



Red Hat Ansible Automation Platform 2.4

Red Hat Ansible Automation Platform operator backup and recovery guide

Safeguard against data loss with backup and recovery of Ansible Automation Platform operator on OpenShift Container Platform

Red Hat Ansible Automation Platform 2.4 Red Hat Ansible Automation Platform operator backup and recovery guide

Safeguard against data loss with backup and recovery of Ansible Automation Platform operator on OpenShift Container Platform

Legal Notice

Copyright © 2024 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[®] is the registered trademark of Linus Torvalds in the United States and other countries.

Java[®] is a registered trademark of Oracle and/or its affiliates.

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

MySQL[®] is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js[®] 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[®] 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

This guide provides procedures and reference information to backup and recover installations of the Red Hat Ansible Automation Platform operator on OpenShift Container Platform.

Table of Contents

PREFACE	3
PROVIDING FEEDBACK ON RED HAT DOCUMENTATION	4
CHAPTER 1. BACKUP AND RECOVERY OF RED HAT ANSIBLE AUTOMATION PLATFORM	5
1.1. ABOUT BACKUP AND RECOVERY	5
1.1.1. Backup recommendations	5
CHAPTER 2. CREATING RED HAT ANSIBLE AUTOMATION PLATFORM BACKUP RESOURCES	6
2.1. BACKING UP THE AUTOMATION CONTROLLER DEPLOYMENT	6
2.2. USING YAML TO BACK UP THE AUTOMATION CONTROLLER DEPLOYMENT	7
2.3. BACKING UP THE AUTOMATION HUB DEPLOYMENT	8
CHAPTER 3. RECOVERING A RED HAT ANSIBLE AUTOMATION PLATFORM DEPLOYMENT	9
3.1. RECOVERING THE AUTOMATION CONTROLLER DEPLOYMENT	9
3.2. USING YAML TO RECOVER THE AUTOMATION CONTROLLER DEPLOYMENT	10
3.3. RECOVERING THE AUTOMATION HUB DEPLOYMENT	11
CHAPTER 4. TROUBLESHOOTING	13
4.1. AUTOMATION CONTROLLER CUSTOM RESOURCE HAS THE SAME NAME AS AN EXISTING DEPLOYMENT	13

PREFACE

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

Use the procedures in this guide to create backup resources that can be used for recovering your Red Hat Ansible Automation Platform deployment in the event of a failure.

PROVIDING FEEDBACK ON RED HAT DOCUMENTATION

If you have a suggestion to improve this documentation, or find an error, you can contact technical support at <https://access.redhat.com> to open a request.

CHAPTER 1. BACKUP AND RECOVERY OF RED HAT ANSIBLE AUTOMATION PLATFORM


To safeguard against unexpected data loss and application errors, it is critical that you perform periodic backups of your Red Hat Ansible Automation Platform deployment. In addition to data loss prevention, backups allow you to fall back to a different deployment state.

1.1. ABOUT BACKUP AND RECOVERY

Red Hat recommends backing up deployments of Red Hat Ansible Automation Platform in your Red Hat OpenShift Container Platform environment to prevent data loss.

A backup resource of your Red Hat Ansible Automation Platform deployment includes the following:

- Custom deployment of specific values in the **spec** section of the Ansible Automation Platform custom resource object
- Back up of the **postgresql** database
- **secret_key**, **admin_password**, and **broadcast_websocket** secrets
- Database configuration



NOTE

Be sure to secure your backup resources because they can include sensitive information.

1.1.1. Backup recommendations

Recovering from data loss requires that you plan for and create backup resources of your Red Hat Ansible Automation Platform deployments on a regular basis. At a minimum, Red Hat recommends backing up deployments of Red Hat Ansible Automation Platform under the following circumstances:

- Before upgrading your Red Hat Ansible Automation Platform deployments
- Before upgrading your Openshift cluster
- Once per week. This is particularly important if your environment is configured for automatic upgrades.

CHAPTER 2. CREATING RED HAT ANSIBLE AUTOMATION PLATFORM BACKUP RESOURCES

Backing up your Red Hat Ansible Automation Platform deployment involves creating backup resources for your deployed automation hub and automation controller instances. Use these procedures to create backup resources for your Red Hat Ansible Automation Platform deployment.

2.1. BACKING UP THE AUTOMATION CONTROLLER DEPLOYMENT

Use this procedure to back up a deployment of the controller, including jobs, inventories, and credentials.

Prerequisites

- You must be authenticated with an OpenShift cluster.
- The Ansible Automation Platform Operator has been installed to the cluster.
- The automation controller is deployed to using the Ansible Automation Platform Operator.

Procedure

1. Log in to **Red Hat OpenShift Container Platform**
2. Navigate to **Operators → Installed Operators**.
3. Select the Ansible Automation Platform Operator installed on your project namespace.
4. Select the **Automation Controller Backup** tab.
5. Click **Create AutomationControllerBackup**.
6. Enter a **Name** for the backup.
7. Enter the **Deployment name** of the deployed Ansible Automation Platform instance being backed up. For example, if your automation controller must be backed up and the deployment name is **aap-controller**, enter 'aap-controller' in the **Deployment name** field.
8. If you want to use a custom, pre-created pvc:
 - a. Optionally enter the name of the **Backup persistent volume claim**
 - b. Optionally enter the **Backup PVC storage requirements** and **Backup PVC storage class**



NOTE

If no pvc or storage class is provided, the cluster's default storage class is used to create the pvc.

- c. If you have a large database, specify your storage requests accordingly under **Backup management pod resource requirements**.

**NOTE**

You can check the size of the existing postgres database data directory by running the following command inside the postgres pod.


```
$ df -h | grep "/var/lib/pgsql/data"
```

9. Click **Create**.

A backup tarball of the specified deployment is created and available for data recovery or deployment rollback. Future backups are stored in separate tar files on the same pvc.

Verification

1. Log in to Red Hat **Red Hat OpenShift Container Platform**
2. Navigate to **Operators → Installed Operators**.
3. Select the Ansible Automation Platform Operator installed on your project namespace.
4. Select the **AutomationControllerBackup** tab.
5. Select the backup resource you want to verify.
6. Scroll to **Conditions** and check that the **Successful** status is **True**.

**NOTE**

If **Successful** is **False**, the backup has failed. Check the automation controller operator logs for the error to fix the issue.

2.2. USING YAML TO BACK UP THE AUTOMATION CONTROLLER DEPLOYMENT

See the following procedure for how to back up a deployment of the automation controller using YAML.

Prerequisites

- You must be authenticated with an OpenShift cluster.
- The Ansible Automation Platform Operator has been installed to the cluster.
- The automation controller is deployed to using the Ansible Automation Platform Operator.

Procedure

1. Create a file named "backup-awx.yml" with the following contents:

```
---
apiVersion: automationcontroller.ansible.com/v1beta1
kind: AWXBackup
metadata:
  name: awxbackup-2024-07-15
```

```
namespace: my-namespace
spec:
  deployment_name: controller
```



NOTE

The "deployment_name" above is the name of the automation controller deployment you intend to backup from. The namespace above is the one containing the automation controller deployment you intend to back up.

2. Use the **oc apply** command to create the backup object in your cluster:

```
$ oc apply -f backup-awx.yml
```

2.3. BACKING UP THE AUTOMATION HUB DEPLOYMENT

Use this procedure to back up a deployment of the hub, including all hosted Ansible content.

Prerequisites

- You must be authenticated with an OpenShift cluster.
- The Ansible Automation Platform Operator has been installed to the cluster.
- The automation hub is deployed to using the Ansible Automation Platform Operator.

Procedure

1. Log in to **Red Hat OpenShift Container Platform**
2. Navigate to **Operators → Installed Operators**.
3. Select the Ansible Automation Platform Operator installed on your project namespace.
4. Select the **Automation Hub Backup** tab.
5. Click **Create AutomationHubBackup**.
6. Enter a **Name** for the backup.
7. Enter the **Deployment name** of the deployed Ansible Automation Platform instance being backed up. For example, if your automation hub must be backed up and the deployment name is **aap-hub**, enter 'aap-hub' in the **Deployment name** field.
8. If you want to use a custom, pre-created pvc:
 - a. Optionally, enter the name of the **Backup persistent volume claim**, **Backup persistent volume claim namespace**, **Backup PVC storage requirements**, and **Backup PVC storage class**.
9. Click **Create**.
A backup of the specified deployment is created and available for data recovery or deployment rollback.

CHAPTER 3. RECOVERING A RED HAT ANSIBLE AUTOMATION PLATFORM DEPLOYMENT

If you lose information on your system or issues with an upgrade, you can use the backup resources of your deployment instances. Use these procedures to recover your automation controller and automation hub deployment files.

3.1. RECOVERING THE AUTOMATION CONTROLLER DEPLOYMENT

Use this procedure to restore a previous controller deployment from an `AutomationControllerBackup`. The deployment name you provide will be the name of the new `AutomationController` custom resource that will be created.



NOTE

The name specified for the new `AutomationController` custom resource must not match an existing deployment.

If the backup custom resource being restored is a backup of a currently running `AutomationController` custom resource the recovery process will fail. See [Troubleshooting](#) for steps to resolve the issue.

Prerequisites

- You must be authenticated with an OpenShift cluster.
- The automation controller has been deployed to the cluster.
- An `AutomationControllerBackup` is available on a PVC in your cluster.

Procedure

1. Log in to **Red Hat OpenShift Container Platform**
2. Navigate to **Operators → Installed Operators**.
3. Select the Ansible Automation Platform Operator installed on your project namespace.
4. Select the **Automation Controller Restore** tab.
5. Click **Create AutomationControllerRestore**.
6. Enter a **Name** for the recovery deployment.
7. Enter a **New Deployment name** for the restored deployment.



NOTE

This should be different from the original deployment name.


8. Select the **Backup source to restore from Backup CR** is recommended.
9. Enter the **Backup Name** of the `AutomationControllerBackup` object.

10. Click **Create**.

A new deployment is created and your backup is restored to it. This can take approximately 5 to 15 minutes depending on the size of your database.

Verification

1. Log in to Red Hat **Red Hat OpenShift Container Platform**
2. Navigate to **Operators → Installed Operators**.
3. Select the Ansible Automation Platform Operator installed on your project namespace.
4. Select the **AutomationControllerRestore** tab.
5. Select the restore resource you want to verify.
6. Scroll to **Conditions** and check that the **Successful** status is **True**.



NOTE

If **Successful** is **False**, the recovery has failed. Check the automation controller operator logs for the error to fix the issue.

3.2. USING YAML TO RECOVER THE AUTOMATION CONTROLLER DEPLOYMENT

See the following procedure for how to restore a deployment of the automation controller using YAML.

Prerequisite

The external database must be a PostgreSQL database that is the version supported by the current release of Ansible Automation Platform.



NOTE

Ansible Automation Platform 2.4 supports PostgreSQL 13.

Procedure

The external postgres instance credentials and connection information must be stored in a secret, which is then set on the automation controller spec.

1. Create a **external-postgres-configuration-secret** YAML file, following the template below:

```
apiVersion: v1
kind: Secret
metadata:
  name: external-restore-postgres-configuration
  namespace: <target_namespace> 1
stringData:
  host: "<external_ip_or_url_resolvable_by_the_cluster>" 2
  port: "<external_port>" 3
  database: "<desired_database_name>"
  username: "<username_to_connect_as>"
```

```
password: "<password_to_connect_with>" 4
sslmode: "prefer" 5
type: "unmanaged"
type: Opaque
```

- 1 Namespace to create the secret in. This should be the same namespace you wish to deploy to.
 - 2 The resolvable hostname for your database node.
 - 3 External port defaults to **5432**.
 - 4 Value for variable **password** should not contain single or double quotes (' ') or backslashes (\) to avoid any issues during deployment, backup or restoration.
 - 5 The variable **sslmode** is valid for **external** databases only. The allowed values are: **prefer**, **disable**, **allow**, **require**, **verify-ca**, and **verify-full**.
2. Apply **external-postgres-configuration-secret.yml** to your cluster using the **oc create** command.

```
$ oc create -f external-postgres-configuration-secret.yml
```

3. When creating your **AutomationControllerRestore** custom resource object, specify the secret on your spec, following the example below:

```
kind: AutomationControllerRestore
apiVersion: automationcontroller.ansible.com/v1beta1
metadata:
  namespace: my-namespace
  name: awxrestore-2024-07-15
spec:
  deployment_name: restored_controller
  backup_name: awxbackup-2024-07-15
  postgres_configuration_secret: 'external-restore-postgres-configuration'
```

3.3. RECOVERING THE AUTOMATION HUB DEPLOYMENT

Use this procedure to restore a previous hub deployment into the namespace. The deployment name you provide will be the name of the new AutomationHub custom resource that will be created.



NOTE

The name specified for the new AutomationHub custom resource must not match an existing deployment or the recovery process will fail.

Prerequisites

- You must be authenticated with an Openshift cluster.
- The automation hub has been deployed to the cluster.
- An AutomationHubBackup is available on a PVC in your cluster.

Procedure

1. Log in to **Red Hat OpenShift Container Platform**
2. Navigate to **Operators → Installed Operators**.
3. Select the Ansible Automation Platform Operator installed on your project namespace.
4. Select the **Automation Hub Restore** tab.
5. Click **Create AutomationHubRestore**.
6. Enter a **Name** for the recovery deployment.
7. Select the **Backup source to restore** from. **Backup CR** is recommended.
8. Enter the **Backup Name** of the AutomationHubBackup object.
9. Click **Create**.
A new deployment is created and your backup is restored to it.

CHAPTER 4. TROUBLESHOOTING

Use this information to diagnose and resolve issues during backup and recovery.

4.1. AUTOMATION CONTROLLER CUSTOM RESOURCE HAS THE SAME NAME AS AN EXISTING DEPLOYMENT

The name specified for the new AutomationController custom resource must not match an existing deployment or the recovery process will fail.

If your AutomationController customer resource matches an existing deployment, perform the following steps to resolve the issue.

Procedure

1. Delete the existing AutomationController and the associated postgres PVC:

```
oc delete automationcontroller <YOUR_DEPLOYMENT_NAME> -n <YOUR_NAMESPACE>
oc delete pvc postgres-13-<YOUR_DEPLOYMENT_NAME>-13-0 -n
<YOUR_NAMESPACE>
```

2. Use AutomationControllerRestore with the same deployment_name in it:

```
oc apply -f restore.yaml
```