Azure Stream Analytics

Creating a Linux Virtual Machine

Creating a Linux Virtual Machine (VM) in Azure Stream Analytics can help us test our jobs by generating dummy data using Python script to pass it to an Azure Event Hub.
Follow these tips to guide your way through this process.

1. In the Azure Dashboard, click **New**.

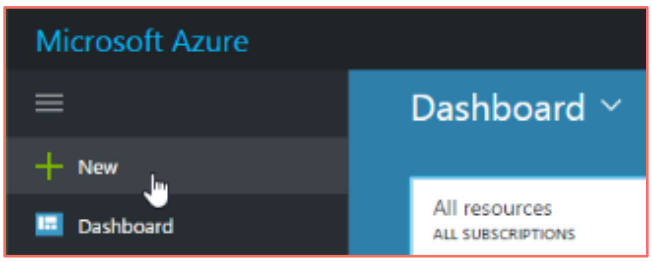

2. Use the **Search** box to look for **Ubuntu** and select your VM option.

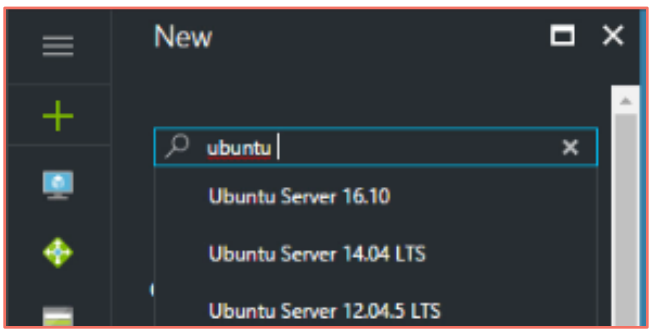

3. After making your selection, click **Create** to start configuring your virtual machine.

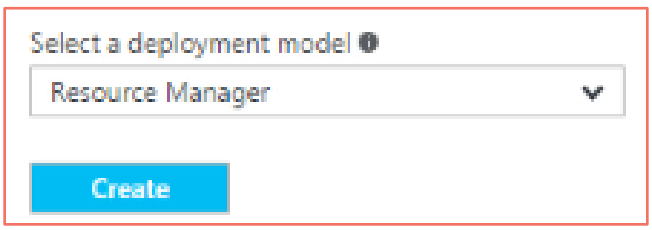

Azure Stream Analytics

Creating a Linux Virtual Machine

4. Complete the information requested in **Step 1**. Then click **OK**.

- Name
- VM Disk Type
- User Name
- Authentication Type
- SSH Public Key
- Subscription
- Resource Group
- Location

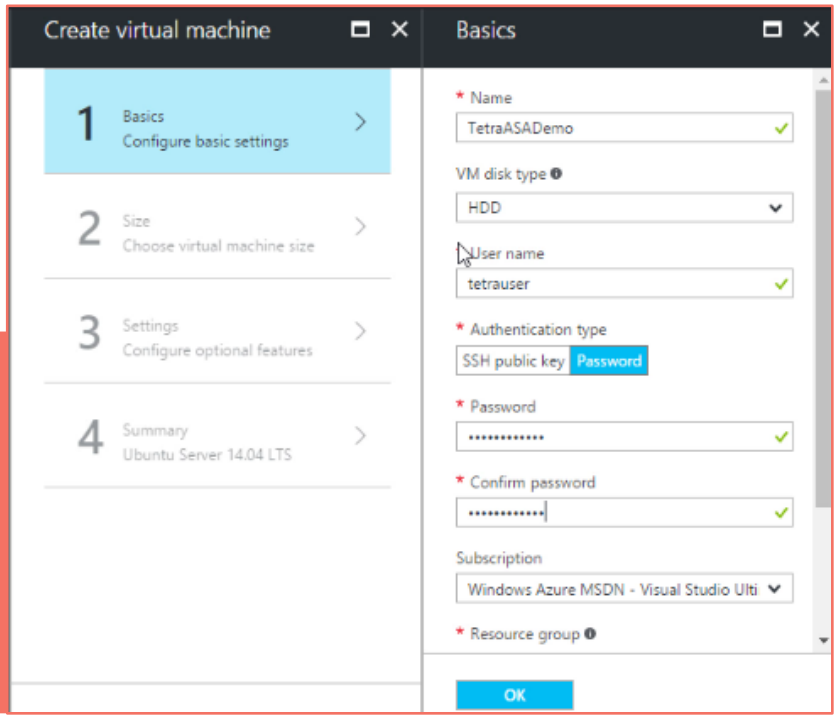

5. Use **Select** your VM's size in **Step 2**.

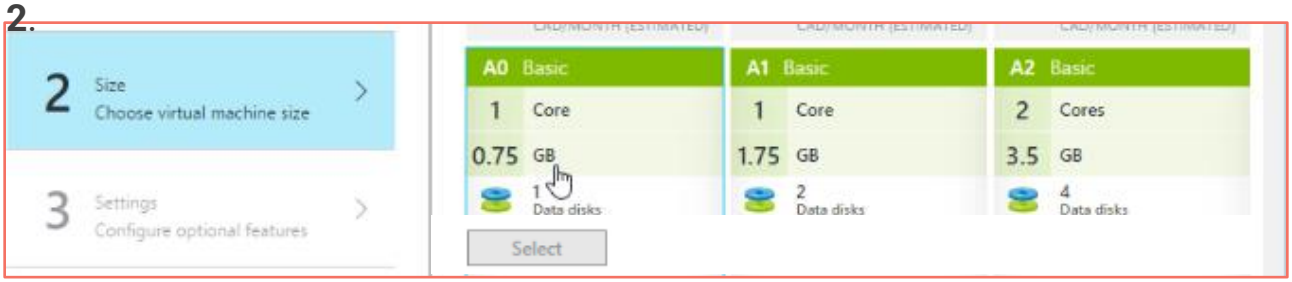

Azure Stream Analytics

Creating a Linux Virtual Machine

6. Complete the information requested in **Step 3**. Then click **OK**.

- Storage
- Network Details
- Extension Details
- High Availability Details
- Monitoring Details

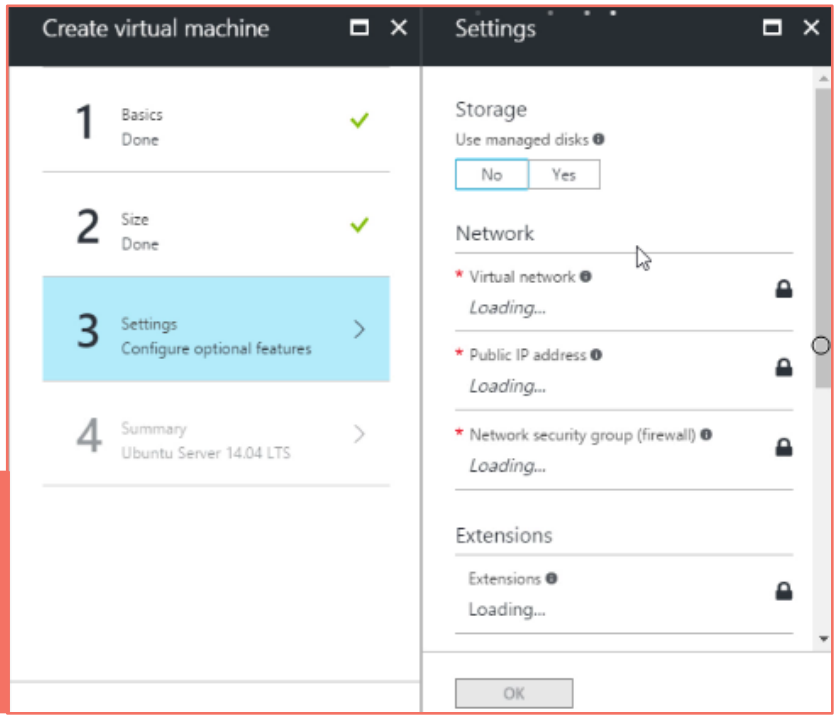

7. Review the summary in **Step 4**. If everything is fine, click **OK**.

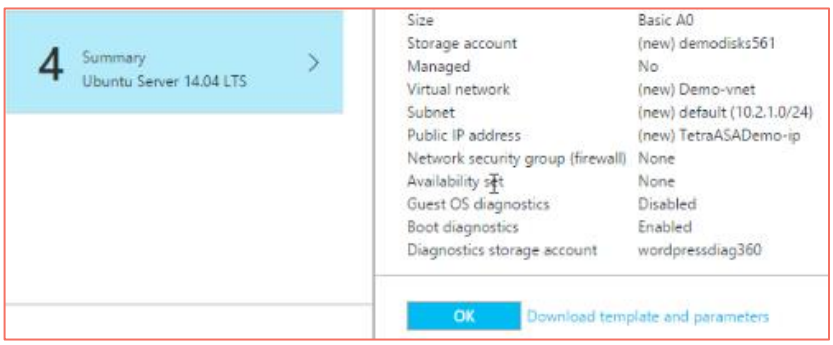

Azure Stream Analytics

Configuring the Linux Virtual Machine

After creating the Linux VM, you need to configure it to send data to your event hub.
Follow these tips to guide your way through this process.

1. Log in to SSH using your credentials and open the **SSH Console**.

2. Run the configuration commands to install the Python libraries I need to push your data to the Event hub.

The commands you need include:

- **Sudo apt-get update** > Starts the configuration
- **Sudo apt-get upgrade** > Upgrades all VM installations
- **Sudo apt-get install** > Installs the required Python version
- **Sudo apt-get install python-pip** > Installs Python's package manager
- **Sudo pip install azure** > Installs the Azure resource library for Python.

NOTE: You may add any other libraries you need here.

- **Wget + your Pyhton script's URL** > Downloads Python's scripts to the VM

Azure Stream Analytics

Pushing Data From the VM to the Event Hub

After creating the Linux VM, you need to configure it to send data to your event hub.
Follow these tips to guide your way through this process.

1. Log in to SSH using your credentials and open the **SSH Console**.

2. Run the configuration commands to install the Python libraries I need to push your data to the Event hub.

The commands you need include:

- **Sudo apt-get update** > Starts the configuration
- **Sudo apt-get upgrade** > Upgrades all VM installations
- **Sudo apt-get install** > Installs the required Python version
- **Sudo apt-get install python-pip** > Installs Python's package manager
- **Sudo pip install azure** > Installs the Azure resource library for Python.

NOTE: You may add any other libraries you need here.

- **Wget + your Pyhton script's URL*** > Downloads Python's scripts to the VM

*The script can be stored in an Azure Storage Account or in a Gdrove, for example.

What You Need to Know: Injecting Data

Azure Stream Analytics

Pushing Data From the VM to the Event Hub

3. Run the script to push your dummy data into your Event Hub. It should include the following structure:

Python **<the script's file name> <service bus name> <event hub policy name> <event hub policy key> <event hub name>**

Please always **keep the Spaces** separating script elements.
Remember to keep your **Event Hub's Shared Access Policies** data handy for this step.

4. Click **Enter** to run your script and start sending dummy data. If you need to stop the script at any time, you can press **ctrl+c**.

Once you have your script running, you will be able to test your event hub and ASA job's performance.

Keep this process handy to stage your ASA jobs and as a best practice before running live data through them.