Windows PowerShell Get-Help on Cmdlet 'New-AzHDInsightCluster'

PS:\>Get-HELP New-AzHDInsightCluster -Full

NAME

New-AzHDInsightCluster

SYNOPSIS

Creates an Azure HDInsight cluster in the specified resource group for the current subscription.

SYNTAX

New-AzHDInsightCluster [-Location] <System.String> [-ResourceGroupName] <System.String> [-ClusterName] <System.String> [-ClusterSizeInNodes] <System.Int32>

[-HttpCredential] <System.Management.Automation.PSCredential> [[-StorageAccountResourceId] <System.String>] [[-StorageAccountKey] <System.String>] [-AadTenantId

<System.Guid>] [-AdditionalStorageAccounts <System.Collections.Generic.Dictionary`2[System.String,System.String]>] [-AmbariDatabase

<Microsoft.Azure.Commands.HDInsight.Models.AzureHDInsightMetastore>] [-ApplicationId <System.Guid>] [-AssignedIdentity <System.String>] [-AutoscaleConfiguration

<Microsoft.Azure.Commands.HDInsight.Models.AzureHDInsightAutoscale>] [-CertificateFileContents <System.Byte[]>] [-CertificatePassword <System.String>] [-ClusterTier

{Standard | Premium}] [-ClusterType <System.String>] [-ComponentVersion <System.Collections.Generic.Dictionary`2[System.String,System.String]>]

[-ComputeIsolationHostSku <System.String>] [-Config <Microsoft.Azure.Commands.HDInsight.Models.AzureHDInsightConfig>] [-Configurations

<System.Collections.Generic.Dictionary`2[System.String,System.Collections.Generic.Dictionary`2[System.String,System.String]]>] [-DefaultProfile

<Microsoft.Azure.Commands.Common.Authentication.Abstractions.Core.IAzureContextContainer>] [-DisksPerWorkerNode <System.Int32>] [-EdgeNodeSize <System.String>]

[-EnableComputeIsolation] [-EnableIDBroker] [-EncryptionAlgorithm {RSA-OAEP | RSA-OAEP-256 | RSA1_5}] [-EncryptionAtHost <System.Nullable`1[System.Boolean]>]

[-EncryptionInTransit <System.Nullable`1[System.Boolean]>] [-EncryptionKeyName <System.String>] [-EncryptionKeyVersion <System.String>] [-EncryptionVaultUri

<System.String>] [-HeadNodeSize <System.String>] [-HiveMetastore <Microsoft.Azure.Commands.HDInsight.Models.AzureHDInsightMetastore>] [-KafkaClientGroupId

<System.String>] [-KafkaClientGroupName <System.String>] [-KafkaManagementNodeSize <System.String>] [-MinSupportedTlsVersion <System.String>] [-ObjectId

<System.Guid>] [-OozieMetastore <Microsoft.Azure.Commands.HDInsight.Models.AzureHDInsightMetastore>] [-OSType {Linux}] [-PrivateLink {Enabled | Disabled}]

[-PrivateLinkConfiguration <Microsoft.Azure.Commands.HDInsight.Models.AzureHDInsightPrivateLinkConfiguration[]>] [-ResourceProviderConnection {Inbound | Outbound}]

[-ScriptActions {HeadNode | WorkerNode | ZookeeperNode | EdgeNode}] [-SecurityProfile <Microsoft.Azure.Commands.HDInsight.Models.AzureHDInsightSecurityProfile>]

[-SshCredential <System.Management.Automation.PSCredential>] [-SshPublicKey <System.String>] [-StorageAccountManagedIdentity <System.String>] [-StorageAccountType

{AzureStorage | AzureDataLakeStore | AzureDataLakeStorageGen2}] [-StorageContainer <System.String>] [-StorageFileSystem <System.String>] [-StorageRootPath

<System.String>] [-SubnetName <System.String>] [-Version <System.String>] [-VirtualNetworkId <System.String>] [-WorkerNodeSize <System.String>] [-Zone

<System.String[]>] [-ZookeeperNodeSize <System.String>] [<CommonParameters>]

New-AzHDInsightCluster [-Location] <System.String> [-ResourceGroupName] <System.String> [-ClusterName] <System.String> [-ClusterSizeInNodes] <System.Int32>

[-HttpCredential] <System.Management.Automation.PSCredential> [[-StorageAccountResourceId] <System.String>] [[-StorageAccountKey] <System.String>] [-AadTenantId

<System.Guid>] [-AdditionalStorageAccounts <System.Collections.Generic.Dictionary`2[System.String,System.String]>] [-AmbariDatabase
<Microsoft.Azure.Commands.HDInsight.Models.AzureHDInsightMetastore>] [-ApplicationId <System.Guid>] [-AssignedIdentity <System.String>] [-AutoscaleConfiguration
<Microsoft.Azure.Commands.HDInsight.Models.AzureHDInsightAutoscale>] [-CertificateFilePath <System.String>] [-CertificatePassword <System.String>] [-ClusterTier
{Standard | Premium}] [-ClusterType <System.String>] [-ComponentVersion <System.Collections.Generic.Dictionary`2[System.String,System.String]>]
[-ComputeIsolationHostSku <System.String>] [-Config <Microsoft.Azure.Commands.HDInsight.Models.AzureHDInsightConfig>] [-Configurations

<System.Collections.Generic.Dictionary`2[System.String,System.Collections.Generic.Dictionary`2[System.String,System.String]]>] [-DefaultProfile
<Microsoft.Azure.Commands.Common.Authentication.Abstractions.Core.IAzureContextContainer>] [-DisksPerWorkerNode <System.Int32>] [-EdgeNodeSize <System.String>]
[-EnableComputeIsolation] [-EnableIDBroker] [-EncryptionAlgorithm {RSA-OAEP | RSA-OAEP-256 | RSA1_5}] [-EncryptionAtHost <System.Nullable`1[System.Boolean]>]
[-EncryptionInTransit <System.Nullable`1[System.Boolean]>] [-EncryptionKeyName <System.String>] [-EncryptionKeyVersion <System.String>] [-EncryptionVaultUri
<System.String>] [-HeadNodeSize <System.String>] [-HiveMetastore <Microsoft.Azure.Commands.HDInsight.Models.AzureHDInsightMetastore>] [-KafkaClientGroupId
<System.String>] [-KafkaClientGroupName <System.String>] [-KafkaManagementNodeSize <System.String>] [-MinSupportedTlsVersion <System.String>] [-ObjectId
<System.Guid>] [-OozieMetastore <Microsoft.Azure.Commands.HDInsight.Models.AzureHDInsightMetastore>] [-OSType {Linux}] [-PrivateLink {Enabled | Disabled}]
[-PrivateLinkConfiguration <Microsoft.Azure.Commands.HDInsight.Models.AzureHDInsightPrivateLinkConfiguration[]>] [-ResourceProviderConnection {Inbound | Outbound}]
[-ScriptActions {HeadNode | WorkerNode | ZookeeperNode | EdgeNode}] [-SecurityProfile <Microsoft.Azure.Commands.HDInsight.Models.AzureHDInsightSecurityProfile>]
[-SshCredential <System.Management.Automation.PSCredential>] [-SshPublicKey <System.String>] [-StorageAccountManagedIdentity <System.String>] [-StorageAccountType
{AzureStorage | AzureDataLakeStore | AzureDataLakeStorageGen2}] [-StorageContainer <System.String>] [-StorageFileSystem <System.String>] [-StorageRootPath

<System.String>] [-SubnetName <System.String>] [-Version <System.String>] [-VirtualNetworkId <System.String>] [-WorkerNodeSize <System.String>] [-Zone
<System.String[]>] [-ZookeeperNodeSize <System.String>] [<CommonParameters>]

DESCRIPTION

The New-AzHDInsightCluster creates an Azure HDInsight cluster by using the specified parameters or by using a configuration object that is created by using the
New-AzHDInsightClusterConfig cmdlet.

The cmdlet may call below Microsoft Graph API according to input parameters:

- GET /servicePrincipals/{id}

PARAMETERS

-AadTenantId <System.Guid>

Specifies the Microsoft Entra tenant ID that will be used when accessing Azure Data Lake Store.

Required? false
Position? named
Default value None
Accept pipeline input? False
Accept wildcard characters? false

-AdditionalStorageAccounts <System.Collections.Generic.Dictionary`2[System.String,System.String]>

Specifies the additional Azure Storage accounts for the cluster. You can alternatively use the Add-AzHDInsightStorage cmdlet.

Required? false
Position? named
Default value None
Accept pipeline input? False

Accept wildcard characters? false

-AmbariDatabase <Microsoft.Azure.Commands.HDInsight.Models.AzureHDInsightMetastore>

Gets or sets the database for ambari.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-ApplicationId <System.Guid>

Gets or sets the Service Principal Application Id for accessing Azure Data Lake.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-AssignedIdentity <System.String>

Gets or sets the assigned identity.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-AutoscaleConfiguration <Microsoft.Azure.Commands.HDInsight.Models.AzureHDInsightAutoscale>

Gets or sets the autoscale configuration

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-CertificateFileContents <System.Byte[]>

Specifies file contents of the certificate that will be used when accessing Azure Data Lake Store.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-CertificateFilePath <System.String>

Specifies the file path to the certificate that will be used to authenticate as the Service Principal. The cluster will use this when accessing Azure Data Lake

Store.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-CertificatePassword <System.String>

Specifies the password for the certificate that will be used to authenticate as the Service Principal. The cluster will use this when accessing Azure Data Lake

Store.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-ClusterName <System.String>

Specifies the name of the cluster.

Required? true

Position? 2

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-ClusterSizeInNodes <System.Int32>

Specifies the number of Worker nodes for the cluster.

Required? true

Position? 3

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-ClusterTier <System.String>

Specifies the HDInsight cluster tier. By default, this is Standard. The Premium tier can only be used with Linux clusters, and it enables the use of some new

features.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-ClusterType <System.String>

Specifies the type of cluster to create. Options are: Hadoop, HBase, Storm, Spark, INTERACTIVEHIVE, Kafka, and RServer

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-ComponentVersion <System.Collections.Generic.Dictionary`2[System.String,System.String]>

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-ComputeIsolationHostSku <System.String>

Gets or sets the dedicated host sku for compute isolation.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-Config <Microsoft.Azure.Commands.HDInsight.Models.AzureHDInsightConfig>

Specifies the cluster object to be used to create the cluster. This object can be created by using the New-AzHDInsightClusterConfig cmdlet.

Required? false

Position? named

Default value None

Accept pipeline input? True (ByValue)

Accept wildcard characters? false

-Configurations <System.Collections.Generic.Dictionary`2[System.String,System.Collections.Generic.Dictionary`2[System.String,System.String]]>

Specifies the configurations of this HDInsight cluster. You can alternatively use the Add-AzHDInsightConfigValues cmdlet.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-DefaultProfile <Microsoft.Azure.Commands.Common.Authentication.Abstractions.Core.IAzureContextContainer>

The credentials, account, tenant, and subscription used for communication with azure

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-DisksPerWorkerNode <System.Int32>

Specifies the number of disks for worker node role in the cluster.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-EdgeNodeSize <System.String>

Specifies the size of the virtual machine for the edge node. Use Get-AzVMSize for acceptable VM sizes, and see HDInsight's pricing page. This parameter is valid

only for RServer clusters.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-EnableComputeIsolation <System.Management.Automation.SwitchParameter>

Enables HDInsight compute isolation feature.

Required? false

Position? named

Default value False

Accept pipeline input? False

Accept wildcard characters? false

-EnableIDBroker <System.Management.Automation.SwitchParameter>

Enables HDInsight Identity Broker feature.

Required? false

Position? named

Default value False

Accept pipeline input? False

Accept wildcard characters? false

-EncryptionAlgorithm <System.String>

Gets or sets the encryption algorithm.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-EncryptionAtHost <System.Nullable`1[System.Boolean]>

Gets or sets the flag which indicates whether enable encryption at host or not.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-EncryptionInTransit <System.Nullable`1[System.Boolean]>

Gets or sets the flag which indicates whether enable encryption in transit or not.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-EncryptionKeyName <System.String>

Gets or sets the encryption key name.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-EncryptionKeyVersion <System.String>

Gets or sets the encryption key version.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-EncryptionVaultUri <System.String>

Gets or sets the encryption vault uri.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-HeadNodeSize <System.String>

Specifies the size of the virtual machine for the Head node. Use Get-AzVMSize for acceptable VM sizes, and see HDInsight's pricing page.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-HiveMetastore <Microsoft.Azure.Commands.HDInsight.Models.AzureHDInsightMetastore>

Specifies the SQL Database to store Hive metadata. You can alternatively use the Add-AzHDInsightMetastore cmdlet.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-HttpCredential <System.Management.Automation.PSCredential>

Specifies the cluster login (HTTP) credentials for the cluster.

Required? true

Position? 4

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-KafkaClientGroupId <System.String>

Gets or sets the client group id for Kafka Rest Proxy access.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-KafkaClientGroupName <System.String>

Gets or sets the client group name for Kafka Rest Proxy access.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-KafkaManagementNodeSize <System.String>

Gets or sets the size of the Kafka Management Node.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-Location <System.String>

Specifies the location for the cluster.

Required? true

Position? 0

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-MinSupportedTlsVersion <System.String>

Gets or sets the minimal supported TLS version.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-ObjectId <System.Guid>

Specifies the Microsoft Entra object ID (a GUID) of the Microsoft Entra service principal that represents the cluster. The cluster will use this when accessing

Azure Data Lake Store.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-OozieMetastore <Microsoft.Azure.Commands.HDInsight.Models.AzureHDInsightMetastore>

Specifies the SQL Database to store Oozie metadata. You can alternatively use the Add-AzHDInsightMetastore cmdlet.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-OSType <System.String>

Specifies the operating system for the cluster. Options are: Windows, Linux

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-PrivateLink <System.String>

Gets or sets the private link type.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-PrivateLinkConfiguration <Microsoft.Azure.Commands.HDInsight.Models.AzureHDInsightPrivateLinkConfiguration[]>

Gets or sets the private link configuration.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-ResourceGroupName <System.String>

Specifies the name of the resource group.

Required? true

Position? 1

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-ResourceProviderConnection <System.String>

Gets or sets the resource provider connection type.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-ScriptActions <System.Collections.Generic.Dictionary`2[Microsoft.Azure.Management.HDInsight.Models.ClusterNodeType,System.Collections.Generic.List`1[Microsoft.Azure.Commands.HDInsight.Models.Management.AzureHDInsightScriptAction]]>

Specifies the script actions to run on the cluster at the end of cluster creation. You can alternatively use Add-AzHDInsightScriptAction.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-SecurityProfile <Microsoft.Azure.Commands.HDInsight.Models.AzureHDInsightSecurityProfile>

Specifies the security related properties used to create a secure cluster. You can alternatively use the Add-AzHDInsightSecurityProfile cmdlet.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-SshCredential <System.Management.Automation.PSCredential>

Specifies the SSH credential to be used for SSH connections. This is only for Linux clusters.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-SshPublicKey <System.String>

Specifies the public key to be used for SSH connections. This is only for Linux clusters.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-StorageAccountKey <System.String>

Gets or sets the Storage Account Access Key for the Storage Account.

Required? false

Position? 6

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-StorageAccountManagedIdentity <System.String>

Gets or sets the storage account managed identity.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-StorageAccountResourceId <System.String>

Gets or sets the Storage Resource Id for the Storage Account.

Required? false

Position? 5

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-StorageAccountType <System.Nullable`1[Microsoft.Azure.Commands.HDInsight.Models.Management.StorageType]>

Gets or sets the type of the storage account.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-StorageContainer <System.String>

Gets or sets the StorageContainer name for the default Azure Storage Account

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-StorageFileSystem <System.String>

Gets or sets the file system for the default Azure Data Lake Storage Gen2 account.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-StorageRootPath <System.String>

Gets or sets the path to the root of the cluster in the default Data Lake Store Account.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-SubnetName <System.String>

Gets or sets the subnet name for this HDInsight cluster.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-Version <System.String>

Specifies the HDI version of the HDInsight cluster.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-VirtualNetworkId <System.String>

Specifies the ID of the virtual network into which to provision the cluster.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-WorkerNodeSize <System.String>

Specifies the size of the virtual machine for the Worker node. Use Get-AzVMSize for acceptable VM sizes, and see HDInsight's pricing page.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-Zone <System.String[]>

Gets or sets the availability zones.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

-ZookeeperNodeSize <System.String>

Specifies the size of the virtual machine for the Zookeeper node. Use Get-AzVMSize for acceptable VM sizes, and see HDInsight's pricing page. This parameter is

valid only for HBase or Storm clusters.

Required? false

Position? named

Default value None

Accept pipeline input? False

Accept wildcard characters? false

<CommonParameters>

This cmdlet supports the common parameters: Verbose, Debug,

ErrorAction, ErrorVariable, WarningAction, WarningVariable,

OutBuffer, PipelineVariable, and OutVariable. For more information, see

about_CommonParameters (https:/go.microsoft.com/fwlink/?LinkID=113216).

INPUTS

Microsoft.Azure.Commands.HDInsight.Models.AzureHDInsightConfig

OUTPUTS

Microsoft.Azure.Commands.HDInsight.Models.AzureHDInsightCluster

NOTES

Keywords: azure, azurerm, arm, resource, management, manager, hadoop, hdinsight, hd, insight

--------- Example 1: Create an Azure HDInsight cluster ---------

```
# Primary storage account info
$storageAccountResourceGroupName = "Group"
$storageAccountResourceId = "yourstorageaccountresourceid"
$storageAccountName = "yourstorageacct001"
$storageAccountKey = Get-AzStorageAccountKey `
    -ResourceGroupName $storageAccountResourceGroupName `
    -Name $storageAccountName | Where-Object {$_.KeyName -eq "key1"} | ForEach-Object{$_.Value}
$storageContainer = "container002"

# Cluster configuration info
$location = "East US 2"
$clusterResourceGroupName = "Group"
$clusterName = "your-hadoop-002"
$clusterCreds = Get-Credential

# If the cluster's resource group doesn't exist yet, run:
# New-AzResourceGroup -Name $clusterResourceGroupName -Location $location

# Create the cluster
New-AzHDInsightCluster `
    -ClusterType Hadoop `
    -ClusterSizeInNodes 4 `
```

```
    -ResourceGroupName $clusterResourceGroupName `
    -ClusterName $clusterName `
    -HttpCredential $clusterCreds `
    -Location $location `
    -StorageAccountResourceId $storageAccountResourceId `
    -StorageAccountKey $storageAccountKey `
    -StorageContainer $storageContainer `
    -SshCredential $clusterCreds
```

This command creates a cluster in the current subscription.

Example 2: Create cluster with customer-managed key disk encryption

```
# Primary storage account info
$storageAccountResourceGroupName = "Group"
$storageAccountResourceId = "yourstorageaccountresourceid"
$storageAccountName = "yourstorageacct001"
$storageAccountKey = Get-AzStorageAccountKey `
    -ResourceGroupName $storageAccountResourceGroupName `
    -Name $storageAccountName | Where-Object {$_.KeyName -eq "key1"} | ForEach-Object{$_.Value}
$storageContainer = "container002"

# Cluster configuration info
$location = "East US 2"
$clusterResourceGroupName = "Group"
$clusterName = "your-cmk-cluster"
$clusterCreds = Get-Credential

# Customer-managed Key info
$assignedIdentity = "your-ami-resource-id"
$encryptionKeyName = "new-key"
$encryptionVaultUri = "https://MyKeyVault.vault.azure.net"
```

```
$encryptionKeyVersion = "00000000000000000000000000000000"

# If the cluster's resource group doesn't exist yet, run:
# New-AzResourceGroup -Name $clusterResourceGroupName -Location $location

# Create the cluster
New-AzHDInsightCluster `
    -ClusterType Spark `
    -ClusterSizeInNodes 4 `
    -ResourceGroupName $clusterResourceGroupName `
    -ClusterName $clusterName `
    -HttpCredential $clusterCreds `
    -Location $location `
    -StorageAccountResourceId $storageAccountResourceId `
    -StorageAccountKey $storageAccountKey `
    -StorageContainer $storageContainer `
    -SshCredential $clusterCreds `
    -AssignedIdentity $assignedIdentity `
    -EncryptionKeyName $encryptionKeyName `
    -EncryptionVaultUri $encryptionVaultUri `
    -EncryptionKeyVersion $encryptionKeyVersion
```

Example 3: Create an Azure HDInsight cluster which enables encryption in transit

```
# Primary storage account info
$storageAccountResourceGroupName = "Group"
$storageAccountResourceId = "yourstorageaccountresourceid"
$storageAccountName = "yourstorageacct001"
$storageAccountKey = Get-AzStorageAccountKey `
    -ResourceGroupName $storageAccountResourceGroupName `
```

```
    -Name $storageAccountName | Where-Object {$_.KeyName -eq "key1"} | ForEach-Object{$_.Value}
$storageContainer = "container002"

# Cluster configuration info
$location = "East US 2"
$clusterResourceGroupName = "Group"
$clusterName = "your-hadoop-002"
$clusterCreds = Get-Credential

# If the cluster's resource group doesn't exist yet, run:
# New-AzResourceGroup -Name $clusterResourceGroupName -Location $location

# Create the cluster
New-AzHDInsightCluster `
    -ClusterType Hadoop `
    -ClusterSizeInNodes 4 `
    -ResourceGroupName $clusterResourceGroupName `
    -ClusterName $clusterName `
    -HttpCredential $clusterCreds `
    -Location $location `
    -StorageAccountResourceId $storageAccountResourceId `
    -StorageAccountKey $storageAccountKey `
    -StorageContainer $storageContainer `
    -SshCredential $clusterCreds `
    -EncryptionInTransit $true
```

Example 4: Create an Azure HDInsight cluster with relay outbound and private link feature

```
# Primary storage account info
$storageAccountResourceGroupName = "Group"
```

```
$storageAccountResourceId = "yourstorageaccountresourceid"
$storageAccountName = "yourstorageacct001"
$storageAccountKey = Get-AzStorageAccountKey `
    -ResourceGroupName $storageAccountResourceGroupName `
    -Name $storageAccountName | Where-Object {$_.KeyName -eq "key1"} | ForEach-Object{$_.Value}
$storageContainer = "container002"

# Cluster configuration info
$location = "East US 2"
$clusterResourceGroupName = "Group"
$clusterName = "your-hadoop-002"
$clusterCreds = Get-Credential

# If the cluster's resource group doesn't exist yet, run:
# New-AzResourceGroup -Name $clusterResourceGroupName -Location $location

# Virtual network info
$virtualNetworkId="yourvnetresourceid"
$subnetName="yoursubnetname"

# Create the cluster
New-AzHDInsightCluster `
    -ClusterType Hadoop `
    -ClusterSizeInNodes 4 `
    -ResourceGroupName $clusterResourceGroupName `
    -ClusterName $clusterName `
    -HttpCredential $clusterCreds `
    -Location $location `
    -StorageAccountResourceId $storageAccountResourceId `
    -StorageAccountKey $storageAccountKey `
    -StorageContainer $storageContainer `
    -SshCredential $clusterCreds `
    -VirtualNetworkId $virtualNetworkId -SubnetName $subnetName `
```

```
    -ResourceProviderConnection Outbound -PrivateLink Enabled
```

Example 5: Create an Azure HDInsight cluster which enables encryption at host

```
# Primary storage account info
$storageAccountResourceGroupName = "Group"
$storageAccountResourceId = "yourstorageaccountresourceid"
$storageAccountName = "yourstorageacct001"
$storageAccountKey = Get-AzStorageAccountKey `
    -ResourceGroupName $storageAccountResourceGroupName `
    -Name $storageAccountName | Where-Object {$_.KeyName -eq "key1"} | ForEach-Object{$_.Value}
$storageContainer = "container002"

# Cluster configuration info
$location = "East US 2"
$clusterResourceGroupName = "Group"
$clusterName = "your-hadoop-002"
$clusterCreds = Get-Credential

# If the cluster's resource group doesn't exist yet, run:
# New-AzResourceGroup -Name $clusterResourceGroupName -Location $location

# Create the cluster
New-AzHDInsightCluster `
    -ClusterType Hadoop `
    -ClusterSizeInNodes 4 `
    -ResourceGroupName $clusterResourceGroupName `
    -ClusterName $clusterName `
    -HttpCredential $clusterCreds `
    -Location $location `
```

```
    -StorageAccountResourceId $storageAccountResourceId `
    -StorageAccountKey $storageAccountKey `
    -StorageContainer $storageContainer `
    -SshCredential $clusterCreds `
    -EncryptionAtHost $true
```

Example 6: Create an Azure HDInsight cluster which enables autoscale.

```
# Primary storage account info
$storageAccountResourceGroupName = "Group"
$storageAccountResourceId = "yourstorageaccountresourceid"
$storageAccountName = "yourstorageacct001"
$storageAccountKey = Get-AzStorageAccountKey `
    -ResourceGroupName $storageAccountResourceGroupName `
    -Name $storageAccountName | Where-Object {$_.KeyName -eq "key1"} | ForEach-Object{$_.Value}
$storageContainer = "container002"

# Cluster configuration info
$location = "East US 2"
$clusterResourceGroupName = "Group"
$clusterName = "your-hadoop-002"
$clusterCreds = Get-Credential

# If the cluster's resource group doesn't exist yet, run:
# New-AzResourceGroup -Name $clusterResourceGroupName -Location $location

# Create autoscale configuration
$autoscaleConfiguration=New-AzHDInsightClusterAutoscaleConfiguration `
    -MinWorkerNodeCount 3 -MaxWorkerNodeCount 5
```

```
# Create the cluster
New-AzHDInsightCluster `
    -ClusterType Hadoop `
    -ClusterSizeInNodes 4 `
    -ResourceGroupName $clusterResourceGroupName `
    -ClusterName $clusterName `
    -HttpCredential $clusterCreds `
    -Location $location `
    -StorageAccountResourceId $storageAccountResourceId `
    -StorageAccountKey $storageAccountKey `
    -StorageContainer $storageContainer `
    -SshCredential $clusterCreds `
    -AutoscaleConfiguration $autoscaleConfiguration
```

Example 7: Create an Azure HDInsight cluster with Kafka Rest Proxy.

```
# Primary storage account info
$storageAccountResourceGroupName = "Group"
$storageAccountResourceId = "yourstorageaccountresourceid"
$storageAccountName = "yourstorageacct001"
$storageAccountKey = Get-AzStorageAccountKey `
    -ResourceGroupName $storageAccountResourceGroupName `
    -Name $storageAccountName | Where-Object {$_.KeyName -eq "key1"} | ForEach-Object{$_.Value}
$storageContainer = "container002"

# Cluster configuration info
$location = "East US 2"
$clusterResourceGroupName = "Group"
$clusterName = "your-hadoop-002"
$clusterCreds = Get-Credential
```

```
# If the cluster's resource group doesn't exist yet, run:
# New-AzResourceGroup -Name $clusterResourceGroupName -Location $location

# Kafka Rest Proxy configuration info
$kafkaClientGroupName = "yourclientgroupname"
$kafkaClientGroupId = "yourclientgroupid"
$kafkaManagementNodeSize = "Standard_D4_v2"
$disksPerWorkerNode = 2

# Create the cluster
New-AzHDInsightCluster `
    -ClusterType Kafka `
    -ClusterSizeInNodes 4 `
    -ResourceGroupName $clusterResourceGroupName `
    -ClusterName $clusterName `
    -HttpCredential $clusterCreds `
    -Location $location `
    -StorageAccountResourceId $storageAccountResourceId `
    -StorageAccountKey $storageAccountKey `
    -StorageContainer $storageContainer `
    -SshCredential $clusterCreds `
    -KafkaClientGroupId  $kafkaClientGroupId -KafkaClientGroupName $kafkaClientGroupName `
    -KafkaManagementNodeSize $kafkaManagementNodeSize -DisksPerWorkerNode $disksPerWorkerNode
```

Example 8: Create an Azure HDInsight cluster with Azure Data Lake Gen2 storage.

```
# Primary storage account info
$storageAccountResourceGroupName = "Group"
$storageAccountResourceId = "yourstorageaccountresourceid"
```

```
$storageManagedIdentity = "yourstorageusermanagedidentity"
$storageFileSystem = "filesystem01"
$storageAccountType = "AzureDataLakeStorageGen2"

# Cluster configuration info
$location = "East US 2"
$clusterResourceGroupName = "Group"
$clusterName = "your-hadoop-002"
$clusterCreds = Get-Credential

# If the cluster's resource group doesn't exist yet, run:
# New-AzResourceGroup -Name $clusterResourceGroupName -Location $location

# Create the cluster
New-AzHDInsightCluster `
    -ClusterType Hadoop `
    -ClusterSizeInNodes 3 `
    -ResourceGroupName $clusterResourceGroupName `
    -ClusterName $clusterName `
    -HttpCredential $clusterCreds `
    -Location $location `
    -StorageAccountResourceId $storageAccountResourceId `
    -StorageAccountManagedIdentity $storageManagedIdentity `
    -StorageFileSystem $storageFileSystem `
    -StorageAccountType $storageAccountType `
    -SshCredential $clusterCreds
```

Example 9: Create an Azure HDInsight cluster with Enterprise Security Package(ESP) and Enable HDInsight ID Broker.

```
# Primary storage account info
```

```
$storageAccountResourceGroupName = "Group"
$storageAccountResourceId = "yourstorageaccountresourceid"
$storageAccountKey = "yourstorageaccountaccesskey"
$storageContainer = "yourcontainer01"

# Cluster configuration info
$location = "East US 2"
$clusterResourceGroupName = "Group"
$clusterName = "your-hadoop-002"
$clusterCreds = Get-Credential

# If the cluster's resource group doesn't exist yet, run:
# New-AzResourceGroup -Name $clusterResourceGroupName -Location $location

# ESP configuration
$domainResourceId = "your Azure AD Domin Service resource id"
$domainUser = "yourdomainuser"
$domainPassword = "yourdoaminpasswd"
$domainPassword = ConvertTo-SecureString $domainPassword -AsPlainText -Force
$domainCredential = New-Object System.Management.Automation.PSCredential($domainUser, $domainPassword)
$clusterUserGroupDns = "dominusergroup"
$ldapUrls = "ldaps://{your domain name}:636"

$clusterTier = "Premium"
$vnetId = "yourvnetid"
$subnetName = "yoursubnetname"
$assignedIdentity = "your user managed assigned identity resourcee id"

#Create security profile
    $config=  New-AzHDInsightClusterConfig|Add-AzHDInsightSecurityProfile  -DomainResourceId  $domainResourceId
-DomainUserCredential $domainCredential -LdapsUrls $ldapUrls
  -ClusterUsersGroupDNs $clusterUserGroupDns
```

```
# Create the cluster
New-AzHDInsightCluster `
    -ClusterTier $clusterTier `
    -ClusterType Hadoop `
    -ClusterSizeInNodes 3 `
    -ResourceGroupName $clusterResourceGroupName `
    -ClusterName $clusterName `
    -HttpCredential $clusterCreds `
    -Location $location `
    -StorageAccountResourceId $storageAccountResourceId `
    -StorageAccountKey $storageAccountKey `
    -StorageContainer $storageContainer `
    -SshCredential $clusterCreds `
    -VirtualNetworkId $vnetId -SubnetName $subnetName `
    -AssignedIdentity $assignedIdentity `
    -SecurityProfile $config.SecurityProfile -EnableIDBroker
```

Example 10: Create an Azure HDInsight cluster which enables compute isolation.

```
# Primary storage account info
$storageAccountResourceGroupName = "Group"
$storageAccountResourceId = "yourstorageaccountresourceid"
$storageAccountName = "yourstorageacct001"
$storageAccountKey = Get-AzStorageAccountKey `
    -ResourceGroupName $storageAccountResourceGroupName `
    -Name $storageAccountName | Where-Object {$_.KeyName -eq "key1"} | ForEach-Object{$_.Value}
$storageContainer = "container002"

# Cluster configuration info
$location = "East US 2"
```

```
$clusterResourceGroupName = "Group"
$clusterName = "your-hadoop-002"
$clusterCreds = Get-Credential
$workerNodeSize="Standard_E16S_V3" # here is just an example
$headNodeSize="Standard_E8S_V3"
$zookeeperNodeSize="Standard_E2S_V3"

# If the cluster's resource group doesn't exist yet, run:
# New-AzResourceGroup -Name $clusterResourceGroupName -Location $location

# Create the cluster
New-AzHDInsightCluster `
    -ClusterType Hadoop `
    -ClusterSizeInNodes 4 `
    -WorkerNodeSize $workerNodeSize `
    -HeadNodeSize $headNodeSize `
    -ZookeeperNodeSize $zookeeperNodeSize `
    -ResourceGroupName $clusterResourceGroupName `
    -ClusterName $clusterName `
    -HttpCredential $clusterCreds `
    -Location $location `
    -StorageAccountResourceId $storageAccountResourceId `
    -StorageAccountKey $storageAccountKey `
    -StorageContainer $storageContainer `
    -SshCredential $clusterCreds `
    -EnableComputeIsolation
```

Example 11: Create an Azure HDInsight cluster with private link configuration feature

```
# Primary storage account info
```

```
$storageAccountResourceGroupName = "Group"
$storageAccountResourceId = "yourstorageaccountresourceid"
$storageAccountName = "yourstorageacct001"
$storageAccountKey = Get-AzStorageAccountKey `
    -ResourceGroupName $storageAccountResourceGroupName `
    -Name $storageAccountName | Where-Object {$_.KeyName -eq "key1"} | ForEach-Object{$_.Value}
$storageContainer = "container002"

# Cluster configuration info
$location = "East US 2"
$clusterResourceGroupName = "Group"
$clusterName = "your-hadoop-002"
$clusterCreds = Get-Credential

# If the cluster's resource group doesn't exist yet, run:
# New-AzResourceGroup -Name $clusterResourceGroupName -Location $location

# Virtual network info
$virtualNetworkId="yourvnetresourceid"
$subnetName="yoursubnetname"

$ipConfigName="ipconfig"
  $privateIPAllocationMethod="dynamic" # the only supported IP allocation method for private link IP configuration is dynamic
$subnetId=$vnetId+"/subnets/"+$subnetName
# Create Private IP configuration
        $ipConfiguration= New-AzHDInsightIPConfiguration -Name $ipConfigName -PrivateIPAllocationMethod $privateIPAllocationMethod -SubnetId $subnetId -Primary

$privateLinkConfigurationName="plconfig"
$groupId="headnode"
# Create private link configuration
  $privateLinkConfiguration= New-AzHDInsightPrivateLinkConfiguration -Name $privateLinkConfigurationName -GroupId
```

```
$groupId -IPConfiguration $ipConfiguration

    # Create the cluster
    New-AzHDInsightCluster `
        -ClusterType Hadoop `
        -ClusterSizeInNodes 4 `
        -ResourceGroupName $clusterResourceGroupName `
        -ClusterName $clusterName `
        -HttpCredential $clusterCreds `
        -Location $location `
        -StorageAccountResourceId $storageAccountResourceId `
        -StorageAccountKey $storageAccountKey `
        -StorageContainer $storageContainer `
        -SshCredential $clusterCreds `
        -VirtualNetworkId $virtualNetworkId -SubnetName $subnetName `
        -ResourceProviderConnection Outbound -PrivateLink Enabled -PrivateLinkConfiguration $privateLinkConfiguration
```

Example 12: Create an Azure HDInsight cluster availability feature

```
    # Primary storage account info
    $storageAccountResourceGroupName = "Group"
    $storageAccountResourceId = "yourstorageaccountresourceid"
    $storageAccountName = "yourstorageacct001"
    $storageAccountKey = Get-AzStorageAccountKey `
        -ResourceGroupName $storageAccountResourceGroupName `
        -Name $storageAccountName | Where-Object {$_.KeyName -eq "key1"} | ForEach-Object{$_.Value}
    $storageContainer = "container002"

    # Cluster configuration info
    $location = "East US 2"
```

```
$clusterResourceGroupName = "Group"
$clusterName = "your-hadoop-002"
$clusterCreds = Get-Credential

# If the cluster's resource group doesn't exist yet, run:
# New-AzResourceGroup -Name $clusterResourceGroupName -Location $location

# Virtual network info
$virtualNetworkId="yourvnetresourceid"
$subnetName="yoursubnetname"

$databaseUserName="yourusername"
$databasePassword="******"
$databasePassword=ConvertTo-SecureString $databasePassword -AsPlainText -Force

$sqlserverCredential=New-Object System.Management.Automation.PSCredential($databaseUserName, $databasePassword)
$sqlserver="yoursqlserver.database.windows.net"
$ambariDatabase="ambaridb"
$hiveDatabase ="hivedb"
$oozieDatabase = "ooziedb"

# availability zone feature requires customer ambari database, hive metatastore and oozie metastore
# Create Ambari database
$config=New-AzHDInsightClusterConfig|Add-AzHDInsightMetastore `
-SqlAzureServerName $sqlserver -DatabaseName $ambariDatabase `
-Credential $sqlserverCredential -MetastoreType AmbariDatabase

# Create Hive metastore
$config=$config|Add-AzHDInsightMetastore `
-SqlAzureServerName $sqlserver -DatabaseName $hiveDatabase `
-Credential $sqlserverCredential -MetastoreType HiveMetastore
```

```
    # Create Oozie metastore
    $config=$config|Add-AzHDInsightMetastore `
    -SqlAzureServerName $sqlserver -DatabaseName $oozieDatabase `
    -Credential $sqlserverCredential -MetastoreType OozieMetastore

    # availability zones
    $zones="1"

    # Create the cluster
    New-AzHDInsightCluster `
        -ClusterType Hadoop `
        -ClusterSizeInNodes 4 `
        -ResourceGroupName $clusterResourceGroupName `
        -ClusterName $clusterName `
        -HttpCredential $clusterCreds `
        -Location $location `
        -StorageAccountResourceId $storageAccountResourceId `
        -StorageAccountKey $storageAccountKey `
        -StorageContainer $storageContainer `
        -SshCredential $clusterCreds `
        -VirtualNetworkId $virtualNetworkId -SubnetName $subnetName `
                    -AmbariDatabase $config.AmbariDatabase -HiveMetastore $config.HiveMetastore -OozieMetastore
$config.OozieMetastore -Zone $zones
```

RELATED LINKS

Online Version: https://learn.microsoft.com/powershell/module/az.hdinsight/new-azhdinsightcluster

New-AzHDInsightClusterConfig