

---

# Burst to Azure Batch with Microsoft HPC Pack 2016 Update 1

Updated: December 13, 2017

Applies To: Microsoft HPC Pack 2016 Update 1

This topic contains information about extending your HPC Pack cluster to include Azure Batch pools as compute resources. Using these Azure Batch pools, you can increase ( “burst” ) the capacity of your HPC cluster on-demand. See the [Azure.com documentation](#) for more about the Azure Batch service.

In this topic:

- [Prerequisites](#)
- [Step 1: Create Azure Batch account](#)
- [Step 2: Create an Azure Batch pool template](#)
- [Step 3: Add an Azure Batch pool](#)
- [Step 4: Start the pool](#)
- [Step 5: Run a job on the pool](#)
- [Step 6: Stop the pool](#)

## Prerequisites

---

- HPC Pack cluster – You must create and configure at least the head node of a cluster. See the [Getting Started Guide for Microsoft HPC Pack 2016 Update 1](#) for steps.
- Azure subscription – If you don’t already have a subscription, sign up for a [free trial](#), use MSDN subscriber benefits, or explore other purchase options.

## Step 1: Create Azure Batch account

---

---

Please follow the steps below to create a Batch account and obtain the required info to burst to Batch. You may also refer to the docs about how to create a Batch account [with the Azure portal](#) or [with the Azure CLI](#).

According to the different Azure Batch resource allocate modes (Batch Service or User Subscription) and client authentication methods (Batch Access Key or Azure AD), you may need the following account information to burst to Batch from HPC Pack.

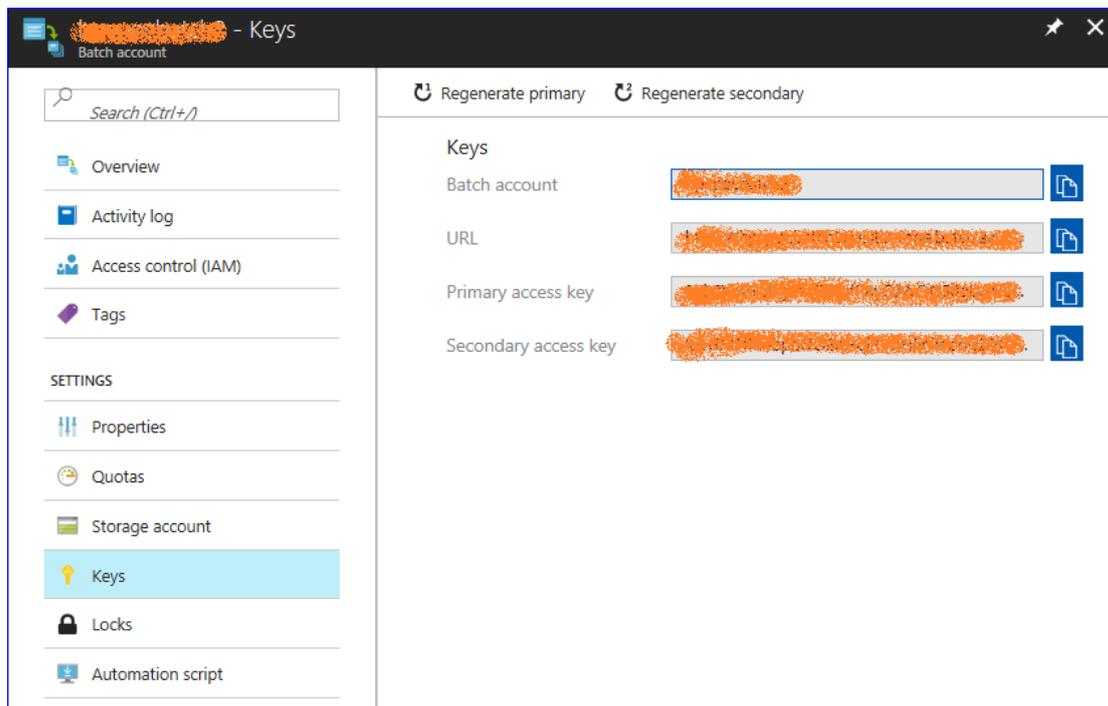
- Batch Service with Access Key
  - Batch account name
  - Batch account URL
  - Batch account key
- Batch Service or User Subscription with Azure AD
  - Batch AAD Instance
  - Batch AAD Tenant Id
  - Batch AAD ClientApp Id
  - Batch AAD ClientApp Key
  - Batch account name
  - Batch account URL

Check the following table to decide which Batch account type and authentication method to choose. You may also check [this blog](#) and [this doc](#) to understand more about User Subscription pool allocation mode and how to use Azure AD authentication for Azure Batch service.

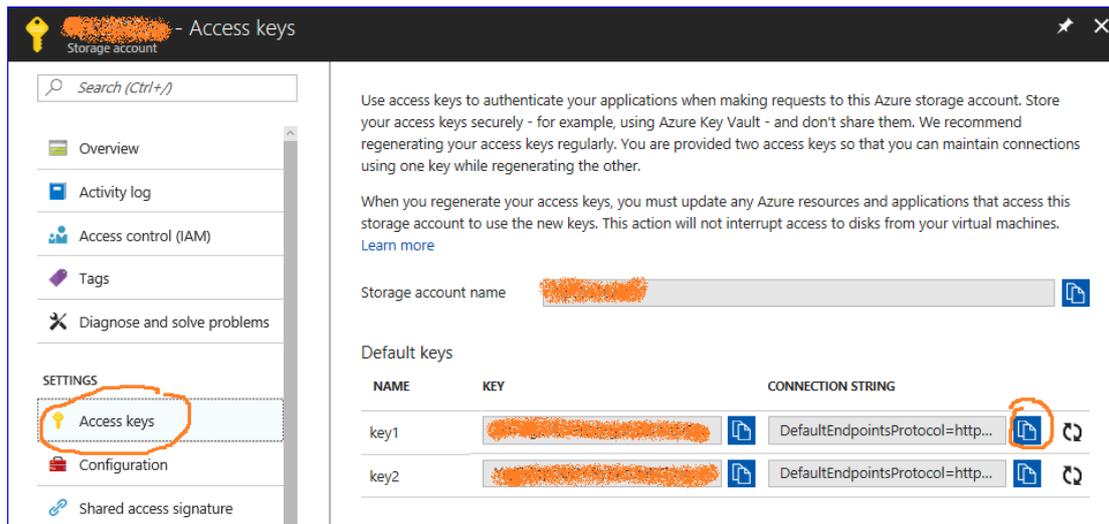
Account Type/Pool Allocation Mode	Authentication Methods	VM image types	Low Priority VM	VNet
Batch Service	Access Key/ Azure AD	PaaS/IaaS (MarketPlaceImage) / IaaS (CustomImage via Azure AD)	Yes	Yes (via Azure AD)
User Subscription	Azure AD	IaaS (MarketPlaceImage)/ IaaS (CustomImage)	No	Yes

To create a Batch account with Batch Service pool allocation mode and use Batch Access Key for authentication

1. Sign in to the [Azure portal](#).
2. Choose **More services**, search and choose **Batch accounts**. You may star it for easy access later.
3. Click on **+Add**
4. In the **New Batch account** form, fill in **Account name**, choose the **Subscription**, create or choose a **Resource group**, choose the **Location**, create or choose a **Storage Account**, and choose **Batch service** as the **Pool allocation mode**.
5. Click on **Create**.
6. After the account is successfully created. Click the account in the list of **Batch accounts**. Choose **Keys**, and obtain the following info for the Batch account:
  - Batch account name
  - Batch account URL
  - Batch account key



7. Choose **More services**, search and choose **Storage accounts**, select the storage account created or chosen for Batch account, click **Access keys**, copy the **CONNECTION STRING** for either **key1** or **key2** and save it for later use.



Alternatively, you may use the Azure CLI to create the account as below,

```
# Authenticate CLI session.
az login

# Select the subscription
az account set -s mysubscription

# Create a resource group.
az group create --name myresourcegroup --location mylocation

# Let's add a storage account reference to the Batch account for use as
'auto-storage'

# for applications. We'll start by creating the storage account.
az storage account create -g myresourcegroup -n mystorageaccount -l
mylocation --sku Standard_LRS

# Create a Batch account.
az batch account create -g myresourcegroup -n mybatchaccount -l
mylocation --storage-account mystorageaccount
```

```
# Now we can display the details of our created account.
az batch account show -g myresourcegroup -n mybatchaccount

# We can view the access keys to the Batch Account for future client
authentication.
az batch account keys list -g myresourcegroup -n mybatchaccount
```

## To create a Batch account with User Subscription pool allocation mode

1. Sign in to the [Azure portal](#).
2. Allow Batch service to access the subscription. (One-time operation)
  - Click **More Services**, search and choose **Subscriptions**, click on the subscription you want to use for the Batch account.
  - In the **Subscription blade**, click **Access control (IAM)** and then click **+Add**.
  - Choose **Contributor Role**, and select the Batch API by one of the following strings:
    1. **MicrosoftAzureBatch**.
    2. **Microsoft Azure Batch**. Newer Azure AD tenants may use this name.
    3. **ddb3205-c6bd-46ae-8127-60eb93363864** is the ID for the Batch API.
  - Select **MicrosoftAzureBatch** and click **Save**.
3. Create a resource group and a key vault for the Batch account.
  - Click **More Services**, search and choose **Resource groups**, click **+Add**, fill/choose the name, subscription and location for the resource group and click **Create**.
  - Click **More Services**, search and choose **Key vaults**, click **+Add**, fill/choose the name, subscription, resource group, location for the key vault.
  - Click on the **Access policies**, and **Add new**, and **Select principal**, select the Batch API **MicrosoftAzureBatch** by one of the strings listed in Step 2, then click **Select**.

- 
- Select all for **Key permissions** and **Secret permissions**. Click **OK**.
- Click on the **Advanced access policy**, choose **Enable access** to Azure Virtual Machine, Azure Resource Manager and Azure Disk Encryption. Click **OK**.
4. Create the Batch account with User Subscription pool allocation mode.
    - Choose **More services**, search and choose **Batch accounts**. You may star it for easy access later.
    - Click on **+Add**
    - In the **New Batch account** form, fill in **Account name**, choose the **Subscription**, choose the **Resource group** created in Step 3, choose the **Location**, create or choose a **Storage Account**, choose **User Subscription** as the **Pool allocation mode**, select the **Key vault** created in Step 3, and click on **Create**.
  5. After the account is successfully created. Click the account in the list of **Batch accounts**. In the **Overview** page obtain the following info for the Batch account:
    - Batch account name
    - Batch account URL

Alternatively, you may use the Azure CLI to create the account as below,

```
# Authenticate CLI session.
az login

# Select the subscription
az account set -s mysubscription

# Allow Azure Batch to access the subscription (one-time operation).
az role assignment create --assignee MicrosoftAzureBatch --role contributor

# Create a resource group.
```

```
az group create --name myresourcegroup --location mylocation

# A Batch account that will allocate pools in the user's subscription
must be configured

# with a Key Vault located in the same region. Let's create this first.
az keyvault create --resource-group myresourcegroup --name mykeyvault -
-location mylocation --enabled-for-deployment true --enabled-for-
disk-encryption true --enabled-for-template-deployment true

# We will add an access-policy to the Key Vault to allow access by the
Batch Service.

az keyvault set-policy --resource-group myresourcegroup --name
mykeyvault --spn ddbf3205-c6bd-46ae-8127-60eb93363864 --key-permissions
all --secret-permissions all

# Now we can create the Batch account, referencing the Key Vault either
by name (if they exist in the same resource group) or by its full
resource ID.

az batch account create --resource-group myresourcegroup --name
mybatchaccount --location mylocation --keyvault mykeyvault
```

To configure Azure AD for Batch authentication and obtain Batch AAD info.

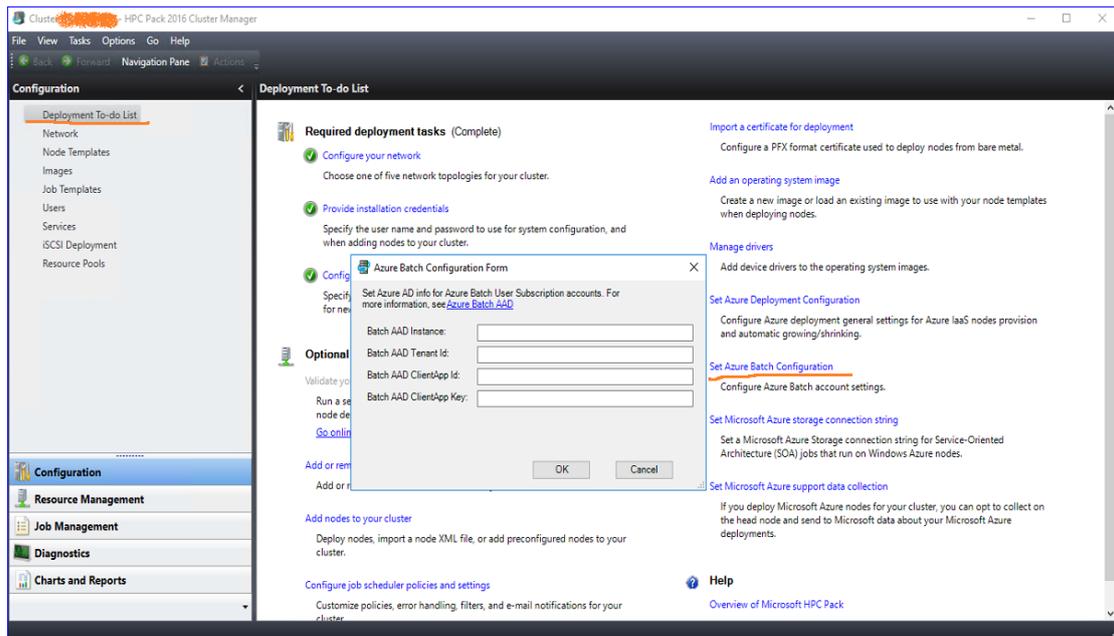
1. Obtain **Batch AAD Instance**. If using global Azure Cloud, the AAD instance is <https://login.microsoftonline.com/>.
2. Obtain **Batch AAD Tenant Id**. In the Azure portal, click **More Services**, search and choose **Azure Active Directory**, select your Active Directory by selecting your account in the top right corner of the page and click **Properties**. Copy the GUID value provided for the **Directory ID**. This value is also called the tenant ID.

- 
3. Register Batch client application and obtain the **Batch AAD ClientApp Id**.
    - In Azure portal, choose your Azure AD tenant by selecting your account in the top right corner of the page.
    - Choose **More services**, search and choose **App Registrations**.
    - Click on **+New application registration**.
    - Fill in the **Name**, choose **Web app / API** as **Application type**, and fill in a value specific to your application e.g. `https://myAppName` for the **Sign-on URI**. Click on **Create**.
    - After the application is successfully created, select the application in the list of **App Registrations**, and click on **Properties**. Copy the GUID value provided for the **Application ID**, it will be used as **Batch AAD ClientApp Id**.
  4. Configure to use a service principal for authentication and obtain the **Batch AAD ClientApp Key**.
    - Request a secret key for the application. Select the application created in the list of **App Registrations**, and click **Keys**, type in **Key description** and choose **Duration**, then click **Save**. Copy the value displayed which will be used as **Batch AAD ClientApp Key**.
    - Assign an RBAC role to the application to authenticate with a service principal. Choose **More services**, search and choose **Batch accounts**, click on the batch account created and select **Access Control (IAM)**. Click on **+ADD**, select **Contributor Role** and the application registered. Click **Save**.

Note that Azure AD authentication is supported by both Batch Service mode and User Subscription mode of Batch accounts.

**Important:** Once configured Azure AD for Batch authentication with following Batch AAD info obtained, open **HPC Cluster Manager**, under **Configuration → Deployment To-do List**, click on **Set Azure Batch Configuration** to fill in the info accordingly in the form, and then click **OK**. This step is **required** when using Azure AD for Batch authentication in HPC Pack.

- Batch AAD Instance
- Batch AAD Tenant Id
- Batch AAD ClientApp Id
- Batch AAD ClientApp Key



Alternatively, you may use HPC Powershell cmdlet to set Batch AAD info as shown below,

```
# HPC PowerShell

Set-HpcClusterRegistry -BatchAADInstance '<AAD Instance e.g.
https://login.microsoftonline.com/>' -BatchAADTenantId <TenantGUID> -
BatchAADClientAppId <AppGUID> -BatchAADClientAppKey '<AppKey>'
```

## Step 2: Create an Azure Batch pool template

To create an Azure Batch pool template, use the Create Node Template Wizard in HPC Cluster Manager.

To create a Batch pool template

1. Start **HPC Cluster Manager**.
2. In **Configuration** navigation pane, click **Node Templates**.
3. Click **New** in the **Actions** pane. The **Create Node Template Wizard** appears.

- 
4. On the **Choose Node Template Type** page, click **Azure Batch pool template**, and then click **Next**.
  5. On the **Specify Template Name** page, type a name for the node template, and optionally type a description for it. Click **Next**.
  6. On the **Provide the Azure Batch account information** page, fill in the Batch account name.

If the Batch account type is Batch Service:

- Choose **BatchService** as Batch account type. If using Batch Access Key for authentication, fill in the key as **Account Key**, if using Azure AD authentication, just leave it blank.

If the Batch account type is User Subscription:

- Choose **UserSubscription** as Batch account type.

Fill in the **Batch account URL** and the **Azure Storage Connection String** we obtained previously, and then click **Next**.

Note, when AAD authentication is required, make sure Batch AAD Instance, Batch AAD Tenant Id, Batch AAD ClientApp Id and Batch AAD ClientApp Key are already set in the **Deployment To-do List**, or the account validation would fail with error “Invalid Azure Batch account. Please check Azure Batch account settings.” when clicking **Next**.

Create Node Template Wizard

### Provide Azure Batch Account Information

Choose Node Template Type

Specify Template Name

**Provide Azure Batch Account Information**

Azure Batch Autoscale configuration

Configure Remote Desktop Credentials and SSH

Specify Startup Script

Review

Azure Batch Account Name: [Redacted]

Azure Batch Account Type: BatchService

When using Azure AD authentication, make sure BatchAADInstance, BatchAADTenantId, BatchAADClientAppId and BatchAADClientAppKey are already configured in To-do list.

Account Key (Optional, blank if using AAD): [Redacted]

Azure Batch Account URL is the Batch service endpoint in the format "https://[account-name].[region].batch.azure.com". You can retrieve it from Azure portal together with account name and key.

Azure Batch Account URL: [Redacted]

Azure Storage Connection String: [Redacted]

[More about Azure Batch information](#)

< Previous   Next >   Create   Cancel

Create Node Template Wizard

### Provide Azure Batch Account Information

Choose Node Template Type

Specify Template Name

**Provide Azure Batch Account Information**

Azure Batch Autoscale configuration

Configure Remote Desktop Credentials and SSH

Specify Startup Script

Review

Azure Batch Account Name: [Redacted]

Azure Batch Account Type: UserSubscription

When using Azure AD authentication, make sure BatchAADInstance, BatchAADTenantId, BatchAADClientAppId and BatchAADClientAppKey are already configured in To-do list.

Azure Batch Account URL is the Batch service endpoint in the format "https://[account-name].[region].batch.azure.com". You can retrieve it from Azure portal together with account name and key.

Azure Batch Account URL: [Redacted]

Azure Storage Connection String: [Redacted]

[More about Azure Batch information](#)

< Previous   Next >   Create   Cancel

7. On the **Azure Batch Autoscale configuration** page, leave **Enable Auto Scale** unchecked, and click **Next**.
8. On the **Configure Remote Desktop Credentials and SSH** page, optionally provide the credentials of a user that will be created on Azure Batch pool compute nodes during deployment. You can use the credentials later to connect to the pool compute nodes. For Linux nodes, you may also specify **SSH Public Key** and **SSH Private Key File(.ppk)** to SSH to the node via [putty.exe](#). Refer [here](#) for how to generate a public key and a private key file for PuTTY. Note if you specify both password and SSH keys, SSH keys will be used for the connection. You also need to copy the generated private key file (.ppk) to %CCP\_HOME%Bin folder to open SSH to the nodes from HPC Cluster Manager. Click on **Next**.

**Create Node Template Wizard**

### Configure Remote Desktop Credentials and SSH

Choose Node Template Type

Specify Template Name

Provide Azure Batch Account Information

Azure Batch Autoscale configuration

**Configure Remote Desktop Credentials and SSH**

Specify Startup Script

Review

You can configure Remote Desktop Credentials or SSH for accessing your Azure Batch compute nodes in Batch pool. These should be different from your domain credentials.

User Name: hpcadmin

Password: .....

Confirm Password: .....

SSH Public Key: ssh-rsa [REDACTED]

SSH Private Key File (.ppk): rd\_rsa.ppk

< Previous   Next >   Create   Cancel

9. On the **Specify Startup Script** page, optionally specify a command line or the name of a startup script to run on all Azure Batch compute nodes in Batch pool. Currently the startup script is only supported for Linux nodes. For example, if you

---

want to run a script named `startup.sh` on all Linux nodes in a Batch pool when they start, you need to use the command line tool `HpcPack.exe` to zip and upload the script to the Azure storage account like below, and then specify `startup.sh` in the **Command Line**.

```
HpcPack.exe create startup.sh.zip startup.sh
HpcPack.exe upload startup.sh.zip /account:<StorageAccountName>
/key:<StorageAccountKey>
```

10. Click **Next** and review all the template settings specified. Click **Create** to generate the Azure Batch pool node template.

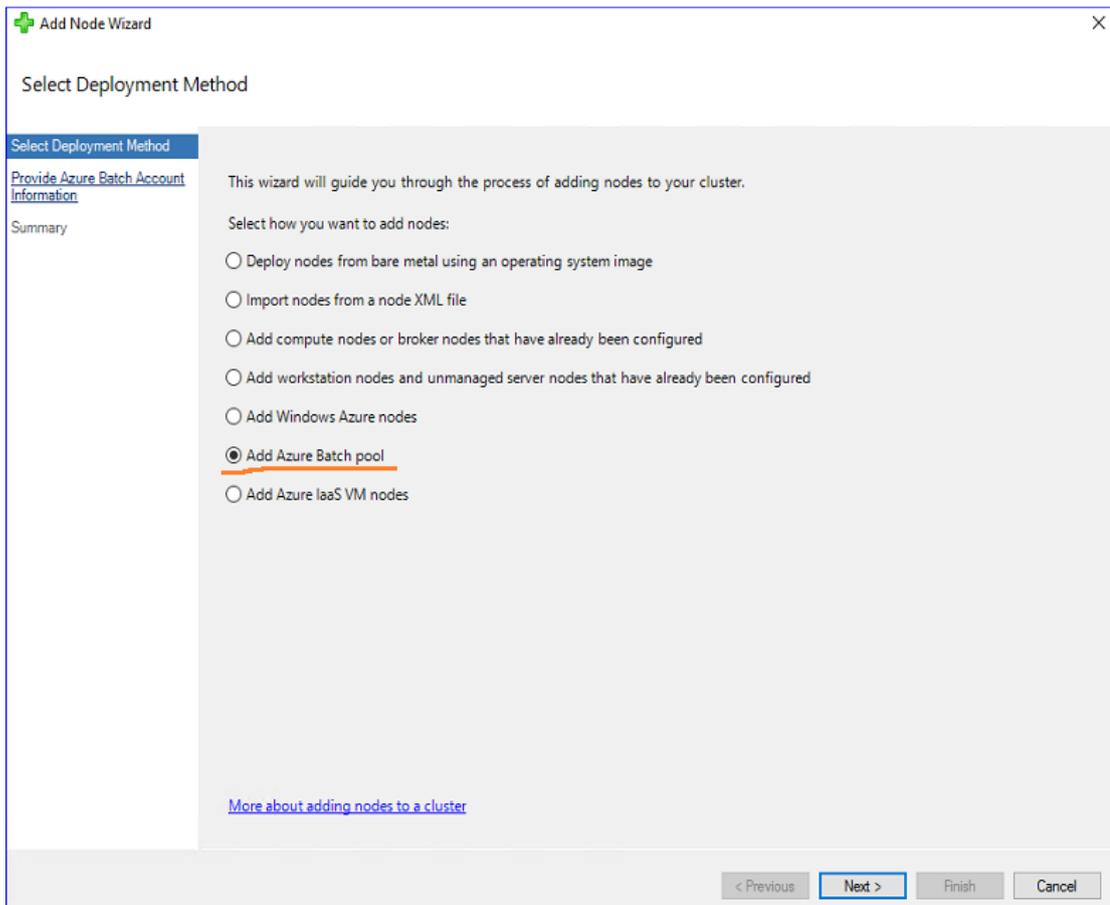
## Step 3: Add an Azure Batch pool

---

Use the Add Node Wizard in HPC Cluster Manager to add the Batch pool compute nodes.

### To add an Azure Batch pool

1. In HPC Cluster Manager, in **Resource Management**, in the **Actions** pane, click **Add Node**. The **Add Node Wizard** appears.
2. On the **Select Deployment Method** page, click **Add Azure Batch pool**, and then click **Next**.



3. On the **Specify Azure Batch Pool Information** page, select an Azure Batch pool template. According to the Batch account type in the template selected, specify the pool information as following:

If Batch account type is Batch Service,

- Number of Compute Nodes – the number of compute nodes (virtual machine instances) in the new Azure Batch pool
- Whether the virtual machines are **Dedicated VM** or **Low Priority VM**
- Choose the image type for either **PaaS** or **IaaSMarketPlace**. For **PaaS**, choose the **OS Family**. E.g. Windows Server 2016. For **IaaSMarketPlace**, choose **Publisher**, **Offer**, and **SKU**. E.g. Canonical, UbuntuServer, 16.04-LTS.
- Size of Compute Nodes – the [role size](#) of each compute node.
- Max Tasks Per Compute Node – the maximum number of concurrent tasks to run on each compute node. The default number is equal to the number of cores in the selected role size. The maximum number is three times larger than the actual number of cores.

Note the Max Tasks Per Compute Node multiplied by the Number of Compute Nodes equals the total cores of the Batch pool node.

- App Packages - optionally specify the application packages that already added in the Batch account in the format of `<Id>:<Version>, <Id>:<Version>, ...`

The screenshot shows the 'Add Node Wizard' dialog box, specifically the 'Specify Azure Batch Pool Information' step. The dialog has a sidebar on the left with 'Specify Azure Batch Pool Information' selected. The main area contains the following fields and options:

- Azure Batch Pool template:** Default AzureBatch Template (dropdown)
- Number of Compute Nodes:** 10 (spin box), with radio buttons for  Dedicated VM and  Low Priority VM
- Image type:** IaaSMarketPlace (dropdown)
- Size of Compute Node:** A3 (4 cores, 7 GB Memory) (dropdown)
- Max Tasks Per Compute Node:** 4 (spin box)
- Publisher:** OpenLogic (dropdown)
- Offer:** CentOS (dropdown)
- Skus:** 7.3 (dropdown)
- App Packages (Id:Version, ... Optional):** myApp:1.0 (text box)

At the bottom, there are navigation buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'. A link for 'More about using Azure Batch pools' is also present.

If Batch account type is User Subscription,

- Number of Compute Nodes - (like above)
- Choose the image type for either **IaaSMarketPlace** or **IaaSCustomImage**. For **IaaSMarketPlace**, choose **Publisher**, **Offer**, and **Skus**. E. g. Canonical, UbuntuServer, 16.04-LTS. For **IaaSCustomImage**, besides **Publisher**, **Offer**, and **Skus**, the **Custom Image Resource Id** is also required in the format of `/subscriptions/{subscriptionId}/resourceGroups/{resourceGroup}/providers/Microsoft.Compute/images/{imageName}`

---

Refer [this doc](#) for how to capture image from a Linux VM and obtain the Image Resource Id.

**Important:** For a custom VM image or a custom VNet (see below), it is required to explicitly assign the Batch client application **Contributor** role to the resource via its **Access control (IAM)**. Or ‘BadRequest’ failure could happen when starting the Batch pool.

- Size of Compute Nodes - (like above)
- Max Tasks Per Compute Node - (like above)
- App Packages - (like above)
- VNet - optionally specify the Subnet Id for the Batch pool. Refer [this doc](#) for how to create a custom VNet and Subnet and obtain the Subnet Id. Also check the Batch requirements for a custom VNet specified in [this doc](#).

The screenshot shows the 'Add Node Wizard' dialog box, specifically the 'Specify Azure Batch Pool Information' step. The dialog has a title bar with a green plus icon and the text 'Add Node Wizard'. Below the title bar, the main heading is 'Specify Azure Batch Pool Information'. On the left side, there is a navigation pane with three items: 'Select Deployment Method', 'Specify Azure Batch Pool Information' (which is highlighted in blue), and 'Summary'. The main area of the dialog contains the following fields and controls:

- Azure Batch Pool template:** A dropdown menu set to 'Default AzureBatch Template'.
- Number of Compute Nodes:** A numeric input field set to '20'.
- Image type:** A dropdown menu set to 'IaaS Custom Image'.
- Size of Compute Node:** A dropdown menu set to 'H8 (8 cores, 448 GB Memory)'.
- Max Tasks Per Compute Node:** A numeric input field set to '8'.
- Publisher:** A dropdown menu set to 'Canonical'.
- Offer:** A dropdown menu set to 'UbuntuServer'.
- Skus:** A dropdown menu set to '16.04-LTS'.
- App Packages (Id:Version,... Optional):** A text input field containing 'myApp:1.0'.
- VNet (Subnet Id, Optional):** A text input field containing '/subscriptions/...'.
- Custom Image Resource Id:** A text input field containing '/subscriptions/...'.

At the bottom of the main area, there is a link: [More about using Azure Batch pools](#). At the bottom right of the dialog, there are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

## Step 4: Start the pool

---

You have to start the pool before running jobs on it.

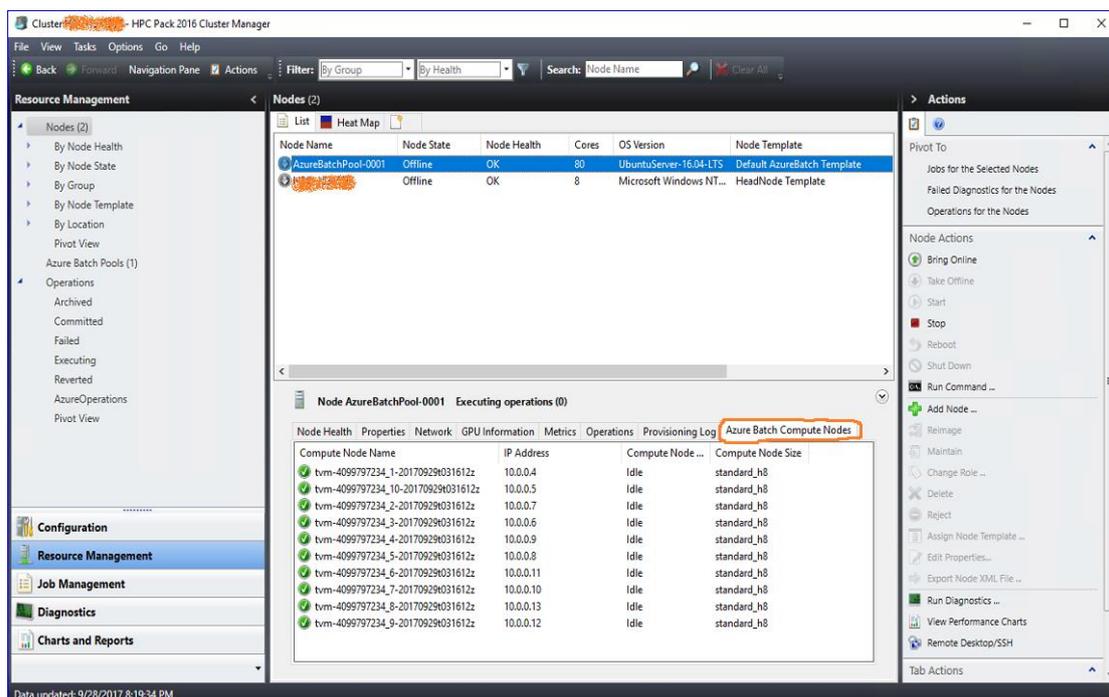
## To start an Azure Batch pool

1. In **Resource Management**, in the **Navigation Pane**, click **Nodes** or **Azure Batch Pools**.
2. In the **List** or **Heat Map** view, select one or more Azure Batch pools.
3. In the **Actions** pane, click **Start**.

The **Start Azure Batch Pools** dialog box appears. Click **Start**.

4. The state of the nodes changes from **Not-Deployed** to **Provisioning**.

If you want to track the provisioning progress, select the pool, and then in the **Details Pane**, click the **Provisioning Log** tab. The Azure Batch pool should be created in less than 1 minute and the state changes to **Offline**.

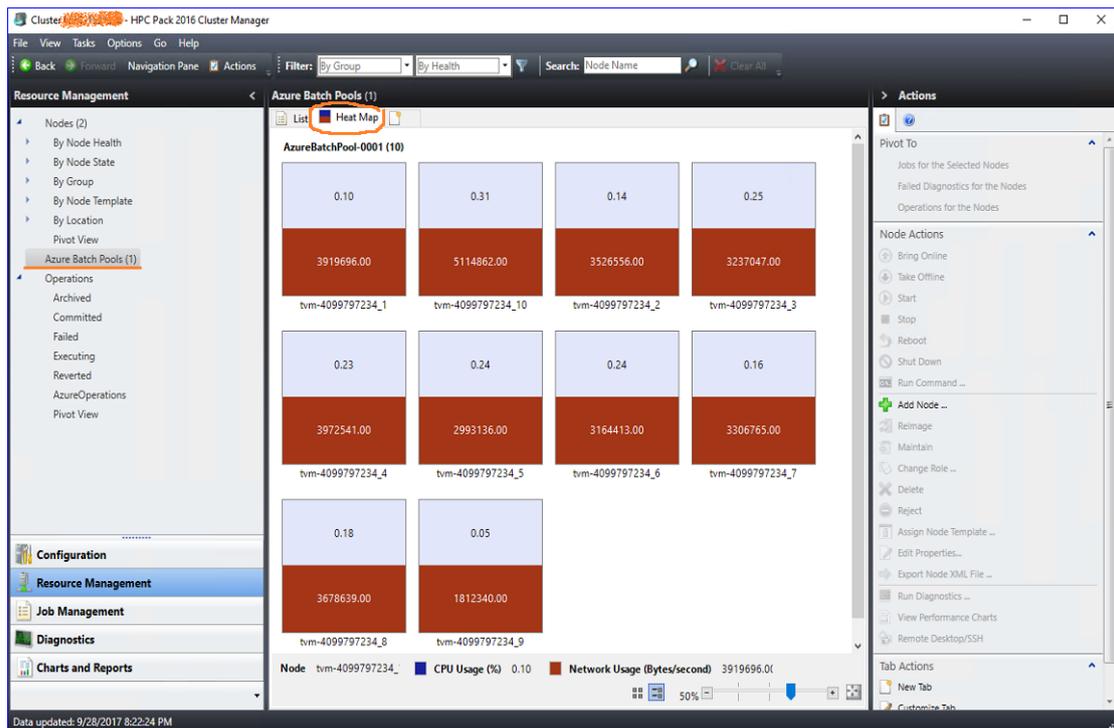


The screenshot displays the HPC Pack 2016 Cluster Manager interface. The main window is titled "Cluster Manager - HPC Pack 2016 Cluster Manager". The "Resource Management" pane on the left shows "Nodes (2)" selected. The main area shows a table of nodes with the following columns: Node Name, Node State, Node Health, Cores, OS Version, and Node Template. Two nodes are listed: "AzureBatchPool-0001" (Offline, OK, 80 Cores, UbuntuServer-16.04-LTS, Default AzureBatch Template) and "HPC-0001" (Offline, OK, 8 Cores, Microsoft Windows NT..., HeadNode Template). The "Actions" pane on the right shows "Start" as an available action. Below the main table, the "Node AzureBatchPool-0001" details are shown, with the "Azure Batch Compute Nodes" tab selected. This tab displays a table of compute nodes with columns: Compute Node Name, IP Address, Compute Node ..., and Compute Node Size. The table lists 12 compute nodes, all with a state of "Idle" and "standard\_h8" size.

Node Name	Node State	Node Health	Cores	OS Version	Node Template
AzureBatchPool-0001	Offline	OK	80	UbuntuServer-16.04-LTS	Default AzureBatch Template
HPC-0001	Offline	OK	8	Microsoft Windows NT...	HeadNode Template

Compute Node Name	IP Address	Compute Node ...	Compute Node Size
tvm-4099797234_1-20170929h031612z	10.0.0.4	Idle	standard_h8
tvm-4099797234_10-20170929h031612z	10.0.0.5	Idle	standard_h8
tvm-4099797234_2-20170929h031612z	10.0.0.7	Idle	standard_h8
tvm-4099797234_3-20170929h031612z	10.0.0.6	Idle	standard_h8
tvm-4099797234_4-20170929h031612z	10.0.0.9	Idle	standard_h8
tvm-4099797234_5-20170929h031612z	10.0.0.8	Idle	standard_h8
tvm-4099797234_6-20170929h031612z	10.0.0.11	Idle	standard_h8
tvm-4099797234_7-20170929h031612z	10.0.0.10	Idle	standard_h8
tvm-4099797234_8-20170929h031612z	10.0.0.13	Idle	standard_h8
tvm-4099797234_9-20170929h031612z	10.0.0.12	Idle	standard_h8



## Additional considerations

- Monitor the status of Azure Batch compute nodes- After the Azure Batch pool is ready, the Azure Batch compute nodes are still being created and starting. To monitor the node status, select the pool and then in the **Details Pane**, click **Azure Batch Compute Nodes**.
- Remote Desktop or SSH to compute nodes - After the compute nodes in Azure Batch pool are started (node state is **Idle**), you can connect by Remote Desktop or SSH to each compute node if you configured template settings to do so, for example, to perform some manual configuration or troubleshooting. To do this, select one or more Azure Batch pools, and then in the **Actions** pane, click **Remote Desktop/SSH**

When connecting to Linux nodes via **SSH**, it is required to copy the generated private key file (.ppk) to the %CCP\_HOME%Bin folder on the client machine.

- View startup tasks - If you specified a startup task in the Azure Batch pool template, after the Azure Batch pool is started, you can view the detailed output of the startup task by running the following HPC PowerShell cmdlet:

```
# HPC PowerShell
```

---

```
Get-HpcBatchPoolStartTask -Name <PoolName>
```

- Heat map view. While the Azure Batch pool is running, you may view the heat map of the pool. In **Resource Management** on navigation pane, click **Nodes**, then choose **Heat Map** view. Currently the following performance counters for a pool are collected:
  - CPU Usage
  - Disk Throughput
  - Free Disk Space
  - Network Usage
  - Available Physical Memory

You may also check the per VM heat map for Linux nodes in the pool. In **Resource Management** on navigation pane, click **Azure Batch Pools**, then choose **Heat Map** view. Currently the following performance counters for a VM are collected:

- CPU Usage
- Disk Throughput
- Free Disk Space
- Network Usage
- Available Physical Memory

## Step 5: Run a job on the pool

---

Currently, HPC Pack supports running [clusrun](#) commands on Azure Batch pools. Please note the following recommendations for using **clusrun** with Batch:

- Run **clusrun** jobs on an Azure Batch pool when no other jobs running on the pool. If there are other jobs or tasks running, the **clusrun** job may need to wait for the running tasks to finish.
- Jobs running on an Azure Batch pool by default don't return the task output to HPC Pack, because of the potential performance impact. You can define node release tasks in the job to retrieve the task output if you want. To retrieve the

---

task output, change the following cluster property through HPC PowerShell:

```
# HPC PowerShell
Set-HpcClusterProperty -GetAzureBatchTaskOutput $true
```

## Step 6: Stop the pool

---

When you are not using the Batch pool, stop the Azure resources. This deprovisions the pool compute nodes, reducing the costs of using a Batch pool.

### To stop the pool

1. In **Resource Management**, in the **Navigation Pane**, click **Nodes or Azure Batch Pools**.
2. In the **List** or **Heat Map** view, select one or more Azure Batch pools that you want to stop.
3. In the **Actions** pane, click **Stop**.

The **Stop Azure Batch Pools** dialog box appears. Click **Stop**.

4. If you want to track the stopping progress, select a node, and then in the **Details Pane**, click the **Provisioning Log** tab.