

## AWS Reference Network Architecture

Date published: 2019-08-22

Date modified: 2025-08-18

# CLOUdera

# Legal Notice

© Cloudera Inc. 2025. 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

<b>Cloudera reference network architecture on AWS.....</b>	<b>4</b>
<b>Taxonomy of network architectures.....</b>	<b>5</b>
Cloudera Management Console to customer cloud network.....	6
Customer on-prem network to cloud network.....	8
<b>Network architecture.....</b>	<b>8</b>
Architecture diagrams.....	8
Component description.....	11
VPC.....	11
Subnets.....	11
Gateways and route tables.....	12
Security groups.....	12
DNS.....	14
DHCP option set.....	16
Determining the CIDR range.....	16
Existing VPC and subnets.....	16
DNS.....	19
Associating additional CIDRs to a VPC.....	21
Cloudera Private Links Network for AWS.....	22
Supported service components.....	24
Setting up Cloudera Private Links Network for AWS environments.....	24
Troubleshooting Cloudera Private Links Network.....	30
References.....	30

# Cloudera reference network architecture on AWS

This topic includes a conceptual overview of the Cloudera on cloud network architecture for AWS, its use cases, and personas who should be using it

## Overview

Cloudera on cloud allows customers to set up cloud Data Lakes and compute workloads in their cloud accounts on AWS, Azure, and Google Cloud. It maps a cloud account to a concept called the environment into which all compute workload clusters (Cloudera Data Hub clusters) and data services (such as Cloudera Data Engineering, Cloudera Data Warehouse, Cloudera AI, Cloudera Operational Database, Cloudera DataFlow) are launched. For these Data Lakes, compute workload clusters, and data services to function correctly, several elements of the cloud architecture need to be configured appropriately: access permissions, networking setup, cloud storage and so on. Broadly, these elements can be configured in one of two ways:

- Cloudera can set up these elements for the customer

Usually, this model helps to set up a working environment quickly and try out Cloudera. However, many enterprise customers prefer or even mandate specific configurations of a cloud environment for Infosec or compliance reasons. Setting up elements such as networking and cloud storage requires prior approvals and they would generally not prefer, or even actively prevent, a third party vendor like Cloudera to set up these elements automatically.

- Cloudera can work with pre-created elements provided by the customer

In this model, the flow for creating the cloud Data Lakes accepts pre-created configurations of the cloud environment and launches workloads within those boundaries. This model is clearly more aligned with enterprise requirements. However, it brings with it the risk that the configuration might not necessarily play well with Cloudera requirements. As a result, customers might face issues launching Cloudera workloads and the turnaround time to get to a working environment might be much longer and involve many tedious interactions between Cloudera and the customer cloud teams.

From our experience in working with several enterprise customers, the most complicated element of the cloud environment setup is the cloud network configuration. The purpose of this document is to clearly articulate the networking requirements needed for setting up a functional Cloudera environment on cloud into which the Data Lakes and compute workloads of different types can be launched. It attempts to establish the different points of access to these workloads and establishes how the given architecture helps to accomplish this access.

Along with this document, you can use the [cloudera-deploy tool](#) to automatically set up a model of this reference architecture, which can then be reviewed for security and compliance purposes.

## Use cases

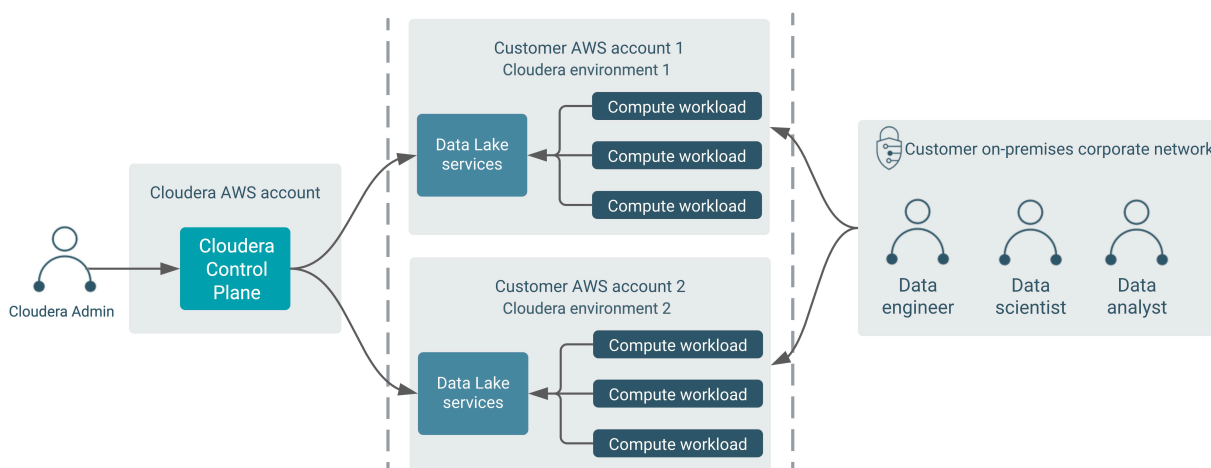
Cloudera on cloud allows customers to process data in the cloud storage under a secure and governed Data Lake using different types of compute workloads that are provisioned via Data Hub or data services. Typically the lifecycle of these workloads is as follows:

- A Cloudera environment is set up by a Cloudera admin using their cloud account. This sets up a cloud Data Lake cluster and FreeIPA cluster with security and governance services and an identity provider for this environment. The Cloudera admin may need to work with a cloud administrator to create all the cloud provider resources (including networking resources) that are required by Cloudera.
- Then one or more compute workload clusters can be launched, linked to the Data Lake. Each of these workload clusters typically serves a specific purpose such as data ingestion, analytics, machine learning and so on.
- These compute workload clusters are accessed by data consumers like data engineers, analysts or scientists. This is the core purpose of using Cloudera on cloud.
- These compute workload clusters can be long-running or ephemeral, depending on the customer needs.

There are two types of users for Cloudera who interact with the product for different purposes:

- **Cloudera admins** - These persons are usually concerned with the launch and maintenance of the cloud environment, and the Data Lake, Cloudera Data Hub clusters, FreeIPA, and Cloudera data services running inside the environment. They use a Cloudera Management Console running in the Cloudera AWS account to perform these operations of managing the environment.
- **Data consumers** - These are the data scientists, data analysts, and data engineers who use the Data Hubs and data services to process data. They mostly interact directly with the compute workloads (Cloudera Data Hub clusters and data services) running in their cloud account. They could access these either from their corporate networks (typically through a VPN) or other cloud networks their corporate owns.

These two types of users and their interaction with Cloudera are represented in the following diagram:



## Taxonomy of network architectures

This topic provides a high-level overview of each type of network architecture that Cloudera supports.

At a high level, there are several types of network architectures Cloudera supports. As can be expected, each type brings a unique trade-off among various aspects, such as ease of setup, security provided, workloads supported, and so on. This section only provides a high level overview of each type. The characteristics of each type are explained under appropriate sections in the rest of the document. The users must review the advantages and disadvantages of each of these taxonomies in detail before making a choice suitable to their needs.

Name	Description	Trade-offs
Publicly accessible networks	Deploys customer workloads to hosts with public IP addresses. Security groups must be used to restrict access only to corporate networks as needed.	Easy to set up for POCs. Low security levels.
Semi-private networks	Deploys customer workloads to private subnets, but exposes services to which data consumers need access over a load balancer with a public IP address. Security groups or allow-lists (of IP addresses or ranges) on load balancers must be used to restrict access to these public services only to corporate networks as needed.	This option is fairly easy to set up too, but it may not solve all the use cases of access (in semi private networks). The surface of exposure is reduced, and it is reasonably secure.
Fully private networks	Deploys customer workloads to private subnets and even services to which data consumers need access are only on private IPs. Requires connectivity to corporate networks to be provided using solutions like VPN gateways, and so on.	Complex to set up depending on prior experience of establishing such connectivity, primarily due to the way the customer has to solve the corporate network peering problem. But it is very secure.

Name	Description	Trade-offs
Fully private outbound restricted networks	This is the same as fully private networks; Except, in addition, Cloudera provides a mechanism for users to configure an outbound proxy or firewall to monitor or restrict the communication outside their networks.	Most complex to set up, mainly considering the varied needs that data consumers would have to connect outside the VPC on an evolving basis. It is also the most secure for an enterprise.

## Cloudera Management Console to customer cloud network

This topic explains the possible ways in which Cloudera Control Plane can communicate with the compute infrastructure in the customer network, in the context of the Cloudera Management Console.

As described previously, the Cloudera admin would typically use the Cloudera Management Console that runs in the Cloudera Control Plane to launch Cloudera environments with Data Lakes, FreeIPA, Cloudera Data Hub clusters, and data services into their cloud accounts. In order to accomplish this, the Cloudera Control Plane and the compute infrastructure in the customer network (such as EC2 instances, EKS clusters) should be able to communicate with each other. Depending on the chosen network architecture, this communication can occur in the ways described below.

### Publicly accessible networks

In this model of publicly accessible networks, the compute infrastructure must be reachable over the public internet from the Cloudera Management Console. While this is fairly easy to set up, it is usually not preferred by enterprise customers, as it implies that the EC2 nodes or EKS nodes are assigned public IP addresses. While the access control rules for these nodes can still be restricted to the IP addresses of the Cloudera Management Console components, it is still considered insecure for each of the network architectures described earlier.

### Semi-private networks

Publicly accessible networks are easy to set up for connectivity, both from the Cloudera Control Plane and the customer on-prem network, but have a large surface area of exposure as all compute infrastructure has public IP addresses. In contrast, fully private networks need special configuration to enable connectivity from the customer on-prem network, due to having no surface area of exposure to any of the compute infrastructure. While very secure, it is more complex to establish.

There is a third configuration supported by Cloudera, semi-private networks, that provides some trade-offs between these two options. In this configuration, the user deploys the worker nodes of the compute infrastructure on fully private networks as described above. However, the user chooses to expose UIs or APIs of the services fronting these worker nodes over a public network load balancer. By using this capability, the data consumers can access the UIs or APIs of the compute infrastructure through these load balancers. It is also possible to restrict the IP ranges from which such access is allowed using security groups.

While this option provides a trade-off between ease of setup and exposure levels, it may not satisfy all use cases related to communication between various endpoints. For example, some compute workloads involving Kafka or NiFi would not benefit from having a simple publicly exposed NLB. It is recommended that customers evaluate their use cases against the trade-off and choose an appropriately convenient and secure model of setup.

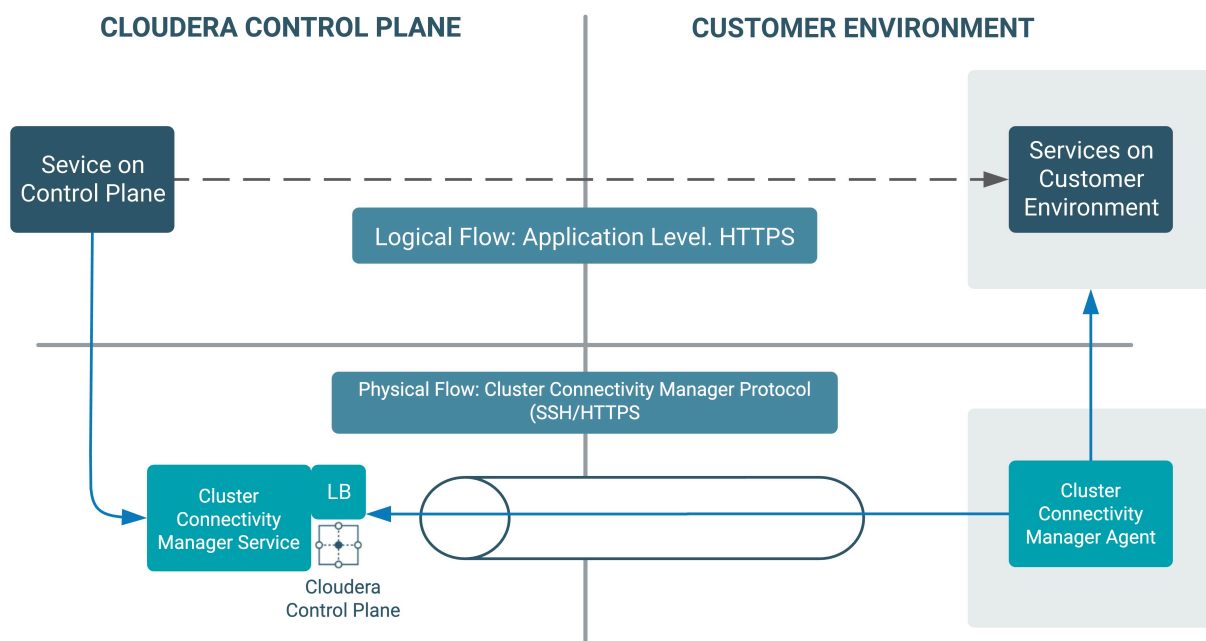
### Fully private networks

In this model of fully private networks, the compute infrastructure is not assigned any public IP addresses. In this case, communication between the Cloudera Control Plane and compute infrastructure is established using a "tunnel" that originates from the customer network to the Cloudera Control Plane. All communication from the Cloudera Control Plane to the compute nodes is then passed through this tunnel. From experience, Cloudera has determined that this is the preferred model of communication for customers.

To elaborate on the tunneling approach, Cloudera uses a solution called [Cluster Connectivity Manager](#). At a high level, the solution uses two components, an agent (Cluster Connectivity Manager agent) that runs on a VM provisioned in the customer network and a service (Cluster Connectivity Manager service) that runs on the Cloudera

Control Plane. The Cluster Connectivity Manager agent, at start-up time, establishes a connection with the Cluster Connectivity Manager service. This connection forms the tunnel. This tunnel is secured by asymmetric encryption. The private key is shared with the agent over cloud specific initialization mechanisms, such as a user-data script in AWS.

When any service on the Cloudera Control Plane wants to send a request to a service deployed on the customer environment (depicted in the below diagram as the “logical flow”), it physically sends a request to the Cluster Connectivity Manager service running in the Cloudera Control Plane. The Cluster Connectivity Manager agent and Cluster Connectivity Manager service collaborate over the established tunnel to accept the request, forward it to the appropriate service, and send a response over the tunnel to be handed over the calling service on the Cloudera Control Plane.



Currently, all EKS clusters provisioned by various Cloudera data services are enabled with public and private cluster endpoints even under fully private network setup (see [Amazon EKS cluster endpoint access control](#)). The EKS public endpoint is needed to facilitate the interactions between Cloudera Control Plane and the EKS cluster while worker nodes and Kubernetes control plane interact over private API endpoints. There are plans to support private EKS endpoints in the future. When this occurs, the documentation will be updated to reflect the same.

### Fully private outbound restricted networks

Fully private outbound restricted networks is a variant of the fully private network where customers would like to pass outbound traffic originating from their cloud account through a proxy or firewall and explicitly allow-list URLs that are allowed to pass through. Cloudera on cloud supports such configuration. If such network architecture is chosen, the customer must ensure the following:

- Users configure a proxy for the environment via Cloudera, as documented in [Use a non-transparent proxy with Cloudera Data Warehouse on AWS environments](#) for Cloudera Data Warehouse and [Using a non-transparent proxy](#) for all other compute workloads and the Data Lake itself.
- Compute resources (such as VMs used by Cloudera Data Hub clusters and data services) can connect to the proxy or firewall via appropriate routing rules.
- The proxy or firewall is set up to allow connections to all hosts, IP ranges, ports, and protocol types that are documented in [Outbound network access destinations for AWS](#).

**Note:**

Given that fully private networks is the recommended option of connectivity in most cases, this document describes the architecture assuming a fully private network setup.

## Customer on-prem network to cloud network

After compute workload clusters are launched in the customer's cloud network, data consumers such as data engineers, data scientists, and data analysts access services running in these Cloudera data services. Sometimes, Cloudera admins who set up and operate these clusters might need this access to diagnose any issues the clusters face.

Examples of these include:

- Web UIs such as:
  - Hue: For running SQL queries in Hive tables
  - Cloudera AI Workbenches: For accessing Cloudera AI projects, models, notebooks, and so on
  - Cloudera Manager: For Cloudera Data Hub clusters and Data Lakes
  - Atlas and Ranger: For metadata, governance, and security in the Data Lake
- JDBC endpoints: Customers can connect tools such as Tableau using a JDBC URL pointing to the Hive server.
- SSH access: Data engineers might log in to nodes on the compute Cloudera data services to run data processing jobs using YARN, Spark, or other data pipeline tools.
- Kube API access: Cloudera data services that run on Amazon EKS (such as Cloudera Data Warehouse and Cloudera AI) also provide admin access to Kubernetes for purposes of diagnosing issues.
- API access: Customers can use APIs for accessing many of the services exposed via the web UIs for purposes of automation and integration with other tools, applications, or other workloads they have. For example, Cloudera AI exposes the “Cloudera AI API v2” to work with Cloudera AI projects and other entities. See [Cloudera AI API v2](#).

These services are accessed by these consumers from within a corporate network inside a VPN. These services typically have endpoints that have a DNS name, the format of which is described more completely in the DNS section of this reference architecture documentation. These DNS names resolve to IP addresses assigned to the nodes, or load balancers fronting the ingest controllers of Kubernetes clusters. Note that these IP addresses are usually private IPs; Therefore, in order to be able to connect to these IPs from the on-premise network within a VPN, some special connectivity setup would be needed, typically accomplished using technologies like VPN peering, DirectConnect, transit gateways, and so on. While there are many options possible here, this document describes one concrete option of achieving this connectivity.

## Network architecture

Cloudera recommends that customers configure their cloud networks as fully private networks, as described in this chapter. This will help on-boarding Data Lakes, Cloudera Data Hub clusters, and data services smoothly.



**Note:** This network architecture only covers the fully private networks and assumes unrestricted outbound access.

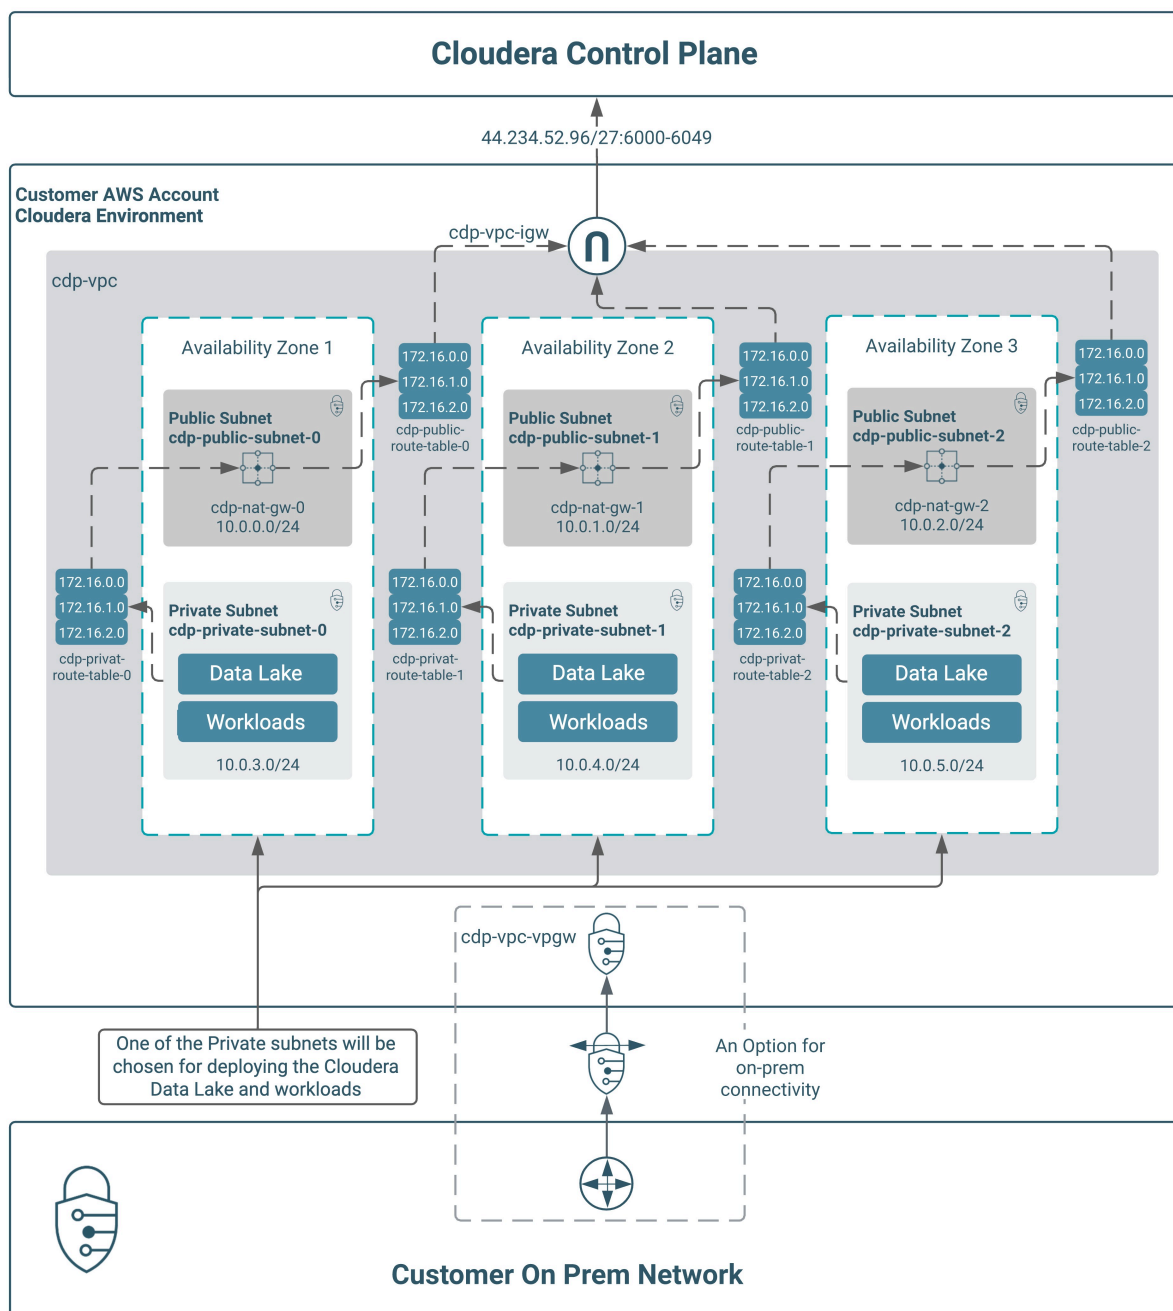
The `cdctl` tool, which is released along with this document can be used to automatically set up a model of this reference architecture, which can then be reviewed for security and compliance purposes.

## Architecture diagrams

This topic includes diagrams illustrating the various elements of the network architecture in the customer's cloud account into which Cloudera environments with Data Lakes, Cloudera Data Hub clusters, and data services will be launched.



Cloudera recommends that customers configure their cloud networks as described in this chapter and illustrated in the following diagrams. This will help onboarding Data Lakes, Cloudera Data Hub clusters, and data services smoothly. The following diagram illustrates the configuration for a fully private network that can be configured by the customer. This configuration can be provided by the Cloudera admins when they are setting up Cloudera environments and workloads which will get launched into this configuration.

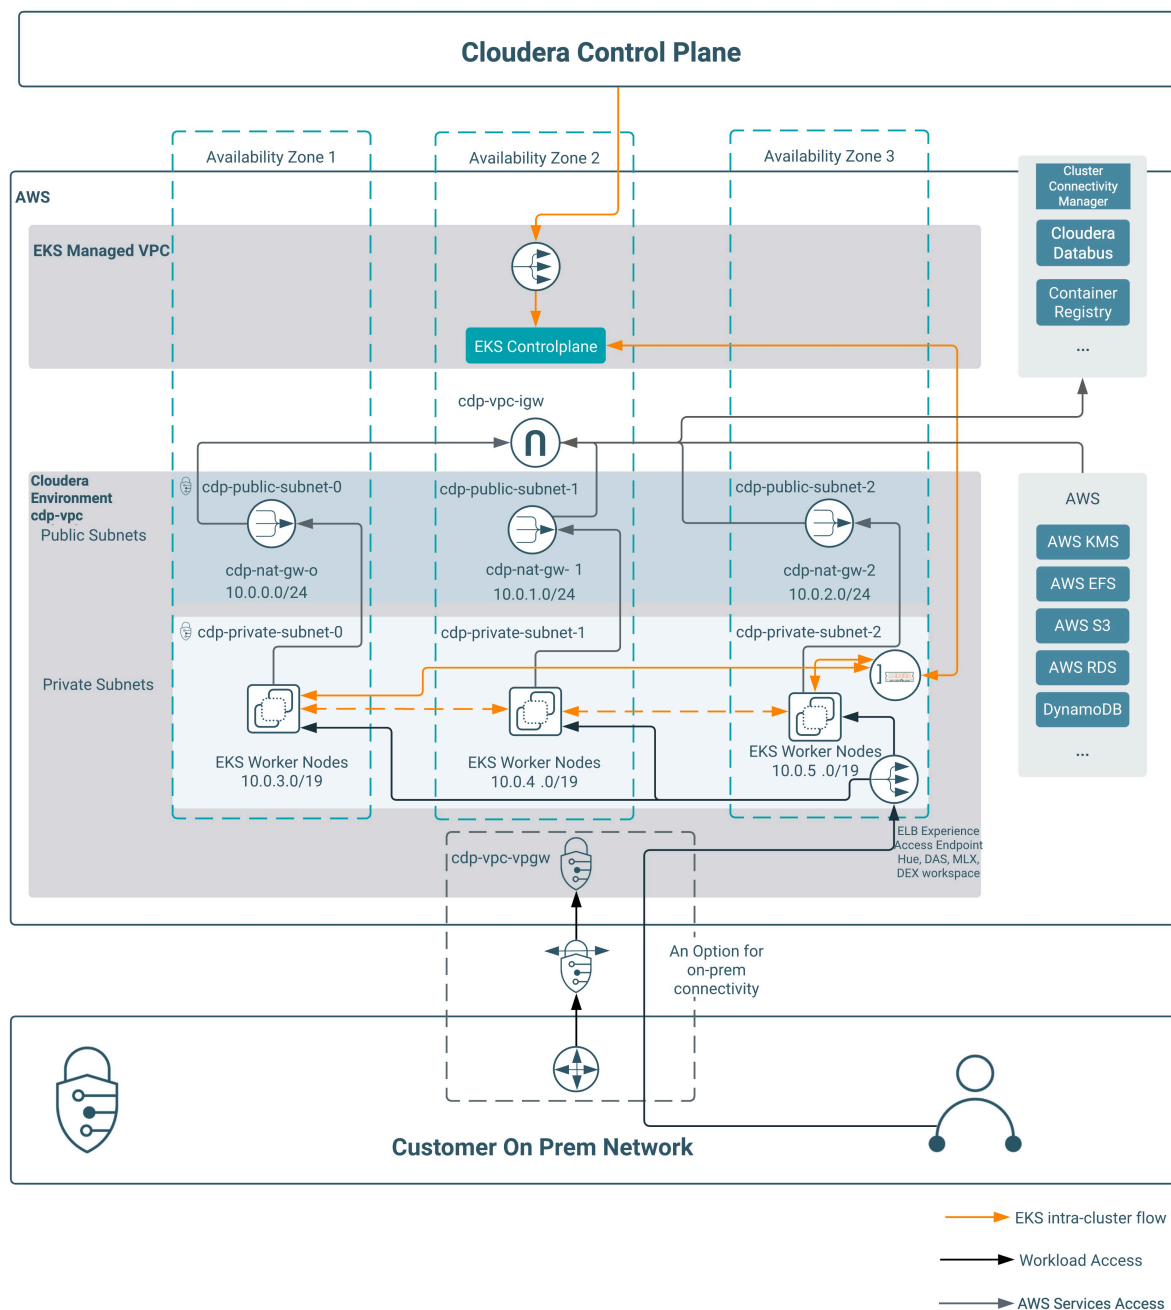


Note the following points about the architecture:

- The configuration is a fully private network configuration - that is, the workloads are launched on nodes that do not have public IP addresses into a private subnet.
- They connect outbound to the Cloudera Control Plane over a fixed IP and port range.
- For users to be able to connect from the customer on-prem network to the Cloudera workloads in the private subnet, some network connectivity setup is required. In this case, a customer's VPN server peered to an AWS virtual private gateway is shown.

Some of the Cloudera data services are based on AWS EKS clusters. Amazon EKS manages the Kubernetes Control Plane while the worker nodes that make up the cluster get provisioned in the customer's VPC. The EKS Control Plane has an API endpoint for administrative purposes which is commonly referred to as "cluster endpoint". The Cloudera data service itself is accessible through a service endpoint ELB.

This is illustrated in the following diagram:



As can be seen in the above diagram, Cloudera workloads have dependencies on some AWS cloud services such as RDS, EFS and so on. A full list of these services, described in the context of each workload is specified in [AWS outbound network access destinations](#).

In the chapters that follow, we detail the elements of this architecture, including specifying the configuration and options in each of the components.

## Component description

This topic provides an overview of the VPC, subnets, gateways and route tables, and security groups required for Cloudera on AWS.

### VPC

An Amazon Virtual Private Cloud (VPC) is needed for deploying Cloudera workloads into the customer's cloud account. Cloudera recommend that the VPC used for Cloudera is configured with properties as specified in this topic.

- The CIDR block for the VPC should be sufficiently large for supporting all the Cloudera Data Hub clusters and data services that you intend to run. Refer to [Determining the CIDR range](#) for understanding how to compute the CIDR block range.
- The VPC properties for DNS hostnames and DNS resolution must be ENABLED. DNS resolution lets Kubernetes pods resolve external host names and also to support DNS hostnames. The DNS hostnames option needs to be enabled as several Cloudera data services rely on EFS (see [Mounting on Amazon EC2 with a DNS name](#)). Enabling these properties is also a requirement (see [Amazon EKS cluster endpoint access control](#)) to enable private access of EKS cluster endpoints.
- VPCs are associated with a DHCP Option Set. The DHCP option set for the VPCs must be set up as per the section described in [DHCP option set](#).

### Subnets

A subnet is a partition of the virtual network in which Cloudera workloads are launched.

It is recommended that the subnets be configured with the following properties:

- It is recommended to have 3 private subnets and 3 public subnets, such that each private-public subnet pair is in a different availability zone (AZ). Even if a region has two AZs instead of three, it's recommended that three private subnets are created, two in the same AZ. This is required to prevent cross AZ routing of traffic and to maintain Quorum-based consistency required by some services.
  - Note that a subnet becomes 'private' or 'public' based on the routing devices it is associated with in the route tables. This is described in [Gateways and route tables](#).
  - The private subnets will be where the compute workloads will be launched by Cloudera. This ensures that these nodes are working in an isolated and secure environment that does not have internet connectivity.
  - The public subnet is needed to host a NAT gateway as this will allow the compute nodes to reach out to the Cloudera Control Plane over the internet. More on this will be described in [Gateways and route tables](#).
- The CIDR block for the subnets should be sufficiently large for supporting all the Cloudera data services you intend to run. Refer to [Determining the CIDR range](#) for understanding how to compute the CIDR block range.
- The CIDR block for the subnets should not overlap with known [AWS EKS ranges for pods/services](#). Several EKS based Cloudera data services in Overlay networks
- In addition, you may want to ensure that the CIDR ranges assigned to the Subnets will not overlap with any of your on-premise network CIDR ranges, as this may be a requirement for setting up connectivity from your on-prem network to the subnets.
- Since Cloudera recommends fully private network configuration, the 'Auto-assign public IPs' option must be disabled for the private subnets.
- A subnet can be associated with a Network ACL (NACL). However, since Cloudera works with a fully private network configuration where communication is always initiated from EC2 nodes within the subnets, a NACL is generally not useful for this configuration.
- Tag private subnets with a tag 'kubernetes.io/role/internal-elb:1'. The key is the string and the value is '1'. Cloud Controller Manager and AWS Load Balancer Controller both require private subnets to have this tag for automatic creation of private ELBs. Private ELBs are created in these subnets by EKS. This is applicable when Cloudera is supporting EKS versions < 1.20 (which is currently the case). See [How can I tag the Amazon VPC subnets in my Amazon EKS cluster](#).

## Gateways and route tables

This topic covers recommended gateway and route table configurations for Cloudera on AWS.

### Connectivity from Cloudera Control Plane to Cloudera workloads

- As described in [Taxonomy of network architectures](#), nodes in the Cloudera workloads need to connect to the Cloudera Control Plane over the internet to establish a ‘tunnel’ over which the Cloudera Control Plane can send instructions to the workloads.
- In order to accomplish this, there are two gateways that need to be configured - a NAT Gateway in each of the public subnets and an Internet Gateway at the VPC level.
- The private subnet hosting the Cloudera workloads should be configured with a route table where the default route (0.0.0.0/0) points to a NAT Gateway in the public subnet of its AZ.
- The public subnet hosting the NAT Gateway should be configured with a route table where the default route (0.0.0.0/0) points to an Internet Gateway the VPC is configured with.
- Each NAT gateway requires an elastic IP address. The VPC should contain as many elastic IP addresses as NAT gateways across the AZs in the VPC.

### Connectivity from customer on-prem to Cloudera workloads

- As described in [Use cases](#), data consumers need to access data processing or consumption services in the Cloudera workloads. Given these are created with private IP addresses in private subnets, the customers will need to arrange for access to these addresses from their on-prem or corporate networks in specific ways.
- There are several possible solutions for achieving this, but one that is depicted in the [Architecture diagram](#), uses a AWS VPN Gateway service.
- In this solution, the customer has to create a Virtual Private Gateway, and connect it to the VPN service on the on-prem network.

## Security groups

During the specification of a VPC to Cloudera, the Cloudera admin specifies the security groups that will be associated with all the Cloudera workloads launched within that VPC. These security groups will be used in allowing the incoming traffic to the hosts.

### Security groups for Data Lakes and Cloudera Data Hub clusters

During the specification of a VPC to Cloudera, the Cloudera admin can either let Cloudera create security groups, taking a list of IP address CIDRs as input; or create them in AWS and then provide them to Cloudera.

When getting started with Cloudera, the Cloudera admin can let Cloudera create security groups, taking a list of IP address CIDRs as input. These will be used in allowing the incoming traffic to the hosts. The list of CIDR ranges should correspond to the address ranges from which the Cloudera workloads will be accessed. In a VPN-peered VPC, this would also include address ranges from customer’s on-prem network. This model is useful for initial testing given the ease of set up.

Alternatively, the Cloudera admin can create security groups on their own and select them during the setup of the VPC and other network configuration. This model is better for production workloads, as it allows for greater control in the hands of the Cloudera admin. However, note that the Cloudera admin must ensure that the rules meet the requirements described below.

For a fully private network, security groups should be configured according to the types of access requirements needed by the different services in the workloads:

- Services accessed only within the VPC must be configured with the following inbound rules:
  - All TCP / UDP / ICMP access is allowed for the CIDRs corresponding to the VPC.
  - Conversely, there is no need to provide any access to these services for any IP CIDRs outside the VPC.

- Endpoint services are the services that can be accessed outside the VPC through the gateway, chiefly by data consumers or Cloudera admins. For example, UIs like Hue, Atlas, Ranger, Cloudera Manager all need to be accessed by data consumers or other administrators. For enabling this, the following in-bound rules are set up:
  - All TCP / UDP / ICMP access is allowed for the CIDRs corresponding to the VPC.
  - All TCP ports that correspond to services like Kafka, HBase, and so on that need to be accessed outside the VPC are to be allowed for the list of CIDR ranges specified at the time of creating the environment/Cloudera data service, or in the security group created by the Cloudera admin. Alternatively, all TCP / UDP / ICMP access may be allowed for these CIDR ranges.
  - SSH access is allowed for the CIDR ranges specified at the time of creating the environment.
  - HTTPS access is allowed for the CIDR ranges specified at the time of creating the environment.
- Note that for a fully private network, even specifying an open access here (such as 0.0.0.0/0) is restrictive because these services are deployed in a private subnet without a public IP address and hence do not have a route to the Internet gateway. However, the list of CIDR ranges may be useful to restrict which private subnets of the customer's on-prem network can access the services. Rules for EKS based workloads are described in the following section.
- Rules for AKS based workloads are described separately in the following section.

### Additional rules for EKS-based workloads

At the time of enabling a Cloudera data service, the Cloudera admin can specify a list of CIDR ranges that will be used in allowing the incoming traffic to the workload Elastic Load Balancer (ELB). This list of CIDR ranges should correspond to the address ranges from which the Cloudera data service workloads will be accessed. In a VPN peered VPC, this would also include address ranges from customer's on-prem network. In a fully private network setup, 0.0.0.0/0 implies access only within the VPC and the peered VPN network which is still restrictive.

Since a public endpoint is enabled by default for all EKS cluster Control Planes at the moment, it is highly recommended to provide a list of outbound public CIDR ranges at the time of provisioning a Cloudera data service to restrict access to the EKS clusters. By default, the public endpoint is always allowed to connect to the Cloudera public CIDR range. The following screenshot is an example configuration section for a Cloudera data service:

☒ Restrict access to Kubernetes API server to authorized IP ranges ⓘ

Restricting access to the API server without providing any authorized IP ranges will result in inaccessible API server. Please specify authorized IP ranges in CIDR notation if you would like to restrict access to the API server to specific IP ranges.

API server Authorized IP Ranges ⓘ

-
+

Specific guidelines for restricting access to Kubernetes API server and workloads are detailed in [Restricting access for Cloudera services that create their own security groups on AWS](#) by each Cloudera data service.

Within the EKS cluster, there are several security groups defined to facilitate EKS control plane-pod communication, inter-pod and inter-worker node communication as well as workload communication through ELBs. These groups are in accordance with AWS documentation (see [Amazon EKS security group considerations](#).)

### Outbound connectivity requirements

Outbound traffic from the worker nodes is unrestricted and is targeted at other AWS services and Cloudera services. The comprehensive list of services that get accessed from a Cloudera environment can be found in AWS documentation (see [Amazon EKS security group considerations](#)).

## DNS

This topic covers recommended DNS configurations for Cloudera on AWS.

The previous sections dealt with how connectivity is established to the workload infrastructure. This section deals with ‘addressability’. The workloads launched by Cloudera contain a few services that need to be accessed by the Cloudera admins or data consumers. These include services like Cloudera Manager, metadata services like the Hive Metastore, Atlas or Ranger, data processing or consumption services such as Oozie server, Hue, and so on. Given the nature of the cloud infrastructure, the IP addresses for the nodes running these services may change (for example, if the infrastructure is restarted or repaired). However, these should have statically addressable DNS names so that users can access them with the same names.

In order to help with this, Cloudera assigns DNS names to these nodes. These naming schemes have the following properties:

- The DNS name is of the following format for each Data Lake node, Cloudera Data Hub node, and the Data Lake/Cloudera Data Hub cluster endpoint: <CLUSTER\_NAME>-{<HOST\_GROUP><i></i></i>}.<ENVIRONMENT\_IDENTIFIER>.<CUSTOMER\_IDENTIFIER>.cloudera.site

An example could be my-dataeng-master0.my-envir.aaaa-1234.cloudera.site

This name has the following components:

- The base domain is cloudera.site. This is a publicly registered DNS suffix (see [Public Suffix List](#)). It is also a registered Route53 hosted zone in a Cloudera owned AWS account.
- The <CUSTOMER\_IDENTIFIER> is unique to a customer account on Cloudera made of alphanumeric characters and "-"
- The <ENVIRONMENT\_IDENTIFIER> is generated based on the environment name and is truncated to 8 characters.
- The <CLUSTER\_NAME> is the cluster name given to the Data Lake or Cloudera Data Hub. It is appended with a <HOST\_GROUP> name such as "gateway", "master", "worker", and so on, depending on the role that the node plays in the cluster. If there are more than one of these nodes playing the same role, they are appended with a serial number, <i></i>.



- The DNS name of the endpoints of the Cloudera data services is of the following format:
  - For a Virtual Warehouse in Cloudera Data Warehouse, it is <VIRTUAL\_WAREHOUSE\_NAME>.<CDW\_ENVIRONMENT\_IDENTIFIERS>.dw.<CUSTOMER\_IDENTIFIERS>.cloudera.site
    - The <VIRTUAL\_WAREHOUSE\_NAME> is the name of the Virtual Warehouse created. There could be multiple virtual warehouses for a given Cloudera environment.
    - The <CDW\_ENVIRONMENT\_IDENTIFIERS> is the identifier for the Cloudera environment.
  - For a Session Terminal in a Cloudera AI workspace, it is <TTY\_SESSION\_ID>.<CML\_WORKSPACE\_ID>.<ENVIRONMENT\_IDENTIFIERS>.<CUSTOMER\_IDENTIFIERS>.cloudera.site
    - The <TTY\_SESSION\_ID> is the ID of the Cloudera AI Terminal Session ID.
    - The <CML\_WORKSPACE\_ID> is the ID of the Cloudera AI workspace created.
    - The <ENVIRONMENT\_IDENTIFIERS> is generated based on the environment name and is truncated to 8 characters. If the 8th character is a "-" (dash), then it is truncated to 7 characters instead.
  - For all the Cloudera data services listed above, the common portions of the DNS include.
    - The base domain is cloudera.site. This is a publicly registered DNS suffix (see [Public Suffix List](#)). It is also a registered Route53 hosted zone in a Cloudera owned AWS account.
    - The <CUSTOMER\_IDENTIFIERS> is unique to a customer account on Cloudera made of alphanumeric characters and a "-" (dash).
  - For a virtual cluster in Cloudera Data Engineering, it is <VIRTUAL\_CLUSTER\_ID>.<CDE\_SERVICE\_ID>.<ENVIRONMENT\_IDENTIFIERS>.<CUSTOMER\_IDENTIFIERS>.cloudera.site
    - The <VIRTUAL\_CLUSTER\_ID> is the 8-character ID of the Cloudera Data Engineering virtual cluster, for example, afg57p98.
    - The <CDE\_SERVICE\_ID> is the ID of the Cloudera Data Engineering service containing the virtual cluster, for example, cde-g6th4kqv.
    - The <ENVIRONMENT\_IDENTIFIERS> is generated based on the Cloudera environment name and is truncated to 8 characters. if the 8th character is a "-" (dash), then it is truncated to 7 characters instead.
  - For a DataFlow service in Cloudera DataFlow, it is
   
dfx.<CDF\_WORKLOAD\_ENDPOINT\_ID>.<CUSTOMER\_IDENTIFIERS>.cloudera.site
    - The <CDF\_WORKLOAD\_ENDPOINT\_ID> is the 8-character ID of the Cloudera DataFlow Service Workload Endpoint, for example, 1bxt50kk.
  - For a database in Cloudera Operational Database, it is <COD\_WORKLOAD\_NAME>-{<HOST\_GROUP>}<ENVIRONMENT\_IDENTIFIERS>.<CUSTOMER\_IDENTIFIERS>.cloudera.site
    - <COD\_Workload\_Name> is the ID of the Cloudera Operational Database, for example, cod-1m6yz9uwqhrq2.
    - The user provides a database name and the environment where they want to create the database. These two entities are hashed together to create the internal <COD\_WORKLOAD\_NAME>, which is set as the Cloudera Data Hub cluster.
    - Except for the <COD\_WORKLOAD\_NAME>, the rest of the DNS name of the endpoint is implemented as per Cloudera Data Hub DNS format as mentioned above.
  - For all the Cloudera data services listed above, the common portions of the DNS include:
    - The base domain is cloudera.site. This is a publicly registered [DNS suffix](#). It is also a registered Route53 hosted zone in a Cloudera owned AWS account.
    - The <CUSTOMER\_IDENTIFIERS> is unique to a customer account on Cloudera made of alphanumeric characters and a "-" (dash).
- The length of the DNS name is restricted to 64 characters due to some limitations with Hue workloads.
- These names are stored as A records in the Route53 hosted zone in the Cloudera managed Cloudera Control Plane AWS account. Hence, you can resolve these names from any location outside of the VPC. However, note that they would still resolve to private IP addresses and hence are constrained by the connectivity setup described in preceding sections.

- Within a Cloudera environment, the DNS resolution happens differently. Every Cloudera environment has a DNS server that is played by a component called FreeIPA. This server is seeded with the hostnames of the nodes of all workload clusters in the environment. Every node in a Data Lake, Cloudera Data Hub, and a Cloudera data service is configured to look up the FreeIPA DNS service for name resolution within the cluster.

## DHCP option set

This topic covers recommended DHCP Option Set configurations for Cloudera on AWS.

- Ensure the DHCP Option Set has the DNS Server set to AmazonProvidedDNS. This is required for EFS (see [Mounting on Amazon EC2 with a DNS name](#)) and EKS (see [Amazon EKS cluster endpoint access control](#)).
- Customers may choose to use their own DNS servers from their corporate network in the DHCP option set. However, the DNS servers need to be configured to conditionally forward DNS queries for the Amazon subdomains to Amazon Provided DNS resolver in the VPC. This also helps ensure the Cloudera connectivity to AWS services. If the Cloudera data services do not allow activation with the custom DNS servers, please contact Cloudera support to get the relevant entitlements added.
- Ensure the DHCP Option Set associated with the VPC has only one domain and it is recommended that it remains default.

## Determining the CIDR range

This topic covers options for determining the CIDR range for Cloudera on AWS. When registering an AWS environment in Cloudera, you are asked to select a VPC and one or more subnets. You can calculate and verify the limits of the VPC and subnets available in your AWS subscription to ensure that you have enough networking resources to create clusters in Cloudera.

## Existing VPC and subnets

When you use an existing VPC, the subnet requirements vary based on the services used. This section is a guide for calculating network requirements per service.

### Subnets for Data Lake and Cloudera Data Hub

Both Data Lake and Cloudera Data Hub share the same subnet, but due to Data Lake's dependency on RDS, Cloudera admins need at least two subnets.

It is recommended to use subnets of /19 CIDR. If you would like to use a smaller subnets, use the following guidelines:

- One IP address is used for each VM.
- Two IPs for your RDS instances.
- One Light Duty Data Lake cluster uses 2 VMs and 1 additional VMs for FreeIPA.
- One Medium Duty Data Lake cluster uses 10 VMs and 3 additional VMs for FreeIPA.
- The exact number of VMs depends on the Cloudera Data Hub's cluster definition, but a typical Cloudera Data Hub cluster uses a minimum of four VMs as a starting point and this number can be dynamically scaled up or down.
- Make sure you allocate enough IPs to handle each cluster running at peak capacity

### Subnets for Cloudera Data Warehouse

The Cloudera Data Warehouse service needs three subnets. You can choose the specific subnets that should be used when you activate Cloudera Data Warehouse for an environment and if the specified subnets are shared with other Cloudera data services then ensure that there are sufficient IPs left for Cloudera Data Warehouse.

An environment can be activated for Cloudera Data Warehouse with Overlay Network or without Overlay Network, the next section outlines the benefits and the pitfalls of one choice over the other.



## Using Cloudera Data Warehouse with Overlay Network vs without Overlay Network

Enabling overlay networks creates two network spaces in your EKS:

- A node network space, which derives per-node IP addresses from the VPC.
- A Kubernetes pod network space, which derives per-pod IP addresses from the CNI plugin's own network space.

The overlay network is bridged into the node network. As a result, one IP address is required per node instead of one IP address needed per pod. It is recommended to enable overlay networks with Cloudera Data Warehouse if the available IPs in your subnets is less than 1024.

Even though using an overlay network for Cloudera Data Warehouse requires less number of IPs it adds more metadata to each network packet which even though are processed in the kernel can have a performance impact. Due to this extra metadata added to each network packet it may not be straightforward to debug incoming and outgoing traffic with some of the network debugging tools. To employ an overlay network for Cloudera Data Warehouse, Weave net CNI (Container Network Interface) plugin is used. Since weave-net is a full-mesh network where each node is connected to every other node in the mesh, the number of nodes the overlay network can span across is limited to 200.

### Cloudera Data Warehouse with Overlay Networks

Use these guidelines to arrive at your desired subnet CIDR if you choose to use overlay networks. It is recommended to enable overlay networks with Cloudera Data Warehouse if the available IPs in your subnets is less than 1024.

VM type	No of VMs	Total no of IPs addresses required
Cloudera Data Warehouse Shared Services - (Shared among all VWs in an environment)	3	3
Per Database Catalog (One catalog is created by default, you can create additional catalogs)	2	2
Shared Services per Virtual Warehouse (HS2, Hue, DAS, coordinators, catalog, statestored, and so on.)	3	3
Per Virtual Warehouse (XS) *	2 to 20	2 to 20
Per Virtual Warehouse (S) *	10 to 100	10 to 100
Per Virtual Warehouse (M) *	20 to 200	20 to 200
Per Virtual Warehouse (L) *	40 to 400	40 to 400
Per Virtual Warehouse (Custom)*	x to 10x ( where x is the initial node count )	x to 10x ( where x is the initial node count)

\* Each autoscaling activity can be treated as deploying a new Virtual Warehouse. For example, when a XS Virtual Warehouse is scaled once, it uses four VMs instead of two.

### Cloudera Data Warehouse without Overlay Network

Use these guidelines to arrive at your desired subnet CIDR if you choose to not to use overlay networks while activating your DW environment.

VM type	No of VMs	Maximum no of pods per VM	Maximum number of IPs per VM (No of pods per VM +1 per VM)	Maximum no of IPs addresses required
Cloudera Data Warehouse Shared Services - (Shared among all VWs in an environment)	3	25	26	78
Per Database Catalog (One catalog is created by default, you can create additional catalogs)	2	25	26	52

VM type	No of VMs	Maximum no of pods per VM	Maximum number of IPs per VM (No of pods per VM +1 per VM)	Maximum no of IPs addresses required
Shared Services per Virtual Warehouse (HS2, Hue, DAS, coordinators, catalog, statestored, and so on.)	3	25	26	78
Per Virtual Warehouse (XS) *	2 to 20	10	11	22 to 220
Per Virtual Warehouse (S) *	10 to 100	10	11	110 to 1100
Per Virtual Warehouse (M) *	20 to 200	10	11	220 to 2200
Per Virtual Warehouse (L) *	40 to 400	10	11	440 to 4400
Per Virtual Warehouse (Custom)*	x to 10 x ( where x is the initial node count )	10	11	11x to 110x

\* Each autoscaling activity can be treated as deploying a new Virtual Warehouse. For example, when a XS Virtual Warehouse is scaled once, it uses four VMs instead of two.

### Query Isolation

If the Query Isolation feature has been enabled for a Virtual Warehouse and a query scans more than the threshold set in the `hive.query.isolation.scan.size.threshold` parameter, the planner runs the query in isolation. This means that an isolated standalone executor group is spawned to run the data-intensive query. The number of executors spawned to run the query does not exceed the default setting for the `hive.query.isolation.max.nodes.per.query` parameter, which default to 2 times the virtual warehouse template size. The number of isolated parallel queries (can go up to 400) and number of nodes per isolated queries (can go up to 400) is configurable for a given virtual warehouse. Each of these nodes can consume up to 11 IPs when overlay networks are not enabled and 1 IP if the overlay network is enabled against your environment.

### Subnets for Cloudera AI

Cloudera AI requires at least two subnets in two different availability zones and you can choose which subnets should be used by a workbench at the time of provisioning. If the specified subnets are shared with other Cloudera data services then ensure that there are sufficient IPs left for Cloudera AI. Cloudera AI uses Calico CNI (Container Network Interface) to run Cloudera AI pods in an overlay network. The formula to calculate IP Addresses per workbench is as follows:

- Each workbench can grow up to 100 CPU worker nodes and 100 GPU workers; each node consumes 1 IP address.
- In addition, you will need to allocate 9 IP addresses for infrastructure nodes (3 IPs for Cloudera AI infra nodes, 4 for Liftie infra nodes and 2 load balancers).

### Subnets for Cloudera Data Engineering

Cloudera Data Engineering requires at least two subnets in two different availability zones. You choose at the time of workspace provisioning which subnets should be used. Ensure that there are adequate IPs left for Cloudera Data Engineering to provide for the maximum expected size of the cluster. Cloudera Data Engineering uses Calico CNI (Container Network Interface) to run pods in an overlay network. A /24 CIDR is recommended for Cloudera Data Engineering subnets, but for a custom range the formula to calculate IP addresses per Cloudera Data Engineering service is as follows:

- Each Cloudera Data Engineering service can scale up to 100 compute nodes, each node consumes one IP address.
- In addition, you need to allocate 5 IP addresses for the infrastructure nodes (1 IP for Cloudera Data Engineering infra node, 4 for Liftie infra nodes) and 2 IP addresses per virtual cluster for the virtual cluster service nodes.

### Subnets for Cloudera DataFlow

Cloudera DataFlow requires at least two subnets in two different availability zones. Cloudera DataFlow by default configures EKS to run in private subnets, if they are available. Cloudera DataFlow uses Calico CNI (Container Network Interface) to run pods in an overlay network. The CIDR block for the subnets must be sized appropriately in each Cloudera DataFlow environment to accommodate the following:

- Each Cloudera DataFlow cluster can grow up to 50 autoscaling compute instances, each of which consumes 1 IP address.
- A fixed overhead of 48 IP addresses for three instances for core Cloudera DataFlow services.

### Subnets for Cloudera Operational Database

Cloudera Operational Database currently leverages Cloudera Data Hub to deploy infrastructure in a private subnet. Clients associated with HBase REST server, Thrift Server, or the Phoenix Thin server can be proxied via the VPN gateway. If Apache HBase Java API or Apache Phoenix Thick JDBC client are used, an edge node must be configured to access the private computing resources. The formula to calculate IP addresses per Cloudera Operational Database database is as follows:

- Each Cloudera Operational Database database defaults to a minimum of 9 nodes (1 leader, 2 master, 1 gateway, and 5 worker nodes), requiring 9 IP addresses.
- Each Cloudera Operational Database database can autoscale the number of worker nodes given various factors/attributes. The range defaults to a minimum of 5 nodes and a maximum of 20 nodes. Users can reduce the minimum to 3 nodes, but there is currently no fixed upper limit. They must consider the behaviour of their database while providing the CIDR range to accommodate for the potential autoscaling growth, where each node takes 1 IP address.
- In addition, configuring an edge node, if required for client applications, takes 1 IP address for 1 accessory node.

## DNS

This topic covers recommended DNS configurations for Cloudera on AWS.

The previous sections dealt with how connectivity is established to the workload infrastructure. This section deals with ‘addressability’. The workloads launched by Cloudera contain a few services that need to be accessed by the Cloudera admins or data consumers. These include services like Cloudera Manager, metadata services like the Hive Metastore, Atlas or Ranger, data processing or consumption services such as Oozie server, Hue, and so on. Given the nature of the cloud infrastructure, the IP addresses for the nodes running these services may change (for example, if the infrastructure is restarted or repaired). However, these should have statically addressable DNS names so that users can access them with the same names.

In order to help with this, Cloudera assigns DNS names to these nodes. These naming schemes have the following properties:

- The DNS name is of the following format for each Data Lake node, Cloudera Data Hub node, and the Data Lake/Cloudera Data Hub cluster endpoint: `<CLUSTER_NAME>-<HOST_GROUP><i><ENVIRONMENT_IDENTIFIER>.<CUSTOMER_IDENTIFIER>.cloudera.site`

An example could be `my-dataeng-master0.my-envir.aaaa-1234.cloudera.site`

This name has the following components:

- The base domain is `cloudera.site`. This is a publicly registered DNS suffix (see [Public Suffix List](#)). It is also a registered Route53 hosted zone in a Cloudera owned AWS account.
- The `<CUSTOMER_IDENTIFIER>` is unique to a customer account on Cloudera made of alphanumeric characters and “-”
- The `<ENVIRONMENT_IDENTIFIER>` is generated based on the environment name and is truncated to 8 characters.
- The `<CLUSTER_NAME>` is the cluster name given to the Data Lake or Cloudera Data Hub. It is appended with a `<HOST_GROUP>` name such as “gateway”, “master”, “worker”, and so on, depending on the role that the node plays in the cluster. If there are more than one of these nodes playing the same role, they are appended with a serial number, `<i>`.

- The DNS name of the endpoints of the Cloudera data services is of the following format:
  - For a Virtual Warehouse in Cloudera Data Warehouse, it is <VIRTUAL\_WAREHOUSE\_NAME>.<CDW\_ENVIRONMENT\_IDENTIFIERS>.dw.<CUSTOMER\_IDENTIFIERS>.cloudera.site
    - The <VIRTUAL\_WAREHOUSE\_NAME> is the name of the Virtual Warehouse created. There could be multiple virtual warehouses for a given Cloudera environment.
    - The <CDW\_ENVIRONMENT\_IDENTIFIERS> is the identifier for the Cloudera environment.
  - For a Session Terminal in a Cloudera AI workspace, it is <TTY\_SESSION\_ID>.<CML\_WORKSPACE\_ID>.<ENVIRONMENT\_IDENTIFIERS>.<CUSTOMER\_IDENTIFIERS>.cloudera.site
    - The <TTY\_SESSION\_ID> is the ID of the Cloudera AI Terminal Session ID.
    - The <CML\_WORKSPACE\_ID> is the ID of the Cloudera AI workspace created.
    - The <ENVIRONMENT\_IDENTIFIERS> is generated based on the environment name and is truncated to 8 characters. If the 8th character is a "-" (dash), then it is truncated to 7 characters instead.
  - For all the Cloudera data services listed above, the common portions of the DNS include.
    - The base domain is cloudera.site. This is a publicly registered DNS suffix (see [Public Suffix List](#)). It is also a registered Route53 hosted zone in a Cloudera owned AWS account.
    - The <CUSTOMER\_IDENTIFIERS> is unique to a customer account on Cloudera made of alphanumeric characters and a "-" (dash).
  - For a virtual cluster in Cloudera Data Engineering, it is <VIRTUAL\_CLUSTER\_ID>.<CDE\_SERVICE\_ID>.<ENVIRONMENT\_IDENTIFIERS>.<CUSTOMER\_IDENTIFIERS>.cloudera.site
    - The <VIRTUAL\_CLUSTER\_ID> is the 8-character ID of the Cloudera Data Engineering virtual cluster, for example, afg57p98.
    - The <CDE\_SERVICE\_ID> is the ID of the Cloudera Data Engineering service containing the virtual cluster, for example, cde-g6th4kqv.
    - The <ENVIRONMENT\_IDENTIFIERS> is generated based on the Cloudera environment name and is truncated to 8 characters. if the 8th character is a "-" (dash), then it is truncated to 7 characters instead.
  - For a DataFlow service in Cloudera DataFlow, it is
   
dfx.<CDF\_WORKLOAD\_ENDPOINT\_ID>.<CUSTOMER\_IDENTIFIERS>.cloudera.site
    - The <CDF\_WORKLOAD\_ENDPOINT\_ID> is the 8-character ID of the Cloudera DataFlow Service Workload Endpoint, for example, 1bxt50kk.
  - For a database in Cloudera Operational Database, it is <COD\_WORKLOAD\_NAME>-{<HOST\_GROUP>}<ENVIRONMENT\_IDENTIFIERS>.<CUSTOMER\_IDENTIFIERS>.cloudera.site
    - <COD\_Workload\_Name> is the ID of the Cloudera Operational Database, for example, cod-1m6yz9uwqhrq2.
    - The user provides a database name and the environment where they want to create the database. These two entities are hashed together to create the internal <COD\_WORKLOAD\_NAME>, which is set as the Cloudera Data Hub cluster.
    - Except for the <COD\_WORKLOAD\_NAME>, the rest of the DNS name of the endpoint is implemented as per Cloudera Data Hub DNS format as mentioned above.
  - For all the Cloudera data services listed above, the common portions of the DNS include:
    - The base domain is cloudera.site. This is a publicly registered [DNS suffix](#). It is also a registered Route53 hosted zone in a Cloudera owned AWS account.
    - The <CUSTOMER\_IDENTIFIERS> is unique to a customer account on Cloudera made of alphanumeric characters and a "-" (dash).
- The length of the DNS name is restricted to 64 characters due to some limitations with Hue workloads.
- These names are stored as A records in the Route53 hosted zone in the Cloudera managed Cloudera Control Plane AWS account. Hence, you can resolve these names from any location outside of the VPC. However, note that they would still resolve to private IP addresses and hence are constrained by the connectivity setup described in preceding sections.

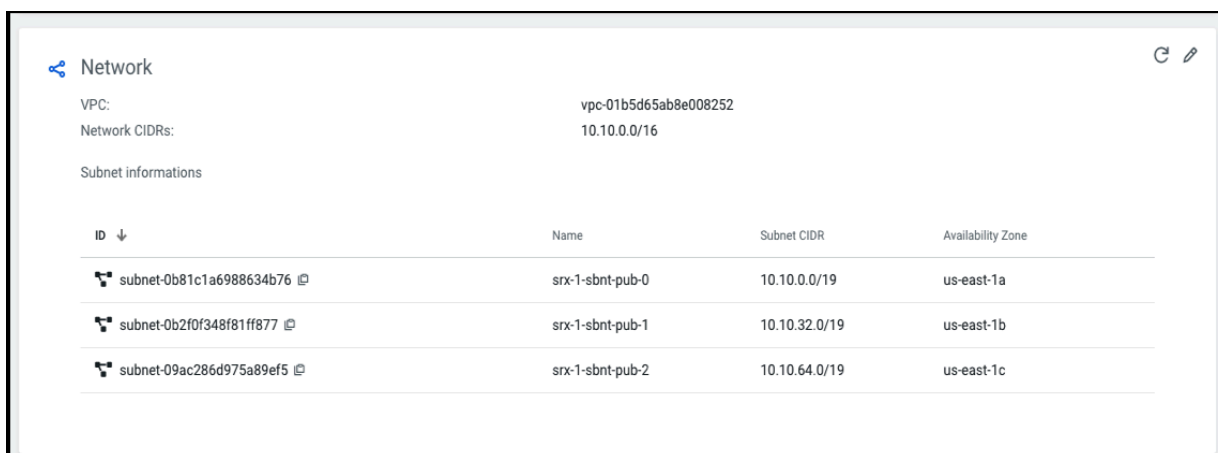
- Within a Cloudera environment, the DNS resolution happens differently. Every Cloudera environment has a DNS server that is played by a component called FreeIPA. This server is seeded with the hostnames of the nodes of all workload clusters in the environment. Every node in a Data Lake, Cloudera Data Hub, and a Cloudera data service is configured to look up the FreeIPA DNS service for name resolution within the cluster.

## Associating additional CIDRs to a VPC

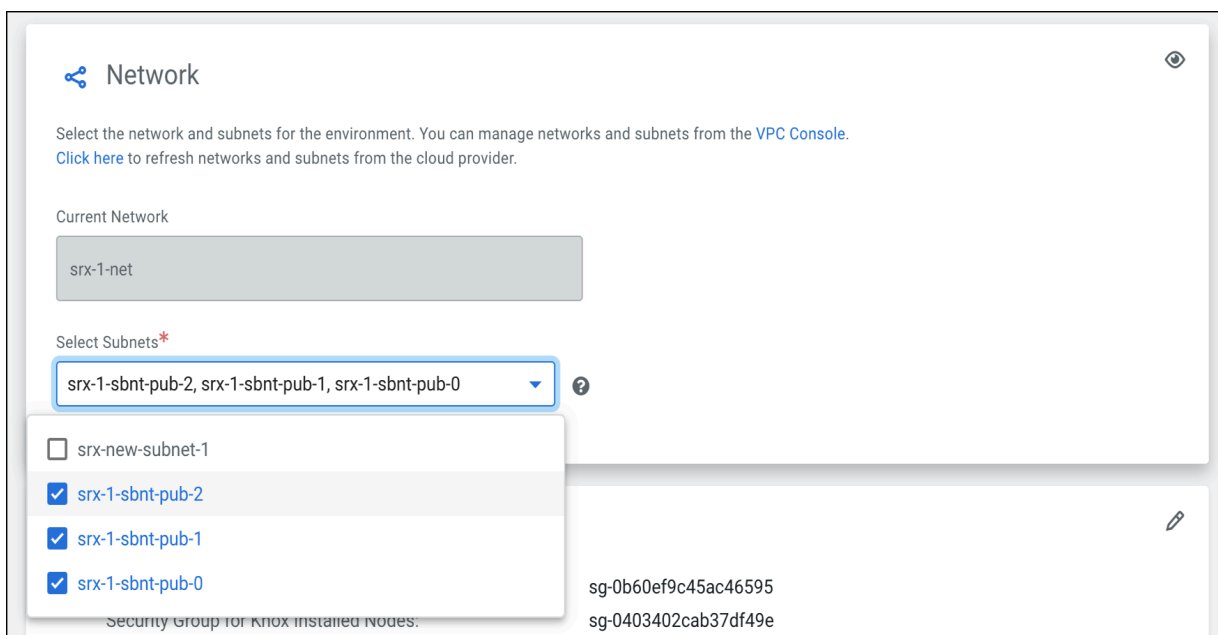
You can choose to add another CIDR to your VPC in case you are close to exhausting the available IPs in your VPC by following the necessary steps from [Associate a secondary IPv4 CIDR block with your VPC](#). Once a new CIDR has been added to your VPC, a created environment not will pick these changes automatically. For the environment to pick up the new CIDR, follow these steps.

### Procedure

1. Go to your environment page and under Summary tab click on the refresh icon in the Network card.



2. To add any new subnets created with the new CIDR, click on the pencil icon and choose the new subnet from the dropdown.



### Results

The user can now use the new subnet for any subsequent workloads that will be created.

## Cloudera Private Links Network for AWS

Cloudera Private Links Network enables you to connect privately and securely to the Cloudera Control Plane without traversing the internet. You can use Cloudera Private Links Network for end-to-end encryption of your workloads between Cloudera Control Plane and AWS VPC endpoints.

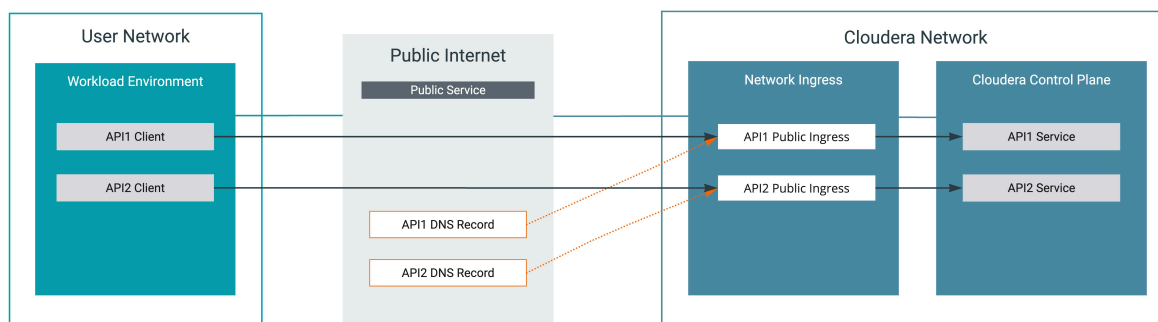
This documentation provides the following details and steps about Cloudera Private Links Network:

- High-level options of VPC endpoint placement
- Configuration options and other considerations for DNS overrides
- Cloudera Private Links Network deployment process
- Instructions of how to set up both Private Link options:
  - VPC: Setup of Cloudera Private Links Network for a workload VPC through CDP CLI
  - Authorization: Authorization with CDP CLI to enable the setup of Cloudera Private Links Network through your automation tools
- References for proxy profile configuration and considerations, and Cloudera Private Links Network commands

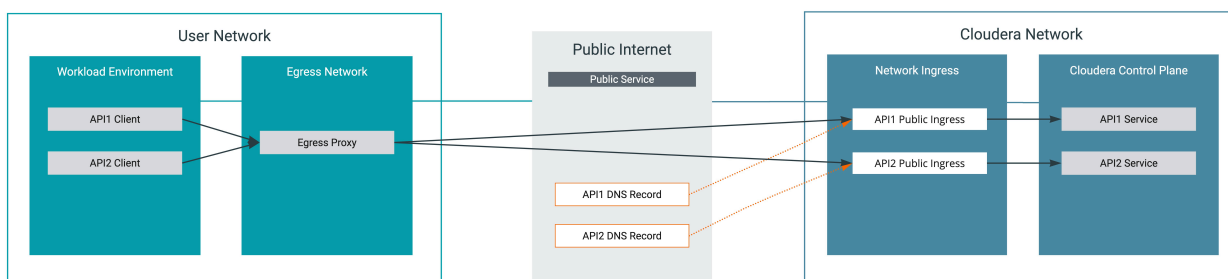
### Comparison of connectivity setup without and with Cloudera Private Links Network

Without Cloudera Private Links Network, your workload environment communicates with the Cloudera Control Plane through the internet. This traffic may optionally flow through a managed egress proxy. The following two diagrams illustrate this:

**Figure 1: Connectivity from workload environment to Cloudera Control Plane through the internet**



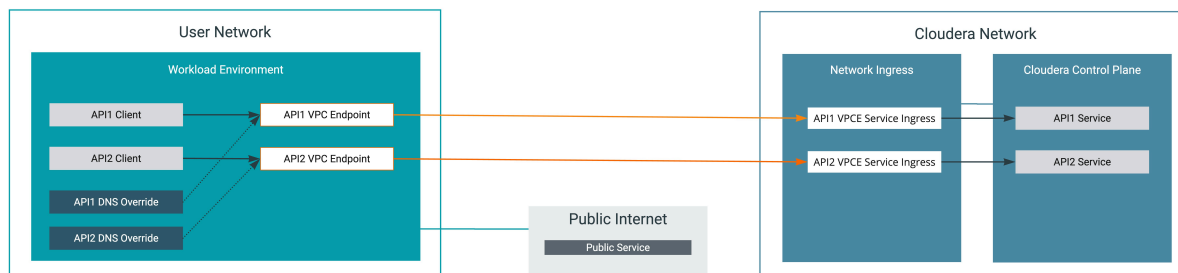
**Figure 2: Connectivity from workload environment to Cloudera Control Plane through the internet and egress proxy**



With Cloudera Private Links Network, the Cloudera Control Plane is accessed as if the Cloudera Control Plane would be on your network. This means that IP addresses are assigned to the Cloudera Control Plane services from your network, and DNS lookups will return your local IP addresses.

To ensure private connectivity through network ingress between the workload environment and Cloudera Control Plane, VPC endpoints can be added. The following illustration details the scenario where the VPC endpoints are in the same VPC as your workload environment. In this case, the VPC endpoints receive IPs from the workload environment VPC subnets:

**Figure 3: VPC endpoint in workload environment VPC**



The following options are available for DNS overrides:

- DNS is a per-VPC view: Deploying a VPC endpoint with private DNS option enabled will automatically install the DNS overrides in the local VPC.
- DNS is a regional or global view: Installing overrides at a regional or global scope will impact DNS resolution for other VPCs, other VPCs will attempt to use the VPC endpoints of the local VPC.

This section does not include an exhaustive list of design options, but should cover most cases. For more information about more advanced use cases, see the [Additional VPC scenarios](#) section.

Cloudera Private Links Network creates an additional layer of security with Private Links and security groups with inbound rule. If security groups are already present in your network, they will be updated with an inbound rule for safe HTTPs traffic. As inbound rules are created or updated within setting up the Cloudera Private Links Network, outbound rules can be removed.

### DNS overrides

The DNS hostnames for Cloudera services remain the same, whether the services are accessed through the internet or through VPC endpoints. DNS resolution in the customer networks perform DNS lookups against the public internet. Responses will contain the public internet IP addresses for the Cloudera services. DNS override records must be set up in your DNS infrastructure, and the hostnames need to be repointed to the VPC endpoint or IP addresses. For more information, see *Setting up DNS overrides*.

### Multi-region support

The AWS multi-region support for Private Links enables you to connect your VPC in one region to a VPC endpoint service in a different region.

This means that if a Cloudera endpoint service is not supported in your region, a fallback region is used to create the VPC endpoint. The fallback region is identified based on the Cloudera Control Plane region:

- US Cloudera Control Plane: us-west-2
- EU Cloudera Control Plane: eu-central-1
- AP Cloudera Control Plane: ap-southeast-2

For example, if you have a VPC hosted in the ca-central-1 region, and requested a Private Link connection for the Cluster Connectivity Manager service in the US Cloudera Control Plane, as Cluster Connectivity Manager is supported in the us-west-2 region, and not in ca-central-1, the connection is established using the us-west-2 region.



**Note:** Private Links can not be established for the Israel region (il-central-1) as this region does not have AWS multi-region support.



For more information about the AWS inter-region support, see [Introducing Cross-Region Connectivity for AWS PrivateLink](#).

## Supported service components

Review the list of supported service components.

Currently, Cloudera supports multiple service components for which Cloudera Private Links Network can be created:

- Cloudera API
- Databus API
- Cluster Connectivity Manager v2
- Console authentication

The following table describes the available service components for each region:

Cloudera Control Plane Region	VPCE service region in AWS	Available VPCE service components
us-west-1 (USA)	eu-west-1	API, DBUSAPI, CCMV2, CONSOLEAUTH
us-west-1 (USA)	ap-southeast-1	API, DBUSAPI, CCMV2, CONSOLEAUTH
us-west-1 (USA)	eu-central-1	API, DBUSAPI, CCMV2, CONSOLEAUTH
us-west-1 (USA)	us-east-1	API, DBUSAPI, CCMV2, CONSOLEAUTH
us-west-1 (USA)	us-east-2	API, DBUSAPI, CCMV2, CONSOLEAUTH
us-west-1 (USA)	us-west-2	API, DBUSAPI, CCMV2, CONSOLEAUTH
eu-1 (Germany)	eu-west-1	API, CCMV2
eu-1 (Germany)	eu-central-1	API, CCMV2
ap-1 (Australia)	ap-southeast-1	API, CCMV2
ap-1 (Australia)	ap-southeast-2	API, CCMV2



**Note:** In case you need to use a region that is not mentioned in the supported region table, reach out to Cloudera Support.

## Setting up Cloudera Private Links Network for AWS environments

Learn about how to set up Cloudera Private Links Network using CDP CLI for both the VPC and Authorization options.

You can configure private links from the workload environment to Cloudera Control Plane using Cloudera Private Links Network. You have the following options to choose from:

- VPC option: Using CDP CLI to set up private links in a Cloudera environment network. Cloudera automation assists you with authorizing and end-to-end setup. Requires cross-account permissions.
- Authorization option: Using your own network automation to create private links in a Cloudera environment. CDP CLI is used to authorize your cloud provider account to connect to Cloudera. Does not require cross-account permissions.

### Prerequisites

Before you begin, review the following prerequisites and requirements.

The following requirements must be met before creating Cloudera Private Links Network:

- Each environment must have its own private links configured.
- The private links should be set up before creating an environment.
- Configuring private links in Cloudera can only be performed through CDP CLI. You need to have CDP CLI installed and configured.
- In order to configure private links in Cloudera through CDP CLI, you must have the EnvironmentCreator or PowerUser role in Cloudera.



**AWS IAM requirements (VPC option only)**

You must have AWS admin privileges to create Cloudera Private Links Network with VPC option. You need to have the ability to update the cross-account access IAM role to include the policy required for private links.



**Note:** This is only required if you create private links with the VPC option. If you use your own network automation to create private links, you do not need to add this additional AWS policy.

Cloudera documents two types of policies: the default policy and the minimum access policies. If you are using the former, you do not need to do anything. If you are using the latter, add the following additional policy that will allow the creation of private links:

```
{
  "Sid": "ElasticComputeCloudFull",
  "Action": [
    "ec2:*"
  ],
  "Effect": "Allow",
  "Resource": [
    "*"
  ]
},
{
  "Sid": "IdentityAccessManagementFull",
  "Action": [
    "iam:AttachRolePolicy",
    "iam:CreateInstanceProfile",
    "iam:CreateRole",
    "iam>DeleteInstanceProfile",
    "iam>DeleteRole",
    "iam:ListRolePolicies",
    "iam:GetInstanceProfile",
    "iam:GetRolePolicy",
    "iam:ListAttachedRolePolicies",
    "iam:ListInstanceProfiles",
    "iam:PutRolePolicy",
    "iam:PassRole",
    "iam:GetRole",
    "iam:AddRoleToInstanceProfile",
    "iam:RemoveRoleFromInstanceProfile",
    "iam:DetachRolePolicy",
    "iam>DeleteRolePolicy",
    "iam:SimulatePrincipalPolicy",
    "iam:ListRoles"
  ],
  "Effect": "Allow",
  "Resource": [
    "*"
  ]
},
{
  "Sid": "IdentityAccessManagementLimited",
  "Action": [
    "iam:CreateServiceLinkedRole"
  ],
  "Effect": "Allow",
  "Resource": [
    "arn:aws:iam::*:role/aws-service-role/*"
  ]
}
```

```
{
```

```

        "Sid": "VisualEditor0",
        "Effect": "Allow",
        "Action": [
            "ec2:AuthorizeSecurityGroupIngress",
            "ec2:ModifyVpcEndpointServicePermissions",
            "ec2:ModifyVpcEndpointServiceConfiguration",
            "ec2:CreateVpcEndpointConnectionNotification",
            "ec2>DeleteVpcEndpoints",
            "ec2:CreateTags",
            "ec2:ModifySecurityGroupRules",
            "ec2:DescribeVpcEndpointServices",
            "ec2:DescribeVpcEndpointServiceConfigurations",
            "ec2>DeleteVpcEndpointServiceConfigurations",
            "ec2:ModifyVpcEndpointConnectionNotification",
            "ec2:DescribeSecurityGroups",
            "ec2:CreateVpcEndpointServiceConfiguration",
            "ec2:DescribeVpcEndpointServicePermissions",
            "ec2:CreateVpcEndpoint",
            "ec2:DescribeVpcs",
            "ec2:CreateSecurityGroup",
            "ec2>DeleteSecurityGroup",
            "route53:AssociateVPCWithHostedZone",
            "ec2:DescribeVpcEndpoints",
            "ec2:ModifyVpcEndpoint"
        ],
        "Resource": "*"
    },
}
]
}

```

This is only required for the Cloudera Private Links Network setup. Once Cloudera Private Links Network is created, you can revert back to the minimum access policy.

### Setting up DNS overrides

If your private DNS is disabled, you need to set up the DNS manually, adding the specified DNS entries for the respective service component.

### About this task



**Note:** If the workload VPC is the same as VPC where the endpoint is deployed, you need to update the existing endpoints with "enablePrivateDns": "true" through CDP CLI and the manual DNS override steps mentioned below are not required.

In the example below, a private hosted zone named cloudera.com is created for the VPC, and an alias record with Route traffic to the Cluster Connectivity Manager v2 VPC endpoint is created.

### Procedure

1. Create private hosted zones in Amazon Route 53 with the names below for the VPC endpoint region.

The following example shows how the endpoints should look like if the Cloudera Control Plane region us-west-1.

Service component	Endpoint
API	us-west-1.cdp.cloudera.com
CCMV2	v2.us-west-1.ccm.cdp.cloudera.com
DBUSAPI	us-west-1.sigma.altus.cloudera.com
CONSOLEAUTH	us-west-1.cdp.cloudera.com
MONITORING	api.monitoring.us-west-1.cdp.cloudera.com

- 2. Create a DNS record in Amazon Route 53 for each of the private hosted zones created in the previous step. You also need to add an alias record and route traffic to the respective endpoint.
  - a) For the API endpoint: \*.us-west-1.cdp.cloudera.com
  - b) For the CCMV2 endpoint: \*.v2.us-west-1.ccm.cdp.cloudera.com
  - c) For the DBUSAPI endpoint: dbusapi.us-west-1.sigma.altus.cloudera.com

The following is an example record routing traffic to the endpoint:

Services53 > Hosted zones > v2.ccm.cdp.mow-dev.cloudera.com > Create record

Create recordInfo

Quick create recordSwitch to wizard

▼ Record 1Delete

Record nameInfo

Record typeInfo

\*.v2.ccm.cdp.mow-dev.cloudera.com

A – Routes traffic to an IPv4 address and some AWS resources

Keep blank to create a record for the root domain.

☒ Alias

Route traffic toInfo

Alias to VPC endpoint

US West (Oregon) [us-west-2]

vpce-040ac0a895750edd0-7750qz5o.vpce-svc-06a1fcb259069ce83.us-west-2.vpce.amazonaws.com

Alias hosted zone ID: Z1YSA3EXCYU9Z

Routing policyInfo

Evaluate target health

Simple routing

☒ Yes

Add another record

CancelCreate records

**Note:** The VPC endpoint name to be routed is available in CDP CLI response of the respective endpoint.

Creating Cloudera Private Links Network with VPC option

Learn about how to create Cloudera Private Links Network using the VPC option.

Required Role: EnvironmentCreator or PowerUser

You need to use the following CDP CLI command to create the Cloudera Private Links Network with VPC option:

```
cdp cloudprivatelinks create-private-link-endpoint
```

This command is used to create a private link endpoint in your workload VPC. This establishes the private link connectivity between the VPC endpoint created and the respective VPC endpoint service present in the Cloudera Control Plane. The private link will be created for all supported components.

The following parameters should be specified:

Parameter	Description
enablePrivateDns	<div>Specifies whether private DNS needs to be enabled on the VPC endpoint.</div> <div><div></div><div><b>Note:</b> As mentioned in <a href="#">Setting up DNS override</a>, if you have a workload VPC that is different from the VPC where the endpoint is deployed, it is recommended that you set this value to false in the request and perform the DNS override.</div></div>

27

Parameter	Description
awsAccountDetails	<p>This should have the following fields</p> <ul style="list-style-type: none"> <li>cloudAccountId - Your AWS account ID where the private endpoints are created. The account ID needs to be provided, because Cloudera needs to authorize the account for private link service access.</li> <li>crossAccountRoleDetails or credentialCrn - The crossAccountRoleDetails can be provided with explicit crossAccountRole and externalId. You can also use credentialCrn that is configured with default policy or reduced access policy for provisioning the VPC endpoint.</li> <li>region - You can provide the AWS region where the private endpoint needs to be created.</li> <li>vpcId - The VPC ID in which private endpoint needs to be created</li> <li>subnets - IDs of the private subnets in which a private endpoint needs to be created</li> <li>subnetIds - It is recommended to provide the private subnet from each of the available zones. Without a specific subnet, there is no way to associate the private endpoint with your VPC resources. This is also required for fault tolerance and high availability.</li> </ul>



**Note:** Enabling private DNS provides DNS lookup of the hostname from within the VPC. You have the option to use endpoints outside of the VPC, and create your own DNS overrides. In this case, you need to disable private DNS.

The following is an example command for creating Cloudera Private Links Network:

```
cdp cloudprivatelinks create-private-link-endpoint --cloud-service-provider
aws
--cli-input-json '{
  "enablePrivateDns": true,
  "awsAccountDetails": {
    "cloudAccountId": "112313717721"
    "crossAccountRoleDetails": {
      "crossAccountRole":
"arn:aws:iam::112313717721:role/test-privatelinkservice-xaccount-role",
      "externalId": "51962fcf-4205-495d-ab4a-c82c5613a99f"
    },
    "region": "us-west-2",
    "vpcId": "vpc-0d006493309f23c3",
    "subnetIds": [
      "subnet-0e5812d895de60c47"
    ]
  }
}
```

The executed command performs the following sequence of steps:

1. Identifying the appropriate VPC endpoint service for the request. Existing vpc-svcs are filtered for the requested Cloudera service component and region.
2. Authorizing access to the VPC endpoint service is performed.
3. Creating a security group for the VPC endpoint with the provided ports.
4. Creating a VPC endpoint with the security group created and the VPC endpoint service filtered.

The command returns a trackingId, which can be used to verify that the domains of the respective Cloudera service components are reachable and resolve to private IPs from your VPC using the following command with the returned tracking ID from creating private link endpoints:

```
cdp cloudprivatelinks list-private-link-endpoint-statuses
--tracking-id [***ID***]
```

You can also check this by accessing the **Monitoring** tab of the endpoints in the AWS Console.

## Creating Cloudera Private Links Network with Authorization option

Learn about how to create Cloudera Private Links Network using the Authorization option.

Required Role: EnvironmentCreator or PowerUser

You need to use the following CLI command to create the Cloudera Private Links Network using your own network automation:

```
cdp cloudprivatelinks authorize-private-link-service-access
```

This command is used to authorize access to the Private Link services for your cloud account.

The following parameters should be specified:

Parameter	Description
cloudAccountId	Your AWS account ID where the private endpoints are created. The account ID needs to be provided, because Cloudera needs to authorize the account for private link service access.
region	Region of the Cloudera Control Plane.
serviceGroup	Name of the Cloudera service group.

The following is an example command for creating Cloudera Private Links Network:

```
cdp cloudprivatelinks authorize-private-link-service-access --cli-input-json
'{
  "cloudAccountId": "112313717721",
  "region": "us-west-2",
  "serviceGroup": "cdp-control-plane",
}'
```

The executed command performs the following sequence of steps:

1. Identifying the appropriate VPC endpoint service for the request. Existing vpc-svcs are filtered for the requested Cloudera service component and region.
2. Authorizing access to the VPC endpoint service is performed.

The command returns the Private Link service name, Cloudera service component and the authorization status as shown in the following example:

```
{
  "authorizePrivateLinkServiceAccessResults": [
    {
      "privateLinkService": "com.amazonaws.vpce.us-west-2.vpce-svc-0ce5a427a56d80603",
      "serviceComponent": "ccmv2",
      "authorizationStatus": "SUCCESS",
      "availabilityZoneList": "[usw2-az2, usw2-az1, usw2-az3]",
      "vpceClientTcpPortList": "[443]",
      "hostname": "[*.v2.ccm.cdp.mow-dev.cloudera.com]"
    }
  ]
}
```

After executing the command, you have to manually complete the following steps:

1. Create a security group for the VPC endpoint with the returned ports.  
For more information about how to create security groups, see the [AWS documentation](#).
2. Create a VPC endpoint with the previously created security group and the VPC endpoint service filtered.  
For more information about how to create VPC endpoints, see the [AWS documentation](#).

You can verify that the domains of the respective Cloudera service components are reachable and resolve to private IPs from your VPC using the following command with the returned tracking ID from creating private link endpoints:

```
cdp cloudprivatelinks list-private-link-endpoint-statuses
--tracking-id [***ID***]
```

You can also check this by accessing the **Monitoring** tab of the endpoints in the AWS Console.

### Deleting Cloudera Private Links Network

You have the option to delete the created Cloudera Private Links Network using CDP CLI. The command first deletes the security groups and then removes the private link endpoints.

Required Role: EnvironmentCreator or PowerUser



**Warning:** In case the Cloudera Private Links Network is deleted, traffic will be routed through the internet.

If you want to delete the Cloudera Private Links Network, use the following command:

```
cdp cloudprivatelinks delete-private-link-endpoint
--cloud-service-provider AWS
--aws-account-info [***ACCOUNT DETAILS***]
```

You can verify that the Cloudera Private Links Network is removed with the following command with the tracking ID returned from deleting the private link endpoints:

```
cdp cloudprivatelinks list-private-link-endpoint-statuses
--tracking-id [***ID***]
```

## Troubleshooting Cloudera Private Links Network

Learn more about how to troubleshoot certain errors when creating Cloudera Private Links Network.

The following error message is displayed if the environment creation is failing with a connectivity issue:

```
Please check your connection and proxy settings and make sure the instance can reach *.v2.ccm.cdp.cloudera.com
```

In this case, you can login to the FreeIPA instance, and confirm if the respective domain resolves to a private IP. The following domains are available for the Cloudera service component:

Service component	Endpoint
API	api.[***CLOUDERA CONTROL PLANE REGION***].cdp.cloudera.com
CCMv2	v2.ccm.[***CLOUDERA CONTROL PLANE REGION***].cdp.cloudera.com
DBUSAPI	dbusapi.[***CLOUDERA CONTROL PLANE REGION***].cdp.cloudera.com
CONSOLEAUTH	consoleauth.altus-[***ENVIRONMENT NAME***].cloudera.com
MONITORING	api.monitoring.[***CLOUDERA CONTROL PLANE REGION***].cdp.cloudera.com

The [\*\*\*CLOUDERA CONTROL PLANE REGION\*\*\*] should be one of the following:

- US control plane: us-west-1
- EU control plane: eu-1
- AP control plane: ap-1

## References

Learn more about additional information related to Cloudera Private Links Network.

## CLI commands for Cloudera Private Links Network

Learn more about the available CLI command for Cloudera Private Links Network.

Command	Description
create-private-link-endpoint	Creating Cloudera Private Links Network with VPC option.
authorize-private-link-service-access	Creating Cloudera Private Links Network with Authorization option. It returns the list of private link service names and their respective authorization status.
list-private-link-endpoint-statuses	This command can be used after creating and deleting Cloudera Private Links Network, or whenever there is a change to the existing setup.  The tracking-id returned from the create or delete command needs to be passed to get the status of request.
list-private-link-services-for-region	Returns a list of the private link services that are supported for the AWS region and the service group. It returns the list of private link service names and the mapped service component.  You can use this command to check which private link services are supported in your AWS region.
delete-private-link-endpoint	Deleting the previously created Cloudera Private Links Network. First the associated security group is deleted, and then deletes the specified endpoints are deleted.

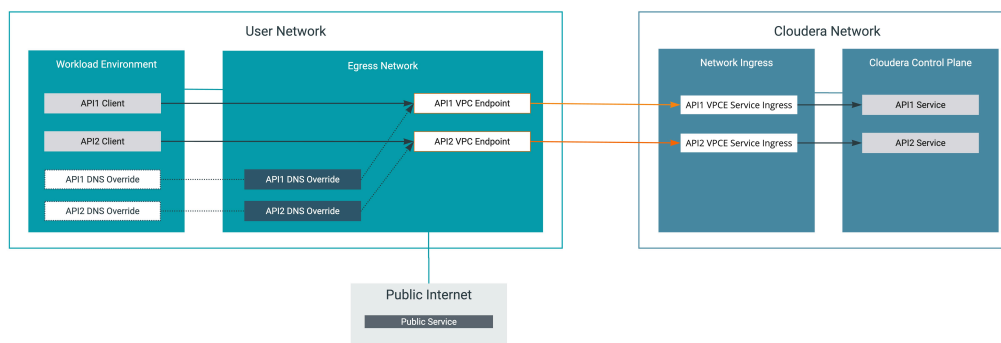
For the full list of command references, see [CDP CLI documentation](#).

## Additional VPC scenarios

Learn more about the additional VPC scenarios that show how VPC can be configured between your workload environment and Cloudera Control Plane.

### Scenario 1: VPC endpoint in your egress VPC with no HTTP proxy

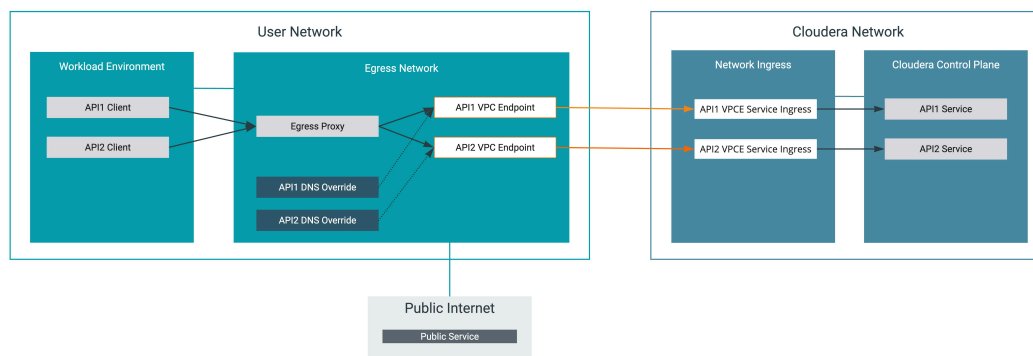
This scenario has the VPC endpoints in your egress VPC, allowing the VPC endpoints to be shared by multiple VPCs in your network. In this setup, VPC endpoints receive IPs from egress VPC.



- DNS overrides where DNS is a per-VPC view:
  - Deploying a VPC endpoint with private DNS option enabled will automatically install the DNS overrides in the egress VPC, but this does not solve the overrides for the workload environment VPC.
  - DNS override zones and records can be deployed in each workload environment VPC, or a single set of DNS override zones and records are deployed and the zones are associated with each workload environment VPC.
- DNS overrides where DNS is a regional or global view:
  - The DNS override will have a regional or global impact to resolution of the Cloudera hostnames, clients in the region/globally will receive these VPC endpoints.

### Scenario 2: VPC endpoint in your egress VPC, HTTP proxy

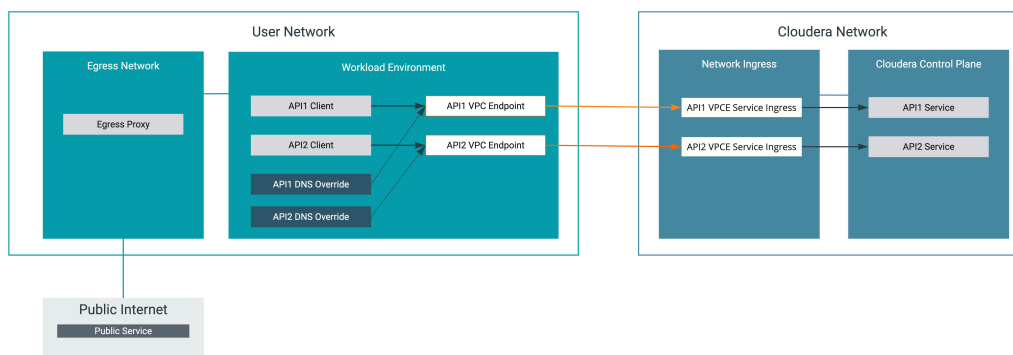
Similar to scenario 1, except you have egress traffic flowing through an HTTP proxy in the egress VPC. In this setup, VPC endpoints receive IPs from egress VPC.



- DNS overrides:
  - Deploying a VPC endpoint with private DNS option enabled will automatically install the DNS overrides in the egress VPC. If the egress proxy is performing the DNS lookup for the destination service, this approach should be sufficient.
  - Transparent proxy or egress firewall policy configurations may require the original destination IP to match the DNS resolution. If this is the case, the override zones/records/VPC associations can be deployed as described in scenario 1.

### Scenario 3: VPC endpoint in your workload environment VPC, HTTP proxy

VPC endpoints deployed in your workload environment network. In this setup, VPC endpoints receive IPs from egress VPC.



- DNS overrides where DNS is a per-VPC view:
  - Deploying a VPC endpoint with “private DNS” option enabled will automatically install the DNS overrides in the egress VPC. This is recommended.
- DNS overrides where DNS is a regional or global view:
  - The overrides will impact resolution for clients elsewhere in the region and globally.
  - Traffic to these hostnames from outside this VPC will attempt to use these VPC endpoints, which may not be a desired configuration
- HTTP forward proxy or non-transparent proxy
  - Workload environment will be configured to use an HTTP proxy profile.
  - The no\_proxy configuration of the profile must include the hostnames of the APIs reachable through VPC endpoint. HTTP requests for destinations in the no\_proxy list will not be forwarded to the proxy, local DNS and therefore the VPC endpoints will be used for that traffic