

vRealize Automation 8.4 Release Notes

vRealize Automation 8.4 | 06 August 2021

- vRA Easy Installer (ISO) build 17879649
- vRA product (appliance) build 17874359

Check regularly for additions and updates to these release notes.

Updates made to this document

Date	Description of update	Type
23 June 2021	NSX-T security tag support for day 0 and day 2 operations	Feature
23 June 2021	Deployments fail for blueprints that contain compute tags longer than 256 characters or a key larger than 128 characters	Known Issue
06 August 2021	vRealize Automation upgrade fails with error code LCMVRAVACONFIG90030 due to password expiration	Known Issue
02 March 2022	Added link to KB workaround used to resolve upgrade failure related to the log4j vulnerabilities.	

What's in the Release Notes

- [About vRealize Automation 8.4](#)
- [What's New](#)
- [Before You Begin](#)
- [Known Issues](#)

About vRealize Automation 8.4

vRealize Automation 8.4 adds to the vRealize Automation 8.3 capabilities to bring it closer in capability to the vRA 7.x release, reintroducing key capabilities like XaaS and adding capabilities such as Powershell support in ABX and python, node.js and Powershell in vRO.

Important

Upgrade failure after performing steps in KB 87120

Performing the instructions used to address the CVE-2021-44228 and CVE-2021-45046 log4j vulnerabilities described in KB 87120 can cause upgrade failures for vRealize Automation and vRealize Orchestrator 8.6.2 or earlier. For a workaround, see [KB 87794](#).

What's New

The many benefits of vRealize Automation 8.4 include:

Federal Information Processing Standard (FIPS) 140-2 compliance - SaltStack Config

SaltStack Config now ships with cryptographic modules that have successfully passed NIST FIPS 140-2 Cryptographic Module Validation Program (CMVP) testing. When these modules are configured to run in 'FIPS-mode', they cover all cryptographic operations in the product that perform a security function and/or process sensitive data.

NOTE: You can choose to enable FIPS-mode **only** at installation time. FIPS mode is currently available for greenfield SaltStack Config environments only. When running with vRealize Automation, mixed FIPS-mode is not supported.

Important change in Access Token API behavior

The behavior of the `/csp/gateway/am/api/login?access_token` API has changed.

This API is used in the first step of the two-step process to obtain an access token for API integrations. The correct way to utilize this API is documented on <https://code.vmware.com/docs/10222/vrealize-automation-api-programming-guide--html-/GUID-AC1E4407-6139-412A-B4AA-1F102942EA94.html> and has been the same since vRA 8.0.1. Previously, this API returned both a refresh token and an access token. The access token was not fully registered in vRA and could not be used with a number of APIs. To avoid confusion, this API now returns only a refresh token to be used in the second step of the process.

Customers who call the vRealize Automation APIs from vRealize Orchestrator should use the vRealize Orchestrator Plug-in for vRealize Automation. This eliminates the need to retrieve or manage the token and its expiration because the plug-in handles this automatically. For information about the plug-in, see [Using the vRealize Orchestrator Plug-in for vRealize Automation](#)

Accessibility enhancements

Significant improvements in accessibility enhancements to follow the Web Content Accessibility Guidelines (WCAG) 2.1 Level A and AA standards. VMware Accessibility Conformance Report for vRA 8.4 is targeted to be published by end of May 2021. For VMware Accessibility Conformance Report for the earlier vRA 8.2 version, refer <https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/product/vpat/Vmware-vrealize-automation-8.2-vpat.pdf>. For more information and to stay up to speed on accessibility efforts at VMware, visit <https://www.vmware.com/help/accessibility.html>.

Policy criteria support for additional Integer/String operators

Enhanced support for Integer and String based operators for policy criteria allow the cloud administrator to define policies with additional granularity.

- Integer operators greater than, less than, greater than or equal and less than or equal have been introduced for criteria clauses 'Total Memory (MB)' and 'CPU Count'.
- String operator 'contains' has been introduced for criteria clauses 'Created By' and 'Owned By'.
- String operator 'Matches regex' has been introduced for criteria clauses
- Boolean values (True/False or On/Off) for operators 'equals' or 'not equals' is now available for resource attributes like 'Has Snapshots' and 'Power State'

Policy criteria support for resource tags across all policy types

Enhanced support for resource based tags as additional criteria enables vRA cloud administrators to define granular policies that can target deployments with resources that have specific tags.

The resource tag policy criteria clause is available across all policy types.

Networking: Reconfigure Existing Security group for vSphere and VMC - Iterative and Day 2

Reconfigure Security Group (Day-2 and Iterative deployment) action allows you to modify, add or remove rules of an existing security group for a running application in vSphere or VMware Cloud on AWS. See [Day 2 Actions](#).

Networking: Change On-Demand and Existing Security groups for VMC - Iterative and Day 2

Change Security Groups (Day-2 and Iterative deployment) action allows you to associate or dissociate a security group (existing/new) that is part of VMware Cloud on AWS deployment to one or more machines in the deployment. You can attach/detach the security group in blueprint to/from respective machine(s), and update deployments with this new topology through iterative development.

If you want to add an additional security group (existing/new) that is not part of a deployment, to one or more machines in the deployment, you can add the additional security group in blueprint and add (attach) it to machine(s), and update deployments with this new topology through iterative development. See [Day 2 Actions](#).

Hostname is updated in Ansible Tower

Previously, when a machine was provisioned by vRA, the IP Address of the machine was added in the Ansible Tower instead of hostname. In this release Hostname is added to the `ansible_host` variable in Ansible Tower.

The Hostname or FQDN string can be passed to Ansible Tower from Cloud Template.

Support for multi-vm/disk configuration

- You can specify creation of multiple VMs with several disks attached to them.
- Support for Day 2 actions on all disks created for the VMs
- Easy identification of the disks attached to the respective VMs

Add disk with different sizes

In this release, vRA cloud templates allow configurations of different size disks.

Changing deployment projects for onboarded deployments

Change project as a day 2 action for onboarded deployments

- Day 2 action is only available for onboarded deployments in this release. With the 8.4 release, only Disks and Machines can be onboarded. If an onboarded deployment is updated to add any provisioned resources, the change project action is not available. If the provisioned resource is deleted, then the change project becomes available again.
- In case of any failure, the action is not automatically rolled back. You can manually initiate the action again.
- The same resource Cloud Zone's should be present in the target Project otherwise subsequent day2 actions might not work as expected. This preconditions are not enforced. This is consistent with the existing Onboarding logic.
- See [Day 2 Actions](#).

Documentation to configure proxy for vRA on premises Terraform environments

Added documentation to configure proxy for Terraform execution environments for vRA on-prem 8.2 and above.

Unregister onboarded machines from vRA

You can now unregister onboarded machines from vRA

- Unregister action is available for "onboarded" machines only.
- This action removes the resource from the deployment and makes it available for onboarding flow again
- When "unregistering" the onboarded machine, any attached disks (that got onboarded along with machine) are unregistered automatically.
- Once you add additional disks to the onboarded machine, the machine is no longer treated as onboarded anymore and the unregister functionality is not available.

Single secret store

Extensibility action secrets are now named "action constants"

Action constants share the same list of project service secrets. There is no action needed for users have existing action constants from a previous release.

Operations center: Custom roles support

Insights, Alerts and Optimizations can now be filtered by custom roles having read only/read write access to Cloud Zones, Projects and Deployments.

Operations center: Cloud zone Insights enhancement

Cloud zone insights now show projects along with their reclaimable capacity.

Operations center: Distinguish optimizable deployments

Optimizable deployments can now be filtered from a deployment list to easily reach them.

Specify order and SCSI controller for vSphere disks

When creating new disks with deployments, you can:

- In the cloud template, you can specify the order in which the disks are created. This allows for better identification of disks for day 2 actions
- In the cloud template, you can specify which SCSI controller needs to be mapped to the disk. vRA supports a total of 4 SCSI controllers per deployment and you can choose among these 4 for each of the disks.

Support for disks which are part of the image template

There can be instances where an image template has disks in addition to the boot disk. In such cases, vRA supports these disks for day 2 actions. You can view these disks under the VM details and take day 2 actions such as resize on these disks. This resize action is on the VM object in the deployment diagram and shows all disks connected to the VM. See [Day 2 Actions](#).

Disk placement should align with the VM in Workload placement\Multi-VM scenario

Previously, when creating multiple VMs in a single deployment (using the count field), there was a possibility that the disk might not always go to the same cluster that hosts the VM. Now, the disk placement is always on the cluster that hosts the VM for optimal performance.

Storage allocation as per full VM size

Previously, when storage was allocated for a template/content library based deployment, it was only allocate based on default capacity and resize later once the full details are known post deployment. Now, storage is allocated for the full deployment size including image data disks so that Workload placement with vROPs is not affected. This also includes the capacity of any data disks which are part of the template.

Simplification of onboarding workflow

The onboarding plan creation workflow is now simplified to make it easier to bring VMs under vRA management. The rules option is deprecated and the workflow now allows direct selection of machines. The machines view shows only those VMs that were explicitly selected by the user.

Support for Azure image gallery

vRA now supports the image gallery to:

- Support provisioning using custom images residing in an image gallery
- Leverage the same image across multiple Azure subscriptions

Snapshot management for Azure disks

You can create and manage disks snapshots with Azure deployments.

- Support for create operations on snapshots
- Support only for managed disks
- See [Day 2 Actions](#).

Support for Azure disk encryption sets

Support for Azure disk encryption sets to:

- Support third-party KMS systems which leverage encryption sets
- Support encrypting VMs and all the attached disks (current and future) with the same key

Enhanced support for Azure availability sets

Enhanced support for availability sets to:

- Support reusing existing availability sets in the cloud template
- Support having the availability set as optional so that the resources are not part of any availability set

Ansible enhancements

- Previously, when a machine is provisioned by vRA, the IP Address of the machine was added in the Ansible Tower instead of hostname. Now, the Hostname is added to ansible_host variable in Ansible Tower. The Hostname or FQDN string can be passed to Ansible Tower from Cloud Template
- New Ansible Tower blueprint property – maxJobRetries which retry Ansible Playbooks
- Ability to call workflow templates from Ansible Tower integration
- Ansible integration with user account execution
- In Ansible open source vRA is creating server using hostname instead of IP Address
- Able to Pass additional variables from blueprint yaml to Ansible tower
- Update the "Prompt on launch / Limit" for Ansible tower integration to use default value

Puppet enhancements

- Pass user defined properties from Blueprint as facts to Puppet master from agent node
- Specify PE master of masters.

Event Broker enhancements

Ability to add subscriptions at post provisioning stage and before power on.

Release of vRA STD + and SaltStack SecOps add-on in rest of world

- With the approval of the Export Compliance in February vRA STD + and SaltStack SecOps, both offerings can be made available outside of the United States

SaltStack Config

- Provides the capability to apply a SaltStack Config license using VMware Lifecycle Manager
- SaltStack Configure is now FIPS compliance
- Determine FIPs (enabled or disabled) mode during deployment

ITSM plug-in

- Support for Catalog Items which has Custom Resource (without for vRO Objects)
- Support for Catalog Items with Custom Day 2 actions
- Ability to customize vRA Catalog by adding Edit Box and Drop down in ServiceNow.
- Ability to add to attach a script to these fields.
- Deployment Details on available in ServicePortal

vRA plug-in

VMware vRealize Orchestrator Plug-in for vRealize Automation allows interaction between vRealize Orchestrator and vRealize Automation.

The out-of-the-box workflows provided with the plug-in help you deploy and manage resources in vRealize Automation in automated way. In addition to the provided workflows, you can create and run custom workflows. Newly provided content in vRO compatible with vRA 8.x, solves the main customer use cases to create and run workflows for the main functions in vRA like managing projects and users, use custom types, manage VMs, etc.

The same plug-in is applicable for vRA on-prem and for vRA cloud.

vRA plug-in phase 1

- Host management and CRUD operations for on-prem and cloud vRA hosts
- Out-of-the-box workflows for host management
- Preserve Authentication to the hosts and dynamic host creation to use it on the fly
- Rest client available allowing requests to vRA

For vRA 8.4 on-prem the plug-in will be pre-installed with the embedded vRO in vRA.

Plug-in is supported for vRA version 8.0 and should be manually downloaded and installed

Plug-in is supported for vRA version 8.3 and should be manually downloaded and installed.

For external vRO it should be manually downloaded and installed.

For vRA cloud manual download and installation of the plug-in from marketplace will be required.

ABX Scale

When running ABX actions, you can reclaim K8s pods to prevent exceeding physical infrastructure limits. Also, ABX actions can be scheduled across the vRA cluster so the number of concurrent ABX action runs also is larger.

GCP Sole Tenancy

You can now set a custom property to take advantage of the GCP Sole Tenancy capability (dedicated host).

IPAM registration for vRA 7.x workloads while onboarding into vRA 8.x

When onboarding resources that are part of vRA 7.x to v8, the IPAM registration is updated for these workloads. This ensures that there is no duplicate assignment with the IPAM provider and also ensures that the IPs are released back to the pool once the workloads are deleted.

Force deleting deployments for the IaaS API endpoint.

We have added a force delete functionality to the IaaS API endpoint for deleting deployments. The option is used with the “forceDelete” query parameter.

If “forceDelete” = true, then best effort is made for deleting the deployment and all related resources. It should be used with caution since in some situations it may leave provisioned infrastructure resources behind which users should then remove manually.

If “forceDelete” = false, a standard delete action will be executed.

OpenShift support for Terraform Integration in vRA

We now support setting up Terraform runtime environment using OpenShift for the Terraform service integration in vRA.

Migration Assessment support for vRA 7.3

The migration assessment service now supports migration from 7.3 environments.

NEW NSX-T security tag support for day 0 and day 2 operations

You can now use NSX-T security tags with day 0 and day 2 operations.

Before You Begin

Familiarize yourself with the supporting documents.

- [Install vRealize Automation with vRealize Easy Installer](#)
- [Administering Users in vRealize Automation](#)
- [vRealize Automation Transition Guide](#)

After installing vRealize Automation and setting up your users, you can use the *Getting Started* and *Using and Managing* guides for each of the included services. The *Getting Started* guides include an end-to-end proof of concept. The *Using and Managing* guides provide more in-depth information that supports your exploration of the available features. Additional information is also available in [vRealize Automation 8.4 product documentation](#).

- [Getting Started with vRealize Automation Cloud Assembly](#)
- [Using and Managing vRealize Automation Cloud Assembly](#)
- [Getting Started with vRealize Automation Code Stream](#)
- [Using and Managing vRealize Automation Code Stream](#)
- [Getting Started with vRealize Automation Service Broker](#)
- [Using and Managing vRealize Automation Service Broker](#)

For information on vRealize Orchestrator 8.4 features and limitations, refer to the [vRealize Orchestrator 8.4 Release Notes](#).

API Documentation and Versioning

API documentation is available with the product. To access all Swagger documents from a single landing page, go to <https://<appliance.domain.com>/automation-ui/api-docs> where *appliance.domain.com* is your vRealize Automation appliance.

Before using the API, consider the latest API updates and changes for this release and note any changes to the API services that you use. If you have not locked your API to a version before, you might encounter a change in an API response. As a best practice, use the `apiVersion` variable to lock your API to the version you want to use. For example:

- To lock your APIs to the vRealize Automation 8.3 APIs, use `apiVersion=2021-02-04`
- To lock your APIs to the vRealize Automation 8.4 APIs, use `apiVersion=2021-04-15`

• To lock your APIs to the vRealize Automation 8.4 APIs, use `isapiVersion=2021-04-15`

If left unlocked, your API requests will default to the latest version which is `isapiVersion=2021-04-15`.

For information on how to lock you APIs to a specific version, see the "API Versioning" section of the [vRealize Automation 8.4 API Programming Guide](#).

Service Name	Service Description	API Updates and Changes
iaas-api	This API holds all functionality specific to Provisioning service including infrastructure setup, validation and provisioning of resources in iterative manner.	No change
project-service	This API holds all functionality specific to creation, management and delete of projects	No change
blueprint-service	This API holds all functionality specific to Blueprint services, including creation, validation, and provisioning.	<p><u>New endpoints</u></p> <p>None</p> <p><u>New params:</u></p> <ul style="list-style-type: none"> • GET /blueprint/api/blueprints/{blueprintId}/inputs-schema • GET /blueprint/api/blueprints/{blueprintId}/versions/{version}/inputs-schema <ul style="list-style-type: none"> ◦ New param : maxProperties ◦ New param : minProperties • POST /blueprint/api/blueprint-validation <ul style="list-style-type: none"> ◦ New request param : blueprintVersion
relocation-service	The relocation service is used to define policy and plans for bringing existing VMs from any cloud under management.	No change
migration-service	This service is used to quickly setup a vRA 8 instance based on information in a configuration file a.k.a Zero-Setup	No change
cgs-service	Content Service APIs are used to connect to your Infrastructure as Code content in external content sources (ex: SCM Providers and VMWare Marketplace).	No change
form-service	Define dynamic form rendering and customization behavior in Service Broker and Cloud Assembly VMware services.	No change
Deployment	This API provides access to deployment objects and platforms/blueprints that have been deployed into the system.	No change
Approvals	Enforce policies which control who must agree to a deployment or day 2 action before the request is provisioned	No change
Resource quota policy - Aggregator service	<p>This is new service running inside approval container</p> <p>These APIs provide access to find the resource usage metrics at org, user and project level</p>	<p>New endpoint</p> <ul style="list-style-type: none"> • /aggregator/api/metrics Returns registered metrics in aggregator service
Snapshot Creation for Block Device - Provisioning Service	This API is used to create snapshots for block devices.	<p>Modifications to the existing API as follows:</p> <ul style="list-style-type: none"> • POST /iaas/api/block-devices/{id}/operations/snapshots <p>Added a new properties Map to take input properties during snapshot creation as there are different properties for snapshots in different cloud accounts</p> <ul style="list-style-type: none"> • GET /iaas/api/block-devices/{id}/snapshots/{id1} <p>Added a properties Map in the snapshot response model. The</p>

Service Name	Service Description	API Updates and Changes
		<p>Added a properties map in the snapshot response model. The API response includes the following changes:</p> <ul style="list-style-type: none"> ◦ snapshotProperties added as a new key-value field ◦ isCurrent field is deprecated
Azure Storage profile creation - Provisioning Service	This API is used to create Azure storage profile	<p>Modification to the existing API:</p> <p>POST /iaas/api/storage-profiles-azure</p> <p>Added a new property diskEncryptionSetId while creating Azure storage profile.</p>
Attach Block Device to an Machine - Provisioning Service	This API is used to attach an existing disk to existing machine	<p>Modification to the existing API:</p> <p>POST /iaas/api/machines/{id}/disks</p> <p>Added two new parameters</p> <ul style="list-style-type: none"> • scsiController :- SCSI controller name to attach the disk. Below are the 4 possible values SCSI_Controller_0, SCSI_Controller_1, SCSI_Controller_2, SCSI_Controller_3 • unitNumber: Any value between 0 to 15
Catalog Service (Policies)	These APIs are used to interact with policies created in Service Broker.	<p>New Field in Policy Model Object:</p> <ul style="list-style-type: none"> • scopeCriteria: Criteria <p>scopeCriteria field in policy object is used to define multi-project scoped policies using project-based expressions in the criteria format.</p> <p>New API:</p> <ul style="list-style-type: none"> • GET /policy/api/policyTypes/{id}/scopeSchema <p>Returns the policy scope schema for the policy type with the specified ID.</p> <p>The new API returns the schema for the scopeCriteria field for a given policy type to allow API users to correctly define scopeCriteria, if desired.</p>
Code stream all pipeline-service	These API provide access to Code Stream services.	No change

At VMware, we value inclusion. To foster this principle within our customer, partner, and internal community, we removed non-inclusive language in our documentation.

Customers that upgraded to vRealize Automation 8.4 using the new upgrade bundle might see errors during scale out (similar to patched environments). As mentioned in [KB 79105](#), the ova bundle is hosted on my.vmware.com.

vIDM 3.3.3 does not support IWA (Integrated Windows Authentication) with an embedded Linux connector. vRA 8.x Customers using LDAP or IWA with the external Windows connector are not impacted. For more details refer to [KB 82013](#).

Support for Azure VMware Solution, Google Cloud VMware Engine, and Oracle Cloud VMware Solution

vRealize Automation Cloud is tested and certified to work with VMware's hosted cloud solutions on Microsoft Azure, Google Cloud Platform and Oracle Cloud, called Azure VMware Solution (AVS), Google Cloud VMware Engine (GCVE), and Oracle Cloud VMware Solution respectively. Workloads running on AVS, GCVE or OCVS are now managed by vRealize Automation Cloud after setting up vCenter and NSX-T cloud accounts. For more information, refer to [Azure VMware Solution documentation](#), Google [Cloud VMware Engine documentation](#) and [Oracle Cloud VMware Solution documentation](#). For more information on using these platforms with vRA, see Using vRealize Automation with [Azure VMware Solution](#), [Google Cloud VMware Engine](#), and [Oracle Cloud VMware Solution](#).

Upgrading to vRealize Automation 8.4

Using VMware vRealize Suite Lifecycle Manager, you can upgrade your vRealize Automation 8.x instance to 8.4. For more information, see [Upgrading vRealize Suite Lifecycle Manager and vRealize Suite Products](#)

Resolved Issues

- **Assessment Service swagger is not available**

The assessment service swagger page is not available.

- **Exception in input dialog if properties not defined in object type schema**

If input property is of object type and properties is not defined in json schema, the input dialog in test or deploy blueprint dialog would not load.

- **Cannot send value while deploying with input array field**

Although users can fill the values in input form, UI is sending array of null to blueprint service in test/deployment dialog.

- **You can create a day2 policy with duplicate actions/authorities using API.**

When you attempt to create a policy with duplicate actions/authorities using the API, the system does not perform validation checks and the policy is created.

Note: This does not happen when you create a policy using the UI because the dropdown does not show or allow duplicate selection of entries.

- **Modify /csp/gateway/am/api/login?access_token to return only a refresh token**

The behavior of the /csp/gateway/am/api/login?access_token API is changed. This API is used in the first step of the two-step process to obtain an access token for API integrations. and has been the same since vRA 8.0.1. Previously, this API was returning an access token that was not fully registered in vRA and couldn't be used with a number of APIs. To avoid confusion, this API now returns only a refresh token to be used in the second step of the process.

Known Issues

The following known issues are present in this release.

- **vRA deployment fails to initialize on new setups from Easy Installer**

vRA deployment (single or clustered) fails to initialize on new setups from Easy Installer or vRealize Suite LCM
The error shown in LCM is *LCMVRAVAVACONFIG590003*

Workaround: Retry cluster initialization from within vRealize Suite LCM.

- **Unable to access the onboarding page.**

Navigating to the onboarding page in Cloud Assembly->Infrastructure might cause a 302 status code. This can happen if you have been logged in for a long time.

Workaround: Log out and then log back in.

- **When a vCenter cloud account is updated to add a data center, the resources from this data center are not immediately available for use.**

Changes made to regions (data centers) for a vCenter cloud account do not take immediate effect and require data collection to run.

Workaround: Wait for the next data collection to complete successfully. Data collection runs approximately every 10 minutes.

- **PowerShell tasks appear to be stuck**

When there is no active session PowerShell tasks appear to be stuck. This behavior is seen because the PowerShell process responsible to run the user script is held by Windows system process WmiPrvSE.

Workaround : Login to the system and keep an active session. Lock the screen instead of completely logging out.

- **vRO represents Array types as complex types with only one column, rather than a field whose "type.isMultiple" is true.**

When adding a workflow which has an array input and consequently customizing its form, do not change the ID of the column in the Values tab of the data grid. The default value must stay set at `_column-0_`. Conversely, you can change the label of the column (which is visible in the UI when adding values to the datagrid).

- **License re-configuring is not supported.**

After configuring vRealize Automation with the Enterprise license, the system can not be re-configured to use the Advanced License.

- **vRealize Automation 8 does not support Internet Explorer 11**

You cannot use Internet Explorer 11 with vRealize Automation 8.

Workaround: Use a different browser instead of Internet Explorer 11.

- **RP Canvas is not refreshed after custom resource has been changed or deleted**

• **Blueprint canvas is not refreshed after custom resource has been changed or deleted.**

If you delete a custom resource, the change is not propagated to the Blueprint canvas immediately.

Workaround: The Canvas has a cache mechanism, which can be updated after using refresh button, next to the search pane.

• **Create different custom resources with the same vRO object type is not supported**

In vRA 7.X it was possible to create different custom resources for the same type. This allowed users to define a different set of create / delete / operate actions for the same vRO type with creating different custom resource types. In vRA 8.x We do not support a case where same vRO_Type can be leveraged from different custom resources.

• **vRO workflow is not executed through catalog when there is empty input with reference type**

Null pointer exception appears on attempt to request vRO Workflow with an empty value for the Workflow input with a reference type.

Workaround: Set a default value for the reference type or make the field mandatory.

• **Unsuccessfully provisioned custom resource can't be deleted from a deployment**

When you request a custom resource, if the workflow run that creates the resource fails, a resource in the deployment service is still created (since we are replying to the initial request with a STARTED status which in turn creates the resource in deployment). This resource cannot be deleted since it doesn't contain the metadata that is added upon successful provisioning of the resource in vRO.

Workaround: Right after the first attempt to delete the custom resource, a dialog appears which asks you whether you want to force deletion. Say yes to force its deletion.

• **Custom Resource Name is not propagated correctly to the deployment view list**

When you create a custom resource based on vRO_Type, you usually use a comprehensive display name. Currently this display name is not available in the Deployment view. The resource, which appears in the deployment is identified only by its type.

• **Available option to set time zone from vCenter Machine Console window**

Undefined behavior when user sets time zone from vCenter Machine Console window

Workaround: Don't change the time zone.

• **Tenant Names with different cases are treated the same way**

A tenant named vmware and another one named VMware are seen as the same.

Workaround: Tenants in vRA 8.x are based on hostnames since hostnames are case insensitive the tenant names are also case insensitive. This means that a tenant named VMware is the same as VMWARE or vmware or any other combination cases. The tenant name capitalization may vary and may not be preserved across the application.

• **vRO Workflow presentation with an OGNL expression does not render properly when used as a custom day2 operation in vRA.**

Custom Resource Actions with workflows that have OGNL constraints in their presentation may not render properly and it may not be possible to populate all required fields.

• **Cost\Price functionality does not work with shared infrastructure multi-tenancy**

The pricing functionality might report inaccurate results when configured to a multi-tenant deployment where tenants can share infrastructure resources. This is because pricing does not recognize multi-tenancy. The price is calculated only for the org for which vROPs is added and deployments are created.

• **Deployments with an existing network fail during allocation on vSphere / NSX-v cloud accounts when DRS is disabled on the vSphere cluster.**

When selecting an NSX-V network in the network profile and requesting a deployment with an existing network, the deployment fails during allocation with the message: "Unable to find a common placement for compute...with the network configuration...". This occurs when the vCenter contains clusters with DRS disabled.

Workaround: Enable DRS on the cluster and include the cluster in the vRA cloud zone, or select a vSphere network in the network profile.

• **Service broker forms do not populate default values set in vRO workflow input**

When vRO workflow has a string input set with default value, it does not get automatically propagated in the request form when starting the workflow from service broker.

Workaround: Set the given default value using service broker Custom forms.

• **Service Broker cannot import vRO workflows that have actions in valueList for a string field**

Schema for string field that contains valueList populated by an action cannot be parsed and imported in Service Broker

• **Pulling Docker Images Behind Proxy requires additional configuration**

The ABX service pulls container images from publicly available Internet repositories. If vRA is deployed on an isolated network that does not allow outbound traffic to public sites, a HTTP proxy must be configured. While vRA 8 enables proxy configuration via its CLI, the workflow does not include an automatic setup for the docker service.

Workaround: Such configuration should be made separately. KB article to be determined.

• **Complex objects with type anyOf are not supported in cloud template request forms**

If the form contains anyOf property for a complex object, anyOf will be visualized as a string dropdown instead of different sets of constraints to validate the input.

Workaround: Use Enum type instead of anyOf values.

- **After upgrading to vRealize Orchestrator or vRealize Automation 8.3, some resource elements in the vRealize Orchestrator Client might appear changed or reverted to an older version.**

After upgrading to vRealize Orchestrator or vRealize Automation 8.3, some resource elements in the vRealize Orchestrator Client might appear changed or reverted to an older version. This problem occurs with resource elements that were previously updated in the vRealize Orchestrator Client by using a different source file. After upgrading your vRealize Orchestrator or vRealize Automation deployment, these resource elements can be replaced by an older version. This is an intermittent issue.

Workaround:

1. Log in to the vRealize Orchestrator Client.
2. Navigate to **Assets>Resources**.
3. Select the resource element affected by the problem.
4. Select the **Version History** tab, and restore the element to the appropriate version.
5. Repeat for all affected resource elements.

- **If vRA is upgraded from vRA 8.0/8.1/8.2 to 8.3 and AD is configured for a project, deployment fails with the error message: "Failed to successfully create Computer object in Active Directory".**

In the vRA 8.3, the AD scripts used to create active directory record are updated to support overriding relative DN from values set in blueprint. User has to re-validate the existing AD integration in vRA after upgrade to deploy the new scripts.

Workaround: Revalidate the AD integration account in UI.

- **When FIPS mode is enabled, Code Stream pod restarts in high load conditions.**

When a high number of concurrent pipelines are run with FIPS mode enabled, Code Stream pods are restarted because the memory consumption exceeds the preset limit of 2.5GB.

Workaround:

With FIPS mode enabled, increase the memory limit of the Code Stream pods to 3GB.

1. SSH into the node. For HA setup, SSH into any one of the nodes.
 2. Check the current pod memory limit: `kubectl -n prelude describe deployment codestream-app`
 3. Verify that the limit is: Limits: memory: 2500M
 4. Edit the deployment yaml: `kubectl -n prelude edit deployment codestream-app`
 5. Increase the memory limit, and verify that the limit is: Limits: memory: 3000M6
 6. Code Stream pods will be recreated.
- **When exporting a package using Mozilla Firefox v84, the generated file has a .zip extension instead of .package and cannot be imported in vRO**

When you export a package with Firefox 84.0.2 on MacOS 10.15 the package is saved as a .zip file.

Workaround:

- Use Google Chrome or a different version of Mozilla Firefox
- Change the file extension from .zip to .package

Note: In macOS, modify the file from the terminal, as the Finder application does not support changing the file format from a known format to an unknown.

- **Execution of local scripts on a virtual machine via a Workflow "Software-Install-Base" triggered by a vRealize Automation Custom Resource can cause the deployment to fail with an error "An Item with the same key has already been added: Key: LinkedView".**

If the vRealize Automation Blueprint (or Cloud Template) is setup to execute local scripts via a Custom Resource that references the vRO Workflow "Software-Install-Base" which has a Dynamic Type: *DynamicTypes:CustomScript.Script* then the deployment fails.

Workaround: Standup a SaltStack Server to run scripts locally on the machine or use another method of local script execution like cloud-init or ABX, or Code Stream.

- **After a single-node installation, the RaaS log shows the error: No such file or directory. Additionally, ctypes.util.find_library() did not manage to locate a library called '/var/lib/raas/unpack/_MEluxtDsP/Cryptodome/Util/./Cipher/_raw_des.so'.**

This error only occurs at the time of installation and only shows up once in the log.

Workaround: There is no impact to RaaS, so no additional action is necessary.

- **Machine creation API ignores scsiController and unitNumber provided to attach the disk to the machine being created.**

The API request `POST /iaas/api/machines` is used to create machines and uses `scsiController` and `unitNumber` to attach the disk. Currently, this API creates the machine and attaches the disk but ignores the user input for `scsiController` and `unitNumber`.

Workaround: Attach disk to the machine separately using the API request `POST /iaas/api/machines/{id}/disks` with the `scsiController` and `unitNumber`.

- **vRealize Automation 8.4 includes upgraded Clarity 3.0.0 designs, causing some product layouts to look different compared to previous releases.**

With the upgrade to Clarity, some UI elements might look different, specifically custom forms. For a list of Clarity version 3 changes, see the [Clarity 3.0.0 Changelog](#).

- **New A long running migration might fail with a 401 error.**

Due to limitations in vRA 8, the migration service uses user tokens to communicate with other vRA services when transferring data. For long-running migrations, the token can expire before the migration completes resulting in migration failure.

Workaround: Temporarily increase the TTL time for a Provisioning client token in the vIDM UI. Login to the vIDM UI using admin credentials. Navigate to Catalog > Settings > Remote App Access. Locate the Provisioning client and increase the Access Token Time-To-Live(TTL) from 8 hrs to 24 hrs.

- **NEW Deployments fail for blueprints that contain compute tags longer than 256 characters or a key larger than 128 characters**

Deployments fail if a blueprint contains compute tags longer than 256 characters or a key longer than 128 characters.

Workaround: Reduce tag or key lengths below the character limits.

- **New vRealize Automation upgrade fails with error code LCMVRAVACONFIG90030 due to password expiration**

During upgrade, when the root password is set to non-expiring `{color:#FF0000}or{color}` has not been changed in over 365 days, the password is updated to expire immediately. As a result LCM can not connect to vRA to check upgrade status and upgrade fails.

Workaround: Update the password before the upgrade.

- **NEW Configuring the IP Address RELEASED period does not work in a multi-tenant environment.**

The task that runs globally to move IP addresses from RELEASED to AVAILABLE is not tenant-aware. In a multi-tenant environment, where one or more tenants has configured an IP address timeout, only one timeout value is applied to all the tenants.

This issue is being addressed and will be resolved in a future release.

Workaround: None.

- **NEW Intermittent failure to deploy machine connected to an NSX-T network and contains tags**

The deployment fails with an error: "SecurityException: : : Failed to query unique virtual machine by external id[UUID]".

This occurs when vRA queries NSX for the machine in order to tag it on NSX and receives multiple records since the machine is migrated during vMotion.

Workaround: Try to deploy again or disable vMotion.

- **NEW NSX-V to NSX-T migration does not update description objects.**

When migrating from vRA 7 -> 8 `load_balancer_description` records have `endpoint_links` set, but greenfield deployments in vRA 8 do not have this field set.

NSX-v to NSX-T migration does not update any description objects as the `endpoint_links` are not expected to be set there. After a V2T migration, the load balancer description records still point to NSX-v and get deleted when the NSX-v cloud account is deleted. This causes problems with Day-2 operations on such deployments.

Workaround: Do not delete the NSX-v cloud account after migration if the vRA 7 to 8 migration was performed.