
Ansible Tower Release Notes

Release Ansible Tower 3.4.2

Red Hat, Inc.

May 05, 2021

CONTENTS

1	Release Notes	2
1.1	Release Notes for 3.4.x	2
1.2	Release Notes for 3.3.x	4
1.3	Release Notes for 3.2.x	7
1.4	Release Notes for 3.1.x	14
1.5	Release Notes for 3.0.x	19
1.6	Release Notes for 2.4.x	23
1.7	Release Notes for 2.3.x	25
1.8	Release Notes for 2.2.x	26
1.9	Release Notes for 2.1.x	27
1.10	Release Notes for 2.0.x	28
1.11	Ansible Tower Version 2.0.2	28
1.12	Release Notes for 1.4.x	29
2	Known Issues	30
2.1	Isolated nodes unsupported in an OpenShift deployment	31
2.2	Tower login via HTTP requires workaround	31
2.3	Job slicing and limit interactions	31
2.4	Default LDAP directory must be configured to use LDAP Authentication	31
2.5	Lost Isolated Jobs	31
2.6	Potential security issue using X_FORWARDED_FOR in REMOTE_HOST_HEADERS	31
2.7	Server error when accessing SAML metadata via hostname	32
2.8	Live events status indicators	32
2.9	VMWare Self-Signed Certs	32
2.10	Tower Database on Disk Becomes Corrupted	33
2.11	Safari unable to establish connection to web socket	33
2.12	Ansible Azure Dependencies	34
2.13	sudo/su not working as expected for local playbooks or playbooks with local_actions	34
2.14	Problems when using SSH customization	34
2.15	Ubuntu unsupported for bundled installations	35
2.16	Reactivating OAuth authentication accounts which have been deleted	35
2.17	Using vaulted variables in inventory sourced from a project	35
2.18	Using Tower credential with Tower 3.3+ Tokens	35
2.19	Saved scheduled and workflow configurations and surveys	36
3	Supported Locales	37
4	Index	39
5	Copyright © 2019 Red Hat, Inc.	40

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

The *Ansible Tower Release Notes* provides release notes, known issues, and related reference materials. This document has been updated to include information for the latest release of Ansible Tower 3.4.2.

We Need Feedback!

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

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

Ansible Tower Version 3.4.2; March 06, 2019; <https://access.redhat.com/>

RELEASE NOTES

1.1 Release Notes for 3.4.x

1.1.1 Ansible Tower Version 3.4.2

- Added a number of notable improvements to translations
- Fixed the Notifications tab to properly appear when viewing a Job Template
- Fixed the total host count to be correct when viewing job output
- Fixed the `awx_workflow_job_name` extra variable to be correct for jobs launched via nested workflows
- Fixed scheduled jobs to no longer run twice in clustered Tower installations
- Fixed the display of playbooks that utilize Ansible's `serial` keyword
- Fixed instance group assignments to restore correctly after Tower is backed up and restored
- Fixed a number of workflow builder user interface issues
- Fixed the Job Template credential selection to display all available options when large numbers of Credential Types are in use
- Fixed the stdout display to no longer delay when Smart Inventories are in use
- Fixed a bug that prevented Ansible to upgrade when Tower is updated with `-e upgrade_ansible_with_tower=1`
- Fixed the installer to properly upgrade Tower dependencies on upgrade
- Fixed a performance issue in Ansible Tower 3.4.0 which caused a notable delay in job start time for large inventories (>1000 hosts).
- Fixed Tower logs to now properly record failed logins
- Updated `settings.SCHEDULE_MAX_JOBS` so that it also limits workflow Job Templates

1.1.2 Ansible Tower Version 3.4.1

- Fixed an issue which could cause a running job to not be marked as “failed” during a database outage
- Fixed an issue where X-Frame-Options was not set for OpenShift installations
- Fixed a potential deadlock when LDAP authentication backends are in use
- Fixed an issue where Insights configuration would be broken on upgrade
- Fixed an issue where the callback receiver would crash with an InterfaceError
- Fixed an issue where the proper hashed PostgreSQL password would not be used when installing on a FIPS-enabled system
- Fixed an issue with OS version detection when used with Ansible 2.8pre
- Updated social-auth-core dependency to fix Google OAuth when Google+ is deactivated

1.1.3 Ansible Tower Version 3.4.0

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

- Improved validation survey specs on save, returning a 400 response in many new cases
- Improved the Workflow job nodes with missing `unified_job_template` to no longer unconditionally fail a workflow. Instead, the workflow job node is treated as a job that ran with status result of failure.
- Improved the Workflow job failure semantics. Before, if any leaf node failed then the workflow job was considered failed, else success. Now, if any node in the workflow fails without a failure handle path, then the workflow job is considered failed; else success.
- Improved workflows to now allow convergence nodes (nodes with multiple parents)
- Updated the environment variables used in job runs, which include the removal of `INVENTORY_HOSTVARS` and the ability for the inventory plugin and inventory unparsed behavior to be configurable
- Updated how project “sync” jobs (as opposed to project updates) work
- Updated the behavior of the Activity Stream associated with deleted inventory and displayed list of credentials used in job creation entry
- Updated the creation of schedules from the API to allow global schedules lists
- Updated GCE instances by moving from `GCE_PEM_FILE_PATH` to `GCE_CREDENTIALS_FILE_PATH`
- Updated Tower to allow use of organization roles with teams in the user interface
- Updated Configure Tower from the Settings menu to contain sub-menus and access to sub-tasks

1.2 Release Notes for 3.3.x

1.2.1 Ansible Tower Version 3.3.4

- Fixed an issue where X-Frame-Options was not set for OpenShift installations
- Fixed a potential deadlock when LDAP authentication backends are in use
- Fixed an issue where the callback receiver would crash with an `InterfaceError`
- Updated social-auth-core dependency to fix Google OAuth when Google+ is deactivated
- Updated `asgi_amqp` dependency to fix an issue with websockets

1.2.2 Ansible Tower Version 3.3.3

- Fixed a RabbitMQ misconfiguration that decreased cluster stability and could allow for unauthorized access (CVE-2018-16879)
- Fixed RabbitMQ HA policy when deployed in OpenShift

1.2.3 Ansible Tower Version 3.3.2

- Fixed external loggers to send activity stream to send changes as raw JSON, instead of JSON-ified string
- Fixed `MANAGE_ORGANIZATION_AUTH` settings to allow superusers to make changes
- Fixed an XSS issue when viewing application tokens
- Fixed permissions to allow project admins to create projects
- Fixed permissions to allow job template admins to delegate permissions to certain users/teams

- Fixed Tower to allow selecting credential types when running in a non-English language
- Fixed the Activity Stream to properly note credentials and custom credentials
- Fixed an error resulting from a credential lookup where there are multiple custom credential types defined
- Fixed certain certificate validation issues
- Fixed templates to allow filtering by template type
- Fixed events with no output to no longer render excess blank lines in the job display
- Fixed detailed events to show hosts that were added during a playbook run
- Fixed the inability to properly mark Tower nodes as disabled when rabbitmq was offline
- Fixed the job output to correctly display certain job events
- Fixed the job events to properly line-wrap in the job output
- Fixed the pagination to no longer redirect to the dashboard when selecting multiple credentials
- Fixed a potential deadlock on inventory deletion
- Fixed the job relaunch operation to use the correct credentials when the defined credentials were changed
- Improved the Schedules view to more clearly denote which resource was being scheduled
- Updated Google Cloud credentials to be passed as `GCE_PEM_FILE_PATH` to `GCE_CREDENTIALS_FILE_PATH` in accordance with recent Ansible versions

1.2.4 Ansible Tower Version 3.3.1

- Fixed event callback error when in-line vaulted variables are used with `include_vars`
- Fixed HSTS and X-Frame-Options to properly be set in nginx configuration
- Fixed isolated node setup to no longer fail when `ansible_host` is used
- Fixed selection of custom virtual environments in job template creation
- Fixed display of `extra_vars` for scheduled jobs
- Fixed websockets for job details to properly work
- Fixed the `/api/v2/authtoken` compatibility shim.
- Fixed page size selection on the jobs screen
- Fixed instances in an instance group to properly be disabled in the user interface
- Fixed the job template selection in workflow creation to properly render
- Fixed `member_attr` to properly set on some LDAP configurations during upgrade, preventing login
- Fixed `PosixUIDGroupType` LDAP configurations
- Improved the RAM requirement in the installer preflight check
- Updated Tower to properly report an error when relaunch was used on a set of failed hosts that is too large
- Updated sosreport configuration to gather more python environment, nginx, and supervisor configuration

1.2.5 Ansible Tower Version 3.3.0

- Added support for container-based clusters using OpenShift or Kubernetes
- Added support for multiple or no credential assignment in Job Templates
- Added support for multiple Vault credential assignment in Job Templates
- Added support for multiple LDAP server configurations
- Added support fact caching for isolated instances
- Added the ability to schedule configurations of a job template using job template prompts
- Added deprecation note to the Tower REST API for Version 1 (`/api/v1/`), which will be removed in a future release of Ansible Tower
- Added the ability to make a copy of existing Tower objects (inventory, project, etc.) as a template for creating a new one
- Added the ability to relaunch jobs on a subset of hosts by status
- Added validation to prevent string "`$encrypted$`" from becoming a literal survey question default
- Added support for more job template prompts at workflow node creation time
- Added `ask_variables_on_launch` to workflow job templates (WFJT)
- Added `diff_mode` and `verbosity` fields to WFJT nodes
- Added Saved Launch-time configurations feature - added WFJT node promptable fields to schedules, added `extra_data` to WFJT nodes, and added "schedule this job" endpoint.
- Added block creation of schedules when variables not allowed are given. Block similar cases for WFJT nodes.
- Added the ability to create instance groups and associate instances at runtime via the user interface and API
- Added the ability to group instances based on policy, such as such as "50% of instances" and "at least three instances"
- Added additional organizational roles for administration of projects, job templates, inventories, workflows and more
- Added support for custom virtual environments for customizing Ansible execution
- Added OAuth2 support for token based authentication
- Added support for OAuth2 applications and access token generation
- Added the ability to forcibly expire sessions through `awx-manage expire_sessions`
- Added support for making inventory parsing errors fatal, and only enable the `script` inventory plugin for job runs and vendored inventory updates
- Added `inventory` field to inventory updates
- Added related credentials endpoint for inventory updates to be more internally consistent with job templates, model changes
- Added the ability to show all teams to organization admins if the `ORG_ADMINS_CAN_SEE_ALL_USERS` setting is enabled
- Added the ability to create schedules and workflow nodes from job templates that use credentials which prompt for passwords if `ask_credential_on_launch` is set.
- Deprecated the `fact_versions` and `fact_view` endpoints from the API, including `OPTIONS`
- Deprecated fact tables

- Deprecated the `awx-manage cleanup_facts` command for fact cleanup
- Deprecated the `/api/v2/authtoken/` endpoint in the API and replaced it with `/api/v2/tokens/`
- Fixed a conflict with Tower credential type by removing `TOWER_HOST` as a default environment variable in job running environments. Playbook authors should replace their use with `AWX_HOST`.
- Fixed a behavior in Tower to prevent it from deleting jobs when event processing is still ongoing
- Fixed a behavior in Tower to disallow relaunching jobs with `execute_role` if another user provided prompts
- Improved project updates so that previously synced git projects do not attempt to contact the server if they are already at the proper revision
- Improved WFJT node `credential` to many-to-many `credentials`
- Improved stricter criteria to admin users where organization admin role now necessary for all organizations of which the target user is a member. Additionally, removed unused `admin_role` associated with users
- Improved logs to consistently catch task exceptions
- Improved external loggers to passively create handler from settings on every log emission, replacing server restart, allowing use in OpenShift deployments
- Improved Tower to automatically run a project update if sensitive fields change like `scm_url`
- Improved queuing logic through setting `execution_node` in task manager and submitting waiting jobs to only the queue for the specific instance job is targeted to run on
- Updated the `auth-token-timeout` header name to `Session-Timeout`
- Updated the `AUTH_TOKEN_EXPIRATION` setting name to change to `SESSION_COOKIE_AGE` and `AUTH_TOKEN_PER_USER` changed to `SESSIONS_PER_USER`
- Updated source-control based inventory to allow for vaulted variable values
- Updated the minimum required version of Red Hat Enterprise Linux to 7.4
- Updated the minimum required RAM for standalone Tower to 4GB
- Updated Ansible Tower to set `ANSIBLE_DISPLAY_ARGS_TO_STDOUT` to `False` by default for all playbook runs to match Ansible's default behavior. See [Jobs](#) for more information.
- Updated all job and tasks to generate consistent output events and make job output available on all cluster nodes
- Updated external logging to default to HTTPS unless `http://` is explicitly specified in the log aggregator hostname
- Updated the behavior of a job template to prohibit configuring callbacks on job templates without an inventory
- Updated the boolean fields for custom credential types to always default `extra_vars` and environment variables to `False` when a value is not provided
- Updated to disallow using HTTP PUT/PATCH methods to modify existing jobs in Job Details API endpoint

1.3 Release Notes for 3.2.x

1.3.1 Ansible Tower Version 3.2.8

- Fixed using `include_vars` with vaulted variables to properly handle `AnsibleVaultEncryptedUnicode` objects in the callback receiver

1.3.2 Ansible Tower Version 3.2.7

- Fixed Smart Inventory filters to no longer filter by the content of sensitive fields

1.3.3 Ansible Tower Version 3.2.6

- Fixed Tower callback plugin handling of `v2_playbook_on_notify` events
- Fixed potential information leakage via websocket
- Fixed a CSRF vulnerability in Tower (CVE-2018-10884)
- Fixed editing a job template to no longer overwrite API-only settings
- Fixed certain cluster topologies to no longer cause duplicate project updates
- Fixed unauthorized credentials to no longer be associated with projects and inventory sources
- Updated oVirt client libraries to work with Ansible 2.5 or later

1.3.4 Ansible Tower Version 3.2.5

- Fixed a RabbitMQ configuration issue that would affect cluster recovery on network interruptions

1.3.5 Ansible Tower Version 3.2.4

- Added `UI_LIVE_UPDATES_ENABLED` setting for disabling websocket updates outside of job output
- Fixed organization admins to no longer be able to modify users by adding them to their organization (CVE-2018-1101)
- Fixed Tower to disable usage of Jinja templates in launch-time variables for security reasons (CVE-2018-1104). This release introduces the `ALLOW_JINJA_IN_EXTRA_VARS` configuration parameter for Tower. This parameter has three values: `template` to allow usage of Jinja saved directly on a job template definition (the default), `never` to disable all Jinja usage (recommended), and `always` to always allow Jinja (strongly discouraged, but an option for prior compatibility). Note that the `always` option is deprecated, and will be removed in a future Tower release.
- Fixed sanitization of module arguments with implicit `no_log`
- Fixed Smart Inventories to no longer run on hosts marked as disabled
- Fixed Fact Caching documentation to no longer refer to memcached
- Updated bundled python-saml for CVE-2017-11427
- Updated memcached to now listen on a local Unix socket instead of a TCP socket

1.3.6 Ansible Tower Version 3.2.3

- Added deprecation warning when installing on certain older operating systems, such as Ubuntu 14.04, which will be removed in a future release
- Fixed Inventory Updates to properly save `group_vars` inside of Tower group variables when used with Ansible 2.5 or later
- Fixed certain Inventory Updates to no longer fail when running against isolated nodes
- Fixed the ability to customize `ANSIBLE_LIBRARY` when Job Template fact caching is enabled
- Fixed fact cache data to no longer prematurely expire for Job Templates with large amounts of fact data
- Fixed isolated job runs to no longer fail when the playbook contained certain Unicode characters
- Fixed the installer to use the correct package version when running isolated Tower nodes
- Fixed Slack notification issues
- Fixed workflow artifacts to no longer periodically go missing in subsequent workflow nodes
- Fixed the Tower web interface to support large numbers of custom Credential Types
- Fixed the “Test” button when configuring UDP-based external logging
- Fixed the database restoration process that affected users with embedded PostgreSQL databases
- Fixed a few XSS vulnerabilities in the Tower web interface
- Fixed the ability to provide the admin password in the MOTD file for the Vagrant and AMI images

1.3.7 Ansible Tower Version 3.2.2

- Added support for Ansible Tower and Red Hat Virtualization credentials.
- Added dynamic inventory scripts for Ansible Tower and Red Hat Virtualization
- Added `awx_*` extra variables to job runs in addition to `tower_*`
- Added a setting for maximum user interface job events to show to Tower configuration
- Added support for setting the Azure Cloud Environment in Azure credentials
- Added retry for cleaning up job artifacts from isolated nodes
- Added python-crypto requirement to RPM packaging for GCE inventory script
- Added rsync requirement to RPM packaging for isolated nodes
- Added error handling in installation for PostgreSQL 9.4 to 9.6 migration failures
- Removed unused `CALLBACK_CONNECTION`, `CALLBACK_QUEUE`, and `JOB_CALLBACK_DEBUG` environment variables from the job environment
- Fixed multiple issues where survey passwords were not properly encrypted in the database
- Fixed an issue where cleanup jobs could run slowly and exhaust system memory when large job output was present
- Fixed an issue where cleanup jobs could fail due to a race condition
- Fixed an issue where use of `remove: True` and `remove_users: True` in LDAP configuration would cause an excessive number of activity stream entries
- Fixed an issue where the GCE inventory script would erroneously cache information

- Fixed an issue when using Ipsilon as a SAML IdP
- Fixed an issue when using SAML authentication behind a load-balancer
- Fixed an issue where '+' in a search string was not handled properly
- Fixed an issue where non-alphanumeric characters were stripped from SAML usernames
- Fixed an issue where credential_type information appeared in api/v1 output
- Fixed a styling issue for Host Config Key in the Job Template display
- Fixed an issue where it was impossible to remove an organization from a credential
- Fixed an issue where overwrite_vars on an inventory source would overwrite inventory toplevel variables
- Fixed an issue where some credential kinds were not properly shown in the user interface
- Fixed calculation of isolated instance capacity
- Fixed an issue where the 'Workflow Editor' and 'Survey Editor' buttons were incorrectly shown in some states
- Fixed navigation to additional pages of hosts in the Smart Inventory view
- Fixed an issue where CloudForms inventory would not work with process isolation
- Fixed an issue where job output would not properly word wrap
- Fixed a migration issue with unicode inventory source names
- Fixed an issue when launching an ad-hoc command with forbidden extra variables
- Fixed an issue with symlinked manual projects when used with process isolation
- Fixed an issue where some host_filter queries could not be removed
- Fixed an issue where non-ascii characters could not be used in a LDAP bind DN
- Fixed sizing of the ad-hoc command launch dialog
- Fixed an issue where <https://github.com/ansible/ansible/issues/30064> would prevent project sync
- Fixed an issue where a Smart Inventory host_filter query would be improperly encoded when saved
- Fixed month name on dashboard chart
- Fixed scheduling error when browser is in UTC timezone
- Fixed autocompletion of SCM inventory file dropdown
- Fixed modal state handling when a modal dialog was closed by clicking outside of it
- Fixed assorted migration errors on upgrade
- Fixed a user interface error when rapidly deleting inventory groups
- Fixed an issue where the system auditor would get a 404 error when viewing job results
- Fixed assorted issues when cascading job cancellation to dependent jobs
- Fixed opacity of disabled 'Run Commands' and 'Smart Inventory' buttons
- Fixed 'total_hosts' field of Smart Inventories
- Fixed virtualenv paths in sosreport plugins
- Fixed installation with Ansible 2.2
- Fixed ownership on ha.py on installation
- Fixed django superuser check in installation

- Fixed setting of custom RabbitMQ AMQP ports during installation
- Fixed an issue where LDAP authentication could timeout or cause a Tower error
- Improved callback worker's ability to deal with idle or disconnected database connections
- Improved activity stream output for Tower configuration changes
- Improved deletion of inventory sources to properly delete imported hosts and groups
- Improved various error messages
- Improved initial zoom setting of workflow view
- Improved inline help popovers for credential types
- Improved configuration for SSH key handling for isolated nodes. This is now configurable during setup
- Improved preflight checks for cluster installation
- Improved backup/restore playbooks to be cluster-aware
- Improved error handling in backup/restore playbooks
- Updated translations for Dutch, French, Japanese, and Spanish

1.3.8 Ansible Tower Version 3.2.1

- Added support to enforce Tower software version consistency across clustered environments
- Fixed an issue where, when using Tower 3.2.0 + Ansible 2.4.0, creating a Job Template that used an inventory with fact caching enabled could cause the job to run against a host which should have been removed
- Fixed a problem where ad-hoc permissions could be used to run commands against the Tower server
- Fixed an issue where the migration of scan jobs failed due to an organization having a unicode character in the name
- Fixed an issue where database migrations failed for upgrades

1.3.9 Ansible Tower Version 3.2.0

- Removed system tracking data (historical facts) feature starting with Ansible Tower 3.2. However, you can collect facts by using the fact caching feature. Refer to [Fact Caching](#) for more detail.
- Removed system tracking views in favor of directly viewing facts on hosts. Comparisons are best done with external data analytics systems.
- Removed Rackspace as a supported inventory source type and credential type.
- Removed the storing of `ansible_env` in job event data.
- Removed Job launching capability from `/api/v2/jobs`. Job template launching and job relaunching are the only support launch options.
- Deprecated the `group` field for InventorySource, which has been renamed to `deprecated_group` and will be removed from InventorySource completely in Tower 3.3. As a result, the related field on Group, `inventory_source` has been renamed `deprecated_inventory_source` and will also be removed in Ansible Tower 3.3.
- Deprecated requirement that inventory sources be associated with a group.
- Deprecated the `/api/v1` heirarchy with the introduction of `/api/v2`. `/api/v1` will be removed in a future Ansible Tower release to be determined.

- Deprecated the `/api/v2/authtoken` endpoint, which will be removed in Ansible Tower 3.3.
- Updated the job environment variables for AWS credentials. Refer to [Amazon Web Services](#) section of the *Ansible Tower User Guide* for new variable names.
- Added support for connecting to external log aggregators via direct TCP and UDP connections.
- Added the ability to test logging configurations through the Configure Tower UI.
- Updated the Ansible Tower Rest API to version 2 which include added endpoints: `instances`, `instance_groups`, `credential_types`, and `inventory_sources`.
- Added ability to create inventory sources and create Smart Inventories.
- Added the ability to access Tower resources via resource-specific human-readable identifiers.
- Added the ability to create and modify credential types.
- Added ability to create and modify instance groups and isolated nodes.
- Added the ability to enable and disable SSL certification verification through the Configure Tower UI. You no longer have to manually set an environment variable in your local `settings.py` file to achieve this.
- Updated upstream Azure libraries will require users who use Ansible Tower with Azure to use Ansible 2.4 or later.
- Fixed an outstanding issue regarding variable precedence so that the variable value is derived from the survey (survey variables take precedence over Job Template variables).
- Added Insights project remediation, which allows you to run the Insights maintenance plan associated with an inventory.
- Added a new API endpoint - `/api/v2/settings/logging/test/` - for testing external log aggregator connectivity.
- Updated passing `-e create_preload_data=False` to skip creating default organization/project/inventory/credential/job_template during Tower installation.
- Added support for sourcing inventory from a file inside of a source control project.
- Added support for custom cloud and network credential types, which give you the ability to modify environment variables, extra vars, and generate file-based credentials (such as file-based certificates or .ini files) at `ansible-playbook` runtime.
- Added support for assigning multiple cloud and network credential types on job templates. Job templates can now prompt for “extra credentials” at launch time in the same manner as promptable machine credentials.
- Updated custom inventory sources to now specify a `Credential`; you can store third-party credentials encrypted within Tower and use their values from within your custom inventory script (for example - by reading an environment variable or a file’s contents).
- Added support for configuring groups of instance nodes to run tower jobs. Instance groups can be assigned to an organization, inventory, or job template.
- Fixed an issue installing Tower on multiple nodes where cluster internal node references are used.
- Updated Tower to now use a modified version of [Fernet](<https://github.com/fernet/spec/blob/master/Spec.md>) for encrypting sensitive fields such as credentials. Our `Fernet256` class uses `AES-256-CBC` instead of `AES-128-CBC` for all encrypted fields.
- Added the ability to set custom environment variables globally for all playbook runs, inventory updates, project updates, and notification sending, via `AWX_TASK_ENV` configuration setting.
- Added `-diff` mode to Job Templates and Ad-Hoc Commands. The diff can be found in the standard out when diff mode is enabled.

- Added support for accessing some Tower resources via their name-related unique identifiers apart from primary keys.
- Added support for authentication to Tower via TACACS+.
- Updated names of tower-mange commands `register_instance` -> `provision_instance`, `deprovision_node` -> `deprovision_instance`, and `instance_group_remove` -> `remove_from_queue`, with backward compatibility support for 3.1 command names.
- Improved handling of workflow logic errors.
- Updated Azure bindings, and therefore, removed support for the old Azure classic modules.
- Fixed system auditor permissions.
- Updated Tower to explicitly prevent non-json bodies from being accepted in the API.
- Improved handling of default values in Tower Configuration.
- Improved handling of sensitive environment variables in job details.
- Added the ability to set the system auditor with `AUTH_LDAP_USER_FLAGS_BY_GROUP`.
- Fixed some minor UTF-8 handling issues.
- Fixed the system to no longer allow using password fields with the `order_by` query parameter in the API.
- Improved censoring of Ansible `no_log` in job output.
- Fixed handling project repository URLs with spaces and special characters.
- Improved explanation when canceling jobs that are dependencies of other jobs.
- Updated the `ansible-playbook` parameters to pass through the `setup.sh` script.
- Added translations for Dutch; updated translations for Japanese, French, and Spanish.
- Improved ability to update org admin/member roles on the user detail page.
- Added force shutdown of cluster nodes that are not at the same version as the rest of the cluster.
- Added configuration options in Tower Configuration UI.
- Updated Postgres to 9.6.
- Updated Tower by separating Vault credentials from machine credentials.
- Added more prompting options to job templates.
- Added the ability to prevent IDP user from assuming a local admin role.
- Improved the display of SCM revision hashes by abbreviating them, and added ability to easily copy revision to clipboard.
- Fixed a potential issue showing encrypted values in the activity stream instead of obfuscation characters.
- Added the ability to set an enabled/disabled flag on all supported cloud inventory sources.
- Added support for vmware `host_filters` and `groupby_patterns`.
- Fixed an issue where Tower wouldn't redirect the user to the right URL after clicking a link and logging in.
- Fixed tower to preserve `stderr` from custom inventory scripts.
- Updated Tower to now act as a fact cache source for jobs.
- Improved handling of related resources when inventories are deleted.
- Added the ability to show an indicator during background inventory delete.

- Updated supported cloud regions for some inventories.
- Improved SAML configurations.
- Improved LDAP settings validation.
- Added support for providing SSL cert for log aggregator service.
- Added the ability to set proxy IP whitelists for trusted vs. untrusted load balancers.
- Improved the efficiency in generating entries in the activity stream.
- Added support for upgrading Ansible during setup playbook run (`-e upgrade_ansible_with_tower=1`).
- Fixed downloading ad-hoc command stdout.
- Fixed job launch dependency handling.
- Fixed some xss vulnerabilities.
- Added runas privilege escalation support.
- Improved handling of instance capacity calculation.
- Fixed SSL certificate handling for LDAP.
- Updated the Job detail event modals to now be resizable.
- Improved yaml/json editor views.
- Improved job list performance.

1.4 Release Notes for 3.1.x

1.4.1 Ansible Tower Version 3.1.8

- Fixed a minor XSS vulnerability in the scheduling page
- Fixed potential information leakage via websocket
- Fixed a CSRF vulnerability in Tower (CVE-2018-10884)

1.4.2 Ansible Tower Version 3.1.7

- Fixed a RabbitMQ configuration issue that would affect cluster recovery on network interruptions

1.4.3 Ansible Tower Version 3.1.6

- Fixed organization admins to no longer be able to modify users by adding them to their organization (CVE-2018-1101)
- Fixed Tower to disable usage of Jinja templates in launch-time variables for security reasons (CVE-2018-1104). This release introduces the `ALLOW_JINJA_IN_EXTRA_VARS` configuration parameter for Tower. This parameter has three values: `template` to allow usage of Jinja saved directly on a job template definition (the default), `never` to disable all Jinja usage (recommended), and `always` to always allow Jinja (strongly discouraged, but an option for prior compatibility). Note that the `always` option is deprecated, and will be removed in a future Tower release.
- Updated memcached to now listen on a local Unix socket instead of a TCP socket

1.4.4 Ansible Tower Version 3.1.5

- Fixed an issue where certain API endpoints were unreasonably slow when jobs contained large amounts of output
- Enhanced Tower to properly show stdout for ad-hoc commands run on other cluster nodes
- Fixed an issue where a user, who could modify playbooks, could potentially compromise Tower via an injection of git hooks in SCM repositories (CVE-2017-12148)
- Fixed an issue where a specially crafted ad-hoc command could compromise Tower (CVE-2017-12148)
- Upgraded the available and bundled versions of RabbitMQ to 3.6.9, which addresses various RabbitMQ CVEs (CVE-2017-4965, CVE-2017-4966, CVE-2017-4967)

1.4.5 Ansible Tower Version 3.1.4

- Added the ability to customize the log format and the `job_events` payload will mirror the API structure with minor differences when upgrading from version 3.1.3 to 3.1.4.
- Added the ability to configure known proxies in order to allow certain load balancers and hosts when setting up proxy server support.

1.4.6 Ansible Tower Version 3.1.3

- Added Spanish translations for Ansible Tower
- Fixed a problem where survey password defaults of a certain length could prevent a job from launching
- Fixed truncation of long job run results causing an excessive number of blank lines
- Fixed a problem where running two different jobs in parallel from the same project could cause an error
- Fixed a problem where Splunk HTTP event collectors did not send `job_event` data
- Fixed a problem where email notifications could send survey password values in plain text in `extra_vars`
- Fixed handling of job output from certain modules
- Fixed a problem where setting “Required” for a Multiple Choice (Single Select) survey question could prevent changing the default
- Fixed a problem where provisioning callbacks configured to prompt for extra variables could return an error
- Fixed a problem where modifying team permissions when using a basic license would return an error
- Fixed a problem in LDAP configuration where adding values to the database caused users to lose permanent access to Tower
- Fixed an issue where a database restore could fail
- Fixed an issue where setup could fail if there was no firewall installed
- Fixed an issue where Insights projects could cause project syncs to fail
- Fixed an issue applying Tower configuration across cluster nodes
- Added the ability to disable SSL certificate verification for hosted Splunk logging
- Improved Ansible Tower to allow for passing `extra_vars` on `ad_hoc` commands
- Updated Ansible Tower so that it can process facts output from Ansible 2.3

1.4.7 Ansible Tower Version 3.1.2

- Added subpackaging for sever, UI, and setup packages
- Added support for Red Hat Insights project type
- Added support for explicitly specifying the host descriptor used for RabbitMQ config via *rabbitmq_host*
- Adjusted search on the Job Details screen to match the behavior across Tower
- Adjusted Tower logging to log asynchronously
- Fixed various and minor UI bugs
- Fixed a callback bug which was causing a *task_args* leak between job events
- Fixed an issue where jobs were not able to be sorted by descending ID
- Fixed an issue where, when working with Splunk, the log aggregator type shows as Logstash instead of Splunk
- Fixed an issue where, when a user has two groups in an inventory (one using a VMware script and one using a custom script), clicking *sync* on the custom script group caused the sync icon to link to the wrong inventory sync
- Fixed a problem where users were not able to put multi-line text in a Text Area-type field in a survey
- Fixed a problem where users who had admin access on Workflows, but were not Org level admins, could not add or remove job templates from Workflows
- Fixed a problem with job templates that include a multiple choice survey response, where, even when multiple selections are required, the job template ran with an empty array
- Fixed a problem where surveys were passing a variable as empty instead of null when they included text or a text area field that had a minimum length >0 and was not filled in
- Fixed a problem where Tower jobs hang and do not run when the Splunk server is unresponsive or unavailable
- Fixed a problem where users with admin level permissions on projects could not modify project details
- Fixed a problem in multiple choice survey inputs where, when selecting a string that had similar characters or words at the beginning or end of the string, a similar but smaller version of that string was rendered as the user's selection (even though the correct value was still passed to extra-vars on launch)
- Fixed an issue around Git project updates failing when the username was specified
- Fixed a problem where job templates from mercurial project updates failed to run
- Fixed a problem with provisioning callbacks where they failed with '400' responses when extra_vars were passed to the API through curl in the callback
- Fixed a problem where running the installer again anytime after successfully creating the rabbitmq user caused the installation program to fail
- Fixed an issue where Windows package scan jobs fail when targeting a Windows 2012R2 host
- Fixed an issue where users with admin access to Workflow Templates could not modify the workflow
- Fixed an issue where a warning was incorrectly displayed for the output of a canceled job
- Fixed an issue where Mercurial project revisions were not read correctly for Projects
- Fixed an issue where Tower upgrades would fail when applying *rabbitmq_user* in a cluster
- Fixed an issue where certain characters in a Project SCM URL would cause updates to fail
- Improved custom inventory scripts support by ensuring that newlines added to the script are not trimmed
- Relaxed the SELinux policy dependency to allow Tower to be installed on older Enterprise Linux 7 releases

- Updated Ansible Tower so that the host config key is marked as required when provisioning callbacks are selected
- Updated Ansible Tower so that PostgreSQL Server is no longer installed on Tower nodes not hosting the database
- Updated Ansible Tower so that Tower shows *extra_vars* for ad-hoc commands in the UI

1.4.8 Ansible Tower Version 3.1.1

- Added a preflight check for password and pre-3.1.0 active/passive (HA) inventory setups prior to installation
- Fixed a problem where, while running a clustered Tower deployment configuration, there were some instances where realtime job event data did not flow through the channel layer
- Fixed a problem with searching where an invalid search term was entered and the error dialog continued to persist
- Fixed a problem with Slack notifications where they were not emitted if only ‘Failure’ was selected
- Fixed a problem where logging out via Tower logout button caused subsequent login attempts to fail
- Fixed an issue where, when logging was enabled, a missing logging UUID setting would cause a startup error, making the system unresponsive

1.4.9 Ansible Tower Version 3.1.0

- Added support for configuring most aspects of Ansible Tower directly from the Tower user interface (and Tower API), rather than editing Tower configuration files
- Added support for “Scale-Out” Clusters, which replaces the HA/Redundancy method from prior Tower releases
- Added support for Workflows, a chain of job templates executed in order
- Added support for sending event and log messages to various logging services (Elastic, Splunk, Sumologic, Loggly, generic REST endpoint)
- Added support for a new Tower Search feature which supports GitHub-style “key:value” searching
- Added support for Ubuntu 16.04
- Added support for a New Project Sync Architecture, where projects are now checked out at job runtime
- Added support for setting timeouts on job runs
- Added support for internationalization and localization (French and Japanese)
- Added support for multi-playbook Workflows
- Added `/api/v1/settings` for Tower managed settings. This corresponds to the in-Tower configuration UI
- Added support for windows scan jobs
- Added support so that the SCM Revision used is now stored on Job
- Added support for API endpoints to now show `__search` filter fields for broader searching of objects
- Added support so that system jobs are now shown in `/api/v1/unified_jobs`
- Added support for the new Ansible `vmware_inventory` script
- Added support for Job stdout downloads, which may generate and cache on the fly
- Added support for `/api/v1/inventory_updates` and `/api/v1/project_updates` to view those specific job types

- Added support for *user_capabilities* API elements in various places to allow API consumers to know if their user can perform the referenced actions on the object
- Added support for *set_stats* for Workflow jobs to persist data between Workflow job runs, support added in Ansible core also
- Added support for Tower callbacks so that they can now resolve *ansible_host* as well as *ansible_ssh_host*
- Added support for Tower callbacks so that they now filter out *ansible_* variables on POST
- Added support for notifications so that they are emitted on jobs marked as failed by the dead job detector
- Added eu-west-2 and ca-central-1 to the list of supported EC2 regions
- Added support for *format=ansi_download* when downloading stdout
- Deprecated support for Rackspace inventories
- Fixed an issue where manual projects could be launched/updated
- Fixed various unicode issues
- Fixed various issues dealing with self signed certificatesvalue.
- Fixed Jobs so that they now show *\$encrypted* for these variables, where they previously did not
- Improved performance for viewing job and job template lists
- Improved Tower virtualenv so that it is purged on upgrade
- Improved setup playbook so that it is more tolerant of various iptables/firewalld configurations
- Improved the optimization of PostgreSQL installation to improve overall performance
- Improved database migrations through consolidation to make upgrades/installs faster
- Improved hardening for web server configuration (SSL, HSTS)
- Removed zeromq as a communications channel between dependent services in favor of rabbitmq
- Removed */api/v1/jobs/n/job_plays* and */api/v1/jobs/n/job_tasks*
- Removed proot in favor of bubblewrap for process isolation
- Removed the ability to make POST requests on the */api/v1/jobs/* endpoint
- Removed *has_schedules* from various endpoints, as it was never populated
- Removed support for Red Hat Enterprise Linux 6/CentOS 6 and Ubuntu 12.04
- Updated surveys so that a blank value for a survey question default value now passes an empty string as a value
- Updated surveys so that previously existing surveys with blank default question values now pass empty strings as an extra variable
- Updated Websockets, moving them from socket.io to django channels and are now served under port 443/80 along with the regular web service. Port 8080 is no longer needed.
- Updated Job results so that they are now driven by job events and thus provides clickable context
- Updated Tower so that it now uses the system time zone by default
- Updated Tower requirements for Ansible–Tower now requires Ansible 2.1 or later
- Updated Ansible inventory plugins to the latest versions
- Updated Web server to NGINX from Apache
- Updated survey passwords so that they are now encrypted when stored in the database

- Updated `request_tower_configuration.sh`

1.5 Release Notes for 3.0.x

1.5.1 Ansible Tower Version 3.0.3

- Added support for new AWS regions, including an update to the boto version included with Tower
- Fixed various minor UI and API related bugs
- Fixed a regression with authentication restrictions
- Fixed an issue where restoring the database failed when using the RHEL6 bundled installation method
- Fixed an issue where, when viewing a host, “extra vars” were not initially formatted properly
- Fixed an issue where users were able to relaunch jobs they did not have permission to initially launch
- Fixed an issue where, after editing a Job Template, retrieving Job Templates failed when filtered
- Fixed an issue where Satellite 6 inventory marked all hosts as disabled
- Fixed an issue where Inventory variables were displayed incorrectly when editing hosts
- Fixed a rendering issue with the Host Event details window
- Fixed an issue where, when launching an inventory update, users were navigated away from the inventory manage view
- Fixed an issue where organization auditors could see the user permissions of other users in their organization
- Fixed an issue where canceling a Windows job in Tower left an orphaned process running on the control machine
- Fixed an issue where empty Host Variable Data produces a 500 error in the API browser after upgrading from 2.4.5
- Fixed an issue when using an Azure Service Principal in conjunction with Microsoft Azure inventory
- Fixed an issue where Inventory syncs fail against a resource group if it contains a non-standard virtual machine size when using Azure
- Fixed an issue where navigating to the admin or users from the organizations view in Tower caused 404 errors
- Fixed an issue where, when updating a Rackspace inventory, TypeError messages appeared
- Improved the run time performance for playbooks in Tower
- Improved support around how YAML is handled with Tower’s variable parser
- Improved the population of manual projects in Tower
- Improved Event Summary status badge counts
- Improved PostgreSQL configuration with regard to authentication (CVE-2016-7070)
- Updated PostgreSQL repository location for installation methods

1.5.2 Ansible Tower Version 3.0.2

- Added support for IAM Roles when configuring an EC2 Inventory Sync
- Added support for backing up and restoring Databases created when installing 3.0.x
- Added the display of a “working” indicator when toggling Tower components on/off
- Added the ability to toggle the view of job labels (view less/view more)
- Added the ability to add skip tags to job templates (which may also be prompted for at launch time)
- Added documentation around resetting the Tower URL provided in Notification links
- Fixed an issue where users could not remove inventory or credentials from job template
- Fixed an issue where admins were not properly allowed to copy or edit to Job Templates via the API
- Fixed an issue where Home/Host column views were not sortable
- Fixed the display of schedules to only show those with future activity
- Fixed an error where clicking to a different page number while editing a resource and making a new selections indicated an item other than the one currently selected/being edited
- Fixed an issue where relaunching a job ignored search filters
- Fixed an issue where searching for a user on an inventory permission page queried a project access list URL instead of the inventory access list URL
- Fixed an issue where pressing the Enter key (instead of clicking ‘Ok’ with your mouse) closes a pop up error message and, unexpectedly, navigates the user back to the Tower home page
- Fixed an issue where system job templates were not being included when viewing unified job template results
- Fixed an issue related to relaunching ad hoc commands
- Fixed an issue preventing projects from being deleted during an SCM update
- Fixed an issue where, when viewing the “Event Summary” field, filtering by task status summary dots returned incorrect tasks information
- Fixed an issue where selecting a host on one page, then going to the next page and selecting another host, did not save the prior selection as expected
- Fixed an issue where processing extra_vars in a survey caused errors
- Fixed an issue regarding how passwords are stored with surveys
- Fixed an issue where, when running a playbook with an ignored task, the ignored task was incorrectly marked as failing
- Fixed an issue so that Webhook notifications properly display the host summary information
- Fixed an issue where provisioning callbacks were running multiple times in a row
- Fixed various minor issues related to RBAC permissions and credentials
- Fixed various minor API bugs
- Fixed various minor UI and tooltips bugs
- Fixed an issue related to SAML logins hanging after multiple authorization attempts
- Fixed an issue where the “start date” header and schedule preview do not match what is set by the browser locale
- Fixed an issue where users could not properly edit their profile
- Fixed an issue related to backup/restoring with the setup.sh script

- Improved Tower installer compatibility with RHUI repos on RHEL non-AWS instances
- Improved upon what the auditor role can view (organization auditors can view inventory script contents in their own organizations, view notification templates in the activity stream, team credentials views)
- Improved the consistency of how scheduling is displayed within the Tower UI
- Improved how credentials are handled in that they should only be shareable when the organization field is not “null”
- Improved how teams are displayed for different organizations when viewing permissions
- Improved support for CloudForms and Red Hat Satellite 6 with Tower 3.0.x
- Reorganized activity stream views/access for organization admins and auditors
- Removed the requirement of needing a password for the network credential when using an SSH key
- Removed the requirement of needing AUTH with Email notifications using SMTP

1.5.3 Ansible Tower Version 3.0.1

- Added a stock schedule job for the ‘Cleanup Fact Details’ management job
- Fixed an issue with inventory syncs using Red Hat Satellite 6 credentials
- Fixed an issue which incorrectly allowed users assigned to a system auditor role to be able to escalate privileges to teams
- Fixed an issue with Webhook notifications where the content-type was being set incorrectly
- Fixed an issue where canceling a new job failed to change state from “new” to “canceled”
- Fixed an upgrade and credential migration issue which involved null inventory fields in job templates
- Fixed an upgrade and migration issue where hosts which had previously been deleted were not skipped during the upgrade process
- Fixed an upgrade and migration issue where job templates linked to deleted inventories caused migrations to fail
- Fixed an upgrade and migration issue where job templates without inventories caused migrations to fail
- Fixed an error related to the logging of RBAC migration data which caused installations to fail
- Fixed an issue related to license checks
- Fixed other various issues related to upgrading and migration
- Fixed the need for elevated permissions to make changes to job templates under some scenarios
- Fixed an issue where Organization-level admins could not edit scan jobs that were created prior to upgrading to Tower 3.0
- Fixed an issue regarding Software Collections (SCL) installation on EL6
- Fixed a problem with subsequent logins after upgrading to Tower 3.0 when using Google OAuth or SAML authentication
- Discovered an issue with MS Azure inventory imports using new-style credentials being unsupported on distributions that ship python-2.7 (e.g. not EL6)
- Updated the UI to display new jobs in the Jobs overview screen and added a cancellation method for these new jobs

1.5.4 Ansible Tower Version 3.0

- Added a notifications system for Tower which supports services like Slack, HipChat, IRC, etc.
- Added support for the new Azure inventory system and the latest Ansible Azure modules (legacy Azure inventory and credentials are still supported)
 - Azure inventory imports using new-style credentials are only supported on distributions that ship python-2.7 (e.g. not EL6)
- Added support for keystone v3 which supports the latest Openstack versions
- Added counts and more detail to Organization endpoints (API)
- Added prompting for Job Templates
- Added labels for Job Templates
- Added support for user customization as Ansible tasks now run in their own environment
- Added support for new Ansible Network Credentials
- Added inventory support for Red Hat Cloudforms and Red Hat Satellite 6
- Added SUSE, OpenSuse, and Debian support for scan jobs
- Added a link to the schedule in the job detail view if the job was started as a result of a schedule
- Added survey spec management without requiring that surveys be enabled on job templates
- Added additional strict extra_vars validation. 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
- Added a deprecation notice for Ubuntu 12 and RHEL 6
- Changed how Projects are linked so that they now tie singularly to an Organization
- Changed how system tracking and scan data are stored—now in postgres. MongoDB dependency removed.
- Discovered an issue with ECDSA credentials—if your Tower server has a version of OpenSSH that predates 5.7, jobs will fail when launched jobs with ECDSA credentials
- Fixed issues with scan jobs on RHEL5
- Fixed an issue with the websocket service when Tower is run on CentOS or RHEL 7.2
- Fixed issues with Ansible’s no_log causing errors or not hiding data when running jobs
- Fixed the way setting a license is done so that it propagates to standby Tower nodes in an HA configuration
- Fixed GCE credential handling and inventory filtering
- Improved (through a complete rewrite to expand and simplify) the Role-Based Access Control system in Tower
- Improved job templates so that multiple invocations of the same job template will only block if the job templates used the same inventory
- Improved the setup playbook so that it now hides potentially sensitive information from stdout and the setup log
- Improved the Setup process now supports installing and configuring postgres on a remote system
- Removed MongoDB and changed view queries to use a Postgres implementation
- Removed soft-deletes: Tower now permanently deletes removed objects and the utilities to manage the cleanup of those soft-deleted objects have been removed

- Removed Munin monitoring
- Updated the look and feel of the entire Tower UI for a more approachable and intuitive user experience
- Updated and simplified the Tower setup process so that new Tower installs are now preloaded with Organization, Inventory, Project, and Job Template demo data
- Updated the setup process to support installing and configuring Postgres on a remote system
- Updated dependencies
- Updated Red Hat Enterprise Linux 6/CentOS 6 to use python 2.7 (for Tower only)
- Updated the minimum open file descriptor check and configuration by raising it from 1024 to 4096

1.6 Release Notes for 2.4.x

1.6.1 Ansible Tower Version 2.4.5

- Corrected an issue where inventory syncs using Rackspace credentials failed
- Corrected an issue where the Host Events display provided different results depending on the version of Ansible used
- Corrected an issue which caused an error when calling the Ansible `yum` module on `ansible-1.9.4` (or newer)
- Improved display for Ansible loops on the job detail page by recognizing new Ansible callback events (`v2_runner_item_on_*`)
- Improved the efficiency of the stdout dump database migration for better memory handling
- Updated the Boto release included with Tower to version 2.39.0

1.6.2 Ansible Tower Version 2.4.4

- Corrected an issue related to Ansible 2.0.0.x job callback events
- Corrected an issue where YAML `extra_vars` were ignored when launching a job template
- Corrected an issue where running scan jobs against Red Hat Enterprise Linux 5 inventory failed
- Corrected an issue where the Services tab was not populating in scan jobs on SLES 11 or RHEL 5
- Corrected an issue with log output filtering
- Corrected an issue where the Rackspace module had caching on by default
- Corrected an issue where Tower was not working properly on Centos 7.2 with Python 2.7.5
- Corrected an issue where OpenStack modules were not running correctly on systems with Python 2.7 (bumping shade and pyrax versions to allow Ansible 2.0 OpenStack modules to run correctly)
- Corrected an issue where the setup/upgrade playbook failed if being run from Ansible 2.X

Note: Ansible 2.0 OpenStack modules will not work on Red Hat Enterprise Linux 6 or CentOS 6.

1.6.3 Ansible Tower Version 2.4.3

- Added sample configurations for LDAP connection options and disable referrals by default, which corrects a problems with queries hanging with AD
- Corrected an issue where the UI does not enable provisioning callbacks properly
- Improved performance of user and group queries though better caching

1.6.4 Ansible Tower Version 2.4.2

- Corrected a problem with EC2 inventories which were not working correctly when instance filters were in use
- Corrected an issue when accessing Tower using IE11 web browsers
- Corrected an issue where clicking on a job in the activities stream did not show the correct job detail page
- Corrected an issue where custom login information was not properly displayed at login
- Corrected an issue with scan jobs against Amazon Linux machines throwing error messages instead of warnings
- Corrected an API-related problem dealing with sparkline data which corrects the ordering of recent jobs as associated with job templates
- Corrected an issue in the UI where cloud credentials associated with an inventory source were not being properly displayed
- Corrected an issue where org admins did not have the proper permissions to delete project updates
- Corrected several small UI issues

1.6.5 Ansible Tower Version 2.4.1

- Resolved a failure that, when not connected to the Internet (such as being behind a restrictive firewall), prevented Tower from functioning

1.6.6 Ansible Tower Version 2.4.0

- Added custom rebranding support
- Added the ability to enable and disable basic authentication
- Added support for authentication via SAML 2.0 servers, Google Apps, GitHub, and RADIUS
- Added support for session limits
- Added support for EC2 STS tokens
- Added default schedules for system jobs on new installs
- Added support to allow multiple scheduled system jobs
- Added an example “request_tower_configuration.ps1” PowerShell for use with Tower’s provisioning callbacks
- Added analytics and data collection for improving the UI experience of Ansible Tower
- Changed the behavior of config.js handling and introduced support for the local_settings.json file for specific variable changes
- Changed the way Job Templates work so that they launch using an [extra variables](#) hierarchy
- Changed session timeout to be set in session.py and no longer in the UI local_config.js file

- Changed the local_config.js file to local_settings.json and made it more flexible to override configuration settings
- Corrected some Tower features when using Ansible 2.0
- Corrected an issue where ‘Overwrite’ in an inventory update would imply ‘Overwrite Variables’
- Corrected an issue where Tower-cli ignored default answers when trying to launch a job with a survey
- Corrected an issue that prevented LDAP logging from working correctly
- Corrected an issue where Null errors were returned after deleting an Organization associated with a Custom Inventory Script with an Inventory
- Incorporated a feature which adds an Auth-Token-Timeout to every responses that include a valid user-supplied token
- Noted a known issue where using the strategies feature of Ansible 2.0 in Ansible Tower causes jobs to not display properly (support for the strategies feature will be added in a future release of Tower)
- Removed the ability to delete the default set Organization for Basic-level license users

1.7 Release Notes for 2.3.x

1.7.1 Ansible Tower 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.

1.7.2 Ansible Tower Version 2.3.0

- 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

1.8 Release Notes for 2.2.x

1.8.1 Ansible Tower 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

1.8.2 Ansible Tower 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

1.8.3 Ansible Tower Version 2.2.0

- 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.

1.9 Release Notes for 2.1.x

1.9.1 Ansible Tower Version 2.1.4

- Corrected Tower's Live Events feature, again. Really.

1.9.2 Ansible Tower 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

1.9.3 Ansible Tower 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

1.9.4 Ansible Tower 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

1.9.5 Ansible Tower Version 2.1.0

- 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

1.10 Release Notes for 2.0.x

1.10.1 Ansible Tower Version 2.0.5

- Ensured websocket connection uses user's RBAC credentials
- Corrected a potential CSRF issue when using the REST API graphical browser

1.10.2 Ansible Tower Version 2.0.4

- Corrected a privilege escalation related to user account levels

1.11 Ansible Tower 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

1.11.1 Ansible Tower 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

1.11.2 Ansible Tower 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

1.12 Release Notes for 1.4.x

1.12.1 Ansible Tower Version 1.4.12

- Corrected an issue handling Unicode output from ansible-playbook
- Corrected an issue displaying job details for some jobs

1.12.2 Ansible Tower 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

1.12.3 Ansible Tower Version 1.4.10

- Correctly handle schedule creation when browser timezone cannot be detected.
- Corrected pagination on job_events page.

1.12.4 Ansible Tower Version 1.4.9

- Corrected a provisioning callback issue on Enterprise Linux.
- Added a sample provisioning callback script.
- Various backend and UI improvements.

1.12.5 Ansible Tower 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

- *Isolated nodes unsupported in an OpenShift deployment*
- *Tower login via HTTP requires workaround*
- *Job slicing and limit interactions*
- *Default LDAP directory must be configured to use LDAP Authentication*
- *Lost Isolated Jobs*
- *Potential security issue using X_FORWARDED_FOR in REMOTE_HOST_HEADERS*
- *Server error when accessing SAML metadata via hostname*
- *Live events status indicators*
- *VMWare Self-Signed Certs*
- *Tower Database on Disk Becomes Corrupted*
- *Safari unable to establish connection to web socket*
- *Ansible Azure Dependencies*
- *sudo/su not working as expected for local playbooks or playbooks with local_actions*
- *Problems when using SSH customization*
- *Ubuntu unsupported for bundled installations*
- *Reactivating OAuth authentication accounts which have been deleted*
- *Using vaulted variables in inventory sourced from a project*
- *Using Tower credential with Tower 3.3+ Tokens*
- *Saved scheduled and workflow configurations and surveys*

2.1 Isolated nodes unsupported in an OpenShift deployment

Isolated nodes are not currently supported when deploying Ansible Tower in OpenShift.

2.2 Tower login via HTTP requires workaround

To access Tower nodes behind your load balancer (in traditional cluster Tower installs) via HTTP, refer to the procedure described in the [Troubleshooting](#) section of the *Ansible Tower Administration Guide*.

2.3 Job slicing and limit interactions

When passing a limit to a Sliced Job, if the limit causes slices to have no hosts assigned, those slices will fail, causing the overall job to fail.

2.4 Default LDAP directory must be configured to use LDAP Authentication

Tower 3.3 introduced the ability to configure multiple LDAP directories for authentication—up to six. On the settings page for LDAP, there is a “Default” LDAP configuration followed by five-numbered configuration slots. If the “Default” is not populated, Tower will not try to authenticate using the other directory configurations.

2.5 Lost Isolated Jobs

After starting an isolated job, if an event impacts the availability of the **controller** which manages the job, Tower will be unable to determine the status of the job, and manual intervention will be necessary. Contact Ansible via the Red Hat Customer portal at <https://access.redhat.com/> for instructions specific to your scenario.

2.6 Potential security issue using X_FORWARDED_FOR in REMOTE_HOST_HEADERS

If placing Tower nodes behind some sort of proxy, this may pose a security issue. This approach assumes traffic is always flowing exclusively through your load balancer, and that traffic that circumvents the load balancer is suspect to X-Forwarded-For header spoofing. See [Proxy Support](#).

2.7 Server error when accessing SAML metadata via hostname

When Tower is accessed via hostname only (e.g. <https://my-little-tower>), trying to read the SAML metadata from `/sso/metadata/saml/` generates a `sp_acs_url_invalid` server error.

A configuration in which uses SAML when accessing Tower via hostname only instead of an FQDN, is not supported. Doing so will generate an error that is captured in the `tower.log` file and in the browser with full traceback information.

2.8 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. They are not seen at all when the system is in a healthy 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.

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.

2.9 VMWare Self-Signed Certs

If you have a VMware instance that uses a self-signed certificate, then you will need to add the following to the *Source Vars* configuration of the Cloud Group:

```
"source_vars": "---\nvalidate_certs: False",
```

You can set this in inventory source for VMware vCenter as follows:

INVENTORIES / Inventory Example with VMWare Source / SOURCES / CREATE INVENTORY SOURCE

CREATE SOURCE

DETAILS | SCHEDULES

* NAME: DESCRIPTION: * SOURCE:

SOURCE DETAILS

* CREDENTIAL: INSTANCE FILTERS: ONLY GROUP BY:

VERBOSITY: UPDATE OPTIONS: OVERWRITE OVERWRITE VARIABLES UPDATE ON LAUNCH

SOURCE VARIABLES:

```

1 ---
2 \nvalidate_certs: False",
    
```


CANCEL SAVE

2.10 Tower Database on Disk Becomes Corrupted

If Tower is not cleanly shutdown, it leaves a `/var/lib/awx/beat.db` file on disk. If that happens, the dispatcher won't start, and you must manually delete the `/var/lib/awx/beat.db` file and restart Tower before the dispatcher will start properly.

2.11 Safari unable to establish connection to web socket

The following connection error displays in Tower:



This error is the result of Safari silently refusing to establish a connection to a web socket that is using a self-signed certificate. To resolve this issue, you must set Safari to always trust the website upon first visiting it:

1. Close the current browser and revisit the site. An error message appears stating Safari can't verify the identity of the website.
2. Click **Show Certificate**.
3. Check the **Always trust ... when connecting to ...** checkbox to allow Safari to accept the connection.

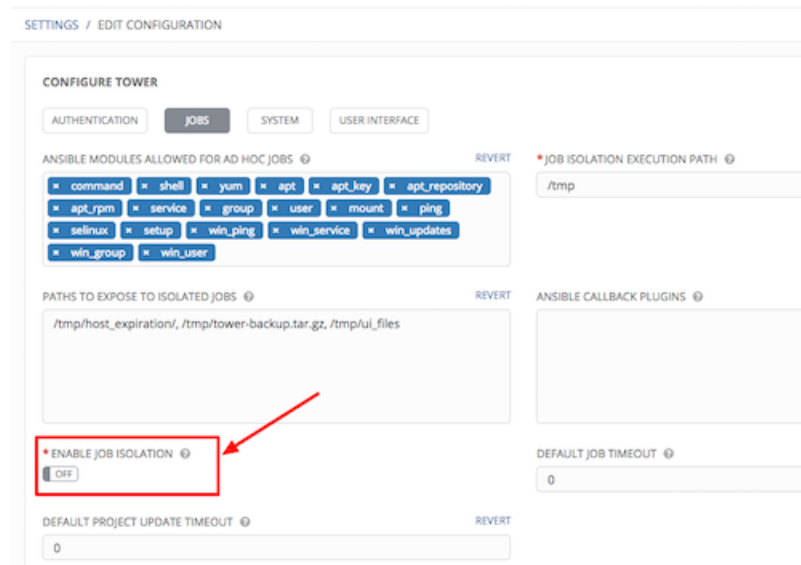
If you click **Continue** without checking the checkbox, this error will persist.

2.12 Ansible Azure Dependencies

Ansible Tower 3.4 contains bindings for Microsoft Azure compatible with Ansible 2.7. However, these bindings will not work with earlier versions of Ansible. If you are using earlier versions of Ansible in a custom virtual environment, you may need to install different versions of Azure dependencies to use Microsoft Azure modules.

2.13 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 job isolation being enabled. To use `sudo/su` commands with local playbooks or those with `local_actions`, job isolation must be disabled. You can disable job isolation through the **Jobs** tab of the Configure Tower screen by setting the **Enable Job Isolation** toggle to **OFF**:

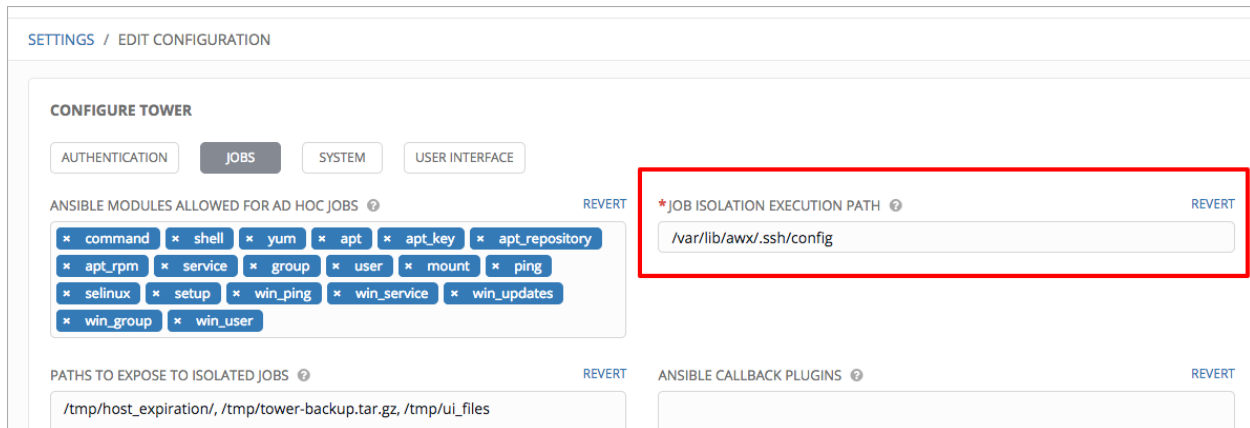


Click **Save** to save your changes and restart services.

2.14 Problems when using SSH customization

The Job Isolation 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 bubblewrap.

For example, to add a custom SSH config in `/var/lib/awx/.ssh/config` and make it available for Tower jobs, you can specify the path in the **Job Execution Isolation Path** field accessed from the **Jobs** tab of the Configure Tower screen:



2.15 Ubuntu unsupported for bundled installations

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

2.16 Reactivating OAuth authentication accounts which have been deleted

Once a user who logs in using social authentication has been deleted, the user will not be able to login again or be recreated until the system administrator runs a `cleanup_deleted` action with `days=0` to allow users to login again. Once `cleanup_deleted` has been run, Tower must be restarted using the `ansible-tower-service restart` command. Accounts which have been deleted prior to having the `cleanup_deleted` action run will receive a “Your account is inactive” message upon trying to login.

2.17 Using vaulted variables in inventory sourced from a project

When using inventory from a source control project, individual vaulted variable values are supported. Vaulted files are not currently supported.

2.18 Using Tower credential with Tower 3.3+ Tokens

The Ansible Tower credential type in Ansible Tower 3.4 does not support use with an OAuth2 Token. Only username + password is currently supported.

2.19 Saved scheduled and workflow configurations and surveys

If a configuration of a job template is scheduled or added to a workflow with answers from a prompted survey, changing the Job Template survey to supply different variable names may cause the saved configuration to not function. The workaround is to delete the saved schedule configuration/workflow node, and recreate it with answers from the updated survey.

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)

- genindex

COPYRIGHT © 2019 RED HAT, INC.

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

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

Third Party Rights

Ubuntu and Canonical are registered trademarks of Canonical Ltd.

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

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

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

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

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

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

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

Safari® is a registered trademark of Apple, Inc.

Firefox® is a registered trademark of the Mozilla Foundation.

All other trademarks are the property of their respective owners.

A

Ansible Azure dependencies
 known issues, 30
 authentication (*reactive user*)
 known issues, 30

B

bundled installer
 known issues, 30

D

database corruption
 known issues, 30

G

green dot
 live event statuses, 30

H

host comparisons
 known issues, 30

I

issues, known, 30

J

job isolation
 known issues, 30

K

known issues, 30
 Ansible Azure dependencies, 30
 authentication (*reactive user*), 30
 bundled installer, 30
 database corruption, 30
 host comparisons, 30
 job isolation, 30
 LDAP authentication, 30
 live event statuses, 30
 login problems with social
 authentication, 30
 login via http, 30

lost isolated jobs, 30
 OAuth account recreation, 30
 proxy support, 30
 self-signed certs (*VMWare*), 30
 session limit, 30
 sosreport, 30
 ssh customization, 30
 sudo/su, 30
 traceback error, 30
 Ubuntu, 30
 upgrades, 30
 user cannot log in using
 authentication, 30
 VMWare self-signed certs, 30
 web sockets in safari, 30
 YAML traceback error, 30

L

LDAP authentication
 known issues, 30
 live event statuses
 green dot, 30
 known issues, 30
 red dot, 30
 locales supported, 37
 login problems with social
 authentication
 known issues, 30
 login via http
 known issues, 30
 lost isolated jobs
 known issues, 30

O

OAuth account recreation
 known issues, 30

P

proxy support
 known issues, 30

R

red dot

- live event statuses, 30
- release notes
 - v1.4.10, 29
 - v1.4.11, 29
 - v1.4.12, 29
 - v1.4.8, 29
 - v1.4.9, 29
 - v2.0, 28
 - v2.0.1, 28
 - v2.0.2, 28
 - v2.0.4, 28
 - v2.0.5, 28
 - v2.1.0, 27
 - v2.1.1, 27
 - v2.1.2, 27
 - v2.1.3, 27
 - v2.1.4, 27
 - v2.2.0, 26
 - v2.2.1, 26
 - v2.2.2, 26
 - v2.3.0, 25
 - v2.3.1, 25
 - v2.4.0, 24
 - v2.4.1, 24
 - v2.4.2, 24
 - v2.4.3, 24
 - v2.4.4, 23
 - v2.4.5, 23
 - v3.0, 19
 - v3.0.1, 19
 - v3.0.2, 19
 - v3.0.3, 19
 - v3.1.0, 14
 - v3.1.1, 14
 - v3.1.2, 14
 - v3.1.3, 14
 - v3.1.4, 14
 - v3.1.5, 14
 - v3.1.6, 14
 - v3.1.7, 14
 - v3.1.8, 14
 - v3.2.0, 7
 - v3.2.1, 7
 - v3.2.2, 7
 - v3.2.3, 7
 - v3.2.4, 7
 - v3.2.5, 7
 - v3.2.6, 7
 - v3.2.7, 7
 - v3.2.8, 7
 - v3.3.0, 4
 - v3.3.1, 4
 - v3.3.2, 4
 - v3.3.3, 4

- v3.3.4, 4
- v3.4.0, 2
- v3.4.1, 2
- v3.4.2, 2

S

- self-signed certs (*VMWare*)
 - known issues, 30
- session limit
 - known issues, 30
- sosreport
 - known issues, 30
- ssh customization
 - known issues, 30
- sudo/su
 - known issues, 30

T

- traceback error
 - known issues, 30

U

- Ubuntu
 - known issues, 30
- upgrades
 - known issues, 30
- user cannot log in using
 - authentication
 - known issues, 30

V

- v1.4.10
 - release notes, 29
- v1.4.11
 - release notes, 29
- v1.4.12
 - release notes, 29
- v1.4.8
 - release notes, 29
- v1.4.9
 - release notes, 29
- v2.0
 - release notes, 28
- v2.0.1
 - release notes, 28
- v2.0.2
 - release notes, 28
- v2.0.4
 - release notes, 28
- v2.0.5
 - release notes, 28
- v2.1.0
 - release notes, 27
- v2.1.1

- release notes, 27
- v2.1.1
 - release notes, 27
- v2.1.2
 - release notes, 27
- v2.1.3
 - release notes, 27
- v2.1.4
 - release notes, 27
- v2.2.0
 - release notes, 26
- v2.2.1
 - release notes, 26
- v2.2.2
 - release notes, 26
- v2.3.0
 - release notes, 25
- v2.3.1
 - release notes, 25
- v2.4.0
 - release notes, 24
- v2.4.1
 - release notes, 24
- v2.4.2
 - release notes, 24
- v2.4.3
 - release notes, 24
- v2.4.4
 - release notes, 23
- v2.4.5
 - release notes, 23
- v3.0
 - release notes, 19
- v3.0.1
 - release notes, 19
- v3.0.2
 - release notes, 19
- v3.0.3
 - release notes, 19
- v3.1.0
 - release notes, 14
- v3.1.1
 - release notes, 14
- v3.1.2
 - release notes, 14
- v3.1.3
 - release notes, 14
- v3.1.4
 - release notes, 14
- v3.1.5
 - release notes, 14
- v3.1.6
 - release notes, 14
- v3.1.7
 - release notes, 14
- v3.1.8
 - release notes, 14
- v3.2.0
 - release notes, 7
- v3.2.1
 - release notes, 7
- v3.2.2
 - release notes, 7
- v3.2.3
 - release notes, 7
- v3.2.4
 - release notes, 7
- v3.2.5
 - release notes, 7
- v3.2.6
 - release notes, 7
- v3.2.7
 - release notes, 7
- v3.2.8
 - release notes, 7
- v3.3.0
 - release notes, 4
- v3.3.1
 - release notes, 4
- v3.3.2
 - release notes, 4
- v3.3.3
 - release notes, 4
- v3.3.4
 - release notes, 4
- v3.4.0
 - release notes, 2
- v3.4.1
 - release notes, 2
- v3.4.2
 - release notes, 2
- VMWare self-signed certs
 - known issues, 30

W

- web sockets in safari
 - known issues, 30

Y

- YAML traceback error
 - known issues, 30