
Ansible Tower Installation and Reference Guide

Release Ansible Tower 2.3.1

Ansible, Inc.

Jun 06, 2017

CONTENTS

1	Licensing, Updates, and Support	2
1.1	Trial Licenses	2
1.2	License Types	2
1.3	License Features	3
1.4	Tower Component Licenses	3
2	Release Notes	4
3	Known Issues	8
3.1	Installation failure related to MongoDB	8
3.2	Errors when editing objects	8
3.3	Host comparisons against a single host	8
3.4	Host comparisons against two hosts	8
3.5	Live events status indicators	8
3.6	Slower than expected performance	9
3.7	<code>sudo/su</code> not working as expected for local playbooks or playbooks with <code>local_actions</code>	9
3.8	Installation program does not enable <code>su</code> command support	9
3.9	Problems when using SSH customization	10
3.10	Ubuntu unsupported for bundled installations	10
3.11	YAML parser errors when installing Tower without <code>python-yaml</code> installed	10
4	General Installation Notes	11
5	Supported Platforms and Requirements	12
6	Platform-Specific Notes	13
6.1	Red Hat Enterprise Linux and CentOS	13
7	Requirements	14
8	Prerequisites	16
8.1	Configuration and Installation for Ansible with Red Hat Enterprise Linux and CentOS	16
8.2	Configuration and Installation for Ansible with Ubuntu	17
9	Tower Installation Scenarios	18
10	Get the Tower Installation Program	19
10.1	Using Vagrant/Amazon AMI Images	19
10.2	Using the Bundled Tower Installation Program	19
11	The Tower Installation Wizard	21

11.1	Installation Arguments	21
11.2	Primary Tower machine configuration	22
11.3	Secondary Installation (if applicable)	23
11.4	Passwords	23
11.5	Connection Information	23
11.6	Review and Confirm	24
11.7	Reviewing the Tower Configuration	25
11.8	The Setup Playbook	25
12	Upgrading an Existing Tower Installation	26
12.1	Requirements	26
12.2	Backing Up Your Tower Installation	26
12.3	Get the Tower Installer	26
12.4	Run the Tower Installation Wizard	27
12.5	The Setup Playbook	28
13	Supported Locales	30
14	Glossary	32
15	Index	34
	Index	35

Thank you for your interest in Ansible Tower, 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.

The *Ansible Tower Installation and Reference Guide* helps you to understand the installation requirements and processes behind installing Ansible Tower. This document has been updated to include information for the latest release of Ansible Tower 2.3.1.

Ansible Tower Version 2.3.1; October 2, 2015; <https://access.redhat.com/>

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>

1.1 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/>

1.2 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.

- **Basic**
 - Manage smaller environments (up to 250 nodes)
 - 30 days of initial web support through <https://access.redhat.com/>, for installation and setup
 - Maintenance and upgrades included
- **Enterprise**
 - Manage any size environment
 - Enterprise 8x5 support and SLA (4 hour critical incident response)
 - Phone and web support
 - Support for Ansible core included
 - Maintenance and upgrades included
- **Premium**
 - Manage any size environment, including mission-critical environments
 - Premium 24x7 support and SLA (4 hour critical incident response, 8 hour non-critical incident response)

- Phone and web support
- Support for Ansible core included
- Maintenance and upgrades included

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/>.

1.3 License Features

Ansible Tower 2.2 introduced a separation of features for Basic versus Enterprise or Premium licenses. The following list of features are available for all new Enterprise or Premium license users:

- Multi-Organization Support
- Activity Streams
- Surveys
- LDAP Support
- High Availability
- System Tracking (*added in Ansible Tower 2.2*)

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

1.4 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.

RELEASE NOTES

- Version 2.3.1
 - Corrected an issue where PRoot being enabled caused jobs to fail on systems using SSH ControlPersist.

Caution: If Ansible’s Customer Support recommended that you disable PRoot to solve the failing jobs problem (setting `AWX_PROOT_ENABLED=False`), consult with [Support](#) to determine if re-enabling PRoot is appropriate for your particular use case.

- Version 2.3
 - Added support for bundled installations
 - Added improvements for preflight free disk space check
 - Added Ansible installation support where the Ansible Tower installer now attempts to install Ansible as part of the installation process
 - Corrected an issue where launching a JT with a Survey attached failed if you had survey data types other than “text” or “text area”
 - Corrected an issue where scan jobs fail on large file scans
 - Corrected an issue where projects were not included in system backups
 - Corrected an issue where downloading stdout in text format would return JSON instead
 - Corrected an issue where downloading stdout in text format would incorrectly escape characters
 - Corrected a performance issue when accessing jobs and job_templates
- Version 2.2.2
 - Corrected an issue where unicode credential passwords caused migrations to fail
 - Corrected a performance issue when Tower redacts sensitive data from job output
- Version 2.2.1
 - Fixed performance issues when job stdout was very large
 - Corrected an issue where stdout display in Tower would fail on some unicode output
 - Corrected an issue where EC2 inventory sync would fail if instances had blank tags
 - Corrected an issue where jobs would not cancel properly on user cancellation (applies to EL6 platforms where PRoot was enabled by default)
 - Corrected an issue when restoring a Tower database backup to a remote PostgreSQL database
 - Added support for newer OpenSSH private key format

- Fixed display of Tower version in ‘About Tower’
- Fixed links to Ansible Github repository in dynamic inventory online help
- Version 2.2
 - Added System Tracking job scan (available for Enterprise and Premium licenses only)
 - Simplified Dashboard and Interface with new Setup Menu
 - Added inventory support for OpenStack
 - Added data cleanup and snapshot retention scheduling
 - Added Ansible Galaxy integration
 - Added support for Remote Command Execution
 - Added Status widget for easily viewing the 10 most recent jobs run on a job template
 - Added integration for easier backups and restorations into the Tower setup playbook
 - Adjusted dates to display in the user’s locale format
 - Simplified password/passphrase entry
 - Added more configurable verbosity levels for job templates
 - Assorted other bugfixes and enhancements
 - API change: Formatting of `extra_vars` attached to Job Template records is preserved. Previously, YAML would be converted to JSON and returned as JSON. In 2.2.0 and newer, YAML is returned as YAML with formatting and comments preserved, and JSON is returned as JSON.
- Version 2.1.4
 - Corrected Tower’s Live Events feature, again. Really.
- Version 2.1.3
 - Corrected an issue where Tower Live Events would attempt to endlessly reconnect
 - Corrected issues when running with Ansible 1.9.0.1
- Version 2.1.2
 - Corrected multiple issues with Tower’s Live Events feature
 - Corrected an issue where Tower would become stuck if a job was killed due to memory exhaustion
 - Improved the response time of Project queries
 - Corrected an error that caused users to be unable to relaunch jobs
- Version 2.1.1
 - Multi-tenancy security enabled by default for new installs
 - Added support for setting VPC id for RDS instances to EC2 dynamic inventory
 - Added the ability for organization admins to create surveys
 - Added support for scheduling of custom inventory scripts
 - Corrected an error when parsing `extra_vars` as YAML
 - Corrected an error when configuring a remote database
 - Added EULA agreement when updating license
 - Corrected the sending of live events in some cases

- Corrected a potential XSS issue
- Version 2.1
 - New simplified Portal Mode view for users, access at `https://<Tower server name>/portal/`
 - New surveys on job templates allow easy prompting of users for job parameters
 - Tower can now use an external PostgreSQL instance as the Tower database, including Amazon's RDS
 - Added support for active/passive High Availability Tower deployments
 - Custom dynamic inventory scripts can be pasted in using the admin user menu
 - Limit Amazon EC2 inventory imports into Tower based on tags, keys, and more
 - Tower data cleanup jobs can now be scheduled and run directly from the Tower interface versus logging into the Tower instance
 - The `/etc/awx` Tower configuration directory has moved to `/etc/tower`
 - Non-admin api users must now use the `/launch` endpoint for a job template and can no longer call a job's `/start` endpoint directly.
 - Many assorted improvements and fixes
- Version 2.0.5
 - Ensured websocket connection uses user's RBAC credentials
 - Corrected a potential CSRF issue when using the REST API graphical browser
- Version 2.0.4
 - Corrected a privilege escalation related to user account levels
- Version 2.0.2
 - Further corrections for job execution with certain Omq library versions
 - Changes to AMI license logic to allow bring-your-own-license usage
- Version 2.0.1
 - Corrected a job execution issue due to Omq library versions on certain platforms
 - Reduced logfile verbosity and retention for some Tower subcomponents
 - Adjusted setup playbook for the release of EPEL 7
- Version 2.0
 - New dashboard that provides at-a-glance status of your Ansible deployment
 - Completely redesigned job status page featuring real-time playbook output and progress updates
 - Added support for multiple new cloud providers - Azure, Google Compute Engine, and VMware vSphere
 - New user interface look and feel
 - Integrated monitoring support for checking the health of your Tower install
 - Tower now requires a license to run. 10 machine free licenses, as well as free large trial licenses, are available at <http://ansible.com/license>
 - Support added for Red Hat Enterprise Linux 7 and CentOS 7
 - Upgrades will reuse password information, not requiring reentry in 'group_vars/all' of setup playbook
 - Many assorted improvements and fixes

- Version 1.4.12
 - Corrected an issue handling Unicode output from ansible-playbook
 - Corrected an issue displaying job details for some jobs
- Version 1.4.11
 - Performance improvements to inventory import and deletion
 - * Groups UI under inventory tab is now paginated
 - * Updated UI options for moving and copying groups (and host contents)
 - Added the ability to optionally prompt for job variables when launching jobs to the job template detail pages
- Version 1.4.10
 - Correctly handle schedule creation when browser timezone cannot be detected.
 - Corrected pagination on job_events page.
- Version 1.4.9
 - Corrected a provisioning callback issue on Enterprise Linux.
 - Added a sample provisioning callback script.
 - Various backend and UI improvements.
- Version 1.4.8
 - Scheduling for Jobs, SCM updates, and Inventory synchronization has been added. The UI for each of these objects has changed to accommodate this new scheduling feature.
 - * The jobs page has been overhauled to show completed, active, queued, and scheduled jobs.
 - * Inventory and project synchronization jobs are now also shown on the jobs page.
 - Added support for Ansible Vault to Credentials. For more information on how to use Ansible Vault, please visit: http://docs.ansible.com/playbooks_vault.html.

KNOWN ISSUES

3.1 Installation failure related to MongoDB

During the releases of Tower 2.x, the upstream URL for the MongoDB repository changed. This change in the upstream URL causes installations to fail.

The Ansible Tower 2.4.5 release has been corrected with the new upstream URL for MongoDB. For those that require MongoDB, use the latest Tower 2.4.5 release. If MongoDB is not a requirement, consider upgrading to the latest version of Ansible Tower.

3.2 Errors when editing objects

Ansible Tower implements a role based access control system. It may appear that you can edit objects that do not belong to you—like being able to pull up an edit dialog of your teammates, which you have been granted permission to view—but when you try to edit something, you will receive a 403 error and you cannot view any information you should not already have access to as defined in the system.

3.3 Host comparisons against a single host

When performing a host comparison against a single host, if there is only one (1) run for the selected date, Tower may display a message saying it could not find any scan job runs in one of the columns.

3.4 Host comparisons against two hosts

When performing a host comparison against two hosts, you can only select from a single page of hosts.

3.5 Live events status indicators

Live events status dots are either seen as a red or orange dot at the top of the Tower Dashboard when something goes wrong, or they are not seen at all to indicate a healthy system state. If you encounter a red or orange live events status indicator, even when your system seems fine, the following suggestions may offer a solution:

- Try manually refreshing/reloading your browser page.
- Try changing web browsers, as Firefox and Safari have been reported to have issues trusting self-signed certificates.

- Try creating a self-signed certificate that matches your DNS and import it into your trust manually.
- Try using an incognito or private browsing session.
- Try disabling your browser plugins to ensure none are blocking the service.
- Try accessing Tower at `https://tower-ip:8080/` which should allow you to accept the certificate for the machine.

Live event status dots are used for troubleshooting problems with your Tower instance and the socketio logs can point out anything problematic. You can collect troubleshooting help by running a `sosreport`. As root, run the command `sosreport` from your system to automatically generate a diagnostic tar file, then contact Ansible's Support team with the collected information for further assistance.

Note: Starting with Ansible Tower 2.2.0, live events status indicators only appear if Tower detects a problem. In earlier releases, a green status dot was shown to indicate a healthy system.

3.6 Slower than expected performance

Slow performance issues have been reported against Ansible Tower beginning with version 2.2.0.

You may experience slower than expected system performance if you have a large number of job templates, users, credentials, permissions, projects, and/or inventories.

If you find that loading the job template page for normal users is unacceptably slow, run the “Cleanup Deleted Data” management job with `days = 0`. Refer to [Management Jobs](#) in the *Ansible Tower Administration Guide* for more information.

3.7 sudo/su not working as expected for local playbooks or playbooks with local_actions

Instances have been reported that `sudo/su` commands do not work when using an entirely local playbook or a playbook with some `local_actions` cases. This is likely due to PRoot being enabled. To use `sudo/su` commands with local playbooks or those with `local_actions`, PRoot must be disabled. You can disable PRoot by navigating to the `/etc/tower/settings.py` file, setting `AWX_PROOT_ENABLED=False`, then restarting services with the `ansible-tower-service restart` command.

3.8 Installation program does not enable su command support

During the testing of Ansible Tower 2.3.0, it was discovered that the installation program does not handle the option to enable `su` command support.

When using the Tower installation program to install Tower to a remote system where `su` access is required, the installation program will not work. To workaround this issue, you should manually run the `ansible-playbook` command with the `--su` parameter.

3.9 Problems when using SSH customization

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. If you are attempting to customize SSH behavior by using a custom SSH configuration in the Tower user's home directory, this directory must be added to the list of directories exposed by PRoot.

For example, to add a custom SSH config in `/var/lib/awx/.ssh/config` and make it available for Tower jobs, edit the `/etc/tower/settings.py` file and add that path to the `AWX_PROOT_SHOW_PATHS` variable:

```
AWX_PROOT_SHOW_PATHS = [ '/var/lib/awx/.ssh/' ]
```

Once the paths you need have been added to `settings.py`, restart Tower using the admin utility script, `ansible-tower-service`:

```
ansible-tower-service restart
```

3.10 Ubuntu unsupported for bundled installations

If you are using the bundled installer for Ansible Tower 2.3.0, note that only Red Hat Enterprise Linux and CentOS are supported at this time. Ubuntu support has not yet been added.

3.11 YAML parser errors when installing Tower without *python-yaml* installed

Instances of YAML parser traceback errors have been reported when users attempt to install Tower without `python-yaml` being installed. In this scenario, the installation program uses the built-in YAML parser, but the built-in YAML parser is unable to parse a YAML variable initialized to an empty dictionary (e.g. `{}`).

The traceback may look like the following:

```
Traceback (most recent call last):
  File "./configure", line 787, in <module>
    if conf['configure_private_vars'].get(field, None):
AttributeError: 'str' object has no attribute 'get'
```

To workaround this issue, remove the offending variable (such as `configure_private_vars: {}`) from the Tower settings file, `tower_setup_conf.yml`.

GENERAL INSTALLATION NOTES

- If you need to access a HTTP proxy to install software from your OS vendor, ensure that the environment variable “HTTP_PROXY” is set accordingly before running `setup.sh`.
- The Tower installer creates a self-signed SSL certificate and keyfile at `/etc/tower/awx.cert` and `/etc/tower/awx.key` for HTTPS communication. These can be replaced after install with your own certificates if you desire, but the filenames are required to be the same.
- If using Ansible 1.8 or later, ensure that fact caching using Redis is not enabled in `ansible.cfg` on the Tower machine.

SUPPORTED PLATFORMS AND REQUIREMENTS

As you prepare for the Tower installation, keep in mind any *Requirements* and *Prerequisites* that may be necessary.

Note that the Tower installation must be run from an internet connected machine that can install software from trusted 3rd-party places such as Ansible's software repository, and your OS vendor's software repositories. In some cases, access to the Python Package Index (PyPI) is necessary as well. If you need to be able to install in a disconnected environment, please contact Ansible support via the Red Hat Customer Portal at <https://access.redhat.com/>.

PLATFORM-SPECIFIC NOTES

6.1 Red Hat Enterprise Linux and CentOS

- PackageKit can frequently interfere with the installation/update mechanism. Consider disabling or removing PackageKit if installed prior to running the setup process.
- Only the “targeted” SELinux policy is supported. The targeted policy can be set to disabled, permissive, or enforcing.

REQUIREMENTS

Ansible Tower has the following requirements:

Note: Tower is a full application and the installation process installs several dependencies such as PostgreSQL, Django, Apache, and others. It is required that you install Tower on a standalone VM or cloud instance and do not co-locate any other applications on that machine (beyond possible monitoring or logging software). Although Tower and Ansible are written in Python, they are not just simple Python libraries. Therefore Tower cannot be installed in a Python virtualenv, a Docker container, or any similar subsystem; you must install it as described in the installation instructions in this guide.

- Supported Operating Systems:
 - Red Hat Enterprise Linux 6 64-bit
 - Red Hat Enterprise Linux 7 64-bit
 - CentOS 6 64-bit
 - CentOS 7 64-bit
 - Ubuntu 12.04 LTS 64-bit
 - Ubuntu 14.04 LTS 64-bit
- An HTML5 compliant web browser
- 2 GB RAM minimum (4+ GB RAM recommended)
 - 2 GB RAM (minimum and recommended for Vagrant trial installations)
 - 4 GB RAM is recommended per 100 forks
- 20 GB hard disk
- 64-bit support required (kernel and runtime)
- For Amazon EC2:
 - Instance size of m3.medium or larger
 - An instance size of m3.xlarge or larger if there are more than 100 hosts
- For HA MongoDB setups:
 - If you plan to run MongoDB, the following guidelines provide a rough estimate for the amount of space required. The basic calculation is:
$$\begin{aligned} &(\text{number of hosts in inventory}) * (\text{number of scans}) * \\ &(\text{average module fact size}) * (\text{number of modules scanning}) \end{aligned}$$

- For example:

hosts = 1,000

number of scans = 365 (1 scan per day for a year)

average module fact size = 100 kb

number of modules = 4 (ansible, packages, services, files)

= 146 GB

The default scan operation has the four (4) modules listed, but you can add your own. Depending on the kinds of modules and the size of the facts you are gathering, that size might be larger.

To help keep the size down, you can use a management job to purge old facts. Refer to [Management Jobs](#) in the *Ansible Tower Administration Guide* for more information

Note: While Ansible Tower depends on Ansible playbooks and requires the installation of the latest stable version of Ansible before installing Tower, manual installations of Ansible are no longer required.

Beginning with Ansible Tower version 2.3, the Tower installation program attempts to install Ansible as part of the installation process. Previously, Tower required manual installations of the Ansible software package before running the Tower installation program. Now, Tower attempts to install the latest stable Ansible package.

If performing a bundled tower installation, the installation program attempts to install Ansible (and its dependencies) from the bundle.

While other operating systems may technically function, currently only the above list is supported to host an Ansible Tower installation. If you have a firm requirement to run Tower on an unsupported operating system, please contact Ansible via the Red Hat Customer Portal at <https://access.redhat.com/>. Management of other operating systems (nodes) is documented by the Ansible project itself and allows for a wider list.

Actual RAM requirements vary based on how many hosts Tower will manage simultaneously (which is controlled by the `forks` parameter in the job template or the system `ansible.cfg` file). To avoid possible resource conflicts, Ansible recommends 4 GB of memory per 100 forks. For example, if `forks` is set to 100, 4 GB of memory is recommended; if `forks` is set to 400, 16 GB of memory is recommended.

A larger number of hosts can of course be addressed, though if the fork number is less than the total host count, more passes across the hosts are required. These RAM limitations are avoided when using rolling updates or when using the provisioning callback system built into Tower, where each system requesting configuration enters a queue and is processed as quickly as possible; or in cases where Tower is producing or deploying images such as AMIs. All of these are great approaches to managing larger environments. For further questions, please contact Ansible via the Red Hat Customer portal at <https://access.redhat.com/>.

The requirements for systems managed by Tower are the same as for Ansible at: http://docs.ansible.com/intro_getting_started.html

PREREQUISITES

Tower is installed using Ansible playbooks. Therefore, you need Ansible to install Tower. Ansible Tower requires Ansible version 1.8 or later.

Ansible can be installed as detailed in the Ansible documentation at: http://docs.ansible.com/intro_installation.html

For convenience, those installation instructions are summarized here:

8.1 Configuration and Installation for Ansible with Red Hat Enterprise Linux and CentOS

The following steps help you configure access to the repository as well as install Ansible.

8.1.1 Configure access to the repository for Ansible with Red Hat Enterprise Linux and CentOS (version 6 or later):

Configure the EPEL repository and any others needed.

As the root user, for Red Hat Enterprise Linux 6 and CentOS 6:

```
root@localhost:~$ yum install http://dl.fedoraproject.org/pub/epel/epel-release-
↳latest-6.noarch.rpm
```

Note:

- For users of Red Hat Enterprise Linux 6, you must enable the optional repository.
-

As the root user, for Red Hat Enterprise Linux 7 and CentOS 7

```
root@localhost:~$ yum install http://dl.fedoraproject.org/pub/epel/epel-release-
↳latest-7.noarch.rpm
```

Note:

- You may also need to enable the extra repository, named `extras` on CentOS 7, `rhel-7-server-extras-rpms` on Red Hat Enterprise Linux 7, and `rhui-REGION-rhel-server-extras` when running in EC2.

- When using the official Red Hat Enterprise Linux 7 marketplace AMI, ensure that the latest `rh-amazon-rhui-client` package that allows enabling the optional repository (named `rhui-REGION-rhel-server-optional` in EC2) is installed.
-

8.1.2 Install Ansible on Red Hat Enterprise Linux and CentOS (version 6 or later):

```
root@localhost:~$ yum install ansible
```

8.2 Configuration and Installation for Ansible with Ubuntu

The following steps help you configure access to the repository as well as install Ansible.

8.2.1 Configure access to the repository for Ansible with Ubuntu 12.04 and Ubuntu 14.04:

As the root user, configure Ansible PPA:

```
root@localhost:~$ apt-get install software-properties-common
root@localhost:~$ apt-add-repository ppa:ansible/ansible
```

8.2.2 Install Ansible on Ubuntu 12.04 and Ubuntu 14.04:

```
root@localhost:~$ apt-get update
root@localhost:~$ apt-get install ansible
```

TOWER INSTALLATION SCENARIOS

Tower can be installed in three scenarios.

- Single Machine integrated installation

This is a single machine install of Tower - the web frontend, REST API backend, and database are all on a single machine. This is the standard installation of Tower. It also installs PostgreSQL from your OS vendor repository, and configures the Tower service to use that as its database.

- Single Machine with an external database

This installs the Tower server on a single machine, and configures it to talk to a remote instance of PostgreSQL as its database. This remote PostgreSQL can be a server you manage, or can be provided by a cloud service such as Amazon RDS.

Tower will not configure replication or failover for the database that it uses, although Tower should work with any replication that you have.

Note: The database server should be on the same network or in the same datacenter as the Tower server for performance reasons.

- High Availability Multi-Machine with an external database

Tower can run in an active-passive high-availability mode. In this mode, Tower will run with one ‘primary’ node active at any time, and any number of passive ‘secondary’ nodes that can be made active if necessary.

Note: Running in a high-availability setup requires any database that Tower uses to be external-Postgres and MongoDB must be installed on a machine that is not one of the primary or secondary tower nodes. When in High Availability mode, the remote Postgres and MongoDB version requirements are *Postgresql 9.4.x* and *mongodb 3.0.x*.

Each of these scenarios can be configured through the Tower Installation Wizard.

GET THE TOWER INSTALLATION PROGRAM

Using the link provided in the email you received based on your interest in Ansible Tower, download and extract the Tower installation tarball.

Note: To obtain a trial version of Ansible Tower, visit: <http://www.ansible.com/tower-trial>

For pricing information, visit: <http://www.ansible.com/pricing>

To download the latest version of Tower directly (note, you must also obtain a license before using this), visit: <http://releases.ansible.com/awx/setup/ansible-tower-setup-latest.tar.gz>

Once extracted, `cd` into the `setup` directory using a command line console. In the following commands, replace the string `VERSION` with the version of Tower that you are installing (e.g., “2.3.0”).

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

10.1 Using Vagrant/Amazon AMI Images

If you use the Vagrant box or Amazon AMI Tower images provided by Ansible, you can find the auto-generated admin password by connecting to the image and reading it from the *message of the day* (MOTD) shown at login.

For Vagrant images, connect via: `vagrant ssh`

For Amazon AMI images, connect via: `ssh root@<your amazon instance>`

You must use the SSH key that you configured the instance to accept at launch time.

10.2 Using the Bundled Tower Installation Program

Beginning in Ansible Tower version 2.3.0, Tower installations can now be performed using a bundled installation program. The bundled installation program is meant for customers who cannot, or would prefer not to, install Tower (and its dependencies) from online repositories. Access to Red Hat Enterprise Linux or Centos repositories is still needed.

To download the latest version of the bundled Tower installation program directly (note, you must also obtain a license before using this), visit: <https://releases.ansible.com/ansible-tower/setup-bundle/>

Note: The bundled installer for Ansible Tower 2.3.0 only supports Red Hat Enterprise Linux and CentOS at this time. Ubuntu support has not yet been added.

Next, select the installation program which matches your distribution (el6 or el7):

```
ansible-tower-setup-bundle-latest.el6.tar.gz
ansible-tower-setup-bundle-latest.el7.tar.gz
```

Note: Red Hat Enterprise Linux customers must enable the following repositories which are disabled by default:

- Red Hat Enterprise Linux 7 users must enable the `extras` repository.
 - Red Hat Enterprise Linux 6 users must enable the `optional` repository.
-

THE TOWER INSTALLATION WIZARD

The Tower setup process consists of two parts—an installation wizard that determines your Tower configuration and a setup playbook that uses that information to install Tower.

The Tower Installation Wizard and the Tower setup playbook do not need to be run from the system that will run Tower, although they can. The Tower Installation Wizard asks for credentials needed to access external systems where necessary.

The Tower Installation Wizard is invoked as `configure` from the path where you unpacked the Tower installation tarball. It writes a file called `tower_setup_conf.yml` which contains the configuration for Tower.

```
username@localhost:~$ ./configure
```

11.1 Installation Arguments

The wizard takes the following arguments:

- `-h, --help` Displays a brief usage summary.
- `-l, --local` Assumes that you are installing Tower on the local machine where you are running `configure`. This implies an internal embedded PostgreSQL database as well. This option skips some questions in the wizard.
- `--no-secondary-prompt` Assumes you are not installing in a high-availability setup. This option skips some questions in the wizard.
- `-A, --no-autogenerate` Do not autogenerate random passwords for PostgreSQL or Redis—prompt the user for them instead.
- `-o FILE, --options-file=FILE` Use the file `FILE` as a source of answers. This can be the `tower_setup_conf.yml` file from a previous run of the wizard. Depending on the contents of the file, this option skips some questions in the wizard. The contents of the `tower_setup_conf.yml` file may look similar to the following:

```
admin_password: password
database: internal
munin_password: password
pg_password: cmTM4eaCpnDS54ReXGv34szoHQiXccFno3atfBi j
primary_machine: localhost
redis_password: 8nG2TRpSDnpr69eWbqFwXTbryCUW64r76VjBqsKx
```

Once you invoke the Tower Installation Wizard, you are asked about the configuration of a few different items.

11.2 Primary Tower machine configuration

First, the Tower wizard asks about where you intend to place the primary (or only) Tower instance.

```
root@localhost:~$ ./configure
-----
Welcome to the Ansible Tower Install Wizard
-----

PRIMARY TOWER MACHINE
Tower can be installed (or upgraded) on this machine, or onto a remote machine
that is reachable by SSH.

Note: If using the High Availability features of Tower, you must use DNS
resolvable hostnames or IP addresses (do not use "localhost").

Enter the hostname or IP to configure Ansible Tower
(default: localhost):
```

If you are installing on the current machine, enter `localhost` or `127.0.0.1` for the current machine. If you are installing on a different machine, enter the IP address or hostname of the machine. This machine must be running and accessible via SSH when running the setup playbook later.

11.2.1 Configuring the Database

Tower can be setup as an internal database installed on the primary Tower machine or as an external PostgreSQL database.

Enter `i` for an internal database on the same machine as Tower, or `e` for an external database. To run Tower in a high-availability configuration, you must use an external database.

```
DATABASE
Tower can use an internal database installed on the Tower machine, or an
external PostgreSQL database. An external database could be a hosted database,
such as Amazon's RDS.

An internal database is fine for most situations. However, to use the High
Availability features of Tower, an external database is required.

If using an external database, the database (but not the necessary tables) must
already exist.

Will this installation use an (i)nternal or (e)ternal database?
```

If you choose to use an external database, the wizard prompts you for the following additional database parameters:

- Database host to connect to
- Database name
- PostgreSQL user to use to access the database
- Password for the above PostgreSQL user
- Port to connect to the PostgreSQL database on (hit enter for the default PostgreSQL port)

The wizard will attempt to verify these parameters if your system has the PostgreSQL client libraries installed.

11.3 Secondary Installation (if applicable)

At this time, if you chose to setup an external database, you can configure any necessary secondary Tower instances.

```
SECONDARY MACHINES
You may optionally elect to add any number of secondary machines, on which
Ansible Tower will also be installed (in secondary mode).
Add secondary machines (y/n)?
```

Enter `y` to configure additional secondary Tower instances.

Enter the hostnames or IP addresses of machines you want to configure as secondary Tower instances, one at a time. Enter a blank line to end the list. These machines must be running and accessible via SSH when running the setup playbook later.

11.4 Passwords

You are then prompted for the passwords you need for various Tower services.

```
PASSWORDS
For security reasons, since this is a new install, you must specify the
following application passwords.
```

The installation wizard prompts you to provide an **Admin Password**. This ‘admin’ password is used for the first user (and superuser) created upon installation. You must have this password for your initial login to Tower.

If you passed the `-A` or `--no-autogenerate` parameters to the Installation Wizard, you are prompted for a PostgreSQL password and a Redis password. These are used internally to Tower and are not needed by the admin at runtime. These additional passwords are normally auto-generated as a random value.

11.4.1 Password Strength Configuration

This feature allows you to define minimum strengths for passwords. Password strength configuration for Tower specifies policies and security mechanisms for providing rules to specify user passwords.

To make use of this feature, set your preferred password strength configuration options in the `local_config.js` file.

- Copy `config.js` in its entirety to the file `local_config.js`.
- Set the `password_*` variables to `true/false` based on the rules you want to enforce.

The password strength configuration feature allows you to create passwords of any combination of upper and lowercase characters, numbers, and special characters that include “!”, “@”, “#”, “\$”, “%”, “^”, “&”, “*”, “(”, and “)”.

You can also set the password’s minimum (1) and maximum (64) length. The recommended minimum password length is eight (8) characters.

11.5 Connection Information

If you chose to install on machines other than the current machine you are running the installation wizard on, you must also enter details on how to connect to those machines.

```
CONNECTION INFORMATION
Enter the SSH user to connect with (default: root):
```

First, you are prompted for the user to SSH to the remote hosts with. If this user is not root, you are prompted for how you will escalate privileges.

```
Root access is required to install Tower.
Will you use (1) sudo or (2) su?
```

Choose either 1 or 2 to configure sudo or su access. If you enter that you need a password for sudo or su access, this will be prompted for during the setup playbook run.

You are then prompted for SSH key information. If you are using a SSH key to access this host, you are prompted for the path to the SSH private key to use.

The same connection and su/sudo information will be used for all machines that are configured by the setup playbook, whether primary or secondary nodes. If you need different connection information for different machines, this can be configured by manually modifying the `inventory` file generated by the Installation Wizard.

11.6 Review and Confirm

You are then asked to review the settings you entered. An example would be:

```
REVIEW
You selected the following options:

The primary Tower machine is: tower.example.com
Tower will operate on an EXTERNAL database.
  host: database.example.com
  database: mydb
  user: db_admin
  password: *****
  port: 5432
Additional secondary machines:
- tower-backup.example.com
- tower-backup2.example.com
Using SSH user: jdoe

Are these settings correct (y/n)?
```

Select `y`, and you will then be given some information on running the setup playbook.

```
FINISHED!
You have completed the setup wizard. You may execute the installation of
Ansible Tower by issuing the following command:

# Add your SSH key to SSH agent.
# You may be asked to enter your SSH unlock key password to do this.
ssh-agent bash
ssh-add ~/.ssh/id_my-example-key
./setup.sh -s
```

11.7 Reviewing the Tower Configuration

The Tower configuration is written into two files by the Tower Installation Wizard.

- `tower_setup_conf.yml` This Tower configuration file contains needed Tower passwords, database connection information, and machine connection information.
- `inventory` This includes the machines that the setup playbook will operate on, grouped into the `primary` and `secondary` groups of nodes.

11.8 The Setup Playbook

The Tower setup playbook is invoked as `setup.sh` from the path where you unpacked the Tower installer tarball. It uses the `tower_setup_conf.yml` and `inventory` files written by the Tower Installation Wizard. The setup script takes the following arguments:

- `-h, --help` Displays a brief usage summary.
- `-c FILE` Use the specified FILE as the Tower configuration file rather than `tower_setup_conf.yml` in the current directory.
- `-i FILE` Use the specified FILE as the inventory for the setup playbook rather than `inventory` in the current directory.
- `-p` Set ansible to prompt for a SSH password when connecting to remote machines
- `-s` Set ansible to prompt for a sudo password on remote machines when installing Tower.
- `-u` Set ansible to prompt for a su password on remote machines when installing Tower.
- `-e` Set additional ansible variables for the playbook to use either in `key=value` or `YAML/JSON` form. This should not be needed in normal operation.
- `-b` Perform a database backup in lieu of installing.
- `-r BACKUP_FILE` Perform a database restore in lieu of installing.

Depending on the configuration you entered when running the Tower Installation Wizard, it may have prompted you to run the setup playbook with some combination of `-p`, `-s`, or `-u`.

After calling `setup.sh` with the appropriate parameters, Tower is installed on the appropriate machines as has been configured.

UPGRADING AN EXISTING TOWER INSTALLATION

You can upgrade your existing Tower installation to the latest version easily.

As with installation, the upgrade process requires that the Tower server be able to access the Internet. The upgrade process takes roughly the same amount of time as a Tower installation, plus any time needed for data migration.

This upgrade procedure assumes that you have a working installation of Ansible and Tower.

Note: You can not convert an embedded-database Tower to a High Availability installation as part of an upgrade. Users who want to deploy Tower in a High Availability configuration should back up their Tower database, install a new HA configuration on a different VM or physical host, and then restore the database. It is possible to add a primary or secondary instance later on to Tower if it is already operating on an external database. Refer to the [High Availability](#) chapter of the *Tower Administration Guide*.

12.1 Requirements

Before upgrading your Tower installation, refer to [Requirements](#) to ensure you have enough disk space and RAM as well as to review any software needs. For example, you should have the latest stable release of Ansible installed before performing an upgrade.

12.2 Backing Up Your Tower Installation

It is advised that you create a backup before upgrading the system. After the backup process has been accomplished, proceed with OS/Ansible/Tower upgrades.

Refer to [Backing Up and Restoring Tower](#) in the *Ansible Tower Administration Guide*.

12.3 Get the Tower Installer

Download the Ansible Tower install/upgrade tool: <http://releases.ansible.com/ansible-tower/setup/>

Extract it, then `cd` into the `setup` directory. Replace the string `VERSION` in the commands below with the version of Tower that you are installing e.g., “2.1.4”.

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

12.4 Run the Tower Installation Wizard

To configure your upgrade, you run the same Tower Installation Wizard as you would for an installation.

12.4.1 Simplified Tower Upgrade

If you are upgrading a Tower instance on a local machine with an internal database, you can bypass many of the questions by invoking the `configure` script as `./configure --local`

```

root@localhost:~$ ./configure --local
-----
Welcome to the Ansible Tower Install Wizard
-----

This wizard will guide you through the setup process.

LOCAL INSTALLATION
You are installing Ansible Tower on this machine, using an internal database.

REVIEW
You are UPGRADING an existing Tower installation on localhost.

Are these settings correct (y/n)?

Confirm that the settings are correct.

```

12.4.2 Upgrade Using an existing Settings File

If you have the `tower_setup_conf.yml` file from when you installed Tower, you can pass it to the `configure` script:

```

root@localhost:~$ ./configure -o tower_setup_conf.yml
-----
Welcome to the Ansible Tower Install Wizard
-----

This wizard will guide you through the setup process.

The configuration provided in /home/tower/ansible-tower-setup-2.1.4/tower_setup_conf.
↪.yml appears complete.

FINISHED!
You have completed the setup wizard. You may execute the installation of
Ansible Tower by issuing the following command:

sudo ./setup.sh

```

12.4.3 Upgrade Interactively

Alternatively, you can walk through the upgrade process. Invoke the `configure` script:

```

root@localhost:~$ ./configure
-----
-----
Welcome to the Ansible Tower Install Wizard
-----

This wizard will guide you through the setup process.

PRIMARY TOWER MACHINE
Tower can be installed (or upgraded) on this machine, or onto a remote machine
that is reachable by SSH.

Note: If using the High Availability features of Tower, you must use DNS
resolvable hostnames or IP addresses (do not use "localhost").

Enter the hostname or IP to configure Ansible Tower
(default: localhost):

```

Once you enter the host, the Tower Installation Wizard contacts it to determine if Tower is installed. If it is a functioning Tower installation, it determines the current Tower configuration, including the location of any secondary nodes that need upgrading.

The Installation Wizard then asks you for connection details for connecting to any remote machines. Enter any needed SSH user, SSH key location, and whether `sudo` or `su` is in use.

12.5 The Setup Playbook

The Tower setup playbook script is invoked as `./setup.sh` from the path where you unpacked the Tower installer tarball. It uses the `tower_setup_conf.yml` and `inventory` files written by the Tower Installation Wizard. The setup script takes the following arguments:

- `-h, --help` Displays a brief usage summary.
- `-c FILE` Use the specified FILE as the Tower configuration file rather than `tower_setup_conf.yml` in the current directory.
- `-i FILE` Use the specified FILE as the inventory for the setup playbook rather than `inventory` in the current directory.
- `-p` Set ansible to prompt for a SSH password when connecting to remote machines
- `-s` Set ansible to prompt for a sudo password on remote machines when upgrading Tower.
- `-u` Set ansible to prompt for a su password on remote machines when upgrading Tower.
- `-e` Set additional ansible variables for the playbook to use either in `key=value` or `YAML/JSON` form. This should not be needed in normal operation.
- `-b` Perform a database backup in lieu of installing.
- `-r BACKUP_FILE` Perform a database restore in lieu of installing.

Depending on the configuration you entered when running the Tower Installation Wizard, it may have prompted you to run the setup playbook with some combination of `-p`, `-s`, or `-u`.

As the root user, call `setup.sh` with the appropriate parameters and Tower is upgraded on the appropriate machines as configured.

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

Note: As part of the upgrade process, database schema migration may be done. Depending on the size of your Tower installation, this may take some time.

If the upgrade of Tower fails or if you need assistance, please contact Ansible via the Red Hat Customer Portal at <https://access.redhat.com/>.

SUPPORTED LOCALES

Ansible Tower supports the following locales for UTC-friendly date and time information.

Tower automatically sets the locale preference based on the user's browser settings. For Safari, Internet Explorer, and older versions of Chrome as well as FireFox, this is handled automatically.

For newer versions of Chrome (v32 and later) and FireFox (v32 and later), Tower uses the language preferences set from your browser's language settings. The browser lists the user's preferred languages and selects the first in the array as the user's top choice, which Tower uses as the preferred locale. This means that you can change your browser's language and change your Tower locale preferences (although you may need to reload/refresh Tower in your browser to see this change.)

- az – Cyrillic
- bg – Bulgarian
- bs – Bosnian
- ca – Catalan
- cs – Czech
- da – Danish
- de – German
- el – Greek
- en-gb – English (United Kingdom)
- es – Spanish
- et – Estonian
- eu – Basque
- fa – Persian
- fi – Finnish
- fo – Faroese
- fr – French
- gl – Galician
- he – Hebrew
- hr – Croatian
- hu – Hungarian
- id – Indonesian

- `is` – Icelandic
- `it` – Italian
- `ja` – Japanese
- `ka` – Georgian
- `lt` – Lithuanian
- `lv` – Latvian
- `mk` – Macedonian
- `nb` – Norwegian
- `nl` – Dutch
- `pl` – Polish
- `pt-br` – Portuguese (Brazil)
- `pt` – Portuguese
- `ro` – Romanian
- `ru` – Russian
- `sk` – Slovak
- `sl` – Slovenian
- `sq` – Albanian
- `sr` – Serbian
- `sv` – Swedish
- `th` – Thai
- `tr` – Turkish
- `uk` – Ukrainian
- `vi` – Vietnamese
- `zh-cn` – Chinese (simplified)
- `zh-tw` – Chinese (traditional)

GLOSSARY

Ad Hoc Refers to running Ansible to perform some quick command, using `/usr/bin/ansible`, rather than the orchestration language, which is `/usr/bin/ansible-playbook`. An example of an ad-hoc command might be rebooting 50 machines in your infrastructure. Anything you can do ad-hoc can be accomplished by writing a playbook, and playbooks can also glue lots of other operations together.

Credentials Authentication details that may be utilized by Tower to launch jobs against machines, to synchronize with inventory sources, and to import project content from a version control system.

Facts Facts are simply things that are discovered about remote nodes. While they can be used in playbooks and templates just like variables, facts are things that are inferred, rather than set. Facts are automatically discovered when running plays by executing the internal **setup** module on the remote nodes. You never have to call the setup module explicitly, it just runs, but it can be disabled to save time if it is not needed. For the convenience of users who are switching from other configuration management systems, the fact module also pulls in facts from the ‘ohai’ and ‘facter’ tools if they are installed, which are fact libraries from Chef and Puppet, respectively.

Forks Ansible and Tower talk to remote nodes in parallel and the level of parallelism can be set several ways—during the creation or editing of a **Job Template**, by passing `--forks`, or by editing the default in a configuration file. The default is a very conservative 5 forks, though if you have a lot of RAM, you can easily set this to a value like 50 for increased parallelism.

Group A set of hosts in Ansible that can be addressed as a set, of which many may exist within a single Inventory.

Host A system managed by Tower, which may include a physical, virtual, cloud-based server, or other device. Typically an operating system instance. Hosts are contained in Inventory. Sometimes referred to as a “node”.

Host Specifier Each Play in Ansible maps a series of tasks (which define the role, purpose, or orders of a system) to a set of systems. This “hosts:” directive in each play is often called the hosts specifier. It may select one system, many systems, one or more groups, or even some hosts that are in one group and explicitly not in another.

Inventory A collection of hosts against which Jobs may be launched.

Inventory Script A very simple program (or a complicated one) that looks up hosts, group membership for hosts, and variable information from an external resource—whether that be a SQL database, a CMDB solution, or something like LDAP. This concept was adapted from Puppet (where it is called an “External Nodes Classifier”) and works more or less exactly the same way.

Inventory Source Information about a cloud or other script that should be merged into the current inventory group, resulting in the automatic population of Groups, Hosts, and variables about those groups and hosts.

Job One of many background tasks launched by Tower, this is usually the instantiation of a Job Template; the launch of an Ansible playbook. Other types of jobs include inventory imports, project synchronizations from source control, or administrative cleanup actions.

Job Detail The history of running a particular job, including its output and success/failure status.

Job Template The combination of an Ansible playbook and the set of parameters required to launch it.

JSON Ansible and Tower use JSON for return data from remote modules. This allows modules to be written in any language, not just Python.

Organization A logical collection of Users, Teams, Projects, and Inventories. The highest level in the Tower object hierarchy is the Organization.

Organization Administrator An Tower user with the rights to modify the Organization's membership and settings, including making new users and projects within that organization. An organization admin can also grant permissions to other users within the organization.

Permissions The set of privileges assigned to Users and Teams that provide the ability to read, modify, and administer Projects, Inventories, and other Tower objects.

Plays A playbook is a list of plays. A play is minimally a mapping between a set of hosts selected by a host specifier (usually chosen by groups, but sometimes by hostname globs) and the tasks which run on those hosts to define the role that those systems will perform. There can be one or many plays in a playbook.

Playbook An Ansible playbook. Refer to <http://docs.ansible.com/> for more information.

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

Roles Roles are units of organization in Ansible and Tower. Assigning a role to a group of hosts (or a set of groups, or host patterns, etc.) implies that they should implement a specific behavior. A role may include applying certain variable values, certain tasks, and certain handlers—or just one or more of these things. Because of the file structure associated with a role, roles become redistributable units that allow you to share behavior among playbooks—or even with other users.

Schedule The calendar of dates and times for which a job should run automatically.

Sudo Ansible does not require root logins and, since it is daemonless, does not require root level daemons (which can be a security concern in sensitive environments). Ansible can log in and perform many operations wrapped in a `sudo` command, and can work with both password-less and password-based sudo. Some operations that do not normally work with `sudo` (like `scp` file transfer) can be achieved with Ansible's *copy*, *template*, and *fetch* modules while running in `sudo` mode.

Superuser An admin of the Tower server who has permission to edit any object in the system, whether associated to any organization. Superusers can create organizations and other superusers.

Survey Questions asked by a job template at job launch time, configurable on the job template.

Team A sub-division 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.

User An Tower operator with associated permissions and credentials.

YAML Ansible and Tower use YAML to define playbook configuration languages and also variable files. YAML has a minimum of syntax, is very clean, and is easy for people to skim. It is a good data format for configuration files and humans, but is also machine readable. YAML is fairly popular in the dynamic language community and the format has libraries available for serialization in many languages (Python, Perl, Ruby, etc.).

- `genindex`

This document is Copyright © 2015 Ansible, Inc. All rights reserved.

Ansible and Ansible Tower are trademarks of Ansible, Inc.

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

Third Party Rights

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

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

Amazon AMI image, 19
 Ansible, configure repository access, 16, 17
 Ansible, installation, 17

B

bundled installer, 19
 known issues, 8

C

CentOS, 16
 components
 licenses, 3
 configuration review
 inventory, 24
 tower_setup_conf.yml, 24

D

DEB files
 licenses, 3
 download Ansible Tower, 19

E

external database
 installation single machine, 18

G

glossary, 32

H

high availability, external database
 installation multi-machine, 18
 host comparisons
 known issues, 8

I

installation
 general notes, 11
 multi-machine high availability, external database,
 18
 platform-specific notes, 13

 requirements, 12
 scenarios, 18
 single machine external database, 18
 single machine integrated, 18
 supported platforms, 12
 installation bundle
 licenses, 3
 installation fail
 known issues, 8
 installation prerequisites, 16
 installation program, 19
 upgrade, 26
 installation requirements, 14
 installation wizard, 21
 arguments, 21
 configuration review, 24
 confirmation, 24
 connection information, 23
 database configuration, 22
 execute Tower, 24
 passwords, 23
 playbook setup, 25
 playbook setup arguments, 25
 primary machine configuration, 22
 secondary installation, 23
 ssh key, 24
 upgrade, 27
 integrated
 installation single machine, 18
 issues, known, 8

K

known issues, 8
 bundled installer, 8
 host comparisons, 8
 installation fail, 8
 live event statuses, 8
 MongoDB, 8
 object editing, 8
 PRoot, 8
 ssh customization, 8
 su command, 8

- sudo/su, 8
- traceback error, 8
- Ubuntu, 8
- YAML traceback error, 8

L

- license, 2
 - features, 3
 - trial, 2
 - types, 2
- licenses
 - components, 3
 - DEB files, 3
 - installation bundle, 3
 - RPM files, 3
- live event statuses
 - known issues, 8
- locales supported, 30

M

- MongoDB
 - known issues, 8
- multi-machine
 - high availability, external database, installation, 18

O

- object editing
 - known issues, 8

P

- passwords
 - strength configuration, 23
- platform-specific notes
 - CentOS, 13
 - Red Hat Enterprise Linux, 13
- playbook setup, 25
 - arguments, 25
 - setup.sh, 25
- prerequisites, 16
- PRoot
 - known issues, 8

R

- Red Hat Enterprise Linux, 16, 17
- release notes, 4
 - release notes, v1.4.10, 7
 - release notes, v1.4.11, 7
 - release notes, v1.4.12, 6
 - release notes, v1.4.8, 7
 - release notes, v1.4.9, 7
 - release notes, v2.0, 6
 - release notes, v2.0.1, 6
 - release notes, v2.0.2, 6
 - release notes, v2.0.5, 6

- release notes, v2.1, 6
 - release notes, v2.1.1, 5
 - release notes, v2.1.2, 5
 - release notes, v2.1.3, 5
 - release notes, v2.1.4, 5
 - release notes, v2.2, 5
 - release notes, v2.2.1, 4
 - release notes, v2.2.2, 4
 - release notes, v2.3, 4
 - release notes, v2.3.1, 4
- requirements, 14
 - installation, 12
- RPM files
 - licenses, 3

S

- single machine
 - external database, installation, 18
 - integrated, installation, 18
- ssh customization
 - known issues, 8
- su command
 - known issues, 8
- sudo/su
 - known issues, 8
- support, 2
- supported platforms
 - installation, 12

T

- traceback error
 - known issues, 8

U

- Ubuntu, 17
 - known issues, 8
- updates, 2
- upgrade, 26
 - existing settings file, 27
 - installation program, 26
 - installation wizard, 27
 - interactive, 27
 - playbook setup, 28

V

- Vagrant image, 19

Y

- YAML traceback error
 - known issues, 8