Pyramid Kubernetes on Azure Guide

Version 1.0

Table of Contents

Overview **3**
Instantiating Kubernetes on Azure **3**
Create Kubernetes cluster **4**
Basics **4**
Node pools, Access, Networking, Integrations, Advanced and Tags 5
Connecting to the Cluster **5**
Generating the Pyramid YAML **6**
Autoscaling the pods: 7
Deploying Pyramid YAML configuration **7**
External IP Access for the Pyramid Kubernetes Instance 8
System Initialization **8**
Finished **9**
Appendix **10**
Deploying an MS-SQL database **10**
Basics **10**
Networking **11**
Creating a new Pyramid Repository Database **12**
Basics 12

Overview

The following guide is provided to customers to set up and launch a Pyramid Kubernetes cluster on Microsoft Azure Cloud platform.
The guide provides a standard walkthrough but is NOT exhaustive and does not cover every available option.

Instantiating Kubernetes on Azure

If you have no prior deployment of Azure Kubernetes engine start here.
Otherwise please start on step 4.
You can either use your existing cluster or choose to deploy one just for Pyramid.

Log into the Azure Admin and search for and then click on “Kubernetes services.”
From the drop down choose “Create a Kubernetes cluster”.

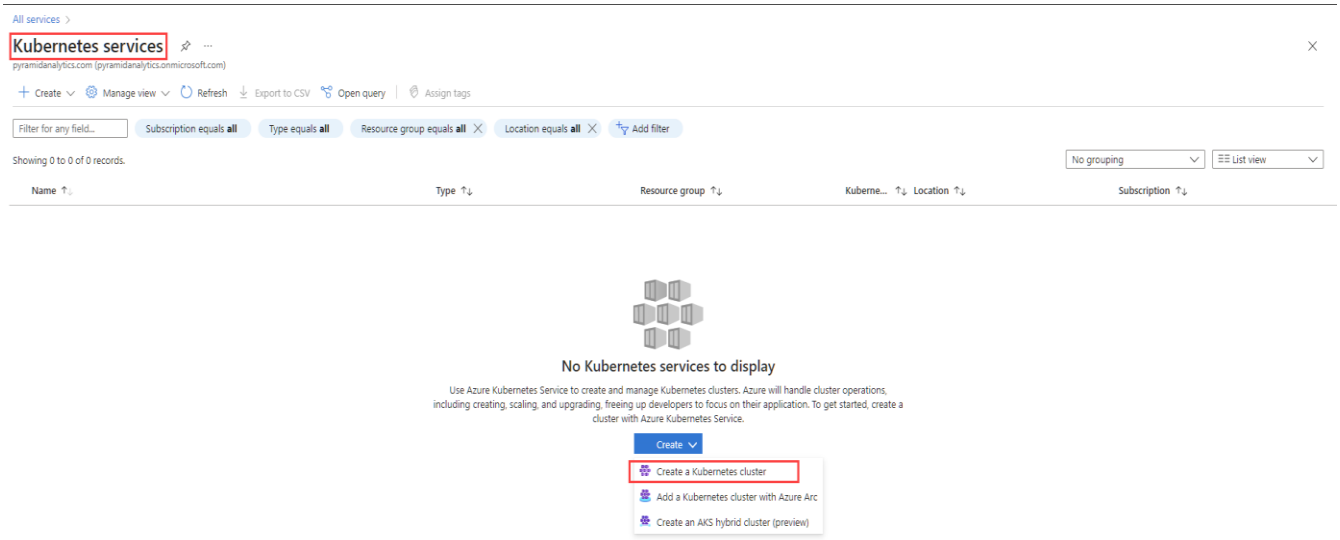

Create Kubernetes cluster

Basics

Settings that are required to be set are:

a. **Subscription and Resource group -** choose your Azure subscription and resource group.
b. **Kubernetes cluster name -** choose a name for your cluster.
c. **Region -** select the region that is best for you.
d. **Kubernetes Version –** This should be left on the default option.
e. **Node size –** depends on expected usage. A suggested start it a machine with 16 cores and 32GB of memory.
f. **Scale method -** Autoscale
g. **Node count range -** depends on expected usage, for production deployments Azure recommended minimum is three nodes.
h. **Review and create -** click on review and create if you do not wish to change any of the other options.

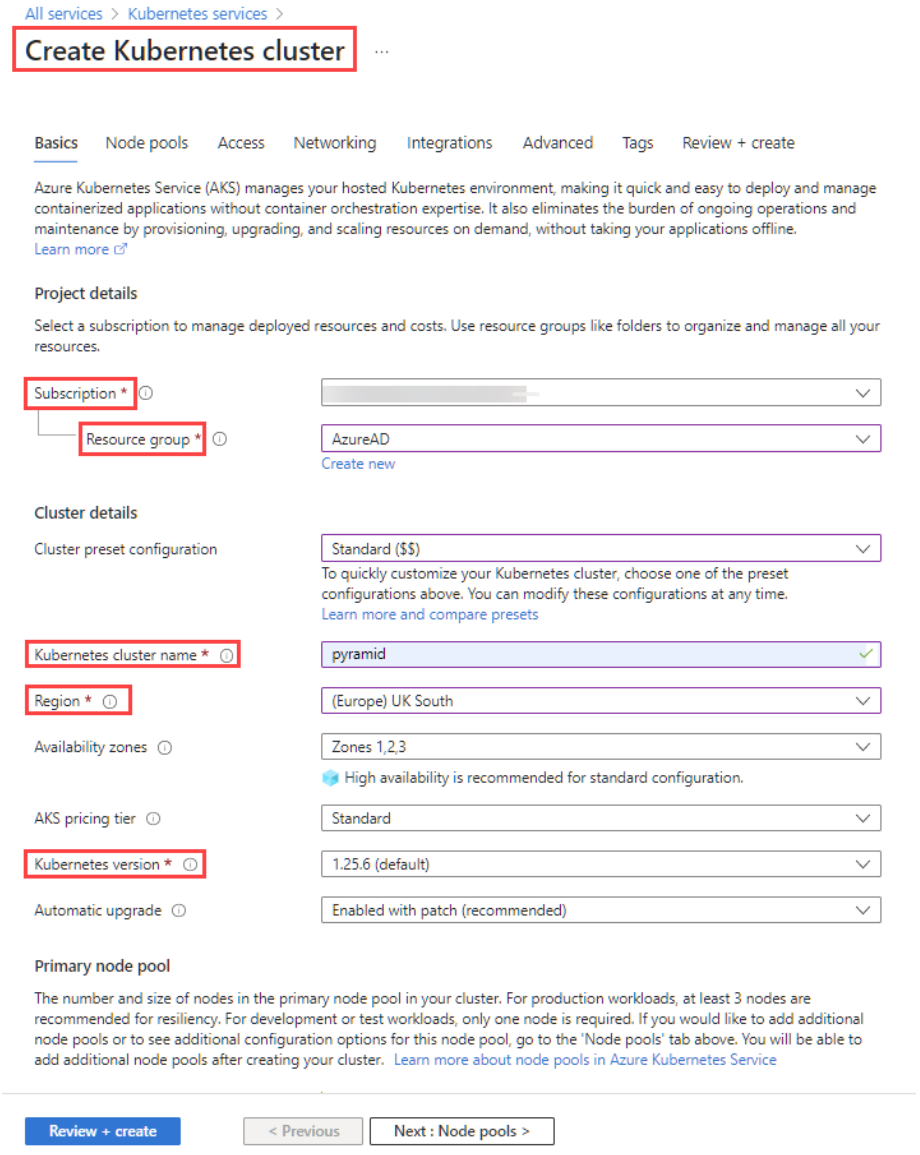

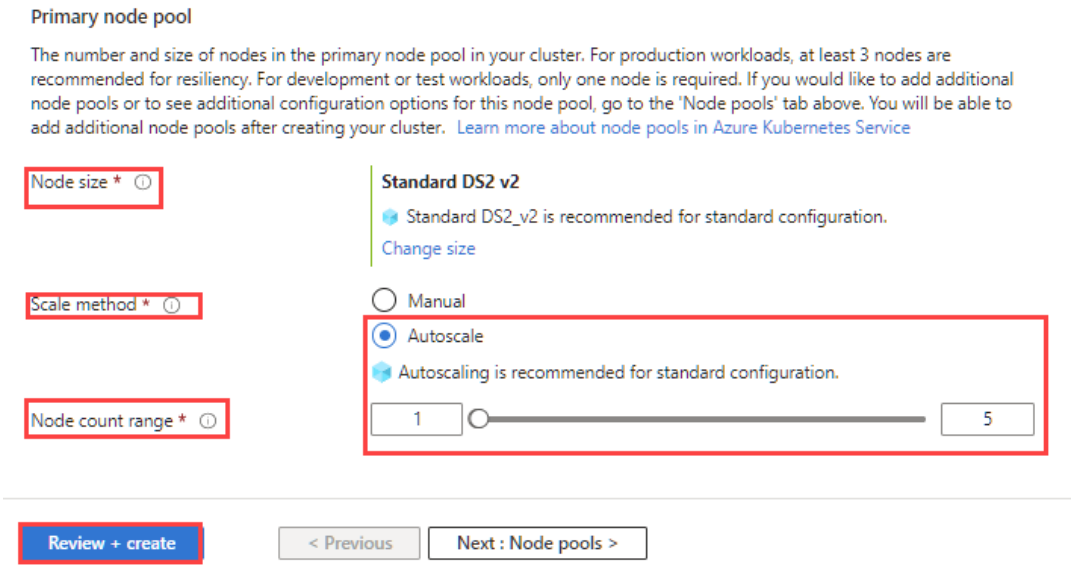

Node pools, Access, Networking, Integrations, Advanced and Tags

All these settings should be customized to confirm with your security and architecture needs. They can all be left as the default option as well.

Connecting to the Cluster

a) Once the cluster has finished being created, click on “connect to cluster” as shown below.

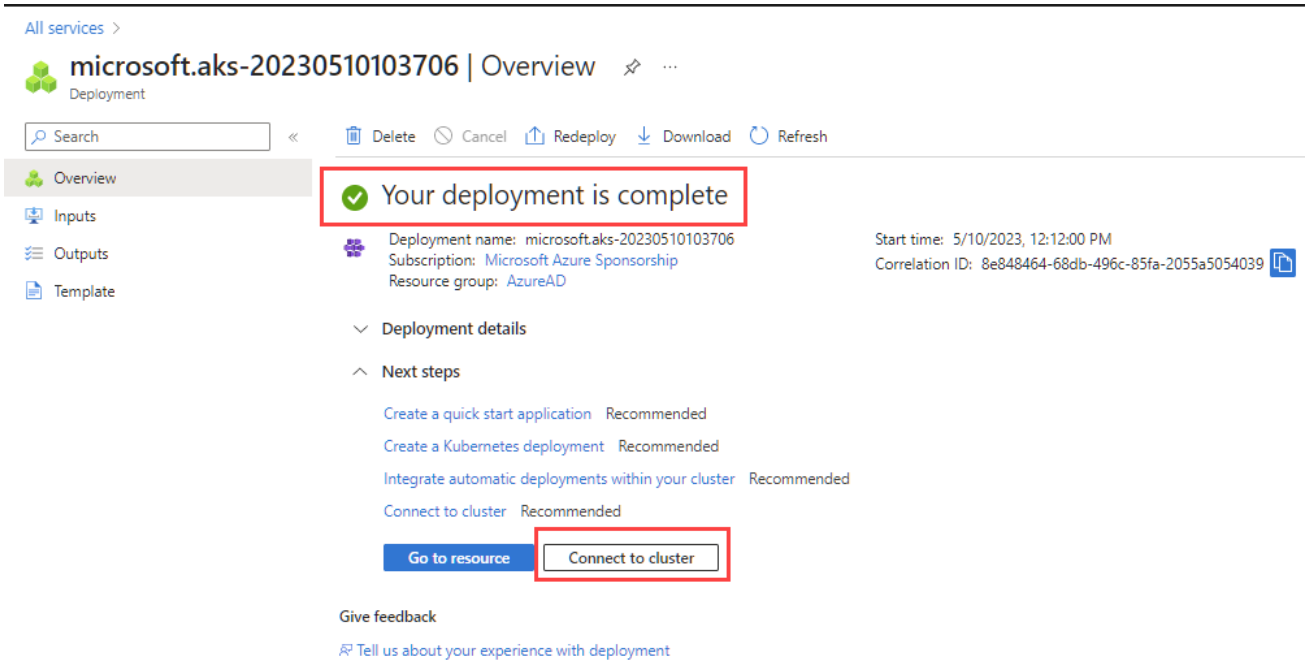

b) Click on Open in Cloud Shell and **make sure to run the two commands** in step number 2 as they are shown for you.

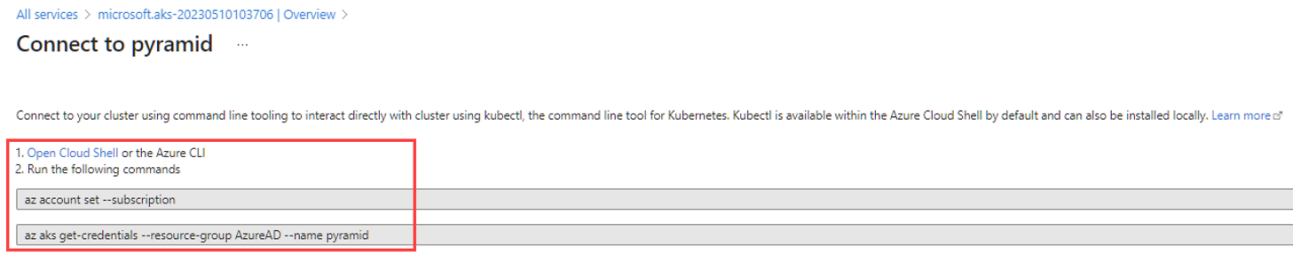

c) Click on either Bash or PowerShell. For this example, we use PowerShell.

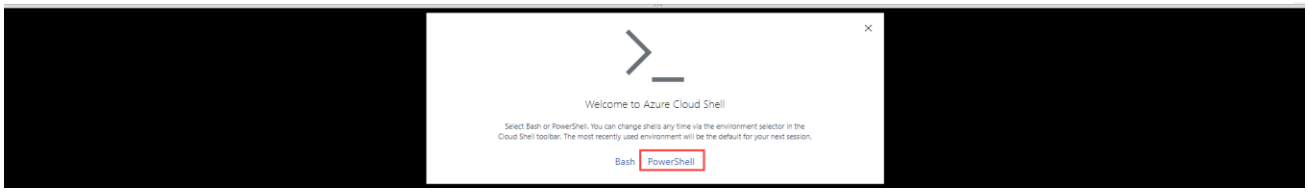

d) If you have no storage mounted, it will prompt you to create one as below. Choose your subscription and then on “create storage”.

e) After running the set account and getting the credentials command, you can test the connectivity to the cluster, by running the below command, if all is working it will bring back the list of nodes as seen in the screenshot.

```
Kubectl get nodes
```

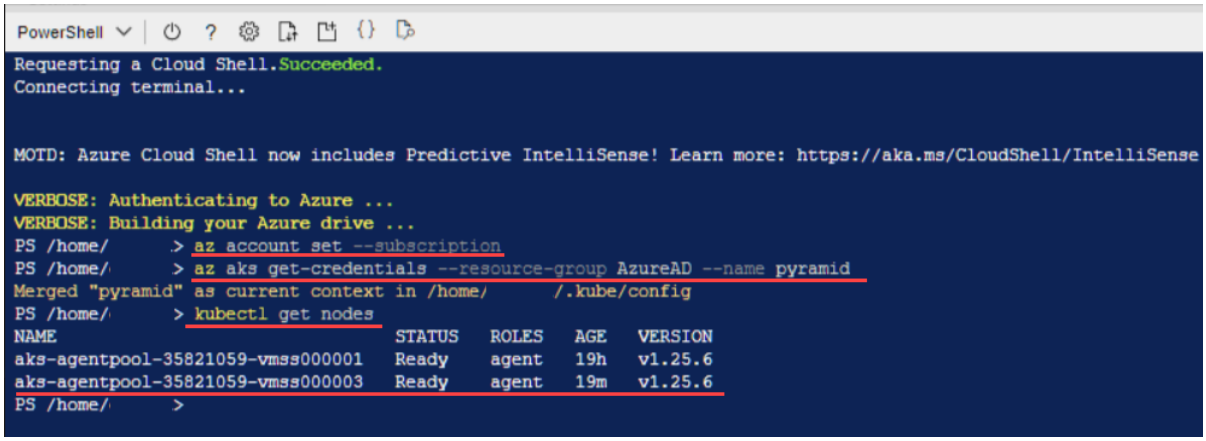

Generating the Pyramid YAML

The setup for Pyramid is *best* driven through a YAML configuration file. If you want, you can also deploy using HELM. For more information on this see this link.
This can be manually created. However, it is simpler to use Pyramid’s YAML configurator.

Login to Pyramid’s customer portal, go to the Kubernetes setup page: https://customers.pyramidanalytics.com/kubernetes/ and generate a YAML file for your Pyramid config.

Choose one of the storage options in the drop-down list.
More info on the configurator can be found here .

Autoscaling the pods:

Pyramid gives you the option of scaling the pods Horizontally.
You can choose the maximum number of replicas(pods) to spawn by ticking the Elastic Scaling option when creating the Pyramid YAML and entering in the max number of pods that can be spawned.

To enable the auto scaling to work, make sure to Install Keda and the Metric Server, these commands are given to you while creating the Pyramid YAML.

Deploying Pyramid YAML configuration

Upload your YAML file (from previous steps) to your cluster as shown below:

Click on the upload file option and upload the Pyramid YAML file.

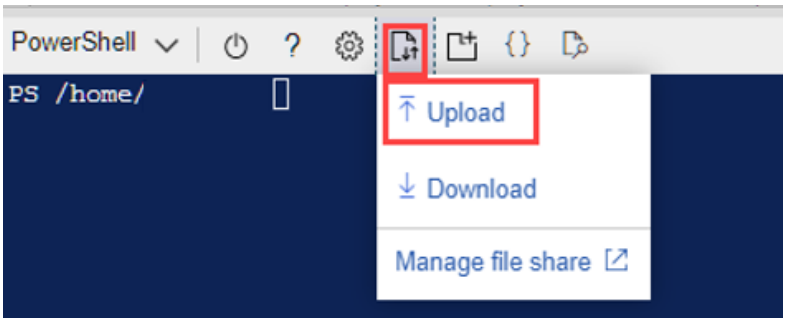

Once you upload the YAML run it as below to pull down the pyramid pods

```
kubectl apply -f pyramid-analytics-config.yaml
```

Then run the below command to see the pods generating or look at the Azure control panel under "Workloads" (it will also show the pods as incomplete until after the full deployment has finished)

```
kubectl -n pyramid get pods -w
```

or

```
kubectl -n pyramid get pods
```

Its normal that only the web-service pod will show 1/1 until the full deployment has finished (until after you have finished the setup in the browser)

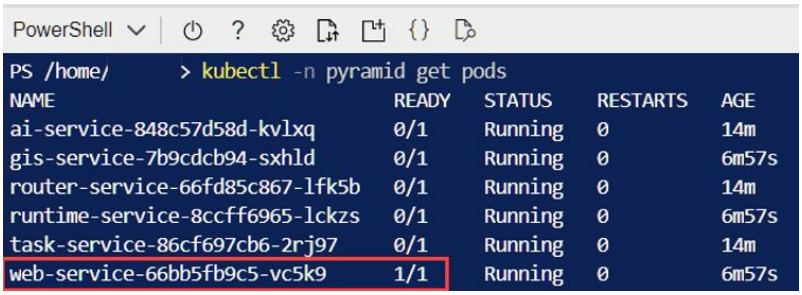

Wait until you see that all pods show as "running." It can take around 10/15 minutes for this to finish.

From the Azure console, it will look as below:

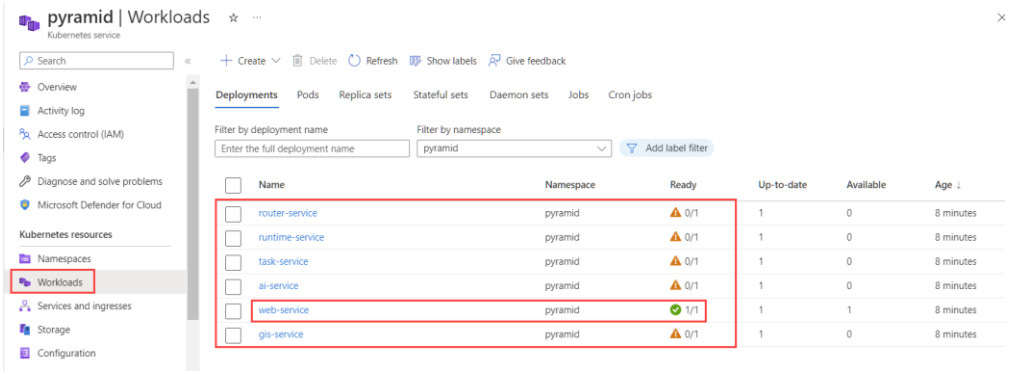

Once you see that the web-service shows as Ready (Green tick 1/1), continue to the next step.

External IP Access for the Pyramid Kubernetes Instance

To get the external IP to access the Pyramid application on, from the Azure Kubernetes portal click on “Services and ingress”. Look for the external IP address and port which is 8181.

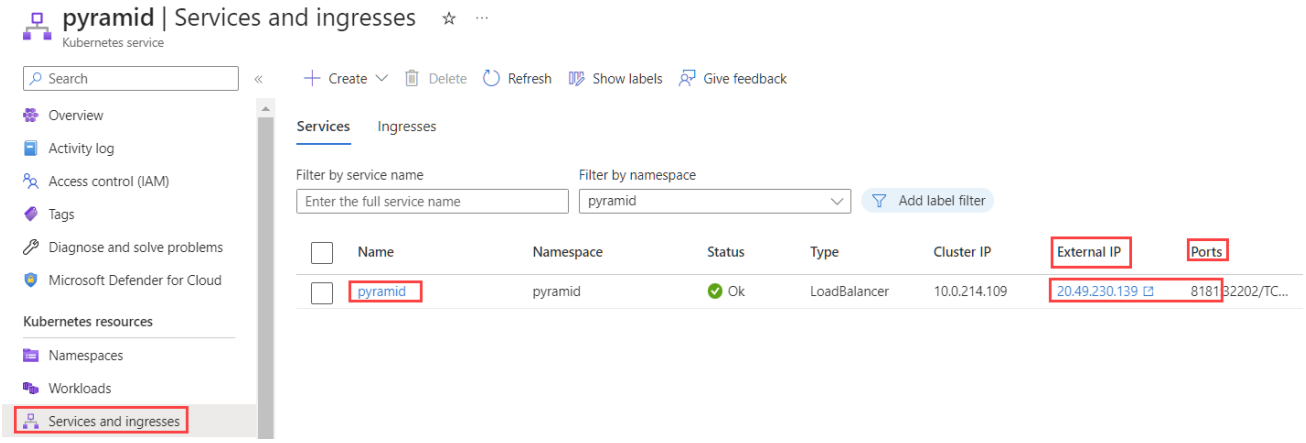

Copy the External IP and paste it into a browser adding :8181 on to the end (e.g., http://204.230.139:8181). This will then take you to the below page, where you can fill out all the needed info to finish the Pyramid deployment.

System Initialization

Once the pods have finished being created, and you click on the link as explained above, you will be prompted with the screen below.
This initializes the system, with persistent storage (this is only if you choose this option when creating the YAML. Otherwise choose one of the other options in the list), the Pyramid repository database and creates the first initial user within Pyramid.
For more information on this stage please see this link.

- See the appendix for details on how to setup a database repository on Azure.
- For more information on this stage please see this link.

For the storage type choose the same persistent storage option that was chosen when you generated the Pyramid YAML.

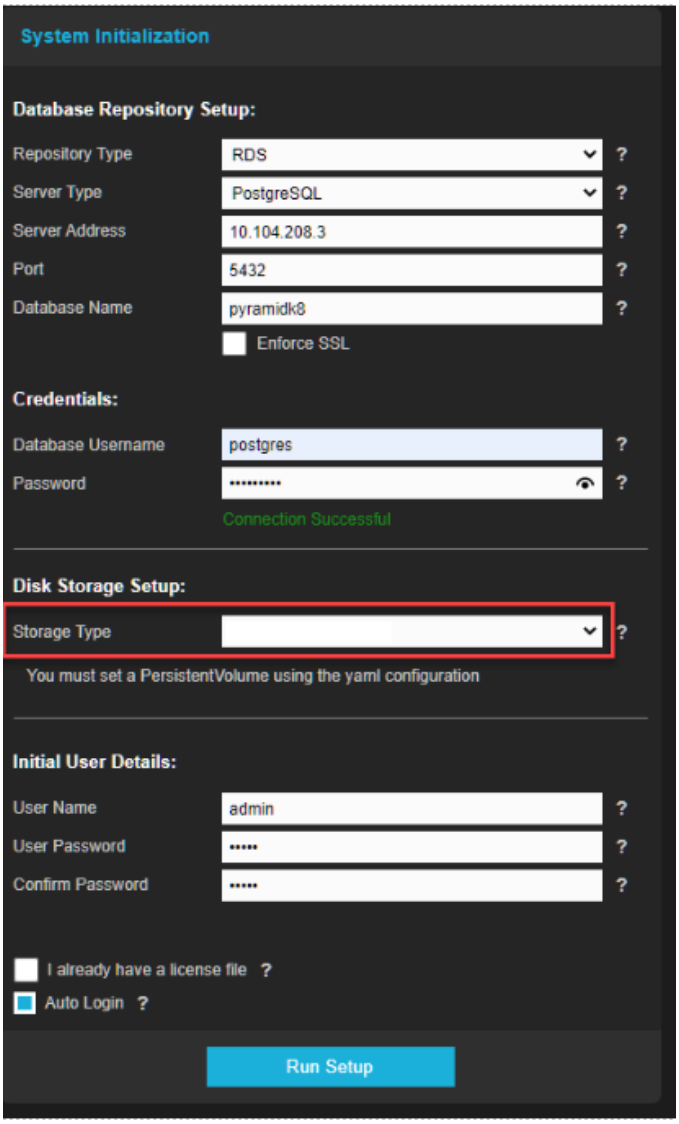

Finished

Once the initialization setup has finished running (normally around 5-10 mins) it will redirect you to the fully installed Pyramid application.

Appendix

Deploying an MS-SQL database

The steps below guide you in the Azure Console for creating an MS-SQL or PostgreSQL instance to host the Pyramid repository.

Notes: it should be a private instance as it does not need to be accessed from outside of your network.
It should be in the same Zone and network as your Kubernetes cluster.

1) Search in the Azure console for “SQL server” and click on “Create”

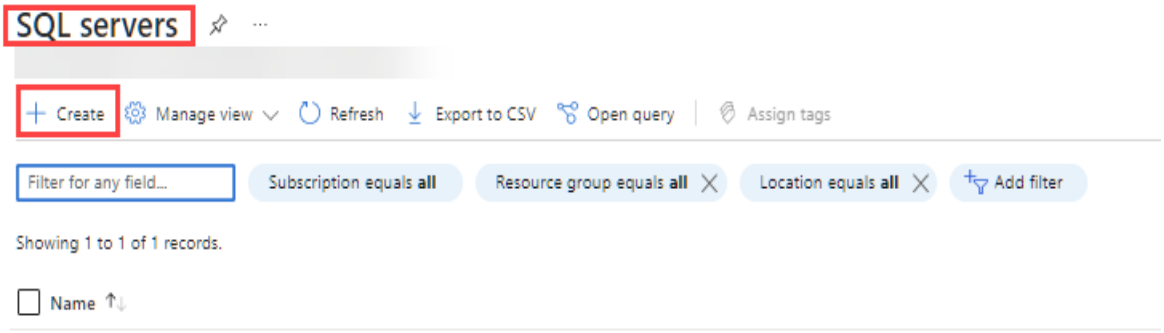

Basics

Settings that are required are:

a. **Subscription and Resource group -** choose your Azure subscription and resource group.
b. **Server name -** give the server a name.
c. **Location –** should be in the same location as your Kubernetes cluster.
d. **Authentication method** must include SQL authentication because this is the only way Pyramid can connect to it as a repository.
e. **Server admin login / password/ confirm password –** give the server an admin login and password.

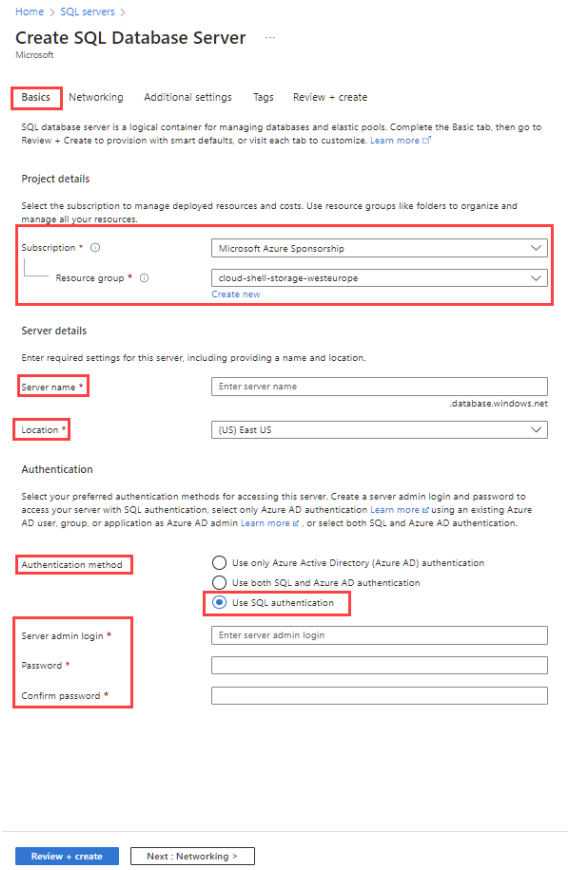

Networking

Settings that are required are:

a. **Allow Azure services and resources to access this server –** this should be turned to “yes” so that the Kubernetes pods can access the SQL instances.

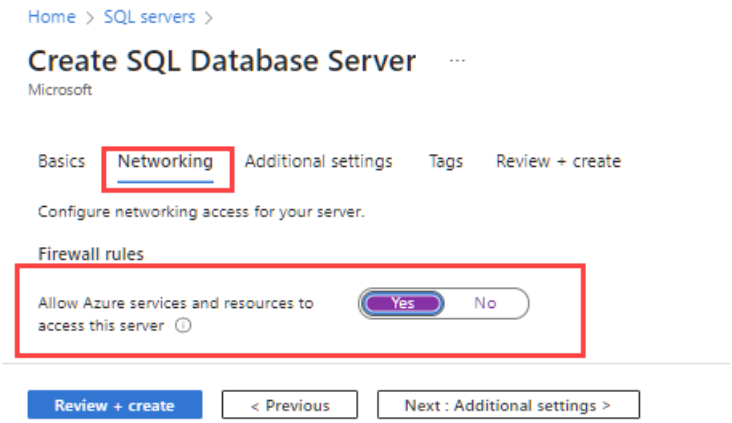

Lastly click on Review + create to create your SQL server as shown below:

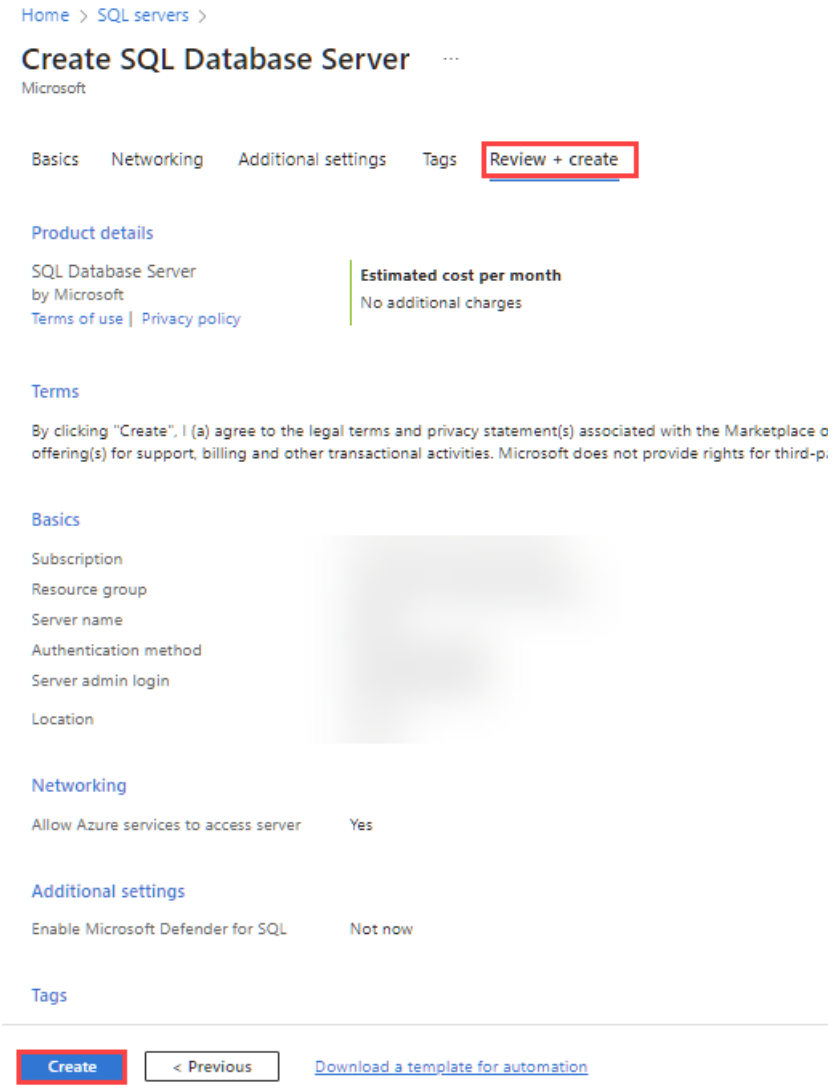

Creating a new Pyramid Repository Database

Once the setup has completed, click on your new SQL instance and create a new blank database. (search for SQL servers and click on the newly created instance).
Click on “Create” as shown in the image below:

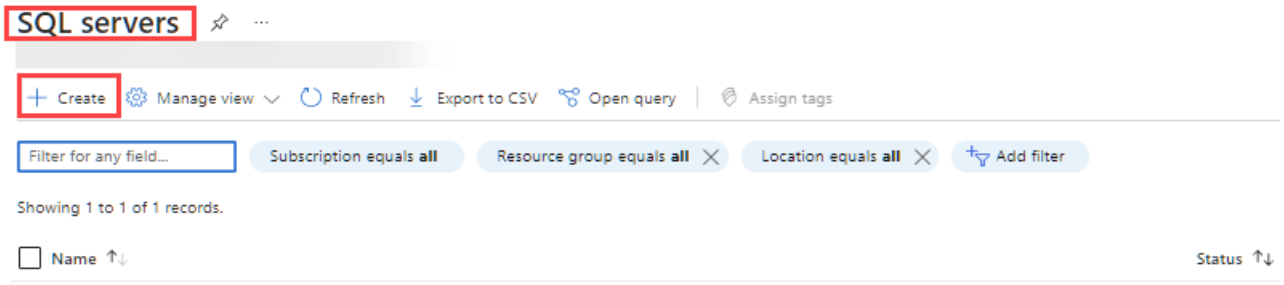

Basics

Settings that are required are:

a. **Database name –** give a name for your pyramid repository database.
b. **Server -** select the SQL server that you created or already have.
c. **Compute + storage -** Ensure that the database is not underpowered. It should not be less than 4 CPU’s (8 is the recommended minimum) and 12-16Gb of Memory.

d. **Review + created –** no other settings need to be changed, click on this to create the database.