Installing Cloudbreak on OpenStack

Date of Publish: 2018-09-14

http://docs.hortonworks.com

Contents

Prerequisites on OpenStack....3
- System requirements....3
  - Supported OpenStack distributions....3
  - Standard modules....3
- Virtual network....3
- Security group....3
- SSH key pair....3
- Browser....4

Preparing the VM....4
- System requirements....4
- Root access....5
- System updates....5
- Iptables....5
- Disable SELINUX....5
- Docker....6

Install Cloudbreak on a VM....6

Configure a self-signed certificate....8

Access Cloudbreak web UI....8

Configure an external Cloudbreak database....9

Next steps....10

Prerequisites on OpenStack

Before installing Cloudbreak, you must meet the following prerequisites:

System requirements

Before launching Cloudbreak on your OpenStack, make sure that your OpenStack deployment fulfills the following requirements.

Supported OpenStack distributions

The following versions of the Red Hat distribution of OpenStack (RDO) are supported:

- Juno
- Kilo
- Liberty
- Mitaka

Standard modules

Cloudbreak requires that the following standard modules are installed and configured on OpenStack:

- Keystone V2 or Keystone V3
- Neutron (Self-service and provider networking)
- Nova (KVM or Xen hypervisor)
- Glance
- Heat
- Cinder (Optional)

Virtual network

You must have a virtual network configured on your cloud provider.

Security group

Ports 22 (SSH), 80 (HTTPS), and 443 (HTTPS) must be open on the security group.

SSH key pair

In order to access the Cloudbreak VM via SSH, you will be required to use your SSH key pair.

Generate a new SSH key pair

All the instances created by Cloudbreak are configured to allow key-based SSH, so you'll need to provide an SSH public key that can be used later to SSH onto the instances in the clusters you'll create with Cloudbreak. You can use one of your existing keys or you can generate a new one.

To generate a new SSH key pair, execute:

```
ssh-keygen -t rsa -b 4096 -C "your_email@example.com"
# Creates a new ssh key, using the provided email as a label
# Generating public/private rsa key pair.
```

```
# Enter file in which to save the key (/Users/you/.ssh/id_rsa): [Press
 enter]
```

You'll be asked to enter a passphrase, but you can leave it empty:

```
# Enter passphrase (empty for no passphrase): [Type a passphrase]
# Enter same passphrase again: [Type passphrase again]
```

After you enter (or not) a passphrase, the key pair is generated. The output should look similar to:

```
# Your identification has been saved in /Users/you/.ssh/id_rsa.
# Your public key has been saved in /Users/you/.ssh/id_rsa.pub.
# The key fingerprint is:
# 01:0f:f4:3b:ca:85:sd:17:sd:7d:sd:68:9d:sd:a2:sd your_email@example.com
```

Later you'll need to pass the content of the .pub file to Cloudbreak and use the private key file to SSH to the instances.

Recover public SSH key

The -y option of ssh-keygen outputs the public key. For example:

```
ssh-keygen -y -f ~/.ssh/id_rsa > ~/.ssh/id_rsa.pub
```

Copy public SSH key

You can use pbcopy to quickly copy your SSH public key. For example: pbcopy < ~/.ssh/id_rsa.pub

Browser

In order to access Cloudbreak web UI, you should use one of the following supported browsers: Chrome, Firefox, or Safari.

Preparing the VM

To install the Cloudbreak deployer and install the Cloudbreak application, you must have an existing VM.

You should launch the VM by using the steps provided in your cloud provider documentation. Once you have the VM ready, perform the steps listed in this section.

System requirements

In order to install Cloudbreak, your system must meet the minimum requirements.

Ensure that your system meets the following requirements:

- Minimum VM requirements: 16GB RAM, 40GB disk, 4 cores
- Supported operating systems: RHEL, CentOS, and Oracle Linux 7 (64-bit)

Note:

You can install Cloudbreak on Mac OS X for evaluation purposes only. Mac OS X is not supported for a production deployment of Cloudbreak.

Root access

Every command mentioned in this documentation must be executed as root.

In order to get root privileges execute:

```
sudo -i
```

System updates

Perform these steps to ensure that your system is up-to-date.

To ensure that your system is up-to-date, run:

```
yum -y update
```

Reboot the VM if necessary.

Iptables

Perform these steps to install and configure iptables.

Steps

1. Install iptables-services:

```
yum -y install net-tools ntp wget lsof unzip tar iptables-services
systemctl enable ntpd && systemctl start ntpd
systemctl disable firewalld && systemctl stop firewalld
```

Note:

Without iptables-services installed the iptables save command will not be available.

2. Configure permissive iptables on your machine:

```
iptables --flush INPUT && \
iptables --flush FORWARD && \
service iptables save
```

Disable SELINUX

Perform these steps to disable SELINUX.

Steps

1. Disable SELINUX:

```
setenforce 0
sed -i 's/SELINUX=enforcing/SELINUX=disabled/g' /etc/selinux/config
```

2. Run the following command to ensure that SELinux is not turned on afterwards:

```
getenforce
```

3. The command should return “Disabled”.

Docker

Perform these steps to install Docker.

The minimum Docker version is 1.13.1. If you are using an older image that comes with an older Docker version, upgrade Docker to 1.13.1 or newer.

Steps

1. Install Docker service:

   CentOS 7

   ```
   yum install -y docker
   systemctl start docker
   systemctl enable docker
   ```

   RHEL 7

   ```
   yum install yum-utils
   yum-config-manager --enable rhui-REGION-rhel-server-extras
   yum install -y docker
   systemctl start docker
   systemctl enable docker
   ```

2. Check the Docker Logging Driver configuration:

   ```
   docker info | grep "Logging Driver"
   ```

3. If it is set to Logging Driver: journald, you must set it to “json-file” instead. To do that:

   a. Open the docker file for editing:

      ```
      vi /etc/sysconfig/docker
      ```

   b. Edit the following part of the file so that it looks like below (showing log-driver=json-file):

      ```
      # Modify these options if you want to change the way the docker daemon
       runs
      OPTIONS='--selinux-enabled --log-driver=json-file --signature-
      verification=false'
      ```

   c. Restart Docker:

      ```
      systemctl restart docker
      systemctl status docker
      ```

Install Cloudbreak on a VM

Install Cloudbreak on your own VM from a Cloudbreak deployer binary.

Steps

1. Install the Cloudbreak deployer and unzip the platform-specific single binary to your PATH. For example:

   ```
   yum -y install unzip tar
   curl -Ls public-repo-1.hortonworks.com/HDP/cloudbreak/cloudbreak-
   deployer_2.8.0_$(uname)_x86_64.tgz | sudo tar -xz -C /bin cbd
   cbd --version
   ```

Once the Cloudbreak deployer is installed, you can set up the Cloudbreak application.

2. Create a Cloudbreak deployment directory and navigate to it:

```
mkdir cloudbreak-deployment
cd cloudbreak-deployment
```

3. In the directory, create a file called Profile with the following content:

```
export UAA_DEFAULT_SECRET=MY-SECRET
export UAA_DEFAULT_USER_PW=MY-PASSWORD
export UAA_DEFAULT_USER_EMAIL=MY-EMAIL
export PUBLIC_IP=MY_VM_IP
```

For example:

```
export UAA_DEFAULT_SECRET=MySecret123
export UAA_DEFAULT_USER_PW=MySecurePassword123
export UAA_DEFAULT_USER_EMAIL=dbialek@hortonworks.com
export PUBLIC_IP=172.26.231.100
```

You will need to provide the email and password when logging in to the Cloudbreak web UI and when using the Cloudbreak CLI. The secret will be used by Cloudbreak for authentication.

You should set the CLOUDBREAK_SMTP_SENDER_USERNAME variable to the username you use to authenticate to your SMTP server. You should set the CLOUDBREAK_SMTP_SENDER_PASSWORD variable to the password you use to authenticate to your SMTP server.

4. Generate configurations by executing:

```
rm *.yml
cbd generate
```

The cbd start command includes the cbd generate command which applies the following steps:

- Creates the docker-compose.yml file, which describes the configuration of all the Docker containers required for the Cloudbreak deployment.
- Creates the uaa.yml file, which holds the configuration of the identity server used to authenticate users with Cloudbreak.

5. Start the Cloudbreak application by using the following commands:

```
cbd pull-parallel
cbd start
```

This will start the Docker containers and initialize the application. The first time you start the Cloudbreak app, the process will take longer than usual due to the download of all the necessary docker images.

6. Next, check Cloudbreak application logs:

```
cbd logs cloudbreak
```

You should see a message like this in the log: Started CloudbreakApplication in 36.823 seconds. Cloudbreak normally takes less than a minute to start.

Related Information

Troubleshooting Cloudbreak

Configure a self-signed certificate

If your OpenStack is secured with a self-signed certificate, you need to import that certificate into Cloudbreak, or else Cloudbreak won't be able to communicate with your OpenStack.

To import the certificate, place the certificate file in the /certs/trusted/ directory by using the following these steps.

Steps

1. Navigate to the certs directory (automatically generated).
2. Create the trusted directory.
3. Copy the certificate to the trusted directory.

Cloudbreak will automatically pick up the certificate and import it into its trust store upon start.

Access Cloudbreak web UI

Log in to the Cloudbreak UI by using the following steps.

Steps

1. You can log into the Cloudbreak application at https://IP_Address. For example https://34.212.141.253. You may use cbd start to obtain the login information. Alternatively, you can obtain the VM's IP address from your cloud provider console.
2. Confirm the security exception to proceed to the Cloudbreak web UI.

   The first time you access Cloudbreak UI, Cloudbreak will automatically generate a self-signed certificate, due to which your browser will warn you about an untrusted connection and will ask you to confirm a security exception.

Browser	**Steps**
Firefox	Click Advanced > Click Add Exception… > Click Confirm Security Exception
Safari	Click Continue
Chrome	Click Advanced > Click Proceed…

3. The login page is displayed:

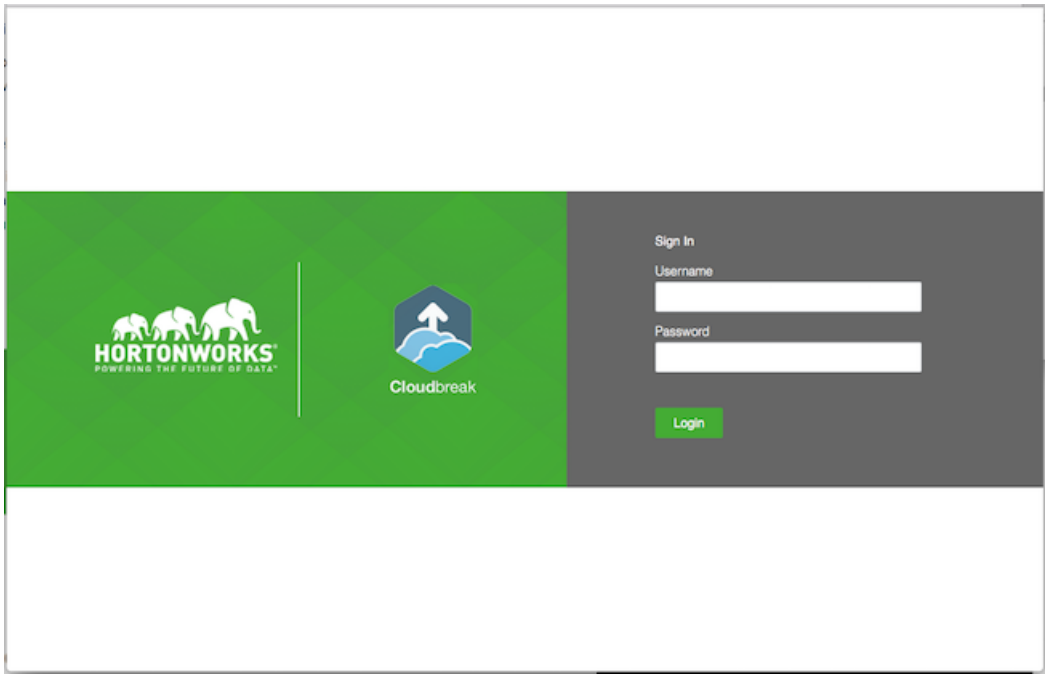

4. Log in to the Cloudbreak web UI using the credentials that you configured in your Profile file:
   - The username is the UAA_DEFAULT_USER_EMAIL
   - The password is the UAA_DEFAULT_USER_PW
5. Upon a successful login, you are redirected to the dashboard:

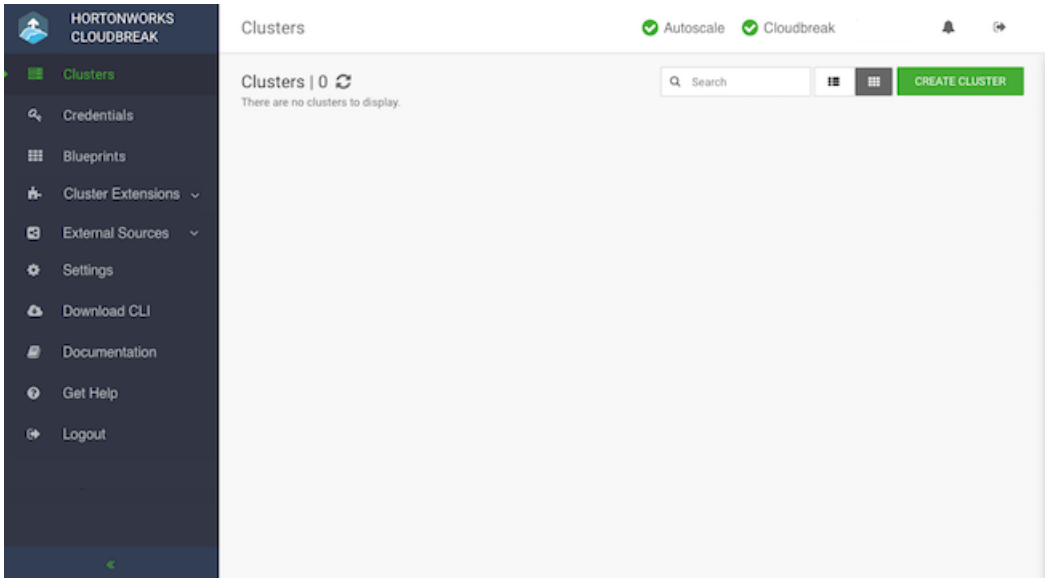

Configure an external Cloudbreak database

By default, Cloudbreak uses an embedded PostgreSQL database to persist data related to Cloudbreak configuration and so on. This database is only suitable for non-production Cloudbreak deployments. For production, you must configure an external Cloudbreak database.

Related Information

External Cloudbreak database

Next steps

After launching Cloudbreak, you must configure a Cloudbreak credential. After you've created one, you can start creating clusters.

Related Information

Create a Cloudbreak credential on OpenStack

Configure Cloudbreak