Cloudbreak Security Overview

Date of Publish: 2019-05-28

https://docs.hortonworks.com/

Contents

Security overview

Cloudbreak utilizes cloud provider security resources such as virtual networks, security groups, and identity and access management:

1. Network isolation is achieved via user-configured virtual networks and subnets.
2. Network security is achieved via out-of-the-box security group settings.
3. Controlled use of cloud resources using IAM roles (AWS, GCP) or Active Directory (Azure).

Virtual networks

Cloud providers use virtual networks which resemble traditional networks. Depending on the options that you selected during deployment, your Cloudbreak instance and clusters are launched into new or existing cloud provider networking infrastructure (virtual networks and subnets). For more information about virtual networks, refer to the cloud-provider documentation:

Cloud provider	External documentation link
AWS	Amazon Virtual Private Cloud (Amazon VPC)
Azure	Microsoft Azure Virtual Network
Google Cloud Platform	Virtual Private Cloud (VPC) network
OpenStack	Network

Network security

Security groups are set up to control network traffic to the instances in the system.

Cloudbreak uses public IP addresses when communicating with cluster nodes. On AWS, you can configure it to use private IPs instead. For instructions, refer to the instructions for Configuring communication via private IPs on AWS.

To learn about ports required for Cloudbreak and Cloudbreak-managed clusters, refer to the following documentation:

Related Information

Configure communication via private IPs on AWS

Cloudbreak instance security group

This section lists ports used by Cloudbreak.

Cloudbreak uses specific ports for access via SSH and HTTP/HTTPS. The following table lists the minimum security group port configuration required for the Cloudbreak instance:

Inbound ports

Port	Description
22	SSH access to the Cloudbreak VM.
80	HTTP access to the Cloudbreak web UI. This is automatically redirected to the HTTPS (443) port.
443	HTTPS access to the Cloudbreak web UI.

Outbound ports

Port	Description
80	Cloudbreak communicates with services via HTTP.
9443	Cloudbreak communicates with the Ambari Server running on its managed clusters.
443	Cloudbreak communicates with the Ambari Server via HTTPS when gateway is not configured for a cluster.
8443	Cloudbreak communicates with the gateway when gateway is configured for a cluster.

Cluster security groups

This section lists ports used by Cloudbreak-manged clusters.

The following tables lists the default and recommended cluster security group settings:

Note:

By default, when creating a cluster, a new network, subnet, and security groups are created automatically. The default experience of creating network resources such as network, subnet and security group automatically is provided for convenience. We strongly recommend that you review these options and for production cluster deployments leverage your existing network resources that you have defined and validated to meet your enterprise requirements.

Note:

Depending on the cluster components that you are planning to use, you may need to open additional ports required by these components.

External ports

Source	Target	Protocol	Port	Description
Cloudbreak	Ambari server	TCP	9443	• This port is used by Cloudbreak to maintain management control of the cluster. • The default security group opens 9443 from anywhere. You should limit this CIDR further to only allow access from the Cloudbreak host. This can be done by default by restricting inbound access from Cloudbreak to cluster.
*	All cluster hosts	TCP	22	• This is an optional port for end user SSH access to the hosts. • You should review and limit or remove this CIDR access.

Source	Target	Protocol	Port	Description
*	Ambari server	TCP	8443	• This port is used to access the gateway (if configured). • You should review and limit this CIDR access. • If you do not configure the gateway, this port does not need to be opened. If you want access to any cluster resources, you must open this port explicitly on the security groups for their respective hosts.
*	Ambari server	TCP	443	• This port is used to access Ambari directly. • If you are configuring the gateway, you should access Ambari through the gateway; In this case you do not need to open this port. • If you do not configure the gateway, to obtain access to Ambari, you can open this port on the security group for the respective host.

Internal ports

In addition to the ports described above, Cloudbreak uses certain ports for internal communication within the subnet. By default, Cloudbreak opens ports 0-65535 to the subnet's internal CIDR (such as 10.0.0.0/16). Use the following table to limit this CIDR:

Source	Target	Protocol	Port	Description
Salt-bootstrap	Gateway instance (Ambari server instance)	TCP	7070	Salt-bootstrap service launches and configures Saltstack.
Salt-master	All hosts in the cluster	TCP	4505, 4506	Salt-minions connect to the Salt-master(s).
Consul server	All hosts in the cluster	TCP, UDP	8300, 8301	Consul agents connect to the Consul server.
Consul agent (all hosts in the cluster)	All hosts in the cluster	TCP, UDP	8300, 8301	Consul agents connect to other Consul agents (Gossip protocol).
Prometheus node exporter	Gateway instance (Ambari server instance)	TCP	9100	Prometheus server scrapes metrics from the node exporters.
Ambari server	All hosts in the cluster	Refer to Default network port numbers for Ambari in Ambari documentation.		Ambari agents connect to the Ambari server.

When creating data lakes and their attached clusters, you must also open the following internal port:

Source	Target	Protocol	Port	Description
Data lake cluster	Clusters attached to the data lake	TCP	6080	Used for communication between the data lake and attached clusters.

Related Information

Restrict inbound access to clusters

Default network port numbers for Ambari (Ambari)

Identity management

To securely control access to cloud resources, cloud providers use identity management services such as IAM roles (AWS and GCP) and Active Directory (Azure).

Cloud provider	External documentation link
AWS	AWS Identity and Access Management (IAM)
Azure	Azure Active Directory (Azure AD)
Google	Google Cloud Identity and Access Management (IAM)
OpenStack	Keystone

Cloudbreak utilizes cloud provider's identity management services via Cloudbreak credential. After launching Cloudbreak on your chosen cloud provider, you must create a Cloudbreak credential, which allows Cloudbreak to authenticate with your cloud provider identity management service. Only after you have completed this step, Cloudbreak can create resources on your behalf.

Refer to the following documentation to learn more about your cloud provider's identity management and how it is used by Cloudbreak:

Authentication with AWS

There are two ways for Cloudbreak to authenticate with and obtain authorization from AWS: be role-based or key-based.

After launching Cloudbreak on AWS, you are required to select one way for Cloudbreak to authenticate with your AWS account and create resources on your behalf: key-based or role-based. While key-based authentication simply uses your AWS access key and secret key, role-based authentication uses IAM roles.

Role-based authentication

If you are using role-based authentication for Cloudbreak on AWS, you will eventually create two IAM roles: one to grant Cloudbreak access to allow Cloudbreak to assume AWS roles (using the AssumeRole policy) and the second one to provide Cloudbreak with the capabilities required for cluster creation (using the CbPolicy policy).

The following diagram and table provide contextual information about the two roles required:

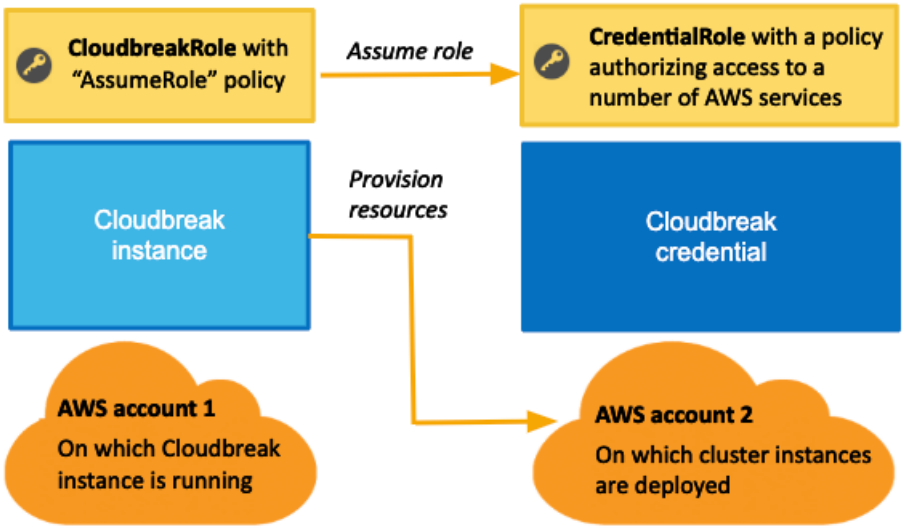

Note:

The AWS account 1 and AWS account 2 presented in the diagram can be the same account.

Role	Purpose	Overview of steps	Where to perform
CloudbreakRole	Allows Cloudbreak to assume other IAM roles - in this case Cloudbreak will assume the CredentialRole.	1. Create a role called CloudbreakRole and attach the AssumeRole policy. 2. As part of Cloudbreak installation process, attach the CloudbreakRole IAM role to the Cloudbreak VM.	1. Create the IAM role and policy in the AWS IAM console. 2. Attach the IAM role to the VM in the EC2 console.
CredentialRole	Authorizes Cloudbreak to create AWS resources, such as VMs, required for clusters.	1. Prior to creating a Cloudbreak credential, you must create an IAM role called CredentialRole and attach the CbPolicy policy to it. 2. When creating a role-based Cloudbreak credential, provide the IAM Role ARN of this role toCloudbreak.	1. Create the IAM role and policy in the AWS IAM console. 2. Create a role-based Cloudbreak credential in the Cloudbreak web UI.

Note:

These role and policy names are just examples. You may use different names when creating your resources.

Alternatively, instead of attaching the CloudbreakRole role during the VM launch, you can assign the CloudbreakRole to an IAM user and then add the access and secret key of that user to your Profile.

Alternatively you can generate the CredentialRole role later once your Cloudbreak VM is running by SSHing to the Cloudbreak VM and running the cbd aws generate-role command. This command creates a role with the name "cbreak-deployer" (equivalent to the CredentialRole). To customize the name of the role, add export AWS_ROLE_NAME=my-cloudbreak-role-name (where "my-cloudbreak-role-name" is your custom role name) as a new line to your Profile. If you choose this option, you must make sure that the CloudbreakRole or the IAM user have a permission not only to assume a role but also to create a role.

Authentication with Azure

There are two ways for Cloudbreak to authenticate with Azure: interactive and app-based.

After launching Cloudbreak on Azure, you are required to create a Cloudbreak credential, which allows Cloudbreak to authenticate with your Azure Active Directory. You have two options:

- Interactive: The app and service principal creation and role assignment are fully automated, so the only input that you need to provide to Cloudbreak is your Subscription ID and Directory ID.
- App-based: The app and service principal creation and role assignment are not automated You must create an Azure Active Directory application registration and then provide its parameters to Cloudbreak, in addition to providing your Subscription ID and Directory ID.

Authentication with GCP

Authentication with GCP is via a service account.

After launching Cloudbreak on GCP, you are required to register a service account in Cloudbreak by creating a Cloudbreak credential. Cloudbreak uses this account to authenticate with the GCP identity management service and to authorize Cloudbreak to provision resources on your behalf.

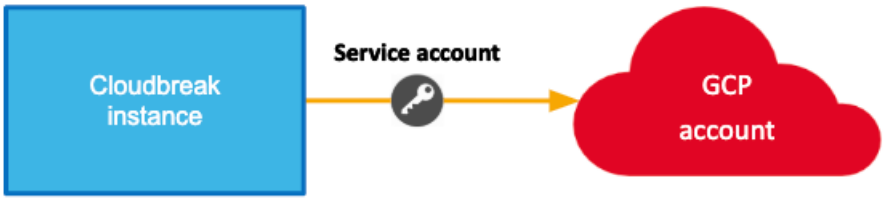

Authentication with OpenStack

Authentication with OpenStack can be configured by providing the required information about your OpenStack account to Cloudbreak.

After launching Cloudbreak on OpenStack, you are required to create a Cloudbreak credential, which allows Cloudbreak to authenticate with keystone.