Outbound Network Access Destinations for CDP

Date published: 2020-10-06

Date modified: 2021-01-20

Legal Notice

CDP Public Cloud endpoints for CDP environment configuration

If you have limited outbound internet access (for example due to using a firewall or proxy), review this document to learn which specific outbound destinations must be available in order to register a CDP environment. This document lists:

- General endpoints applicable to all CDP environments
- AWS-specific endpoints
- Azure-specific endpoints

General endpoints

ID	Description/Usage	CDP Service	Destination	Protocol & Authentication	IP Protocol / Port	Comments
1	**Cloudera CCM** Persistent Control Plane connection	All services	IP: 44.234.52.96/27 Ports: 6000-6049 Hostname pattern: *.ccm.cdp.cloudera.com	SSH public/private key authentication	TCP/6000-6049	One connection per cluster configured; persistent.
2	**Cloudera Databus** Telemetry, billing and metering data	All services	dbusapi.us-west-1.altus.cloudera.com dbusapi.us-west-1.sigma.altus.cloudera.com	HTTPS with Cloudera-generated access key	TCP/443	Regular interval for telemetry, billing, metering services, and used for Workload Manager if enabled.
3	**Control Plane API**	Machine Learning	api.us-west-1.cdp.cloudera.com	HTTPS with Cloudera-generated access key	TCP/443	Cloudera's control plane REST API.
4	**Cloudera Manager parcels** Software distribution	Data Lake, Data Hub, Operational Database	archive.cloudera.com	HTTPS	TCP/443	Cloudera's public software repository. CDN backed service; IP range not predictable.
5	**Docker Images** Software Distribution	Data Engineering Machine Learning	container.repository.cloudera.com	HTTPS	TCP/443	Cloudera's public docker registry. CDN backed service; IP range not predictable.
6	**Docker Images** Software Distribution	Data Warehouse	auth.docker.io* cloudera-docker-dev.jfrog.io* docker-images-prod.s3.amazonaws.com* gcr.io* k8s.gcr.io* quay-registry.s3.amazonaws.com* quay.io* quayio-production-s3.s3.amazonaws.com* docker.io* production.cloudflare.docker.com* storage.googleapis.com* container.repo.cloudera.com **required only for old/existing CDW environments*	HTTPS	TCP/443	Moved to container.repo.cloudera.com

AWS specific endpoints

ID	Description/Usage	CDP Service	Destination	Protocol & Authentication	IP Protocol / Port	Comments
1	AWS STS	Data Lake	sts.amazonaws.com sts.*.amazonaws.com	HTTPS (one way) IAM authentication	TCP/443	CDP 7.1.1+ required before can be made internal with VPC endpoints.
2	AWS S3	Data Lake, Data Hub, Data Engineering, Data Warehouse, Machine Learning, Operational Database	*.s3.amazonaws.com *.s3.*.amazonaws.com s3.amazonaws.com	HTTPS (one way) IAM authentication	TCP/443	Can be made internal with VPC endpoints.
3	AWS DynamoDB	Data Lake, Data Hub, Data Engineering, Data Warehouse, Machine Learning, Operational Database	dynamodb.*.amazonaws.com	HTTPS (one way) IAM authentication	TCP/443	Can be made internal with VPC endpoints.
4	AWS RDS	Data Lake, Data Hub, Data Engineering	*.*.rds.amazonaws.com	JDBC / Postgres binary protocol / MySQL	TCP 5432 / 3306	VPC Internal. Only Data Engineering uses MySQL and requires port 3306 to be open.
5	AWS ECR	Data Warehouse, Machine Learning	api.ecr.*.amazonaws.com *.dkr.ecr.*.amazonaws.com	HTTPS (one way) IAM authentication	TCP/443	Can be made internal with VPC endpoints.
6	AWS EC2	Data Warehouse, Machine Learning, Operational Database	ec2.*.amazonaws.com	HTTPS (one way) IAM authentication	TCP/443	Can be made internal with VPC endpoints.
7	AWS EKS	Data Engineering, Data Warehouse, Machine Learning	eks.*.amazonaws.com	HTTPS (one way) IAM authentication	TCP/443	AWS does not support EKS VPC endpoints at this time.
8	AWS Cloudformation	Data Warehouse, Machine Learning	cloudformation.*.amazonaws.com	HTTPS (one way) IAM authentication	TCP/443	Can be made internal with VPC endpoints.
9	AWS Autoscaling	Data Engineering, Data Warehouse, Machine Learning	autoscaling.*.amazonaws.com	HTTPS (one way) IAM authentication	TCP/443	Can be made internal with VPC endpoints.
10	AWS EFS	Data Engineering, Data Warehouse, Machine Learning	elasticfilesystem.*.amazonaws.com	HTTPS (one way) IAM authentication	TCP/443	Can be made internal with VPC endpoints.
11	AWS EKS k8s cluster api	Data Warehouse	UNIQUEID.*.eks.amazonaws.com	HTTPS (one way) IAM authentication	TCP/443	Optional for new clusters.
12	AWS ELB	Data Engineering, Data Warehouse	elasticloadbalancing.*.amazonaws.com	HTTPS (one way) IAM authentication	TCP/443	Can be made internal with VPC endpoints.

13	AWS RDS API	Data Warehouse	rds.*.amazonaws.com	HTTPS (one way) IAM authentication	TCP/443	AWS does not support RDS API VPC endpoints at this time. This requirement is under further evaluation. Data Warehouse uses Amazon RDS for PostgreSQL.
14	AWS Service Quotas	Data Warehouse	servicequotas.*.amazonaws.com	HTTPS (one way) IAM authentication	TCP/443	AWS does not support Service Quota via VPC endpoints. Used to check limits and warn prior to hitting the limits.
15	AWS Price List Service	Data Warehouse	pricing.*.amazonaws.com	HTTPS (one way) IAM authentication	TCP/443	AWS Price List Service uses us-east-1 or ap-south-1 as the region.

Azure specific endpoints

ID	Description/Usage	CDP Service	Destination	Protocol & Authentication	IP Protocol / Port	Comments
0	General Azure guidelines	All	See Safelist the Azure portal URLs on your firewall or proxy server for Azure egress best practices.			
1	Azure Kubernetes Services (AKS)	Data Warehouse, Machine Learning	See Control egress traffic for cluster nodes in Azure Kubernetes Service (AKS). CDP uses AKS and has the same requirements.			
2	Azure Data Lake Storage Gen 2	Data Lake, Data Hub, Operational Database	<storage account name>.dfs.core.windows.net	HTTPS Azure authentication	TCP/443	Azure Storage VPC endpoint is required (*Microsoft.Storage*).
3	Azure Database for Postgres	Data Lake, Data Hub, Data Warehouse, Machine Learning	*.postgres.database.azure.com	JDBC / Postgres binary protocol	TCP/5432	Azure SQL VPC endpoint is required (*Microsoft.Sql*).
4	ARM to manage User Assigned Managed Identities	Data Lake	management.azure.com	HTTPS Azure authentication	TCP/443	This can be allowed by using the AzureResourceManager Azure service tag. Additionally IP addresses to whitelist are available to download.
5	Microsoft Log Analytics	All	*.agentsvc.azure-automation.net *.ods.opinsights.azure.com *.oms.opinsights.azure.com *.blob.core.windows.net	HTTPS Azure authentication	TCP/443	Optional; but may cause issues with Azure approved images if blocked.