
ansible_system_capabilities_enforced	True		True	
ansible_bios_version	VirtualBox		VirtualBox	
ansible_virtualization_type	virtualbox		virtualbox	
ansible_memtotal_mb	1840		1840	
ansible_architecture	x86_64		x86_64	
ansible_swapfree_mb	2047		1691	
ansible_system_vendor	innotek GmbH		innotek GmbH	

Please note that if the scans found for the selected date are identical, Tower displays a single result of all facts scanned. As an example, say that a user selects “7/7/2015” for both dates and selects the “packages” module. And say that two runs occurred on this date, but there were no changes to packages on the selected host. The user sees a message indicating the scans were identical as well as a single column containing all package versions, instead of a two-column listing of differences.

SYSTEM TRACKING
DEMO INVENTORY

COMPARE LATEST FACTS
COLLECTED ON OR BEFORE

06/30/2016

TO LATEST FACTS COLLECTED ON
OR BEFORE

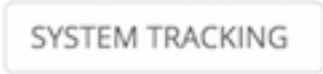
06/30/2016

PACKAGES
SERVICES
ANSIBLE

THE SELECTED FACT SCANS WERE IDENTICAL FOR THIS MODULE. SHOWING ALL FACTS FROM THE LATEST SELECTED SCAN INSTEAD.

COMPARING FACTS FROM:	HOST-00000001.EXAMPLE.COM - 6/30/2016 1:20:11 PM
acl	2.2.51-12.el7.x86_64
aic94xx-firmware	30-6.el7.noarch
alsa-firmware	1.0.28-2.el7.noarch
alsa-lib	1.0.28-2.el7.x86_64

11.4.2 Host to Host Workflow

To compare two hosts, select the hosts and click the  button.

Copyright © 2016 Red Hat, Inc.

90

TOWER

PROJECTS

INVENTORIES

JOB TEMPLATES

JOBS

admin

INVENTORIES / WEB SERVERS

GROUPS 1

RUN COMMANDS

+ ADD GROUP

NAME

SEARCH

GROUPS

ACTIONS

CMS Web Gro
up West

22

ITEMS 1-1 OF 1

HOSTS 101

SYSTEM TRACKING

+ ADD HOST

NAME

SEARCH

HOSTS

ACTIONS

127.0.0.1

☒

host-00000000.example.com

☒

host-00000001.example.com

host-00000002.example.com

host-00000003.example.com

host-00000004.example.com

host-00000005.example.com

host-00000006.example.com

host-00000007.example.com

host-00000008.example.com

host-00000009.example.com

host-00000010.example.com

host-00000011.example.com

host-00000012.example.com

host-00000013.example.com

host-00000014.example.com

host-00000015.example.com

host-00000016.example.com

host-00000017.example.com

host-00000018.example.com

<

1

2

3

4

5

6

>

PAGE 1 OF 6

ITEMS 1-20 OF 101

Copyright © 2016 Red Hat, Inc.

Although Tower only supports picking a single date for both hosts, you may notice different dates in the results. Remember that Tower is designed to make every attempt to find your data. If a date is selected without any scan runs, Tower gathers the previous year's worth of scan runs to verify possible data to include. Note that differences among the “ansible” and “files” modules changes are highlighted in red, while only changes for “packages” and “services” are shown.

SYSTEM TRACKING | DEMO INVENTORY

COMPARE LATEST FACTS
COLLECTED ACROSS BOTH
HOSTS ON OR BEFORE

06/30/2016

- PACKAGES
- SERVICES
- ANSIBLE

COMPARING FACTS FROM:	HOST-00000001.EXAMPLE.COM - 6/30/2016 1:20:11 PM	HOST-00000003.EXAMPLE.COM - 6/30/2016 1:20:11 PM
ansible_product_serial	NA	NA
ansible_form_factor	Other	Other
ansible_product_version	1.2	1.2
ansible_fips	false	false
ansible_service_mgr	systemd	systemd
ansible_user_id	awx	awx
ansible_user_dir	/var/lib/awx	/var/lib/awx
ansible_userspace_bits	64	64
ansible_ssh_host_key_rsa_public	AAAAB3NzaC1yc2EAAAADAQABAAQDU0nk47tYOcejwWgW6/LaRA3Xwfo2j6U+h5G49o6kUwQikAectGhAEukYmhlRfGFwaEzpEI5Q3iynFop1VskyGHdTq9i0BSDvbbar2xM6Me/p0MNxKaYaloiH6X1t6DI4aQeo3RYHncPh7dm/W2KHs79hRhqaqucYZ4YmgJhQ32TeureeHq6Pqj8DIOFHH+r6UgwZS2Z2+x5mGja427yMLqnD82cz3aFRu3pcErdfxWJ181q+Dg2KE1K7wjpuHEFeioaSrejxVDJ4clUQyb6MMDQkvNwORFcl7K48LeJOSrcldIMrGisPF4cnn1IK74Bip1yMeCoWe1JQhAOPGPfZT	AAAAB3NzaC1yc2EAAAADAQABAAQDU0nk47tYOcejwWgW6/LaRA3Xwfo2j6U+h5G49o6kUwQikAectGhAEukYmhlRfGFwaEzpEI5Q3iynFop1VskyGHdTq9i0BSDvbbar2xM6Me/p0MNxKaYaloiH6X1t6DI4aQeo3RYHncPh7dm/W2KHs79hRhqaqucYZ4YmgJhQ32TeureeHq6Pqj8DIOFHH+r6UgwZS2Z2+x5mGja427yMLqnD82cz3aFRu3pcErdfxWJ181q+Dg2KE1K7wjpuHEFeioaSrejxVDJ4clUQyb6MMDQkvNwORFcl7K48LeJOSrcldIMrGisPF4cnn1IK74Bip1yMeCoWe1JQhAOPGPfZT
ansible_ssh_host_key_ecdsa_public	AAAAE2VjZHNhLXNoYTItbmlzdHAyNTYAAAAIbmlzdHAyNTYAAABBBcSd4s1dE22hUtoxQmFt9N+ZQFcxllhRQIYSAfTC2Nlc+Q8Iok/okGphfE9coUyUD3lWF9M9TZar7JfI87pGcj4=	AAAAE2VjZHNhLXNoYTItbmlzdHAyNTYAAAAIbmlzdHAyNTYAAABBBcSd4s1dE22hUtoxQmFt9N+ZQFcxllhRQIYSAfTC2Nlc+Q8Iok/okGphfE9coUyUD3lWF9M9TZar7JfI87pGcj4=
ansible_distribution_version	7.2.1511	7.2.1511
ansible_domain	localdomain	localdomain
ansible_ssh_host_key_ed25519_public	AAAAC3NzaC1lZDI1NTE5AAAAILSAqff3xMyHiFxsRZUN680fFOMcmYYRD3h1FXxn1joY	AAAAC3NzaC1lZDI1NTE5AAAAILSAqff3xMyHiFxsRZUN680fFOMcmYYRD3h1FXxn1joY
ansible_processor_cores	2	2
ansible_virtualization_role	guest	guest
ansible_processor_vcpus	2	2
ansible_system_capabilities_enforced	True	True
ansible_bios_version	VirtualBox	VirtualBox
ansible_virtualization_type	virtualbox	virtualbox
ansible_memtotal_mb	1840	1840
ansible_architecture	x86_64	x86_64
ansible_swapfree_mb	1691	1691
ansible_system_vendor	innotek GmbH	innotek GmbH

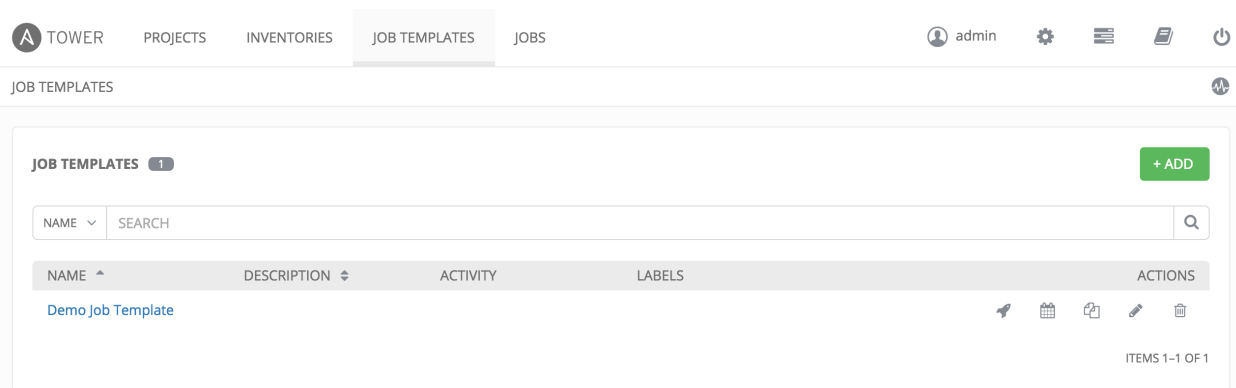
11.4. System Tracking


ansible_os_family	RedHat	RedHat
ansible_distribution_release	Core	Core
ansible_userspace_architecture	x86_64	x86_64

JOB TEMPLATES

A **job template** is a definition and set of parameters for running an Ansible job. Job templates are useful to execute the same job many times. Job templates also encourage the reuse of Ansible playbook content and collaboration between teams. While the REST API allows for the execution of jobs directly, Tower requires that you first create a job template.

This menu opens a list of the job templates that are currently available. The job template list may be sorted and searched by **Name** or **Description**. The **Job Templates** tab also enables the user to launch, schedule, modify, and remove a job template.



To create a new job template click the  button.

NEW JOB TEMPLATE

DETAILS
COMPLETED JOBS
PERMISSIONS
NOTIFICATIONS

* NAME

DESCRIPTION

* JOB TYPE

Run

☐ Prompt on launch

* INVENTORY

Q

Demo Inventory

☐ Prompt on launch

* PROJECT

Q

Demo Project

* PLAYBOOK

Choose a playbook

* MACHINE CREDENTIAL

Q

Demo Credential

☐ Prompt on launch

CLOUD CREDENTIAL

Q

NETWORK CREDENTIAL

Q

FORKS

0

LIMIT

☐ Prompt on launch

* VERBOSITY

0 (Normal)

JOB TAGS

☐ Prompt on launch

SKIP TAGS

☐ Prompt on launch

OPTIONS

☐ Enable Privilege Escalation
☐ Allow Provisioning Callbacks

LABELS

EXTRA VARIABLES

YAML
JSON

1

☐ Prompt on launch

ADD SURVEY
CANCEL
SAVE

Enter the appropriate details into the following fields:

- **Name:** Enter a name for the job (this is required).
- **Description:** Enter an arbitrary description as appropriate.
- **Job Type:**
 - **Run:** Execute the playbook when launched, running Ansible tasks on the selected hosts.
 - **Check:** Setting the type to Check does not execute the playbook, but does check the syntax, test the environment setup, and report problems. Think of this as running the playbook in dry-run mode and having it report “changed” when an item would be changed, but not actually making changes.
 - **Scan:** Gather system tracking information. Only Superusers and Admins have permission to create scan jobs. A default playbook has been created for your use. Custom written scan playbooks can use scan modules.
 - **Prompt on Launch** – If selected, even if a default value is supplied, you will be prompted upon launch to choose a job type of run or check (scan job types cannot be changed at the time of launch).

Note: More information on job types can be found in the [Playbooks: Special Topics](#) section of the Ansible

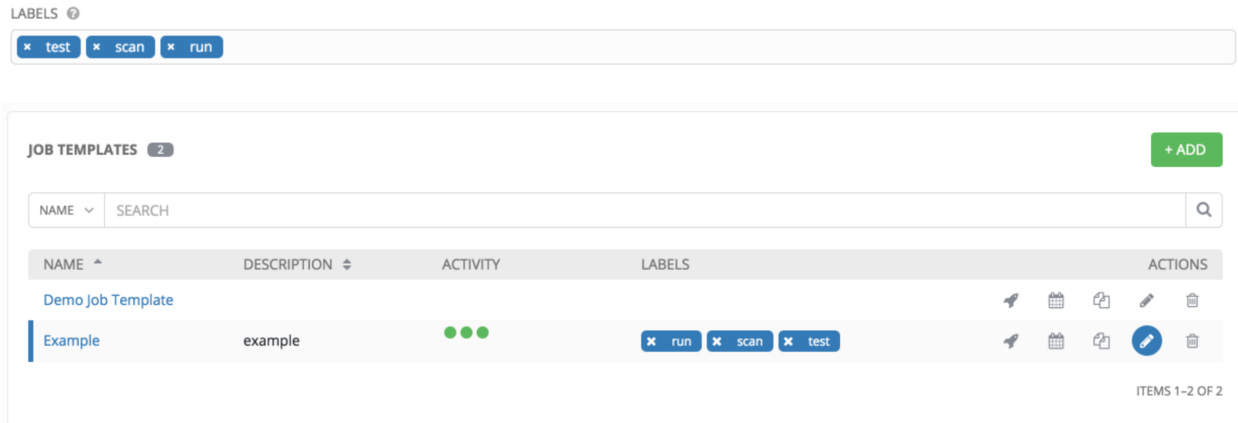
documentation.

- **Inventory:** Choose the inventory to be used with this job template from the inventories available to the currently logged in Tower user. - **Prompt on Launch** – If selected, even if a default value is supplied, you will be prompted upon launch to choose an inventory to run this job template against.
- **Project:** Choose the project to be used with this job template from the projects available to the currently logged in Tower user.
- **Playbook:** Choose the playbook to be launched with this job template from the available playbooks. This menu is automatically populated with the names of the playbooks found in the project base path for the selected project. For example, a playbook named “jboss.yml” in the project path appears in the menu as “jboss”.
- **Machine Credential:** Choose the machine credential to be used with this job template from the credentials available to the currently logged in Tower user.
- **Cloud Credential:** Choose the credential to be used with this job template from the credentials available to the currently logged in Tower user.
- **Network Credential:** Choose the network credential to be used with this job template from the credentials available to the currently logged in Tower user.
- **Forks:** The number of parallel or simultaneous processes to use while executing the playbook. A value of zero uses the Ansible default setting, which is 5 parallel processes unless overridden in `/etc/ansible/ansible.cfg`.
- **Limit:**
 - A host pattern to further constrain the list of hosts managed or affected by the playbook. Multiple patterns can be separated by colons (":"). As with core Ansible, “a:b” means “in group a or b”, “a:b:&c” means “in a or b but must be in c”, and “a:!b” means “in a, and definitely not in b”.
 - **Prompt on Launch:** If selected, even if a default value is supplied, you will be prompted upon launch to choose a limit.

Note: For more information and examples refer to [Patterns](#) in the Ansible documentation.

- **Verbosity:** Control the level of output Ansible produces as the playbook executes. Set the verbosity to any of Default, Verbose, or Debug. This only appears in the “details” report view. Verbose logging includes the output of all commands. Debug logging is exceedingly verbose and includes information on SSH operations that can be useful in certain support instances. Most users do not need to see debug mode output.
- **Job Tags:**
 - Provide a comma-separated list of playbook tags to specify what parts of the playbooks should be executed.
 - For more information and examples refer to [Tags](#) in the Ansible documentation.
 - **Prompt on Launch** – If selected, even if a default value is supplied, you will be prompted upon launch to choose a job tag.
- **Skip Tags:**
 - Provide a comma-separated list of playbook tags to skip certain tasks or parts of the playbooks to be executed.
 - For more information and examples refer to [Tags](#) in the Ansible documentation.
 - **Prompt on Launch** – If selected, even if a default value is supplied, you will be prompted upon launch to choose a job tag.

- **Labels:** Supply optional labels that describe this job template, such as “dev” or “test”. Labels can be used to group and filter job templates and completed jobs in the Tower display.
 - Labels are created when they are added to the Job Template. Labels are associated to a single Organization using the Project that is provided in the Job Template. Members of the Organization can create labels on a Job Template if they have edit permissions (such as an admin role).
 - Once the Job Template is saved, the labels appear in the Job Templates overview.
 - Click on the “x” beside a label to remove it. When a label is removed, and is no longer associated with a Job or Job Template, the label is permanently deleted from the list of Organization labels.
 - Jobs inherit labels from the Job Template at the time of launch. If a label is deleted from a Job Template, it is also deleted from the Job.



- **Extra Variables:**
 - Pass extra command line variables to the playbook. This is the “-e” or “--extra-vars” command line parameter for ansible-playbook that is documented in the Ansible documentation at [Passing Variables on the Command Line](#).
 - Provide key/value pairs using either YAML or JSON. These variables have a maximum value of precedence and overrides other variables specified elsewhere. An example value might be:


```
git_branch: production
release_version: 1.5
```
 - **Prompt on Launch** – If selected, even if a default value is supplied, you will be prompted upon launch to choose command line variables.
- **Enable Privilege Escalation:** If enabled, run this playbook as an administrator. This is the equivalent of passing the `--become` option to the `ansible-playbook` command.
- **Allow Provisioning Callbacks:** Enable a host to call back to Tower via the Tower API and invoke the launch of a job from this job template. Refer to [Provisioning Callbacks](#) for additional information.
- **Add Survey:** Creates a survey.
- **Edit Survey:** Edits the existing survey for this job template.

When you have completed configuring the job template, select **Save**.

When editing an existing job template, by clicking the job template name or the **Edit** button, the bottom of the screen displays a list of all of the jobs that have been launched from this template. Refer to the section [Jobs](#) for more information about this interface.

Note: Starting with Ansible Tower 3.0, Job Templates now have an option to be copied. If you choose to **Copy** a Job Template, you must understand that this **does not** copy any associated schedule, notifications, or permissions. Schedules and notifications must be recreated by the user or admin creating the copy of the Job Template. The user copying the Job Template will be granted the admin permission, but no permissions are assigned (copied) to the Job Template.

The **Details** view of a saved job allows you to review, edit, and add a survey (if the job type is not a scan).

EXAMPLE

DETAILS | COMPLETED JOBS | PERMISSIONS | NOTIFICATIONS

* NAME: Example

DESCRIPTION: example

* JOB TYPE: Scan

* INVENTORY: Demo Inventory

* PROJECT: Default

* PLAYBOOK: Default

* MACHINE CREDENTIAL: Demo Credential

CLOUD CREDENTIAL:

NETWORK CREDENTIAL:

☐ Prompt on launch

FORKS: 0

LIMIT:

* VERBOSITY: 0 (Normal)

☐ Prompt on launch

JOB TAGS: scan, test

OPTIONS:

- ☐ Enable Privilege Escalation
- ☐ Allow Provisioning Callbacks

☐ Prompt on launch

LABELS: scan, test

EXTRA VARIABLES: ☒ YAML ☐ JSON

1

☐ Prompt on launch

CANCEL SAVE

The **Completed Jobs** view provides details of how this job has been run. It provides you with the ID, Name, Job Type, when it completed, and allows you to relaunch or delete the job. You can filter the list of completed jobs using the job ID, Name, Type, or if the Job Failed.

EXAMPLE

DETAILS **COMPLETED JOBS** PERMISSIONS NOTIFICATIONS

STATUS ▾ FILTER ▾

ID ▾	NAME ▾	TYPE ▾	FINISHED ▾	ACTIONS
10	Example	Playbook Run	6/30/2016 1:23:06 PM	
7	Example	Playbook Run	6/30/2016 1:16:19 PM	
4	Example	Playbook Run	6/30/2016 12:52:21 PM	

ITEMS 1-3 OF 3

Clicking on **Permissions** allows you to review, grant, edit, and remove associated permissions for users as well as team members.

EXAMPLE

DETAILS COMPLETED JOBS **PERMISSIONS** NOTIFICATIONS

USER SEARCH

USER ▾	ROLE	TEAM ROLES
admin	ADMIN	
cdoge		
gdoge	EXECUTE	
jdoge	ADMIN	

ITEMS 1-4 OF 4

Click the button to create new permissions for this Job Template.

In this example, two users and one team have been selected and each have been granted permissions for this Job Template.

EXAMPLE | ADD PERMISSIONS

1

Please select Users / Teams from the lists below.

USERS

TEAMS

NAME

SEARCH

NAME ^

☒ Production Operations

ITEMS 1-1 OF 1

2

Please assign roles to the selected users/teams

KEY

Ginger Doge USER

Admin

Josie Doge USER

Execute

Production Operations TEAM

Admin

SAVE

CANCEL

Note that you do not have to choose between teams or users, and that you can assign permissions to both at the same time.

Clicking on **Notifications** allows you to review any notification integrations you have setup.

Click on the  button to create a notification.

Refer to [Notifications](#) for more information.

EXAMPLE

DETAILS

COMPLETED JOBS

PERMISSIONS

NOTIFICATIONS

+ ADD NOTIFICATION TEMPLATE

THIS LIST IS POPULATED BY NOTIFICATION TEMPLATES ADDED FROM THE [NOTIFICATIONS](#) SECTION

12.1 Utilitizing Cloud Credentials

Cloud Credentials can be used when syncing a respective cloud inventory. Cloud Credentials may also be associated with a Job Template and included in the runtime environment for use by a playbook. The use of Cloud Credentials was introduced in Ansible Tower version 2.4.0.

12.1.1 OpenStack

The sample playbook below invokes the `nova_compute` Ansible OpenStack cloud module and requires credentials to do anything meaningful, and specifically requires the following information: `auth_url`, `username`, `password`, and `project_name`. These fields are made available to the playbook via the environmental variable `OS_CLIENT_CONFIG_FILE`, which points to a YAML file written by Tower based on the contents of the cloud credential. This sample playbook loads the YAML file into the Ansible variable space.

`OS_CLIENT_CONFIG_FILE` example:

```
clouds:
  devstack:
    auth:
      auth_url: http://devstack.yoursite.com:5000/v2.0/
      username: admin
      password: your_password_here
      project_name: demo
```

Playbook example:

```
- hosts: all
  gather_facts: false
  vars:
    config_file: "{{ lookup('env', 'OS_CLIENT_CONFIG_FILE') }}"
    nova_tenant_name: demo
    nova_image_name: "cirros-0.3.2-x86_64-uec"
    nova_instance_name: autobot
    nova_instance_state: 'present'
    nova_flavor_name: ml.nano

    nova_group:
      group_name: antarctica
      instance_name: deceptacon
      instance_count: 3
  tasks:
    - debug: msg="{{ config_file }}"
    - stat: path="{{ config_file }}"
      register: st
    - include_vars: "{{ config_file }}"
      when: st.stat.exists and st.stat.isreg

    - name: "Print out clouds variable"
      debug: msg="{{ clouds|default('No clouds found') }}"

    - name: "Setting nova instance state to: {{ nova_instance_state }}"
      local_action:
        module: nova_compute
        login_username: "{{ clouds.devstack.auth.username }}"
        login_password: "{{ clouds.devstack.auth.password }}"
```

12.1.2 Amazon Web Services

Amazon Web Services cloud credentials are exposed as the following environment variables during playbook execution (in the job template, choose the cloud credential needed for your setup):

- `AWS_ACCESS_KEY`
- `AWS_SECRET_KEY`

All of the AWS modules will implicitly use these credentials when run via Tower without having to set the `aws_access_key` or `aws_secret_key` module options.

12.1.3 Rackspace

Rackspace cloud credentials are exposed as the following environment variables during playbook execution (in the job template, choose the cloud credential needed for your setup):

- `RAX_USERNAME`
- `RAX_API_KEY`

All of the Rackspace modules will implicitly use these credentials when run via Tower without having to set the `username` or `api_key` module options.

12.1.4 Google

Google cloud credentials are exposed as the following environment variables during playbook execution (in the job template, choose the cloud credential needed for your setup):

- `GCE_EMAIL`
- `GCE_PROJECT`
- `GCE_PEM_FILE_PATH`

All of the Google modules will implicitly use these credentials when run via Tower without having to set the `service_account_email`, `project_id`, or `pem_file` module options.

12.1.5 Azure

Azure cloud credentials are exposed as the following environment variables during playbook execution (in the job template, choose the cloud credential needed for your setup):

- `AZURE_SUBSCRIPTION_ID`
- `AZURE_CERT_PATH`

All of the Azure modules implicitly use these credentials when run via Tower without having to set the `subscription_id` or `management_cert_path` module options.

12.1.6 VMware

VMware cloud credentials are exposed as the following environment variables during playbook execution (in the job template, choose the cloud credential needed for your setup):

- `VMWARE_USER`
- `VMWARE_PASSWORD`

- VMWARE_HOST

The sample playbook below demonstrates usage of these credentials:

```
- vsphere_guest:
  vcenter_hostname: "{{ lookup('env', 'VMWARE_HOST') }}"
  username: "{{ lookup('env', 'VMWARE_USER') }}"
  password: "{{ lookup('env', 'VMWARE_PASSWORD') }}"
  guest: newvm001
  from_template: yes
  template_src: centosTemplate
  cluster: MainCluster
  resource_pool: "/Resources"
  vm_extra_config:
    folder: MyFolder
```

12.2 Surveys

Surveys set extra variables for the playbook similar to ‘Prompt for Extra Variables’ does, but in a user-friendly question

and answer way. Surveys also allows for validation of user input. Click the



button to create a survey.

Use cases for surveys are numerous. An example might be if operations wanted to give developers a “push to stage” button they could run without advanced Ansible knowledge. When launched, this task could prompt for answers to questions such as, “What tag should we release?”

Many types of questions can be asked, including multiple-choice questions.

Note: Surveys are only available to those with Enterprise-level licenses.

12.2.1 Creating a Survey



Clicking on the button brings up the **Add Survey** window.

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

A survey can consist of any number of questions. For each question, enter the following information:

- **Name:** The question to ask the user
- **Description:** (optional) A description of what’s being asked of the user.
- **Answer Variable Name:** The Ansible variable name to store the user’s response in. This is the variable to be used by the playbook. Variable names cannot contain spaces.
- **Answer Type:** Choose from the following question types.
 - *Text:* A single line of text. You can set the minimum and maximum length (in characters) for this answer.
 - *Textarea:* A multi-line text field. You can set the minimum and maximum length (in characters) for this answer.

NEW JOB TEMPLATE | SURVEY ON

ADD SURVEY PROMPT

* PROMPT

Which group(s) should include this user?

DESCRIPTION

Enter groups, one per line.

* ANSWER VARIABLE NAME ?

group_name

* ANSWER TYPE

Text

MINIMUM LENGTH

0

MAXIMUM LENGTH

1024

DEFAULT ANSWER

☒ REQUIRED

CANCEL

ADD

PREVIEW

PLEASE ADD A SURVEY PROMPT ON THE LEFT.

CANCEL

SAVE

- *Password*: Responses are treated as sensitive information, much like an actual password is treated. You can set the minimum and maximum length (in characters) for this answer.
- *Multiple Choice (single select)*: A list of options, of which only one can be selected at a time. Enter the options, one per line, in the **Multiple Choice Options** box.
- *Multiple Choice (multiple select)*: A list of options, any number of which can be selected at a time. Enter the options, one per line, in the **Multiple Choice Options** box.
- *Integer*: An integer number. You can set the minimum and maximum length (in characters) for this answer.
- *Float*: A decimal number. You can set the minimum and maximum length (in characters) for this answer.
- **Default Answer**: The default answer to the question. This value is pre-filled in the interface and is used if the answer is not provided by the user.
- **Required**: Whether or not an answer to this question is required from the user.

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

A stylized version 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

Click **Save** to save the survey.

12.2.2 Optional Survey Questions

The **Required** setting on a survey question determines whether the answer is optional or not for the user interacting with it.

Behind the scenes, optional survey variables can be passed to the playbook in `extra_vars`, even when they aren't filled in.

- If a non-text variable (input type) is marked as optional, and is not filled in, no survey `extra_var` is passed to the playbook.
- If a text input or text area input is marked as optional, is not filled in, and has a minimum `length > 0`, no survey `extra_var` is passed to the playbook.
- If a text input or text area input is marked as optional, is not filled in, and has a minimum `length == 0`, that survey `extra_var` is passed to the playbook, with the value set to an empty string (`""`).

12.2.3 Extra Variables

Note: Additional strict `extra_vars` validation was added in Ansible Tower 3.0.0. `extra_vars` passed to the job launch API are only honored if one of the following is true:

- They correspond to variables in an enabled survey
- `ask_variables_on_launch` is set to `True`

When you pass survey variables, they are passed as extra variables (`extra_vars`) within Tower. This can be tricky, as passing extra variables to a job template (as you would do with a survey) can override other variables being passed from the inventory and project.

For example, say that you have a defined variable for an inventory for `debug = true`. It is entirely possible that this variable, `debug = true`, can be overridden in a job template survey.

To ensure that the variables you need to pass are not overridden, ensure they are included by redefining them in the survey. Keep in mind that extra variables can be defined at the inventory, group, and host levels.

Note: Beginning with Ansible Tower version 2.4, the behavior for Job Template extra variables and Survey variables has changed. Previously, variables set using a Survey overrode any extra variables specified in the Job Template. In 2.4 and later, the Job Template extra variables dictionary is merged with the Survey variables. This may result in a change of behavior upon upgrading to 2.4.

The following table notes the behavior (hierarchy) of variable precedence in Ansible Tower as it compares to variable precedence in Ansible.

Ansible Tower Variable Precedence Hierarchy (last listed wins)

Ansible	Tower
role defaults	
dynamic inventory variables	
inventory variables	Tower inventory variables
inventory group_vars	Tower group variables
inventory host_vars	Tower host variables
playbook group_vars	
playbook host_vars	
host facts	
registered variables	
set_facts	
play variables	
play vars_prompt	(not supported in Tower)
play vars_files	
role variables and include variables	
block variables	
task variables	
extra variables	Job Template extra variables Job Template Survey (defaults) Job Launch extra_vars

12.2.4 Relaunching Job Templates

Another change for Ansible Tower version 2.4 introduced a `launch_type` setting for your jobs. Instead of manually relaunching a job, a relaunch is denoted by setting `launch_type` to `relaunch`. The relaunch behavior deviates from the launch behavior in that it **does not** inherit `extra_vars`.

Job relaunching does not go through the inherit logic. It uses the same `extra_vars` that were calculated for the job being relaunched.

For example, say that you launch a Job Template with no `extra_vars` which results in the creation of a Job called **j1**. Next, say that you edit the Job Template and add in some `extra_vars` (such as adding `"{ \"hello\": \"world\" }"`).

Relaunching **j1** results in the creation of **j2**, but because there is no inherit logic and **j1** had no `extra_vars`, **j2** will not have any `extra_vars`.

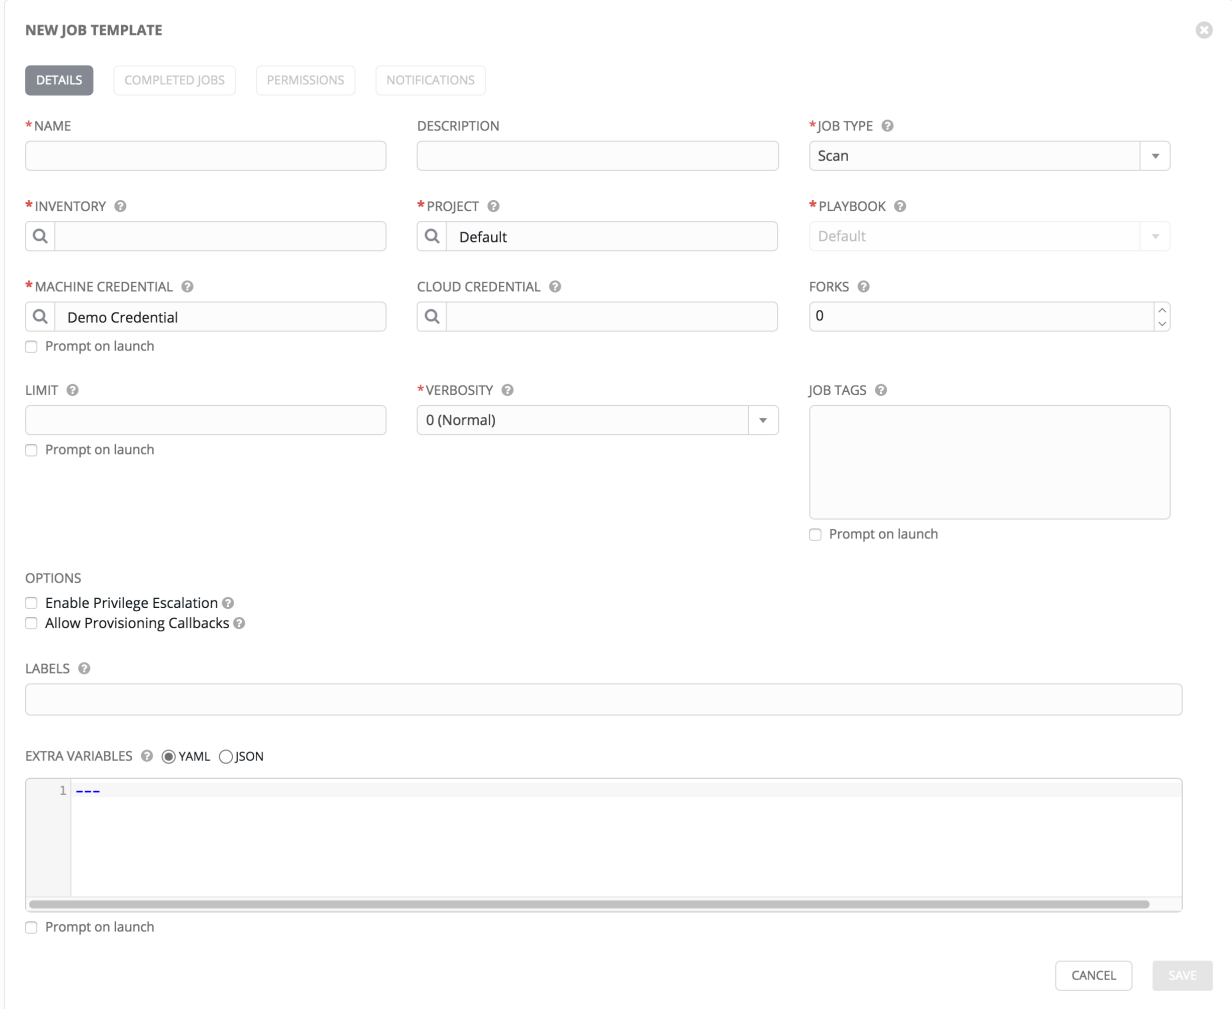
To continue upon this example, if you launched the Job Template with the `extra_vars` you added after the creation of **j1**, the relaunch job created (**j3**) will include the `extra_vars`. And relaunching **j3** results in the creation of **j4**, which would also include `extra_vars`.

12.3 Scan Job Templates

Scan jobs are special Job Templates that only collect information about the host on which the job is running.

Click on **Job Templates** to create this special type of job template.

To create a new job template click the  button, which opens the **Create Job Templates** window.



NEW JOB TEMPLATE

DETAILS | COMPLETED JOBS | PERMISSIONS | NOTIFICATIONS

*NAME: DESCRIPTION: *JOB TYPE: Scan

*INVENTORY: *PROJECT: Default *PLAYBOOK: Default

*MACHINE CREDENTIAL: Demo Credential CLOUD CREDENTIAL: FORKS: 0

☐ Prompt on launch

LIMIT: *VERBOSITY: 0 (Normal) JOB TAGS:

☐ Prompt on launch

OPTIONS

☐ Enable Privilege Escalation ☐ Allow Provisioning Callbacks

LABELS:

EXTRA VARIABLES: YAML JSON

1

☐ Prompt on launch

CANCEL SAVE

Enter values for the following fields and select **Save**.

- **Name:** Enter a name appropriate for this inventory. (Required.)
- **Description:** Enter an arbitrary description as appropriate.
- **Job Type:** Jobs can be of type Run, Check, or **Scan**. (Required.)
- **Inventory:** Select the inventory containing the hosts you want this job to manage. (Required.)

- **Project:** Select the project containing the playbook you want this job to execute. Use the default project included with Tower unless you have a specific project to scan. (Required.)
- **Playbook:** Select the playbook to be executed by this job. Use the default playbook included with Tower unless you have a specific playbook to scan. (Required.)
- **Machine Credential:** Select the credential you want the job to use when accessing the remote hosts. Choose the credential containing the username and SSH key or password that Ansible will need to log into the remote hosts.
- **Cloud Credential:** Selecting an optional cloud credential in the job template will pass along the access credentials to the running playbook, allowing provisioning into the cloud without manually passing parameters to the included modules.
- **Forks:** The number of parallel or simultaneous processes to use while executing the playbook.
- **Limit:** Provide a host pattern to further constrain the list of hosts that will be managed or affected by the playbook.
- **Verbosity:** Control the level of output Ansible produces as the playbook executes (this is required).
- **Job Tags:** Provide a comma separated list of tags to run a specific part of a play or task.
- **Enable Privilege Escalation:** If enabled, run this playbook as an administrator. This is the equivalent of passing the `--become` option to the `ansible-playbook` command.
- **Allow Provisioning Callbacks:** Enable a host to call back to Tower via the Tower API and invoke the launch of a job from this job template. Refer to [Provisioning Callbacks](#) for additional information.
- **Labels:** Supply optional labels that describe this job template, such as “dev” or “test”. Labels can be used to group and filter job templates and completed jobs in the Tower display.
 - Labels are created when they are added to the Job Template. Labels are associated to a single Organization using the Project that is provided in the Job Template. Members of the Organization can create labels on a Job Template if they have edit permissions (such as an admin role).
 - Once the Job Template is saved, the labels appear in the Job Templates overview.
 - Click on the “x” beside a label to remove it. When a label is removed, and is no longer associated with a Job or Job Template, the label is permanently deleted from the list of Organization labels.
 - Jobs inherit labels from the Job Template at the time of launch. If a label is deleted from a Job Template, it is also deleted from the Job.

LABELS ?

x test
x scan
x run

JOB TEMPLATES 2 + ADD

NAME ▾ SEARCH Q

NAME ^	DESCRIPTION ▴	ACTIVITY	LABELS	ACTIONS
Demo Job Template				📌 📅 📄 ✎ 🗑
Example	example	● ● ●	x run x scan x test	📌 📅 📄 ✎ 🗑

ITEMS 1-2 OF 2

- **Extra Variables:** Variable definitions and values to be applied to all hosts in this job template. Enter variables using either JSON or YAML syntax. Use the radio button to toggle between the two.

Extra variables can be passed as command line variables to the playbook. This is the “-e” or “--extra-vars” command line parameter for ansible-playbook that is documented in the Ansible documentation at [Passing Variables on the Command Line](#). Example commands might include:

```
scan_file_paths:
- /root/
- /root/
scan_use_checksum: true
scan_use_recursive: true
```

Extra variables can also be provided by key/value pairs using either YAML or JSON. These variables have a maximum value of precedence and overrides other variables specified elsewhere.

NEW JOB TEMPLATE

DETAILS | COMPLETED JOBS | PERMISSIONS | NOTIFICATIONS

*NAME: Test Scan | DESCRIPTION: | *JOB TYPE: Scan

*INVENTORY: Demo Inventory | *PROJECT: Default | *PLAYBOOK: Default

*MACHINE CREDENTIAL: Demo Credential | CLOUD CREDENTIAL: | FORKS: 0

☐ Prompt on launch

LIMIT: | *VERBOSITY: 0 (Normal) | JOB TAGS: scan

☐ Prompt on launch

OPTIONS

☐ Enable Privilege Escalation

☐ Allow Provisioning Callbacks

LABELS: test

EXTRA VARIABLES: ☒ YAML ☐ JSON

```
1 scan_file_paths:
2 - /root/
3 - /root/
4 scan_use_checksum: true
5 scan_use_recursive: true
```

☐ Prompt on launch

CANCEL SAVE

Note: You cannot assign a new inventory at the time of launch to a scan job. Scan jobs must be tied to a fixed inventory.

Note: You cannot change the Job Type at the time of launch to or from the type of “scan”. The `ask_job_type_on_launch` option only enables you to toggle “run” versus “check” at launch time.

12.3.1 Supported OSES for Scan Job Templates

The following operating systems are supported for Scan Jobs:

- Red Hat Enterprise Linux 5, 6, & 7
- CentOS 5, 6, & 7
- Ubuntu 12.04, 14.04, 16.04
- OEL 6 & 7
- SLES 11 & 12
- Debian 6 & 7
- Fedora 22, 23, 24
- Amazon Linux 2016.03

Note that some of these operating systems may require initial configuration in order to be able to run python and/or have access to the python packages that the scan modules depend on.

Pre-scan Setup

The following are examples of playbooks that configure certain distributions so that scan jobs can be run against them.

Bootstrap Ubuntu (16.04)

```
---
- name: Get Ubuntu 15, 16, and on ready
  hosts: all
  sudo: yes
  gather_facts: no

  tasks:

- name: install python-simplejson
  raw: sudo apt-get -y update
  raw: sudo apt-get -y install python-simplejson
  raw: sudo apt-get install python-apt
```

Bootstrap Fedora (23, 24)

```
---
- name: Get Fedora ready
  hosts: all
  sudo: yes
  gather_facts: no

  tasks:

- name: install python-simplejson
  raw: sudo dnf -y update
  raw: sudo dnf -y install python-simplejson
  raw: sudo dnf -y install rpm-python
```

CentOS 5 or Red Hat Enterprise Linux 5 may also need the simplejson package installed.

12.3.2 Launching a Scan Job Template

Job Templates

3

+ Add


NAME

SEARCH

NAME	DESCRIPTION	ACTIVITY	LABELS	ACTIONS
Demo Job Template				<div><div></div><div></div><div></div><div></div><div></div></div>
Hello World				<div><div></div><div></div><div></div><div></div><div></div></div>
Hello World 2 (Scan)		<div></div>		<div><div></div><div></div><div></div><div></div><div></div></div>

Items 1-3 of 3

You can **Launch**, **Schedule**, **Edit**, **Delete**, or **Copy** the scan job template using the buttons to the right.

Click on the  button. Enter any necessary credentials, passwords, passphrases, etc. that were setup for this job template.

The Jobs page shows details of all the tasks and events for that playbook run.

TOWER

PROJECTS

INVENTORIES

JOB TEMPLATES

JOB

admin

JOBS / 5 - HELLO WORLD 2 (SCAN)

RESULTS

STATUS

Successful

TEMPLATE

[Hello World 2 \(Scan\)](#)

STARTED

6/27/2016 10:32:36 PM

FINISHED

6/27/2016 10:32:45 PM

LAUNCHED BY

admin

ELAPSED

00:00:09

INVENTORY

Demo Inventory

MACHINE CREDENTIAL

Demo Credential

VERBOSITY

Default

EXTRA VARIABLES

1

STANDARD OUT

SSH password:

PLAY [all] *****

TASK [setup] *****

ok: [localhost]

TASK [scan_packages] *****

ok: [localhost]

TASK [scan_services] *****

ok: [localhost]

TASK [scan_files] *****

skipping: [localhost]

PLAY RECAP *****

localhost : ok=3 changed=0 unreachable=0 failed=0

DETAILS

1 Please select from a play below to view its associated tasks.

PLAY NAME

ALL FAILED

PLAYS

STARTED

ELAPSED

all

22:32:41

00:00:04

2 Please select a task below to view its associated hosts

TASK NAME

ALL FAILED

TASKS

STARTED

ELAPSED

HOST STATUS

setup

22:32:41

00:00:01

1

scan_packages

22:32:42

00:00:01

1

scan_services

22:32:44

00:00:01

1

scan_files

22:32:45

00:00:00

1

3 Please select a host below to view associated task details.

HOST NAME

ALL FAILED

HOSTS

ITEM

MESSAGE

localhost

EVENT SUMMARY

4 Please select a host below to view a summary of all associated tasks.

HOST NAME

ALL FAILED

HOSTS


COMPLETED TASKS

localhost

3 1

12.3.3 Scheduling a Scan Job Template



To access scheduling for your scan job, click the  button (most easily accessible from the **Job Templates** navigational link).


HELLO WORLD 2 (SCAN) | SCHEDULES 0 + ADD

PLEASE ADD ITEMS TO THIS LIST

Click the + ADD button to add a schedule.

DAILY SCAN ✕

* NAME
Daily Scan

* START DATE (MM/DD/YYYY)
 06/27/2016

* START TIME (HH24:MM:SS)
22 : 45 : 00


* LOCAL TIME ZONE
America/New_York

* REPEAT FREQUENCY
Day

FREQUENCY DETAILS

* EVERY
2 DAYS

* END
On Date

* END DATE (MM/DD/YYYY)
 7/19/2016

* END TIME (HH24:MM:SS)
23 : 45 : 00

SCHEDULE DESCRIPTION

every 2 days until July 19, 2016
OCCURRENCES (Limited to first 10) DATE FORMAT ☒ LOCAL TIME ☐ UTC

06/27/2016 22:45:00 EDT
06/29/2016 22:45:00 EDT
07/01/2016 22:45:00 EDT
07/03/2016 22:45:00 EDT
07/05/2016 22:45:00 EDT
07/07/2016 22:45:00 EDT
07/09/2016 22:45:00 EDT
07/11/2016 22:45:00 EDT
07/13/2016 22:45:00 EDT
07/15/2016 22:45:00 EDT

EXTRA VARIABLES ☒ YAML ☐ JSON

1

CANCEL SAVE

Enter the appropriate details into the following fields and select **Save**:

- Name (required)
- Start Date (required)
- Start Time (required)
- Local Time Zone (the entered Start Time should be in this timezone)
- Repeat Frequency (the appropriate options display as the update frequency is modified.)

The **Details** tab displays a description of the schedule and a list of the scheduled occurrences in the selected Local Time Zone.

Note: Jobs are scheduled in UTC. Repeating jobs that runs at a specific time of day may move relative to a local timezone when Daylight Saving Time shifts occur.

NAME	FIRST RUN	NEXT RUN	FINAL RUN	ACTIONS
ON Daily Scan	6/27/2016 10:45:00 PM	6/27/2016 10:45:00 PM	7/19/2016 11:45:00 PM	[edit] [delete]

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

There are several actions available for schedules, under the Actions column:

- Edit Schedule
- Delete schedule

12.3.4 Custom Scan Job Templates

Custom scan jobs are normal scan job templates which use a custom scan playbook with customized fact modules. Additional Ansible fact modules can be included easily via custom scan playbooks.

Fact Scan Playbooks

The default scan job playbook bundled with Tower, `scan_facts.yml`, contains invocations of three fact scan modules - packages, services, and files, along with Ansible's standard fact gathering. The `scan_facts.yml` playbook file looks like the following:

```
- hosts: all
  vars:
    scan_use_checksum: false
    scan_use_recursive: false
  tasks:
    - scan_packages:
    - scan_services:
    - scan_files:
      paths: '{{ scan_file_paths }}'
      get_checksum: '{{ scan_use_checksum }}'
      recursive: '{{ scan_use_recursive }}'
      when: scan_file_paths is defined
```

The `scan_files` fact module is the only module that accepts parameters, passed via `extra_vars` on the scan job template.

```
scan_file_paths: '/tmp/'
scan_use_checksum: true
scan_use_recursive: true
```

- The `scan_file_paths` parameter may have multiple settings (such as `/tmp/` or `/var/log`).
- The `scan_use_checksum` and `scan_use_recursive` parameters may also be set to false or omitted. An omission is the same as a false setting.

Custom Fact Scans

A playbook for a custom fact scan is similar to the example of the Fact Scan Playbook above. As an example, a playbook that only uses a custom `scan_foo` Ansible fact module would look like this:

scan_custom.yml:

```
- hosts: all
  gather_facts: false
  tasks:
    - scan_foo:
```

scan_foo.py:

```
def main():
    module = AnsibleModule(
        argument_spec = dict()

    foo = [
        {
            "hello": "world"
        },
        {
            "foo": "bar"
        }
    ]
    results = dict(ansible_facts=dict(foo=foo))
    module.exit_json(**results)

main()
```

To use a custom fact module, ensure that it lives in the `/library/` subdirectory of the Ansible project used in the scan job template. This fact scan module is very simple, returning a hard-coded set of facts:

```
[
  {
    "hello": "world"
  },
  {
    "foo": "bar"
  }
]
```

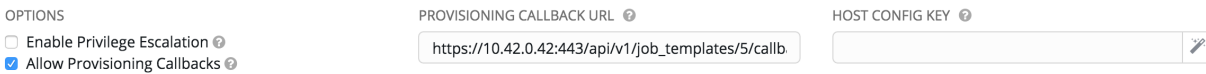
12.4 Provisioning Callbacks

Provisioning callbacks are a feature of Tower that allow a host to initiate a playbook run against itself, rather than waiting for a user to launch a job to manage the host from the tower console. Please note that provisioning callbacks are *only* used to run playbooks on the calling host. Provisioning callbacks are meant for cloud bursting, ie: new instances with a need for client to server communication for configuration (such as transmitting an authorization key), not to run a job against another host. This provides for automatically configuring a system after it has been provisioned by another system (such as AWS auto-scaling, or a OS provisioning system like kickstart or preseed) or for launching a job programmatically without invoking the Tower API directly. The Job Template launched only runs against the host requesting the provisioning.

Frequently this would be accessed via a firstboot type script, or from cron.


To enable callbacks, check the *Allow Provisioning Callbacks* checkbox in the Job Template. This displays the **Provisioning Callback URL** for this job template.

Note: If you intend to use Tower’s provisioning callback feature with a dynamic inventory, Update on Launch should be set for the inventory group used in the Job Template.




Options: ☐ Enable Privilege Escalation, ☒ Allow Provisioning Callbacks

Provisioning Callback URL:

Host Config Key: 

Callbacks also require a Host Config Key, to ensure that foreign hosts with the URL cannot request configuration.

Click the  button to create a unique host key for this callback, or enter your own key. The host key may be reused across multiple hosts to apply this job template against multiple hosts. Should you wish to control what hosts are able to request configuration, the key may be changed at any time.

To callback manually via REST, look at the callback URL in the UI, which is of the form:

```
http://<Tower server name>/api/v1/job_templates/1/callback/
```

The ‘1’ in this sample URL is the job template ID in Tower.

The request from the host must be a POST. Here is an example using curl (all on a single line):

```
root@localhost:~$ curl --data "host_config_key=5a8ec154832b780b9bdef1061764ae5a" \
    http://<Tower server name>/api/v1/job_templates/1/callback/
```

The requesting host must be defined in your inventory for the callback to succeed. If Tower fails to locate the host either by name or IP address in one of your defined inventories, the request is denied. When running a Job Template in this way, the host initiating the playbook run against itself must be in the inventory. If the host is missing from the inventory, the Job Template will fail with a “No Hosts Matched” type error message.

Note: If your host is not in inventory and Update on Launch is set for the inventory group, Tower attempts to update cloud based inventory source before running the callback.

Successful requests result in an entry on the Jobs tab, where the results and history can be viewed.

While the callback can be accessed via REST, the suggested method of using the callback is to use one of the example scripts that ships with Tower - /usr/share/awx/request_tower_configuration.sh (Linux/UNIX) or /usr/share/awx/request_tower_configuration.ps1 (Windows). Usage is described in the source code

of the file. This script is intelligent in that it knows how to retry commands and is therefore a more robust way to use callbacks than a simple curl request. As written, the script retries once per minute for up to ten minutes.

Note: Please note that this is an example script. You should edit this script if you need more dynamic behavior when detecting failure scenarios, as any non-200 error code may not be a transient error requiring retry.

Most likely you will use callbacks with dynamic inventory in Tower, such as pulling cloud inventory from one of the supported cloud providers. In these cases, along with setting *Update On Launch*, be sure to configure an inventory cache timeout for the inventory source, to avoid hammering of your Cloud's API endpoints. Since the `request_tower_configuration.sh` script polls once per minute for up to ten minutes, a suggested cache invalidation time for inventory (configured on the inventory source itself) would be one or two minutes.

While we recommend against running the `request_tower_configuration.sh` script from a cron job, a suggested cron interval would be perhaps every 30 minutes. Repeated configuration can be easily handled by scheduling in Tower, so the primary use of callbacks by most users is to enable a base image that is bootstrapped into the latest configuration upon coming online. To do so, running at first boot is a better practice. First boot scripts are just simple init scripts that typically self-delete, so you would set up an init script that called a copy of the `request_tower_configuration` script and make that into an autoscaling image.

12.4.1 Passing Extra Variables to Provisioning Callbacks

Just as you can pass `extra_vars` in a regular Job Template, you can also pass them to provisioning callbacks. To pass `extra_vars`, the data sent must be part of the body of the POST request as application/json (as the content type). Use the following JSON format as an example when adding your own `extra_vars` to be passed:

```
'{"extra_vars": {"variable1": "value1", "variable2": "value2", ...}}'
```

(Added in Ansible Tower version 2.2.0.)

You can also pass extra variables to the Job Template call using `curl`, such as is shown in the following example:

```
root@localhost:~$ curl -f -H 'Content-Type: application/json' -XPOST \
    -d '{"host_config_key": "5a8ec154832b780b9bdef1061764ae5a", "extra_
    ↪vars": {"foo": "bar"}}' \
    http://<Tower server name>/api/v1/job_templates/1/callback
```

For more information, refer to [Launching Jobs with Curl](#).

12.5 Launching Jobs

A major benefit of Ansible Tower is the push-button deployment of Ansible playbooks. You can easily configure a template within Tower to store all parameters you would normally pass to the ansible-playbook on the command line—not just the playbooks, but the inventory, credentials, extra variables, and all options and settings you can specify on the command line.

Easier deployments drive consistency, by running your playbooks the same way each time, and allow you to delegate responsibilities—even users who aren't Ansible experts can run Tower playbooks written by others.

To launch a job template, click the  button.

A job may require additional information to run. The following data may be requested at launch:

- Credentials that were setup

- Passwords or passphrases that have been set to **Ask**
- A survey, if one has been configured for the job templates
- Extra variables, if requested by the job template

Here is an example job launch that prompts for Job Tags, and runs the example survey created in [Surveys](#).

LAUNCH JOB | HELLO WORLD

OTHER PROMPTS

SURVEY

JOB TAGS

INVENTORY
Demo Inventory

CREDENTIAL
Demo Credential

CANCEL

NEXT

LAUNCH JOB | HELLO WORLD

OTHER PROMPTS

SURVEY

* WHICH GROUP(S) SHOULD INCLUDE THIS USER?

Enter groups, one per line.

INVENTORY
Demo Inventory

CREDENTIAL
Demo Credential

CANCEL

LAUNCH

Along with any extra variables set in the job template and survey, Tower automatically adds the following variables to the job environment:

- `tower_job_id`: The Job ID for this job run
- `tower_job_launch_type`: One of *manual*, *callback*, or *scheduled* to indicate how the job was started
- `tower_job_template_id`: The Job Template ID that this job run uses

- `tower_job_template_name`: The Job Template name that this job uses
- `tower_user_id`: The user ID of the Tower user that started this job. This is not available for callback or scheduled jobs.
- `tower_user_name`: The user name of the Tower user that started this job. This is not available for callback or scheduled jobs.

Upon launch, Tower automatically redirects the web browser to the Job Status page for this job under the **Jobs** tab.

12.6 Scheduling



Launching job templates may also be scheduled via the button. Clicking this button opens the **Schedules** page.

Job Templates

3

+ Add

NAME

SEARCH

Q

NAME	DESCRIPTION	ACTIVITY	LABELS	ACTIONS
Demo Job Template				<div><div></div><div></div><div></div><div></div><div></div></div>
Hello World				<div><div></div><div></div><div></div><div></div><div></div></div>
Hello World 2 (Scan)		<div><div></div><div></div><div></div></div>		<div><div></div><div></div><div></div><div></div><div></div></div>

ITEMS 1-3 OF 3

This page displays a list of the schedules that are currently available for the selected **Job Template**. The schedule list may be sorted and searched by **Name**.

The list of schedules includes:

- **Name**: Clicking the schedule name opens the **Edit Schedule** dialog
- **First Run**: The first scheduled run of this task
- **Next Run**: The next scheduled run of this task
- **Final Run**: If the task has an end date, this is the last run of the task

Buttons located in the upper right corner of the **Schedules** screen provide the following actions:

- Create a new schedule
- Refresh this view
- View Activity Stream

12.6.1 Add a new schedule

+ ADD

To create a new schedule click the button.

DAILY SCAN

* NAME

* START DATE (MM/DD/YYYY)

* START TIME (HH24:MM:SS)

 : :

* LOCAL TIME ZONE

* REPEAT FREQUENCY

FREQUENCY DETAILS

* EVERY

 DAYS

* END

* END DATE (MM/DD/YYYY)

* END TIME (HH24:MM:SS)

 : :

SCHEDULE DESCRIPTION

every 2 days until July 19, 2016
OCCURRENCES (Limited to first 10) DATE FORMAT: ☒ LOCAL TIME ☐ UTC
06/27/2016 22:45:00 EDT
06/29/2016 22:45:00 EDT
07/01/2016 22:45:00 EDT
07/03/2016 22:45:00 EDT
07/05/2016 22:45:00 EDT
07/07/2016 22:45:00 EDT
07/09/2016 22:45:00 EDT
07/11/2016 22:45:00 EDT
07/13/2016 22:45:00 EDT
07/15/2016 22:45:00 EDT

EXTRA VARIABLES ☒ YAML ☐ JSON

1

CANCEL

SAVE

Enter the appropriate details into the following fields and select **Save**:

- Name (required)
- Start Date (required)
- Start Time (required)
- Local Time Zone (the entered Start Time should be in this timezone)
- Repeat Frequency (the appropriate options display as the update frequency is modified.)

HELLO WORLD 2 (SCAN) | SCHEDULES 1

+ ADD

NAME	SEARCH	
NAME ^	FIRST RUN ^	NEXT RUN ^
FINAL RUN ^	ACTIONS	
ON Daily Scan	6/27/2016 10:45:00 PM	6/27/2016 10:45:00 PM
7/19/2016 11:45:00 PM		

ITEMS 1-1 OF 1



The **Details** tab displays a description of the schedule and a list of the scheduled occurrences in the selected Local Time Zone.

Note: Jobs are scheduled in UTC. Repeating jobs that runs at a specific time of day may move relative to a local timezone when Daylight Saving Time shifts occur.

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

There are several actions available for schedules, under the Actions column:

- Edit Schedule
- Delete schedule

HELLO WORLD 2 (SCAN) SCHEDULES 1				+ ADD
NAME	SEARCH			Q
NAME ^	FIRST RUN ↕	NEXT RUN ↕	FINAL RUN ↕	ACTIONS
ON Daily Scan	6/27/2016 10:45:00 PM	6/27/2016 10:45:00 PM	7/19/2016 11:45:00 PM	 
ITEMS 1-1 OF 1				

JOBS

A **job** is an instance of Tower launching an Ansible playbook against an inventory of hosts.

The Jobs link displays a list of jobs and their status—shown as completed successfully or failed, or as an active (running) job.













JOBS

JOBS

SCHEDULES

STATUS

FILTER

ID	NAME	TYPE	FINISHED	ACTIONS
6	Hello World 2 (Scan)	Playbook Run	6/27/2016 10:45:37 PM	 
5	Hello World 2 (Scan)	Playbook Run	6/27/2016 10:32:45 PM	 
4	CMS Web group (Database Servers)	Inventory Sync	6/27/2016 9:39:21 PM	 
3	Hello World 2 (Scan)	Playbook Run	6/27/2016 9:35:13 PM	 
2	Demo Project	SCM Update	6/27/2016 7:40:21 PM	 
1	Demo Project	SCM Update	6/27/2016 7:31:52 PM	 

ITEMS 1-6 OF 6

- Jobs can be searched by **Job Failed?**, job **ID**, or **Status**.
- Jobs can be filtered by job **ID** or **Finished**.

From this screen, you can relaunch jobs, remove jobs, or view the standard output of a particular job. Clicking on a **Name** for an **SCM Update** job opens the **Job Results** screen, while clicking on a manually created **Playbook Run** job takes you to the **Job Status** page for this job (also accessible after launching a job from the **Job Templates** link in the main Tower navigational menu).

The **Job Results** window displays information about jobs of type **Inventory Sync**, **Playbook Run**, and **SCM Update**.

13.1 Job Results - Inventory Sync

The Job Results page for an inventory sync consists of two sections:

- Results
- Standard Out

13.1.1 Results

The **Results** area shows the basic status of the job—*Running*, *Pending*, *Successful*, or *Failed*—and its start time. The buttons in the top right of the Results view allow you to relaunch or delete the job.

The **Results** display includes details on the job execution:

- **Name:** The name of the associated inventory group.
- **Status:** Can be any of *Pending*, *Running*, *Successful*, or *Failed*.
- **License Error:** Only shown for **Inventory Sync** jobs. If this is *True*, the hosts added by the inventory sync caused Tower to exceed the licensed number of managed hosts.
- **Started:** The timestamp of when the job was initiated by Tower.
- **Finished:** The timestamp of when the job was completed.
- **Elapsed:** The total time the job took.
- **Launch Type:** *Manual* or *Scheduled*.
- **Group:** The group being synced.
- **Source:** The type of cloud inventory.
- **Overwrite:** The value of *Overwrite* for this Inventory Sync. Refer to [Inventories](#) for details.
- **Overwrite Vars:** The value of *Overwrite Vars* for this Inventory Sync. Refer to [Inventories](#) for details.

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

13.1.2 Standard Out

The **Standard Out** display shows the full results of running the SCM Update or Inventory Sync playbook. This shows the same information you would see if you ran the Ansible playbook using Ansible from the command line, and can be useful for debugging. The buttons in the top right corner of the Standard Out display allow you to toggle the output as a main view or to download the output.

The screenshot shows the Ansible Tower web interface. At the top, there's a navigation bar with 'TOWER', 'PROJECTS', 'INVENTORIES', 'JOB TEMPLATES', and 'JOBS'. The 'JOBS' tab is selected, and the breadcrumb 'JOBS / CMS WEB GROUP' is visible. The main content area is divided into two panels. The left panel, titled 'RESULTS', contains a table with job details: NAME (CMS Web group), STATUS (Successful), LICENSE ERROR (False), STARTED (6/27/2016 9:39:14 PM), FINISHED (6/27/2016 9:39:21 PM), ELAPSED (6.536 seconds), LAUNCH TYPE (Manual), GROUP (CMS Web group), SOURCE (Custom Script), OVERWRITE (False), and OVERWRITE VARS (False). The right panel, titled 'STANDARD OUT', displays the job's output logs, showing a series of 'INFO' messages for adding various groups to the inventory, such as 'cms web group', 'eights.example.com', 'evens.example.com', 'fives.example.com', 'fours.example.com', and many others.

Copyright © 2016 Red Hat, Inc.

13.2 Job Results - SCM

The Job Results page for an SCM update consists of two sections:

- Results
- Standard Out

13.2.1 Results

The **Results** area shows the basic status of the job—*Running*, *Pending*, *Successful*, or *Failed*—and its start time. The buttons in the top right of the Results view allow you to relaunch or delete the job.

The **Results** display includes details on the job execution:

- **Name:** The name of the associated inventory group.
- **Status:** Can be any of *Pending*, *Running*, *Successful*, or *Failed*.
- **Started:** The timestamp of when the job was initiated by Tower.
- **Finished:** The timestamp of when the job was completed.
- **Elapsed:** The total time the job took.
- **Launch Type:** *Manual* or *Scheduled*.
- **Project:** The name of the project.

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

13.2.2 Standard Out

The **Standard Out** display shows the full results of running the SCM Update or Inventory Sync playbook. This shows the same information you would see if you ran the Ansible playbook using Ansible from the command line, and can be useful for debugging. The buttons in the top right corner of the Standard Out display allow you to toggle the output as a main view or to download the output.

The screenshot shows the Ansible Tower web interface. At the top, there's a navigation bar with 'TOWER', 'PROJECTS', 'INVENTORIES', 'JOB TEMPLATES', and 'JOBS'. On the right, there's a user profile 'admin' and several icons. Below the navigation bar, the breadcrumb 'JOBS / DEMO PROJECT' is visible. The main content area is split into two panels. The left panel, titled 'RESULTS', contains a table with the following data:

RESULTS	
NAME	Demo Project
STATUS	Successful
STARTED	6/27/2016 7:40:18 PM
FINISHED	6/27/2016 7:40:21 PM
ELAPSED	2.961 seconds
LAUNCH TYPE	Manual
PROJECT	Demo Project

The right panel, titled 'STANDARD OUT', displays the raw output of the playbook run. It starts with 'Using /etc/ansible/ansible.cfg as config file' and 'PLAY [all]'. The output shows several tasks, all of which were skipped because a conditional check failed. The tasks are:

- TASK [delete project directory before update]
- TASK [update project using git and accept hostkey]
- TASK [update project using git]
- TASK [update project using hg]
- TASK [update project using svn]
- TASK [update project using svn with auth]

Each task output shows 'skipping: [localhost] => {"changed": false, "skip_reason": "Conditional check failed", "skipped": true}'. The 'git' task also shows an 'ok' status with 'after' and 'before' hashes. At the bottom right of the screenshot, there is a copyright notice: 'Copyright © 2016 Red Hat, Inc.'

13.3 Job Results - Playbook Run

The Job Results page for a playbook run consists of several sections:

- Results
- Standard Out
- Details
- Event Summary

13.3.1 Results

The **Results** area shows the basic status of the job—*Running*, *Pending*, *Successful*, or *Failed*—and its start time. The buttons in the top right of the Results view allow you to relaunch or delete the job.

The **Results** display includes details on the job execution:

- **Status:** Can be any of *Pending*, *Running*, *Successful*, or *Failed*.
- **Template:** The name of the job template from which this job was launched.
- **Started:** The timestamp of when the job was initiated by Tower.
- **Finished:** The timestamp of when the job was completed.
- **Elapsed:** The total time the job took.
- **Launch By:** The name of the user, job, or scheduled scan job which launched this job.
- **Inventory:** The inventory selected to run this job against.
- **Machine Credential:** The name of the credential used in this job.
- **Verbosity:** The level of verbosity set when creating the job template.
- **Extra Variables:** Any extra variables passed when creating the job template are displayed here.

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

13.3.2 Standard Out

The **Standard Out** display shows the full results of running the SCM Update or Inventory Sync playbook. This shows the same information you would see if you ran the Ansible playbook using Ansible from the command line, and can be useful for debugging. The buttons in the top right corner of the Standard Out display allow you to toggle the output as a main view or to download the output.

TOWER

PROJECTS

INVENTORIES

JOB TEMPLATES

JOB

admin

JOBS

/

6 - HELLO WORLD 2 (SCAN)

RESULTS

STATUS

Successful

TEMPLATE

[Hello World 2 \(Scan\)](#)

STARTED

6/27/2016 10:45:28 PM

FINISHED

6/27/2016 10:45:37 PM

ELAPSED

00:00:08

LAUNCHED BY

[Daily Scan](#)

INVENTORY

[Demo Inventory](#)

MACHINE CREDENTIAL

[Demo Credential](#)

VERBOSITY

Default

EXTRA VARIABLES

1

STANDARD OUT

SSH password:

PLAY [all] *****

TASK [setup] *****

ok: [localhost]

TASK [scan_packages] *****

ok: [localhost]

TASK [scan_services] *****

ok: [localhost]

TASK [scan_files] *****

skipping: [localhost]

PLAY RECAP *****

localhost : ok=3 changed=0 unreachable=0 failed=0

DETAILS

1 Please select from a play below to view its associated tasks.

PLAY NAME

Q

ALL

FAILED

PLAYS

STARTED

ELAPSED

all

22:45:32

00:00:04

2 Please select a task below to view its associated hosts

TASK NAME

Q

ALL

FAILED

TASKS

STARTED

ELAPSED

HOST STATUS

setup

22:45:32

00:00:01

1

scan_packages

22:45:34

00:00:01

1

scan_services

22:45:35

00:00:01

1

scan_files

22:45:36

00:00:00

1

3 Please select a host below to view associated task details.

HOST NAME

Q

ALL

FAILED

HOSTS

ITEM

MESSAGE

localhost

EVENT SUMMARY

4 Please select a host below to view a summary of all associated tasks.

HOST NAME

Q

ALL

FAILED

HOSTS

COMPLETED TASKS

localhost

3 1

13.3.3 Job Details View

DETAILS

1 Please select from a play below to view its associated tasks.

PLAY NAME

ALL FAILED

PLAYS	STARTED	ELAPSED
 Hello World Sample	09:12:42	00:00:01

2 Please select a task below to view its associated hosts

TASK NAME

ALL FAILED

TASKS	STARTED	ELAPSED	HOST STATUS
 setup	09:12:42	00:00:01	1
 Hello Message	09:12:44	00:00:00	1

3 Please select a host below to view associated task details.

HOST NAME

ALL FAILED

No matching host events

Plays

The **Plays** area shows the plays that were run as part of this playbook. The displayed plays can be filtered by **Play Name**, and can be limited to show only failed plays (using the **ALL/FAIL** view toggle).

For each play, Tower shows the **Play Name**, start time for the play, the elapsed time of the play, the play **Name**, and whether the play succeeded or failed (indicated by the status dot to the left of the **pPlay Name**). Clicking on a specific play filters the **Tasks** and **Host Events** area to only display tasks and hosts relative to that selected play.

1 Please select from a play below to view its associated tasks.

PLAY NAME	Q	ALL	FAILED
PLAYS	STARTED	ELAPSED	
● Hello World Sample	09:12:42	00:00:01	

Tasks

The **Tasks** area shows the tasks run as part of plays in the playbook. The displayed tasks can be filtered by **Task Name**, and can be limited to only failed tasks.

For each task, Tower shows the task **Name**, the start time for the task, the elapsed time of the task, whether the task succeeded or failed. Clicking on a specific task filters the **Host Events** area to only display hosts relative to that task.

2 Please select a task below to view its associated hosts

TASK NAME	Q	ALL	FAILED
TASKS	STARTED	ELAPSED	HOST STATUS
● setup	09:12:42	00:00:01	1
● Hello Message	09:12:44	00:00:00	1

Host Events

The **Host Events** area shows hosts affected by the selected play and task. For each host, Tower shows the host's status, its name, and any **Item** or **Message** set by that task.

3 Please select a host below to view associated task details.

HOSTS	ITEM	MESSAGE
● localhost		

Clicking on the linked hostname brings up the **Host Event** dialog for that host and task.

The **Host Event** dialog shows the events for this host and the selected play and task:

- the **Host**
- the **Status**
- a unique **ID**
- a **Created** time stamp
- the name of the **Play**
- the name of the **Task**
- if applicable, the Ansible **Module** for the task, and any *arguments* for that module
- the **Results** of the task

There is also a **JSON** tab which displays the result in JSON format.

HOST EVENT

EVENT		RESULTS	
HOST	localhost	MSG	Hello World!
STATUS	● ok	_ANSIBLE_VER	
ID	6	BOSE_ALWAYS	true
CREATED	2016-07-07T13:12:44.215Z	_ANSIBLE_NO	
PLAY	Hello World Sample	_LOG	false
TASK	Hello Message		
MODULE	No result found		

There is also a **JSON** tab which displays the result in JSON format.

The screenshot shows a modal window titled "HOST EVENT" with a close button in the top right corner. Below the title are two tabs: "DETAILS" (selected) and "JSON". The main content area displays a JSON object representing an event, with line numbers 1 through 18 on the left. The JSON data is as follows:

```

1 {
2   "id": 6,
3   "created": "2016-07-07T13:12:44.215Z",
4   "modified": "2016-07-07T13:12:44.215Z",
5   "job": 4,
6   "event": "runner_on_ok",
7   "counter": 6,
8   "event_display": "Host OK",
9   "event_level": 3,
10  "failed": false,
11  "changed": false,
12  "host": 1,
13  "host_name": "localhost",
14  "parent": 4,
15  "play": "Hello World Sample",
16  "task": "Hello Message",
17  "role": ""
18 }

```

At the bottom of the modal are three buttons: "PREV HOST", "NEXT HOST", and "CLOSE".

13.3.4 Event Summary

The **Event Summary** area shows a summary of events for all hosts affected by this playbook as well as the **Host Status Summary**.

Hosts can be filtered by their hostname, and can be limited to showing only changed, failed, OK, and unreachable hosts.

EVENT SUMMARY ▾

4 Please select a host below to view a summary of all associated tasks.

HOST NAME

Q


ALL

FAILED

HOSTS	COMPLETED TASKS
localhost	2

HOST STATUS SUMMARY

● OK: 100%



For each host, the **Event Summary** area shows the hostname and the number of completed tasks for that host, sorted by status.

EVENT SUMMARY ▾

4 Please select a host below to view a summary of all associated tasks.

HOSTS	COMPLETED TASKS
localhost	2

Clicking on the hostname brings up a **Host Events** dialog, displaying all tasks that affected that host.

This dialog can be filtered by the status of the tasks, as well as by the hostname.

For each event, Tower displays the status, the play name, and the task name.

Host status can be one of the following:

- **Changed:** the playbook task actually executed. Since Ansible tasks should be written to be idempotent, tasks may exit successfully without executing anything on the host. In these cases, the task would return Ok, but not Changed.
- **Failure:** the task failed. Further playbook execution was stopped for this host.
- **OK:** the playbook task returned “Ok”.
- **Unreachable:** the host was unreachable from the network or had another fatal error associated with it.
- **Skipped:** the playbook task was skipped because no change was necessary for the host to reach the target state.

HOST EVENTS | LOCALHOST

STATUS	PLAY	TASK
ok	Hello World Sample	setup
ok	Hello World Sample	Hello Message

OK

The **Host Summary** area shows a graph summarizing the status of all hosts affected by this playbook run.

HOST STATUS SUMMARY

● OK: 100%



13.4 Job Concurrency

Tower limits the number of simultaneous jobs that can run based on the amount of physical memory and the complexity of the playbook.

If the “Update on Launch” setting is checked, job templates that rely on the inventory or project *also* trigger an update on them if it is within the cache timeout. If multiple jobs are launched within the cache timeout that depend on the same project or inventory, only one of those project or inventory updates is created (instead of one for each job that relies on it).

If you are having trouble, try setting the cache timeout on the project or inventory source to 60 seconds.

The restriction related to the amount of RAM on your Tower server and the size of your inventory equates to the total number of machines from which facts can be gathered and stored in memory. The algorithm used is:

```
50 + ((total memory in megabytes) / 1024) - 2) * 75
```

With 50 as the baseline.

Each job that runs is:

```
(number of forks on the job) * 10
```

Which defaults to 50 if the limit is set Tower’s default value of 0.

Forks determine the default number of parallel processes to spawn when communicating with remote hosts. The fork number is automatically limited to the number of possible hosts, so this is really a limit of how much network and CPU load you can handle. Many users set this to 50, while others set it to 500 or more. If you have a large number of hosts, higher values will make actions across all of those hosts complete faster. You can edit the `ansible.cfg` file to change this value.

The Ansible fork number default is extremely conservative and is set to five (5). When you do not pass a forks value in Tower (leaving it as 0), Ansible uses 5 forks (the default). If you set your forks value to one (1) in Tower, Ansible uses the value entered and one (1) fork is created. Non-zero inputs are used as instructed.

As an example, if you have a job with 0 forks (the Tower default) on a system with 2 GB of memory, your algorithm would look like the following:

```
50 + ((2048 / 1024) - 2) * 75 = 50
```

If you have a job with 0 forks (the Tower default) on a system with 4 GB of memory then you can run four (4) tasks simultaneously which includes callbacks.

```
50 + ((4096 / 1024) - 2) * 75 = 200
```

This can be changed by setting a value in the Tower settings file (`/etc/tower/settings.py`) such as:

```
SYSTEM_TASK_CAPACITY = 300
```

If you want to override the setting, use caution, as you may run out of RAM if you set this value too high. You can determine what the calculated setting is by reviewing `/var/log/tower/task_system.log` and looking for a line similar to:

Running Nodes: []; Capacity: 50; Running Impact: 0; Remaining Capacity: 50

The Capacity: 50 is the current calculated setting.

As long as you have the capacity to do so, Tower attempts to reorder and run the most number of jobs possible. There are some blockers and exceptions worth noting, however.

- A Job Template will block the same instance of another Job Template launched with the same inventory (this is a change in behavior – prior to Ansible Tower 3.0, a Job Template would block the same instance of another Job Template).
- A project update will block for another project requiring the same update.
- Job Templates which launch via provisioning callbacks can run, just not as an instance on the same host. This allows running two (2) templates on the same inventory. However, if the inventory requires an update, they will not run. Callbacks are special types of job templates which receive “push requests” from the host to the inventory. They run on one host only and can run parallel with other jobs as long as they are different callbacks and different hosts.
- System Jobs can only run one at a time. They block all other jobs and must be run on their own to avoid conflict. System jobs will finish behind jobs schedule ahead of them, but will finish ahead of those jobs scheduled behind it.
- Ad hoc jobs are blocked by any inventory updates running against the inventory for that ad hoc job as specified.

NOTIFICATIONS

A **Notifier** is an instance of a **Notification** type (Email, Slack, Webhook, etc.) with a name, description, and a defined configuration.

For example:

- A username, password, server, and recipients are needed for an Email notifier
- The token and a list of channels are needed for a Slack notifier
- The URL and Headers are needed for a Webhook notifier

A Notification is a manifestation of the notifier; for example, when a job fails, a notification is sent using the configuration defined by the Notifier.

At a high level, the typical flow for the notification system works as follows:

- A user creates a notifier to the Tower REST API at the `/api/v1/notifiers` endpoint (either through the API or through the Tower UI).
- A user assigns the notifier to any of the various objects that support it (all variants of job templates as well as organizations and projects) and at the appropriate trigger level for which they want the notification (error, success, or any). For example a user may wish to assign a particular Notifier to trigger when Job Template 1 fails. In which case, they will associate the notifier with the job template at `/api/v1/job_templates/n/notifiers_error` API endpoint.

14.1 Notifier Hierarchy

Notifiers assigned at certain levels will inherit notifiers defined on parent objects as such:

- Job Templates will use notifiers defined on it as well as inheriting notifiers from the Project used by the Job Template and from the Organization that it is listed under (via the Project).
- Project Updates will use notifiers defined on the project and will inherit notifiers from the Organization associated with it
- Inventory Updates will use notifiers defined on the Organization that it is listed under
- Ad-hoc commands will use notifiers defined on the Organization that the inventory is associated with


14.2 Workflow


When a job succeeds or fails, the error or success handler will pull a list of relevant notifiers using the procedure defined above. It will then create a Notification object for each one containing relevant details about the job and then sends it to the destination (email addresses, slack channel(s), sms numbers, etc). These Notification objects are available as

related resources on job types (jobs, inventory updates, project updates), and also at `/api/v1/notifications`. You may also see what notifications have been sent from a notifier by examining its related resources.

If a notification fails, it will not impact the job associated to it or cause it to fail. The status of the notification can be viewed at its detail endpoint (`/api/v1/notifications/<n>`).

14.3 Create a Notification Template

To create a Notification Template, click the  button and select **Notifications**.



14.4 Notification Types

Topics:

- *Email*
- *Slack*
- *Twilio*
- *PagerDuty*
- *HipChat*
- *Webhook*
- *IRC*

Each of these have their own configuration and behavioral semantics and testing them may need to be approached in different ways. The following sections will give as much detail as possible.

14.4.1 Email

The email notification type supports a wide variety of SMTP servers and has support for TLS/SSL connections.

You must provide the following details to setup an email notification: - Username - Host - Sender email - Recipient list - Password - Port

NEW NOTIFICATION TEMPLATE

* NAME: Tell Me

DESCRIPTION: notification test

* ORGANIZATION: Honey Dog, Inc.

* TYPE: Email

TYPE DETAILS

* USERNAME: luckydog

* HOST: honeydog.com

* SENDER EMAIL: info@honeydog.com

* RECIPIENT LIST: engineering@honeydog.com, release-managers@honeydog.com

* PASSWORD: SHOW *****

* PORT: 1

OPTIONS:

☐ Use TLS

☐ Use SSL

CANCEL SAVE

Caution: TLS and SSL connections are mutually exclusive and should not be selected at the same time. Be sure to only select one—checking both causes the notification to silently fail.

14.4.2 Slack

Slack, a collaborative team communication and messaging tool, is pretty easy to configure.

You must supply the following to setup Slack notifications:

- A token (which you can obtain from creating a bot in the integrations settings for the Slack team at <https://api.slack.com/bot-users>)
- Destination channel(s)

You must also invite the notification bot to join the channel(s) in question. Note that private messages are not supported.

NEW NOTIFICATION TEMPLATE

* NAME: Tell Me

DESCRIPTION: notification test

* ORGANIZATION: Honey Dog, Inc.

* TYPE: Slack

TYPE DETAILS

* DESTINATION CHANNELS: #engineering, #helpdesk, #support

* TOKEN: SHOW *****

CANCEL SAVE

14.4.3 Twilio

Twilio service is an Voice and SMS automation service. Once you are signed in, you must create a phone number from which the message will be sent. You can then define a “Messaging Service” under Programmable SMS and associate the number you created before with it.

Note that you may need to verify this number or some other information before you are allowed to use it to send to any numbers. The Messaging Service does not need a status callback URL nor does it need the ability to Process inbound messages.

Under your individual (or sub) account settings, you will have API credentials. Twilio uses two credentials to determine which account an API request is coming from. The “Account SID”, which acts as a username, and the “Auth Token” which acts as a password.

To setup Twilio, provide the following details:

- Account Token
- Source phone number (this is the number associated with the messaging service above and must be given in the form of “+15556667777”)
- Destination phone number (this will be the list of numbers to receive the SMS and should be the 10-digit phone number)
- Account SID

NEW NOTIFICATION TEMPLATE

* NAME: Tell Me

DESCRIPTION: notification test

* ORGANIZATION: Honey Dog, Inc.

* TYPE: Twilio

TYPE DETAILS

* ACCOUNT TOKEN: SHOW [REDACTED]

* SOURCE PHONE NUMBER: 9109876555

* DESTINATION SMS NUMBER: 9109676565

* ACCOUNT SID: Aa54e6dg4345d3t676f60

CANCEL SAVE

14.4.4 PagerDuty

PagerDuty is a fairly straightforward integration. The user must first create an API Key in the pagerduty system (this is the token that is given to Tower) and then create a “Service” which provides an “Integration Key” that will also be given to Tower. The other options of note are:

- API Token: The user must first create an API Key in the PagerDuty system (this is the token that is given to Tower).
- PagerDuty Subdomain: When you sign up for the PagerDuty account, you receive a unique subdomain to communicate with. For instance, if you signed up as “towertest”, the web dashboard will be at `towertest.pagerduty.com` and you will give the Tower API `towertest` as the subdomain (not the full domain).
- API Service/Integration Key

- **Client Identifier:** This will be sent along with the alert content to the pagerduty service to help identify the service that is using the api key/service. This is helpful if multiple integrations are using the same API key and service.

The screenshot shows a 'NEW NOTIFICATION TEMPLATE' form. It has the following fields:

- * NAME:** A text input field containing 'Tell Me'.
- DESCRIPTION:** A text input field containing 'notification test'.
- * ORGANIZATION:** A search input field containing 'Honey Dog, Inc.' with a magnifying glass icon.
- * TYPE:** A dropdown menu with 'Pagerduty' selected.
- TYPE DETAILS:** A section containing:
 - * API TOKEN:** A text input field with a 'SHOW' button.
 - * PAGERDUTY SUBDOMAIN:** A text input field.
 - * API SERVICE/INTEGRATION KEY:** A text input field.
 - * CLIENT IDENTIFIER:** A text input field.

At the bottom right, there are 'CANCEL' and 'SAVE' buttons.

14.4.5 HipChat

There are several ways to integrate with HipChat. The Tower implementation uses HipChat “Integrations”. Currently you can find this at the bottom right of the main HipChat webview. From there, you will select “Build your own Integration”. After creating that, it will list the `auth_token` that needs to be supplied to Tower. Some other relevant details on the fields accepted by Tower for the HipChat notification type:

- **Destination Channels:** Channels which should receive the notification (“engineering” or “#support”, for example).
- **Token:** The token listed after building your own HipChat integration.
- **Label to be shown with notification:** Along with the integration name itself this will put another label on the notification (which could be helpful if multiple services are using the same integration to distinguish them from each other).
- **API URL:** The URL of the Hipchat API service. If you create a team hosted by them it will be something like: `https://team.hipchat.com`. For a self-hosted integration, use a base URL similar to `https://hipchat.yourcompany.com/` and add in appropriate Destination Channels without the # leading them (“engineering” rather than “#engineering”).
- **Notification Color:** This will highlight the message as the given color. If set to something HipChat does not expect, then the notification will generate an error in the given color.
- **Notify Channel:** Selecting this will cause the bot to “notify” channel members. Normally it will just be stuck as a message in the chat channel without triggering anyone’s notifications. This option will notify users of the channel respecting their existing notification settings (browser notification, email fallback, etc.).

NEW NOTIFICATION TEMPLATE

*NAME: Tell Me DESCRIPTION: notification test *ORGANIZATION: Honey Dog, Inc.

*TYPE: HipChat

TYPE DETAILS

*DESTINATION CHANNELS: #engineering, #support *TOKEN: SHOW ***** *LABEL TO BE SHOWN WITH NOTIFICATION: hipchat-notifier-tell-me

*API URL: https://mycompany.hipchat.com *NOTIFICATION COLOR: red ☐ NOTIFY CHANNEL

CANCEL SAVE

14.4.6 Webhook

The webhook notification type in Ansible Tower provides a simple interface to sending POSTs to a predefined web service. Tower will POST to this address using application/json content type with the data payload containing all relevant details in json format.

The parameters are pretty straightforward:

- Target URL: The full URL that will be POSTed to
- HTTP Headers: Headers in JSON form where the keys and values are strings. For example:

```
{ "Authentication": "988881adc9fc3655077dc2d4d757d480b5ea0e11", "MessageType": "Test" }
```

NEW NOTIFICATION TEMPLATE

*NAME: Tell Me DESCRIPTION: notification test *ORGANIZATION: Honey Dog, Inc.

*TYPE: Webhook

TYPE DETAILS

*TARGET URL: http://www.honeydog.com/web/db/notification

*HTTP HEADERS:

```
{ "Authentication": "988881adc9fc3655077dc2d4d757d480b5ea0e11", "MessageType": "Test" }
```

CANCEL SAVE

14.4.7 IRC

The Tower IRC notification takes the form of an IRC bot that will connect, deliver its messages to channel(s) or individual user(s), and then disconnect. The Tower notification bot also supports SSL authentication. The Tower bot does not currently support Nickserv identification. If a channel or user does not exist or is not on-line then the Notification will not fail; the failure scenario is reserved specifically for connectivity.

Connectivity information is straightforward:

- **IRC Server Password:** IRC servers can require a password to connect. If the server does not require one, leave blank
- **IRC Server Port:** The IRC server Port
- **IRC Server Address:** The host name or address of the IRC server
- **IRC Nick:** The bot's nickname once it connects to the server
- **Destination Channels or Users:** A list of users and/or channels to which to send the notification.
- **SSL Connection:** Should the bot use SSL when connecting

NEW NOTIFICATION TEMPLATE

* NAME: Tell Me

DESCRIPTION: notification test

* ORGANIZATION: Honey Dog, Inc.

* TYPE: IRC

TYPE DETAILS

* IRC SERVER PASSWORD: SHOW

* IRC SERVER PORT: 6667

* IRC SERVER ADDRESS: irc.testirc.net

* IRC NICK: helpbot

* DESTINATION CHANNELS OR USERS: #engineering, #release-engineers

☐ SSL CONNECTION

CANCEL SAVE

14.5 Configuring the towerhost hostname

In `/etc/tower/settings.py`, you can modify `TOWER_URL_BASE='https://tower.example.com'` to change the notification hostname, replacing `https://tower.example.com` with your preferred hostname. You must restart Tower services after saving your changes with `ansible-tower-service restart`.


Refreshing your Tower license also changes the notification hostname. New installations of Ansible Tower 3.0 should not have to set the hostname for notifications.

14.5.1 Resetting the TOWER_URL_BASE

The primary way that Tower determines how the base URL (`TOWER_URL_BASE`) is defined is by looking at an incoming request and setting the server address based on that incoming request.

Tower takes settings values from the database first. If no settings values are found, Tower falls back to using the values from the settings files. If a user posts a license by navigating to the Tower host's IP address, the posted license is written to the settings entry in the database.

To change the `TOWER_URL_BASE` if the wrong address has been picked up, navigate to the license from the Tower

Settings () Menu's 'VIEW YOUR LICENSE' link using the DNS entry you wish to appear in notifications, and re-add your license.

BEST PRACTICES

15.1 Use Source Control

While Tower supports playbooks stored directly on the Tower server, best practice is to store your playbooks, roles, and any associated details in source control. This way you have an audit trail describing when and why you changed the rules that are automating your infrastructure. Plus, it allows for easy sharing of playbooks with other parts of your infrastructure or team.

15.2 Ansible file and directory structure

Please review the Ansible best practices from the Ansible documentation at http://docs.ansible.com/playbooks_best_practices.html. If creating a common set of roles to use across projects, these should be accessed via source control submodules, or a common location such as `/opt`. Projects should not expect to import roles or content from other projects.

Note: Playbooks should not use the `vars_prompt` feature, as Tower does not interactively allow for `vars_prompt` questions. If you must use `vars_prompt`, refer to and make use of the *Surveys* functionality of Tower.

Jobs run in Tower use the playbook directory as the current working directory, although jobs should be coded to use the `playbook_dir` variable rather than relying on this.

15.3 Use Dynamic Inventory Sources

If you have an external source of truth for your infrastructure, whether it is a cloud provider or a local CMDB, it is best to define an inventory sync process and use Tower's support for dynamic inventory (including cloud inventory sources and *custom inventory scripts*). This ensures your inventory is always up to date.

Note: With the release of Ansible Tower 2.4.0, edits and additions to Inventory host variables now persist beyond an inventory sync as long as `--overwrite_vars` is **not** set. To have inventory syncs behave as they did before, it is now required that both `--overwrite` and `--overwrite_vars` are set.

15.4 Variable Management for Inventory

Keeping variable data along with the objects in Tower (see the inventory editor) is encouraged, rather than using `group_vars/` and `host_vars/`. If you use dynamic inventory sources, Tower can sync such variables with the database as long as the **Overwrite Variables** option is not set.

15.5 Autoscaling

Using the “callback” feature to allow newly booting instances to request configuration is very useful for auto-scaling scenarios or provisioning integration.

15.6 Larger Host Counts

Consider setting “forks” on a job template to larger values to increase parallelism of execution runs. For more information on tuning Ansible, see [the Ansible blog](#).

15.7 Continuous integration / Continuous Deployment

For a Continuous Integration system, such as Jenkins, to spawn an Tower job, it should make a curl request to a job template, or use the [Tower CLI](#) tool. The credentials to the job template should not require prompting for any particular passwords. Using the API to spawn jobs is covered in the [Tower API guide](#).

SECURITY

The following sections will help you gain an understanding of how Ansible Tower handles and lets you control file system security.

All playbooks are executed via the `awx` file system user. For running jobs, Ansible Tower defaults to offering job isolation via Linux namespacing and `chroots`. This projection ensures jobs can only access playbooks and roles from the Project directory for that job template and common locations such as `/opt`. Playbooks are not able to access roles, playbooks, or data from other Projects by default.

If you need to disable this protection (not recommended), you can edit `/etc/tower/settings.py` and set `AWX_PROOT_ENABLED` to `False`.

Note: In this scenario, playbooks have access to the file system and all that that implies; therefore, users who have access to edit playbooks **must** be trusted.

For credential security, users may choose to upload locked SSH keys and set the unlock password to “ask”. You can also choose to have the system prompt them for SSH credentials or sudo passwords rather than having the system store them in the database.

16.1 Playbook Access and Information Sharing

By default, Tower’s multi-tenant security prevents playbooks from reading files outside of their project directory. To share information between playbooks or to read files on the file system outside of their project directory, you must edit `/etc/tower/settings.py` and add the directories that are available to the `AWX_PROOT_SHOW_PATHS` setting.

The following paths, plus any user specified paths, are hidden by `AWX_PROOT_HIDE_PATHS`:

- `/etc/tower`
- `/var/lib/awx`
- `/var/log`
- `/tmp`
- `/var/lib/awx/projects`
- `/var/lib/awx/job_status`

The following paths, plus any user specified paths, are shown by `AWX_PROOT_SHOW_PATHS`:

- `/var/lib/awx/projects/<current_project>`
- `/tmp/ansible_tower_XXXXX`

The primary file you may want to add to `AWX_PROOT_SHOW_PATHS` is `/var/lib/awx/.ssh`, if your playbooks need to use keys or settings defined there.

16.2 PRoot functionality and variables

The PRoot functionality in Ansible Tower limits which directories on the Tower file system are available for playbooks to see and use during playbook runs. You may find that you need to customize your PRoot settings in some cases. To fine tune your usage of PRoot, there are certain variables that can be set:

```
# Enable proot support for running jobs (playbook runs only).
AWX_PROOT_ENABLED = False

# Command/path to proot.
AWX_PROOT_CMD = 'proot'

# Additional paths to hide from jobs using proot.
AWX_PROOT_HIDE_PATHS = []

# Additional paths to show for jobs using proot.
AWX_PROOT_SHOW_PATHS = []
```

To customize your PRoot settings, navigate to the `/etc/tower/settings.py` file. Once your changes have been saved, restart services with the `ansible-tower-service restart` command.

16.3 Role-Based Access Controls

Role-Based Access Controls (RBAC) are built into Tower and allow Tower administrators to delegate access to server inventories, organizations, and more. Administrators can also centralize the management of various credentials, allowing end users to leverage a needed secret without ever exposing that secret to the end user. RBAC controls allow Tower to help you increase security and streamline management.

RBACs are easiest to think of in terms of Roles which define precisely who or what can see, change, or delete an “object” for which a specific capability is being set. In releases prior to Ansible Tower version 3.0, RBAC was thought of in terms of granting permissions to users or teams. Starting with Tower 3.0, RBAC is best thought of as granting roles to users or teams, which is a more intuitive approach.


There are a few main concepts that you should become familiar with regarding Tower’s RBAC design—roles, resources, and users. Users can be members of a role, which gives them certain access to any resources associated with that role, or any resources associated with “descendant” roles.

A role is essentially a collection of capabilities. Users are granted access to these capabilities and Tower’s resources through the roles to which they are assigned or through roles inherited through the role hierarchy.

Roles associate a group of capabilities with a group of users. All capabilities are derived from membership within a role. Users receive capabilities only through the roles to which they are assigned or through roles they inherit through the role hierarchy. All members of a role have all capabilities granted to that role. Within an organization, roles are relatively stable, while users and capabilities are both numerous and may change rapidly. Users can have many roles.

16.3.1 Role Hierarchy and Access Inheritance

Imagine that you have an organization named “SomeCompany” and want to allow two people, “Josie” and “Carter”, access to manage all the settings associated with that organization. You should make both people members of the organization’s `admin_role`.

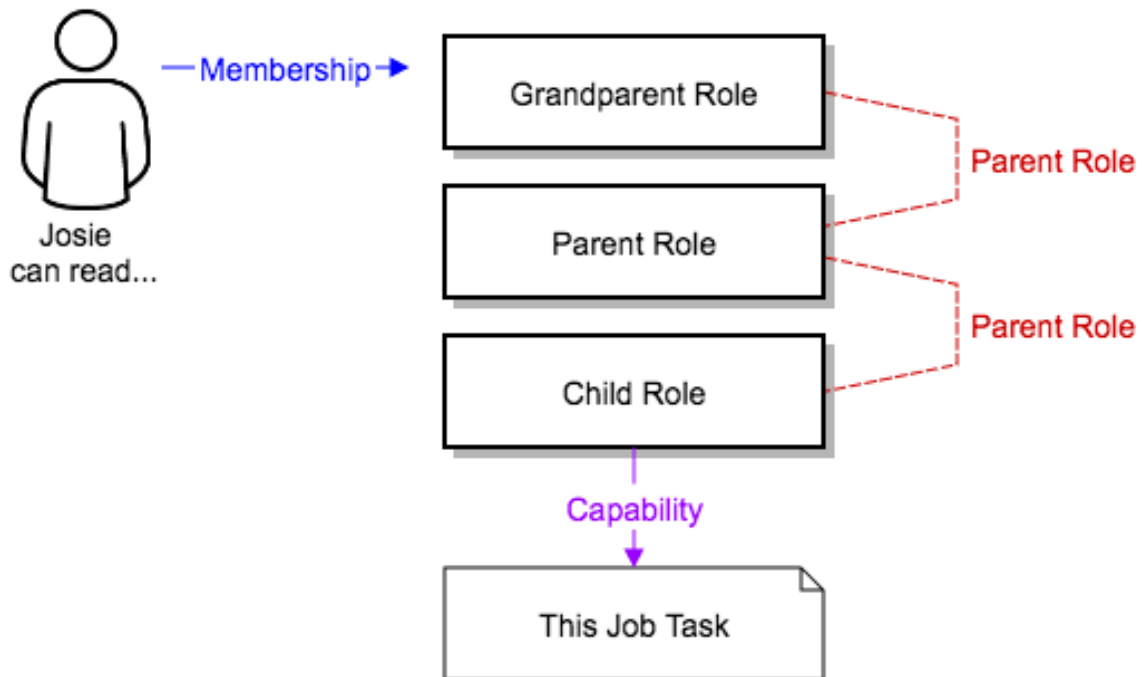


Often, you will have many Roles in a system and you will want some roles to include all of the capabilities of other roles. For example, you may want a System Administrator to have access to everything that an Organization Administrator has access to, who has everything that a Project Administrator has access to, and so on.

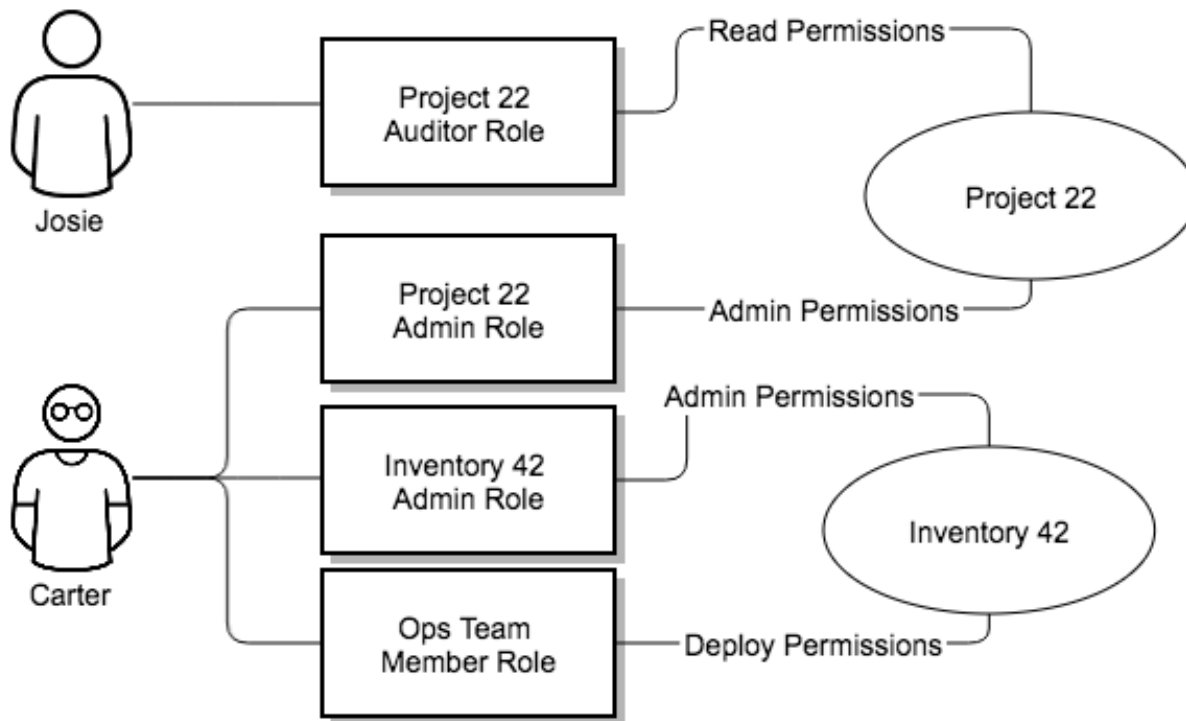
This concept is referred to as the ‘Role Hierarchy’:

- Parent roles get all capabilities bestowed on any child roles
- Members of roles automatically get all capabilities for the role they are a member of, as well as any child roles.

The Role Hierarchy is represented by allowing Roles to have “Parent Roles”. Any capability that a Role has is implicitly granted to any parent roles (or parents of those parents, and so on).



Often, you will have many Roles in a system and you will want some roles to include all of the capabilities of other roles. For example, you may want a System Administrator to have access to everything that an Organization Administrator has access to, who has everything that a Project Administrator has access to, and so on. We refer to this concept as the 'Role Hierarchy' and it is represented by allowing Roles to have "Parent Roles". Any capability that a Role has is implicitly granted to any parent roles (or parents of those parents, and so on). Of course Roles can have more than one parent, and capabilities are implicitly granted to all parents.



RBAC controls also give you the capability to explicitly permit User and Teams of Users to run playbooks against certain sets of hosts. Users and teams are restricted to just the sets of playbooks and hosts to which they are granted capabilities. And, with Tower, you can create or import as many Users and Teams as you require—create users and teams manually or import them from LDAP or Active Directory.

RBACs are easiest to think of in terms of who or what can see, change, or delete an “object” for which a specific capability is being determined.

16.3.2 Applying RBAC

The following sections cover how to apply Tower’s RBAC system in your environment.

Editing Users

When editing a user, a Tower system administrator may specify the user as being either a *System Administrator* (also referred to as the Superuser) or an *Auditor*.

- System administrators implicitly inherit all capabilities for all objects (read/write/execute) within the Tower environment.
- Auditors implicitly inherit the read-only capability for all objects within the Tower environment.

Editing Organizations

When editing an organization, system administrators may specify the following roles:

- One or more users as organization administrators

- One or more users as organization auditors
- And one or more users (or teams) as organization members

Users/teams that are members of an organization can view their organization administrator.

Users who are organization administrators implicitly inherit all capabilities for all objects within that Tower organization.

Users who are organization auditors implicitly inherit the read-only capability for all objects within that Tower organization.

Editing Projects in an Organization

When editing a project in an organization for which they are the administrator, system administrators and organization administrators may specify:

- One or more users/teams that are project administrators
- One or more users/teams that are project members
- And one or more users/teams that may update the project from SCM, from among the users/teams that are members of that organization.

Users who are members of a project can view their project administrators.

Project administrators implicitly inherit the capability to update the project from SCM.

Creating Inventories and Credentials within an Organization

All access that is granted to use, read, or write credentials is now handled through roles. 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.

System administrators and organization administrators may create inventories and credentials within organizations under their administrative capabilities.

Whether editing an inventory or a credential, System administrators and organization administrators may specify one or more users/teams (from those that are members of that organization) to be granted the usage capability for that inventory or credential.

System administrators and organization administrators may specify one or more users/teams (from those that are members of that organization) that have the capabilities to update (dynamic or manually) an inventory. Administrators can also execute ad hoc commands for an inventory.

Editing Job Templates

System administrators, organization administrators, and project administrators, within a project under their administrative capabilities, may create and modify new job templates for that project.

When editing a job template, administrators (Tower, organization, and project) can select among the inventory and credentials in the organization for which they have usage capabilities or they may leave those fields blank so that they will be selected at runtime.

Administrators can also specify one or more users/teams (from those that are members of that project) that can use that project in a job template.

Additionally, they may specify one or more users/teams (from those that are members of that project) that have execution capabilities for that job template. The execution capability is valid regardless of any explicit capabilities the user/team may have been granted against the inventory or credential specified in the job template.

User View

A user can:

- See any organization or project for which they are a member
- Create their own credential objects which only belong to them
- See and execute any job template for which they have been granted execution capabilities

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.

16.3.3 Roles

As stated earlier in this documentation, all access that is granted to use, read, or write credentials is now handled through roles, and roles are defined for a resource.

Built-in roles

The following table lists the RBAC system roles and a brief description of the how that role is defined with regard to privileges in Tower.

System Role	What it can do
System Administrator - System wide singleton	Manages all aspects of the system
System Auditor - System wide singleton	Views all aspects of the system
Ad Hoc Role - Inventory	Runs ad hoc commands on an Inventory
Admin Role - Organizations, Projects, Inventory, Projects, Job Templates	Manages all aspects of a defined Organization, Project, Inventory, or Job Template
Auditor Role - Organizations, Projects, Inventory, Projects, Job Templates	Views all aspects of a defined Organization, Project, Inventory, or Job Template
Execute Role - Job Templates	Runs assigned Job Template
Member Role - Organization, Team	User is a member of a defined Organization or Team
Read Role - All	Views settings for a defined Organization, Project, Inventory, or Job Template
SCM Update Role - Project	Updates the Project from the configured source control management system
Update Role - Inventory	Updates the Inventory using the cloud source update system
Owner Role - Credential	Owns and manages all aspects of this Credential
Use Role - Credential, Inventory, Project	Uses the Credential, Inventory, or Project in a Job Template

A Singleton Role is a special role that you can create and is intended for system wide roles.

INDEX

- `genindex`

COPYRIGHT © 2016 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

- about Tower
 - settings menu, 13
- activity streams, 17
- ad hoc commands, 84
 - inventories, 84
- add new
 - inventories, 63
 - scan job, 106
- adding new
 - credentials, 42
- admin menu, 12
- administrators
 - organization, 26
- Amazon Web Services
 - credential types, 47
 - inventories, 69
- Ansible Galaxy, 62
- Ansible Galaxy integration
 - features, 3
- automation
 - features, 2
- autoscaling
 - best practices, 144
- autoscaling flexibility
 - features, 3
- AWS
 - cloud credentials, 101

B

- backup and restore
 - features, 3
- best practices, 143
 - autoscaling, 144
 - deployment, continuous, 144
 - dynamic inventory sources, 143
 - file and directory structure, 143
 - host counts, larger, 144
 - integration, continuous, 144
 - source control, 143
 - variable inventory management, 144

C

- callbacks
 - extra variables, 116
- check
 - job types, 93
- cloud credentials
 - AWS, 101
 - Google, 101
 - job templates, 100
 - MS Azure, 101
 - OpenStack, 100
 - Rackspace, 101
 - VMware, 101
- cloud flexibility
 - features, 3
- CloudForms
 - credential types, 50
- components
 - licenses, 7
- concurrency
 - jobs, 133
- create template
 - notifications, 136
- credential types, 43
 - Amazon Web Services, 47
 - CloudForms, 50
 - Google Compute Engine, 51
 - machine, 44
 - Microsoft Azure Classic, 52
 - Microsoft Azure Resource Manager, 52
 - network, 45
 - OpenStack, 54
 - rackspace, 48
 - Red Hat Satellite, 49
 - source control, 46
 - VMware, 49
- credentials, 39
 - adding new, 42
 - getting started, 39
 - how they work, 39
 - settings menu, 13
 - types, 43

- custom
 - fact scan job, 114
 - scan job, 113
- custom fact scans
 - playbook, 114
 - system tracking, 114
- custom script
 - inventories, 77

D

- dashboard, 16
 - host count, 16
 - job status, 16
 - jobs tab, 16
 - main menu, 11
 - schedule status, 16
- DEB files
 - licenses, 7
- deployment, continuous
 - best practices, 144
- dynamic inventory sources
 - best practices, 143

E

- Email
 - notifications types, 136
- extra variables
 - callbacks, 116
 - provisioning callbacks, 116
 - surveys, 104
- extra_vars, 105

F

- fact scan job
 - custom, 114
 - playbook, 113
- fact scan playbook
 - system tracking, 113
- features, 4
 - Ansible Galaxy integration, 3
 - automation, 2
 - autoscaling flexibility, 3
 - backup and restore, 3
 - cloud flexibility, 3
 - inventory sources, Red Hat CloudForms, 4
 - inventory sources, Red Hat Satellite 6, 4
 - notifications, 4
 - OpenStack inventory support, 3
 - overview, 2
 - real-time playbook, 2
 - remote command execution, 3
 - RESTful API, 3
 - role-based access control, 2
 - run-time job customization, 4

- system tracking, 4
- file and directory structure
 - best practices, 143
- forks
 - jobs, 133
- functionality
 - PRoot, 146

G

- Galaxy support, 62
- getting started
 - credentials, 39
- Google
 - cloud credentials, 101
- Google Compute Engine
 - credential types, 51
 - inventories, 71
- granted permissions
 - teams, 37
- groups
 - notifications, 135

H

- Hipchat
 - notifications types, 136
- host count
 - dashboard, 16
- host counts, larger
 - best practices, 144
- host to host
 - scan job, 89
 - system tracking, 89
- hostname configuration
 - notifications, 141
- hosts
 - inventories, 83
- hosts, add new
 - inventories, 84
- how they work
 - credentials, 39

I

- installation bundle
 - licenses, 7
- integration, continuous
 - best practices, 144
- inventories, 63
 - ad hoc commands, 84
 - add new, 63
 - Amazon Web Services, 69
 - custom script, 77
 - Google Compute Engine, 71
 - groups, 65
 - add new, 66

- groups and hosts, 65
- hosts, 83
- hosts, add new, 84
- Microsoft Azure Classic (deprecated), 71
- Microsoft Azure Resource Manager, 72
- OpenStack, 76
- Rackspace Cloud Servers, 68
- Red Hat CloudForms, 75
- Red Hat Satellite 6, 74
- scan job, 86
- scan job templates, creating, 106
- scan job templates, custom, 113
- scan job templates, launching, 110
- scan job templates, scheduling, 112
- scheduling, 78
- scheduling, add new, 79
- system tracking, 86
- VMware vCenter, 73
- inventory scripts
 - settings menu, 13
- inventory sources
 - notifications, 135
- inventory sources, Red Hat CloudForms
 - features, 4
- inventory sources, Red Hat Satellite 6
 - features, 4
- inventory sync
 - job results, 121
- IRC
 - notifications types, 136

J

- job results, 121
 - inventory sync, 121
- job status
 - dashboard, 16
- job templates, 93
 - cloud credentials, 100
 - job variables, 105
 - jobs, launching, 116
 - portal mode, 15
 - provisioning callbacks, 115
 - relaunch, 105
 - scheduling, 118
 - survey creation, 102
 - survey extra variables, 104
 - survey optional questions, 104
 - surveys, 102
- job templates, hierarchy, 105
- job templates, overview, 105
- job types
 - check, 93
 - run, 93
 - scan, 93

- job variables
 - job templates, 105
- jobs, 121
 - concurrency, 133
 - event summary, 130
 - forks, 133
 - host events, 128
 - host summary, 130
 - job details, 127
 - notifications, 135
 - plays, 127
 - portal mode, 15
 - results, 121
 - tasks, 128
- jobs results
 - playbook run, 124
 - SCM, 123
- jobs tab
 - dashboard, 16
- jobs, launching
 - job templates, 116

L

- launching
 - scan job, 110
- license, 4, 5
 - features, 6
 - nodes, 6
 - trial, 5
 - troubleshooting, 10
 - types, 5
- license features, 4
- license, add manually, 10
- license, import, 9
- license, viewing, 13
- licenses
 - components, 7
 - DEB files, 7
 - installation bundle, 7
 - RPM files, 7
- logging in, 8

M

- machine
 - credential types, 44
- main menu
 - dashboard, 11
- management jobs
 - settings menu, 13
- Microsoft Azure Classic
 - credential types, 52
- Microsoft Azure Classic (deprecated)
 - inventories, 71
- Microsoft Azure Resource Manager

- credential types, 52
- inventories, 72
- MS Azure
 - cloud credentials, 101
- my view, 14

N

- network
 - credential types, 45
- new schedule addition
 - projects, 61
- notifications
 - create template, 136
 - features, 4
 - groups, 135
 - hostname configuration, 141
 - inventory sources, 135
 - jobs, 135
 - notifier, 135
 - notifier hierarchy, 135
 - notifier workflow, 135
 - organizations, 24
 - resetting the TOWER_URL_BASE, 141
 - template, 136
 - troubleshooting TOWER_URL_BASE, 141
 - types, 136
 - types Email, 136
 - types Hipchat, 136
 - types IRC, 136
 - types pagerduty, 136
 - types Slack, 136
 - types Twilio, 136
 - types Webhook, 136
- notifier
 - notifications, 135
- notifier hierarchy
 - notifications, 135
- notifier workflow
 - notifications, 135

O

- OpenStack
 - cloud credentials, 100
 - credential types, 54
 - inventories, 76
- OpenStack inventory support
 - features, 3
- organization
 - administrators, 26
 - settings menu, 13
- organizations, 19
 - notifications, 24
 - permissions, 22
 - users, 25, 32

- overview
 - features, 2

P

- pagerduty
 - notifications types, 136
- permissions
 - organizations, 22
 - teams, 37
 - users, 33
- playbook
 - custom fact scans, 114
 - fact scan job, 113
- playbook run
 - jobs results, 124
- playbooks
 - manage manually, 58
 - projects, 58
 - PRoot settings, 145
 - sharing access, 145
 - sharing content, 145
 - source control, 58
- portal mode, 14
 - job templates, 15
 - jobs, 15
- projects, 56
 - add new, 57
 - new schedule addition, 61
 - playbooks, 58
 - source control update, 59
- PRoot
 - functionality, 146
 - troubleshooting, 146
 - variables, 146
- PRoot settings
 - playbooks, 145
- provisioning callbacks
 - extra variables, 116
 - job templates, 115

R

- Rackspace
 - cloud credentials, 101
- rackspace
 - credential types, 48
- Rackspace Cloud Servers
 - inventories, 68
- RBAC
 - security, 146
- real-time playbook
 - features, 2
- Red Hat CloudForms
 - inventories, 75
- Red Hat Satellite

- credential types, 49
- Red Hat Satellite 6
 - inventories, 74
- relaunch
 - job templates, 105
- remote command execution
 - features, 3
- resetting the TOWER_URL_BASE
 - notifications, 141
- RESTful API
 - features, 3
- role-based access control
 - features, 2
- role-based access controls, 146
- RPM files
 - licenses, 7
- run
 - job types, 93
- run-time job customization
 - features, 4

S

- scan
 - job types, 93
- scan job
 - add new, 106
 - custom, 113
 - host to host, 89
 - inventories, 86
 - launching, 110
 - scheduling, 112
 - single host, 87
- scan job templates, creating
 - inventories, 106
- scan job templates, custom
 - inventories, 113
- scan job templates, launching
 - inventories, 110
- scan job templates, scheduling
 - inventories, 112
- schedule status
 - dashboard, 16
- scheduling
 - add new, 118
 - inventories, 78
 - job templates, 118
 - scan job, 112
- scheduling, add new
 - inventories, 79
- SCM
 - jobs results, 123
- security, 145
 - RBAC, 146
- settings menu

- about Tower, 13
- credentials, 13
- inventory scripts, 13
- management jobs, 13
- organization, 13
- teams, 13
- users, 13
- view license, 13
- sharing access
 - playbooks, 145
- sharing content
 - playbooks, 145
- single host
 - scan job, 87
 - system tracking, 87
- Slack
 - notifications types, 136
- source control
 - best practices, 143
 - credential types, 46
- source control update
 - projects, 59
- support, 4, 5
- survey extra variables
 - job templates, 104
- surveys
 - creation, 102
 - extra variables, 104
 - job templates, 102
 - optional questions, 104
- system tracking
 - custom fact scans, 114
 - fact scan playbook, 113
 - features, 4
 - host to host, 89
 - inventories, 86
 - scan job, 93
 - single host, 87

T

- teams, 34
 - granted permissions, 37
 - permissions, 37
 - settings menu, 13
 - users, 32, 35
- template
 - notifications, 136
- Tower admin menu, 12
- Tower settings menu, 13
- troubleshooting
 - license, 10
 - PRoot, 146
- troubleshooting TOWER_URL_BASE
 - notifications, 141

Twilio

notifications types, 136

types

Email, notifications, 136

Hipchat, notifications, 136

IRC, notifications, 136

notifications, 136

pagerduty, notifications, 136

Slack, notifications, 136

Twilio, notifications, 136

Webhook, notifications, 136

U

updates, 5

user menu, 14

users, 29

organizations, 25, 32

permissions, 33

settings menu, 13

teams, 32, 35

V

variable inventory management

best practices, 144

variable precedence, 105

variables

PRoot, 146

view license

settings menu, 13

VMware

cloud credentials, 101

credential types, 49

VMware vCenter

inventories, 73

W

Webhook

notifications types, 136