



# vRealize Automation 7.3.1 Release Notes

Updated on: 27 FEB 2019

vRealize Automation | 15 MAR 2018 | Build 8004225

Check for additions and updates to these release notes.

## What's in the Release Notes

The release notes cover the following topics:

- [What's New](#)
- [System Requirements](#)
- [Documentation](#)
- [Before You Upgrade](#)
- [Resolved Issues](#)
- [Known Issues](#)

## What's New

This release contains 73 resolutions for known issues in vRealize Automation 7.3.

**New** Updated SLES 11 Linux kernel to version 3.0.101-108.35-default to mitigate CVE-2017-5753, CVE-2017-5715 (Spectre), and CVE-2017-5754 (Meltdown)

For more information, see [Resolved Issues](#).

## System Requirements

For information about supported host operating systems, databases, and Web servers, see the [vRealize Automation Support Matrix](#).

## Documentation

The process is changed for upgrade and migration to vRealize Automation 7.3.1. These changes are added to the vRealize Automation 7.3 Upgrade Guide and Migration Guide. For vRealize Automation 7.3 documentation, see [VMware vRealize Automation](#).

## Before You Upgrade

For general guidance, see [Upgrading vRealize Automation](#).

The vRealize Production Test Upgrade Assist Tool analyzes your vRealize Automation 6.2.x environment for any feature configuration that can cause upgrade problems and checks that your environment is ready for upgrade. To download this tool and related documentation, go to the [Download Product](#) page.

## Using Your License to Enable vRealize Code Stream

You can now use your vRealize Automation license to enable vRealize Code Stream for your vRealize Automation environments. The vRealize Automation license allows you to unlock vRealize Code Stream so that you can use it with the vRealize Code Stream Management Pack for IT DevOps.

To enable vRealize Code Stream, you must have either vRealize Suite Enterprise, vRealize Automation Advanced, or vRealize Automation Enterprise, and a new license for vRealize Automation 7.3.

You must install the vRealize Code Stream Management Pack on a separate and dedicated appliance that has vRealize Automation and vRealize Code Stream enabled in a non-HA (High Availability) mode. For more information, see the [vRealize Code Stream Reference Architecture Guide](#). You cannot install the vRealize Code Stream Management Pack on your existing vRealize Automation production instance for the following reasons:

- The extra load might negatively impact the performance of your vRealize Automation production instance.
- vRealize Code Stream does not support installation on an appliance that is configured for HA. Do not enable vRealize Code Stream on a vRealize Automation instance that is in HA mode. If you do, your vRealize Automation environment might become unstable.

To install the management pack, see the VMware vRealize Code Stream Management Pack for IT DevOps Installation Guide. You can download this guide when you download the management pack.

To enable vRealize Code Stream in the vRealize Automation Installation Wizard or the vRealize Automation Appliance Management Interface, select the **Enable vRealize Code Stream** check box.

If you apply your license without enabling vRealize Code Stream, you can enable vRealize Code Stream later. Apply the license again, and select **Enable vRealize Code Stream**. After you apply the license, you must restart the vRealize Automation appliance.

For more information, see:

- vRealize Code Stream Information Center. See [vRealize Code Stream Installation](#) and [Apply a vRealize Code Stream License to an Appliance](#)
- vRealize Automation Information Center. See [Licensing vRealize Code Stream](#).

## Resolved Issues

The resolved issues are grouped as follows.

- [Upgrade](#)
- [Configuring and Provisioning](#)
- [Maintenance](#)

## Upgrade

- **Migration from vRA 6.2.x to vRA 7.2 or 7.3 leaves a large number of mismatched and empty resources in Deployments or fails with "504 Gateway time-out" or "413 Request Entity Too Large" errors**  
This issue occurs because, migration of environments that contain a large number of machines or blueprints can result in a timeout between vRA services resulting in incomplete migration of machines and blueprints.
- **vRealize Automation 6.2.3 to 7.3 migration fails**  
Migration of environments that contain a large number of machines or blueprints can result in a timeout between vRealize Automation services resulting in incomplete migration of machines and blueprints.
- **Software Agent migration fails because of missing software agent queues**  
After migration, the software agent topics and the agent's underlying queues are not present in the new appliance.
- **Governance upgrade can restart and fail if it runs for more than 100 seconds**  
vRealize Automation upgrade performs migration of the governance content from the infrastructure setup to the virtual appliance using HTTP protocol and has a 100 second timeout set by default for HTTP connections. Exceeding the timeout restarts the operation and causes multiple updates of the same data on the virtual appliance to fail. This version of vRealize Automation increases the timeout to 1 hour and makes timeout configurable.
- **A Distributed Execution Manager (DEM) or Distributed Execution Manager Orchestrator (DEO) does not update when you upgrade to vRealize Automation 7.3.x.**  
The DEM or DEO IaaS component must be installed in the default location at c:\program files (x86)\vmware\vcac when you upgrade to vRealize Automation 7.3.x. If these components are not installed in the default location, they do not update during upgrade.
- **Endpoints are missing after upgrading to vRealize Automation 7.3**  
After a successful upgrade to vRealize Automation 7.3, the Endpoints page in the vRealize Automation console does not display all the endpoints.
- **After upgrading from 6.2.x to 7.x, extra disks are created on the existing blueprints**  
Upgrade was not setting "userCreated" flag in upgraded blueprints, so any disks added by the blueprint creator to those already on the template could not be deleted from a blueprint after the upgrade.
- **When upgrading from vRealize Automation 6.x to 7.3, migration of machine blueprints fails with the error "Cannot add X disks for Component Machine. Maximum number of volumes is set at 0."**  
In vRealize Automation 6.x, a value ZERO for ReconfigureMaxDisks means that you can not add more disks during request time than are already present in the blueprint. In 7.x,

the same limitation is achieved by setting max\_volumes to the total disks that you want to limit. Migration from 6.x to 7.3 fails because max\_volumes of ZERO means no disk can be added, contrary to expectation. If the 6.x value for ReconfigureMaxDisks is ZERO, calculate the number of disks in the associated template and set max\_volumes to that count.

- **After migrating to vRealize Automation 7.x, custom properties defined at multi-machine blueprint level are missing**  
Custom properties on multi-machine blueprints were not migrated during blueprint upgrade.
- **After upgrade from vRealize Automation 6.2.x to 7.x, the blueprint property does not include "encrypted" and "show in request"**  
Upgrade of a multi-machine blueprint does not handle encrypted properties. If a 6.2.x blueprint has a property that is marked "encrypted" and "show in request," that property does not retain "encrypted" and "show in request" in the upgraded blueprint after you upgrade to 7.x.
- **New When you reboot or start clustered appliances at the same time in an upgraded or newly installed vRealize Automation 7.3 environment with three appliances in the cluster, the licenses are missing and services do not register on the secondary nodes**  
See [Knowledge Base article 2150833](#).

## Configuring and Provisioning

- **Occasionally, some vApp templates are marked as missing and any associated blueprints are changed to unpublished**  
This happens when there is a problem reading a template from vCloud Director during data collection. vRealize Automation breaks the relationship between the parent and child template in the virtual machine table, and marks the template as missing.
- **If you reconfigure a virtual machine by increasing or decreasing the CPU or memory or by extending an existing disk, the change fails unless you add a new disk**  
Now when you make a single reconfigure request to change both the CPU and memory of a virtual machine, only one of the changes, either to the CPU or memory, succeeds unless you add a new disk.
- **When users submit multiple post provision actions concurrently on the same resource, the requests get stuck in the IN-PROGRESS state.**  
Sometimes requests get stuck if 2 actions are submitted on a single resource at same time. Now only one request on any resource can be IN-PROGRESS at any given time. An error message similar to this appears if another request is in progress "Request [5] must be completed before we accept new request on resource [bp1-31560851]."
- **When you submit an XaaS blueprint, this error message appears: "The data specified within the request is invalid. The '~' field is required."**  
This error message appears when you create an XaaS blueprint based on the vRealize Orchestrator work flow with mandatory constraint inputs, publish the XaaS blueprint, wrap

it in a composition blueprint (CBP), entitle the CBP, and request it. The request process is changed to prevent this error.

- **When user requests an XaaS catalog item, the XaaS data bindings are reset after user submits the request form or changes a page**  
The validation check creates an inconsistent state by redrawing the whole form instead just the affected fields.
- **When a user requests an XaaS catalog item, some fields are colored with a red border and marked as invalid**  
Previously, a custom validation check for the whole form took place when the form was initially rendered. If the form had custom validation for some fields and the initial value was invalid, those fields were marked as invalid before any user interaction. Now the validation check only takes place after a user changes the field values.
- **Unable to enable or disable storage for a specific reservation**  
Storage in a reservation can be disabled on the Edit Reservation page under the **Resources** tab. When you click the storage icon to edit a storage path and select the **Disabled** check box, the change is not seen in the **Disabled** column after you save your work. When you edit the storage again, the Disabled check box is deselected.
- **In the advanced search under Reservations, the drop-down menus are empty**  
User is unable to search reservations by filtering for the tenant, business group, or reservation type.
- **Folders, also called namespaces in vSAN, are not deleted when virtual machines with user-created disks are provisioned and deleted.**  
Improved the handling of the DestroyVM workflow in vSAN installations when additional disks are added to cloned virtual machines.
- **Openstack data collection fails for networks with missing tenant information**  
vRealize Automation collects network details from an OpenStack provider to support provisioning. In OpenStack, networks are bound to specific tenants or shared across many. While the restricted networks always have information about the parent tenants, shared networks may not provide such details. This causes a problem when vRealize Automation collects shared networks from the provider, since vRealize Automation always expects a valid tenant associated with a network, regardless of its visibility. In addition to this, OpenStack HA installations silently generate hidden networks not bound to a tenant either, thus contributing to the problem from another perspective. vRealize Automation now assigns default tenant details to the networks lacking such data at data collection time.
- **Degradation in vIDM performance because of large number of OAuth Clients**  
Health Service generates new OAuth Clients on every test run. On scheduled execution, this doesn't scale well.
- **When you reprovision an existing virtual machine, you are required to choose optional properties**  
The system presented optional custom properties as required properties in the reprovision request form.

- When system contains 400 or more provisioned VRealize Automation resources, the resources are not all visible in the CAFE plugin inventory in vRealize Orchestrator**  
 Server.findAllForType(); does not return all resources and the result set contains duplicate objects. This happens because of a CAFE plugin paging limitation and because resources are not sorted when they are obtained.
- In a vRealize Automation clustered environment, the container service might not respond when Docker hosts are removed after Xenon service is running for several days and when one or more nodes are restarted, you see some inconsistencies**  
 This issue occurs because of problems during the setup of the Xenon cluster and in the clustering implementation.
- Data collection does not display a virtual machine under the Managed Machines tab when its vCenter instance ID is changed but its BIOS universally unique identifier (UUID) is the same**  
 Now the system provides an option for data collection to recognize a virtual machine when only the BIOS UUID is not changed. In the ManagerService.exe.config file, the data collection option "MatchBiosUuid" is set to false by default, for example, add key="MatchBiosUuid" value="false." This option can be set to true when the BIOS UUID and the vCenter instance ID have not changed.
- Users see different outputs in 6.x and 7.x for the same API call /requests/request\_id/forms/details**  
 Changes to the vRealize Automation API in 7.x added new optional fields to the output. We have updated the API functions to be backward compatible with 6.x. Now all 6.x API functions are available in 7.x if you pass a header with (version=6.2). If no header is passed, the 7.x API is called.
- Some vRealize Automation 6.x API functions do not work in 7.x**  
 Previously some vRealize Automation 6.x API functions were not backward compatible and would not work as expected in 7.x. Now all 6.x API functions are available in 7.x if you pass a header with (version=6.2).
- The root partition quickly fills up with log files in the /var/lib/vrhb/logs folder**  
 Logs were placed in the root partition instead of the /var/log partition and proper log rotation was not in place.
- When a form is reset after setting a value in Manual User Action, the values of some of the fields are reset**  
 Fields are reset if the attribute values in the vRealize Orchestrator workflow that are mapped to the user action are set during the execution. Reset behavior is added to these fields by XaaS, so every time that a form is reset by setting a refresh-on-change field, the values are lost.
- Internal error appears when opening an XaaS request details page that contains vRealize Orchestrator workflow outputs with array of composite type or vRealize Orchestrator workflow outputs with array of string type that are displayed in JSON format.**  
 The mapping of vRealize Orchestrator output of array type was improperly set by default to string representation.

- **The list of property groups applied to a blueprint changes when saved and reopened with the design canvas.**

After saving property groups in a composite blueprint, reopening that blueprint in the design canvas and viewing the list of saved property groups can show that some groups have changed their names to the names of other property groups in the system. This is likely to happen when customers use naming schemes that have similar string patterns throughout all the property groups saved in vRealize Automation. The design canvas can choose approximate matches instead of exact matches when rendering the list of property groups. This can result in an incorrect list of property groups being saved in the blueprint if users save after viewing the list.

- **The message board text box is not resizing**

The message board text box on the home screen does not resize.

- **The vertical order on the component tree for composite blueprint requests has an unpredictable sort that cannot be configured by users**

Previously when you requested a composite blueprint, you could not modify the inconsistent sorting of the blueprint components on the left-side tree panel. Now the vertical order follows the left-to-right and top-to-bottom order of blueprint components that you see in the design canvas. Blueprints present or upgraded before this update are not affected, nor are blueprints that were not created or edited on the design canvas.

- **Notification emails do not work as expected**

Notification emails are not generated for business group managers for virtual machines migrated from vRealize Automation 6.2.x to 7.2. Lease changes on resources with business group and reservation changes send notification emails to wrong business group managers.

- **When manually installing a 7.3 IaaS Website component, the IaaS installer displays a certificate validation error.**

The error message appears when you click Next on the IaaS Server Custom Install page with the Website component selected. This error message is a false negative and appears even when you select the right option.

- **The Edit button on the disk grid on Storage tab is disabled for users who are not business group managers**

Users who are not business group managers are unable to modify the VM disk at request time.

- **Search by owner for deployments or machines under the Items tab doesn't work correctly**

When there was a case difference between the ref field of the user we were searching and the same field in the database, no results were returned. Search is now case insensitive.

- **Users assigned to the Support role in a business group have access to edit custom properties when reprovisioning a machine.**

In vRealize Automation 7.3, the Support role included a privilege typically granted only to Business Group admins. This fix removes that privilege. Users with the Support role will notice that the machine's property grid no longer appear when requesting a reprovision of a machine. Custom properties exposed to requests using the "Show in Request Form" still

appear.

- **Provisioning or data collection requests are stuck in the "In progress" status. The logs indicate that more than one Manager Service is active at the same time.**  
The passive Manager Service is not able to stop one or more of the scheduled operations it manages (e.g., Data Collection) and the operations continue to be executed after the Manager Service node has entered the passive state. This results in multiple nodes racing for the same tasks, which leaves some of those tasks in an inconsistent state. Example: When examining a passive Manager Service log the similar messages could be observed after the node has entered the passive state: [UTC:2017-09-19 21:00:35 Local:2017-09-19 17:00:35] [Info]: [sub-thread-Id="6" context="" token=""] Successfully marked current node as passive in the database.... [UTC:2017-09-19 21:02:38 Local:2017-09-19 17:02:38] [Debug]: [sub-thread-Id="12" context="" token=""] DC: Created data collection item, WorkflowInstanceId 195158, Task state, EntityID 6e939736-aedf-4eb9-89e2-e983389a9a59, StatusID = d7a22ba4-444e-4028-8663-3ee66f332033
- **Software agent is not supported on Windows 7, 8, or 10 clients**  
Prior releases did not support Windows client 32 bit operating systems, nor did they support 32 bit operating systems for client and server. In 7.3.1, Windows 7, 8 and 10 software agents are now supported in both 64 and 32 bit operating systems.
- **When the replica node has more than one NIC, the current Replica VAMI might show the current node status as N/A**  
This issue occurs when the NIC IP is resolved by the external DNS server by using the machine FQDN, which is not the first NIC configuration as listed by the ifconfig command. Then the external nodes use the IP that is resolved by the DNS, but the replica uses the IP of the other NIC. In the 7.3.1 release, the issue is resolved by matching the current IP to each of the configured NICs using the method for defining local address.
- **Locations such as Canada Central, Canada East, and US West2 are not listed during blueprint creation and vRealize Orchestrator client inventory browsing**  
Unable to see locations such as Canada Central, Canada East, Us West2 during blueprint creation and vRealize Orchestrator client inventory browsing because they were not enabled in 7.2.
- **Provisioning on vCloud Director or vCloud Air can fail with the message: "vCloud Machine: Multiple templates with name XXXXX were found."**  
Provisioning on vCloud Director or vCloud Air could fail if there are multiple templates with the same name. It should also be looking at the vApp template name to identify potential duplicate templates.
- **When configured with Global catalog over SSL (e.g., with port 3269), after a restart of the service, sync attempts to remove the existing user/group**  
When the directory is configured with Global catalog over SSL (e.g., with port 3269, after a restart of the service, the flag that checks the directory type in config-state.json, "isGlobalCatalog" reset to 'false', but should be set to true. As a result, the directory sync does not work properly and if you do a sync it attempts to remove the existing user/group.
- **User wants to set a custom directory search attribute for groups about OpenLDAP IdentityStore that is different from the custom directory search attribute**

A new property, `customDirectorySearchAttributeForGroups`, was added in the `IdentityStore` object, which allows you to set an attribute for OpenLDAP `IdentityStore`. If the attribute `customDirectorySearchAttributeForGroups` is equal to `customDirectorySearchAttribute`, then `customDirectorySearchAttributeForGroups` will not be set in the created/updated OpenLDAP `IdentityStore`.

- **In a high-latency 3-node HA environment with automatic failover enabled, reverse proxy configuration of the appliance's configuration is broken intermittently**

A broken reverse proxy configuration creates nodes that connect to the wrong database. This broken configuration causes service request problems. When you enable automatic failover and there is a high latency between nodes or if the master node fails, the automatic failover agent attempts to promote a new master node. If the promotion process failed while the agent reconfigured the reverse proxies for the other nodes, the agent was unable to roll back and restore the previous reverse proxy configurations. The auto failover agent now detects and repairs a reverse proxy configuration that does not match the actual database state. Now if there is an intermittent promotion failure, the agent detects a bad configuration and repairs the configuration on demand.

- **When a machine provisioning event topic contains multiple blocking workflow subscriptions on the same state, the master workflow continues even if a workflow goes into an error state**

When there are multiple blocking workflow subscriptions on the same state, vRealize Orchestrator workflow errors are not sent to the IaaS. Provisioning continues even though the states are configured to go into an error state. Error codes returned by vRealize Orchestrator workflow executions were not merged before being sent to the next subscription with a lower priority.

## Maintenance

- **Error message similar to this appears in the vRealize Orchestrator logs:  
ch.dunes.model.type.ConvertorException: Unable to convert object,  
'com.microsoft.azure.management.compute.models.VirtualMachine@3c68691'  
plugin exception reason : convertToResult() --> Finder 'Azure:VirtualMachine' :  
unexpected error 'ch.dunes.model.sdk.SDKFinderException: convertToResult() -->  
Finder 'Azure:VirtualMachine' : unable to invoke read method : 'osDiskUri'**

Azure inventory fails in a resource group when any of the virtual machines in the group has Azure managed disks. If a resource group does not have a complete inventory, the virtual machines located in that resource group do not appear. Now the system skips reading the `osDiskUri` property if the disk is managed. This property remains empty for Azure virtual machines with managed disks so that a complete inventory can be taken.

- **The vIDM service cannot start**

On certain occasions the vIDM service failed to start due to the presence of a leftover lock. This is now resolved.

- **The vRealize Automation Appliance Management interface shows status N/A for particular Replicas**

If a network failure occurred during the promote operation causing some of the replica

virtual appliances to fail to confirm a successful reconfiguration to the newly promoted master virtual appliance, the master would delete a part of the information for these replicas. Now the master does not delete unsuccessfully configured replicas but just logs information about them.

- **Name of a dynamic type XaaS output is missing in the request form details in catalog requests and in REST response in catalog-service/api/consumer/requests/<requestID>/forms/details**

When you deployed a custom dynamic type, you could only see the vRealize Orchestrator ID of the resource in the request form details. There was no way to retrieve the created object.

- **Virtual machines show their blueprint component to be unassigned after upgrade or migration from vRealize Automation 6.2.x to 7.3**

The upgrade or migration process reports success migrating virtual machines created in 6.2.x but the logs contain errors similar to:

```
Exception thrown while upgrading machine blueprint. [ID : c5006949-0df0-4e6d-9dca-577791522966] [Name : MyWindowsBlueprint] Value cannot be null. Parameter name: g
```

- **When vRealize Automation virtual appliances in a distributed environment are restarted in the same time, the services in the second and third node do not start**  
When you restart two or more vRealize Automation virtual appliances in the same environment at the same time, no services appear under the **Services** tab in vRealize Automation Appliance Management on nodes two and three. The server log reports that the component-registry service can not start due to a licensing service issue that blocks other services.
- **Users see different outputs in vRealize Automation 6.x and 7.x returned for same API call catalog-service/api/consumer/resources/<resource\_id>"**  
Updates to this API call added catalog item details to parent deployment and removed the details from the child resource. Now the 7.x API is backward compatible with the 6.x API as long as a header with (version=6.2) is passed. If no header is passed, the 7.x API is called.
- **RabbitMQ cluster can occasionally fail to be responsive after network partitioning**  
The watchdog has been improved to detect and unblock more cases where RabbitMQ seems stuck.
- **vRealize Automation limits display of endpoints to 30**  
This issue occurs due to a pagination logic in endpoint listing. This restricts number of endpoints being displayed to 30 only.
- **When the CAFE plugin is used to retrieve a result in a JSON format from vRealize Automation, all values are retrieved as a string, although some values are actually a Boolean or number data type**  
This happens because all values enclosed in parentheses ("" ) are interpreted as strings.
- **After you encrypt a password field using property definitions, the variable is no longer able to pass to vRealize Orchestrator.**  
The system displays this error when you select the SECURE\_STRING option to encrypt a

password field: "The data specified within the request is invalid. The current data type SECURE\_STRING for the 'xxx' field must be STRING." Property definitions does not process the "Password that requires confirmation" type in a drop-down menu when you select the SECURE\_STRING data type to encrypt the password field.

- **vRealize Automation 7.3 API call adds {?businessGroupId,requestedFor} at the end of the href attribute**

The href has optional parameters enclosed in {} for vRealize Automation 7.3. In 7.2, the href just contained the link. This causes Jenkins workflows to fail.

- **After upgrading to vRealize Automation 7.3, vRealize Automation nodes are missing in vRealize Automation Appliance Management**

Post upgrade to vRealize Automation 7.3, when you log in to vRealize Automation Appliance Management using `https://hostname.domain.name:5480` and click **vRA Settings > Cluster**, nodes are missing. However, when you log in to the vRealize Automation host with an SSH client and run this command, `vra-command list-nodes`, you see that all the nodes are present. This issue occurs if the virtual appliance Hostname/FQDN is defined in the system with uppercase characters.

- **In a vRealize Automation clustered environment, the container service might not respond when you remove Docker hosts after Xenon service is running for several days**

Communication between the container service and the Xenon service requires the container service to retrieve a token. Under some circumstances, this operation can leak a connection. After a period of time, the maximum number of connections is reached, and the **Containers** tab disappears.

- **The string type property definitions under the Administration tab are missing regular expression attributes once the definitions are migrated from vRealize Automation 6.2.5 to 7.3**

Regular expression attributes are included during the migration of the definitionsChanges property.

- **Config state JSON file becomes corrupted**

System now backs up file daily.

- **PowerShell scripts fail during execution and system displays error message: "Cannot add a member with the name"**

Properties needed by a PowerShell script are generated in duplicate. This causes a collision during the PowerShell script execution. The bug is fixed by resolving the conflict during the execution of PowerShell scripts.

- **OpenLDAP directory synchronization fails**

When either the user or group object in an OpenLDAP directory does not have the custom search attribute specified on the directory config page, the object cannot be searched and synchronized. Now you can give two custom search attributes for users and groups to enable synchronization of both using different attributes.

- **When you use**

`https://vra_node:5480/config#!/SSL_Certificates/put_ssl_import_certificates` to

replace the IaaS certificates, the task fails and the system displays: **ERROR com.vmware.vcac.cli.configurator.ConfiguratorImpl.processExceptionResult:154 - Command execution failed with unexpected error: java.lang.NullPointerException**  
The certificate changing cluster commands are executed in parallel instead of serial order. This causes multiple attempts to change the certificates which leads to the observed error.

- **When you change a vRealize Automation certificate, the system displays this message in vRealize Automation Appliance Management: error calling OpenSSL tools**

The message appears when the total number of characters in the certificate is divisible by 64 without a remainder, that is, the final line of the certificate is complete. In this case, two new line characters (0DH) are added. One is for the end of the line, and the other marks the beginning of the final certificate delimiter, (-----END CERTIFICATE-----). These additions break the certificate, which is not supposed to have two new line characters in a row.

- **Property with specified display order shows "orderIndex":null in API call**

When you use this API call with the specified display order, it returns the correct orderIndex value.

<https://{{url}}/catalog-service/api/consumer/entitledCatalogItems/{catalogItemId}/requests/schema>

For example:

```
{
  "id": "prop2",
  "label": "prop2",
  "description": null,
  "dataType": {
    "type": "primitive",
    "typeId": "STRING"
  },
  "displayAdvice": "TEXTBOX",
  "orderIndex": 3
}
```

- **Composition service post-registration initialization tasks fail with NullPointerException**

Some earlier deployments could be put into a bad state when service was interrupted by a power outage or taken down during a provisioning operation. These events could remove some blueprint parameter data which was not restored by any of the initialization upgrade tasks.

- **Resources displayed in vRealize Automation are missing in the CAFE plug-in in vRealize Orchestrator**

When retrieving objects from vRealize Automation, the CAFE plugin did not send the request with sorting information. As a result, some resources returned multiple pages while other resources were not returned at all.

- **String input named \_\_asd\_catalogRequestId used to create an XaaS blueprint fails**

### **during the blueprint request because its value is an empty string**

When a client has workflows with `__asd_catalogRequestId` as an input for their vRealize Orchestrator workflows instead of getting the input from the system context, the value of the field is not correctly passed from vRealize Automation to vRealize Orchestrator.

- **When user requests XaaS catalog item, XaaS data bindings are reset after request form submit validation check or after a page change and causes an inconsistent state**  
The validation check redraws the whole form instead only the affected fields.
- **Fields in a XaaS form are marked with red border even when a valid value is selected**  
This happens because string input contains predefined content.

## Known Issues

- **When the primary domain controller is unavailable, the login is very slow or fails**  
When an attempt to contact the primary domain controller fails, vIDM contacts the secondary domain controller. Because vIDM always contacts the primary domain controller before contacting the secondary domain controller, there is a delay in processing the login requests. This causes the requests to pile up and slow down the system.

Workaround: See [Knowledge Base article 52840](#).

- **NEW Custom updates in `setenv.sh` file for vRO are overridden after upgrading.**  
Post upgrade, manually update the values as needed post upgrade and restart the VCO server to apply the changes.
- **You can experience .NET compatibility issues when deploying on Windows 7, 8 and 10 operating systems**  
You must run a script to update the agent bootstrap zip file with appropriate changes required to avoid the .NET error while deploying on Windows 7, 8, or 10.

Workaround: See [Knowledge Base article 53149](#).

- **Endpoints are missing after upgrade to vRealize Automation 7.3 or 7.3.1**  
If the vRealize Orchestrator custom property `VMware.VCenterOrchestrator.Priority` is present on the vCenter endpoint, the endpoint upgrade fails.

Workaround: Prior to upgrade, remove `VMware.VCenterOrchestrator.Priority` from the **Infrastructure > Endpoints > Edit Endpoint** screen.

- **NEW vRA portal is down (services are unavailable) after upgrade or after changing the SSO admin password**  
If the user has configured a custom vRO authentication provider in the vRO Control Center, this value gets overridden by the default value (`vsphere.local/vcoadmins`) after upgrade or after changing the SSO admin password.

Workaround: See Knowledge Base article [56627](#).

- **New Java update 1.8 (versions u192, u201, and u202) are incompatible with 7.3.1**

**installations.**

The latest Java update 1.8 (versions u192, u201, and u202) can cause problems with the IaaS Database installation if the automatic prerequisite fixes have not been applied.

Workaround: Use Java version 1.8 u191. Alternatively, you can install 7.3.1 with the automatic prerequisite fixes enabled to install the older Java version 1.8 u191 on the IaaS machines. After the installation completes, the newer Java version 1.8 u201 or u202 can be installed manually.

Copyright © 2020 VMware, Inc. All rights reserved.