Cloudbursting: How to configure Cluster Extension to AWS without OpenVPN (7.3+)

This article describes how to configure Cluster Extension cloudbursting to AWS on a Bright 7.3 cluster without relying on OpenVPN for communication between the headnode and the cloud director, but instead relying on an existing IP connectivity to/from the VPC (e.g. Amazon Direct Connect).

(For older Bright versions (pre 7.3), for settting up Cluster Extension with Direct Connect (or simply with OpenVPN-based bursting to a pre-existing VPC) check out this FAQ entry

Prerequisites:

- installed Bright 7.3 Cluster
- VPC created in AWS
- working IP connectivity from your head node, to your VPC (e.g. via Amazon Direct Connect) https://aws.amazon.com/directconnect/ Nodes within your VPC should be able to reach (via IP) the head node of your cluster, and vice versa
- ICMP working on the link between the cloud director, and the head node. During cloud director boot process, the director will try to ping the head node's external IP before proceeding with installation. If the director is not capable of doing that (e.g. because of Firewall config on Direct Connect), it will not be able to enter the node installer.

Configuration steps:

- Start cm-cluster-extension CLI tool, select "Add cloud provider" menu entry.
- Follow on-screen instructions up until the "Summary" screen. When selecting "region" select the region which contains the existing VPC you want to use.
- On the last ("Summary") screen go to 'Advanced' -> clusterextension:
    - Create tunnel networks: **No**.
    - VPC for region <region>: Select the VPC you want to burst info.
    - Security groups: Pick security groups within the VPC to use for cloud director and cloud nodes, otherwise new SGs will be created.
- Go back to "Summary" screen and finish deployment.

If all goes well, you should now be able to power on a cloud director and provision it over DirectConnect.

Here is an alternative procedure for Bright versions older than 7.3-19, which don't have menus for selecting an existing VPC.

Cloudbursting: How to configure Cluster Extension to AWS without OpenVPN (7.3+)

Overview

The entire process involves two major steps:

- configuring regular (OpenVPN - based) Cluster extension to AWS,

- reconfigure the so-obtained configuration to not use OpenVPN.

Configure to burst to pre-existing VPC

Before you can reconfigure the cluster to not use OpenVPN for cluster extension, you first will need to follow the same steps as needed to configure regular Cluster Extension to a pre-existing VPC with OpenVPN. You can find those steps here.

Once that's set up, and before you power on your cloud director, you will need to modify your cluster configuration to not use OpenVPN, and instead use your pre-existing IP connectivity to the VPC. This is covered in the remainder of this document.

Cloud Director

Start the cmsh command line tool on your headnode. Follow these steps to reconfigure the cloud director to not use OpenVPN.

```
[head]% device
[head->device]% use vpc-director # the hostname of your cloud director
[head->device*[vpc-director*]->cloudsettings*]% ..
[head->device*[vpc-director*]]% set managementnetwork vpc-sn1  # name
of your VPC subnet
# Note, that the network you assign to the eth0 interface determines
# inside which subnet the cloud node will be created
# for cloud director it's imnportant that routing rules allow for a wo
rking Public (Elastic) IP to be assigned to the node on that subnet, a
s the head node will have to be able to connect to the cloud director,
 shortly after director gets powered on
[head->device*[vpc-director*]]% interfaces
[head->device*[vpc-director*]->interfaces]% remove tun0
[head->device*[vpc-director*]->interfaces*]% set eth0 network vpc-sn1
```

```
[head->device*[vpc-director*]->interfaces]% commit
```

Cloud Nodes

Now, modify the network interface configuration on the cloud nodes.

[head->device]% use cnode001
[head->device*[cnode001*]]% interfaces
[head->device*[cnode001*]->interfaces]% remove tun0
[head->device*[cnode001*]->interfaces*]% set eth0 network vpc-sn2
[head->device*[cnode001*]->interfaces*]% ..
[head->device*[cnode001*]]% set managementnetwork vpc-sn2
[head->device*[cnode001*]]% commit

Cluster configuration

Now, you will need to remove the tunnel network, and the netmap network (along with any interfaces on that network). Only physical nodes and the head node should have interface on that network.

For all compute nodes and the head node

[head->device]% use master
[head->device[head]]% interfaces
[head->device*[head]->interfaces]% remove tun0
[head->device*[head]]% commit
[head->device]% use node001
[head->device*[node001*]]% interfaces
[head->device*[node001*]->interfaces]% remove nmap0
[head->device*[node001*]]% commit

[head->network]% remove <name of the tunnel network>
[head->network*]% remove netmap
[head->network*]% commit

Cloudbursting: How to configure Cluster Extension to AWS without OpenVPN (7.3+)

Other Configuration

At this point you should be able to successfully power on the cloud director, and have it go through the node-installer phase.

Depending on your network config (if the VPC is reached via the external network), and your security policies, you might want to reconfigure the head node to autosign certificates sign requests which come in via the external network (from the VPC). That's because by default all cert sign requests comming in via external network need to by confirmed manually.

```
[pw-trunk-openstack->network[externalnet]]% show
Parameter                      Value

------------------------------ ------------------------------------
-----------
Allow autosign                 Always
```

You might also need to configure the management network in the base partition to be the external network so that when the node-installer falls back to using the base partition management network it will be able to reach the head node.

Unique solution ID: #1325
Author: Piotr Wachowicz
Last update: 2017-08-28 17:08