

IBM Cloud Pak for Data on Azure RedHat OpenShift

Deployment Guide

June 2025

Overview.....	3
Cost and licenses.....	3
Architecture	4
Planning the deployment.....	5
Specialized knowledge	5
Technical requirements	5
Deployment options	6
Pre-requisites	6
Step 1. IBM Cloud Pak for Data Subscription.....	6
Step 2. Red Hat Subscription	7
Step 3. Storage Subscription	7
Step 4. Sign in to your Azure portal account	7
Launch the Deployment	7
PARAMETERS FOR DEPLOYING CLOUD PAK FOR DATA INTO A NEW OR EXISTING VIRTUAL NETWORK	8
<i>Basic configuration:</i>	8
<i>Infrastructure configuration:</i>	9
<i>Virtual Network configuration:</i>	9
<i>OpenShift hosts configuration</i>	9
<i>Storage Configuration:</i>	11
<i>IBM Cloud Pak for Data configuration:</i>	11
<i>IBM Cloud Pak for Data Base Services</i>	12
Generating the Portworx Spec URL for ARO cluster	15
Manage your cluster using the OpenShift Console	19
Login to Cloud Pak for Data web client	21
(Optional) Provide Boot Node SSH access	24
Scaling up your cluster by adding compute nodes	25

Cloud Pak for Data services	25
Cloud Pak for Data minimum System and Services requirements	27
Install a Cloud Pak for Data Service	27
Upgrade Cloud Pak for Data Services	27
Limitations	28
Additional resources	28
Document revisions	28

Overview

This deployment guide provides step-by-step instructions for deploying IBM Cloud Pak for Data on a [Red Hat OpenShift Container Platform cluster on Azure VMs](#) or on an [Azure RedHat Managed Openshift \(ARO\)](#) Azure Marketplace. You can automatically deploy a production instance of Cloud Pak for Data.

Cloud Pak for Data is a data and 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.

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 Cloud Pak for Data on Azure.

This reference deployment provides Azure ARM templates to deploy Cloud Pak for Data on a new OpenShift cluster. This cluster includes:

- A Red Hat OpenShift Container Platform cluster created in a new or existing Vnet 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.
- A highly available storage infrastructure with Red Hat OpenShift Container Storage.
- Scalable OpenShift compute nodes running Cloud Pak for Data services. For more information about Cloud Pak for Data, see the [IBM Documentation](#).

Cost and licenses

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

The Azure ARM template for this deployment includes configuration parameters that you can customize. You can use it to build a new Vnet for your Cloud Pak for Data on Azure cluster or deploy on an existing Azure Vnet. Some of these settings, such as instance type, will affect the cost of deployment. For cost estimates, see the pricing pages for each Azure service you will be using. Prices are subject to change.

For more information about licensing terms, see the 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 <https://www-03.ibm.com/software/sla/sladb.nsf/search?OpenForm>, execute a search for 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 4.0.5 and there is no LI for this version, you will need to review the LI for version 4.0)

For more information on License and Agreement for Cloud Pak for Data Base and Cartridge Services:

https://cp4dazuremarketplacedoc.blob.core.windows.net/ibmcp4dazuremarketplace/CPD_CTG_ARO_EULA_v48.pdf

Architecture

Deploying the Azure Marketplace template for a new VNet with **default parameters** builds the following Cloud Pak for Data environment in the Azure Cloud.




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

The ARM template sets up the following:

- A highly available architecture that spans up to three Availability Zones.*
- A Virtual network configured with public and private subnets.*
- In the public subnets:
 - a bastion host to allow inbound Secure Shell (SSH) access to compute instances in private subnets.
- In the private subnets:
 - OpenShift Container Platform master instances.
 - OpenShift Container Platform (OCP) compute nodes that combined, contain Cloud Pak for Data Collect, Organize, and Analyse services.
- An Azure Load Balancer spanning the public subnets for accessing Cloud Pak for Data from a web browser.
- Storage disks with Azure Managed Disk mounted on compute nodes for ODF (OpenShift Data Foundation) v4.17 or Portworx Enterprise.
- An Azure domain as your public Domain Name System (DNS) zone for resolving domain names of the IBM Cloud Pak for Data management console and applications deployed on the cluster.

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

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




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 Azure services.

Technical requirements

For Cloud Pak for Data requirements, see [Planning 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 [Azure documentation](#).

	This deployment uses	Resource
Virtual Network	1	
Public IP addresses	3	
Network Load Balancers	2	
Standard_D4s_v3 virtual machines (Bootnode)	1	
Standard_D8s_v3 virtual machines (Master nodes)	3	
Standard_D16s_v3 virtual machines (Compute nodes)	3	
Standard_D16s_v3 virtual machines (Compute nodes)	3(ODF)	
Public IP addresses for ODF storage	2	
<u>Regions</u>	This deployment includes 3 Availability Zones, which isn't currently supported in all Azure Regions.	
<u>IAM permissions</u>	To deploy the template, you must log in to the Azure portal with Azure Identity and Access Management (IAM) permissions for the resources and actions the templates will deploy. The <i>AdministratorAccess</i> managed policy within IAM provides sufficient permissions, although your organization may choose to use a custom policy with more restrictions.	

Deployment options

This template provides the following deployment options:

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

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


Pre-requisites

Ensure the following pre-requisites are in place with your existing Cloud Pak for Data entitlements.

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 ARM Template.

Note: Please make sure you have valid entitlement key before using the offering, you can check your entitlement key is valid or not [here](#). Login to myibm.ibm.com and refer to the name(s) below when accessing your container software as shown in the below screenshot.



Step 2. Red Hat Subscription

Ensure that you have the Red Hat [OpenShift pull secret](#) with your purchase. You should keep it handy as this json is required parameter in the ARM Template. If you are deploying through Quickstart Template or from Azure Marketplace, you can directly enter the pull secret's json value.

You can select one of the three container storages while installing this Quick Start.

Step 3. Storage Subscription

- OpenShift Data Foundation (ODF): The Red Hat ODF license is linked as a separate entitlement to your RedHat subscription. This is the preferred option on Azure Marketplace.
- Portworx: When you select [Portworx](#) as the persistent storage layer, you will need to specify the install spec from your [Portworx account](#). You can generate a new spec using the [Spec Generator](#). Note that the Portworx trial edition expires in 30 days after which you need to upgrade to an Enterprise Edition. [Here are the steps to generate Portworx Spec.](#)

Step 4. Sign in to your Azure portal account

1. Sign in to your Azure account at <https://portal.azure.com/> with an Azure 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 Azure 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 Azure Region where you want to deploy Cloud Pak for Data on Azure. An IBM Cloud Pak for Data high availability deployment is restricted to Azure Regions with at least three Availability Zones.
4. The following resources should be made available for Cloud Pak for Data deployment
 - A new or an existing Key Vault location with Red Hat pull secret

Launch the Deployment

Note: The instructions in this section reflect the current version of the Azure portal. If you're using the redesigned portal, some of the user interface elements might be different.

You are responsible for the cost of the Azure services used while running this deployment. For full details, see the pricing pages for each Azure service you will be using for this deployment. Prices are subject to change.

1. Launch the Azure ARM template into your Azure account from the [IBM Cloud Pak for Data on ARO - Azure Marketplace](#) or 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, and a high availability (HA) deployment takes about 4 hours to complete.
2. Create a new Resource Group or specify existing one. This is where all the deployment resources will be stored. There will be two resource group created through ARM templates as specified in the parameters section below. **Note:** Resource group should be unique and same name shouldn't be specified it earlier. The ARM template may fail if the same name resource group created earlier in the past.
3. Check the Region that's displayed in the Region field 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 'West US 2' Region by default.
4. Specify Service Principal App ID (Client ID) and Azure AD Client Secret. For more information check [how to create Service Principal](#) on Azure. It has steps to create service principal and get the client id and client secret.
5. For the Infrastructure settings, OpenShift settings, and Cloud Pak settings pages, 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.

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

PARAMETERS FOR DEPLOYING CLOUD PAK FOR DATA INTO A NEW OR EXISTING VIRTUAL NETWORK

Basic configuration:

Parameter label (name)	Default	Description
Subscription	Microsoft Azure Enterprise	The subscription is necessary to deploy cluster and create resources on the Azure portal. Make sure you have required administrator permission to create cluster resources, secrets, users, etc.

Resource Group	<i>Requires input</i>	Resource group is a container that holds related resource for an Azure Solution. Create a new resource group with the unique name (never used it earlier). There will be two Resource group created one with the name of the Resource Group and another one with the cluster name. The Cluster name Resource Group will hold all master and worker nodes related resources. E.g., Virtual Machines, Network Security Gateway (NSG), etc.
Region	West US 2	This is where the network infrastructure for Cloud Pak for Data will be built. The template is launched in the 'West US 2' Region by default.
Service Principal Azure App ID	—	This client ID will be used to create resources on the Azure portal. Check how to create Service Principal Azure App ID
Azure AD Client Secret	—	This client secret will be used to create resources on the Azure portal. Check how to get Azure Client Secret

Infrastructure configuration:

Parameter label (name)	Default	Description
Bootnode Public IP (Attach Public IP to BootnodeVM)	true	Bootnode can be accessible through SSH connection if bootnode has the public IP.
SSH public key	<i>Requires input</i>	Your machine's SSH public key will be added in the authorized_key in the bootnode so that it can be accessible through public IP.
Number of master nodes (NumberOfMaster)	3	The desired capacity for the OpenShift master instances. Must be an odd number. For a development deployment, 1 is sufficient; for production deployments, a minimum of 3 is required.
Number of worker nodes (NumberOfNodes)	3	The desired capacity for the OpenShift node instances. Minimum of 3 nodes required. Warning If the number of node instances exceeds your Red Hat entitlement limits or Azure virtual machine quotas, the stack will fail. Choose a number that is within your limits.
Bootnode VM size (BootnodeType)	Standard_D4s_v3	The virtual machine type for the OpenShift bootnode VM.
Master VM size (MasterInstanceType)	Standard_D8s_v3	The virtual machine type for the OpenShift master VMs.
Compute VM size (NodesInstanceType)	Standard_D16s_v3	The virtual machine type for the OpenShift compute VMs.

Virtual Network configuration:

Parameter label (name)	Default	Description
Virtual Network	<i>Requires input</i>	Create a new Virtual Network or select existing VNet
BootNode Subnet	<i>Requires input</i>	Subnets for virtual network
(Bootnode Subnet CIDR)		
Master Subnet	<i>Requires input</i>	Subnets for virtual network
(Master Subnet CIDR)		
Worker Subnet	<i>Requires input</i>	Subnets for virtual network
(Worker Subnet CIDR)		

Single or Multi Zone	Multi Zone	Deploy VMs to Single Zone region or Multiple Zone (Availability Zones). Recommended option is MultiZone. Please make sure that selected region supports AvailabilityZones: https://learn.microsoft.com/en-us/azure/reliability/availability-zones-service-support#azure-regions-with-availability-zone-support
-----------------------------	------------	--

OpenShift hosts configuration:

Parameter label (name)	Default	Description
RedHat subscription pull secret. (RedhatPullSecret)	<i>Requires input</i>	Enter json value of the RedHat Openshift Pull Secret, if you are deploying cluster through Azure Marketplace or Azure Quickstart template. Note: If you are using internal Github quickstart/marketplace repo to deploy the cluster then use Key vault path of OpenShift Installer Provisioned Infrastructure pull secret. e.g., <pre>"reference": { "keyVault": { "id": "/subscriptions/<YOUR-SUBSCRIPTION-ID>/resourceGroups/<RESOURCE-GROUP-OF-YOUR-KEY-VAULT>/providers/Microsoft.KeyVault/vaults/<YOUR-KEY-VALUE-NAME>" }, "secretName": "pullsecret" }</pre>
Cluster prefix (ClusterName)	<i>Requires input</i>	Custom cluster name for kubernetes.io/cluster/tags. The cluster name should be <i>unique</i> and should never be used earlier because this name will be also assigned it to another cluster resource group. This cluster resource group will hold all node related resources as mentioned in the Resource Group parameter.

ARO Resource Provider Principal ID (Only for IBM CPD on ARO - BYOL)	ARO cluster deployment needs Principal ID of the ARO Service Principal exist in your account in order to create ARO cluster in the given Azure account. Here is the command to fetch the Principal ID (ObjectId) of the ARO Service Principal, Azure CLI command: (Prereqs: Azure CLI installation in your computer.) <pre>`az ad sp list --filter "displayname eq 'Azure Red Hat OpenShift RP'" --query "[?appDisplayName=='Azure Red Hat OpenShift RP'].{name: appDisplayName, objectId: id}"`</pre> E.g.: <pre>[user]@[hostname] % az ad sp list --filter "displayname eq 'Azure Red Hat OpenShift RP'" --query "[?appDisplayName=='Azure Red Hat OpenShift RP'].{name: appDisplayName, objectId: id}"</pre> This command or command group has been migrated to Microsoft Graph API. Please carefully review all breaking changes introduced during this migration: https://docs.microsoft.com/cli/azure/microsoft-graph-migration <pre>[{ "name": "Azure Red Hat OpenShift RP", "objectId": "464114b9-XXXX-XXXX-a068-35ab76dXXXXX" }]</pre> Please make sure you have “Network Contributor” role assigned to ARO Resource Provider Principal ID Note: If this Azure CLI command doesn’t work for you then you may need to register the required resource provider. Please follow instructions on this page and execute “<i>Get the service principal object ID for the Openshift resource provider – Azure CLI</i>” command: https://learn.microsoft.com/en-us/azure/openshift/quickstart-openshift-arm-bicep-template?pivot=aro-arm#register-the-required-resource-providers---azure-cli
Use Private or Public Endpoints (PublicCluster)	public To Deploy a private cluster, select “private” and “public” for public cluster.
Enable Machine Autoscaler	false Enable Machine Autoscaler to automate the scale up and down the cluster based on requirement dynamically.
Egress Outbound Type	Load Balancer Choose value of Egress Outbound type values either Load Balancer and User Defined Routing.

Requires input

Storage Configuration:

Parameter label (name)	Default	Description
Storage type for Cluster (StorageType)	<i>ODF or Portworx</i>	OpenShift Data Foundation (ODF), and Portworx storage options are available. ODF is recommended storage option.
IBM Cloud Pak for Data on ARO for Azure Marketplace Standard_Ops_t December 2024		
Cloud Instance Type (ODFInstanceType)		Update this value if Storage type selected is ODF. The Virtual Machine type for the ODF instances.
Number of ODF nodes (NumberOfODF)	3	The desired capacity for the ODF instances. Minimum of 3 is required. You don't need to add these 3 ODF nodes in worker nodes count as our template will add 3 ODF nodes for you. Note: These 3 nodes will only be created for ODF but not for the Portworx storage option.
Portworx Spec URL	<i>Requires input (if Portworx is selected as a storage option.)</i>	Required field Only if you chose Portworx as storage option. Generated Spec URL for ARO cluster from Portworx spec generator: https://central.portworx.com/specGen/wizard . Here are the steps to generate Portworx Spec.
Portworx CSI	true	Required field Only if you chose Portworx as storage option. Portworx CSI option to enable CSI environment. Default value for Portworx CSI is true. It is recommended to enable Portworx CSI option.
<i>Cloud Pak for Data configuration:</i>		
Parameter label (name)	Default	Description
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>Reject</i>	I have read and agreed to the license terms for IBM Cloud Pak for Data that were provided to me at time of purchase. <i>You must accept the license to install Cloud Pak for Data services.</i>
OpenShift project (NameSpace)	zen	The OpenShift project that will be created for deploying Cloud Pak for Data. It can be any lowercase string.
IBM Cloud Pak for Data version (CPDVersion)	5.2.x	The default version of Cloud Pak for Data to be deployed.
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

IBM Cloud Pak for Data Base and Cartridge Services:

Parameter label (name)	Default	Description
IBM Knowledge Catalog (WKC)	False	Choose True to install the IBM Knowledge Catalog service.
Watson Machine Learning service (WML)	False	Choose True to install the Watson Machine Learning service.
Data Management Console (DMC)	False	Choose True to install the Data Management Console service.
Watson Query (DV)	False	Choose True to install the Watson Query (Data Virtualization) service.
Watson Studio (WSL)	False	Choose True to install the Watson Studio Local service.
Watson OpenScale and Watson Machine Learning services (WOS+WML)	False	Choose True to install the Watson OpenScale and Watson Machine Learning services.
Analytics Engine (AE)	False	Choose True to install the Analytics Engine (Apache Spark) service.
Cognos Dashboard (CDE)	False	Choose True to install the Cognos Dashboard Engine service
DataStage (DS)	False	Choose True to install DataStage service
Db2 (Db2oltp)	False	Choose True to install Db2 service
Cognos Analytics (CA)	False	Choose True to install Cognos Analytics service
Db2 Warehouse (Db2wh)	False	Choose True to install Db2 Warehouse service
Decision Optimization (DO)	False	Choose True to install Decision Optimization service
Match360 with Watson (MDM)	False	Choose True to install Match360 with Watson service
Planning Analytics (PA)	False	Choose True to install Planning Analytics service
OpenPages (OpenPages)	False	Choose True to install OpenPages service
RStudio Server Runtimes (RStudio)	False	Choose True to install RStudio server runtime service
Db2 BigSQL (BigSQL)	False	Choose True to install Db2 BigSQL service
Watson Pipelines (WS Pipelines)	False	Choose True to install Watson Pipelines service
AI Factsheets (Factsheet)	False	Choose True to install AI Factsheet service
Watson Discovery (WD)	False	Choose True to install Watson Discovery service
Watsonx Assistant (WA)	False	Choose True to install Watsonx Assistant service (Supports only in Single Zone cluster)
Watson Speech (Watson Speech)	False	Choose True to install Watson Speech Service
Manta with IBM Knowledge Catalog (Mantaflow)	False	Choose True to install Manta with IBM Knowledge Catalog service

1. On the Openshift and Cloud Pak Settings page, select appropriate values from the above table.

2. On the **Review + Create** page, review and confirm the template settings. Make the necessary changes based on your requirement before deploying template.
3. Choose **Create** to deploy the templates.
4. On Azure Portal, monitor the status of the templates *Resource Group > Deployments*. When the status is succeeded for the CPD deployments, the Cloud Pak for Data cluster is ready. 5. Use the URLs displayed in the *Resource Group > Deployments > AzureRMSamples > Outputs*. The URL for the “**Cloud Pak for Data URL**” output key will navigate to the console login page.

Steps to Create Azure Service Principal Client ID and Client Secret

- Create an Azure Service Principal with Contributor and User Access Administrator roles.
 - Create a Service Principal, using your Azure Subscription ID, named with a valid SP_NAME (e.g. AROClusterServicePrincipal) and save the returned json:

```
az login
```

```
az ad sp create-for-rbac --role="Contributor" --name="<SP_NAME>" --  
scopes="/subscriptions/<subscription_id>"
```

- Assign the User Access Administrator role, using the AppId:

```
az role assignment create --role "User Access Administrator" --assignee "<app_id>" --  
scopes="/subscriptions/<subscription_id>"
```

Note: Save “*appId*” as a *ClientID* and “*password*” as a *Client Secret*. These values are needed for the deployment.

If you face any issues while creating Azure Service Principal Client ID, Please [refer](#) Microsoft Official Documentation to create Azure Service Principal Client ID and Client Secret

(Optional) Edit the Network Security Group

Optional: You might need to edit the Azure network security group to add IP addresses that can access the Cloud Pak for Data web client.

Navigate to Load Balancers on your Azure portal and filter on tags, for example *kubernetes.io/service-name: openshift-ingress/router-default*.

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

Create load balancer

Subscription * Microsoft Azure Enterprise

Resource group * vc-rg [Create new](#)

Instance details

Name * vc-lb

Region * West US 2

SKU * Standard Basic

Microsoft recommends Standard SKU load balancer for production workloads. Learn more about pricing differences between Standard and Basic SKU

Type * Public Internal

Tier * Regional Global

[Next : Frontend IP configuration >](#)

[Review + create](#) [< Previous](#) [Next : Frontend IP configuration >](#) [Download a template for automation](#) [Give feedback](#)

2. Select Security Group and modify the Inbound rules

vc-az-nsg Network security group

Search (Cmd+J)

Activity log Access control (IAM) Tags Diagnose and solve problems

Settings

- Inbound security rules
- Outbound security rules
- Network interfaces
- Subnets
- Properties
- Locks

Monitoring

- Alerts
- Diagnostics settings
- Logs
- NSG flow logs

→ Move Delete Refresh Give feedback

Essentials [JSON View](#)

Resource group (Move) vc-rg

Location West US 2

Subscription (Move) Microsoft Azure Enterprise

Subscription ID 6f046466-4b41-405b-b34c-c8992178e4fc


Tags (Edit) [Click here to add tags](#)

Filter by name

Port == all Protocol == all Source == all Destination == all Action == all

Priority	Name	Port	Protocol	Source	Destination
Inbound Security Rules					
Outbound Security Rules					
65000	AllowVnetOutBound	Any	Any	VirtualNetwork	VirtualNetwork
65001	AllowInternetOutBound	Any	Any	Any	Internet
65500	DenyAllOutBound	Any	Any	Any	Any


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.




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

Generating the Portworx Spec URL for ARO cluster


1. Launch the [spec generator](#)
2. Select **Portworx Enterprise** and press *Continue*:



3. Select **Portworx Essentials** or **Portworx Enterprise** and press *Continue*:



4. In generate Spec page, Select Portworx version **3.1**, Select Platform **Azure** and Select Kubernetes Distribution Name **Azure Red Hat OpenShift (ARO)**, Change Namespace name to **kube-system** and press **Customize**



5. Check Use the Portworx Operator box, select the Portworx version as **3.1**, enter Namespace as **kube-system**, select **Built-in** and then press **Next**. (Note: Please make sure all selected options in Step 4 are reflected on this page)

The screenshot shows the 'Basic' configuration tab for the IBM Cloud Pak for Data on ARO for Azure Marketplace. The 'Use the Portworx Operator' checkbox is checked. The 'Portworx Version' is set to 2.13, 'Kubernetes Version' is 1.25.0, and the 'Namespace' is kube-system. The 'etcd' section shows 'Built-in' is selected. A note states: 'Portworx will create and manage an internal key-value store (etcd) cluster. You can restrict the nodes that will run the key-value store by labelling your nodes with the label px/metadata-node=true. Only the nodes with the label will participate in the etcd cluster. This allows you to use nodes with dedicated hardware for the key-value store. For example: kubectl label nodes node1 node2 node3 px/metadata-node=true'. The 'Next' button is highlighted in orange.


6. Select *Cloud* as your environment. Click on *Azure* and select *Create Using a Spec* option for Select type of disk. Enter value for Size (GB) as **500** and then press *Next*.

The screenshot shows the 'Storage' configuration tab. The 'Select your environment' dropdown is set to 'Cloud'. The 'Select Cloud Platform' dropdown is set to 'Azure'. Under 'Configure storage devices', the 'Select type of disk' dropdown is set to 'Create Using a Spec'. The 'Select volume type' dropdown is set to 'Premium', and the 'Size (GB)' is entered as 500. The 'Encryption' dropdown is set to 'None'. The 'Auto create journal device' checkbox is unchecked. The 'Next' button is highlighted in orange.


7. Keep the default option *auto* for the network interfaces and press *Next*:

The screenshot shows the 'Network' configuration tab. The 'Interface(s)' section has 'Data Network Interface' and 'Management Network Interface' both set to 'auto'. The 'Advanced Settings' section is collapsed. The 'Next' button is highlighted in orange.


8. Select *Azure Red Hat OpenShift (ARO)* as Openshift environment, go to *Advanced Settings*:
(Note: if you don't find the ARO option in the Customize tab, please verify you have selected Portworx 2.12 or later versions as it is only available for Portworx 2.12 or later versions).



9. In the Advanced Settings tab: *Enable Stork, CSI, Monitoring and Telemetry* and press *Finish*:




10. Copy the Spec URL and Save Spec in your Portworx account. (Note: Please copy only spec URL but not the whole command or single quote in order to avoid any errors.)




Manage your cluster using the OpenShift Console

1. To access the OpenShift console, go to the overview section of the root stack.
2. In the overview section you can get **Openshift URL**, **username** and **password**, please select the Azure Redhat openshift cluster deployment from the **overview** section of root stack as shown in the below screenshot:


Home > Resource Groups > Cluster's Resource groups > select Deployments > **ibm-usa-ny- armonk-hq-6275750-ibmcloud-asperia.ibm-20231207151241 > overview**



3. Click on the “Connect” and get the **openshift URL**, **username & password** from the right panel.



4. Open the OpenShift Console URL in a browser and Login with the **username** and **password** from the previous step.



Login to Cloud Pak for Data web client


When the Azure ARM template has successfully created the stack, all server nodes will be running with the software installed in your Azure portal. 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 output for **key Cloud Pak for Data URL**


Home > Resource Groups > Cluster's Resource groups > select Deployments > ibm- alliance-global-1560886.cloud_pak_for_data-20211118184734 > Outputs




- To get the password of Cloud Pak for Data web client, please login in to openshift console as mentioned in the previous step-6 and click on the workloads in the left pane **Workloads > secrets > click on platform-auth-idp-credentials > copy the admin_username & admin_password under Data.**



- Log in to the Cloud Pak for Data web client by using the user “cpadmin” and the password from the previous step.



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




See [resources](#) on platform features and capabilities. For a list of supported browsers, see [Supported browsers](#).

(Optional) Provide Boot Node SSH access


The boot node 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 network security group (nsg) inbound rule.

This section describes the steps to modify the network security group (nsg) inbound rules.

1. In the Azure cluster's **Resource Groups** page, select name containing **bastion-nsg**




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 ***** in the “**Source port ranges**” field.



3. In the rule editor window, click on **Add**.

Scaling up your cluster by adding compute nodes

Login to your OpenShift Console, navigate Compute Machine Sets, each machine set can be scaled up.



- An Azure instance will be created, and Desired count and current count will get updated to replica value.
- After few mins once the node joins the cluster ready and available count will be updated to replica value

Note

1. 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.
2. Cluster auto scaler can overrule the scaling activity to maintain the required threshold.

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.




Figure: Services catalog page in Cloud Pak for Data

As part of this installation, the control plane is installed by default, and the following services can be enabled while creating the stack:

- Watson Studio
- IBM Knowledge Catalog
- Watson Machine Learning
- Data Management Console
- Watson Query (Data Virtualization)
- Watson OpenScale
- Cognos Dashboard
- Analytics Engine for Apache Spark
- DataStage
- Db2
- Cognos Analytics
- Db2 Warehouse
- Decision Optimization
- Match360 with Watson
- OpenPages
- Planning Analytics
- Db2 BigSQL
- RStudio server runtime
- Watson Pipelines
- AI Factsheets
- Watson Discovery
- Watsonx Assistant
- Watson Speech
- Manta with IBM Knowledge Catalog

Cloud Pak for Data minimum System and Services requirements

The overall Cloud Pak for Data platform minimum hardware requirements are available at [Cloud Pak for Data platform hardware requirements](#).

Cloud Pak for Data includes some shared cluster components/services which provide underlying functionality for the Cloud Pak for Data control plane and services. The minimum system requirements for these components/services can be found at [Shared cluster component requirements](#).

The minimum system requirements for the services above can be found at [Service hardware requirements](#) in the Cloud Pak for Data product documentation. Please note that the services list includes all Cloud Pak for Data services, and not just the ones currently supported by Azure deployments. Please use browser search to find the services above in the services list.

Steps to Install any other Cloud Pak for Data supported Service (Not available on Azure Marketplace template)

- Login into your bootnode server
 - Select your {CLUSTER_NAME} Resource Group from Azure Portal > Click on “Overview” page from the left panel > Click on “bastionNode” > Click on “Connect” from Overview page > Follow SSH command to bastionNode.
- Follow the [Prerequisite and Install steps for the selected service](#)
 - Downloaded CPD-CLI file location in the bootnode: `“/mnt/.cpd/templates/cpd-cli”`
 - Here are necessary params value which might be useful for the CPD-CLI commands:
 - \${OCP_USERNAME}: Select your {CLUSTER_NAME} Resource Group from Azure Portal > Click on “Overview” page from the left panel > Click on your Azure Red Hat OpenShift Cluster > Click on “Connect” from the Overview page > Kubeadmin credentials.
 - \${OCP_PASSWORD}: Select your {CLUSTER_NAME} Resource Group from Azure Portal > Click on “Overview” page from the left panel > Click on your Azure Red Hat OpenShift Cluster > Click on “Connect” from the Overview page > Kubeadmin credentials.
 - \${OCP_URL}: Select your {CLUSTER_NAME} Resource Group from Azure Portal > Click on “Overview” page from the left panel > Click on your Azure Red Hat OpenShift Cluster > “API Server URL” from the Overview Page
 - \${VERSION}: “cpd-version” from “/mnt/.cpd/templates/config.json” file path
 - \${PROJECT_CPD_INST_OPERATORS}: value “cpd-operator”
 - \${PROJECT_CPD_INST_OPERANDS}: value “zen” (Default value is “zen”, unless it is being changed during deployment time.)
 - \${STG_CLASS_BLOCK}: For ODF storage, value “ocs-storagecluster-ceph-rbd”
 - \${STG_CLASS_FILE}: For ODF storage, value “ocs-storagecluster-cephfs”

To get information on various other services that are available, see the [Cloud Pak for Data Service Catalog](#).

Note: These steps require knowledge of CPD CLI commands and Linux command line. If you need assistance with deployment, please create a Support Ticket if you require assistance for any question. Support Links can be found in the [Learn more section of the Overview Page](#).

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.
 - Follow the [Operator Subscriptions upgrade instructions](#) for the services you are interested to upgrade.
 - Apply upgrades to the custom resources for the [service](#) that you are interested in.

Limitations

- Review the [known issues and limitations](#) for Cloud Pak for Data.

Additional resources

Azure resources

- [Getting Started Resource Center](#)
- [Azure General Reference](#)

Azure services

- [Azure VM](#)
- [Azure DNS](#)
- [Azure Resource groups](#)

IBM Cloud Pak for Data documentation

- [IBM Documentation](#)
- [Red Hat OpenShift Container Platform](#)

Document revisions

Date	Change
June	Updated ODF v4.17 and cp4d 5.2.x
May 2025	Updated ARO version to 4.16
December 2024	Updated ARO version and IBM Cloud Pak for Data documentation with all latest v5.1.x references.
July 2024	Updated the ARO version, IBM CPD and all reference links to version 5.0.x
March 2024	Updated the Architecture diagram and updated the doc to IBM CPD v4.8.x
December 2023	Updated IBM Cloud Pak for Data version to 4.8.x and updated Links, Updated the process of extracting username and password for CPD web console. Added new Watson services (Watson Discovery, Watson Speech Service, Watosnx Assistant) and Manta service
September 2023	Updated the Portworx version and Bootnode VM type
July 2023	Updated IBM Cloud Pak for Data version 4.7.x Links, Added Portworx storage support, Added support for services: RStudio, Watson Pipelines, BigSQL, and AI Factsheets

February 2023	Updated IBM Cloud Pak for Data & Openshift Versions, and 4.6.x Links
November 2022	Updated IBM Cloud Pak for Data & Openshift Versions, and 4.5.x Links
May 2022	Added support for services: Planning Analytics, Watson Discovery
April 2022	Added support for services: DataStage, Db2, Cognos Analytics, Db2 Warehouse, Decision Optimization, Match360 with Watson
March 2022	Changed Deployment Guide Title Updated CPD version
February 2022	Updated Cloud Pak for Data minimum System and Services requirements
December 2021	Initial version