



Red Hat Service Interconnect 1.5

Installation

Installing the CLI or the Operator

Red Hat Service Interconnect 1.5 Installation

Installing the CLI or the Operator

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 describes how to install Red Hat Service Interconnect 1.5. Red Hat Service Interconnect is a Red Hat build of the open source Skupper project.

Table of Contents

CHAPTER 1. INSTALLING THE CLI	3
1.1. DOWNLOADING BINARIES	3
1.2. USING RED HAT PACKAGES	3
CHAPTER 2. INSTALLING THE OPERATOR	5
2.1. INSTALLING THE OPERATOR FOR ALL NAMESPACES USING THE CLI	5
2.2. INSTALLING THE OPERATOR FOR A SINGLE NAMESPACE USING THE CLI	6
2.3. INSTALLING THE OPERATOR USING THE OPENSIFT CONSOLE	7
CHAPTER 3. UPGRADING THE CLI AND SITES	9
CHAPTER 4. UPGRADING THE RED HAT SERVICE INTERCONNECT OPERATOR AND SITES	10
CHAPTER 5. INSTALLING THE POLICY SYSTEM	11
CHAPTER 6. USING YOUR SUBSCRIPTION	12
6.1. ACCESSING YOUR ACCOUNT	12
6.2. ACTIVATING A SUBSCRIPTION	12
6.3. REGISTERING YOUR SYSTEM FOR PACKAGES	12
APPENDIX A. YAML FOR THE SKUPPER POLICY CRD	13
CHAPTER 7. DEPLOYMENT OPTIONS ON KUBERNETES	15
7.1. SCALING FOR INCREASED TRAFFIC	15
7.2. CREATING A HIGH AVAILABILITY SITE	16
APPENDIX B. ABOUT SERVICE INTERCONNECT DOCUMENTATION	17
MAKING OPEN SOURCE MORE INCLUSIVE	17

CHAPTER 1. INSTALLING THE CLI

The Skupper CLI provides a method to create both Kubernetes and Podman sites.

There are two methods to install the CLI:

- [Section 1.1, “Downloading binaries”](#)
- [Section 1.2, “Using Red Hat packages”](#)

1.1. DOWNLOADING BINARIES

Downloading the Skupper CLI binaries is a quick way to get started with Red Hat Service Interconnect. However, consider using Red Hat packages on Linux to receive the latest updates.

Procedure

1. Download binary files for Linux, macOS or Windows, choose the latest **Version** for 1.5 at [Software Downloads](#).
For a Mac with Apple silicon, use [Rosetta 2](#) and the **Skupper CLI for Mac on x86-64** download.
2. Unzip the downloaded file and place the Skupper executable on your PATH.
3. Verify installation:

```
$ skupper version
client version      1.5.3-rh-5
```

1.2. USING RED HAT PACKAGES

Installing Red Hat packages on Linux makes sure you receive the latest updates to the Skupper CLI.

Prerequisites

- Make sure your subscription is activated and your system is registered. For more information about using the Customer Portal to activate your Red Hat subscription and register your system for packages, see [Chapter 6, Using your subscription](#).

Procedure

1. Use the **subscription-manager** command to subscribe to the required package repositories. Replace **<version>** with **1** for the main release stream or **1.4** for the long term support release stream.



NOTE

Replacing **<version>** with **1** installs **1.5**, while **1.5** is the main release stream and changes after further releases.

Red Hat Enterprise Linux 8

```
$ sudo subscription-manager repos --enable=service-interconnect-__<version>__for-rhel-8-x86_64-rpms
```

Red Hat Enterprise Linux 9

```
$ sudo subscription-manager repos --enable=service-interconnect-__<version>__for-rhel-9-x86_64-rpms
```

- a. Use the **yum** or **dnf** command to install the **skupper** command:

```
$ sudo dnf install skupper-cli
```

Additional information

- See [Examples](#) for the 'Hello world' tutorial.
- Use **man containers.conf** to view more information about podman configuration.

CHAPTER 2. INSTALLING THE OPERATOR

The Red Hat Service Interconnect Operator creates and manages sites in OpenShift.



NOTE

The Red Hat Service Interconnect Operator is supported only on OpenShift 4. Installing an Operator requires administrator-level privileges for your cluster.

2.1. INSTALLING THE OPERATOR FOR ALL NAMESPACES USING THE CLI

The steps in this section show how to use the **oc** command to install and deploy the latest version of the Red Hat Service Interconnect Operator in a given OpenShift cluster. Installing the operator for all namespaces allows you create a site in any namespace. See [Getting started with the OpenShift CLI](#) for more information about the **oc** command.

Prerequisites

- Access to an OpenShift cluster using a **cluster-admin** account. See [Release Notes](#) for supported OpenShift versions.


Procedure

1. Log in to OpenShift as a cluster administrator. For example:

```
$ oc login -u system:admin
```

2. Complete the steps described in [Red Hat Container Registry Authentication](#) .
3. Create a file named **subscription-all.yaml** with the following:

```
apiVersion: operators.coreos.com/v1alpha1
kind: Subscription
metadata:
  name: skupper-operator
  namespace: openshift-operators
spec:
  channel: stable-1
  installPlanApproval: Automatic
  name: skupper-operator
  source: redhat-operators
  sourceNamespace: openshift-marketplace
  startingCSV: skupper-operator.v1.5.3-rh-1
```

**NOTE**

If you want to keep updates confined to 1.5.x releases, set the value of **channel** to **stable-1.5**.

If you do not specify **startingCSV**, the subscription defaults to the latest operator version.

If you specify **installPlanApproval** as **Manual**, sites are not automatically upgraded to the latest version of Service Interconnect. See [Chapter 4, Upgrading the Red Hat Service Interconnect Operator and sites](#) for information on manually upgrading sites.

4. Apply the subscription YAML:

```
$ oc apply -f subscription-all.yaml
```

Additional information

- See [Using Skupper](#) for instructions about using YAML to create sites.

2.2. INSTALLING THE OPERATOR FOR A SINGLE NAMESPACE USING THE CLI

The steps in this section show how to use the **oc** command to install and deploy the latest version of the Red Hat Service Interconnect Operator in a given OpenShift cluster. Installing the operator for a single namespace allows you create a site in the specified namespace. See [Getting started with the OpenShift CLI](#) for more information about the **oc** command.

Prerequisites

- Access to an OpenShift cluster using a **cluster-admin** account. See [Release Notes](#) for supported OpenShift versions.

Procedure

1. Log in to OpenShift as a cluster administrator. For example:

```
$ oc login -u system:admin
```

2. Complete the steps described in [Red Hat Container Registry Authentication](#) .
3. Create an Operator group in the namespace where you want to create a site:
 - a. Create a file named **operator-group.yaml** with the following:

```
kind: OperatorGroup
apiVersion: operators.coreos.com/v1
metadata:
  name: skupper-operator
  namespace: my-namespace
spec:
  targetNamespaces:
    - my-namespace
```

■

where **my-namespace** is the name of the namespace you want to create the site.

- b. Apply the Operator group YAML:

```
$ oc apply -f operator-group.yaml
```

4. Create a file named **subscription-myns.yaml** with the following:

```
apiVersion: operators.coreos.com/v1alpha1
kind: Subscription
metadata:
  name: skupper-operator
  namespace: my-namespace
spec:
  channel: stable-1
  installPlanApproval: Automatic
  name: skupper-operator
  source: redhat-operators
  sourceNamespace: openshift-marketplace
  startingCSV: skupper-operator.v1.5.3-rh-1
```

where **my-namespace** is the name of the namespace you want to create the site.



NOTE

If you want to keep updates confined to 1.5.x releases, set the value of **channel** to **stable-1.5**.

If you do not specify **startingCSV**, the subscription defaults to the latest operator version.

If you specify **installPlanApproval** as **Manual**, sites are not automatically upgraded to the latest version of Service Interconnect. See [Chapter 4, Upgrading the Red Hat Service Interconnect Operator and sites](#) for information on manually upgrading sites.

5. Apply the subscription YAML:

```
$ oc apply -f subscription-myns.yaml
```

Additional information

- See [Using Skupper](#) for instructions about using YAML to create sites.

2.3. INSTALLING THE OPERATOR USING THE OPENSIFT CONSOLE

The procedures in this section show how to use the OperatorHub from the OpenShift console to install and deploy the latest version of the Red Hat Service Interconnect Operator in a given OpenShift namespace.

Prerequisites

- Access to an OpenShift cluster using a **cluster-admin** account. See [Release Notes](#) for supported OpenShift versions.

Procedure

1. In the OpenShift web console, navigate to **Operators** → **OperatorHub**.
2. Choose **Red Hat Service Interconnect Operator** from the list of available Operators, and then click **Install**.
3. On the **Operator Installation** page, two **Installation mode** options are available:
 - **All namespaces on the cluster**
 - **A specific namespace on the cluster**
For this example, choose **A specific namespace on the cluster**
4. Choose an **Update approval** option.
By default, **Automatic** approval is selected, and sites will upgrade to the latest version of Service Interconnect. If you choose **Manual** approval, sites will not be automatically upgraded to the latest version of Service Interconnect. See [Chapter 4, Upgrading the Red Hat Service Interconnect Operator and sites](#) for information on manually upgrading sites.
5. Select the namespace into which you want to install the Operator, and then click **Install**.
The **Installed Operators** page appears displaying the status of the Operator installation.
6. Verify that the Red Hat Service Interconnect Operator is displayed and wait until the **Status** changes to **Succeeded**.
7. If the installation is not successful, troubleshoot the error:
 - a. Click **Red Hat Service Interconnect Operator** on the **Installed Operators** page.
 - b. Select the **Subscription** tab and view any failures or errors.

For more information about installing Operators, see the [OpenShift Documentation](#)

Additional information

- See [Using Skupper](#) for instructions about using YAML to create sites.

CHAPTER 3. UPGRADING THE CLI AND SITES

Upgrading sites requires the latest version of the Skupper CLI.



NOTE

Update all sites to ensure the same version of Service Interconnect is running across your service network. You can expect some minimal downtime during the update process.

Procedure

1. Upgrade the **skupper** CLI:

```
$ dnf upgrade skupper-cli
```

2. Upgrade each site:

- a. Set the context to the site you want to upgrade.
- b. Run the upgrade command:

```
$ skupper update
```

CHAPTER 4. UPGRADING THE RED HAT SERVICE INTERCONNECT OPERATOR AND SITES

This release changes the Operator channel name from **alpha**. You can now choose from the following channels:

- **stable-1.5** - Limit updates to 1.5.x.
- **stable-1** - Limit updates to 1.x.
- **stable** - All releases.

You must install the Operator again as described in [Installing the Operator](#) to use version 1.5.

If you chose automatic updates when installing the Red Hat Service Interconnect Operator, sites are upgraded to future releases whenever a new version of Service Interconnect is available for the channel you chose. If you chose manual updates, complete this procedure to upgrade your sites.

Procedure

1. Log into the OpenShift console.
2. Navigate to the **Installed Operators** page.
3. If **Upgrade available** is displayed as the **Status** for the Red Hat Service Interconnect Operator, click that text.
4. On the **InstallPlan details** page, click **Preview InstallPlan**.
5. Click **Approve** to upgrade sites.

CHAPTER 5. INSTALLING THE POLICY SYSTEM

Installing the Skupper policy system on a cluster allows you control how Skupper is used on the cluster.



NOTE

Applying the policy system in a cluster without specific policy rules prohibits site linking and service exposure. If you are installing the policy system on a cluster where there are existing sites, you must create policies before installing the policy system to avoid disruption.

Prerequisites

- Access to a Kubernetes cluster with **cluster-admin** privileges.
- The Red Hat Service Interconnect Operator is installed

Procedure

1. Log into your cluster.
2. Deploy the policy CRD:

```
$ kubectl apply -f skupper_cluster_policy_crd.yaml
```

```
customresourcedefinition.apiextensions.k8s.io/skupperclusterpolicies.skupper.io created  
clusterrole.rbac.authorization.k8s.io/skupper-service-controller created
```

where the contents of **skupper_cluster_policy_crd.yaml** is specified in the [Appendix A, YAML for the Skupper policy CRD](#) appendix.

Additional information

See [Securing a service network using policies](#) for more information about using policies.

CHAPTER 6. USING YOUR SUBSCRIPTION

Red Hat Service Interconnect is provided through a software subscription. To manage your subscriptions, access your account at the Red Hat Customer Portal.

6.1. ACCESSING YOUR ACCOUNT

Procedure

1. Go to access.redhat.com.
2. If you do not already have an account, create one.
3. Log in to your account.

6.2. ACTIVATING A SUBSCRIPTION

Procedure

1. Go to access.redhat.com.
2. Navigate to **My Subscriptions**.
3. Navigate to **Activate a subscription** and enter your 16-digit activation number.

6.3. REGISTERING YOUR SYSTEM FOR PACKAGES

To install RPM packages for this product on Red Hat Enterprise Linux, your system must be registered. If you are using downloaded release files, this step is not required.

Procedure

1. Go to access.redhat.com.
2. Navigate to **Registration Assistant**.
3. Select your OS version and continue to the next page.
4. Use the listed command in your system terminal to complete the registration.

For more information about registering your system, see one of the following resources:

- [Red Hat Enterprise Linux 8 - Registering the system and managing subscriptions](#)
- [Red Hat Enterprise Linux 9 - Registering the system and managing subscriptions](#)

APPENDIX A. YAML FOR THE SKUPPER POLICY CRD

The policy system allows a cluster administrator restrict Skupper usage on a cluster. It is not required for typical Skupper usage.

The following YAML applies the Skupper policy CRD to a cluster.

```

---
apiVersion: apiextensions.k8s.io/v1
kind: CustomResourceDefinition
metadata:
  name: skupperclusterpolicies.skupper.io
spec:
  group: skupper.io
  versions:
  - name: v1alpha1
    served: true
    storage: true
    schema:
      openAPIV3Schema:
        type: object
        properties:
          spec:
            type: object
            properties:
              namespaces:
                type: array
                items:
                  type: string
              allowIncomingLinks:
                type: boolean
              allowedOutgoingLinksHostnames:
                type: array
                items:
                  type: string
              allowedExposedResources:
                type: array
                items:
                  type: string
              allowedServices:
                type: array
                items:
                  type: string
      scope: Cluster
      names:
        plural: skupperclusterpolicies
        singular: skupperclusterpolicy
        kind: SkupperClusterPolicy
---
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRole
metadata:
  labels:
    application: skupper-service-controller
  name: skupper-service-controller
rules:

```

- apiGroups:
 - skupper.io
- resources:
 - skupperclusterpolicies
- verbs:
 - get
 - list
 - watch
- apiGroups:
 - ""
- resources:
 - namespaces
- verbs:
 - get

CHAPTER 7. DEPLOYMENT OPTIONS ON KUBERNETES

When you create a site on Kubernetes, there are many options you can use. For example, you can set the number of pods and the resources allocated to each pod. This guide focusses on the following goals:

- [Section 7.1, "Scaling for increased traffic"](#)
- [Section 7.2, "Creating a high availability site"](#)

7.1. SCALING FOR INCREASED TRAFFIC

For optimal network latency and throughput, you can adjust the CPU allocation for the router using the **router-cpu** option. Router CPU is the primary factor governing Skupper network performance.



NOTE

Increasing the number of routers does not improve network performance. An incoming router-to-router link is associated with just one active router. Additional routers do not receive traffic while that router is responding

1. Determine the router CPU allocation you require.
By default, the router CPU allocation is **BestEffort** as described in [Pod Quality of Service Classes](#).

Consider the following CPU allocation options:

Router CPU	Description
1	Helps avoid issues with BestEffort on low resource clusters
2	Suitable for production environments
5	Maximum performance

2. If you are using the Skupper CLI, set the CPU allocation for the router using the **--router-cpu** option. For example:

```
$ skupper init --router-cpu 2
```

3. If you are using YAML, set the CPU allocation for the router by setting a value for the **router-cpu** attribute. For example:

```
apiVersion: v1
kind: ConfigMap
metadata:
  name: "skupper-site"
data:
  name: "my-site"
  router-cpu: 2
```

7.2. CREATING A HIGH AVAILABILITY SITE

By default, Kubernetes restarts any router that becomes unresponsive. (If you encounter router restarts, consider [Section 7.1, "Scaling for increased traffic"](#) in order to improve responsiveness.)

If the cluster where you are running Skupper is very busy, it may take time for Kubernetes to schedule a new router pod. You can "preschedule" a backup router by deploying two routers in a site.

1. If you are using the Skupper CLI, set the number of routers to **2** using the **--routers** option:

```
$ skupper init --routers 2
```

2. If you are using YAML, set the number of routers to **2** by setting the **routers** attribute:

```
apiVersion: v1
kind: ConfigMap
metadata:
  name: "skupper-site"
data:
  name: "my-site"
  routers: 2
```

Setting the number of routers to more than two does not provide increased availability and can adversely affect performance.

Note: Clients must reconnect when a router restarts or traffic is redirected to a backup router.

APPENDIX B. ABOUT SERVICE INTERCONNECT DOCUMENTATION

MAKING OPEN SOURCE MORE INCLUSIVE

Red Hat is committed to replacing problematic language in our code, documentation, and web properties. We are beginning with these four terms: master, slave, blacklist, and whitelist. Because of the enormity of this endeavor, these changes will be implemented gradually over several upcoming releases. For more details, see [our CTO Chris Wright's message](#).

Revised on 2024-06-24 15:01:15 UTC