



# Red Hat OpenShift Data Science 1

## Release notes

Features, enhancements, resolved issues, and known issues associated with this release



# Red Hat OpenShift Data Science 1 Release notes

---

Features, enhancements, resolved issues, and known issues associated with this release

## Legal Notice

Copyright © 2023 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, the Red Hat logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux<sup>®</sup> is the registered trademark of Linus Torvalds in the United States and other countries.

Java<sup>®</sup> is a registered trademark of Oracle and/or its affiliates.

XFS<sup>®</sup> is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL<sup>®</sup> is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js<sup>®</sup> is an official trademark of Joyent. Red Hat is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack<sup>®</sup> Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

## Abstract

These release notes provide an overview of new features, enhancements, resolved issues, and known issues in this release of Red Hat OpenShift Data Science. OpenShift Data Science is currently available in Red Hat OpenShift Dedicated and Red Hat OpenShift Service on Amazon Web Services (ROSA).

---

## Table of Contents

<b>CHAPTER 1. OVERVIEW OF OPENSIFT DATA SCIENCE</b> .....	<b>3</b>
<b>CHAPTER 2. SUPPORT REMOVALS</b> .....	<b>4</b>
2.1. NVIDIA GPU OPERATOR REPLACES NVIDIA GPU ADD-ON	4
2.2. KUBEFLOW NOTEBOOK CONTROLLER REPLACES JUPYTERHUB	4
<b>CHAPTER 3. RESOLVED ISSUES</b> .....	<b>5</b>
<b>CHAPTER 4. KNOWN ISSUES</b> .....	<b>14</b>
<b>CHAPTER 5. PRODUCT FEATURES</b> .....	<b>24</b>



# CHAPTER 1. OVERVIEW OF OPENSIFT DATA SCIENCE

Using Red Hat OpenShift Data Science, users can integrate data, artificial intelligence and machine learning software to execute end-to-end machine learning workflows. OpenShift Data Science is supported in two configurations:

1. Installed as an Add-on to a Red Hat managed environment such as Red Hat OpenShift Dedicated and Red Hat OpenShift Service on Amazon Web Services (ROSA).
2. Installed as a self-managed Operator on a self-managed environment, such as Red Hat OpenShift Container Platform.

For data scientists, OpenShift Data Science includes Jupyter and a collection of default notebook images optimized with the tools and libraries required for model development, and the TensorFlow and PyTorch frameworks. Deploy and host your models, integrate models into external applications, and export models to host them in any hybrid cloud environment. You can also accelerate your data science experiments through the use of graphics processing units (GPUs).

For administrators, OpenShift Data Science enables data science workloads in an existing Red Hat OpenShift Dedicated or ROSA environment. Manage users with your existing OpenShift identity provider, and manage the resources available to notebook servers to ensure data scientists have what they require to create, train, and host models.

To learn more, see [Introduction to Red Hat OpenShift Data Science](#) .

## CHAPTER 2. SUPPORT REMOVALS

This section describes major changes in support for user-facing features in Red Hat OpenShift Data Science.

### 2.1. NVIDIA GPU OPERATOR REPLACES NVIDIA GPU ADD-ON

Previously, to enable graphics processing units (GPUs) to help with compute-heavy workloads, you installed the NVIDIA GPU add-on. OpenShift Data Science no longer supports this add-on.

Now, to enable GPU support, you must install the NVIDIA GPU Operator. To learn how to install the GPU Operator, see [NVIDIA GPU Operator on OpenShift](#) (external).

### 2.2. KUBEFLOW NOTEBOOK CONTROLLER REPLACES JUPYTERHUB

In OpenShift Data Science 1.15 and earlier, JupyterHub was used to create and launch notebook server environments. In OpenShift Data Science 1.16 and later, JupyterHub is no longer included, and its functionality is replaced by Kubeflow Notebook Controller.

This change provides the following benefits:

- Users can now immediately cancel a request, make changes, and retry the request, instead of waiting 5+ minutes for the initial request to time out. This means that users do not wait as long when requests fail, for example, when a notebook server does not start correctly.
- The architecture no longer prevents a single user from having more than one notebook server session, expanding future feature possibilities.
- The removal of the PostgreSQL database requirement allows for future expanded environment support in OpenShift Data Science.

However, this update also creates the following behavior changes:

- For IT Operations administrators, the notebook server administration interface does not currently allow login access to data scientist users' notebook servers. This is planned to be added in future releases.
- For data scientists, the JupyterHub interface URL is no longer valid. Update your bookmarks to point to the OpenShift Data Science Dashboard.

The JupyterLab interface is unchanged and data scientists can continue to use JupyterLab to work with their notebook files as usual.



## CHAPTER 3. RESOLVED ISSUES

This section describes notable issues that have been resolved in Red Hat OpenShift Data Science.

### [ODH-DASHBOARD-1639](#) - Wrong TLS value in dashboard route

Previously, when a route was created for the OpenShift Data Science dashboard on OpenShift, the **tls.termination** field had an invalid default value of **Reencrypt**. This issue is now resolved. The new value is **reencrypt**.

### [ODH-DASHBOARD-1638](#) - Name placeholder in Triggered Runs tab shows Scheduled run name

Previously, when you clicked **Pipelines > Runs** and then selected the **Triggered** tab to configure a triggered run, the example value shown in the **Name** field was **Scheduled run name**. This issue is now resolved.

### [ODH-DASHBOARD-1547](#) - "We can't find that page" message displayed in dashboard when pipeline operator installed in background

Previously, when you used the **Data Science Pipelines** page of the dashboard to install the OpenShift Pipelines Operator, when the Operator installation was complete, the page refreshed to show a "We can't find that page" message. This issue is now resolved. When the Operator installation is complete, the dashboard redirects you to the **Pipelines** page, where you can create a pipeline server.

### [ODH-DASHBOARD-1545](#) - Dashboard keeps scrolling to bottom of project when Models tab is expanded

Previously, on the **Data Science Projects** page of the dashboard, if you clicked the **Deployed models** tab to expand it and then tried to perform other actions on the page, the page automatically scrolled back to the **Deployed models** section. This affected your ability to perform other actions. This issue is now resolved.

### [NOTEBOOKS-156](#) - Elyra included an example runtime called Test

Previously, Elyra included an example runtime configuration called **Test**. If you selected this configuration when running a data science pipeline, you could see errors. The **Test** configuration has now been removed.

### [RHODS-8939](#) - Default shared memory for a Jupyter notebook created in a previous release causes a runtime error

Starting with release 1.3.1, this issue is fixed and any new notebook's shared memory is set to the size of the node.

For a Jupyter notebook created in a release earlier than 1.3.1, the default shared memory for a Jupyter notebook is set to 64 Mb and you cannot change this default value in the notebook configuration.

To fix this issue, you must recreate the notebook or follow the process described in the [Known Issues section](#) of these release notes.

### [RHODS-8932](#) - Incorrect cron format was displayed by default when scheduling a recurring pipeline run

When you scheduled a recurring pipeline run by configuring a cron job, the OpenShift Data Science interface displayed an incorrect format by default. It now displays the correct format.

### **RHODS-9374 - Pipelines with non-unique names did not appear in the data science project user interface**

If you launched a notebook from a Jupyter application that supported Elyra, or if you used a workbench, when you submitted a pipeline to be run, pipelines with non-unique names did not appear in the **Pipelines** section of the relevant data science project page or the **Pipelines** heading of the data science pipelines page. This issue has now been resolved.

### **RHODS-9329 - Deploying a custom model-serving runtime could result in an error message**

Previously, if you used the OpenShift Data Science dashboard to deploy a custom model-serving runtime, the deployment process could fail with an **Error retrieving Serving Runtime** message. This issue is now resolved.

### **RHODS-9064 - After upgrade, the Data Science Pipelines tab was not enabled on the OpenShift Data Science dashboard**



When you upgraded from OpenShift Data Science 1.26 to OpenShift Data Science 1.28, the **Data Science Pipelines** tab was not enabled in the OpenShift Data Science dashboard. This issue is resolved in OpenShift Data Science 1.29.

### **RHODS-9443 - Exporting an Elyra pipeline exposed S3 storage credentials in plain text**

In OpenShift Data Science 1.28.0, when you exported an Elyra pipeline from JupyterLab in Python DSL format or YAML format, the generated output contained S3 storage credentials in plain text. This issue has been resolved in OpenShift Data Science 1.28.1. However, after you upgrade to OpenShift Data Science 1.28.1, if your deployment contains a data science project with a pipeline server and a data connection, you must perform the following additional actions for the fix to take effect:

1. Refresh your browser page.
2. Stop any running workbenches in your deployment and restart them.

Furthermore, to confirm that your Elyra runtime configuration contains the fix, perform the following actions:

1. In the left sidebar of JupyterLab, click **Runtimes** (  ).
2. Hover the cursor over the runtime configuration that you want to view and click the **Edit** button (  ).  
The **Data Science Pipelines runtime configuration** page opens.
3. Confirm that **KUBERNETES\_SECRET** is defined as the value in the **Cloud Object Storage Authentication Type** field.
4. Close the runtime configuration without changing it.

### **RHODS-8460 - When editing the details of a shared project, the user interface remained in a loading state without reporting an error**

When a user with permission to edit a project attempted to edit its details, the user interface remained in a loading state and did not display an appropriate error message. Users with permission to edit projects cannot edit any fields in the project, such as its description. Those users can edit only components belonging to a project, such as its workbenches, data connections, and storage.

The user interface now displays an appropriate error message and does not try to update the project description.

### **RHODS-8482 - Data science pipeline graphs did not display node edges for running pipelines**

If you ran pipelines that did not contain Tekton-formatted **Parameters** or **when** expressions in their YAML code, the OpenShift Data Science user interface did not display connecting edges to and from graph nodes. For example, if you used a pipeline containing the **runAfter** property or **Workspaces**, the user interface displayed the graph for the executed pipeline without edge connections. The OpenShift Data Science user interface now displays connecting edges to and from graph nodes.

### **RHODS-8923 - Newly created data connections were not detected when you attempted to create a pipeline server**

If you created a data connection from within a Data Science project, and then attempted to create a pipeline server, the **Configure a pipeline server** dialog did not detect the data connection that you created. This issue is now fixed.

### **RHODS-8461 - When sharing a project with another user, the OpenShift Data Science user interface text was misleading**

When you attempted to share a Data Science project with another user, the user interface text misleadingly implied that users could edit all of its details, such as its description. However, users can edit only components belonging to a project, such as its workbenches, data connections, and storage. This issue is now fixed and the user interface text no longer misleadingly implies that users can edit all of its details.

### **RHODS-8462 - Users with "Edit" permission could not create a Model Server**

Users with "Edit" permissions can now create a Model Server without token authorization. Users must have "Admin" permissions to create a Model Server with token authorization.

### **RHODS-8796 - OpenVINO Model Server runtime did not have the required flag to force GPU usage**

OpenShift Data Science includes the OpenVINO Model Server (OVMS) model-serving runtime by default. When you configured a new model server and chose this runtime, the **Configure model server** dialog enabled you to specify a number of GPUs to use with the model server. However, when you finished configuring the model server and deployed models from it, the model server did not actually use any GPUs. This issue is now fixed and the model server uses the GPUs.

### **RHODS-8861 - Changing the host project when creating a pipeline ran resulted in an inaccurate list of available pipelines**

If you changed the host project while creating a pipeline run, the interface failed to make the pipelines of the new host project available. Instead, the interface showed pipelines that belong to the project you initially selected on the **Data Science Pipelines > Runs** page. This issue is now fixed. You no longer select a pipeline from the **Create run** page. The pipeline selection is automatically updated when you click the **Create run** button, based on the current project and its pipeline.

### **RHODS-8249 - Environment variables uploaded as ConfigMap were stored in Secret instead**

Previously, in the OpenShift Data Science interface, when you added environment variables to a workbench by uploading a **ConfigMap** configuration, the variables were stored in a **Secret** object instead. This issue is now fixed.

### **RHODS-7975 - Workbenches could have multiple data connections**

Previously, if you changed the data connection for a workbench, the existing data connection was not released. As a result, a workbench could stay connected to multiple data sources. This issue is now fixed.

#### **RHODS-7948 - Uploading a secret file containing environment variables resulted in double-encoded values**

Previously, when creating a workbench in a data science project, if you uploaded a YAML-based secret file containing environment variables, the environment variable values were not decoded. Then, in the resulting OpenShift secret created by this process, the encoded values were encoded again. This issue is now fixed.

#### **RHODS-6429 - An error was displayed when creating a workbench with the Intel OpenVINO or Anaconda Professional Edition images**

Previously, when you created a workbench with the Intel OpenVINO or Anaconda Professional Edition images, an error appeared during the creation process. However, the workbench was still successfully created. This issue is now fixed.

#### **RHODS-6372 - Idle notebook culler did not take active terminals into account**

Previously, if a notebook image had a running terminal, but no active, running kernels, the idle notebook culler detected the notebook as inactive and stopped the terminal. This issue is now fixed.

#### **RHODS-5700 - Data connections could not be created or connected to when creating a workbench**

When creating a workbench, users were unable to create a new data connection, or connect to existing data connections.

#### **RHODS-6281 - OpenShift Data Science administrators could not access Settings page if an admin group was deleted from cluster**

Previously, if a Red Hat OpenShift Data Science administrator group was deleted from the cluster, OpenShift Data Science administrator users could no longer access the **Settings** page on the OpenShift Data Science dashboard. In particular, the following behavior was seen:

- When an OpenShift Data Science administrator user tried to access the **Settings → User management** page, a "Page Not Found" error appeared.
- Cluster administrators *did not* lose access to the **Settings** page on the OpenShift Data Science dashboard. When a cluster administrator accessed the **Settings → User management** page, a warning message appeared, indicating that the deleted OpenShift Data Science administrator group no longer existed in OpenShift. The deleted administrator group was then removed from **OdhDashboardConfig**, and administrator access was restored.

This issue is now fixed.

#### **RHODS-1968 - Deleted users stayed logged in until dashboard was refreshed**

Previously, when a user's permissions for the Red Hat OpenShift Data Science dashboard were revoked, the user would notice the change only after a refresh of the dashboard page.

This issue is now fixed. When a user's permissions are revoked, the OpenShift Data Science dashboard locks the user out within 30 seconds, without the need for a refresh.

#### **RHODS-6384 - A workbench's data connection was incorrectly updated when creating a duplicated data connection**

When creating a data connection that contained the same name as an existing data connection, the data connection creation failed, but the associated workbench still restarted and connected to the wrong data connection. This issue has been resolved. Workbenches now connect to the correct data connection.

### **RHODS-6370 - Workbenches failed to receive the latest toleration**

Previously, to acquire the latest toleration, users had to attempt to edit the relevant workbench, make no changes, and save the workbench again. Users can now apply the latest toleration change by stopping and then restarting their data science project's workbench.

### **RHODS-6779 - Models failed to be served after upgrading from OpenShift Data Science 1.20 to OpenShift Data Science 1.21**

When upgrading from OpenShift Data Science 1.20 to OpenShift Data Science 1.21, the **modelmesh-serving** pod attempted to pull a non-existent image, causing an image pull error. As a result, models were unable to be served using the model serving feature in OpenShift Data Science. The **odh-openvino-servingruntime-container-v1.21.0-15** image now deploys successfully.

### **RHODS-5945 - Anaconda Professional Edition could not be enabled in OpenShift Data Science**

Anaconda Professional Edition could not be enabled for use in OpenShift Data Science. Instead, an **InvalidImageName** error was displayed in the associated pod's **Events** page. Anaconda Professional Edition can now be successfully enabled.

### **RHODS-5822 - Admin users were not warned when usage exceeded 90% and 100% for PVCs created by data science projects.**

Warnings indicating when a PVC exceeded 90% and 100% of its capacity failed to display to admin users for PVCs created by data science projects. Admin users can now view warnings about when a PVC exceeds 90% and 100% of its capacity from the dashboard.

### **RHODS-5889 - Error message was not displayed if a data science notebook was stuck in "pending" status**

If a notebook pod could not be created, the OpenShift Data Science interface did not show an error message. An error message is now displayed if a data science notebook cannot be spawned.

### **RHODS-5886 - Returning to the Hub Control Panel dashboard from the data science workbench failed**

If you attempted to return to the dashboard from your workbench Jupyter notebook by clicking on **File** → **Log Out**, you were redirected to the dashboard and remained on a "Logging out" page. Likewise, if you attempted to return to the dashboard by clicking on **File** → **Hub Control Panel**, you were incorrectly redirected to the **Start a notebook server** page. Returning to the Hub Control Panel dashboard from the data science workbench now works as expected.

### **RHODS-6101 - Administrators were unable to stop all notebook servers**

OpenShift Data Science administrators could not stop all notebook servers simultaneously. Administrators can now stop all notebook servers using the **Stop all servers** button and stop a single notebook by selecting **Stop server** from the action menu beside the relevant user.

### **RHODS-5891 - Workbench event log was not clearly visible**

When creating a workbench, users could not easily locate the event log window in the OpenShift Data Science interface. The **Starting** label under the **Status** column is now underlined when you hover over it, indicating you can click on it to view the notebook status and the event log.

#### **RHODS-6296 - ISV icons did not render when using a browser other than Google Chrome**

When using a browser other than Google Chrome, not all ISV icons under **Explore** and **Resources** pages were rendered. ISV icons now display properly on all supported browsers.

#### **RHODS-3182 - Incorrect number of available GPUs was displayed in Jupyter**

When a user attempts to create a notebook instance in Jupyter, the maximum number of GPUs available for scheduling was not updated as GPUs are assigned. Jupyter now displays the correct number of GPUs available.

#### **RHODS-5890 - When multiple persistent volumes were mounted to the same directory, workbenches failed to start**

When mounting more than one persistent volume (PV) to the same mount folder in the same workbench, creation of the notebook pod failed and no errors were displayed to indicate there was an issue.

#### **RHODS-5768 - Data science projects were not visible to users in Red Hat OpenShift Data Science**

Removing the **[DSP]** suffix at the end of a project's **Display Name** property caused the associated data science project to no longer be visible. It is no longer possible for users to remove this suffix.

#### **RHODS-5701 - Data connection configuration details were overwritten**

When a data connection was added to a workbench, the configuration details for that data connection were saved in environment variables. When a second data connection was added, the configuration details are saved using the same environment variables, which meant the configuration for the first data connection was overwritten. At the moment, users can add a maximum of one data connection to each workbench.

#### **RHODS-5252 - The notebook Administration page did not provide administrator access to a user's notebook server**

The notebook **Administration** page, accessed from the OpenShift Data Science dashboard, did not provide the means for an administrator to access a user's notebook server. Administrators were restricted to only starting or stopping a user's notebook server.

#### **RHODS-2438 - PyTorch and TensorFlow images were unavailable when upgrading**

When upgrading from OpenShift Data Science 1.3 to a later version, PyTorch and TensorFlow images were unavailable to users for approximately 30 minutes. As a result, users were unable to start PyTorch and TensorFlow notebooks in Jupyter during the upgrade process. This issue has now been resolved.

#### **RHODS-5354 - Environment variable names were not validated when starting a notebook server**

Environment variable names were not validated on the **Start a notebook server** page. If an invalid environment variable was added, users were unable to successfully start a notebook. The environmental variable name is now checked in real-time. If an invalid environment variable name is entered, an error message displays indicating valid environment variable names must consist of alphabetic characters, digits, `_`, `-`, or `.`, and must not start with a digit.

#### **RHODS-4617 - The Number of GPUs drop-down was only visible if there were GPUs available**

Previously, the **Number of GPUs** drop-down was only visible on the **Start a notebook server** page if GPU nodes were available. The **Number of GPUs** drop-down now also correctly displays if an autoscaling machine pool is defined in the cluster, even if no GPU nodes are currently available, possibly resulting in the provisioning of a new GPU node on the cluster.

#### **RHODS-5420 - Cluster admin did not get administrator access if it was the only user present in the cluster**

Previously, when the cluster admin was the only user present in the cluster, it did not get Red Hat OpenShift administrator access automatically. Administrator access is now correctly applied to the cluster admin user.

#### **RHODS-4321 - Incorrect package version displayed during notebook selection**

The **Start a notebook server** page displayed an incorrect version number (11.4 instead of 11.7) for the CUDA notebook image. The version of CUDA installed is no longer specified on this page.

#### **RHODS-5001 - Admin users could add invalid tolerations to notebook pods**

An admin user could add invalid tolerations on the **Cluster settings** page without triggering an error. If a invalid toleration was added, users were unable to successfully start notebooks. The toleration key is now checked in real-time. If an invalid toleration name is entered, an error message displays indicating valid toleration names consist of alphanumeric characters, `-`, `_`, or `.`, and must start and end with an alphanumeric character.

#### **RHODS-5100 - Group role bindings were not applied to cluster administrators**

Previously, if you had assigned cluster admin privileges to a group rather than a specific user, the dashboard failed to recognize administrative privileges for users in the administrative group. Group role bindings are now correctly applied to cluster administrators as expected.

#### **RHODS-4947 - Old Minimal Python notebook image persisted after upgrade**

After upgrading from OpenShift Data Science 1.14 to 1.15, the older version of the Minimal Python notebook persisted, including all associated package versions. The older version of the Minimal Python notebook no longer persists after upgrade.

#### **RHODS-4935 - Excessive "missing x-forwarded-access-token header" error messages displayed in dashboard log**

The **rhods-dashboard** pod's log contained an excessive number of "missing x-forwarded-access-token header" error messages due to a readiness probe hitting the **/status** endpoint. This issue has now been resolved.

#### **RHODS-2653 - Error occurred while fetching the generated images in the sample Pachyderm notebook**

An error occurred when a user attempted to fetch an image using the sample Pachyderm notebook in Jupyter. The error stated that the image could not be found. Pachyderm has corrected this issue.

#### **RHODS-4584 - Jupyter failed to start a notebook server using the OpenVINO notebook image**

Jupyter's **Start a notebook server** page failed to start a notebook server using the OpenVINO notebook image. Intel has provided an update to the OpenVINO operator to correct this issue.

#### **RHODS-4923 - A non-standard check box displayed after disabling usage data collection**

After disabling usage data collection on the **Cluster settings** page, when a user accessed another area of the OpenShift Data Science dashboard, and then returned to the **Cluster settings** page, the **Allow collection of usage data** check box had a non-standard style applied, and therefore did not look the same as other check boxes when selected or cleared.

#### **RHODS-4938 - Incorrect headings were displayed in the Notebook Images page**

The **Notebook Images** page, accessed from the **Settings** page on the OpenShift Data Science dashboard, displayed incorrect headings in the user interface. The **Notebook image settings** heading displayed as **BYON image settings**, and the **Import Notebook images** heading displayed as **Import BYON images**. The correct headings are now displayed as expected.

#### **RHODS-4818 - Jupyter was unable to display images when the NVIDIA GPU add-on was installed**

The **Start a notebook server** page did not display notebook images after installing the NVIDIA GPU add-on. Images are now correctly displayed, and can be started from the **Start a notebook server** page.

#### **RHODS-4797 - PVC usage limit alerts were not sent when usage exceeded 90% and 100%**

Alerts indicating when a PVC exceeded 90% and 100% of its capacity failed to be triggered and sent. These alerts are now triggered and sent as expected.

#### **RHODS-4366 - Cluster settings were reset on operator restart**

When the OpenShift Data Science operator pod was restarted, cluster settings were sometimes reset to their default values, removing any custom configuration. The OpenShift Data Science operator was restarted when a new version of OpenShift Data Science was released, and when the node that ran the operator failed. This issue occurred because the operator deployed ConfigMaps incorrectly. Operator deployment instructions have been updated so that this no longer occurs.

#### **RHODS-4318 - The OpenVINO notebook image failed to build successfully**

The OpenVINO notebook image failed to build successfully and displayed an error message. This issue has now been resolved.

#### **RHODS-3743 - Starburst Galaxy quick start did not provide download link in the instruction steps**

The Starburst Galaxy quick start, located on the **Resources** page on the dashboard, required the user to open the **explore-data.ipynb notebook**, but failed to provide a link within the instruction steps. Instead, the link was provided in the quick start's introduction.

#### **RHODS-1974 - Changing alert notification emails required pod restart**

Changes to the list of notification email addresses in the Red Hat OpenShift Data Science Add-On were not applied until after the **rhods-operator** pod and the **prometheus-\*** pod were restarted.

#### **RHODS-2738 - Red Hat OpenShift API Management 1.15.2 add-on installation did not successfully complete**

For OpenShift Data Science installations that are integrated with the Red Hat OpenShift API Management 1.15.2 add-on, the Red Hat OpenShift API Management installation process did not successfully obtain the SMTP credentials secret. Subsequently, the installation did not complete.

#### **RHODS-3237 - GPU tutorial did not appear on dashboard**

The "GPU computing" tutorial, located at [Gtc2018-numba](#), did not appear on the **Resources** page on the dashboard.



### **RHODS-3069 - GPU selection persisted when GPU nodes were unavailable**

When a user provisioned a notebook server with GPU support, and the utilized GPU nodes were subsequently removed from the cluster, the user could not create a notebook server. This occurred because the most recently used setting for the number of attached GPUs was used by default.

### **RHODS-3181 - Pachyderm now compatible with OpenShift Dedicated 4.10 clusters**

Pachyderm was not initially compatible with OpenShift Dedicated 4.10, and so was not available in OpenShift Data Science running on an OpenShift Dedicated 4.10 cluster. Pachyderm is now available on and compatible with OpenShift Dedicated 4.10.

### **RHODS-2160 - Uninstall process failed to complete when both OpenShift Data Science and OpenShift API Management were installed**

When OpenShift Data Science and OpenShift API Management are installed together on the same cluster, they use the same Virtual Private Cluster (VPC). The uninstall process for these Add-ons attempts to delete the VPC. Previously, when both Add-ons are installed, the uninstall process for one service was blocked because the other service still had resources in the VPC. The cleanup process has been updated so that this conflict does not occur.

### **RHODS-2747 - Images were incorrectly updated after upgrading OpenShift Data Science**

After the process to upgrade OpenShift Data Science completed, Jupyter failed to update its notebook images. This was due to an issue with the image caching mechanism. Images are now correctly updating after an upgrade.

### **RHODS-2425 - Incorrect TensorFlow and TensorBoard versions displayed during notebook selection**

The **Start a notebook server** page displayed incorrect version numbers (2.4.0) for TensorFlow and TensorBoard in the TensorFlow notebook image. These versions have been corrected to TensorFlow 2.7.0 and TensorBoard 2.6.0.

### **RHODS-24339 - Quick start links did not display for enabled applications**

For some applications, the **Open quick start** link failed to display on the application's card on the **Enabled** page. As a result, users did not have direct access to the quick start tour for the relevant application.

### **RHODS-2215 - Incorrect Python versions displayed during notebook selection**

The **Start a notebook server** page displayed incorrect versions of Python for the TensorFlow and PyTorch notebook images. Additionally, the third integer of package version numbers is now no longer displayed.

### **RHODS-1977 - Ten minute wait after notebook server start fails**

If the Jupyter leader pod failed while the notebook server was being started, the user could not access their notebook server until the pod restarted, which took approximately ten minutes. This process has been improved so that the user is redirected to their server when a new leader pod is elected. If this process times out, users see a 504 Gateway Timeout error, and can refresh to access their server.

## CHAPTER 4. KNOWN ISSUES

This section describes known issues in Red Hat OpenShift Data Science and any known methods of working around these issues.

### [DATA-SCIENCE-PIPELINES-OPERATOR-294](#) - Scheduled pipeline run that uses data-passing might fail to pass data between steps, or fail the step entirely

A scheduled pipeline run that uses an S3 object store to store the pipeline artifacts might fail with an error such as the following:

```
Bad value for --endpoint-url "cp": scheme is missing. Must be of the form http://<hostname>/ or https://<hostname>/
```

This issue occurs because the S3 object store endpoint is not successfully passed to the pods for the scheduled run.

#### Workaround

You can work around this issue by applying a custom artifact-passing script and then restarting the pipeline server. To do this, perform the following actions:

1. In a text editor, paste the following YAML-based artifact-passing script. The script defines a **ConfigMap** object.

```
apiVersion: v1
data:
  artifact_script: |-
    #!/usr/bin/env sh
    push_artifact() {
      workspace_dir=$(echo ${context.taskRun.name} | sed -e
"s/${context.pipeline.name}-//g")

      workspace_dest=/workspace/${workspace_dir}/artifacts/${context.pipelineRun.name}/${cont
ext.taskRun.name}
      artifact_name=$(basename $2)
      if [ -f "$workspace_dest/$artifact_name" ]; then
        echo sending to: ${workspace_dest}/${artifact_name}
        tar -cvzf $1.tgz -C ${workspace_dest} ${artifact_name}
        aws s3 --endpoint <Endpoint> cp $1.tgz
s3://<Bucket>/artifacts/$PIPELINERUN/$PIPELINETASK/$1.tgz
      elif [ -f "$2" ]; then
        tar -cvzf $1.tgz -C $(dirname $2) ${artifact_name}
        aws s3 --endpoint <Endpoint> cp $1.tgz
s3://<Bucket>/artifacts/$PIPELINERUN/$PIPELINETASK/$1.tgz
      else
        echo "$2 file does not exist. Skip artifact tracking for $1"
      fi
    }
  push_log() {
    cat /var/log/containers/$PODNAME*$NAMESPACE*step-main*.log > step-main.log
    push_artifact main-log step-main.log
  }
  strip_eof() {
    if [ -f "$2" ]; then
      awk 'NF' $2 | head -c -1 > $1_temp_save && cp $1_temp_save $2
```

```

    fi
  }
  kind: ConfigMap
  metadata:
    name: custom-script

```

2. In the script, replace any occurrences of *<Endpoint>* with your S3 endpoint (for example, <https://s3.amazonaws.com>), and occurrences of *<Bucket>* with your S3 bucket name.
3. Save the YAML file for the **ConfigMap** object.
4. Apply the YAML file.

```
$ oc apply -f <configmap-file-name>.yaml
```

5. Restart the pipeline server.

```

$ oc project <data-science-project-name>
$ oc delete pod $(oc get pods -l app=ds-pipeline-pipelines-definition --no-headers | awk
{print $1})

```

When you apply the script and restart the pipeline server, scheduled runs should successfully pass data between steps and into your S3 object store.

### ODH-DASHBOARD-1741 - Cannot create a workbench whose name begins with a number

If you try to create a workbench whose name begins with a number, the workbench does not start.

#### Workaround

Delete the workbench and create a new one with a name that begins with a letter.

### ODH-DASHBOARD-1699 - Workbench does not restart after configuration update

If you update a workbench configuration (for example, by changing an environment variable value), the workbench does not automatically restart and you do not see the changes.

#### Workaround

Manually restart the workbench.

### KUBEFLOW-157 - Logging out of JupyterLab does not work if you are already logged out of the OpenShift Data Science dashboard

If you log out of the OpenShift Data Science dashboard before you log out of JupyterLab, then logging out of JupyterLab is not successful. For example, if you know the URL for a Jupyter notebook, you are able to open this again in your browser.

#### Workaround

Log out of JupyterLab before you log out of the OpenShift Data Science dashboard.

### RHODS-9789 - Pipeline servers fail to start if they contain a custom database that includes a dash in its database name or username field

When you create a pipeline server that uses a custom database, if the value that you set for the **dbname** field or **username** field includes a dash, the pipeline server fails to start.

## Workaround

Edit the pipeline server to omit the dash from the affected fields.

### [RHODS-9764](#) - Data connection details get reset when editing a workbench

When you edit a workbench that has an existing data connection and then select the **Create new data connection** option, the edit page might revert to the **Use existing data connection** option before you have finished specifying the new connection details.

## Workaround

To work around this issue, perform the following actions:

1. Select the **Create new data connection** option again.
2. Specify the new connection details and click **Update workbench** before the page reverts to the **Use existing data connection** option.

### [RHODS-9583](#) - Data Science dashboard does not detect an existing OpenShift Pipelines installation

When the OpenShift Pipelines operator is installed as a global operator on your cluster, the Data Science dashboard does not properly detect it.

An alert icon appears next to the **Data Science Pipelines** option in the left navigation bar. When you open **Data Science Pipelines**, you see the message: "To use pipelines, first install the Red Hat OpenShift Pipelines Operator." However, when you view the list of installed operators in the **openshift-operators** project, you see that OpenShift Pipelines is installed as a global operator on your cluster.

## Workaround

Follow these steps as a user with **cluster-admin** permissions:

1. Log in to your cluster using the **oc** client.
2. Enter the following command to update **OdhDashboardConfig** in the **redhat-ods-applications** application namespace:

```
$ oc patch OdhDashboardConfig odh-dashboard-config -n redhat-ods-applications --type=merge -p '{"spec": {"dashboardConfig": {"disablePipelines": false}}}'
```

### [RHODS-9412](#) - Elyra pipeline fails to run if workbench is created by a user with edit permissions

If a user who has been granted edit permissions for a project creates a project workbench, that user sees the following behavior:

- During the workbench creation process, the user sees an **Error creating workbench** message related to the creation of Kubernetes role bindings.
- Despite the preceding error message, OpenShift Data Science still creates the workbench. However, the error message means that the user will not be able to use the workbench to run Elyra data science pipelines.
- If the user tries to use the workbench to run an Elyra pipeline, Jupyter shows an **Error making request** message that describes failed initialization.

## Workaround

A user with administrator permissions (for example, the project owner) must create the workbench on behalf of the user with edit permissions. That user can then use the workbench to run Elyra pipelines.

### **RHODS-9030 - Uninstall process for OpenShift Data Science might become stuck when removing `kfdefs` resources**

The steps for uninstalling the OpenShift Data Science managed service are described in [Uninstalling OpenShift Data Science](#).

However, even when you follow this guide, you might see that the uninstall process does not finish successfully. Instead, the process stays on the step of deleting `kfdefs` resources that are used by the Kubeflow Operator. As shown in the following example, `kfdefs` resources might exist in the `redhat-ods-applications`, `redhat-ods-monitoring`, and `rhods-notebooks` namespaces:

```
$ oc get kfdefs.kfdef.apps.kubeflow.org -A
```

NAMESPACE	NAME	AGE
redhat-ods-applications	rhods-anaconda	3h6m
redhat-ods-applications	rhods-dashboard	3h6m
redhat-ods-applications	rhods-data-science-pipelines-operator	3h6m
redhat-ods-applications	rhods-model-mesh	3h6m
redhat-ods-applications	rhods-nbc	3h6m
redhat-ods-applications	rhods-osd-config	3h6m
redhat-ods-monitoring	modelmesh-monitoring	3h6m
redhat-ods-monitoring	monitoring	3h6m
rhods-notebooks	rhods-notebooks	3h6m
rhods-notebooks	rhods-osd-config	3h5m

Failed removal of the `kfdefs` resources might also prevent later installation of a newer version of OpenShift Data Science.

#### **Workaround**

To manually delete the `kfdefs` resources so that you can complete the uninstall process, see the "**Force individual object removal when it has finalizers**" section of the following Red Hat solution article: <https://access.redhat.com/solutions/4165791>.

### **RHODS-8939 - For a Jupyter notebook created in a previous release, default shared memory might cause a runtime error**

For a Jupyter notebook created in a release earlier than the current release, the default shared memory for a Jupyter notebook is set to 64 Mb and you cannot change this default value in the notebook configuration.

For example, PyTorch relies on shared memory and the default size of 64 Mb is not enough for large use cases, such as when training a model or when performing heavy data manipulations. Jupyter reports a "no space left on device" message and `/dev/smh` is full.

Starting with release 1.31, this issue is fixed and any new notebook's shared memory is set to the size of the node.

#### **Workaround**

For a Jupyter notebook created in a release earlier than 1.31, either recreate the Jupyter notebook or follow these steps:

1. In your data science project, create a workbench as described in [Creating a project workbench](#).
2. In the data science project page, in the **Workbenches** section, click the **Status** toggle for the workbench to change it from **Running** to **Stopped**.
3. Open your OpenShift Console and then select **Administrator**.
4. Select **Home** → **API Explorer**.
5. In the **Filter by kind** field, type **notebook**.
6. Select the **kubeflow v1** notebook.
7. Select the **Instances** tab and then select the instance for the workbench that you created in Step 1.
8. Click the **YAML** tab and then select **Actions** → **Edit Notebook**.
9. Edit the YAML file to add the following information to the configuration:

- For the container that has the name of your Workbench notebook, add the following lines to the **volumeMounts** section:

```
- mountPath: /dev/shm
  name: shm
```

For example, if your workbench name is **myworkbench**, update the YAML file as follows:

```
spec:
  containers:
    - env
      ...
      name: myworkbench
      ...
      volumeMounts:
        - mountPath: /dev/shm
          name: shm
```

- In the volumes section, add the lines shown in the following example:

```
volumes:
  name: shm
  emptyDir:
    medium: Memory
```

**Note:** Optionally, you can specify a limit to the amount of memory to use for the **emptyDir**.

10. Click **Save**.
11. In the data science dashboard, in the **Workbenches** section of the data science project, click the **Status** toggle for the workbench. The status changes from **Stopped** to **Starting** and then **Running**.
12. Restart the notebook.

**WARNING**

If you later edit the notebook's configuration through the Data Science dashboard UI, your workaround edit to the notebook configuration will be erased.

**RHODS-8921 - You cannot create a pipeline server when cumulative character limit is exceeded**

When the cumulative character limit of a data science project name and a pipeline server name exceeds 62 characters, you are unable to successfully create a pipeline server.

**Workaround**

Rename your data science project so that it does not exceed 30 characters.

**RHODS-8865 - A pipeline server fails to start unless you specify an Amazon Web Services (AWS) Simple Storage Service (S3) bucket resource**

When you create a data connection for a data science project, the `AWS_S3_BUCKET` field is not designated as a mandatory field. However, if you do not specify a value for this field, and you attempt to configure a pipeline server, the pipeline server fails to start successfully.

**RHODS-7718 - User without dashboard permissions is able to continue using their running notebooks and workbenches indefinitely**

When a Red Hat OpenShift Data Science administrator revokes a user's permissions, the user can continue to use their running notebooks and workbenches indefinitely.

**Workaround**

When the OpenShift Data Science administrator revokes a user's permissions, the administrator should also stop any running notebooks and workbenches for that user.

**RHODS-6907 - Attempting to increase the size of a Persistent Volume (PV) fails when it is not connected to a workbench**

Attempting to increase the size of a Persistent Volume (PV) that is not connected to a workbench fails. When changing a data science project's storage, users can still edit the size of the PV in the user interface, but this action does not have any effect.

**RHODS-6950 - Unable to scale down a workbench's GPUs when all GPUs in the cluster are being used**

It is not possible to scale down a workbench's GPUs if all GPUs in the cluster are being used. This issue applies to GPUs being used by one workbench, and GPUs being used by multiple workbenches.

**Workaround**

To work around this issue, perform the following steps:

1. Stop all active workbenches that are using GPUs.
2. Wait until the relevant GPUs are available again.
3. Edit the workbench and scale down the GPU instances.

### **RHODS-6539 - Anaconda Professional Edition cannot be validated and enabled in OpenShift Data Science**

Anaconda Professional Edition cannot be enabled as the dashboard's key validation for Anaconda Professional Edition is inoperable.

### **RHODS-6346 - Unclear error message displays when using invalid characters to create a data science project**

When creating a data science project's data connection, workbench, or storage connection using invalid special characters, the following error message is displayed:

```
the object provided is unrecognized (must be of type Secret): couldn't get version/kind; json parse error: unexpected end of JSON input ({"apiVersion":"v1","kind":"Sec ...)
```

The error message fails to clearly indicate the problem.

### **RHODS-6955 - An error can occur when trying to edit a workbench**

When editing a workbench, an error similar to the following can occur:

```
Error creating workbench
Operation cannot be fulfilled on notebooks.kubeflow.org "workbench-name": the object has been modified; please apply your changes to the latest version and try again
```

### **RHODS-6913 - When editing the configuration settings of a workbench, a misleading error message appears**

When you edit the configuration settings of a workbench, a warning message appears stating the workbench will restart if you make any changes to its configuration settings. This warning is misleading, as if you change the values of its environment variables, the workbench does not automatically restart.

### **RHODS-6383 - An ImagePullBackOff error message is not displayed when required during the workbench creation process**

Pods can experience issues pulling container images from the container registry. If an error occurs, the relevant pod enters into an **ImagePullBackOff** state. During the workbench creation process, if an **ImagePullBackOff** error occurs, an appropriate message is not displayed.

#### **Workaround**

Check the event log for further information on the **ImagePullBackOff** error. To do this, click on the workbench status when it is starting.

### **RHODS-6373 - Workbenches fail to start when cumulative character limit is exceeded**

When the cumulative character limit of a data science project's title and workbench title exceeds 62 characters, workbenches fail to start.

### **RHODS-6356 - The notebook creation process fails for users who have never logged in to the dashboard**

The dashboard's notebook **Administration** page displays users belonging to the user group and admin group in OpenShift. However, if an administrator attempts to start a notebook server on behalf of a user who has never logged in to the dashboard, the server creation process fails and displays the following error message:



Request invalid against a username that does not exist.

### Workaround

Request that the relevant user logs into the dashboard.

### RHODS-6216 - The ModelMesh oauth-proxy container is intermittently unstable

ModelMesh pods do not deploy correctly due to a failure of the ModelMesh **oauth-proxy** container. This issue occurs intermittently and only if authentication is enabled in the ModelMesh runtime environment. It is more likely to occur when additional ModelMesh instances are deployed in different namespaces.

### RHODS-5906 - The NVIDIA GPU Operator is incompatible with OpenShift 4.11.12

Provisioning a GPU node on a OpenShift 4.11.12 cluster results in the **nvidia-driver-daemonset** pod getting stuck in a CrashLoopBackOff state. The NVIDIA GPU Operator is compatible with OpenShift 4.11.9 and 4.11.13.

### RHODS-5763 - Incorrect package version displayed during notebook selection

The **Start a notebook server** page displays an incorrect version number for the Anaconda notebook image.

### RHODS-5543 - When using the NVIDIA GPU Operator, more nodes than needed are created by the Node Autoscaler

When a pod cannot be scheduled due to insufficient available resources, the Node Autoscaler creates a new node. There is a delay until the newly created node receives the relevant GPU workload. Consequently, the pod cannot be scheduled and the Node Autoscaler's continuously creates additional new nodes until one of the nodes is ready to receive the GPU workload. For more information about this issue, see [When using the NVIDIA GPU Operator, more nodes than needed are created by the Node Autoscaler](#).

### Workaround

Apply the **cluster-api/accelerator** label in **machineset.spec.template.spec.metadata**. This causes the autoscaler to consider those nodes as unready until the GPU driver has been deployed.

### RHODS-5251 - Notebook server administration page shows users who have lost permission access

If a user who previously started a notebook server in Jupyter loses their permissions to do so (for example, if an OpenShift Data Science administrator changes the user's group settings or removes the user from a permitted group), administrators continue to see the user's notebook servers on the server **Administration** page. As a consequence, an administrator is able to restart notebook servers that belong to the user who's permissions were revoked.

### RHODS-4769 - GPUs on nodes with unsupported taints cannot be allocated to notebook servers

GPUs on nodes marked with any taint other than the supported *nvidia.com/gpu* taint cannot be selected when creating a notebook server. To avoid this issue, use only the *nvidia.com/gpu* taint on GPU nodes used with OpenShift Data Science.

### RHODS-4799 - Tensorboard requires manual steps to view

When a user has TensorFlow or PyTorch notebook images and wants to use TensorBoard to display data, manual steps are necessary to include environment variables in the notebook environment, and to import those variables for use in your code.

## Workaround

When you start your notebook server, use the following code to set the value for the `TENSORBOARD_PROXY_URL` environment variable to use your OpenShift Data Science user ID.

```
import os
os.environ["TENSORBOARD_PROXY_URL"] = os.environ["NB_PREFIX"] + "/proxy/6006/"
```

## RHODS-4718 - The Intel® oneAPI AI Analytics Toolkits quick start references nonexistent sample notebooks

The Intel® oneAPI AI Analytics Toolkits quick start, located on the **Resources** page on the dashboard, requires the user to load sample notebooks as part of the instruction steps, but refers to notebooks that do not exist in the associated repository.

## RHODS-4627 - The CronJob responsible for validating Anaconda Professional Edition's license is suspended and does not run daily

The CronJob responsible for validating Anaconda Professional Edition's license is automatically suspended by the OpenShift Data Science operator. As a result, the CronJob does not run daily as scheduled. In addition, when Anaconda Professional Edition's license expires, Anaconda Professional Edition is not indicated as disabled on the OpenShift Data Science dashboard.

## RHODS-4502 - The NVIDIA GPU Operator card on the dashboard displays button unnecessarily

GPUs are automatically available in Jupyter after the NVIDIA GPU Operator is installed. The **Enable** button, located on the NVIDIA GPU Operator card on the **Explore** page, is therefore redundant. In addition, clicking the **Enable** button moves the NVIDIA GPU Operator card to the **Enabled** page, even if the Operator is not installed.

## RHODS-3985 - Dashboard does not display \*Enabled page content after ISV operator uninstall

After an ISV operator is uninstalled, no content is displayed on the **Enabled** page on the dashboard. Instead, the following error is displayed:

```
Error loading components
HTTP request failed
```

## Workaround

Wait 30-40 seconds and then refresh the page in your browser.

## RHODS-3984 - Incorrect package versions displayed during notebook selection

In the OpenShift Data Science interface, the **Start a notebook server page** displays incorrect version numbers for the JupyterLab and Notebook packages included in the oneAPI AI Analytics Toolkit notebook image. The page might also show an incorrect value for the Python version used by this image.

## Workaround

When you start your oneAPI AI Analytics Toolkit notebook server, you can check which Python packages are installed on your notebook server and which version of the package you have by running the **!pip list** command in a notebook cell.

## RHODS-2956 - Error can occur when creating a notebook instance

When creating a notebook instance in Jupyter, a **Directory not found** error appears intermittently. This error message can be ignored by clicking **Dismiss**.

### **RHODS-2881 - Actions on dashboard not clearly visible**

The dashboard actions to re-validate a disabled application's license, and to remove a disabled application's card are not clearly visible to the user. These actions only appear when the user clicks on the application card's **Disabled** label. As a result, the intended workflows may not be clear to the user.

### **RHODS-2879 - License re-validation action appears unnecessarily**

The dashboard action to re-validate a disabled application's license appears unnecessarily for applications that do not have a license validation or activation system. In addition, when a user attempts to re-validate a license that cannot be re-validated, feedback is not displayed to state why the action cannot be completed.

### **RHODS-2650 - Error can occur during Pachyderm deployment**

When creating an instance of the Pachyderm operator, a webhook error appears intermittently, preventing the creation process from starting successfully. The webhook error is indicative that, either the Pachyderm operator failed a health check, causing it to restart, or that the operator process exceeded its container's allocated memory limit, triggering an Out of Memory (OOM) kill.

#### **Workaround**

Repeat the Pachyderm instance creation process until the error no longer appears.

### **RHODS-2096 - IBM Watson Studio not available in OpenShift Data Science**

IBM Watson Studio is not available when OpenShift Data Science is installed on OpenShift Dedicated 4.9 or higher, because it is not compatible with these versions of OpenShift Dedicated. Contact [Marketplace support](#) for assistance manually configuring Watson Studio on OpenShift Dedicated 4.9 and higher.

### **RHODS-1888 - OpenShift Data Science hyperlink still visible after uninstall**

When the OpenShift Data Science Add-on is uninstalled from an OpenShift Dedicated cluster, the link to the OpenShift Data Science interface remains visible in the application launcher menu. Clicking this link results in a "Page Not Found" error because OpenShift Data Science is no longer available.

## CHAPTER 5. PRODUCT FEATURES

Red Hat OpenShift Data Science provides a rich set of features for data scientists and IT operations administrators. To learn more, see [Introduction to Red Hat OpenShift Data Science](#) .