IBM Cloud Pak for Data With Managed OpenShift on AWS Marketplace

Deployment Guide

July 2024

Contents

Overview

This deployment guide provides step-by-step instructions for deploying IBM Cloud Pak for Data on ROSA (Red Hat OpenShift Service on AWS) Container Platform cluster on the AWS Cloud. You can automatically deploy a multi-master, production instance of Cloud Pak for Data.

IBM Cloud Pak for Data enables enterprise users to connect, catalog, govern, transform, and analyze the data regardless of where the data is.

IBM Cloud Pak for Data on AWS

IBM Cloud Pak for Data is an analytics platform that helps prepare data for artificial intelligence (AI). It enables data engineers, data stewards, data scientists, and business analysts to collaborate using an integrated multiple-cloud platform.

IBM Cloud Pak for Data uses IBM's deep analytics portfolio to help organizations meet data and analytics challenges. The required building blocks (collect, organize, analyze, infuse) for information architecture are available using IBM Cloud Pak for Data on AWS.

IBM Cloud Pak for Data uses AWS services and features, including virtual private clouds (VPCs), Availability Zones, security groups, Elastic Block Storage, and Elastic Load Balancing to build a reliable and scalable cloud platform.

This reference deployment provides AWS CloudFormation templates to deploy Cloud Pak for Data onto a new OpenShift cluster. This cluster includes:

- A ROSA(Red Hat OpenShift Service on AWS) Container Platform cluster created in a new or existing VPC on Red Hat CoreOS (RHCOS) instances, using the Redhat OpenShift Installer Provisioned Infrastructure. See the OpenShift Container Platform Installation overview for details about the underlying OpenShift deployment architecture.
- Storage: EFS or Amazon FSx for NetApp ONTAP storage
- Scalable OpenShift compute nodes running Cloud Pak for Data services.

For more information about Cloud Pak for Data, see the IBM Cloud Pak for Data Documentation

This product requires an internet connection to deploy properly.

The following packages are downloaded on deployment: git, wget httpd-tools, python3, python3-pip, jq, unzip, podman, terraform, awscliv2, openshift-client-linux-4.14.17 , helm-v3.9.4 , quickstart-linux-utilities, amazon-ssm-agent from below repos:
https://github.com/aws-quickstart/quickstart-linux-utilities

https://s3-us-west-1.amazonaws.com/amazon-ssm-us-west-1/latest/linux_amd64/amazon-ssm-agent.rpm

Cost and licenses

The Cloud Pak for Data environment is deployed by using AWS CloudFormation template. You are responsible for the cost of the AWS services used for the infrastructure.

The AWS CloudFormation template for these deployments includes configuration parameters that you can customize. You can use it to build a new VPC for your Data Express solution on AWS cluster or deploy on an existing AWS VPC. Some of these settings, such as instance type, will affect the cost of deployment. For cost estimates, see the pricing pages for each AWS service you will be using. Prices are subject to change.
Pricing:
ROSA Service Fee: https://aws.amazon.com/rosa/pricing/
Infrastructure Fee:

- Master Nodes: 3 * m5.2xlarge
- Infra Nodes: 2 * r5.xlarge
- Worker Nodes: Count * m5.4xlarge
- Bastion Node: 1 * i3.large

Storage: EFS or Amazon FSx for NetApp ONTAP storage

For more information about licensing terms, Cloud Pak for Data software license agreement.

Upgrading to the latest version of Cloud Pak for Data indicates your acceptance of any new terms that may be applicable for the new version. To determine if new terms apply and to review them, please visit our IBM Terms,
To execute a search for License Information (LI) documents, and locate the LI applicable for the version you wish to upgrade to. To locate the LI for the Cloud Pak for Data Standard edition, type, ‘IBM Cloud Pak for Data Standard Edition’ and for the Enterprise edition type, ‘IBM Cloud Pak for Data Enterprise Edition.’

In the event there are no results for the version you are upgrading to, review the LI associated to the previous version which would apply in this scenario (for example, if you are upgrading to version 4.8.x and there is no LI for this version, you will need to review the LI for version 4.0.x)

Architecture

Deploying the AWS Marketplace template for a Cloud Pak for Data on new VPC with **default parameters** builds the following Cloud Pak for Data environment in the AWS Cloud to deliver the solution:

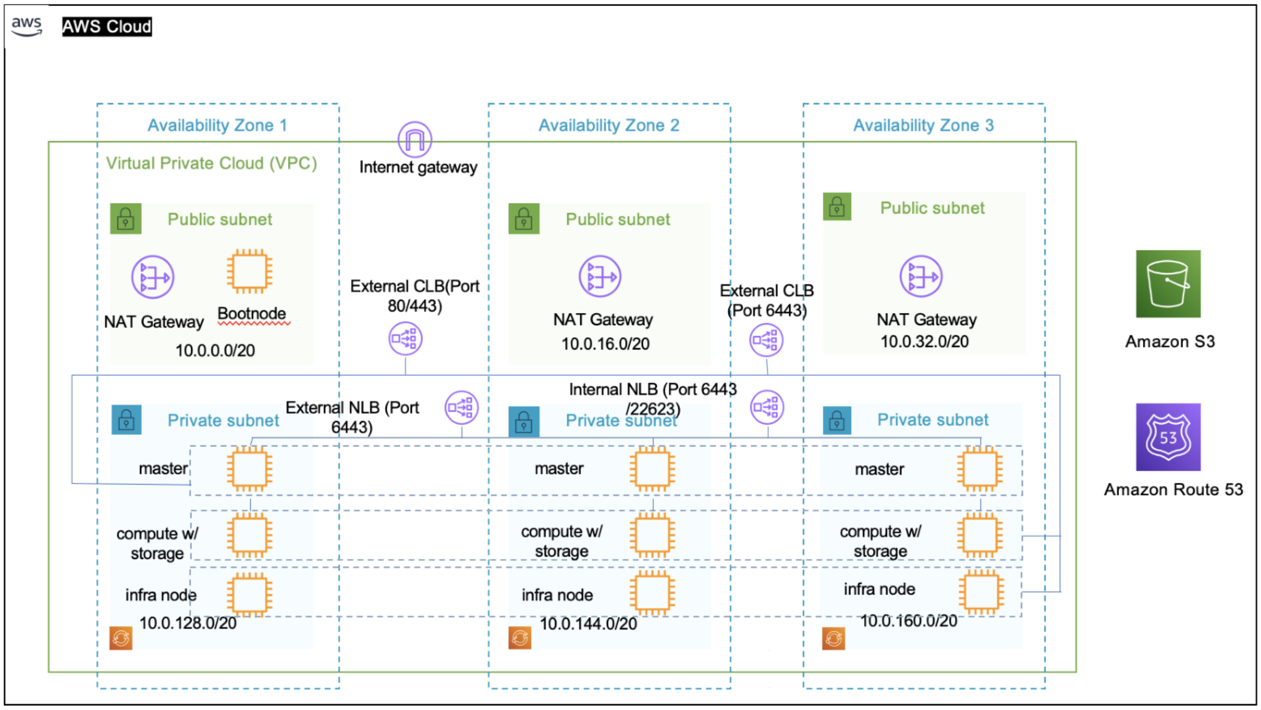

Figure 2: Deployment architecture for Cloud Pak for Data on AWS

The Cloud Formation Stack sets up the following:

- An architecture that spans one or three Availability Zones.* A VPC configured with public and private subnet according to AWS best practices, to provide you with your own virtual network on AWS. *
- In the public subnets:
  - Managed network address translation (NAT) gateway to allow outbound internet access for resources in the private subnet.*
  - A boot node Amazon Elastic Compute Cloud (Amazon EC2) instance that also serves as a bastion host to allow inbound Secure Shell (SSH) access to EC2 instances in private subnets.
- In the private subnets:
  - OCP master instances up to three Availability Zones.
  - OpenShift Container Platform (OCP) compute nodes that combined, contain Cloud Pak for Data Collect, Organize, and Analyze services.
  - Elastic Block Storage disks that are mounted on the compute nodes for container persistent data
- A Classic Load Balancer spanning the public subnets for accessing Cloud Pak for Data from a web browser.
- A Classic Load Balancer spanning the public subnets for accessing the OCP master instances.
- A Network Load Balancer spanning the private subnets for routing internal OpenShift application programming interface (API) traffic to the OCP master instances.

* The template that deploys into an existing VPC skips the components marked by asterisks and prompts you for your existing VPC configuration.

Cloud Pak for Data microservices are preconfigured on compute nodes. The following diagram shows the platform architecture.

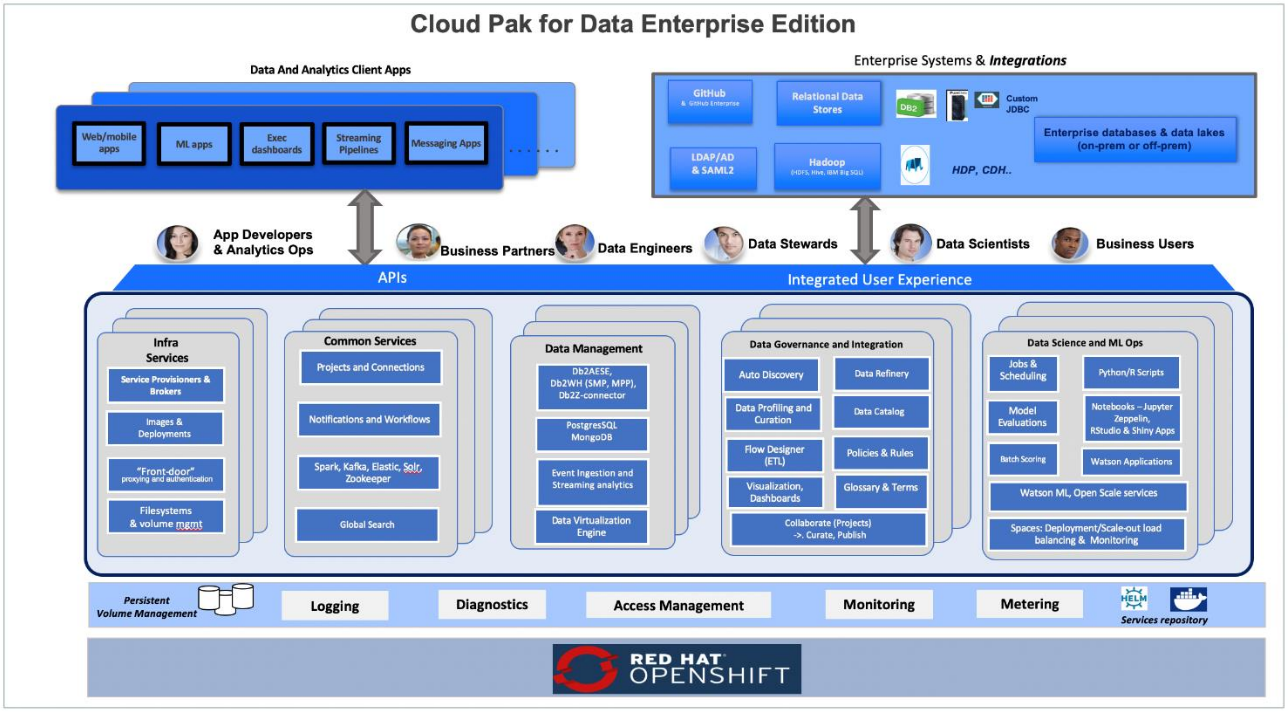
Cloud Pak for Data Enterprise Edition
Data And Analytics Client Apps
Web/mobile apps
ML apps
Exec dashboards
Streaming Pipelines
Messaging Apps
Enterprise Systems & Integrations
GitHub
& GitHub Enterprise
Relational Data Stores
DB2
Custom JDBC
Enterprise databases & data lakes (on-prem or off-prem)
LDAP/AD & SAML2
Hadoop
(HDFS, Hive, IBM Big SQL)
HDP, CDH..
App Developers & Analytics Ops
Business Partners
Data Engineers
Data Stewards
Data Scientists
Business Users
APIs
Integrated User Experience
Infra Services
Service Provisioners & Brokers
Images & Deployments
"Front-door" proxying and authentication
Filesystems & volume mgmt
Common Services
Projects and Connections
Notifications and Workflows
Spark, Kafka, Elastic, Solr, Zookeeper
Global Search
Data Management
Db2AESE, Db2WH (SMP, MPP), Db2Z-connector
PostgresSQL MongoDB
Event Ingestion and Streaming analytics
Data Virtualization Engine
Data Governance and Integration
Auto Discovery
Data Refinery
Data Profiling and Curation
Data Catalog
Flow Designer (ETL)
Policies & Rules
Visualization, Dashboards
Glossary & Terms
Collaborate (Projects) ->. Curate, Publish
Data Science and ML Ops
Jobs & Scheduling
Python/R Scripts
Model Evaluations
Notebooks – Jupyter Zeppelin, RStudio & Shiny Apps
Batch Scoring
Watson Applications
Watson ML, Open Scale services
Spaces: Deployment/Scale-out load balancing & Monitoring
Persistent Volume Management
Logging
Diagnostics
Access Management
Monitoring
Metering
HELM
Services repository
RED HAT OPENSHIFT

Figure 3: Cloud Pak for Data services

Planning the deployment

Specialized knowledge

This deployment assumes basic familiarity with Cloud Pak for Data components and services. If you're new to Cloud Pak for Data and Red Hat OpenShift, see the Additional resources section.
This deployment also assumes familiarity with the OpenShift command line interface and Linux, in addition to a moderate level of familiarity with AWS services.

Technical requirements

For Cloud Pak for Data requirements, see System requirements for Cloud Pak for Data.
Red Hat Enterprise Linux CoreOS (RHCOS) is used for the OpenShift compute node instances in this deployment.
Before you launch the template, your account must have resource quotas as specified in the following table.

Resources

If necessary, request service quota increases for the following resources. You might need to do this if an existing deployment uses these resources, and you might exceed the default quotas with this deployment. The Service Quotas console displays your usage and quotas for some aspects of some services. For more information, see the AWS documentation.

Resource	This deployment uses
VPCs	1
Elastic IP addresses	1
Network Load Balancers	2
Classic Load Balancers	2
i3.large instances	1
m5.2xlarge instances	3
m5.4xlarge instances	3

Regions

This deployment uses 1 or 3 Availability Zones.

Key pair

Make sure that at least one Amazon EC2 key pair exists in your AWS account in the Region where you are planning to deploy the template. Make note of the key pair name. You'll be prompted for this information during deployment. To create a key pair, follow the instructions in the AWS documentation.

If you're deploying for testing or proof-of-concept purposes, we recommend that you create a new key pair instead of specifying a key pair that's already being used by a production instance.

IAM permissions

To deploy the template, you must log in to the AWS Management Console with AWS Identity and Access Management (IAM) permissions for the resources and actions the templates will deploy.

The permissions required for the creation of ROSA (Redhat OpenShift Service on AWS) in sts mode. Please review:
https://docs.openshift.com/rosa/rosa_architecture/rosa-sts-about-iam-resources.html#rosa-minimum-scp_rosa-sts-about-iam-resources

Deployment options

This template provides the following deployment options:

- **Deploy Cloud Pak for Data into a new VPC** (end-to-end deployment). This option builds a new AWS environment consisting of the VPC, subnets, NAT gateways, security groups, bastion hosts, and other infrastructure components, and then deploys Cloud Pak for Data into this new VPC.
- **Deploy Cloud Pak for Data into an existing VPC**. This option provisions Cloud Pak for Data in your existing VPC infrastructure.

The template also lets you configure CIDR blocks, instance types, and Cloud Pak for Data settings, as discussed later in this guide.

Pre-requisites

Ensure the following subscriptions are in place with your existing Cloud Pak for Data entitlements or obtained from AWS Marketplace

Step 1. IBM Cloud Pak for Data Subscription

When you purchase Cloud Pak for Data from Marketplace, you will get the Cloud Pak for Data entitlement Username and API Key. You should keep it handy as it's a required parameter in the Cloud Formation Template

Step 2. ROSA Subscription

Enable ROSA
* Enable ROSA [here]
* Get RedHat ROSA token [here]

Step 3. Sign in to your AWS account

1. Sign in to your AWS account at https://aws.amazon.com with an AWS Identity and Access Management (IAM) user role that has the necessary permissions. For details, see Planning the deployment, earlier in this guide.
2. Make sure that your AWS account is configured correctly, as discussed in Technical requirements, earlier in this guide.
3. Use the Region selector in the navigation bar to choose the AWS Region where you want to deploy Cloud Pak for Data on AWS.
4. The following resources should be made available for Cloud Pak for Data deployment
   - A key pair to SSH to the cluster nodes. In the navigation pane of the Amazon EC2 console, choose **Key Pairs**, and then choose your key pair from the list.
   - An existing S3 Bucket to store deployment log files.

Step 4. Launch the Deployment

> **Notes** The instructions in this section reflect the current version of the AWS CloudFormation console. If you're using the redesigned console, some of the user interface elements might be different.
>
> You are responsible for the cost of the AWS services used while running this deployment. For full details, see the pricing pages for each AWS service you will be using for this deployment. Prices are subject to change.

5. Launch the AWS CloudFormation template into your AWS account from the AWS Marketplace directly or you can download the template and launch it separately from your account.

   A Cloud Pak for Data standard deployment takes about 3 hours.

   1. Check the Region that's displayed in the upper-right corner of the navigation bar and change it if necessary. This is where the network infrastructure for Cloud Pak for Data will be built. The template is launched in the US East (N. Virginia) Region by default.
   2. On the **Create stack** page, keep the default setting for the template URL, and then choose **Next**.

3. On the **Specify stack details** page, change the stack name if needed. Review the parameters for the template. Provide values for the parameters that require input. For all other parameters, review the default settings and customize them as necessary.

In the following tables, parameters are listed by category and described separately for these deployment options:

- Parameters for deploying Cloud Pak for Data into a new VPC
- Parameters for deploying Cloud Pak for Data into an existing VPC

When you finish reviewing and customizing the parameters, choose **Next**.

Option 1: Parameters for deploying Cloud Pak for Data into a new VPC

VPC network configuration:

Parameter label (name)	Default	Description
Number of Availability Zones (NumberOfAZs)	1	The number of Availability Zones to be used for the deployment. Allowed values: 1 or 3.
Availability Zones (AvailabilityZones)	*Requires input*	The list of Availability Zones to use for the subnets in the VPC. The template uses one/ three Availability Zones.
VPC CIDR (VPCCIDR)	10.0.0.0/16	The CIDR block for the VPC to be created.
Private subnet 1 CIDR (PrivateSubnet1CIDR)	10.0.0.0/19	The CIDR block for the private subnet located in Availability Zone 1.
Private subnet 2 CIDR (PrivateSubnet2CIDR)	10.0.32.0/19 The CIDR block for the private subnet located in Availability Zone 2.	The CIDR block for the private subnet located in Availability Zone 2.
Private subnet 3 CIDR (PrivateSubnet3CIDR)	10.0.64.0/19 The CIDR block for the private subnet located in Availability Zone 3.	The CIDR block for the private subnet located in Availability Zone 3.
Public subnet 1 CIDR (PublicSubnet1CIDR)	10.0.128.0/20	The CIDR block for the public subnet located in Availability Zone 1.
Public subnet 2 CIDR (PublicSubnet2CIDR)	10.0.144.0/20	The CIDR block for the public subnet located in Availability Zone 2.
Public subnet 3 CIDR (PublicSubnet3CIDR)	10.0.160.0/20	The CIDR block for the public subnet located in Availability Zone 3.
BootNodeAccessCIDR (**BootNodeAccessCIDR**)	*Requires input*	The CIDR IP range that is permitted to access the Boot Node instance. We recommend that you set this value to a trusted IP range. The value \`0.0.0.0/0\` permits all IP

Parameter label (name)	Default	Description
		addresses to access. Additional values can be added post-deployment from the Amazon EC2 console.
Cluster Network CIDR (**ClusterNetworkCIDR**)	10.128.0.0/14	The Cluster Network CIDR IP range that is used as IP address pools for pods.
ServiceNetworkCIDR	172.30.0.0/16	The Service Network CIDR IP range that is used as block of IP addresses for services.

Amazon EC2 configuration:

Parameter label (name)	Default	Description
Key pair name (KeyPairName)	*Requires input*	The name of an existing public/private key pair, which allows you to securely connect to your instance after it launches.

OpenShift hosts configuration:

Parameter label (name)	Default	Description
Compute instance type (NodesInstanceType)	m5.4xlarge	The EC2 instance type for the OpenShift node instances.
Number of compute nodes (NumberOfNodes)	3(default)	The number of nodes depends on the numbers of cp4d services selected for the deployment. For more info check here. **Warning** If the number of node instances exceeds your Red Hat entitlement limits or AWS instance quotas, the stack will fail. Choose a number that is within your limits.
Cluster name (ClusterName)	*Requires input*	Custom cluster name for kubernetes.io/cluster/tags. If left blank, will use the stack name suffixed with the AWS Region.
PrivateCluster (PrivateCluster)	false	To Deploy a Private cluster select true and false for Public cluster
Choose OpenShift Version (OpenShiftVersion)	4.14.17	Choose a supported OpenShift Version.

Storage Configuration:

Parameter label (name)	Default	Description
Storage type for Cluster (StorageType)	*EFS-EBS*	EFS-EBS and fsx options are available.
FSx Admin password (FSxAdminPassword)	Netapp1!	Default fsx filesystem fsxadmin user password. Applicable only when Storage Type chosen is fsx
FSx Storage Capacity (FSxStorageCapacity)	1TB	FSx file system Storage Capacity. Applicable only when Storage Type chosen is fsx
FSx ThroughPut Capacity (FSxThroughputCapacity)	128MBs	FSx file system throughput Capacity. Applicable only when Storage Type chosen is fsx
FSx Endpoint CIDR (FSxEndpointCIDR)	10.0.255.192/26	The IP address range in which the endpoints to access your fsx file system will be created, required only for multi-AZ deployment. Applicable only when Storage Type chosen is fsx

ROSA subscription information:

Parameter label (name)	Default	Description
ROSA Token	*Requires input*	Update the value with the ROSA Token.

IBM Cloud Pak for Data configuration:

Parameter label (name)	Default	Description
IBM Cloud Pak for Data Entitled Registry User (APIUsername)	cp	Enter the IBM Cloud Pak for Data Username value to access IBM Container Registry.
IBM Cloud Pak for Data Entitled Registry API key Value (APIKey)	—	Enter the IBM Cloud Pak for Data API key to access IBM Container Registry
License agreement (LicenseAgreement)	—	I have read and agreed to the license terms for IBM Cloud Pak for Data that were provided to me at time of purchase.
Output S3 bucket name (ICPDDeploymentLogsBucketName)	*Requires Input*	The Existing S3 bucket name where the zip file output should be placed.

IBM Cloud Pak for Data services:

Parameter label (name)	**Default**	**Description**
IBMKnowledgeCatalog	no	Installs Watsosn Knowledge Catalog service.
Watson Query	**no**	Installs the Watson Query(Data Virtualization) service.
AnalyticsEngine	**no**	Installs the Analytics Engine service.
WatsonStudio	**no**	Installs the Watson Studio service.
WatsonMachineLearning	**no**	Installs the Watson Machine Learning service.
WatsonAIOpenscale	**no**	Installs the Watson AI Openscale service.
SPSSModeler	**no**	Installs the SPSS Modeler service.
CognosDashboardEmbedded	**no**	Installs the Cognos Dashboard Enbedded service.
DataStage	**no**	Installs the DataStage service.
Db2Warehouse	**no**	Installs the Db2 Warehouse service.
Db2OLTP	**no**	Installs the Db2OLTP service.
CognosAnalytics	**no**	Installs the Cognos Analytics service.
IBM Match 360 with Watson	**no**	Installs the IBM Match 360 with Watson service.
DecisionOptimization	**no**	Installs the Decision Optimization service.
BigSQL	**no**	Installs the BigSQL service.
PlanningAnalytics	**no**	Installs the Planning Analytics service.

Parameter label (name)	Default	Description
Watsonx Assistant	**no**	Installs Watsonx Assistant. ***Note: Only single Availability zone is supported for a successful Watsonx Assistant installation. Multi Aavailiability zone is not supported.***
Watson Discovery	**no**	Installs Watson Discovery service.
Watson Speech	**no**	Installs Watson Speech service.
Manta	**no**	Installs Manta Service. ***Note: For a successful Manta service installation, IBM Knowledge Catalog service is the pre requisite. Set both parameters Manta and IBM Knowledge Catalog to 'yes' for a successful installation.***
OpenPages	**no**	Installs the OpenPages service.
IBM Db2 Data Management Console	**no**	Installs the Data Management Console service.
Factsheets	**no**	Installs the Factsheets service. ***Note: For a successful installation of the Factsheets service, IBM Knowledge Catalog service or Watson Studio service must have been installed on the cluster.***
RStudio	**no**	Installs the RStudio service.
WatsonPipelines	**no**	Installs the Watson Pipelines service.

Option 2: Parameters for deploying Cloud Pak for Data into an existing VPC

Network configuration:

Parameter label (name)	Default	Description
Number of Availability Zones (NumberOfAZs)	1	The number of Availability Zones to be used for the deployment. Allowed values: 1 or 3.
VPC ID (VPCID)	*Requires input*	The ID of your existing VPC for deployment.
VPC CIDR (VPCCIDR)	*Requires input*	The CIDR block for the VPC to be created.
Private subnet 1 ID (PrivateSubnet1ID)	*Requires input*	The ID of the private subnet in Availability Zone 1 for the workload (e.g., subnet-a0246dcd).

Parameter label (name)	Default	Description
Private subnet 2 CIDR (PrivateSubnet2CIDR)	*Requires input*	The CIDR block for the private subnet located in Availability Zone 2.
Private subnet 3 CIDR (PrivateSubnet3CIDR) 1	*Requires input*	The CIDR block for the private subnet located in Availability Zone 3.
Public subnet 1 ID (PublicSubnet1ID)	*Requires input*	The ID of the public subnet in Availability Zone 1 for the ELB load balancer (e.g., subnet-9bc642ac).
Public subnet 2 CIDR (PublicSubnet2CIDR)	*Requires input*	The CIDR block for the public subnet located in Availability Zone 2.
Public subnet 3 CIDR (PublicSubnet3CIDR)	*Requires input*	The CIDR block for the public subnet located in Availability Zone 3.
BootNodeAccessCIDR (**BootNodeAccessCIDR**)	*Requires input*	The CIDR IP range that is permitted to access the Boot Node instance. We recommend that you set this value to a trusted IP range. The value \`0.0.0.0/0\` permits all IP addresses to access. Additional values can be added post-deployment from the Amazon EC2 console.
Cluster Network CIDR (**ClusterNetworkCIDR**)	*Requires input*	The Cluster Network CIDR IP range that is used as IP address pools for pods.
ServiceNetworkCIDR	*Requires input*	The Service Network CIDR IP range that is used as block of IP addresses for services.

Amazon EC2 configuration:

Parameter label (name)	Default	Description
Key pair name (KeyPairName)	*Requires input*	The name of an existing public/private key pair, which allows you to securely connect to your instance after it launches.

OpenShift hosts configuration:

Parameter label (name)	Default	Description
Compute instance type (NodesInstanceType)	m5.4xlarge	The EC2 instance type for the OpenShift node instances.
Number of compute nodes (NumberOfNodes)	3(default)	The number of nodes depends on the numbers of cp4d services selected for the deployment. For more info, check here . **Warning** If the number of node instances exceeds your Red Hat entitlement limits or AWS instance quotas, the stack will fail. Choose a number that is within your limits.

Parameter label (name)	Default	Description
Cluster name (ClusterName)	*Requires input*	Custom cluster name for kubernetes.io/cluster/tags. If left blank, will use the stack name suffixed with the AWS Region.
Private Cluster (PrivateCluster)	false	To Deploy a Private cluster select true and false for Public cluster.
Choose OpenShift Version (OpenShiftVersion)	4.14.17	Choose a supported OpenShift Version.

Storage Configuration:

Parameter label (name)	Default	Description
Storage type for Cluster(StorageType)	*EFS-EBS*	EFS-EBS and fsx options are available.
FSx Admin password (FSxAdminPassword)	Netapp1!	Default fsx filesystem fsxadmin user password. Applicable only when Storage Type chosen is fsx
FSx Storage Capacity (FSxStorageCapacity)	1TB	FSx file system Storage Capacity. Applicable only when Storage Type chosen is fsx
FSx ThroughPut Capacity (FSxThroughputCapacity)	128MBs	FSx file system throughput Capacity. Applicable only when Storage Type chosen is fsx
FSx Endpoint CIDR (FSxEndpointCIDR)	10.0.255.192/26	The IP address range in which the endpoints to access your fsx file system will be created, required only for multi-AZ deployment. Applicable only when Storage Type chosen is fsx

ROSA subscription information:

Parameter label (name)	Default	Description
ROSA Token	*Requires input*	Update the value with the ROSA Token.

IBM Cloud Pak for Data configuration:

Parameter label (name)	Default	Description
IBM Cloud Pak for Data Entitled Registry User (APIUsername)	cp	Enter the IBM Cloud Pak for Data Username value to access IBM Container Registry.
IBM Cloud Pak for Data Entitled Registry API	—	Enter the IBM Cloud Pak for Data API key to access IBM Container Registry

Parameter label (name)	Default	Description
key Value (APIKey)		
License agreement (LicenseAgreement)	—	I have read and agreed to the license terms for IBM Cloud Pak for Data that were provided to me at time of purchase.
OpenShift project (NameSpace)	zen	The OpenShift project that will be created for deploying Cloud Pak for Data. It can be any lowercase string.
Output S3 Bucket **(**ICPDDeploymentLogsBucketName)		The name of the S3 bucket where IBM Cloud Pak for Data deployment logs are to be exported. The deployment logs provide a record of the boot strap scripting actions and are useful for problem determination if the deployment fails in some way

IBM Cloud Pak for Data services:

Parameter label (name)	Default	Description
IBM Knowledge Catalog	no	Installs IBM Knowledge Catalog(Watson Knowledge Catalog) service.
Watson Query	**no**	Installs the Watson Query(Data Virtualization) service.
AnalyticsEngine	**no**	Installs the Analytics Engine service.
WatsonStudio	**no**	Installs the Watson Studio service.
WatsonMachineLearning	**no**	Installs the Watson Machine Learning service.
WatsonAIOpenscale	**no**	Installs the Watson AI Openscale service.
SPSSModeler	**no**	Installs the SPSS Modeler service.
CognosDashboardEmbedded	**no**	Installs the Cognos Dashboard Enbedded service.
DataStage	**no**	Installs the DataStage service.
Db2Warehouse	**no**	Installs the Db2 Warehouse service.
Db2OLTP	**no**	Installs the Db2OLTP service.
CognosAnalytics	**no**	Installs the Cognos Analytics service.

Parameter label (name)	Default	Description
IBM Match 360 with Watson	**no**	Installs the IBM Match 360 with Watson service.
DecisionOptimization	**no**	Installs the Decision Optimization service.
BigSQL	**no**	Installs the BigSQL service.
PlanningAnalytics	**no**	Installs the Planning Analytics service.
Watsonx Assistant	**no**	Installs Watsonx Assistant service. Installs Watsonx Assistant. *Note: Only single Availability zone is supported for a successfulWatsonx Assistant installation. Multi Aavailiability zone is not supported.*
Watson Discovery	**no**	Installs Watson Discovery service.
Watson Speech	**no**	Installs Watson Speech service.
Manta	**no**	Installs Manta service. *Note: For a successful Manta service installation, IBM Knowledge Catalog service is the pre requisite. Set both parameters Manta and IBM Knowledge Catalog to 'yes' for a successful installation.*
OpenPages	**no**	Installs the OpenPages service.
IBM Db2 DataManagementConsole	**no**	Installs the IBM Db2 Data Management Console service.
Factsheets	**no**	Installs the Factsheets service. ***Note: For a successful installation of the Factsheets service, IBM Knowledge Catalog service or Watson Studio service must have been installed on the cluster.***
RStudio	**no**	Installs the RStudio service.
WatsonPipelines	**no**	Installs the Watson Pipelines service.

4. On the options page, you can specify tags (key-value pairs) for resources in your stack and set advanced options. When you're done, choose **Next**.

   5. On the **Review** page, review and confirm the template settings. Under **Capabilities**, select the check box to acknowledge that the template will create IAM resources.
   6. Choose **Create stack** to deploy the stack.
   7. Monitor the status of the stack. When the status is **CREATE_COMPLETE**, the Cloud Pak for Data cluster is ready.

6. Use the URLs displayed in the **Outputs** tab for the stack to view the resources that were created. The URL for the **ICPDWebClientURL** output key will navigate to the console login page.

Step 5. (Optional) Specifying Services to deploy a Data Fabric Solution

IBM Data Fabric Solutions (Data Governance & Privacy, Multicloud Data Integration and Data Science and MLOps) are prescribed combinations of Cloud Pak for Data services which solve particular common use cases. You can deploy these solutions using either option 1 or 2 above by specifying the required services within the applicable Cloud Formation Template.

To configure each solution, simply select "yes" for the following CPD services in your Cloud Formation Template.

Data Fabric Solution	Services to Select
Data Governance and Privacy	IBM Knowledge Catalog(WatsonKnowledgeCatalog) WatsonQuery IBM Match 360 with Watson
Multicloud Data Integration	IBM Knowledge Catalog(WatsonKnowledgeCatalog) Watson Query DataStage
Data Science and MLOps	WatsonStudio WatsonMachineLearning WatsonAIOpenscale

Step 6. (Optional) Edit the application security group

Optional: You might need to edit the AWS application group to add IP addresses that can access the Cloud Pak for Data web client.
Navigate to Load Balancers on your AWS console and filter on tags, for example *kubernetes.io/service-name: openshift-ingress/router-default*.

1. In Load Balancers, filter and select the security group.

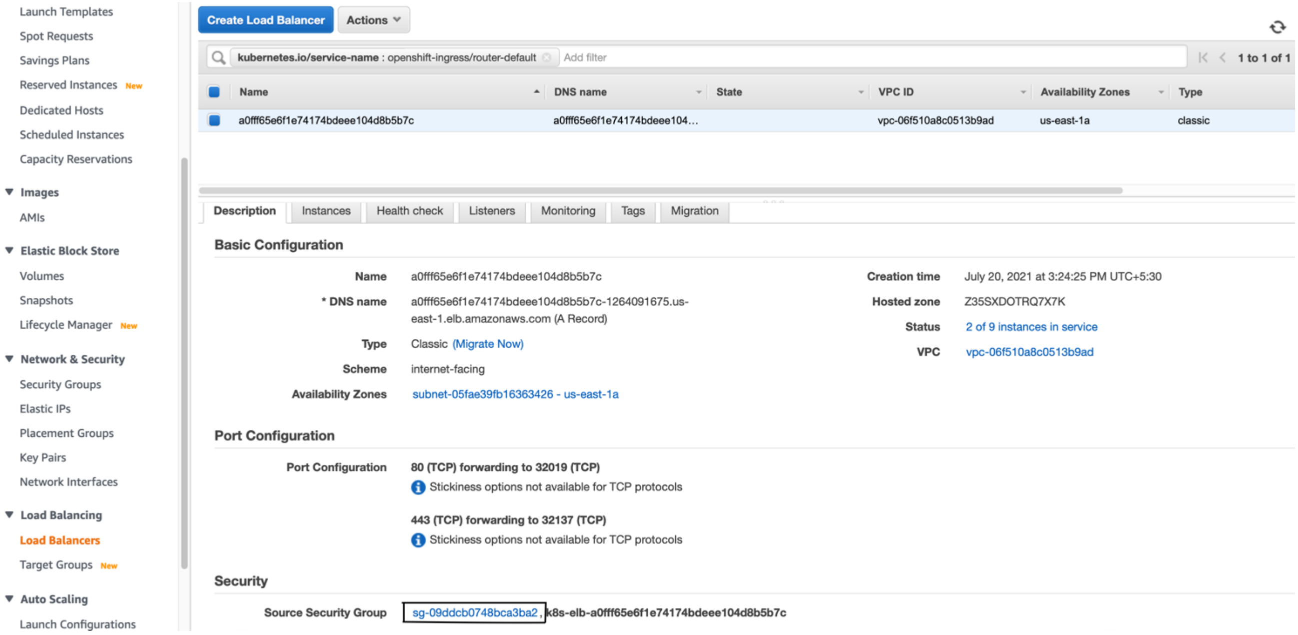

kubernetes.io/service-name : openshift-ingress/router-default

Name	DNS name	State	VPC ID	Availability Zones	Type
a0fff65e6f1e74174bdeee104d8b5b7c	a0fff65e6f1e74174bdeee104...		vpc-06f510a8c0513b9ad	us-east-1a	classic

Basic Configuration

Name: a0fff65e6f1e74174bdeee104d8b5b7c

* DNS name: a0fff65e6f1e74174bdeee104d8b5b7c-1264091675.us-east-1.elb.amazonaws.com (A Record)

Type: Classic (Migrate Now)

Scheme: internet-facing

Availability Zones: subnet-05fae39fb16363426 - us-east-1a

Creation time: July 20, 2021 at 3:24:25 PM UTC+5:30

Hosted zone: Z35SXDOTRQ7X7K

Status: 2 of 9 instances in service

VPC: vpc-06f510a8c0513b9ad

Port Configuration

Port Configuration: 80 (TCP) forwarding to 32019 (TCP)
Stickiness options not available for TCP protocols

443 (TCP) forwarding to 32137 (TCP)
Stickiness options not available for TCP protocols

Security

Source Security Group: sg-09ddcb0748bca3ba2 , k8s-elb-a0fff65e6f1e74174bdeee104d8b5b7c

2. Select Security Group and modify the Inbound rules

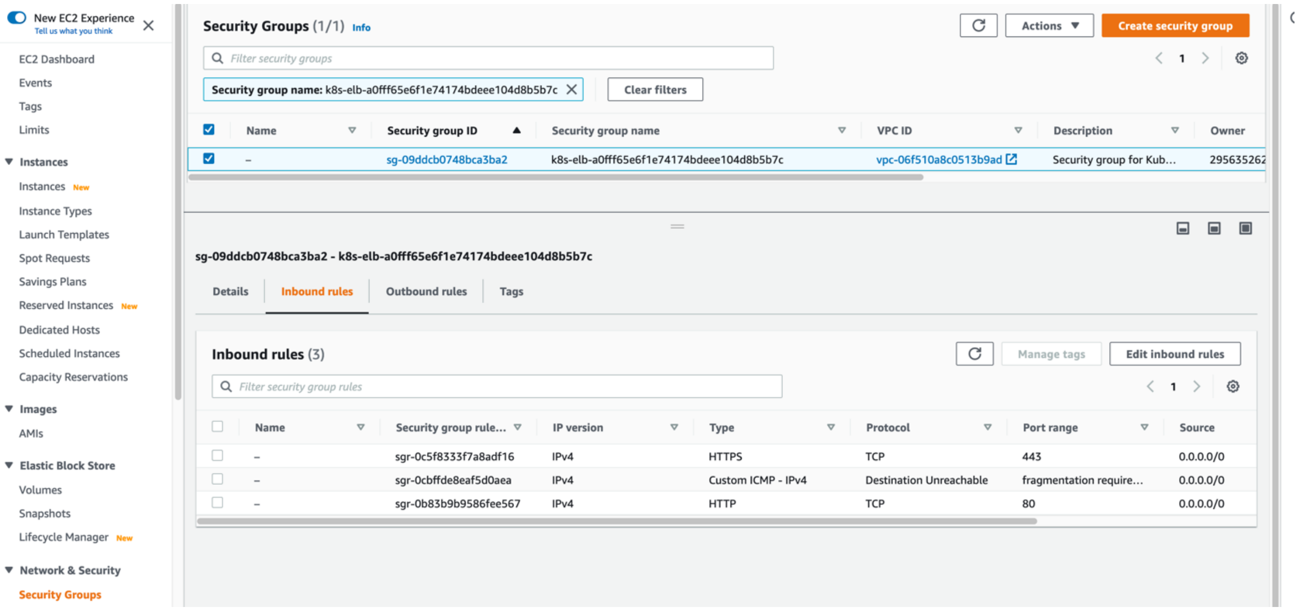

Security Groups (1/1)

Security group name: k8s-elb-a0fff65e6f1e74174bdeee104d8b5b7c

Name	Security group ID	Security group name	VPC ID	Description	Owner
–	sg-09ddcb0748bca3ba2	k8s-elb-a0fff65e6f1e74174bdeee104d8b5b7c	vpc-06f510a8c0513b9ad	Security group for Kub...	295635262

sg-09ddcb0748bca3ba2 - k8s-elb-a0fff65e6f1e74174bdeee104d8b5b7c

Inbound rules (3)

Name	Security group rule...	IP version	Type	Protocol	Port range	Source
–	sgr-0c5f8333f7a8adf16	IPv4	HTTPS	TCP	443	0.0.0.0/0
–	sgr-0cbffde8eaf5d0aea	IPv4	Custom ICMP - IPv4	Destination Unreachable	fragmentation require...	0.0.0.0/0
–	sgr-0b83b9b9586fee567	IPv4	HTTP	TCP	80	0.0.0.0/0

Figure 5: OpenShift node security group

3. Choose **Add Rule**, and fill in the rule details. For the rule **Type**, select either HTTP or HTTPS in the drop-down menu. Port 80 or 443 is filled in automatically. Add the network CIDR for the group of IP addresses that you want to permit HTTP or HTTPS access to the proxy nodes. To allow any IP address, use 0.0.0.0/0.

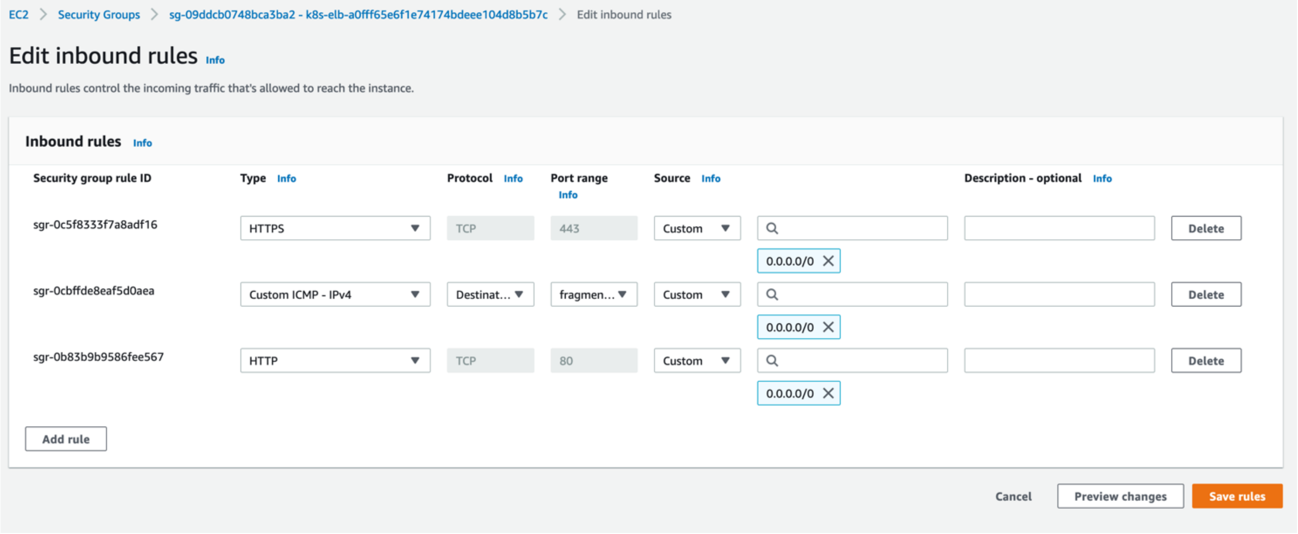

EC2 > Security Groups > sg-09ddcb0748bca3ba2 - k8s-elb-a0fff65e6f1e74174bdeee104d8b5b7c > Edit inbound rules

Edit inbound rules

Inbound rules control the incoming traffic that's allowed to reach the instance.

Inbound rules

Security group rule ID	Type	Protocol	Port range	Source		Description - optional
sgr-0c5f8333f7a8adf16	HTTPS	TCP	443	Custom	0.0.0.0/0	
sgr-0cbffde8eaf5d0aea	Custom ICMP - IPv4	Destinat…	fragmen…	Custom	0.0.0.0/0	
sgr-0b83b9b9586fee567	HTTP	TCP	80	Custom	0.0.0.0/0	

Figure 7: Supplying rule details

4. In the rule editor window, choose **Save**.

Step 7. Login to Cloud Pak for Data web client

When the AWS CloudFormation template has successfully created the stack, all server nodes will be running with the software installed in your AWS account. In the following steps, connect to Cloud Pak for Data web client to verify the deployment, and then use the web client to explore Cloud Pak for Data features.

1. To access the Cloud Pak for Data web client, first get the console URL from the Cloud Formation Stack output for key ICPDWebClientURL

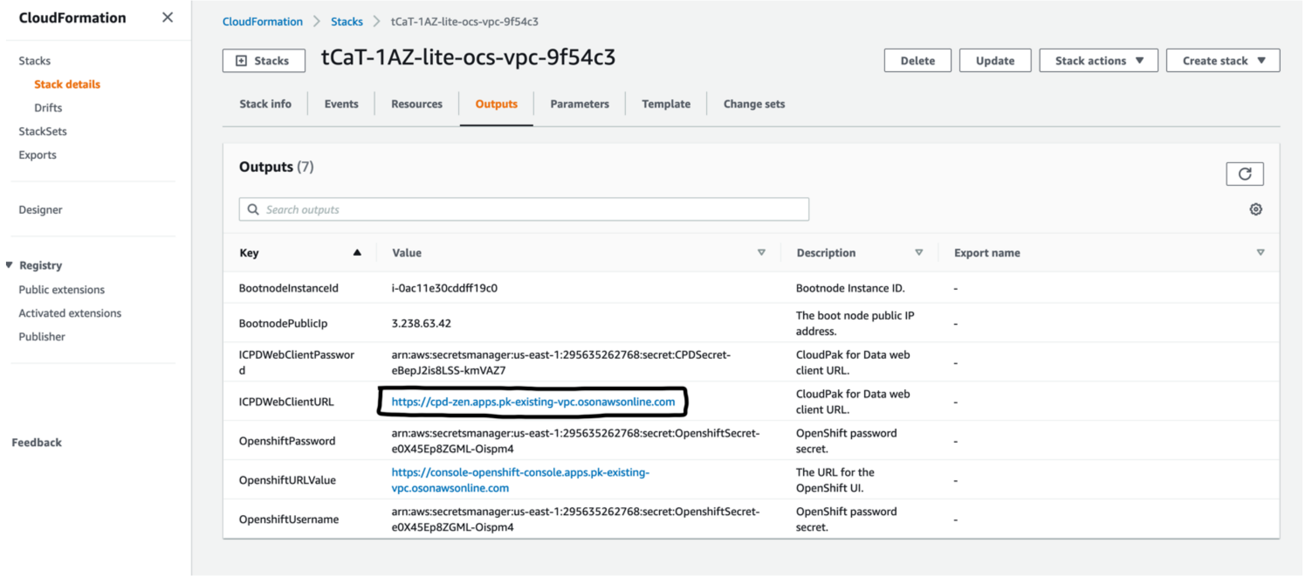

CloudFormation > Stacks > tCaT-1AZ-lite-ocs-vpc-9f54c3

tCaT-1AZ-lite-ocs-vpc-9f54c3

Outputs (7)

Key	Value	Description	Export name
BootnodeInstanceId	i-0ac11e30cddff19c0	Bootnode Instance ID.	-
BootnodePublicIp	3.238.63.42	The boot node public IP address.	-
ICPDWebClientPassword	arn:aws:secretsmanager:us-east-1:295635262768:secret:CPDSecret-eBepJ2is8LSS-kmVAZ7	CloudPak for Data web client URL.	-
ICPDWebClientURL	https://cpd-zen.apps.pk-existing-vpc.osonawsonline.com	CloudPak for Data web client URL.	-
OpenshiftPassword	arn:aws:secretsmanager:us-east-1:295635262768:secret:OpenshiftSecret-e0X45Ep8ZGML-Oispm4	OpenShift password secret.	-
OpenshiftURLValue	https://console-openshift-console.apps.pk-existing-vpc.osonawsonline.com	The URL for the OpenShift UI.	-
OpenshiftUsername	arn:aws:secretsmanager:us-east-1:295635262768:secret:OpenshiftSecret-e0X45Ep8ZGML-Oispm4	OpenShift password secret.	-

2. Get the Cloud Pak for Data web client user name and password from the Secret Manager with names <stack name>-cp4d-username and <stack name>-cp4d-password. Get the Openshift username

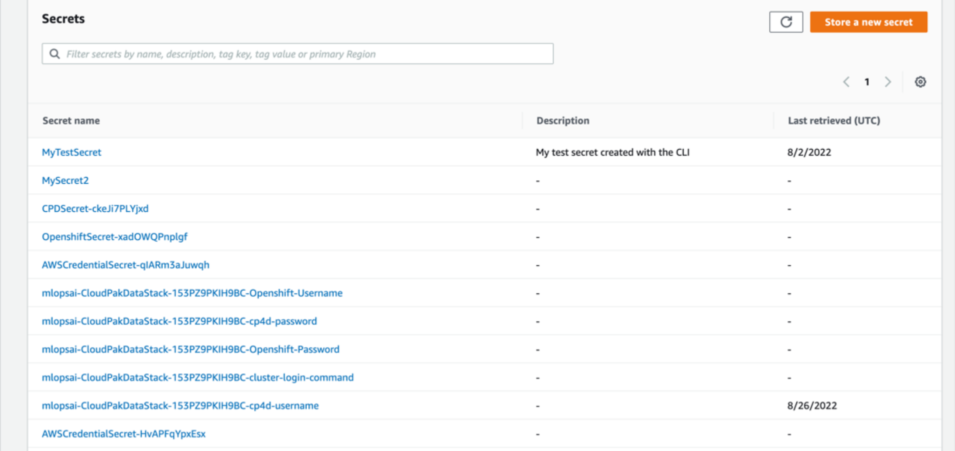
Secrets
Store a new secret
Filter secrets by name, description, tag key, tag value or primary Region
1

Secret name	Description	Last retrieved (UTC)
MyTestSecret	My test secret created with the CLI	8/2/2022
MySecret2	-	-
CPDSecret-ckeJi7PLYjxd	-	-
OpenshiftSecret-xadOWQPnplgf	-	-
AWSCredentialSecret-qlARm3aJuwqh	-	-
mlopsai-CloudPakDataStack-153PZ9PKIH9BC-Openshift-Username	-	-
mlopsai-CloudPakDataStack-153PZ9PKIH9BC-cp4d-password	-	-
mlopsai-CloudPakDataStack-153PZ9PKIH9BC-Openshift-Password	-	-
mlopsai-CloudPakDataStack-153PZ9PKIH9BC-cluster-login-command	-	-
mlopsai-CloudPakDataStack-153PZ9PKIH9BC-cp4d-username	-	8/26/2022
AWSCredentialSecret-HvAPFqYpxEsx	-	-

Figure 8: OpenShift Secret in Secret Manager

3. Click the secret suffixed with “OpenShift-Username” and retrieve the value of the secret. Take a note of the value “OpenShift-Username”. Click the secret suffixed with “OpenShift-password” and retrieve the value of the secret. Take a note of the value of the “OpenShift-password”. Click the secret suffixed with “cp4d-username” and retrieve the value of the secret. Take a note of the value of the “cp4d-username”. Click the secret suffixed with “cp4d-password” and retrieve the value of the secret. Take a note of the value of the “cp4d-password”. Click the secret suffixed with cluster-login-command and retrieve the value of the secret. Take a note of the value of Openshift cluster login command.

Openshift Username:

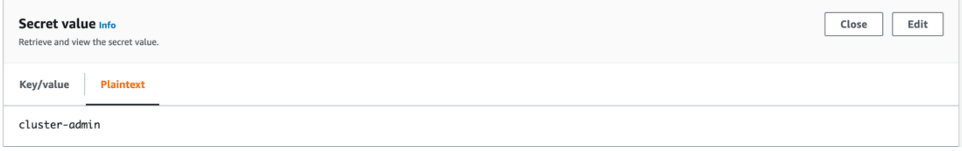

```
cluster-admin
```

Openshift password:

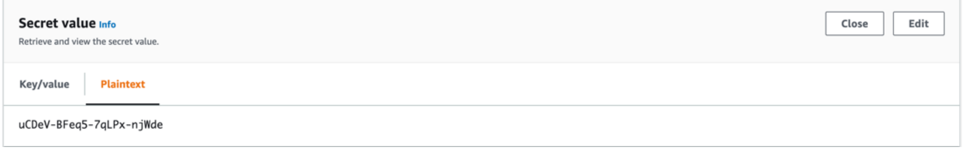

```
uCDeV-BFeq5-7qLPx-njWde
```

Cloud Pak for Data username:

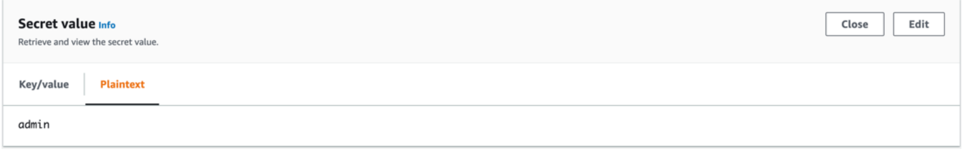

```
admin
```

Cloud Pak for Data password:

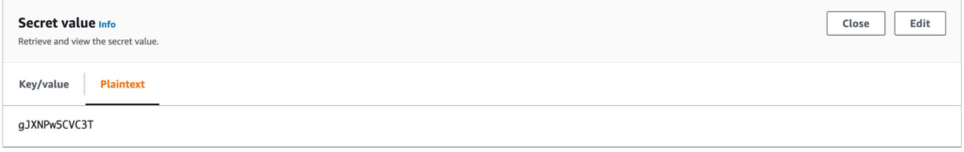

```
gJXNPw5CVC3T
```

Openshift cluster login command:

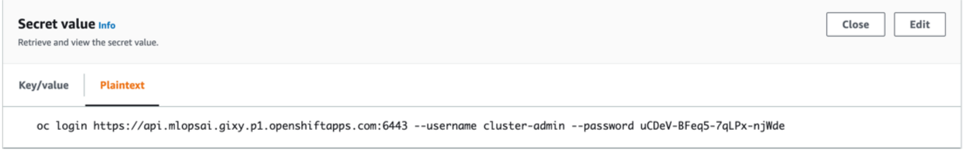

```
oc login https://api.mlopsai.gixy.p1.openshiftapps.com:6443 --username cluster-admin --password uCDeV-BFeq5-7qLPx-njWde
```

4. Log in to the Cloud Pak for Data web client by using the username and the password from the previous step.

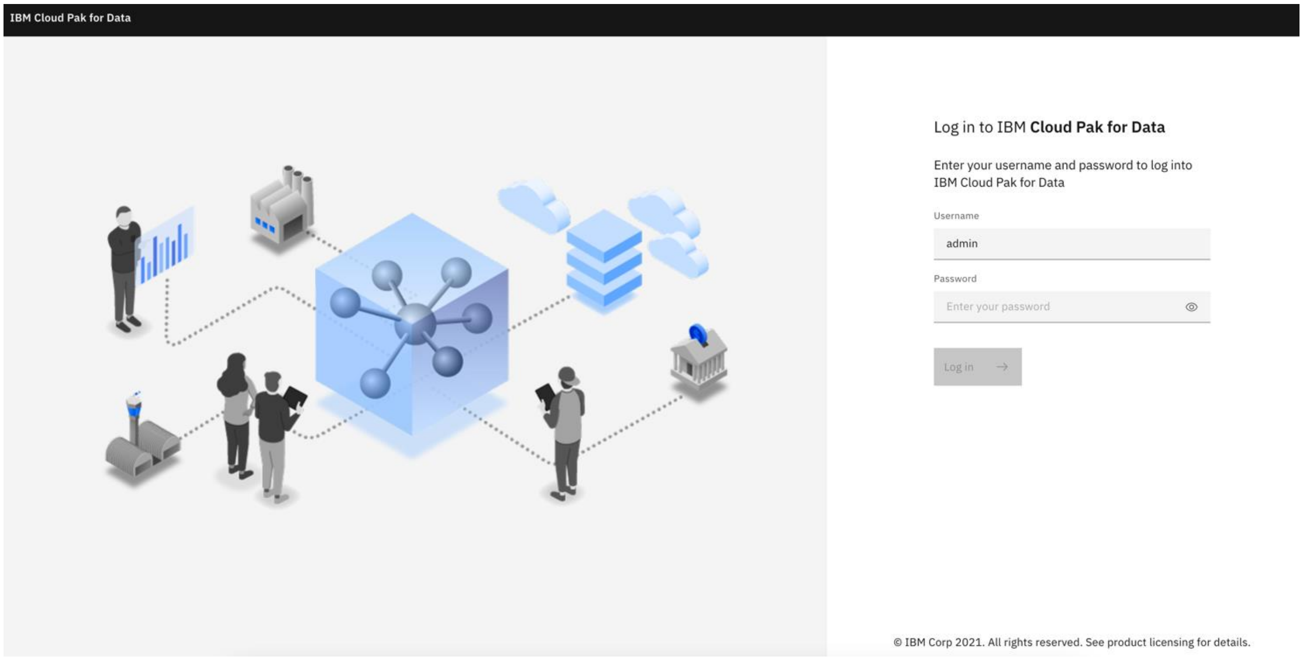
IBM Cloud Pak for Data
Log in to IBM Cloud Pak for Data
Enter your username and password to log into IBM Cloud Pak for Data
Username
Password
Enter your password
Log in →
© IBM Corp 2021. All rights reserved. See product licensing for details.

Username: admin

5. Once you log in, the welcome page is displayed.

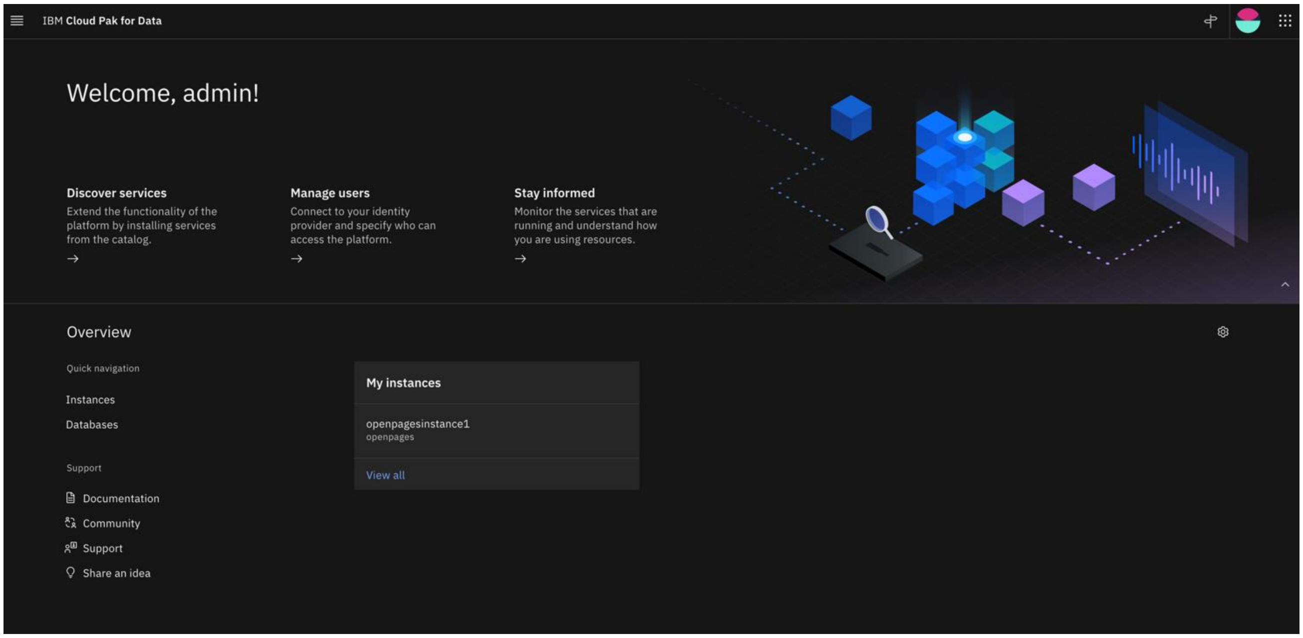
IBM Cloud Pak for Data
Welcome, admin!
Discover services
Extend the functionality of the platform by installing services from the catalog.
→
Manage users
Connect to your identity provider and specify who can access the platform.
→
Stay informed
Monitor the services that are running and understand how you are using resources.
→
Overview
Quick navigation
Instances
Databases
Support
Documentation
Community
Support
Share an idea
My instances
View all

openpagesinstance1
openpages

Figure 9: Welcome page for Cloud Pak for Data web client

See resources on platform features and capabilities. For a list of supported browsers, see Supported browsers.

Step 8. Manage your cluster using the OpenShift Console

1. To access the Openshift Console, go to the OpenShiftURLValue output of the root stack

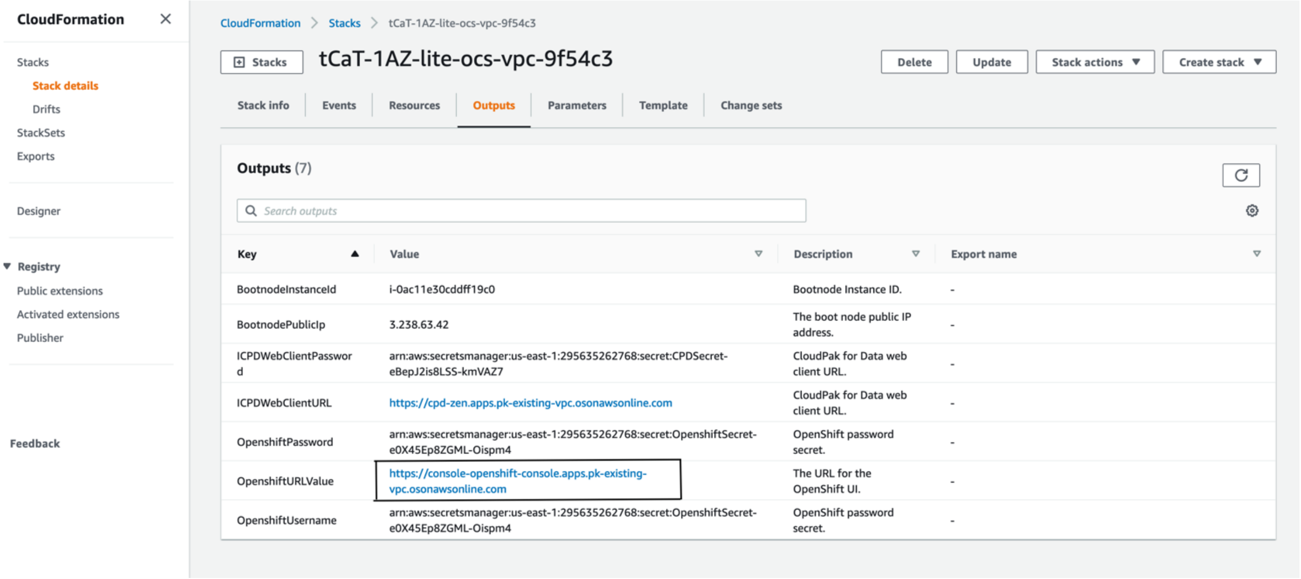

tCaT-1AZ-lite-ocs-vpc-9f54c3

Outputs (7)

Key	Value	Description	Export name
BootnodeInstanceId	i-0ac11e30cddff19c0	Bootnode Instance ID.	-
BootnodePublicIp	3.238.63.42	The boot node public IP address.	-
ICPDWebClientPassword	arn:aws:secretsmanager:us-east-1:295635262768:secret:CPDSecret-eBepJ2is8LSS-kmVAZ7	CloudPak for Data web client URL.	-
ICPDWebClientURL	https://cpd-zen.apps.pk-existing-vpc.osonawsonline.com	CloudPak for Data web client URL.	-
OpenshiftPassword	arn:aws:secretsmanager:us-east-1:295635262768:secret:OpenshiftSecret-e0X45Ep8ZGML-Oispm4	OpenShift password secret.	-
OpenshiftURLValue	https://console-openshift-console.apps.pk-existing-vpc.osonawsonline.com	The URL for the OpenShift UI.	-
OpenshiftUsername	arn:aws:secretsmanager:us-east-1:295635262768:secret:OpenshiftSecret-e0X45Ep8ZGML-Oispm4	OpenShift password secret.	-

The default OpenShift administrative user and password value is retrieved from the previous steps from the secrets value.

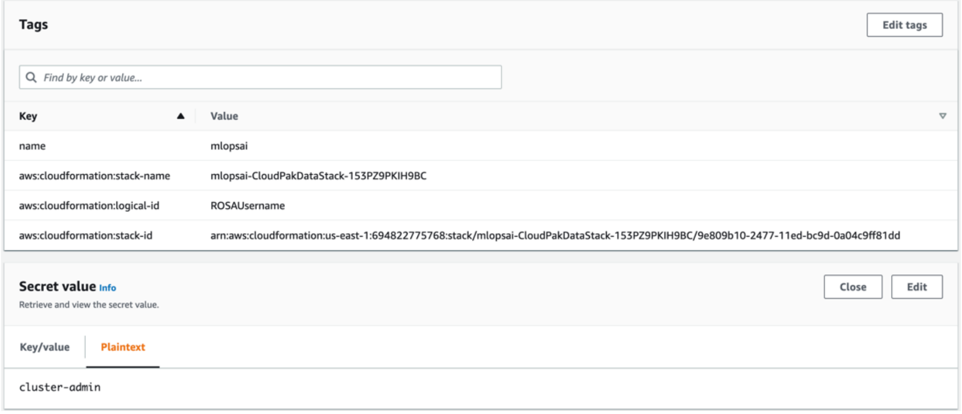

Tags

Key	Value
name	mlopsai
aws:cloudformation:stack-name	mlopsai-CloudPakDataStack-153PZ9PKIH9BC
aws:cloudformation:logical-id	ROSAUsername
aws:cloudformation:stack-id	arn:aws:cloudformation:us-east-1:694822775768:stack/mlopsai-CloudPakDataStack-153PZ9PKIH9BC/9e809b10-2477-11ed-bc9d-0a04c9ff81dd

Secret value Info

Retrieve and view the secret value.

```
cluster-admin
```

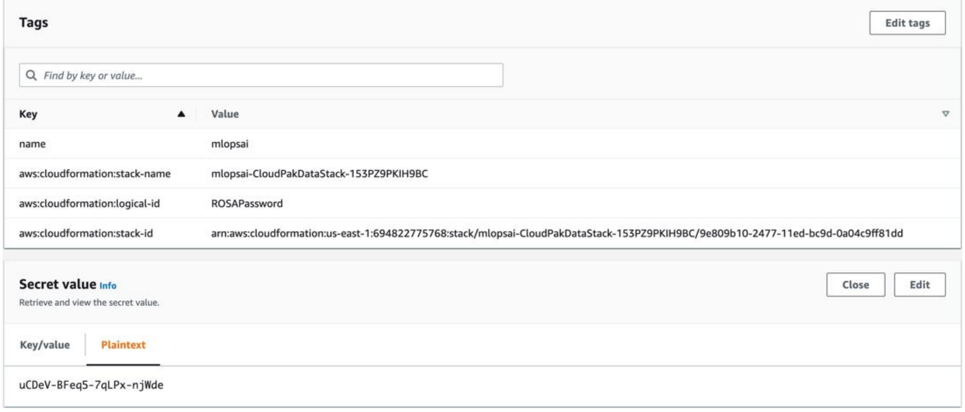

Key	Value
name	mlopsai
aws:cloudformation:stack-name	mlopsai-CloudPakDataStack-153PZ9PKIH9BC
aws:cloudformation:logical-id	ROSAPassword
aws:cloudformation:stack-id	arn:aws:cloudformation:us-east-1:694822775768:stack/mlopsai-CloudPakDataStack-153PZ9PKIH9BC/9e809b10-2477-11ed-bc9d-0a04c9ff81dd

```
uCDeV-BFeq5-7qLPx-njWde
```

Figure 10.1: Retrieve secret value for console password

2. Open the OpenShift Console URL in a browser, select htpasswdProvider at Log in and Login with the username and password from the previous step.

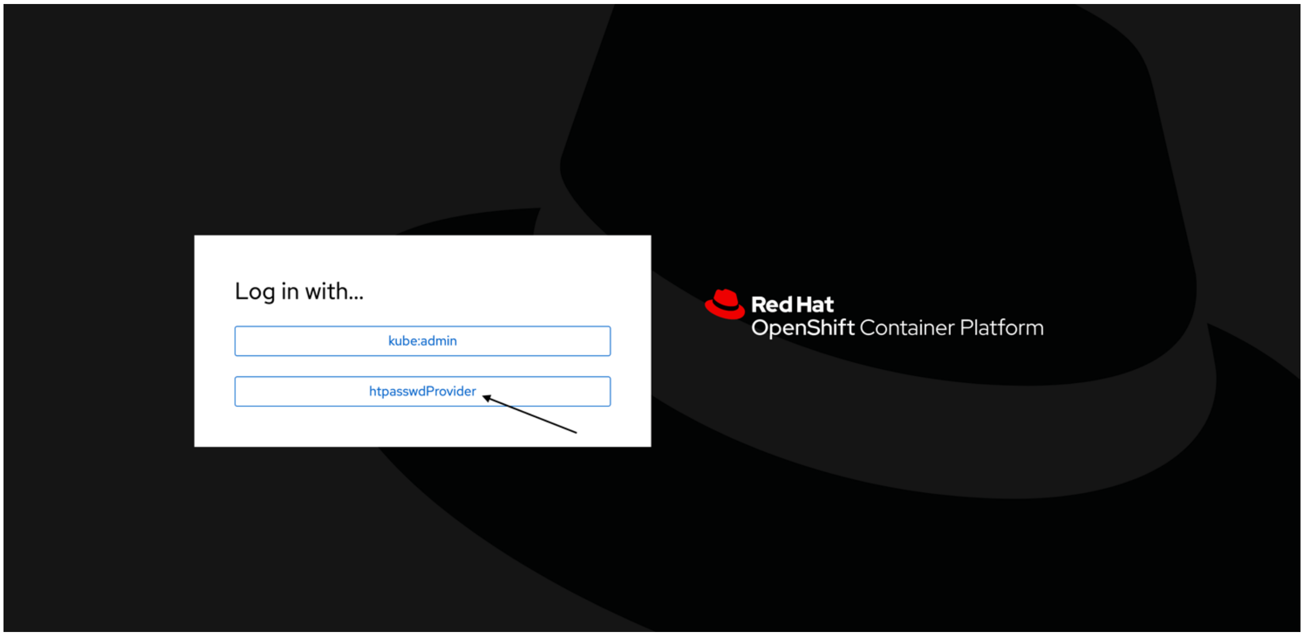

Step 9. (Optional) Provide Boot Node SSH access

The boot node EC2 instance is used for certain command-line cluster administration tasks, such as adding compute nodes. SSH access to the boot node is required for some cluster administrators.

After deployment, you only have access to the boot node. Provide the workstation IP address CIDR as the value of the BootNodeSecurityGroup rule.

This section describes the steps to modify the BootNodeSecurityGroup inbound rules.

> **Note** : These steps assume access to the AWS CloudFormation console for the IBM Cloud Pak for Data deployment.

1. In the AWS Security Groups page, filter name containing BootNodeSecurityGroup

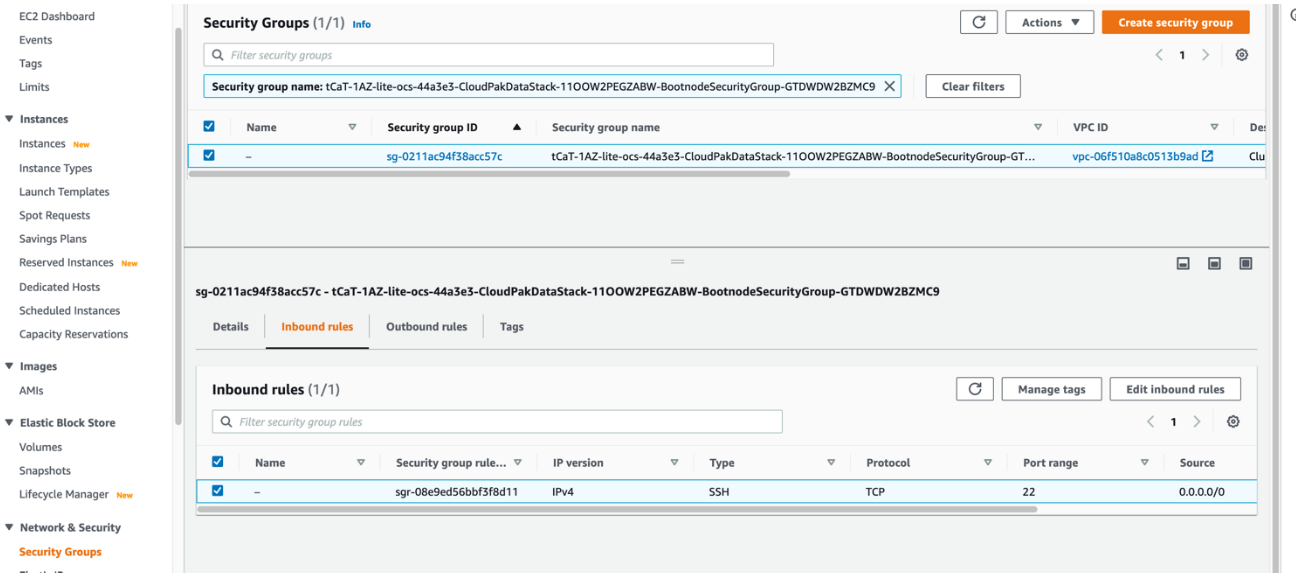

Security Groups (1/1) Info

Security group name: tCaT-1AZ-lite-ocs-44a3e3-CloudPakDataStack-11OOW2PEGZABW-BootnodeSecurityGroup-GTDWDW2BZMC9

☑	Name	Security group ID	Security group name	VPC ID	De
☑	–	sg-0211ac94f38acc57c	tCaT-1AZ-lite-ocs-44a3e3-CloudPakDataStack-11OOW2PEGZABW-BootnodeSecurityGroup-GT…	vpc-06f510a8c0513b9ad	Clu

sg-0211ac94f38acc57c - tCaT-1AZ-lite-ocs-44a3e3-CloudPakDataStack-11OOW2PEGZABW-BootnodeSecurityGroup-GTDWDW2BZMC9

Inbound rules (1/1)

☑	Name	Security group rule…	IP version	Type	Protocol	Port range	Source
☑	–	sgr-08e9ed56bbf3f8d11	IPv4	SSH	TCP	22	0.0.0.0/0

Figure 11: Security Groups

2. The security group window displays the ingress rules. Select the **Inbound** tab, and choose **Edit** to bring up the rule editor, choose **Add Rule**, and fill in the rule details. Add the network CIDR for the group of IP addresses that you want to permit SSH access to the boot nodes. To allow any IP address, use 0.0.0.0/0.

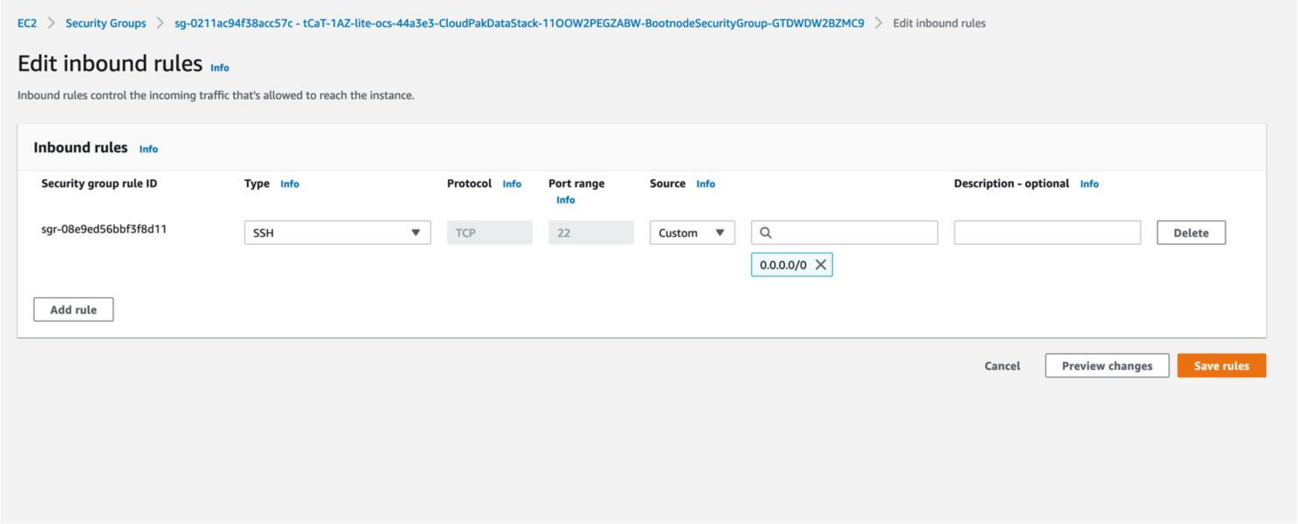

EC2 > Security Groups > sg-0211ac94f38acc57c - tCaT-1AZ-lite-ocs-44a3e3-CloudPakDataStack-11OOW2PEGZABW-BootnodeSecurityGroup-GTDWDW2BZMC9 > Edit inbound rules

Edit inbound rules Info

Inbound rules control the incoming traffic that's allowed to reach the instance.

Inbound rules Info

Security group rule ID	Type	Protocol	Port range	Source		Description - optional	
sgr-08e9ed56bbf3f8d11	SSH	TCP	22	Custom	0.0.0.0/0		Delete

Figure 14: Supplying rule details

3. In the rule editor window, choose **Save Rules**.

Accessing the control plane through the Boot Node

The recommended method of SSH access to the OpenShift cluster instances via the bastion host is by using SSH agent forwarding, as in the following Bash instructions:

7. Run the command `ssh-add -K <your-key.pem>` to store the key in your keychain. On Linux, you might need to omit the -K flag.

8. Retrieve the host name of the Boot Node from the Amazon EC2 console.

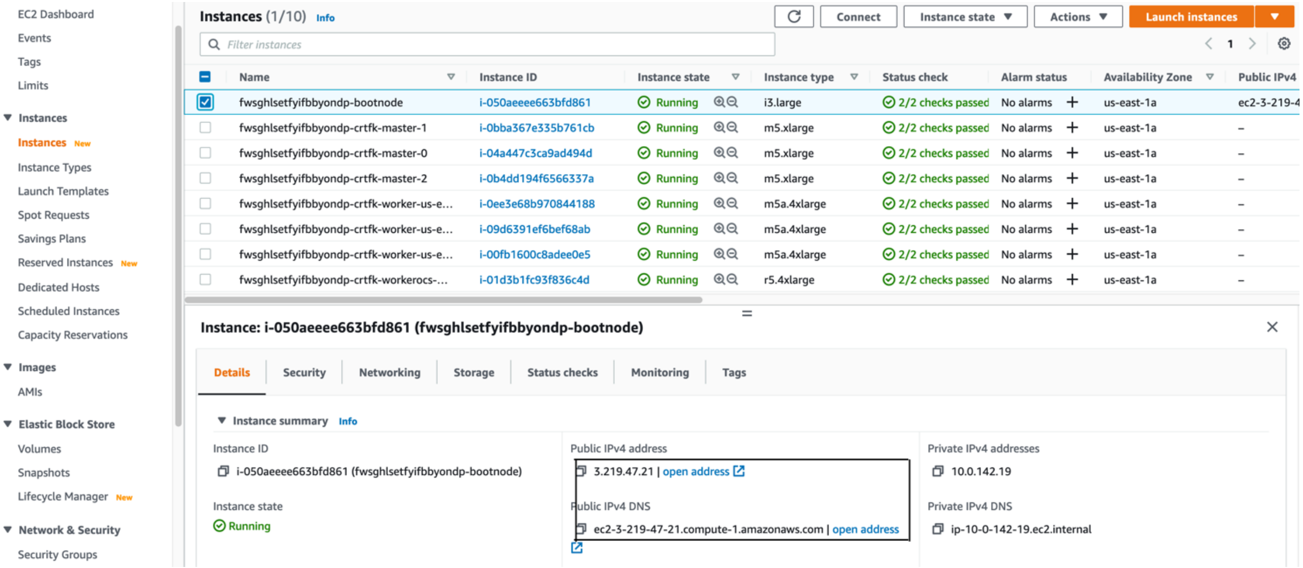

Instances (1/10) Info

	Name	Instance ID	Instance state	Instance type	Status check	Alarm status	Availability Zone	Public IPv4
☑	fwsghlsetfyifbbyondp-bootnode	i-050aeeee663bfd861	Running	i3.large	2/2 checks passed	No alarms	us-east-1a	ec2-3-219-4
□	fwsghlsetfyifbbyondp-crtfk-master-1	i-0bba367e335b761cb	Running	m5.xlarge	2/2 checks passed	No alarms	us-east-1a	–
□	fwsghlsetfyifbbyondp-crtfk-master-0	i-04a447c3ca9ad494d	Running	m5.xlarge	2/2 checks passed	No alarms	us-east-1a	–
□	fwsghlsetfyifbbyondp-crtfk-master-2	i-0b4dd194f6566337a	Running	m5.xlarge	2/2 checks passed	No alarms	us-east-1a	–
□	fwsghlsetfyifbbyondp-crtfk-worker-us-e...	i-0ee3e68b970844188	Running	m5a.4xlarge	2/2 checks passed	No alarms	us-east-1a	–
□	fwsghlsetfyifbbyondp-crtfk-worker-us-e...	i-09d6391ef6bef68ab	Running	m5a.4xlarge	2/2 checks passed	No alarms	us-east-1a	–
□	fwsghlsetfyifbbyondp-crtfk-worker-us-e...	i-00fb1600c8adee0e5	Running	m5a.4xlarge	2/2 checks passed	No alarms	us-east-1a	–
□	fwsghlsetfyifbbyondp-crtfk-workerocs-...	i-01d3b1fc93f836c4d	Running	r5.4xlarge	2/2 checks passed	No alarms	us-east-1a	–

Instance: i-050aeeee663bfd861 (fwsghlsetfyifbbyondp-bootnode)

Instance ID
i-050aeeee663bfd861 (fwsghlsetfyifbbyondp-bootnode)

Instance state
Running

Public IPv4 address
3.219.47.21 | open address

Public IPv4 DNS
ec2-3-219-47-21.compute-1.amazonaws.com | open address

Private IPv4 addresses
10.0.142.19

Private IPv4 DNS
ip-10-0-142-19.ec2.internal

Figure 15: Host name of the Boot Node

9. To log in to the bastion host, run

   $ ssh -I <path-to-private-ssh-key> ec2-user@<bootnode-host-name>

10. Run oc login to authenticate with OpenShift and and verify that services are in a running state:

    $ oc get pods

Scaling up your cluster by adding compute nodes:

RedHat Cluster management url:

rosa describe cluster --cluster <cluster_name> | grep Details

Scale up worker nodes:

$ rosa list machinepools --cluster=<cluster_name>
$ rosa edit machinepool --cluster=<cluster_name> <machinepool_ID> --replicas=<number>

> **Note :** If you choose to scale down your cluster or reduce the number of compute nodes, there is a risk of the cluster becoming unstable because pods will need to be rescheduled. Scaling down the worker nodes is not a recommended option.

Cloud Pak for Data services

You can browse the various services that are available for use by navigating to the Cloud Pak for Data Service Catalog in Cloud Pak for Data.

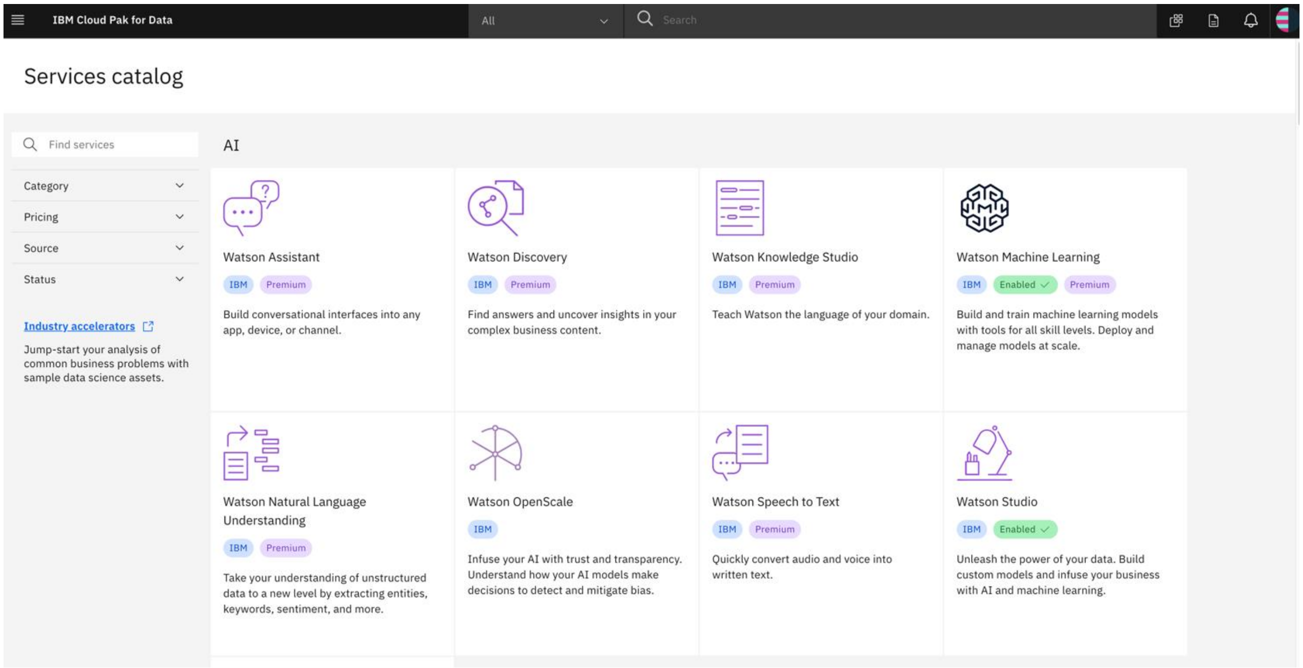

Jump-start your analysis of common business problems with sample data science assets.

Watson Assistant — IBM, Premium
Build conversational interfaces into any app, device, or channel.

Watson Discovery — IBM, Premium
Find answers and uncover insights in your complex business content.

Watson Knowledge Studio — IBM, Premium
Teach Watson the language of your domain.

Watson Machine Learning — IBM, Enabled ✓, Premium
Build and train machine learning models with tools for all skill levels. Deploy and manage models at scale.

Watson Natural Language Understanding — IBM, Premium
Take your understanding of unstructured data to a new level by extracting entities, keywords, sentiment, and more.

Watson OpenScale — IBM
Infuse your AI with trust and transparency. Understand how your AI models make decisions to detect and mitigate bias.

Watson Speech to Text — IBM, Premium
Quickly convert audio and voice into written text.

Watson Studio — IBM, Enabled ✓
Unleash the power of your data. Build custom models and infuse your business with AI and machine learning.

Figure 16: Services catalog page in Cloud Pak for Data

As part of this installation, the control plane is installed by default.

System requirements for the Cloud Pak for Data:
Please refer system requirements from here.

Note: The resources listed are only for the services and does not include the resources for Cloud Pak for Data control plane and foundational services.

Upgrade Cloud Pak for Data Services

See what new features and improvements are available in the latest release of IBM® Cloud Pak for Data.

- Login to your bootnode server as mentioned in access section.
- Follow the upgrade instructions and upgrade corresponding services

Limitations

- Review the known issues and limitations for Cloud Pak for Data.

Troubleshooting

Q. I encountered a CREATE_FAILED error when I launched the CF Template

A. If AWS CloudFormation fails to create the stack, we recommend that you relaunch the template with **Rollback on failure** set to **No**. (This setting is under **Advanced** in the AWS CloudFormation console, **Options** page.) With this setting, the stack's state will be retained and the instance will be left running, so you can troubleshoot the issue. (Look at the log files in %ProgramFiles%\Amazon\EC2ConfigService and C:\cfn\log.)

> **Important** When you set **Rollback on failure** to **No**, you will continue to incur AWS charges for this stack. Make sure to delete the stack when you finish troubleshooting.

For additional information, see [Troubleshooting AWS CloudFormation]() on the AWS website.

Q. I encountered a size limitation error when I deployed the AWS CloudFormation templates.

A. We recommend that you launch the CF templates from the links in this guide or from another S3 bucket. If you deploy the templates from a local copy on your computer or from a non-S3 location, you might encounter template size limitations when you create the stack. For more information about AWS CloudFormation limits, see the [AWS documentation]().

Q. I am not able to delete the stack that I created from the CF Template

A. If any manual changes are done to the resources created from the automation, such as adding new rules to the security group, they must be manually removed for the automation to uninstall successfully. To clean up a deletion failed cluster: a) Storage Cleanup: Based on the storage chosen, either EFS-EBS or FSXdelete the corresponding storage The file system name prefix would be same as the cluster name. For EFS – Delete the file system with the name as _cpd_efs

For NetApp FSX: a) Delete the volumes corresponding to the file system.The FSX file system name would be your cluster name

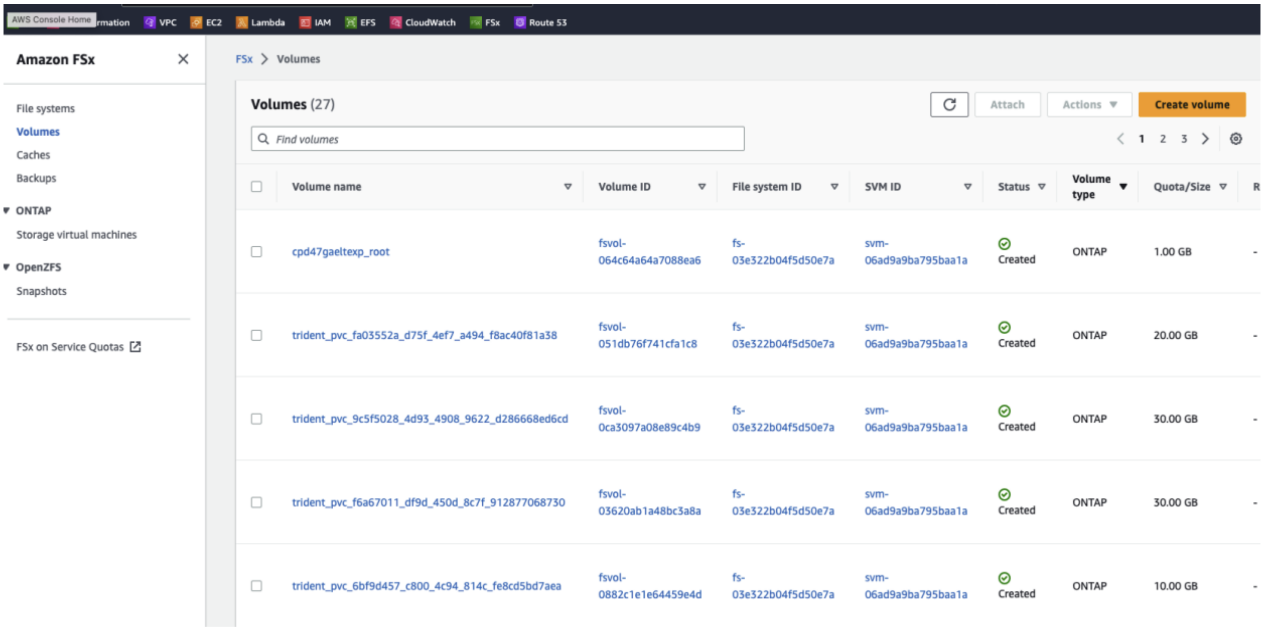

Volumes (27)

□	Volume name	Volume ID	File system ID	SVM ID	Status	Volume type	Quota/Size	R
□	cpd47gaeltexp_root	fsvol-064c64a64a7088ea6	fs-03e322b04f5d50e7a	svm-06ad9a9ba795baa1a	✓ Created	ONTAP	1.00 GB	-
□	trident_pvc_fa03552a_d75f_4ef7_a494_f8ac40f81a38	fsvol-051db76f741cfa1c8	fs-03e322b04f5d50e7a	svm-06ad9a9ba795baa1a	✓ Created	ONTAP	20.00 GB	-
□	trident_pvc_9c5f5028_4d93_4908_9622_d286668ed6cd	fsvol-0ca3097a08e89c4b9	fs-03e322b04f5d50e7a	svm-06ad9a9ba795baa1a	✓ Created	ONTAP	30.00 GB	-
□	trident_pvc_f6a67011_df9d_450d_8c7f_912877068730	fsvol-03620ab1a48bc3a8a	fs-03e322b04f5d50e7a	svm-06ad9a9ba795baa1a	✓ Created	ONTAP	30.00 GB	-
□	trident_pvc_6bf9d457_c800_4c94_814c_fe8cd5bd7aea	fsvol-0882c1e1e64459e4d	fs-03e322b04f5d50e7a	svm-06ad9a9ba795baa1a	✓ Created	ONTAP	10.00 GB	-

b) Delete the Storage volume machine

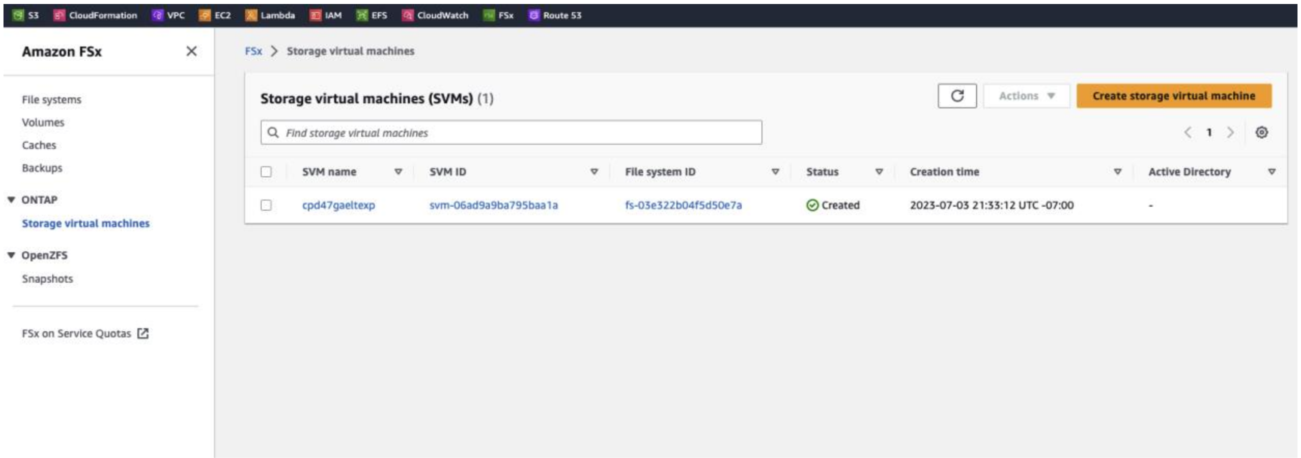

Storage virtual machines (SVMs) (1)

	SVM name	SVM ID	File system ID	Status	Creation time	Active Directory
□	cpd47gaeltexp	svm-06ad9a9ba795baa1a	fs-03e322b04f5d50e7a	Created	2023-07-03 21:33:12 UTC -07:00	-

c) Delete the NetApp FSX File system

b) ROSA cluster deletion: Navigate to the rosa console, Select View only my clusters & identify the cluster with name same as the cluster name value supplied in the cloud formation cluster name parameter and uninstall the cluster.

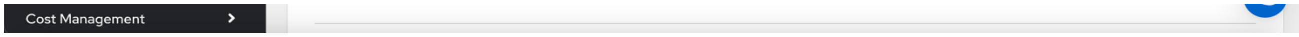

c) Once the cluster uninstall is completed, attempt deletion of cloud formation stack.

Q. The stack creation fails with error Failed to retrieve attribute [Value] for resource [CPDURL]
A. If the Stack creation fails with the error “Failed to retrieve attribute [Value] for resource [CPDURL]: The security token included in the request is invalid (Service: AmazonSSM; Status Code: 400; Error Code: UnrecognizedClientException;”)

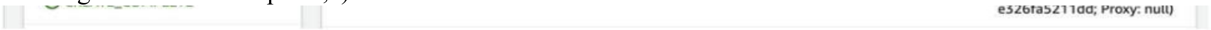
e526fa5211dd; Proxy: null)

The cluster creation is successful, even though the stack creation failed. You can fetch the cluster details OpenshiftURLValue, OpenshiftUsername, OpenshiftPassword , Openshiftlogincommand, ICPDWebClientURL, ICPDWebClientUsername, ICPDWebClientPassword from the Resources

edHandle

Additional resources

AWS resources

- Getting Started Resource Center
- AWS General Reference
- AWS Glossary

AWS services

- Amazon EC2
- Amazon S3

- Amazon VPC
- AWS CloudFormation

IBM Cloud Pak for Data documentation

- IBM Documentation

Redhat OpenShift Service on AWS:

https://docs.openshift.com/rosa/rosa_getting_started/rosa-sts-getting-started-workflow.html

Document revisions

Date	Change	In sections
February 2024	AWS Marketplace CP4D on ROSA BYOL refresh	
July 2024	AWS Marketplace CP4D on ROSA BYOL v5.0.x refresh	