CDP One Prerequisites

Date published: 2022-06-03
Date modified: 2022-08-15

https://docs.cloudera.com/

Legal Notice

© Cloudera Inc. 2024. All rights reserved.

The documentation is and contains Cloudera proprietary information protected by copyright and other intellectual property rights. No license under copyright or any other intellectual property right is granted herein.

Unless otherwise noted, scripts and sample code are licensed under the Apache License, Version 2.0.

Copyright information for Cloudera software may be found within the documentation accompanying each component in a particular release.

Cloudera software includes software from various open source or other third party projects, and may be released under the Apache Software License 2.0 (“ASLv2”), the Affero General Public License version 3 (AGPLv3), or other license terms. Other software included may be released under the terms of alternative open source licenses. Please review the license and notice files accompanying the software for additional licensing information.

Please visit the Cloudera software product page for more information on Cloudera software. For more information on Cloudera support services, please visit either the Support or Sales page. Feel free to contact us directly to discuss your specific needs.

Cloudera reserves the right to change any products at any time, and without notice. Cloudera assumes no responsibility nor liability arising from the use of products, except as expressly agreed to in writing by Cloudera.

Cloudera, Cloudera Altus, HUE, Impala, Cloudera Impala, and other Cloudera marks are registered or unregistered trademarks in the United States and other countries. All other trademarks are the property of their respective owners.

Disclaimer: EXCEPT AS EXPRESSLY PROVIDED IN A WRITTEN AGREEMENT WITH CLOUDERA, CLOUDERA DOES NOT MAKE NOR GIVE ANY REPRESENTATION, WARRANTY, NOR COVENANT OF ANY KIND, WHETHER EXPRESS OR IMPLIED, IN CONNECTION WITH CLOUDERA TECHNOLOGY OR RELATED SUPPORT PROVIDED IN CONNECTION THEREWITH. CLOUDERA DOES NOT WARRANT THAT CLOUDERA PRODUCTS NOR SOFTWARE WILL OPERATE UNINTERRUPTED NOR THAT IT WILL BE FREE FROM DEFECTS NOR ERRORS, THAT IT WILL PROTECT YOUR DATA FROM LOSS, CORRUPTION NOR UNAVAILABILITY, NOR THAT IT WILL MEET ALL OF CUSTOMER’S BUSINESS REQUIREMENTS. WITHOUT LIMITING THE FOREGOING, AND TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, CLOUDERA EXPRESSLY DISCLAIMS ANY AND ALL IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, QUALITY, NON-INFRINGEMENT, TITLE, AND FITNESS FOR A PARTICULAR PURPOSE AND ANY REPRESENTATION, WARRANTY, OR COVENANT BASED ON COURSE OF DEALING OR USAGE IN TRADE.

Contents

Configuring a site-to-site IPSec VPN to CDP One

This guide describes how to establish a site-to-site IPSec VPN (Virtual Private Network) connection from your firewall to CDP One. These steps are written for generic firewall configuration, so the concepts apply to all firewalls, but the configuration may vary depending on your exact firewall.

About this task

This configuration consists of a single tunnel using static routes. If you have any issues with the settings below, please create a case with Cloudera Support.

IPSec Tunnel Configuration

1. Internet Key Exchange (IKE) Configuration

   Configure the IKE SA as follows:

   - IKE version: IKEv2
   - Authentication Method: Pre-Shared Key
   - Pre-Shared Key: <xxxx >
   - Authentication Algorithm: sha256
   - Encryption Algorithm: aes-256-cbc
   - Lifetime: 24 hours
   - Phase 1 Negotiation Mode: main
   - Diffie-Hellman: Group 14
2. IPSec Configuration

   Configure the IPSec SA as follows:

   - Protocol: esp
   - Authentication Algorithm: sha256
   - Encryption Algorithm: aes-256-cbc or aes-256-gcm
   - Lifetime: 8 hours
   - Mode: tunnel
   - Perfect Forward Secrecy: Diffie-Hellman Group 14
3. Tunnel Interface Configuration

   Outside IP Addresses:

   - Cloudera VPN Gateway: <x.x.x.x>
   - Customer Gateway: (customer provided)
4. Routing Configuration

   Static routes will be used within the tunnel to send traffic between CDP One and your firewall.

Create a Security Policy Rule Set

1. Outbound Rule – This will allow customer enterprise access to CDP One.
   a. Create a new policy, and for the following services:
      1. Knox: Proxy gateway for access to cluster services (TCP 443)
      2. SSH: (TCP 22)

   b. Add a second outbound Policy for ping (testing only).

2. Inbound Rule (optional) – This will allow CDP One to access specific enterprise resources.
    - Create a new Security Policy and add services to allow CDP One to access specific enterprise data endpoints for providing data into the service, or for the service to push data into.

Configuring gateway nodes

This guide describes how to configure gateway nodes to be compatible with CDP One automation.

About this task

Gateway nodes are the primary mechanism to interact with the command line utilities for CDP One endpoints. The command line utilities are most commonly used when automating processes. They are commonly used for orchestration purposes.

CDP One automation controls gateway nodes. It is important that the gateway nodes are used in a manner that does not interfere with this automation.

Cloudera strongly recommends that Gateway Nodes are not used for any data persistence. All data should be persisted to one of the core storage mechanisms that support redundancy and resilience.

Installing .cer certs

Copy certs to /opt/apps/customer/certs, and run the import_customer_certs command to add the cert to the default truststore. Any .cer files in this location are preserved by automation. After upgrades, automation automatically imports all .cer files from this location to the default truststore.

Installing additional OS packages

There are orchestration scenarios that require additional OS packages to be installed to enable an orchestration process. In these scenarios;

- Install packages using yum, and be sure to record the package information (name and version) in the /opt/apps/customer/meta/yum_packages.txt** file.
- After upgrades, automation reads this metadata file and reinstalls all of the listed packages.

Installing Python packages

- Use venv to create a virtual environment under the /opt/apps/customer/ directory, then install Python packages in this virtual environment.
- For more information about venv, see Installing packages using pip and virtual environments.

Storing data files

Use the /opt/apps/, /home, or /data directories to store data files. Data stored outside these locations will not be persisted through the upgrade process.

Using a DNS configuration file

Configure DNS forwarders using /etc/unbound/conf.d/customer_dns.conf. Automated backup and restore ensures that this file is backed up before running the CDP OS upgrade, and restored after upgrade.

Using a sudoers configuration file

Configure customer sudoers groups using /etc/sudoers.d/customer_sudoers. Automated backup and restore ensures that this file is backed up before running the CDP OS upgrade, and restored after upgrade.