



Red Hat Ansible Automation Platform 2.1

Deploying a high availability automation hub

Overview of the requirements and procedures for a high availability deployment of automation hub.

Red Hat Ansible Automation Platform 2.1 Deploying a high availability automation hub

Overview of the requirements and procedures for a high availability deployment of automation hub.

Legal Notice

Copyright © 2022 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

Providing Feedback: If you have a suggestion to improve this documentation, or find an error, please contact technical support at [to create an issue on the Ansible Automation Platform Jira project](#) using the Docs component.

Table of Contents

PREFACE	3
CHAPTER 1. REQUIREMENTS FOR A HIGH AVAILABILITY AUTOMATION HUB	4
1.1. REQUIRED SHARED FILESYSTEM	4
1.2. NETWORK STORAGE INSTALLATION REQUIREMENTS	4
1.3. REQUIRED CHANGES TO THE INVENTORY FILE	4

PREFACE

This guide provides an overview of the requirements and procedures for a high availability deployment of your automation hub.

A high availability (HA) configuration prepares your system for damage mitigation, where a failure in a primary cluster prompts a standby cluster to take over, resulting in no downtime for the customer. HA provides a reliable and scalable solution that minimizes downtime for your Ansible automation environment. In the case of automation hub, you can set up high availability by configuring your automation hub deployment to do so using the Red Hat Ansible Automation Platform (AAP) installer.

Please note that this guide covers deployment of a HA automation hub application stack only. Other HA components, such as database and file system HA, or setting up DNS load balancing, are out of scope for this guide.

CHAPTER 1. REQUIREMENTS FOR A HIGH AVAILABILITY AUTOMATION HUB

Ensure that the following requirements are met as you configure the Ansible Automation Platform installer. Before you run the Ansible Automation Platform installer, you will need to make several changes to the inventory file found in the Ansible Automation Platform installer, as well as set up other components necessary to deploy in HA. For more information about installing automation hub using the Ansible Automation Platform installer, please see [the Ansible Automation Platform installation guides](#).

1.1. REQUIRED SHARED FILESYSTEM

A high availability automation hub requires you to have a shared file system already set up in your environment. Before you run the Red Hat Ansible Automation Platform installer, verify that the `/var/lib/pulp` directory exists across your cluster as a result of your shared file system installation. The Red Hat Ansible Automation Platform installer will return an error if `/var/lib/pulp` is not detected in one of your nodes, causing your HA automation hub setup to fail.

1.2. NETWORK STORAGE INSTALLATION REQUIREMENTS

If you intend to install a HA automation hub using a network storage on the automation hub nodes itself, you must first install and use **firewalld** to open the necessary ports as required by your shared storage system before running the Ansible Automation Platform installer.

Install and configure **firewalld** by executing the following commands:

1. Install the **firewalld** daemon:

```
$ dnf install firewalld
```

2. Add your network storage under `<service>` using the following command:

```
$ firewall-cmd --permanent --add-service=<service>
```



NOTE

For a list of supported services, use the **\$ firewall-cmd --get-services** command

3. Reload to apply the configuration:

```
$ firewall-cmd --reload
```

1.3. REQUIRED CHANGES TO THE INVENTORY FILE

Before you run the Ansible Automation Platform installer using `./setup.sh`, ensure that you have made the following configuration changes on your inventory file:

Specify database host IP

Specify the IP address for your database host, using the **automation_pg_host** and **automation_pg_port** fields. For example:


```
automationhub_pg_host='192.0.2.10'  
automationhub_pg_port='5432'
```

also specify the IP address for your database host in the [database] section, using the value in the **automationhub_pg_port** field:

```
[database]  
192.0.2.10
```

List all instances in a clustered setup

If installing a clustered setup, replace **localhost ansible_connection=local** in the [automationhub] section with the hostname or IP of all instances. For example:

```
[automationhub]  
automationhub1.testing.ansible.com ansible_user=cloud-user ansible_host=192.0.2.18  
automationhub2.testing.ansible.com ansible_user=cloud-user ansible_host=192.0.2.20  
automationhub3.testing.ansible.com ansible_user=cloud-user ansible_host=192.0.2.22
```

Red Hat Single Sign-On requirements

If you are implementing Red Hat Single Sign-On on your automation hub environment, specify the main automation hub URL that clients will connect to, using the **automationhub_main_url** field. For example:

```
automationhub_main_url = 'https://automationhub.ansible.com'
```



NOTE

If **automationhub_main_url** is not specified, the first node in the [automationhub] group will be used as default.