Red Hat OpenShift Data Science 1

Installing OpenShift Data Science

Install Red Hat OpenShift Data Science as an Add-on to an OpenShift cluster

Red Hat OpenShift Data Science 1 Installing OpenShift Data Science

Install Red Hat OpenShift Data Science as an Add-on to an OpenShift cluster

Legal Notice

Abstract

Install Red Hat OpenShift Data Science 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).

Table of Contents

CHAPTER 1. ARCHITECTURE OF OPENSHIFT DATA SCIENCE

Red Hat OpenShift Data Science is a fully Red Hat managed cloud service that is available as an Add-on to Red Hat OpenShift Dedicated and to Red Hat OpenShift Service on Amazon Web Services (ROSA).

OpenShift Data Science integrates the following components and services:

- At the service layer:

  ### OpenShift Data Science dashboard

  A customer-facing dashboard that shows available and installed applications for the OpenShift Data Science environment as well as learning resources such as tutorials, quick start examples, and documentation. You can also access administrative functionality from the dashboard, such as user management, cluster settings, and notebook image settings. In addition, data scientists can create their own projects from the dashboard. This enables them to organize their data science work into a single project.

  ### Model serving

  Data scientists can deploy trained machine-learning models to serve intelligent applications in production. After deployment, applications can send requests to the model using its deployed API endpoint.

  ### Data science pipelines

  Data scientists can build portable machine learning (ML) workflows with data science pipelines, using Docker containers. This enables your data scientists to automate workflows as they develop their data science models.

  ### Jupyter (Red Hat managed)

  A Red Hat managed application that allows data scientists to configure their own notebook server environment and develop machine learning models in JupyterLab.

- At the management layer:

  ### The Red Hat OpenShift Data Science operator

  A meta-operator that deploys and maintains all components and sub-operators that are part of OpenShift Data Science.

  ### Monitoring services

  Alertmanager, OpenShift Telemetry, and Prometheus work together to gather metrics from OpenShift Data Science and organize and display those metrics in useful ways for monitoring and billing purposes. Alerts from Alertmanager are sent to PagerDuty, responsible for notifying Red Hat of any issues with your managed cloud service.

When you install the OpenShift Data Science Add-on in the Cluster Manager, the following new projects are created:

- The **redhat-ods-operator** project contains the OpenShift Data Science operator.

- The **redhat-ods-applications** project installs the dashboard and other required components of OpenShift Data Science.

- The **redhat-ods-monitoring** project contains services for monitoring and billing.

- The **rhods-notebooks** project is where notebook environments are deployed by default.

You or your data scientists must create additional projects for the applications that will use your machine learning models.

Do not install independent software vendor (ISV) applications in namespaces associated with OpenShift Data Science Add-ons unless you are specifically directed to do so on the application's card on the dashboard.

CHAPTER 2. OVERVIEW OF INSTALLING AND DEPLOYING OPENSHIFT DATA SCIENCE

Red Hat OpenShift Data Science is a platform for data scientists and developers of artificial intelligence (AI) applications. It provides a fully supported environment that lets you rapidly develop, train, test, and deploy machine learning models on-premises and/or in the public cloud.

OpenShift Data Science is provided as a managed cloud service add-on for Red Hat OpenShift or as self-managed software that you can install on-premise or in the public cloud on OpenShift.

For information on installing OpenShift Data Science as self-managed software, see Installing OpenShift Data Science self-managed.

There are two deployment options for Red Hat OpenShift Data Science as a managed cloud service add-on:

- **OpenShift Dedicated with a Customer Cloud Subscription on Amazon Web Services or Google Cloud Platform**
  OpenShift Dedicated is a complete OpenShift Container Platform cluster provided as a cloud service, configured for high availability, and dedicated to a single customer. OpenShift Dedicated is professionally managed by Red Hat and hosted on Amazon Web Services (AWS) or Google Cloud Platform (GCP). The Customer Cloud Subscription (CCS) model allows Red Hat to deploy and manage clusters into a customer's AWS or GCP account. Contact your Red Hat account manager to get OpenShift Dedicated through a CCS.

- **Red Hat OpenShift Service on AWS (ROSA)**
  ROSA is a fully-managed, turnkey application platform that allows you to focus on delivering value to your customers by building and deploying applications. You subscribe to the service directly from your AWS account.

Installing OpenShift Data Science as a managed cloud service involves the following high-level tasks:

1. Confirm that your OpenShift Dedicated cluster meets all requirements.

2. Configure an identity provider for your OpenShift cluster.

3. Add administrative users for your OpenShift cluster.

4. Subscribe to the Red Hat OpenShift Data Science Add-on.
   For OpenShift Dedicated with a CCS for AWS or GCP, get a subscription through Red Hat.

   For ROSA, get a subscription through the AWS Marketplace.

5. Install the OpenShift Data Science Add-on.

6. Access the OpenShift Data Science dashboard.

7. Optionally, enable graphics processing units (GPUs) in OpenShift Data Science to ensure that your data scientists can use compute-heavy workloads in their models.

CHAPTER 3. REQUIREMENTS FOR OPENSHIFT DATA SCIENCE

You must meet the following requirements before you can install OpenShift Data Science on your Red Hat OpenShift Dedicated or Red Hat OpenShift Service on Amazon Web Services (ROSA) cluster.

- **A subscription for Red Hat OpenShift Dedicated**or **a subscription for ROSA**
  You can deploy Red Hat OpenShift Dedicated on your Amazon Web Services (AWS) or Google Cloud Platform (GCP) account by using the Customer Cloud Subscription on AWS or Customer Cloud Subscription on GCP model. Note that while Red Hat provides an option to install OpenShift Dedicated on a Red Hat cloud account, if you want to install OpenShift Data Science then you must install OpenShift Dedicated on your own cloud account.

  Contact your Red Hat account manager to purchase a new Red Hat OpenShift Dedicated subscription. If you do not yet have an account manager, complete the form at https://cloud.redhat.com/products/dedicated/contact/ to request one.

  You can subscribe to Red Hat OpenShift Service on AWS (ROSA) directly from your AWS account or by contacting your Red Hat account manager.

- **A Red Hat customer account**
  Go to OpenShift Cluster Manager (http://console.redhat.com/openshift) and log in or register for a new account.

- **Cluster administrator access to your OpenShift cluster**
  Use an existing cluster or create a new cluster by following the steps in the relevant documentation:

  - Creating an OpenShift Dedicated cluster

  - Creating a ROSA cluster with STS

- **An OpenShift Dedicated or ROSA cluster configuration that meets the following configuration requirements.**
  At least 2 worker nodes with at least 8 CPUs and 32 GiB RAM available for OpenShift Data Science to use when you install the Add-on. If this requirement is not met, the installation process fails to start and an error is displayed.

  When you create a new cluster, select **m6a.2xlarge** for the computer node instance type to satisfy the requirements.

  For an existing ROSA cluster, you can get the compute node instance type by using this command:

  rosa list machinepools --cluster=cluster-name

  You cannot alter a cluster's compute node instance type, but you can add an additional machine pool or modify the default pool to meet the minimum requirements. However, the minimum resource requirements must be met by a single machine pool in the cluster.

  For more information, see the relevant documentation:

  - Creating a machine pool in OpenShift Dedicated

  - OpenShift Data Science Service Definition

- Creating a machine pool in ROSA

- Prepare your environment (ROSA)

- **For a ROSA cluster, select an access management strategy**
  For installing OpenShift Data Science on a ROSA cluster, decide whether you want to install on a ROSA cluster that uses AWS Security Token Service (STS) or one that uses AWS Identity and Access Management (IAM) credentials. See Install ROSA Classic clusters for advice on deploying a ROSA cluster with or without AWS STS.

- **Install the Red Hat OpenShift Pipelines Operator**
  OpenShift Data Science supports data science pipelines. A pipeline is a collection of Task resources that are arranged in a specific order of execution. By using Red Hat OpenShift Data Science pipelines, you can standardize and automate machine learning workflows to automate the build and deployment of your data science models. Before you can use pipelines with OpenShift Data Science, install the Red Hat OpenShift Pipelines Operator as described in Installing OpenShift Pipelines.

CHAPTER 4. CONFIGURING AN IDENTITY PROVIDER FOR YOUR OPENSHIFT CLUSTER

Configure an identity provider for your OpenShift Dedicated or Red Hat OpenShift Service on Amazon Web Services (ROSA) cluster to manage users and groups.

Red Hat OpenShift Data Science supports the same authentication systems as Red Hat OpenShift Dedicated and ROSA. Check the appropriate documentation for your cluster for more information.

- Supported identity providers on OpenShift Dedicated

- Supported identity providers on ROSA
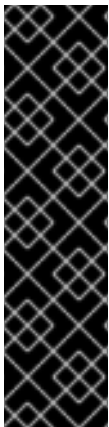
> **IMPORTANT**
>
> Adding more than one OpenShift Identity Provider can create problems when the same user name exists in multiple providers.
>
> When **mappingMethod** is set to **claim** (the default mapping method for identity providers) and multiple providers have credentials associated with the same user name, the first provider used to log in to OpenShift is the one that works for that user, regardless of the order in which identity providers are configured.
>
> Refer to Identity provider parameters in the OpenShift Dedicated documentation for more information about mapping methods.

Prerequisites

- Credentials for OpenShift Cluster Manager (https://console.redhat.com/openshift/).

- An existing OpenShift Dedicated cluster.

Procedure

1. Log in to OpenShift Cluster Manager (https://console.redhat.com/openshift/).

2. Click **Clusters**. The **Clusters** page opens.

3. Click the name of the cluster to configure.

4. Click the **Access control** tab.

5. Click **Identity providers**.

6. Click **Add identity provider**.

   a. Select your provider from the **Identity Provider** list.

   b. Complete the remaining fields relevant to the identity provider that you selected. See Configuring identity providers for more information.

7. Click **Confirm**.

Verification

- The configured identity providers are visible on the **Access control** tab of the **Cluster details** page.

Additional resources

- [Configuring identity providers](#)

- [Syncing LDAP groups](#)

CHAPTER 5. ADDING ADMINISTRATIVE USERS

Before you can install and configure OpenShift Data Science for your data scientist users, you must define administrative users. Only administrative users can install and configure OpenShift Data Science.

Prerequisites

- Credentials for OpenShift Cluster Manager (https://console.redhat.com/openshift/).

- An existing OpenShift Dedicated cluster with an identity provider configured.

Procedure

1. Log in to OpenShift Cluster Manager (https://console.redhat.com/openshift/).

2. Click **Clusters**. The **Clusters** page opens.

3. Click the name of the cluster to configure.

4. Click the **Access control** tab.

5. Click **Cluster Roles and Access**.

6. Under **Cluster administrative users** click the **Add user** button.
   The **Add cluster user** popover appears.

7. Enter the user name in the **User ID** field.

8. Select an appropriate **Group** for the user.
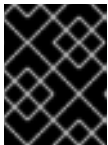
   > **IMPORTANT**
   >
   > If this user needs to use existing groups in an identity provider to control OpenShift Data Science access, select **cluster-admins**.

   Check Cluster administration in the OpenShift Dedicated documentation for more information about these user types.

9. Click **Add user**.

Verification

- The user name and selected group are visible in the list of **Cluster administrative users**.

Additional resources

- OpenShift Dedicated Cluster administration

CHAPTER 6. SUBSCRIBING TO THE RED HAT OPENSHIFT DATA SCIENCE MANAGED CLOUD SERVICE

You can subscribe to the Red Hat OpenShift Data Science managed cloud service in the following ways:

- Subscribe through Red Hat if you have a Red Hat OpenShift Dedicated cluster deployed with a Customer Cloud Subscription (CCS) on Amazon Web Services (AWS) or Google Cloud Platform (GCP).

- Subscribe through the AWS Marketplace if you have a Red Hat OpenShift Service on AWS (ROSA) cluster.

> **NOTE**
>
> You can also purchase Red Hat OpenShift Data Science as self-managed software. To purchase a new subscription, contact your Red Hat account manager. If you do not yet have an account manager, complete the form at https://www.redhat.com/en/contact to request one.

6.1. SUBSCRIBING TO THE OPENSHIFT DATA SCIENCE MANAGED CLOUD SERVICE ON AWS OR GCP

For a Red Hat OpenShift Dedicated cluster that is deployed on AWS or GCP, contact your Red Hat account manager to purchase a new subscription. If you do not yet have an account manager, complete the form at https://cloud.redhat.com/products/dedicated/contact/ to request one.

Prerequisite

- You have worked with Red Hat Sales to enable a private offer of OpenShift Data Science, follow these steps to accept your offer and deploy the solution.

Procedure

1. Visit your Private Offer with the URL link provided by your Red Hat Sales representative.

2. Click **Accept Terms** to subscribe to the AMI Private Offer named **OpenShift Data Science from AWS Marketplace**.

3. After accepting the offer terms, click **Continue to Configuration**.

6.2. SUBSCRIBING TO THE OPENSHIFT DATA SCIENCE MANAGED CLOUD SERVICE ON RED HAT OPENSHIFT SERVICE ON AWS (ROSA)

For a ROSA cluster, you can subscribe to the OpenShift Data Science managed cloud service through the Amazon Web Services (AWS) Marketplace.

Prerequisites

- Access to a ROSA cluster, including permissions to view and install add-ons.

- An AWS account with permission to view and subscribe to offerings in the AWS marketplace.

Procedure

1. In the AWS Console, navigate to the AWS Marketplace. For example:

   a. Click the help icon and then select Getting Started Resource Center.

   b. Select AWS Marketplace > Browse AWS Marketplace.

2. In the top **Search** field, type: **Red Hat OpenShift Data Science**.

3. Select one of the two options depending on the geographical location of the billing address for your AWS account (note that this location might differ from the geographical location of the cluster):

   - Europe, the Middle East, and Africa (EMEA region)

   - North America and regions outside EMEA

4. Click **Continue to Subscribe**.

5. Click **Continue to Configuration** and then select the appropriate fulfillment options. Note that some selectors might have only one option.

6. Click **Continue to Launch**.

7. Link your AWS account with your Red Hat account to complete your registration:

   a. In the AWS Marketplace console, navigate to the **Manage Subscriptions** page.

   b. In the **Red Hat OpenShift Data Science** card, click **Set up product**.

   c. In the top banner, click **Set up account**.
      This link takes you to the Red Hat Hybrid console.

   d. If you are not already logged in, log in.

   e. Review and then accept the terms and agreements.

   f. Click **Connect accounts**.

Verification

The Data Science product page opens.

# CHAPTER 7. INSTALLING OPENSHIFT DATA SCIENCE ON YOUR OPENSHIFT CLUSTER

You can use Red Hat OpenShift Cluster Manager to install Red Hat OpenShift Data Science as an Add-on to your Red Hat OpenShift cluster.

Prerequisites

- A subscription to the Red Hat OpenShift Data Science Add-on, as described in Subscribing to the Red Hat OpenShift Data Science Add-on.

- If you purchased the Red Hat OpenShift Data Science Add-on for ROSA by using the AWS Marketplace, you have associated your AWS account with your Red Hat account as described in Subscribing to the Red Hat OpenShift Data Science Add-on through the AWS Marketplace (ROSA only).

- Credentials for OpenShift Cluster Manager (https://console.redhat.com/openshift/).

- Administrator access to the OpenShift cluster.

Procedure

1. Log in to OpenShift Cluster Manager (https://console.redhat.com/openshift/).

2. Click **Clusters**.
   The **Clusters** page opens.

3. Click the name of the cluster you want to install OpenShift Data Science on.
   The **Details** page for the cluster opens.

4. Click the **Add-ons** tab and locate the **Red Hat OpenShift Data Science** card.
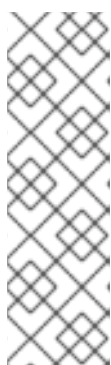
   > **NOTE**
   >
   > If there is a **Prerequisites not met** warning message, click the **Prerequisites** tab. Note down the error message. If the error message states that you require a new machine pool, or that more resources are required, take the appropriate action to resolve the problem. You might need to add more resources to your cluster, or increase the size of your default machine pool. To increase your cluster's resources, contact your infrastructure administrator. For more information about increasing the size of your machine pool, see Allocating additional resources to OpenShift Data Science users.

5. Select a **Subscription type**:
   If you obtained your RHODS subscription through your Red Hat account manager, select **Standard** and then skip to Step 7.

   If you obtained your RHODS subscription directly from the AWS Marketplace, select **Marketplace** and then continue to Step 6.

6. For a Marketplace subscription, select your AWS account number from the list.

NOTE

If your AWS account number is not in the list, you might need to link your Red Hat and AWS accounts, as described in Subscribing to the Red Hat OpenShift Data Science Add-on through the AWS Marketplace (ROSA only).

7. Click **Install**. The **Configure Red Hat OpenShift Data Science** pane appears.

8. In the **Notification email** field, enter any email addresses that you want to receive important alerts about the state of Red Hat OpenShift Data Science, such as outage alerts.

9. Click **Install**.

Verification

- In OpenShift Cluster Manager, under the **Add-ons** tab for the cluster, confirm that the OpenShift Data Science card shows one of the following states:

  - **Installing** - installation is in progress; wait for this to change to **Installed**. This takes around 30 minutes.

  - **Installed** - installation is complete; verify that the **View in console** button is visible.

- In OpenShift Dedicated, click **Home → Projects** and confirm that the following project namespaces are visible and listed as **Active**:

  - **redhat-ods-applications**

  - **redhat-ods-monitoring**

  - **redhat-ods-operator**

  - **rhods-notebooks**

CHAPTER 8. ACCESSING THE OPENSHIFT DATA SCIENCE DASHBOARD

After you have installed OpenShift Data Science and added users, you can access the URL for your OpenShift Data Science console and share the URL with the users to let them log in and work on their data models.

Prerequisites

- You have installed OpenShift Data Science on your OpenShift Dedicated or Red Hat OpenShift Service on Amazon Web Services (ROSA) cluster.

- You have added at least one user to the user group for OpenShift Data Science as described in Adding users for OpenShift Data Science .

Procedure

1. Log in to OpenShift web console.

2. Click the application launcher ( ).

3. Right-click on **Red Hat OpenShift Data Science** and copy the URL for your OpenShift Data Science instance.

4. Provide this instance URL to your data scientists to let them log in to OpenShift Data Science.

Verification

- Confirm that you and your users can log in to OpenShift Data Science by using the instance URL.

Additional resources

- Logging in to OpenShift Data Science

- Adding users for OpenShift Data Science

CHAPTER 9. ENABLING GPU SUPPORT IN OPENSHIFT DATA SCIENCE

Optionally, to ensure that your data scientists can use compute-heavy workloads in their models, you can enable graphics processing units (GPUs) in OpenShift Data Science. To enable GPUs on OpenShift, you must install the NVIDIA GPU Operator. As a prerequisite to installing the NVIDIA GPU Operator, you must install the Node Feature Discovery Operator. For information about how to install these operators, see GPU Operator on OpenShift.

> **IMPORTANT**
>
> The NVIDIA GPU Add-on is no longer supported. Instead, enable GPUs by installing the NVIDIA GPU Operator. If your deployment has a previously-installed NVIDIA GPU Add-on, before you install the NVIDIA GPU Operator, use OpenShift Cluster Manager to uninstall the NVIDIA GPU Add-on from your cluster.

CHAPTER 10. TROUBLESHOOTING COMMON INSTALLATION PROBLEMS

If you are experiencing difficulties installing the Red Hat OpenShift Data Science Add-on, read this section to understand what could be causing the problem, and how to resolve the problem.

If you cannot see the problem here or in the release notes, contact Red Hat Support.

10.1. THE OPENSHIFT DATA SCIENCE OPERATOR CANNOT BE RETRIEVED FROM THE IMAGE REGISTRY

Problem

When attempting to retrieve the OpenShift Data Science operator from the image registry, an **Failure to pull from quay** error message appears. The OpenShift Data Science operator might be unavailable for retrieval in the following circumstances:

- The image registry is unavailable.

- There is a problem with your network connection.

- Your cluster is not operational and is therefore unable to retrieve the image registry.

Diagnosis

Check the logs in the **Events** section in OpenShift Dedicated for further information about the **Failure to pull from quay** error message.

Resolution

- To resolve this issue, contact Red Hat support.

10.2. OPENSHIFT DATA SCIENCE CANNOT BE INSTALLED DUE TO INSUFFICIENT CLUSTER RESOURCES

Problem

When attempting to install OpenShift Data Science, an error message appears stating that installation prerequisites have not been met.

Diagnosis

1. Log in to OpenShift Cluster Manager (https://console.redhat.com/openshift/).

2. Click **Clusters**.
   The **Clusters** page opens.

3. Click the name of the cluster you want to install OpenShift Data Science on.
   The **Details** page for the cluster opens.

4. Click the **Add-ons** tab and locate the **Red Hat OpenShift Data Science** card.

5. Click **Install**. The **Configure Red Hat OpenShift Data Science** pane appears.

6. If the installation fails, click the **Prerequisites** tab.

7. Note down the error message. If the error message states that you require a new machine pool, or that more resources are required, take the appropriate action to resolve the problem.

Resolution

- You might need to add more resources to your cluster, or increase the size of your machine pool. To increase your cluster's resources, contact your infrastructure administrator. For more information about increasing the size of your machine pool, see Nodes and Allocating additional resources to OpenShift Data Science users.

10.3. THE DEDICATED-ADMINS ROLE-BASED ACCESS CONTROL (RBAC) POLICY CANNOT BE CREATED

Problem

The Role-based access control (RBAC) policy for the dedicated-admins group in the target project cannot be created. This issue occurs in unknown circumstances.

Diagnosis

1. In the OpenShift Dedicated web console, change into the **Administrator** perspective.

2. Click **Workloads → Pods**.

3. Set the **Project** to **All Projects** or **redhat-ods-operator**.

4. Click the **rhods-operator-<random string>** pod.
   The **Pod details** page appears.

5. Click **Logs**.

6. Select **rhods-deployer** from the drop-down list

7. Check the log for the **ERROR: Attempt to create the RBAC policy for dedicated admins group in $target_project failed.** error message.

Resolution

- Contact Red Hat support.

10.4. OPENSHIFT DATA SCIENCE DOES NOT INSTALL ON UNSUPPORTED INFRASTRUCTURE

Problem

Customer deploying on an environment not documented as being supported by the RHODS operator.

Diagnosis

1. In the OpenShift Dedicated web console, change into the **Administrator** perspective.

2. Click **Workloads → Pods**.

3. Set the **Project** to **All Projects** or **redhat-ods-operator**.

4. Click the **rhods-operator-<random string>** pod.
   The **Pod details** page appears.

5. Click **Logs**.

6. Select **rhods-deployer** from the drop-down list

7. Check the log for the **ERROR: Deploying on $infrastructure, which is not supported. Failing Installation** error message.

Resolution

Before proceeding with a new installation, ensure that you have a fully supported environment on which to install OpenShift Data Science. For more information, see Requirements for OpenShift Data Science .

10.5. THE CREATION OF THE OPENSHIFT DATA SCIENCE CUSTOM RESOURCE (CR) FAILS

Problem

During the installation process, the OpenShift Data Science Custom Resource (CR) does not get created. This issue occurs in unknown circumstances.

Diagnosis

1. In the OpenShift Dedicated web console, change into the **Administrator** perspective.

2. Click **Workloads → Pods**.

3. Set the **Project** to **All Projects** or **redhat-ods-operator**.

4. Click the **rhods-operator-<random string>** pod.
   The **Pod details** page appears.

5. Click **Logs**.

6. Select **rhods-deployer** from the drop-down list

7. Check the log for the **ERROR: Attempt to create the ODH CR failed.** error message.

Resolution

Contact Red Hat support.

10.6. THE CREATION OF THE OPENSHIFT DATA SCIENCE NOTEBOOKS CUSTOM RESOURCE (CR) FAILS

Problem

During the installation process, the OpenShift Data Science Notebooks Custom Resource (CR) does not get created. This issue occurs in unknown circumstances.

Diagnosis

1. In the OpenShift Dedicated web console, change into the **Administrator** perspective.

2. Click **Workloads → Pods**.

3. Set the **Project** to **All Projects** or **redhat-ods-operator**.

4. Click the **rhods-operator-<random string>** pod.
   The **Pod details** page appears.

5. Click **Logs**.

6. Select **rhods-deployer** from the drop-down list

7. Check the log for the **ERROR: Attempt to create the RHODS Notebooks CR failed.** error message.

Resolution

Contact Red Hat support.

10.7. THE DEAD MAN'S SNITCH OPERATOR'S SECRET DOES NOT GET CREATED

Problem

An issue with Managed Tenants SRE automation process causes the Dead Man's Snitch operator's secret to not get created.

Diagnosis

1. In the OpenShift Dedicated web console, change into the **Administrator** perspective.

2. Click **Workloads → Pods**.

3. Set the **Project** to **All Projects** or **redhat-ods-operator**.

4. Click the **rhods-operator-<random string>** pod.
   The **Pod details** page appears.

5. Click **Logs**.

6. Select **rhods-deployer** from the drop-down list

7. Check the log for the **ERROR: Dead Man Snitch secret does not exist.** error message.

Resolution

Contact Red Hat support.

10.8. THE PAGERDUTY SECRET DOES NOT GET CREATED

Problem

An issue with Managed Tenants SRE automation process causes the PagerDuty's secret to not get created.

Diagnosis

Diagnosis

1. In the OpenShift Dedicated web console, change into the **Administrator** perspective.

2. Click **Workloads → Pods**.

3. Set the **Project** to **All Projects** or **redhat-ods-operator**.

4. Click the **rhods-operator-<random string>** pod.
   The **Pod details** page appears.

5. Click **Logs**.

6. Select **rhods-deployer** from the drop-down list

7. Check the log for the **ERROR: Pagerduty secret does not exist** error message.

Resolution

Contact Red Hat support.

10.9. THE SMTP SECRET DOES NOT EXIST

Problem

An issue with Managed Tenants SRE automation process causes the SMTP secret to not get created.

Diagnosis

1. In the OpenShift Dedicated web console, change into the **Administrator** perspective.

2. Click **Workloads → Pods**.

3. Set the **Project** to **All Projects** or **redhat-ods-operator**.

4. Click the **rhods-operator-<random string>** pod.
   The **Pod details** page appears.

5. Click **Logs**.

6. Select **rhods-deployer** from the drop-down list

7. Check the log for the **ERROR: SMTP secret does not exist** error message.

Resolution

Contact Red Hat support.

10.10. THE ODH PARAMETER SECRET DOES NOT GET CREATED

Problem

An issue with the OpenShift Data Science add-on's flow could result in the ODH parameter secret to not get created.

Diagnosis

1. In the OpenShift Dedicated web console, change into the **Administrator** perspective.

2. Click **Workloads → Pods**.

3. Set the **Project** to **All Projects** or **redhat-ods-operator**.

4. Click the **rhods-operator-<random string>** pod.
   The **Pod details** page appears.

5. Click **Logs**.

6. Select **rhods-deployer** from the drop-down list

7. Check the log for the **ERROR: Addon managed odh parameter secret does not exist.** error message.

Resolution

Contact Red Hat support.