



## **Securing your container workloads in Kubernetes**

# Table of Contents

---

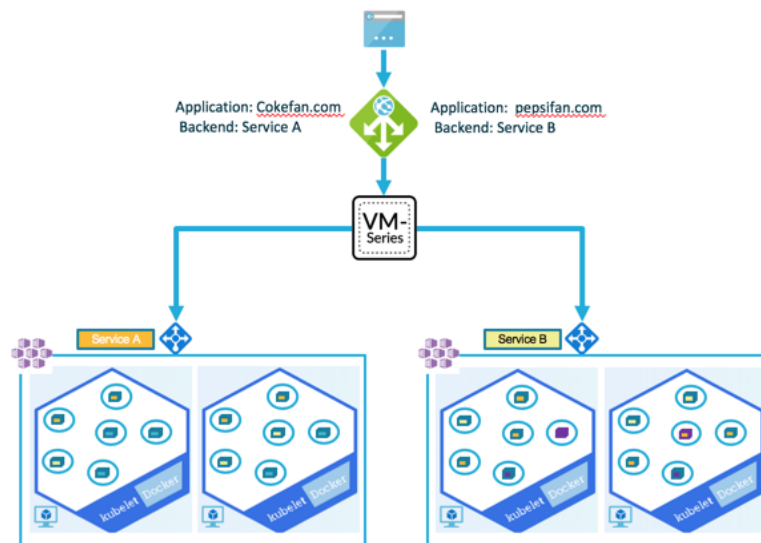
<b>About the Azure Kubernetes Service Terraform Template .....</b>	<b>3</b>
<b>Support Policy.....</b>	<b>4</b>
<b>Instances Used .....</b>	<b>4</b>
<b>Prerequisites.....</b>	<b>4</b>
<b>Download GitHub files .....</b>	<b>5</b>
<b>Azure Service Principal Creation.....</b>	<b>6</b>
<b>Bootstrap storage account creation.....</b>	<b>8</b>
<b>SSH keys.....</b>	<b>16</b>
<b>Deploy the Terraform Template .....</b>	<b>18</b>
<b>Review what was deployed .....</b>	<b>21</b>
Task 1 – Look around Azure console.....	21
Task 2 – Review the Kubernetes Cluster.....	26
Task 3 – Connect to the Kubernetes Cluster.....	27
Task 3 – Log into the firewall .....	28
<b>Launch a two tiered WordPress application .....</b>	<b>32</b>
Task 1 – WordPress Application Deployment YAML file.....	32
Task 2 – Launch the Application .....	35
<b>Launch a two tiered Guestbook application .....</b>	<b>37</b>
Task 1 – Guestbook Application Deployment YAML file.....	37
Task 2 – Launch the Application .....	40
<b>Explore the newly deployed applications .....</b>	<b>41</b>
<b>Securing Inbound Traffic .....</b>	<b>44</b>
Task 1 – Azure Application Gateway IP Address .....	44
Task 2 – Update the Firewall’s Address Objects .....	45
Task 3 – Connect to the Guestbook Frontend .....	47
<b>Securing Outbound Traffic.....</b>	<b>51</b>
Task 1 – Add Outbound Route .....	51
<b>Lab Termination.....</b>	<b>54</b>
<b>Conclusion .....</b>	<b>56</b>

# About the Azure Kubernetes Service Terraform Template

Azure Kubernetes Service (AKS) Terraform Templates are files that can deploy, configure, and launch AZURE resources such as Resource Groups, VNETS, subnets, security groups, application gateways, route tables, Kubernetes clusters, and more. These templates are used for ease of deployment and are key to any cloud deployment model.

For more information on Templates refer to Google's documentation <https://docs.microsoft.com/en-us/azure/terraform/>

This document will walk through the setup and deployment of a Terraform template that deploys the AKS infrastructure and a Palo Alto Networks VM-Series firewall that provides advanced protection for the Kubernetes cluster North/South traffic. During the deployment the template will create two Azure resource groups. One that has the infrastructure including the bootstrapped VM-Series Firewall and another with the k8s cluster resources. The guide also walks through the deployment of two separate applications. Each 2-tier application consists of database and web pods. After completing this guide, the following infrastructure will be instantiated:



# Support Policy

This template is released under an as-is, best effort, support policy. These scripts should be seen as community supported and Palo Alto Networks will contribute our expertise as and when possible. We do not provide technical support or help in using or troubleshooting the components of the project through our normal support options such as Palo Alto Networks support teams, or ASC (Authorized Support Centers) partners and backline support options. The underlying product used (the VM-Series firewall) by the scripts or templates are still supported, but the support is only for the product functionality and not for help in deploying or using the template or script itself.

# Instances Used

When deploying this Terraform template the following machine types are used:

Instance	Machine Type	QTY
PayGo Bundle 1 – VM-Series Firewall	Standard_D3_v2	1
Kubernetes Ubuntu Cluster Nodes	Standard_D3_v2	2
Internal Load Balancer		1
Application Gateway		1

Note: There are Azure costs associated with each machine type launched, please refer to the Microsoft instance pricing page <https://azure.microsoft.com/en-us/pricing/details/virtual-machines/windows/>

# Prerequisites

Here are the prerequisites required to successfully launch this template:

- Terraform application - Instructions on the installation can be found here: <https://www.terraform.io/intro/getting-started/install.html>
- Azure account- Account creation instructions can be found here: <https://azure.microsoft.com/en-us/resources/videos/sign-up-for-microsoft-azure/>
- Azure command-line tool – Instructions for doing this can be found here: <https://docs.microsoft.com/en-us/cli/azure/install-azure-cli?view=azure-cli-latest>
- Kubernetes command-line tool – Instructions for doing this can be found here: <https://kubernetes.io/docs/tasks/tools/install-kubectl/>

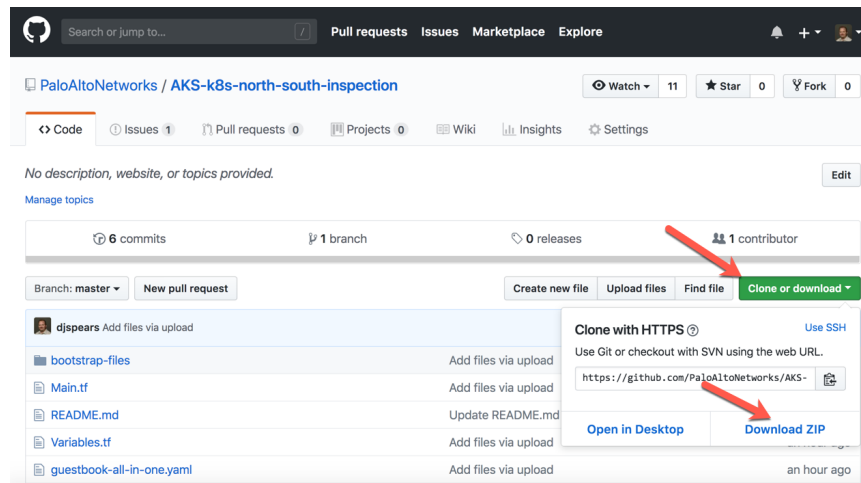


# Download GitHub files


In this activity, you will:

Download a zip copy of the GitHub files used for this lab

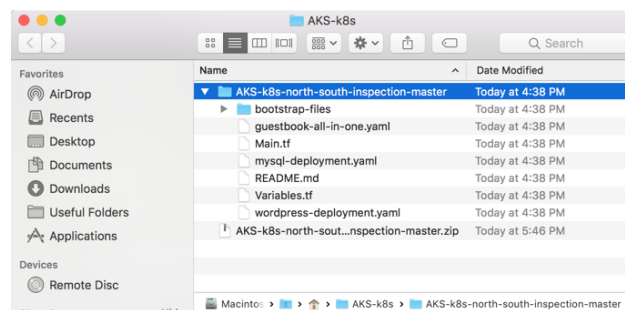
During this lab, the Terraform templates and Kubernetes (k8s) command will be executed from a local computer. This lab requires some customization of the terraform files. To download the files from GitHub, click on the Clone or download drop down and select Download ZIP.



Save the zip file to a new directory. This directory will be used to deploy the Terraform template and will automatically keep the Terraform state files so the deployment can be managed in the future:



Unzip the files:



# Azure Service Principal Creation

*In this activity, you will:*


*Authenticate to an Azure subscription via the Azure command line tool*

*Create a Service Principal with the appropriate RBAC to deploy a kubernetes (k8s) cluster*

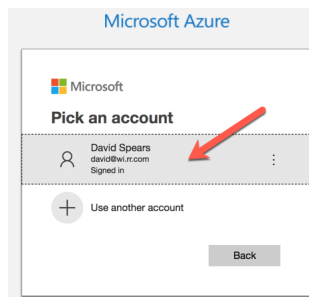
*Update the Terraform Variables.tf file with the Service Principal information needed to execute*

Microsoft has documented the steps to create a service principal that can be used to deploy a k8s cluster. That document can be found here: <https://docs.microsoft.com/en-us/azure/container-service/kubernetes/container-service-kubernetes-service-principal>

This guide assumes that the prerequisites have been completed and the Azure command line tool has been installed. Open a terminal window and type the command **az login** to authenticate the command line tool to the appropriate subscription:



Next a browser window should open that will give the option to select the Azure account associate with the subscription that will get the deployment:

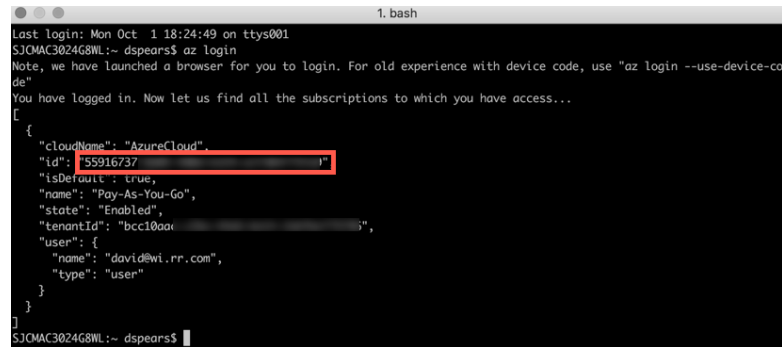


Once the account has been selected, the following message will appear:

**You have logged into Microsoft Azure!**

You can close this window, or we will redirect you to the [Azure CLI documents](#) in 10 seconds.

Check the terminal window. There should be confirmation that the login process was a success:

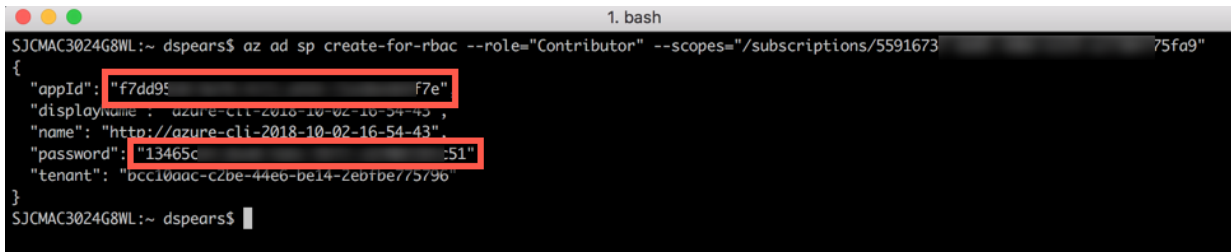


```
1. bash
Last login: Mon Oct 1 18:24:49 on ttys001
SJCMAC3024G8WL:~ dspears$ az login
Note, we have launched a browser for you to login. For old experience with device code, use "az login --use-device-code"
You have logged in. Now let us find all the subscriptions to which you have access...
[
  {
    "cloudName": "AzureCloud",
    "id": "55916737",
    "isDefault": true,
    "name": "Pay-As-You-Go",
    "state": "Enabled",
    "tenantId": "bcc10aac",
    "user": {
      "name": "david@wi.rr.com",
      "type": "user"
    }
  }
]
SJCMAC3024G8WL:~ dspears$
```

Copy the “id” from the output. This is the subscription id for the service principal. To be able to deploy a k8s cluster in Azure the service principal must have the “contributor” role. Use the following command to create the service principal:

**\$ az ad sp create-for-rbac --role="Contributor" --scopes="/subscriptions/<id>"**

where “id” is the subscription id copied from the last step:



```
1. bash
SJCMAC3024G8WL:~ dspears$ az ad sp create-for-rbac --role="Contributor" --scopes="/subscriptions/559167375fa9"
{
  "appId": "f7dd95f7e",
  "displayName": "azure-cli-2018-10-02-16-54-43",
  "name": "http://azure-cli-2018-10-02-16-54-43",
  "password": "13465c51",
  "tenant": "bcc10aac-c2be-44e6-be14-2e01be77579b"
}
SJCMAC3024G8WL:~ dspears$
```

The Terraform deployment files consist of a main, variables, and output files. The Variables.tf file contains information that is easily modified and commonly changed for various situations. The variables in the Variables.tf file are used by the Main.tf file during deployment. Deploying this Terraform template in Azure does require modification of the Variable.tf file to include deployment-specific information.

Copy the “appId” and “password” fields from the service principal creation output. These are needed for the terraform script and need to be added to the Variables.tf file. Open an editor of your choice and update these fields and save the file:

```
1 // PROJECT Variables
2 variable "client_id" {
3   default = "<appId>"
4 }
5 variable "client_secret" {
6   default = "<password>"
7 }
8
9 variable "agent_count" {
10  default = 2
11 }
12
13 variable "ssh_public_key" {
14   default = "<path to public ssh key>"
15 }
16
17 variable "dns_prefix" {
18   default = "k8s-AZURE-HOW"
19 }
20
21 variable cluster_name {
22   default = "k8s-Cluster-MGMT"
23 }
24
25 variable resource_group_name {
26   default = "k8s-RG"
27 }
28
29 variable location {
```

```
1 // PROJECT Variables
2 variable "client_id" {
3   default = "f7dd95f7e"
4 }
5
6 variable "client_secret" {
7   default = "13465c:5j"
8 }
9
10 variable "agent_count" {
11  default = 2
12 }
13
14 variable "ssh_public_key" {
15   default = "/AKS/AKS-PANFW-k8s/djs-gcp-keyb.pub"
16 }
17
18 variable "dns_prefix" {
19   default = "k8s-AZURE-HOW"
20 }
21
22 variable cluster_name {
23   default = "k8s-Cluster-MGMT"
24 }
25
26 variable resource_group_name {
27   default = "k8s-RG"
28 }
29
30 variable location {
31   default = "Central US"
32 }
33 }
```

## Bootstrap storage account creation

*In this activity, you will:*


*Create an Azure Resource Group and deploy a storage account*

*Create a file share with the folder structure needed to bootstrap the VM-Series Firewall*

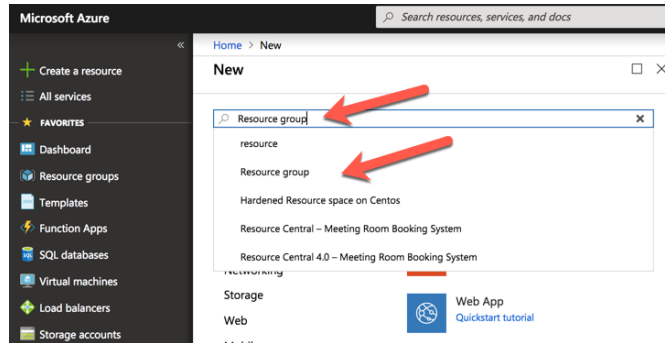
*Copy the files to the Azure file share needed for bootstrapping*

*Update the Terraform Variables.tf file with the Azure storage access key that will allow the VM-Series Firewall to bootstrap*

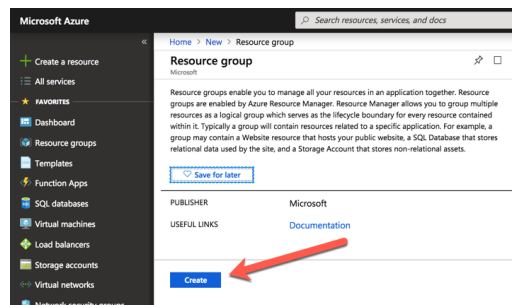
The terraform template is going to bootstrap the initial VM-Series firewall configuration. To accomplish this an Azure storage account will be created with the appropriate files. To start, open the Azure Portal and create a new resource group. Click on the “+ Create a resource” link:



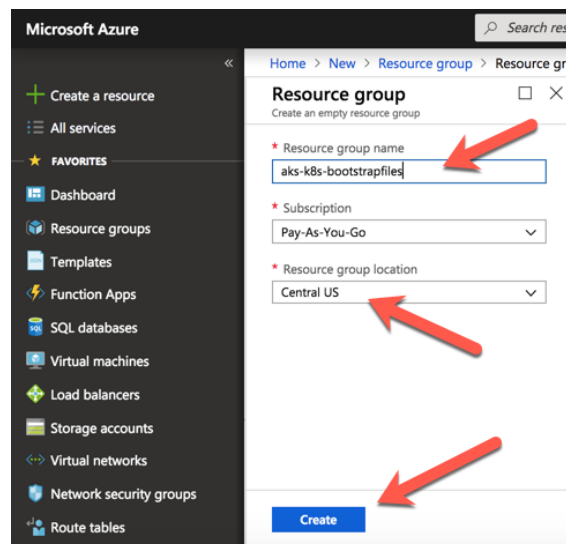
Next enter “Resource group” in the search and select Resource group:



Next select “Create” to create:



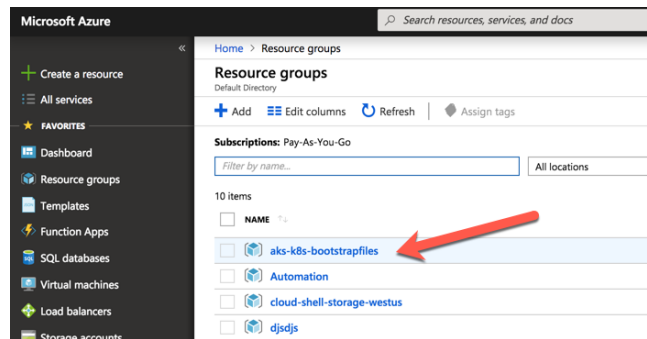
In the next window, create a resource group name and select the Resource group location. It is recommended for this lab to use the same location that the terraform script deploys in. The default setting is Central US. Click “Create” to create the Resource group.



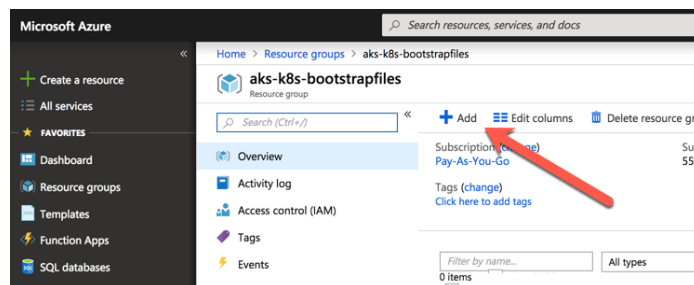
Navigate to the new Resource group. If a favorite is not available, click the “All Services” option on the left Nav and type “resource” in the All services search window. Click on Resource groups to open all the resources.




Now click the newly created Resource group:




Once in the resource group the next step is to create a storage account. Click on the plus sign to add a resource in the resource group:




Type storage account in the search field:



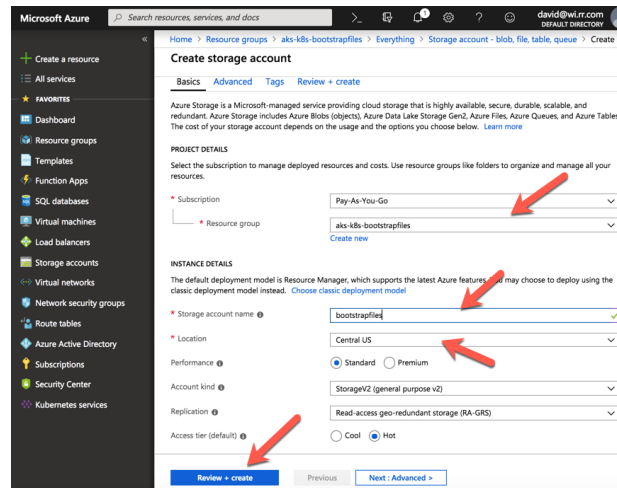
Select the Storage account published by Microsoft:



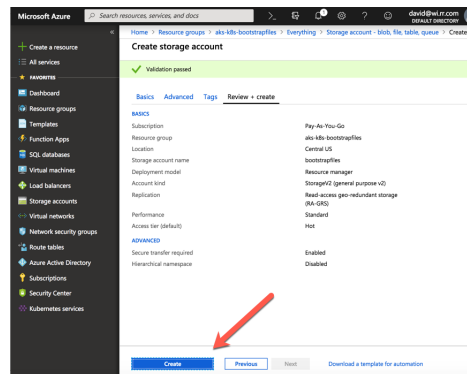
Next click “Create”:



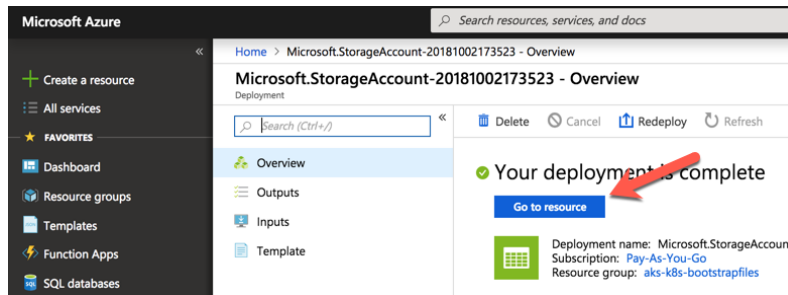
Make sure the Resource group is correct. Enter a Storage account name and select the same location as the rest of the deployment. Finally click “Review and create”



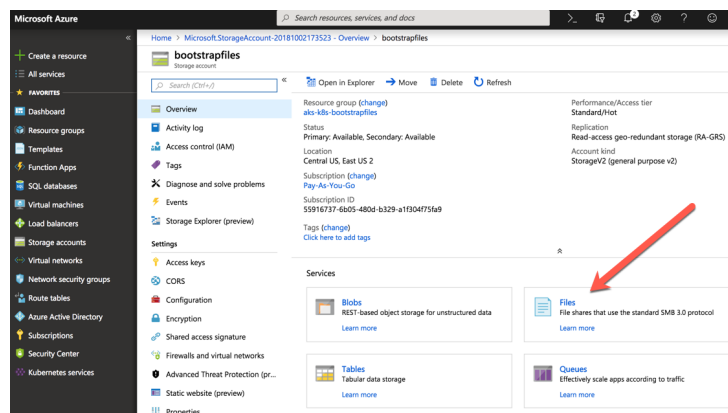
Once the validation is complete, select Create:



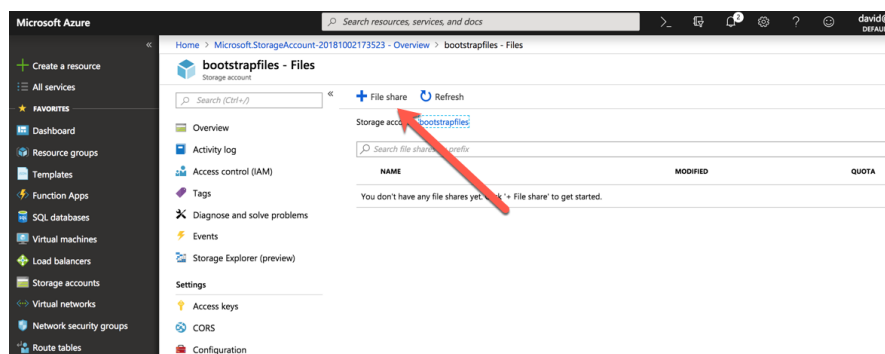
After the deployment is complete, click on the go to resource button:



Once the storage account is open. Click on the Files section. This is where the folders and files to bootstrap the firewall will be placed.




Next click the plus sign to create a new File Share:






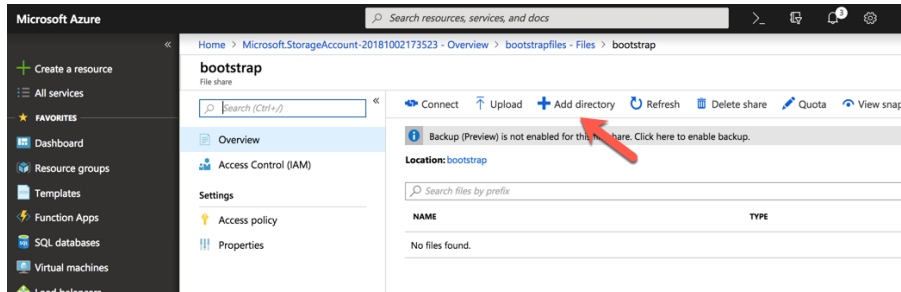
When the dialogue window opens, enter the file share information and click create. Note: The Name will be used to update the Variables.tf file in a few steps:



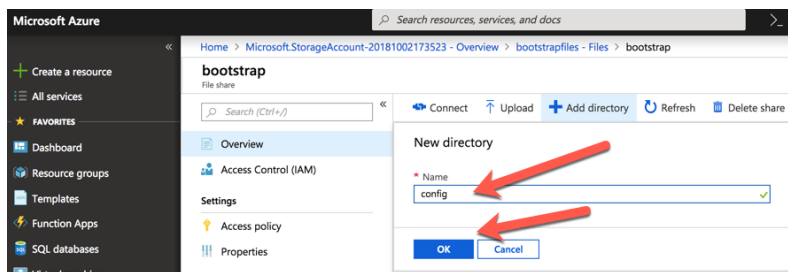
Click on the newly created file share:



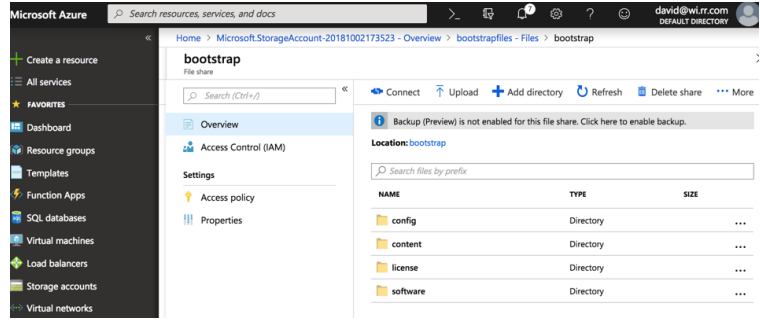
Click on the “Add directory” to create a directory:



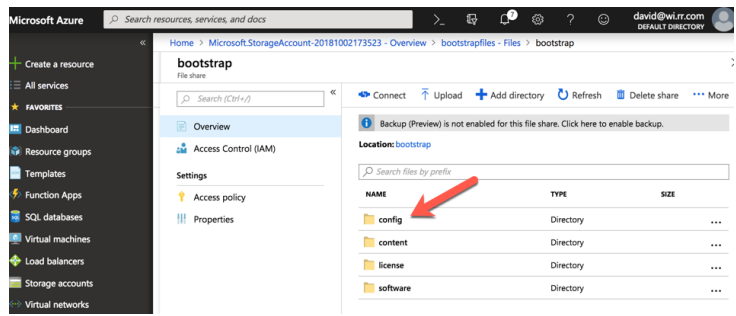
Enter config and click ok:




Repeat this step to create a content, license, and software directory. It is important that all 4 directories are present:



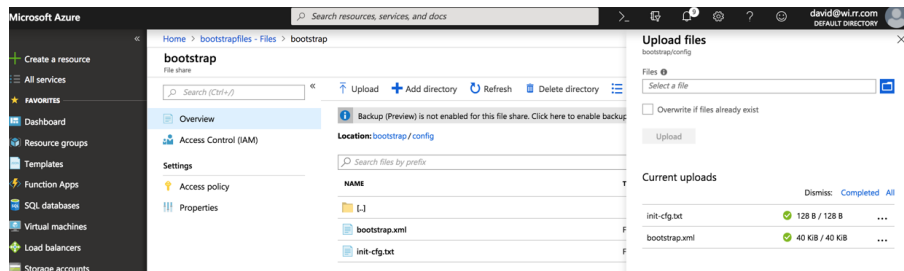
Click on the config folder:



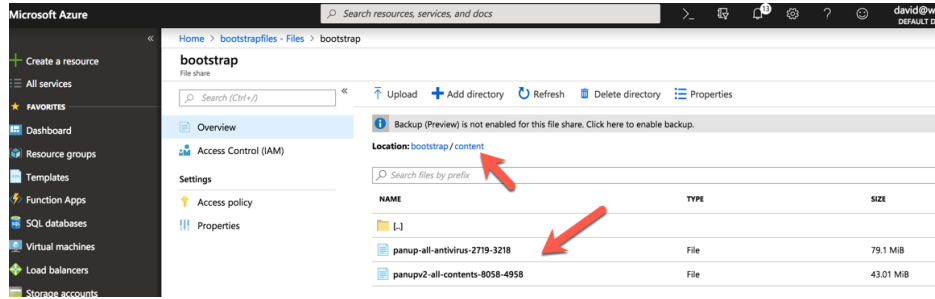
Click "Upload". When the upload blade opens, select the folder browse and navigate to the files previously downloaded from GitHub. Select the bootstrap.xml and init-cfg.txt. Then click "Upload":



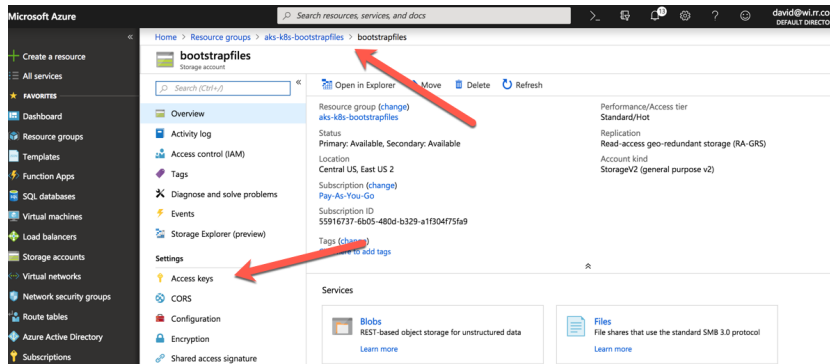
Once the files have been uploaded, they should be visible in the directory:



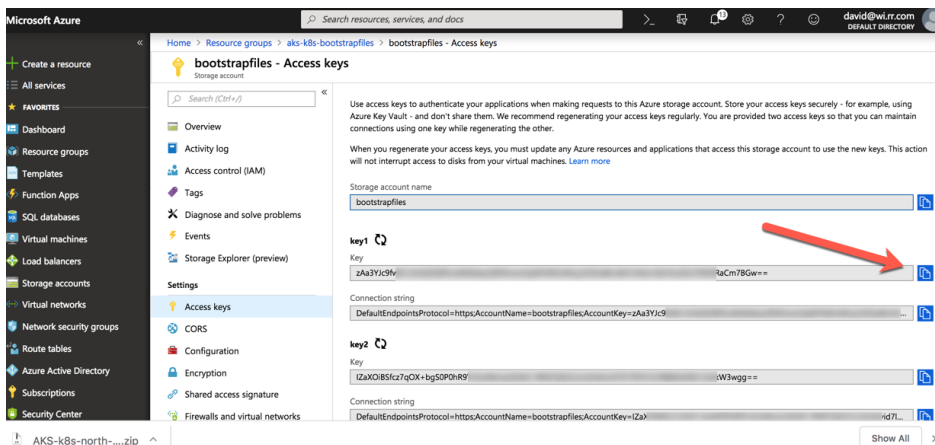
It is also possible to add content updates to the content directory that will get loaded into the firewall during the bootstrapping process. The follow figure shows some content files uploaded to the content directory:



The next step is to identify the Access Key and update the Terraform Variables.tf file. Navigate to the Storage account and click Access keys:



Next click the copy button to copy the access key for the storage account:



Open the Variables.tf file in an editor and update the custom data variable. The access key, storage account name, and share name need to be added:

```
3 variable "fwOffer" {
4   default = "vmseries1"
5 }
6
7 variable "fwPublisher" {
8   default = "paloaltonetworks"
9 }
10
11 variable "adminUsername" {
12   default = "paloalto"
13 }
14 variable "adminPassword" {
15   default = "Pal0Alt0@123"
16 }
17 variable "gvmSize" {
18   default = "Standard_A1"
19 }
20 variable "customdata" {
21   default = "storage-account=<insert storage account name>,access-key=<insert file key>,file-share=<insert share name>,share-directory=None"
22 }
23 }
24
25
26
```

This is a screen shot of the file with the updated information:

```
93 variable "fwOffer" {
94   default = "vmseries1"
95 }
96
97 variable "fwPublisher" {
98   default = "paloaltonetworks"
99 }
100
101 variable "adminUsername" {
102   default = "paloalto"
103 }
104 variable "adminPassword" {
105   default = "Pal0Alt0@123"
106 }
107 variable "gvmSize" {
108   default = "Standard_A1"
109 }
110 variable "customdata" {
111   default = "storage-account=bootstrapfiles,access-key=za3YJc5n78Gw==,file-share=bootstrap,share-directory=None"
112 }
113 }
114
```

## SSH keys

*In this activity, you will:*

**Generate SSH Keys – if needed**

**Update the Terraform Variables.tf with the path to the SSH keys**

The Terraform Variables.tf file has an option for supplying ssh keys that can be used to log into the Kubernetes nodes after deployment.

If you do not already have an SSH key, the follow example shows how to create an SSH key on a Mac using the `ssh-keygen -t rsa` command:

```

3. bash
SJC3024G8WL:AKS-k8s dspears$ ssh-keygen -t rsa
Generating public/private rsa key pair.
Enter file in which to save the key (/Users/dspears/.ssh/id_rsa): /Users/dspears/AKS-k8s/djs-aks-key
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /Users/dspears/AKS-k8s/djs-aks-key.
Your public key has been saved in /Users/dspears/AKS-k8s/djs-aks-key.pub.
The key fingerprint is:
SHA256:svhFXNTZISKgk7SsudAsf17Mmzwh+4sh0KIYR8d2c dspears@SJC3024G8WL
The key's randomart image is:
+---[RSA 2048]-----+
|.. . . E      |
|. + . + .    |
|. * . . . . .|
|+. = .       |
|.| o . S .   |
|+ oo . o . . |
|o+o*.o o    |
|..++X.o .   |
| .B=*..     |
+---[SHA256]-----+
SJC3024G8WL:AKS-k8s dspears$

```


In the previous example the keys were generated and stored in the same directory as the other lab files. The public and private keys can be seen using the `ls -la` command.

```

SJC3024G8WL:AKS-k8s dspears$ ls -la
total 64
drwxr-xr-x@ 7 dspears PALOALTONETWORK\Domain Users 224 Oct 2 20:42 .
drwxr-xr-x@ 59 dspears PALOALTONETWORK\Domain Users 1888 Oct 2 17:47 ..
-rw-r--r--@ 1 dspears PALOALTONETWORK\Domain Users 6148 Oct 2 17:49 .DS_Store
drwxr-xr-x@ 10 dspears PALOALTONETWORK\Domain Users 320 Oct 2 20:33 AKS-k8s-north-south-inspection-master
-rw-r--r--@ 1 dspears PALOALTONETWORK\Domain Users 12850 Oct 2 17:46 AKS-k8s-north-south-inspection-master.zip
-rw----- 1 dspears PALOALTONETWORK\Domain Users 1679 Oct 2 20:42 djs-aks-key
-rw-r--r-- 1 dspears PALOALTONETWORK\Domain Users 404 Oct 2 20:42 djs-aks-key.pub
SJC3024G8WL:AKS-k8s dspears$ pwd
/Users/dspears/AKS-k8s
SJC3024G8WL:AKS-k8s dspears$

```

Next edit the Terraform Variables.tf file to include the path to the public SSH key. The following diagram shows the field that needs to be updated and the field after it has been updated:

<pre> 1 // PROJECT Variables 2 variable "client_id" { 3   default = "&lt;appId&gt;" 4 } 5 variable "client_secret" { 6   default = "&lt;password&gt;" 7 } 8 9 variable "agent_count" { 10  default = 2 11 } 12 13 variable "ssh_public_key" { 14   default = "&lt;path to public ssh key&gt;" 15 } 16 17 variable "dns_prefix" { 18   default = "k8s-AZURE-HOW" 19 } 20 21 variable cluster_name { 22   default = "k8s-Cluster-MGMT" </pre>		<pre> 1 // PROJECT Variables 2 variable "client_id" { 3   default = "&lt;appId&gt;" 4 } 5 variable "client_secret" { 6   default = "&lt;password&gt;" 7 } 8 9 variable "agent_count" { 10  default = 2 11 } 12 13 variable "ssh_public_key" { 14   default = "/Users/dspears/AKS-k8s/djs-aks-key.pub" 15 } 16 17 variable "dns_prefix" { 18   default = "k8s-AZURE-HOW" 19 } 20 21 variable cluster name { </pre>
---	---	---

# Deploy the Terraform Template

*In this activity, you will:*

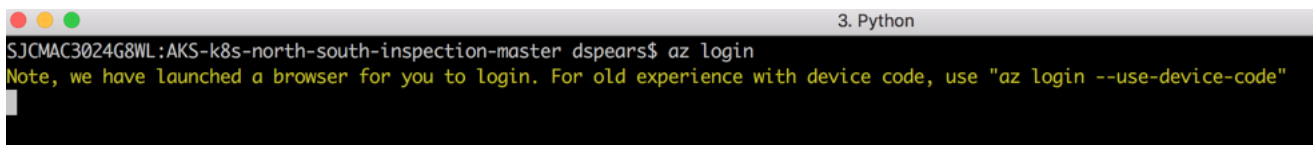
*Authenticate to Azure via the Azure command line tool*

*Initialize Terraform and download the appropriate plugins*

*Apply the Terraform template*

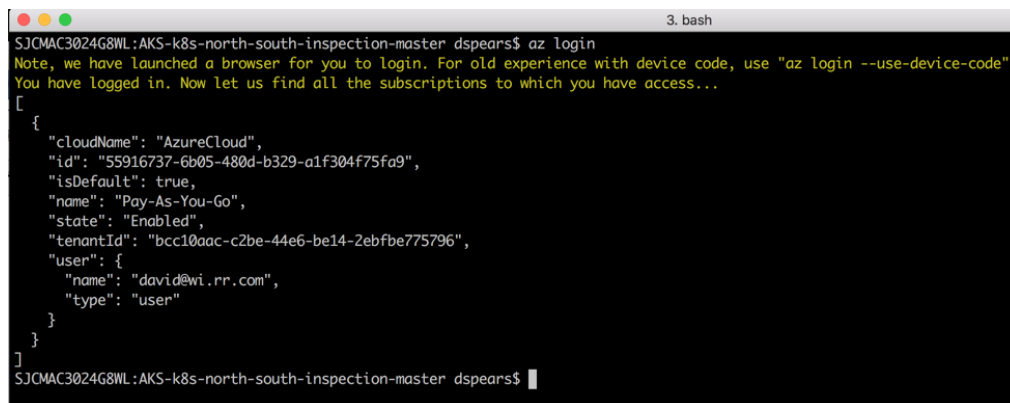
Open a terminal shell and navigate to the directory containing the Terraform template files.

The Azure cli tool token obtained earlier has most likely expired. Use the “**az login**” login command to get a new token:



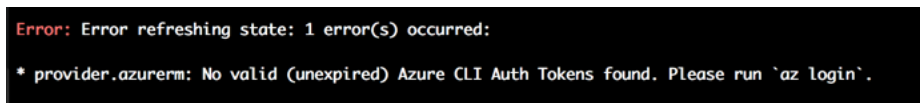
```
SJCMAC3024G8WL:AKS-k8s-north-south-inspection-master dspears$ az login
Note, we have launched a browser for you to login. For old experience with device code, use "az login --use-device-code"
```

After getting redirected to the Microsoft Azure Login and completing the login process successfully, the following prompt will be displayed:



```
SJCMAC3024G8WL:AKS-k8s-north-south-inspection-master dspears$ az login
Note, we have launched a browser for you to login. For old experience with device code, use "az login --use-device-code"
You have logged in. Now let us find all the subscriptions to which you have access...
[
  {
    "cloudName": "AzureCloud",
    "id": "55916737-6b05-480d-b329-a1f304f75fa9",
    "isDefault": true,
    "name": "Pay-As-You-Go",
    "state": "Enabled",
    "tenantId": "bcc10aac-c2be-44e6-be14-2ebf75796",
    "user": {
      "name": "david@wi.rr.com",
      "type": "user"
    }
  }
]
SJCMAC3024G8WL:AKS-k8s-north-south-inspection-master dspears$
```

As a note, the following error message is displayed when the azure cli tool token has expired:



```
Error: Error refreshing state: 1 error(s) occurred:
* provider.azurearm: No valid (unexpired) Azure CLI Auth Tokens found. Please run `az login`.
```

Ensure you are in the directory with the Main.tf and Variables.tf files and execute the “**terraform init**” command which will initialize terraform and ensure all the provider plugins are download and up to date:

```
3. bash
SJCMAC3024G8WL:AKS-k8s-north-south-inspection-master dspears$ terraform init

Initializing provider plugins...

Terraform has been successfully initialized!

You may now begin working with Terraform. Try running "terraform plan" to see
any changes that are required for your infrastructure. All Terraform commands
should now work.

If you ever set or change modules or backend configuration for Terraform,
rerun this command to reinitialize your working directory. If you forget, other
commands will detect it and remind you to do so if necessary.
SJCMAC3024G8WL:AKS-k8s-north-south-inspection-master dspears$
```

Once the terraform init has completed run the **terraform plan** command. This will show what changes will be implemented with the terraform script. This will also identify if there are any errors detected with the terraform files:

```
3. bash
SJCMAC3024G8WL:AKS-k8s-north-south-inspection-master dspears$ terraform plan
Refreshing Terraform state in-memory prior to plan...
The refreshed state will be used to calculate this plan, but will not be
persisted to local or remote state storage.

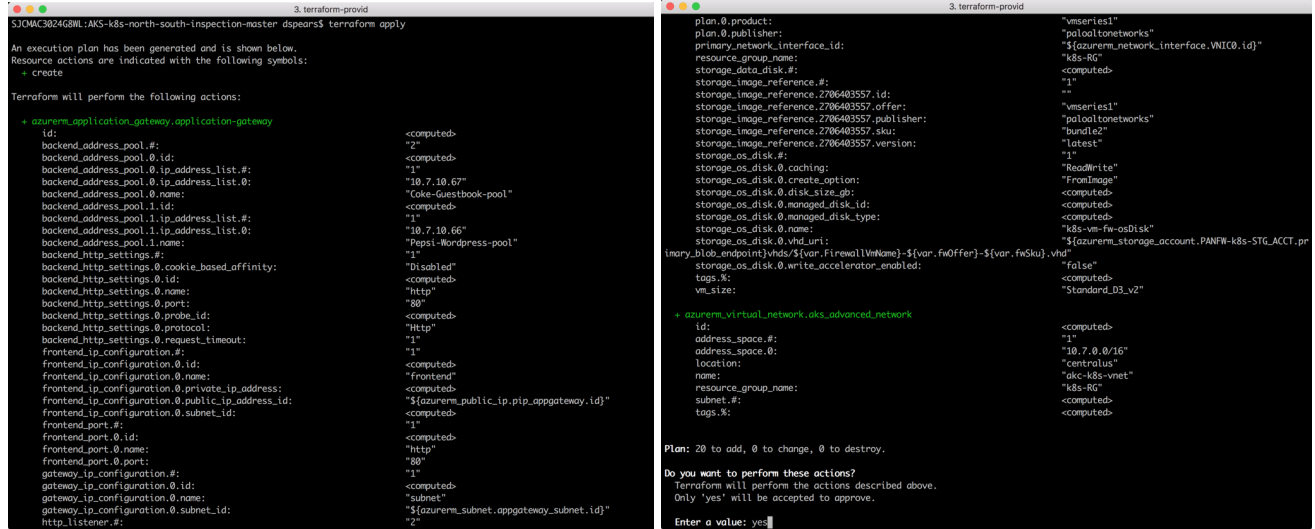
-----
An execution plan has been generated and is shown below.
Resource actions are indicated with the following symbols:
+ create

Terraform will perform the following actions:

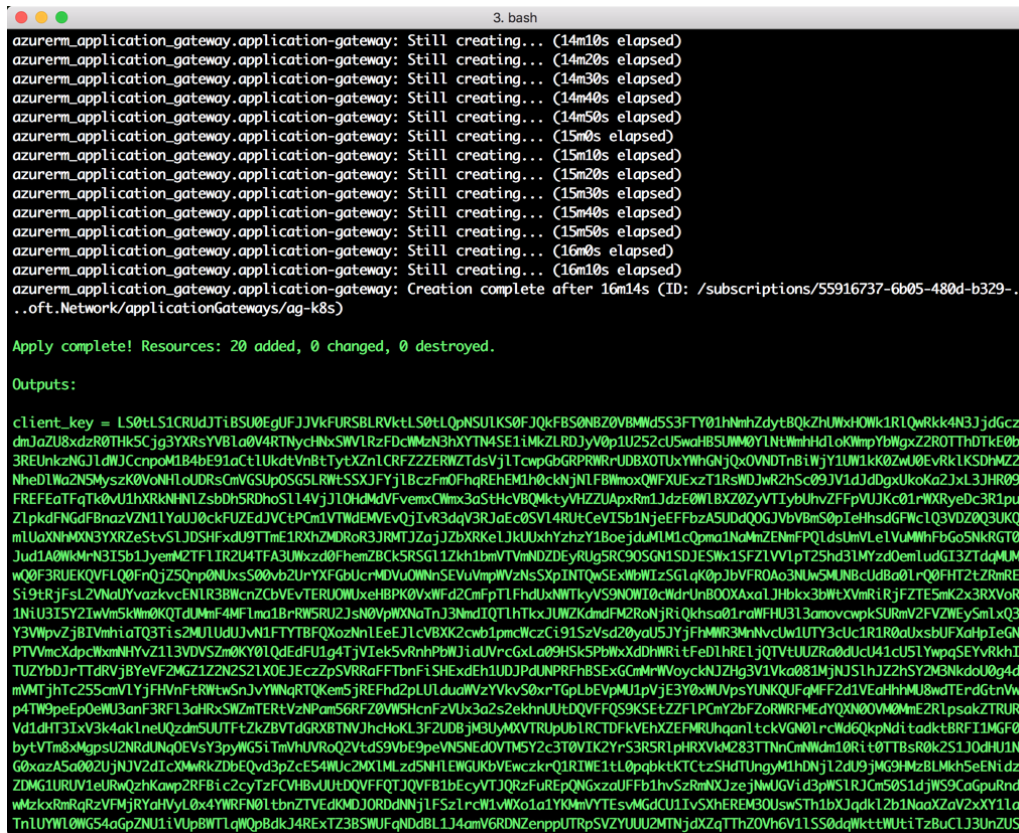
+ azurem_application_gateway.application-gateway
  id: <computed>
  backend_address_pool.#: "2" <computed>
  backend_address_pool.0.id: <computed>
  backend_address_pool.0.ip_address_list.#: "1"
  backend_address_pool.0.ip_address_list.0: "10.7.10.67"
  backend_address_pool.0.name: "Coke-Guestbook-pool"
  backend_address_pool.1.id: <computed>
  backend_address_pool.1.ip_address_list.#: "1"
  backend_address_pool.1.ip_address_list.0: "10.7.10.66"
  backend_address_pool.1.name: "Pepsi-WordPress-pool"
  backend_http_settings.#: "1"
  backend_http_settings.0.cookie_based_affinity: "Disabled"
  backend_http_settings.0.id: <computed>
  backend_http_settings.0.name: "http"
  backend_http_settings.0.port: "80"
  backend_http_settings.0.protocol: <computed>
  backend_http_settings.0.probe_id: "http"
  backend_http_settings.0.request_timeout: "1"
  frontend_ip_configuration.#: "1"
  frontend_ip_configuration.0.id: <computed>
  frontend_ip_configuration.0.name: "frontend"
  frontend_ip_configuration.0.private_ip_address: <computed>
  frontend_ip_configuration.0.public_ip_address_id: "${azurem_public_ip.pip_appgateway.id}"
  frontend_ip_configuration.0.subnet_id: <computed>
  frontend_port.#: "1"
  frontend_port.0.id: <computed>
  frontend_port.0.name: "http"
```



Now run the **terraform apply** command to deploy the template. At the action prompt enter **yes**.



It will take a few minutes to complete. If all goes well, Terraform will output; “Apply Complete!” and provide some additional output information about the resources deployed:





# Review what was deployed

*In this activity, you will:*

*Review the resources that have been launched*

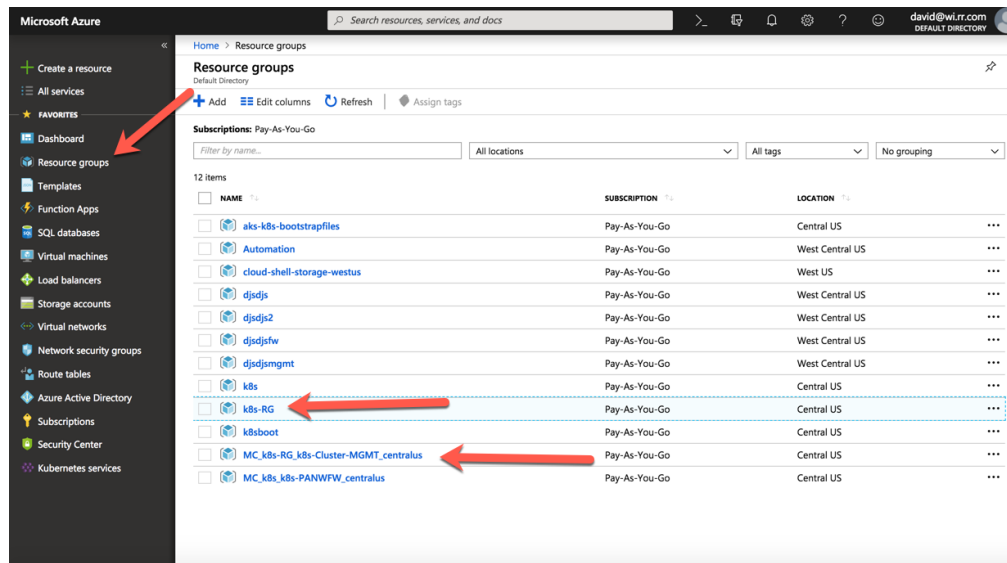
*Inspect k8s cluster*

*Log into the VM-Series firewall*

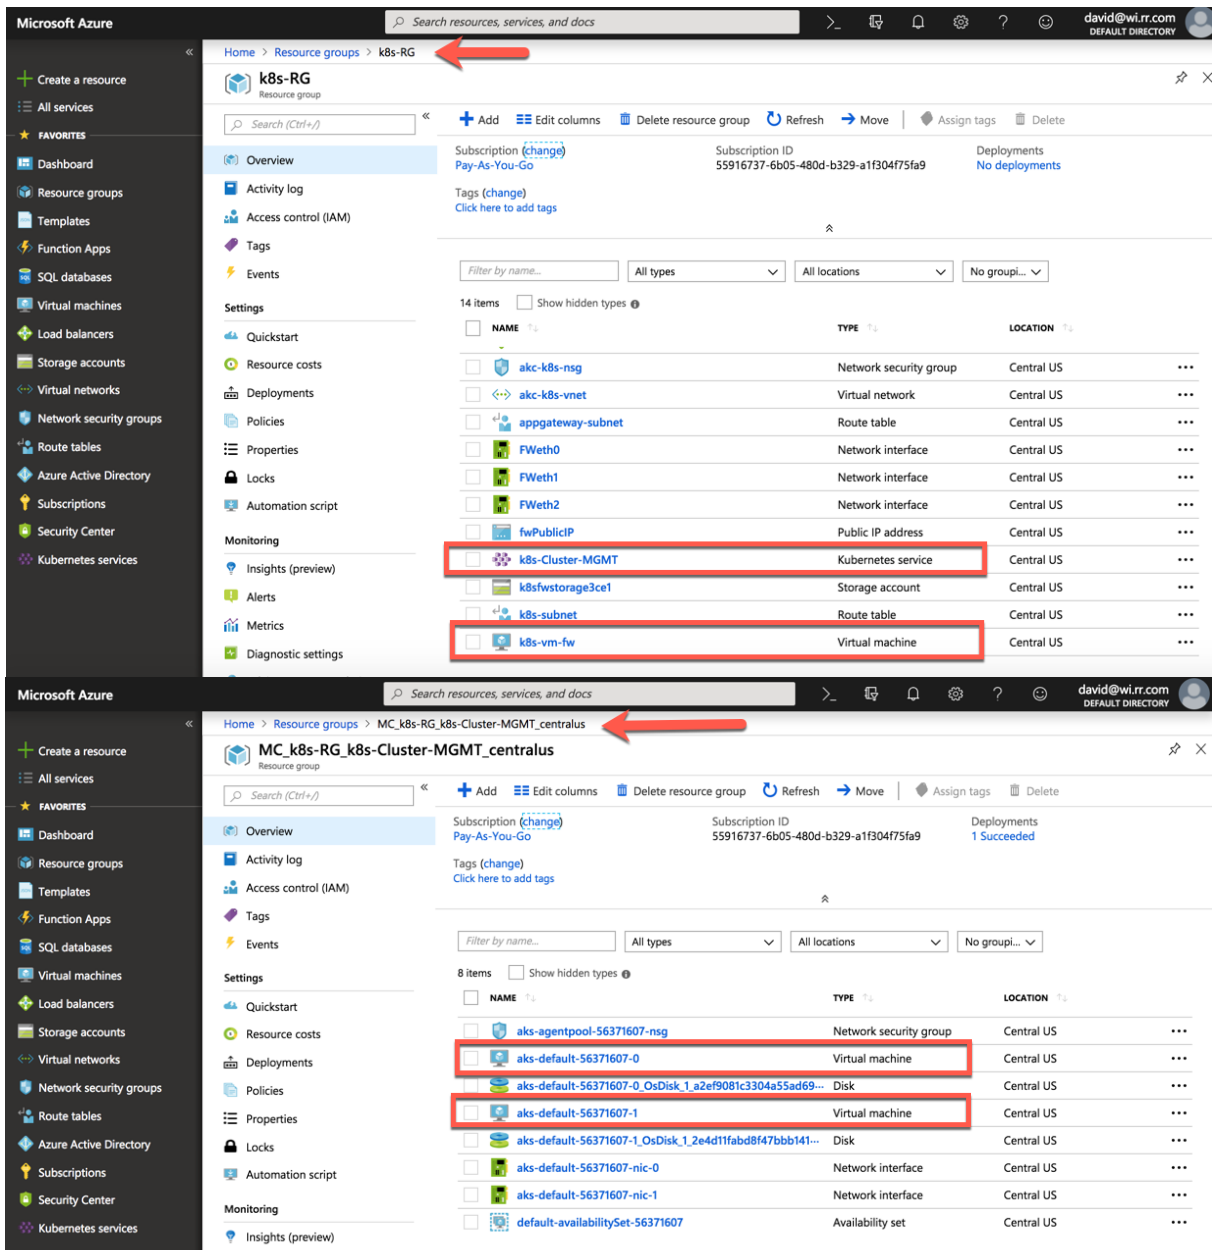
*Confirm bootstrap success*

## Task 1 – Look around Azure console

Navigate to Resource Groups. Notice that there are two resource groups that were deployed. The first one, k8s-RG, has the infrastructure that was defined in the Terraform template. The second has the k8s nodes and associated resources.

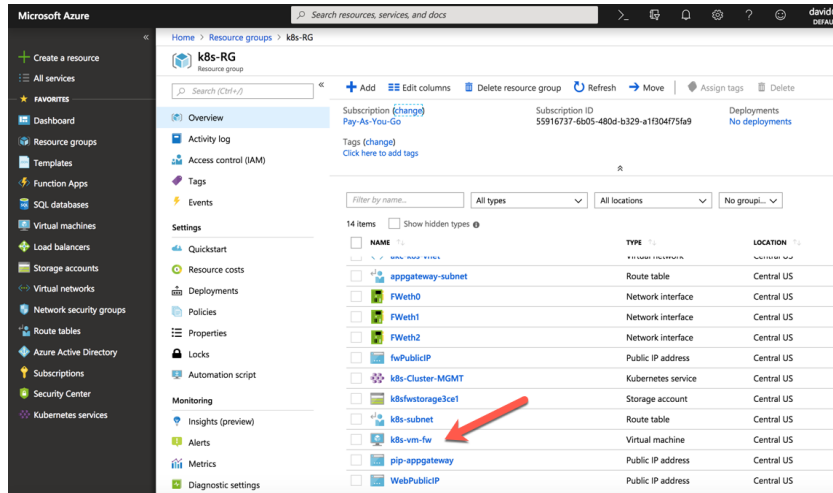


Open the two resource groups to view what has been deployed:

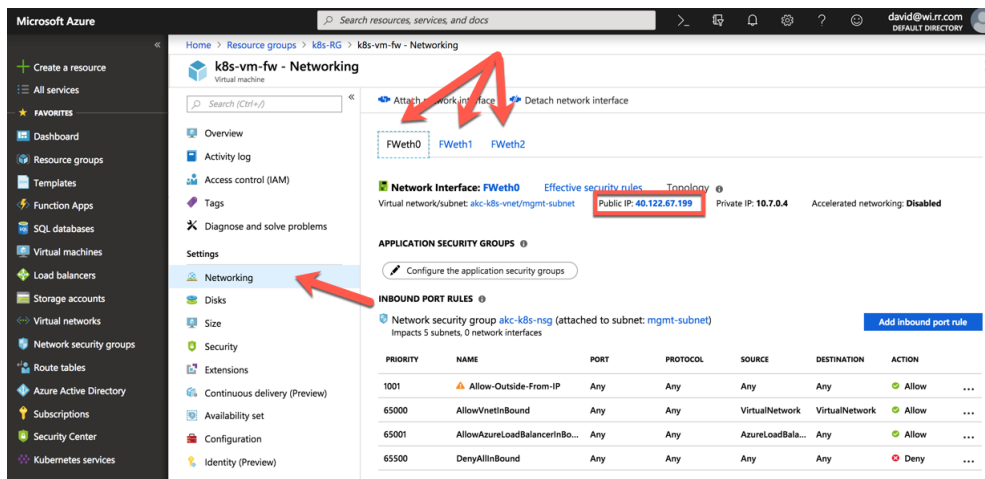


There should be 1 firewall, 1 k8s service master, and two k8s nodes displayed.

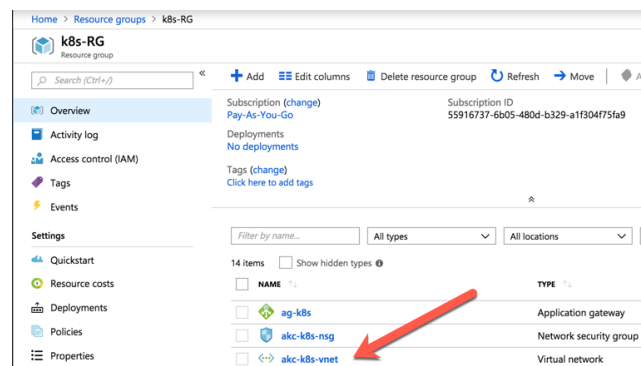
Click on the firewall to open a detailed view of the deployed firewall:



Explore the options on the firewall. One interesting area to review is the Networking section. The IP address and security information for each interface can be identified:

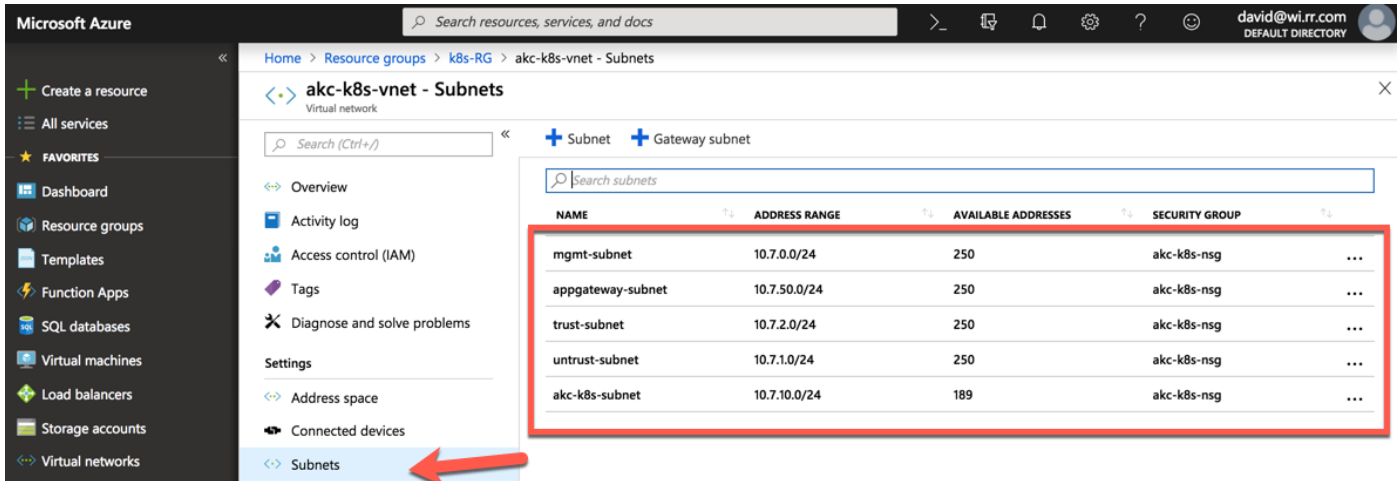


Navigate to akc-k8s-vnet virtual network in the k8s-RG resource group to see the different networks that have been created as part of the lab.

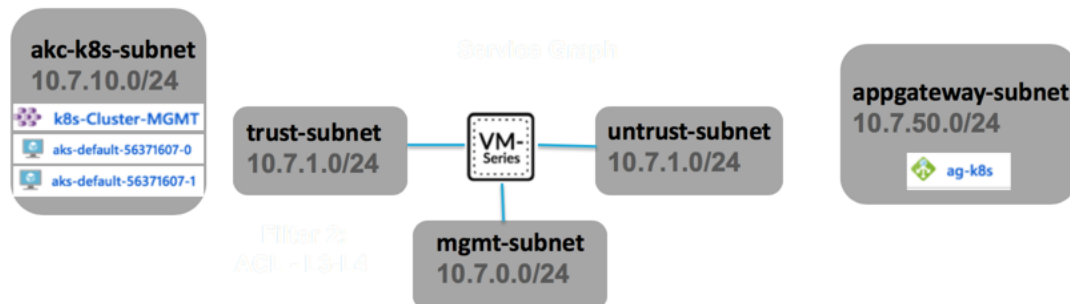


Click Subnets on the left Nav. You should see 5 subnets:

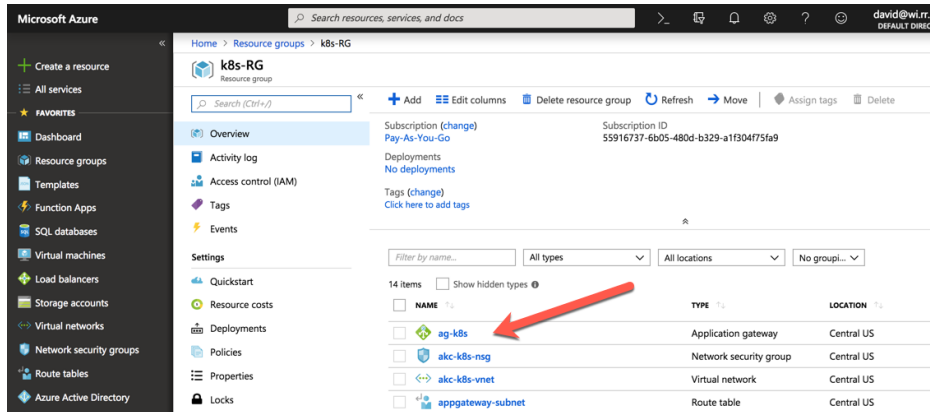
- mgmt-subnet, trust and untrust are used by the firewall
- appgateway-subnet is used by the application gateway
- akc-k8s-subnet is where the k8s nodes and load balancing services are deployed



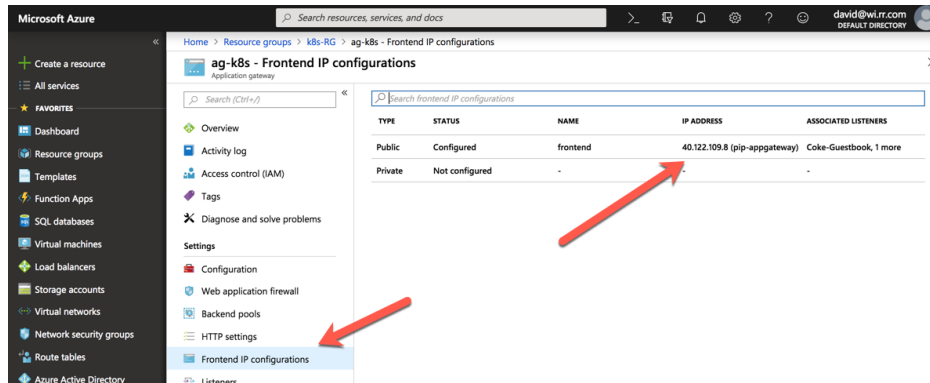
The following diagram describes the network topology of what has been deployed:



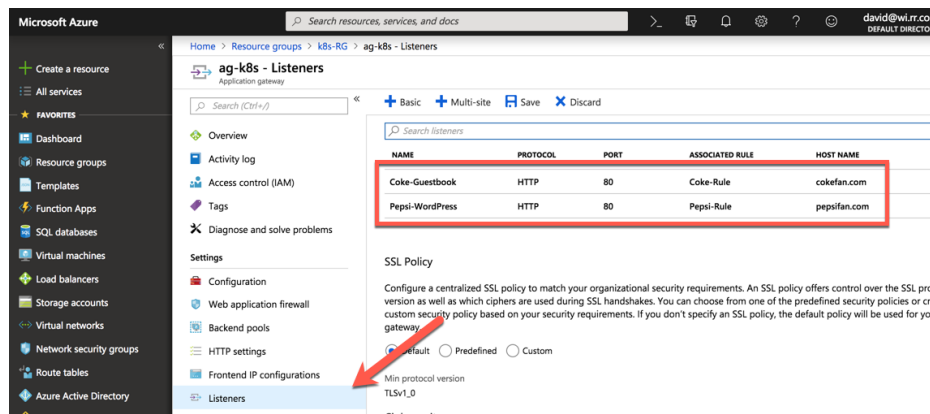
Next Navigate to the k8s-RG resource group and open the application gateway:



Click on the Frontend IP configurations options on the left Nav and notice that there is a single front-end IP address. The application gateways only support a single front-end IP address. This address will be needed later in the lab.



Next, go to Listeners on the left Nav. Notice that there are two listeners. This lab will leverage the Applications Gateway's ability to do host header redirection to send traffic to the correct internal load balancer address based on the http request.

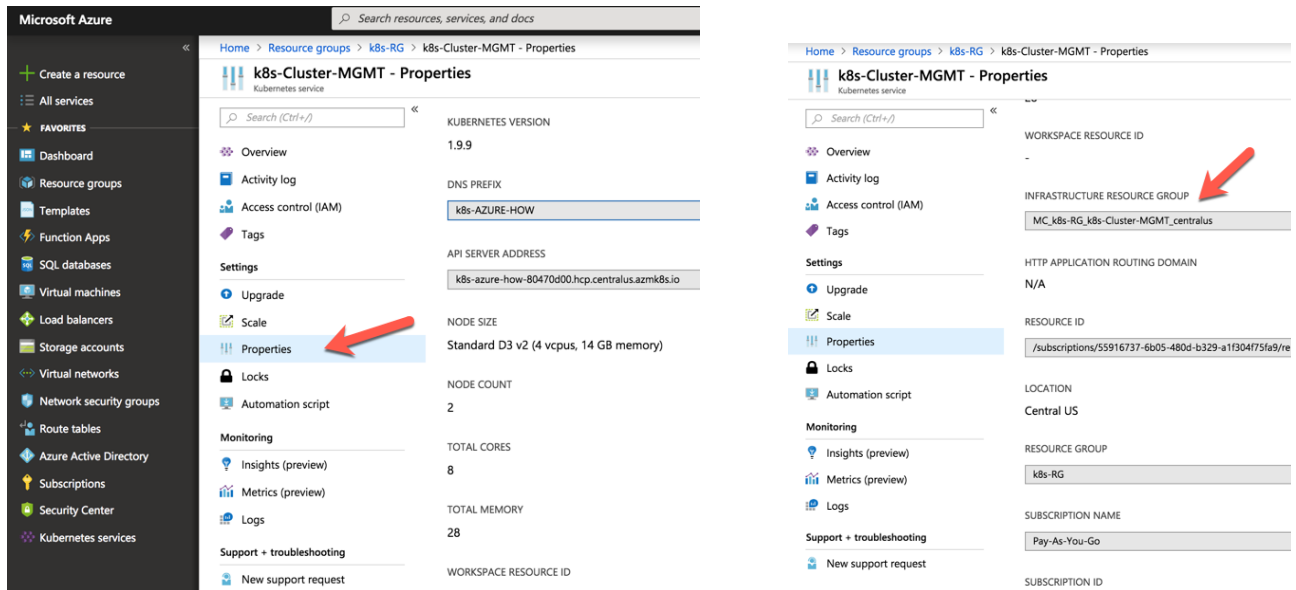


Feel free to navigate through other parts of the Azure Console. This will come in handy in activities later on.

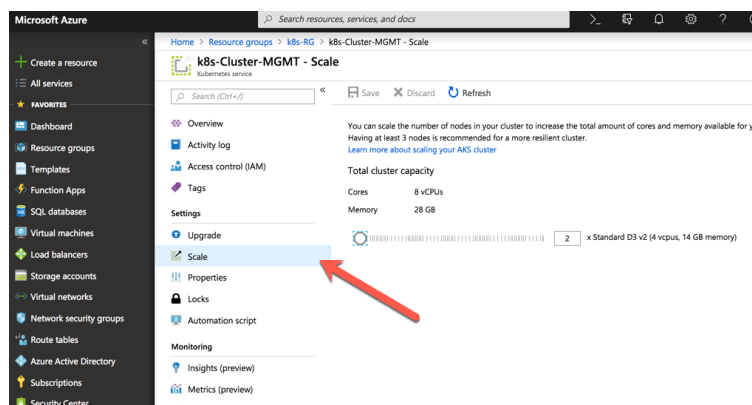
## Task 2 – Review the Kubernetes Cluster

Kubernetes is a portable, extensible, open-source orchestrator that is used to manage containerized workloads. Kubernetes has a large and rapidly growing ecosystem. The portability of Kubernetes allows for workloads to be migrated between various clouds (public or private). Further documentation is available at: <https://kubernetes.io/docs/concepts/overview/what-is-kubernetes/>

Navigate to the k8s-RG resource group and click on the k8s-Cluster-MGMT resource. Click on Properties in the k8s-Cluster-MGMT blade. This will show the k8s version, number of nodes deployed, and the infrastructure resource group that was created to deploy k8s resources. This is where the k8s nodes get deployed.



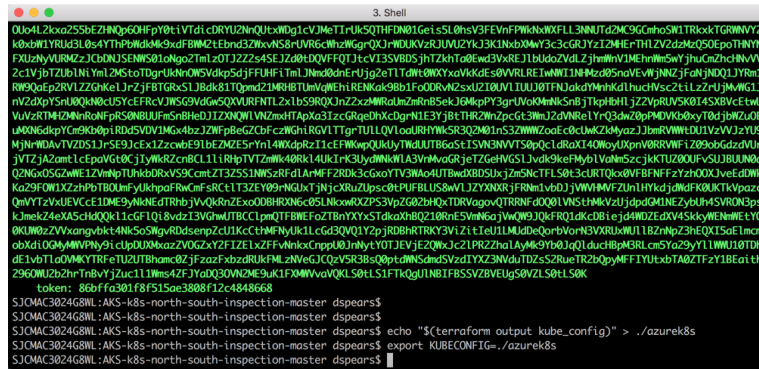
Clicking on the Scale link in the left Navigation displays the current number of nodes. From here the number of nodes deployed in the cluster can be increased or decreased.



## Task 3 – Connect to the Kubernetes Cluster

Navigate back to the terminal window used to deploy the Terraform script. In order to run Kubectl commands, the Kubernetes config from the Terraform state need to be captured and stored in a file that kubectl can read. Execute the following commands in the same directory that the terraform files are in:

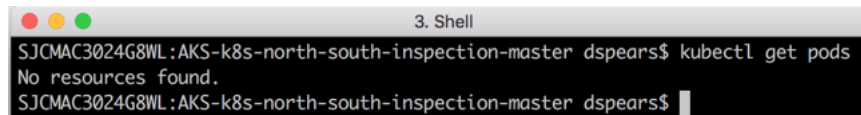
```
$ echo "$(terraform output kube_config)" > ./azurek8s
$ export KUBECONFIG=./azurek8s
```



```
3. Shell
0Uo4L2kxa255bEzInQp60HfPY0tiVTdi cDRYU2NnQUlxW0glcVJMeTIJuk5QTHFDN01Gei sSL0hsV3FEVnFPkNdiXFL3NNUTd2MC9GcnohSWiTRkxkTGRWVY2N
k0aW1YRUd3L0s4YTHPwJdKk9xdFBM2LEbnd3Z0xvNS8rUVRGWhzWggrQJrW0UKVzRUVU2YK3K1KxbMwY3C3cGRJYzI2MErTHLZVZdzZQ50ep0THNYN1
FXUzNyVURMzZjCjDNDSEMS0lOnqozTmlz0TJZ22s45EjZd0LQVFFQJtEcV13SVBDSjhtZkht0e0wd3VxRE1lBudoZVdLzjhmWVMEHrhmSwY7huChndhNwVU
ZcIVjPZT0bWmTm2K6toTdyUkHmW5Vd95QJFJBF1m1Jm00dne4Ujg28TlTant6WYxv0KkGE50RRLRE1wMT1Mkz05v0EwVnJNZJF0nJNQL1YrmlG
R9PQz0z2R1VZZ0KkL3rZrFBTGR51J8d8L1TQmZ2UW48B1Ublqf0h1REK0a9891F00BvM2exdZJ0U1L1U10TTLJ0k0MhKglhucd1F5czL1Lzr1JmMG1JG
WZ0zYRLURQ0k0U5YcEFRVJMSG9VdG65QVURFNTL2x1b59RQJnZz0MWRd0z0r0B0k30MkpY3grUv0k0NkS08jT0p0H1jZzVpRUS001450BVEt0wU
VwzRTH4ZMh0k0PFR50NBUJfms0BhdJ1Z2XW1VNZmHTApXg31zcGrqeDhXcdgrN1E3Yj8tTRZ2Wnzpc63WmJ2dVARE1YQ3d0z0p0M0Wk0k0yT0j0z0w0E5
M0N6d0pYc0K0k0PFR50NBUJfms0BhdJ1Z2XW1VNZmHTApXg31zcGrqeDhXcdgrN1E3Yj8tTRZ2Wnzpc63WmJ2dVARE1YQ3d0z0p0M0Wk0k0yT0j0z0w0E5
MjN0W0pYVZ0S1JrSE9Jc0x1Zz0w0E5Yn14W0k0Rz11cF0W0p0k0yT0W0U0T0B0S0E1SVN0WVTS0p0Qc1dR0X140W0yU0p0R0R0V0F1Z090b0d0z0Urb
jVTZjAZ0m1cEpaVg0cJjY0k0R2cnBCL111R0pVTZ0Wk0Rk14UK1rK3Uy0dWk0W1A3VhM0v0GRjeTZG0H0V0S1Jvd0K0eF0y0lV0N0S0c0jK0T0Z00F0V0U0J0UN0d3
Q2N0x0S0G0W0E1ZV0h0pT0U0k0R0xV59C0mLT3Z551NWSzRFd1ArMFFZ0R03c0G0YV3W0A04T0W0d0X0D0S0x0jZm5NcTFLS03cUR0T0k00F0BF0NFzYz00X0Jve0d0W0K5
Ka29F0W1XZ0hP0B0U0mfYUk0p0R0wCmF0R0c1T3ZEY09jN0G0X0TjNj0cR0Z0p0c0EPU0F0L0S0W1JZY0X0R0jFR0N1v0DjYVW0M0FZ0Jn1H0Yk0j0d0F0K0K0T0Vp0z03
0mVTZV0U0Vc0c1IM0S0k0N0d0R0b0jV0k0R0z0x000B0R0M0c05L0k0W0Z0P5VpZ0G0Z0H0Q0TDR0V0g0V0TR0NF0Q0L0VNS0h0k0Vz0j0p0d0M0E0Z0B0H45VR0N03psZ
0j0m0Z040AS0d0R0Q01Lc0r1Q10v0z0E0W0T0C0p0Q0T0B0E0F0Z0B0Y0K0S0T0k0h0R0Z00R0E0S0k0c0jW0W0J0Q0R010k00B1e0j040D0Z00V450k0y0B0W0E010W0
0K0M0Z0Y0W0x0p0k0tM0S0S0p0R0D0s0mp0Z0U1K0c0M0E0Y0K0L0c0d0Q0Q1Y2p0R0B0R0T0R0Y0N1Z1t0t011M0d0e0q0r0v0r0N0V0R0U0W0U11Bz0w0Z0h0X150e10nc0d
0b0d010Q0M0W0P0y0i0p0D0U0k0az0V0G0z0YF0Z0L0ZFF0W0k0C0pp0U0h0N0Y0T0E0V0E0Z0N0c0Z1PR0Z0h0A0M0S0Y001Q01d0c0B0M0R0L0m050z0y0Y11W0U010D0hZ
d0E1L0d0M0KY0TR0E0T0U0T0B0m00Z0jF0az0F0x0z0R0K0F0M0c0W0C0Q0z0R0S00p0d0MNS0d0Sv0z0d1YX0Z0N0W0U0T0Z0S02R0e0R0T0z0p0MFF0Y0L0Y0b0T00Z0F0Z10E0t0hL
2960W02b0z0h0Tn0b0YjZuc111Mms4ZF0Y0d0Q0W0M0E09uK1F0M0W0V0QKLS0HLS1FTk0gU1N01FB55VZ0V0E0S0VZLS0HLS0K
token: 86bffa301f8f515ae3808f12c4848668
SJCMA3024G8WL:AKS-k8s-north-south-inspection-master dspears$
SJCMA3024G8WL:AKS-k8s-north-south-inspection-master dspears$
SJCMA3024G8WL:AKS-k8s-north-south-inspection-master dspears$ echo "$(terraform output kube_config)" > ./azurek8s
SJCMA3024G8WL:AKS-k8s-north-south-inspection-master dspears$ export KUBECONFIG=./azurek8s
SJCMA3024G8WL:AKS-k8s-north-south-inspection-master dspears$
```

Let's explore some pods and services that have deployed. Run this command in the cloud shell:

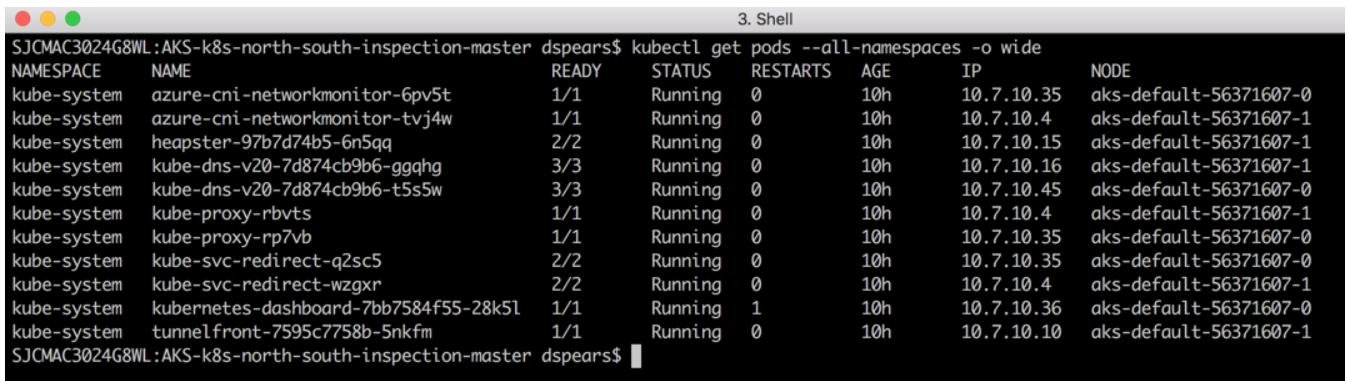
```
$ kubectl get pods
```



```
3. Shell
SJCMA3024G8WL:AKS-k8s-north-south-inspection-master dspears$ kubectl get pods
No resources found.
SJCMA3024G8WL:AKS-k8s-north-south-inspection-master dspears$
```

Since we have not deployed any resources this is normal. Now let us see what system pods have been deployed. Run this command in the shell:

```
$ kubectl get pods --all-namespaces -o wide
```



```
3. Shell
SJCMA3024G8WL:AKS-k8s-north-south-inspection-master dspears$ kubectl get pods --all-namespaces -o wide
NAMESPACE      NAME                                READY   STATUS    RESTARTS   AGE   IP              NODE
kube-system    azure-cni-networkmonitor-6pv5t     1/1     Running  0           10h   10.7.10.35      aks-default-56371607-0
kube-system    azure-cni-networkmonitor-tvj4w     1/1     Running  0           10h   10.7.10.4       aks-default-56371607-1
kube-system    heapster-97b7d74b5-6n5qq          2/2     Running  0           10h   10.7.10.15      aks-default-56371607-1
kube-system    kube-dns-v20-7d874cb9b6-ggqhg      3/3     Running  0           10h   10.7.10.16      aks-default-56371607-1
kube-system    kube-dns-v20-7d874cb9b6-t5s5w      3/3     Running  0           10h   10.7.10.45      aks-default-56371607-0
kube-system    kube-proxy-rbvts                   1/1     Running  0           10h   10.7.10.4       aks-default-56371607-1
kube-system    kube-proxy-rp7vb                   1/1     Running  0           10h   10.7.10.35      aks-default-56371607-0
kube-system    kube-svc-redirect-q2sc5            2/2     Running  0           10h   10.7.10.35      aks-default-56371607-0
kube-system    kube-svc-redirect-wzgxr            2/2     Running  0           10h   10.7.10.4       aks-default-56371607-1
kube-system    kubernetes-dashboard-7bb7584f55-28k5l 1/1     Running  1           10h   10.7.10.36      aks-default-56371607-0
kube-system    tunnelfront-7595c7758b-5nkfm       1/1     Running  0           10h   10.7.10.10      aks-default-56371607-1
SJCMA3024G8WL:AKS-k8s-north-south-inspection-master dspears$
```



**Note:** If the output does not show all the pods in a running state, wait and rerun the **kubectl get pods --all-namespaces -o wide** command until they do. An example of this is state is in the following screen-print:

```
davejspears@cloudshell:~ (djs-gcp-2018)$ kubectl get pods -o wide
NAMESPACE   NAME                                     READY   STATUS    RESTARTS   AGE   IP           NODE
kube-system  heapster-v1.4.3-6644cc4b46-dwn8s      0/2     Pending   0          11m   <none>       gke-cluster-1-default-pool-2df01519-n3w8
kube-system  heapster-v1.4.3-7875d9f9ff-mlv7g      2/2     Terminating 0      12m   10.8.0.4     gke-cluster-1-default-pool-2df01519-6sgf
kube-system  kube-dns-778977457c-2zwfg             3/3     Running   0          13m   10.8.0.3     gke-cluster-1-default-pool-2df01519-6sgf
kube-system  kube-dns-778977457c-144zf             0/3     ContainerCreating 0      12m   <none>       gke-cluster-1-default-pool-2df01519-n3w8
kube-system  kube-dns-autoscaler-7db47cb9b7-j8n7n  1/1     Running   0          13m   10.8.1.3     gke-cluster-1-default-pool-2df01519-n3w8
kube-system  kube-proxy-gke-cluster-1-default-pool-2df01519-6sgf  1/1     Running   0          12m   10.5.2.2     gke-cluster-1-default-pool-2df01519-6sgf
kube-system  kube-proxy-gke-cluster-1-default-pool-2df01519-n3w8  1/1     Running   0          12m   10.5.2.3     gke-cluster-1-default-pool-2df01519-n3w8
kube-system  kubernetes-dashboard-6bb875b5bc-n7wj8  1/1     Running   0          13m   10.8.0.2     gke-cluster-1-default-pool-2df01519-6sgf
davejspears@cloudshell:~ (djs-gcp-2018)$
```

Now let us see what services have been deployed as part of the system:  
Run the following in the shell:

**\$ kubectl get svc**

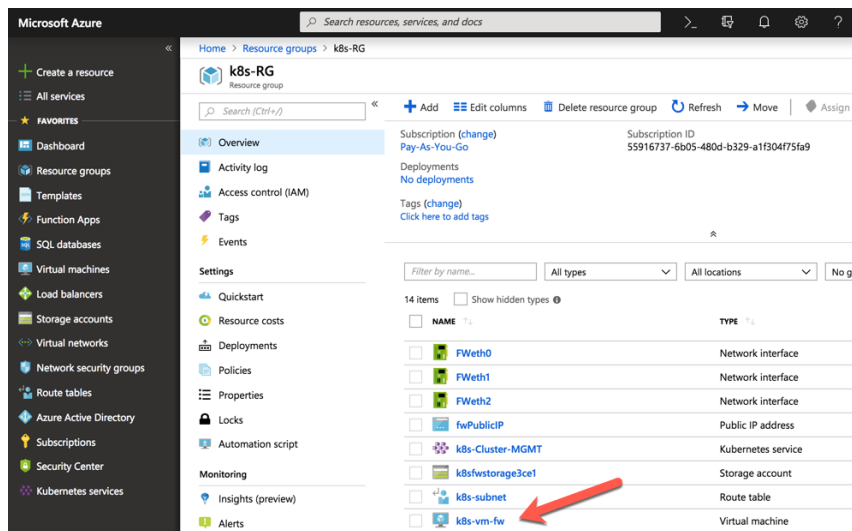
```
3. Shell
SJCMAC3024G8WL:AKS-k8s-north-south-inspection-master dspears$ kubectl get svc
NAME          TYPE          CLUSTER-IP   EXTERNAL-IP   PORT(S)   AGE
kubernetes    ClusterIP    10.21.0.1    <none>        443/TCP   10h
SJCMAC3024G8WL:AKS-k8s-north-south-inspection-master dspears$
```

As you can see no services besides the system cluster have been deployed.

### Task 3 - Log into the firewall

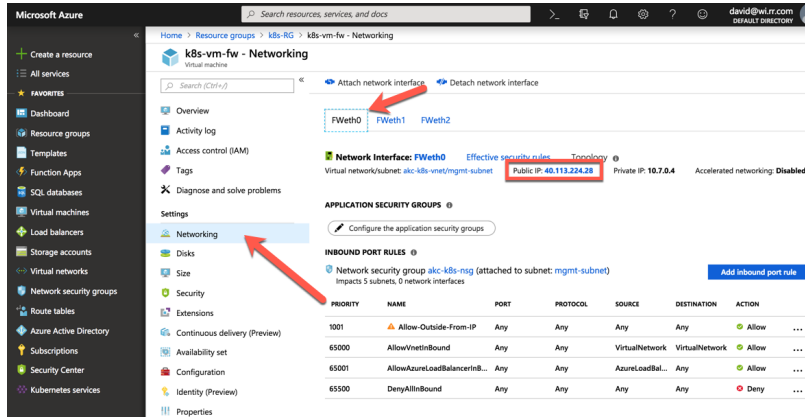
The VM-Series firewall deployed as part of the lab has been bootstrapped. Bootstrapping is a feature of the VM-Series firewall that allows you to load a pre-defined configuration into the firewall during boot-up. This ensures that the firewall is configured and ready at initial boot-up, thereby removing the need for manual configuration. The bootstrapping feature also enables automated deployment of the VM-Series.

Navigate to the k8s-RG resource group and click on the VM-Series firewall Virtual machine:






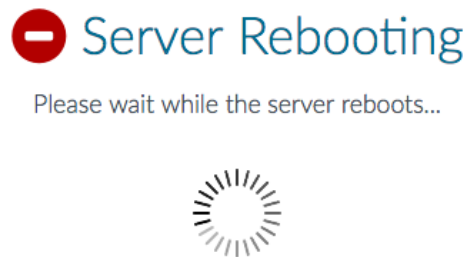
Click on Networking in the left Nav. Copy the Public IP of FWeth0 which is the mgmt interface of the VM-Series firewall:



Open another browser tab and navigate to the firewall management interface:




If you get a security exception, please ignore for this lab and proceed to the firewall login page. The VM-Series firewall by default uses a self-signed certificate which causes the exception. Depending on how quickly you do this, you might see the following message. It is normal and part of the bootup process:



If you wish to SSH into the FW, the following syntax can be used:

**\$ ssh admin@<ip address of firewall>**

**The password is: Pal0Alt0@123      -yes, those are zeros.**



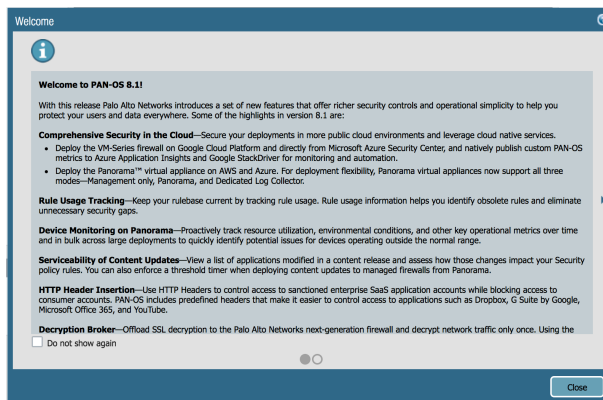
When presented with the login screen you should be able to login to the firewall using (Hint: It's a good idea to jot this password down or save it to a notepad as you will regularly need it):

username: **admin**

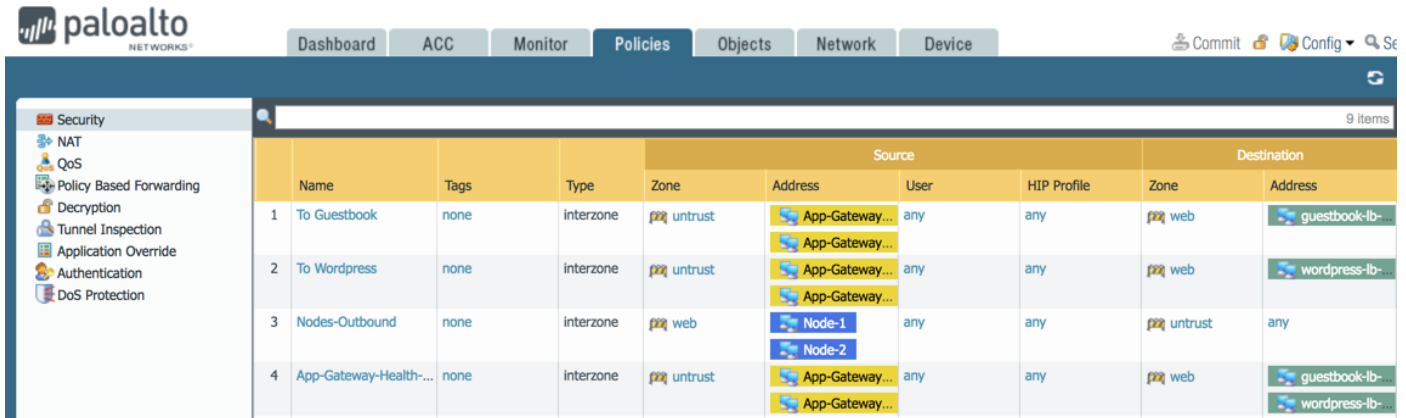
password: **Pal0Alt0@123      -yes, those are zeros.**



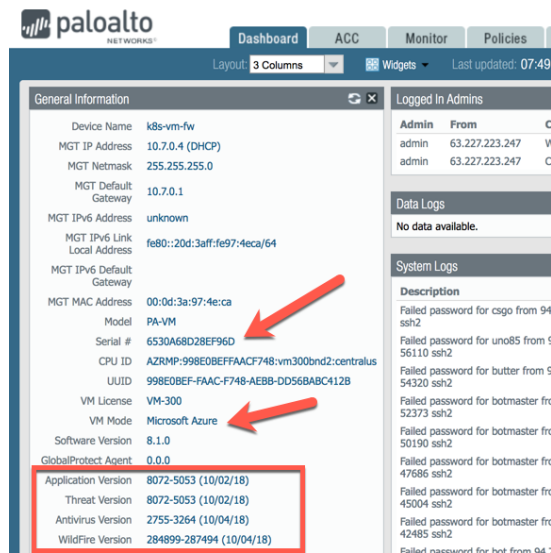
Once logged in you will see a welcome screen, dismiss the welcome dialog box by clicking Close.



Click the Policies tab and you will notice a predefined security policy which was imported using the bootstrapping feature. There are also some predefined NAT policies:



Click on the Dashboard tab, check to verify that the firewall has a serial number. The image defined in the terraform template is a Pay as you Go bundle2. This was used because a license will be required to view the logs later in the lab. If you added content files in the bootstrap folder, you should also see that these have been uploaded.



# Launch a two tiered WordPress application

*In this activity, you will:*

*Optionally: Explore the application's manifest file*

*Launch a two-tier WordPress application within your cluster*

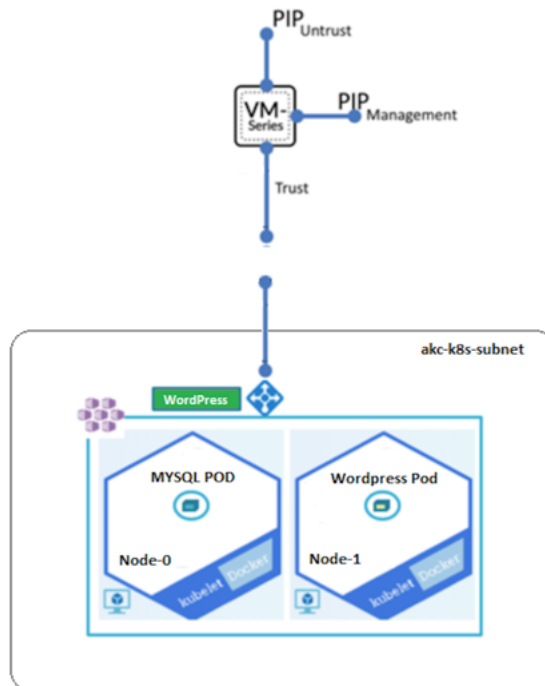
In this activity we will start using Kubernetes specific terms such as Pods, Services, etc. Here is a good primer: <https://kubernetes.io/docs/concepts/workloads/pods/pod-overview/>

## Task 1 – WordPress Application Deployment YAML file

WordPress is a piece of software which has become one of the most widely used content management systems. It is open source, licensed under the GPL, and written in PHP.

WordPress allows users to create and edit websites through a central administrative dashboard, which includes a text editor for modifying content, menus and various design elements. WordPress provides plugins which provide additional functionality through WordPress Plugin Directory. Plugins can be installed through either upload or by one-click installation through the WordPress Plugin Library.

This lab will deploy the following simple WordPress application on the cluster nodes created during the Terraform template deployment:



As you can see this is a two-tiered application with Pods that are dedicated to front-end WordPress services and backend MYSQL DB services.

If interested, the following section dives a bit deeper into the templates being used to create this application. There are two application manifests for this deployment. The first is for the MYSQL DB and the second is for the WordPress frontend. Optionally, open the links below it in a browser of your choice to view the files.

<https://github.com/PaloAltoNetworks/AKS-k8s-north-south-inspection/blob/master/mysql-deployment.yaml>  
and

<https://github.com/PaloAltoNetworks/AKS-k8s-north-south-inspection/blob/master/wordpress-deployment.yaml>

The manifest file declares various aspects of the application. For instance, it tells the orchestrator what type of resources you intend to deploy. In this case we will first deploy a MYSQL DB server and then a WordPress Frontend.

MYSQL Service:

```
apiVersion: v1
kind: Service
metadata:
  name: wordpress-mysql
  labels:
    app: wordpress
spec:
  ports:
    - port: 3306
  selector:
    app: wordpress
    tier: mysql
  clusterIP: None
---
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: mysql-pv-claim
  labels:
    app: wordpress
spec:
  accessModes:
    - ReadWriteOnce
  resources:
    requests:
      storage: 20Gi
---
apiVersion: apps/v1 # for versions before 1.9.0 use apps/v1beta2
kind: Deployment
metadata:
  name: wordpress-mysql
  labels:
    app: wordpress
spec:
  selector:
    matchLabels:
      app: wordpress
      tier: mysql
  strategy:
    type: Recreate
  template:
    metadata:
      labels:
        app: wordpress
        tier: mysql
    spec:
      containers:
        - image: djspears/panw:wordpress-mysql
          name: mysql
          - name: MYSQL_ROOT_PASSWORD
            valueFrom:
              secretKeyRef:
                name: mysql-pass
                key: password
          - containerPort: 3306
            name: mysql
      volumeMounts:
        - name: mysql-persistent-storage
          mountPath: /var/lib/mysql
      volumes:
        - name: mysql-persistent-storage
          persistentVolumeClaim:
            claimName: mysql-pv-claim
```

Some things to notice are the listening port, 3306, the container image, and the credentials that will be used during the deployment.

## Wordpress-Frontend :

```
apiVersion: v1
kind: Service
metadata:
  name: wordpress
  annotations:
    service.beta.kubernetes.io/azure-load-balancer-internal: "true"
  labels:
    app: wordpress
spec:
  ports:
    - port: 80
  selector:
    app: wordpress
    tier: frontend
  type: LoadBalancer

apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: wp-pv-claim
  labels:
    app: wordpress
spec:
  accessModes:
    - ReadWriteOnce
  resources:
    requests:
      storage: 20Gi
---
apiVersion: apps/v1 # for versions before 1.9.0 use apps/v1beta2
kind: Deployment
metadata:
  name: wordpress
  labels:
    app: wordpress
spec:
  replicas: 1
  selector:
    matchLabels:
      app: wordpress
      tier: frontend
  strategy:
    type: Recreate
  template:
    metadata:
      labels:
        app: wordpress
        tier: frontend
    spec:
      containers:
        - image: djspears/panw:wordpress
          name: wordpress
        - name: WORDPRESS_DB_HOST
          value: wordpress-mysql
        - name: WORDPRESS_DB_PASSWORD
          valueFrom:
            secretKeyRef:
              name: mysql-pass
              key: password
      ports:
        - containerPort: 80
          name: wordpress
      volumeMounts:
        - name: wordpress-persistent-storage
          mountPath: /var/www/html
      volumes:
        - name: wordpress-persistent-storage
          persistentVolumeClaim:
            claimName: wp-pv-claim
```

Highlighted in this file are the area that specifies the load balancer service and also the container image. Even though we have two tiers in our application, only one (the frontend service) is exposed to the outside world via a load balancer. The annotation listed above tells AKS and Kubernetes that the load balancer would be of type: Internal.

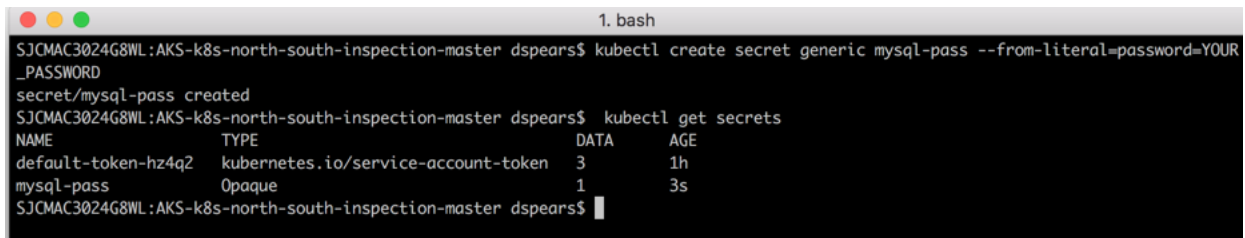
## Task 2 – Launch the Application

As mentioned previously, the application deployment will be done in two steps. The first step will be to deploy the MYSQL DB server. One of the parameters that needs to be passed to the DB server is a root password. To do this securely, the kubectl secrets command will be used. Kubectl secrets are objects intended to hold sensitive information, such as passwords, OAuth tokens, and ssh keys. Putting this information in a secret is safer and more flexible than putting it verbatim in a pod definition or in a docker image. To create a secret, execute the following commands in the terminal window:

```
$ kubectl create secret generic mysql-pass --from-literal=password=YOUR_PASSWORD
```

And the following command will verify that the secrets have been stored

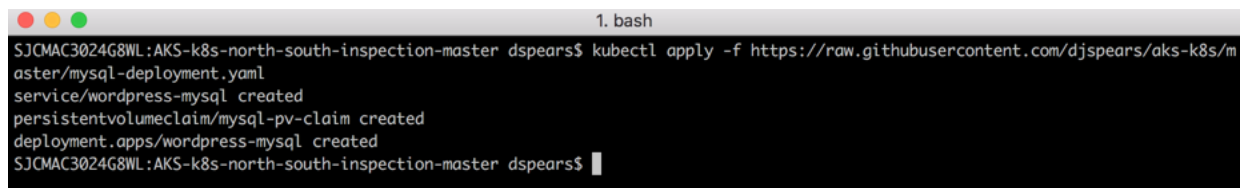
```
$ kubectl get secrets
```



```
1. bash
SJCMA3024G8WL:AKS-k8s-north-south-inspection-master dspears$ kubectl create secret generic mysql-pass --from-literal=password=YOUR_PASSWORD
secret/mysql-pass created
SJCMA3024G8WL:AKS-k8s-north-south-inspection-master dspears$ kubectl get secrets
NAME                                TYPE                                DATA    AGE
default-token-hz4q2                 kubernetes.io/service-account-token 3        1h
mysql-pass                           Opaque                               1        3s
SJCMA3024G8WL:AKS-k8s-north-south-inspection-master dspears$
```

Now the MYSQL pod can be deployed. To do this, execute the following command:

```
$ kubectl apply -f https://raw.githubusercontent.com/PaloAltoNetworks/AKS-k8s-north-south-inspection/master/mysql-deployment.yaml
```



```
1. bash
SJCMA3024G8WL:AKS-k8s-north-south-inspection-master dspears$ kubectl apply -f https://raw.githubusercontent.com/djspears/aks-k8s/master/mysql-deployment.yaml
service/wordpress-mysql created
persistentvolumeclaim/mysql-pv-claim created
deployment.apps/wordpress-mysql created
SJCMA3024G8WL:AKS-k8s-north-south-inspection-master dspears$
```

You should see the services and deployments being created. Next, validate the new pods in your cluster have been created. In your terminal execute:

```
$ kubectl get pods -o wide
```

You may see the status as Pending or ContainerCreating. This is usually a normal situation:

```
1. bash
SJCMAC3024G8WL:AKS-k8s-north-south-inspection-master dspears$ kubectl get pods -o wide
NAME                                READY    STATUS    RESTARTS  AGE    IP          NODE
wordpress-mysql-795bf5f54c-jqbgk    0/1     Pending  0          15s   <none>     <none>
SJCMAC3024G8WL:AKS-k8s-north-south-inspection-master dspears$
```

```
SJCMAC3024G8WL:AKS-k8s-north-south-inspection-master dspears$ kubectl get pods -o wide
NAME                                READY    STATUS    RESTARTS  AGE    IP          NODE
wordpress-mysql-795bf5f54c-jqbgk    0/1     ContainerCreating  0          1m    <none>     aks-default-56371607-0
SJCMAC3024G8WL:AKS-k8s-north-south-inspection-master dspears$
```

By executing the **kubectl get pods -o wide** again, you start seeing that the Ready and Status of pods change as they start up. Verify that the pod gets to a running status.

```
1. bash
SJCMAC3024G8WL:AKS-k8s-north-south-inspection-master dspears$ kubectl get pods -o wide
NAME                                READY    STATUS    RESTARTS  AGE    IP          NODE
wordpress-mysql-795bf5f54c-jqbgk    1/1     Running   0          2m    10.7.10.48  aks-default-56371607-0
SJCMAC3024G8WL:AKS-k8s-north-south-inspection-master dspears$
```

With the MYSQL DB Running, create the WordPress frontend by executing the following command:

**\$ kubectl apply -f <https://raw.githubusercontent.com/PaloAltoNetworks/AKS-k8s-north-south-inspection/master/wordpress-deployment.yaml>**

```
1. bash
SJCMAC3024G8WL:AKS-k8s-north-south-inspection-master dspears$ kubectl apply -f https://raw.githubusercontent.com/djspears/aks-k8s/master/wordpress-deployment.yaml
service/wordpress created
persistentvolumeclaim/wp-pv-claim created
deployment.apps/wordpress created
```

Next, validate the new pods in your cluster have been created. In your terminal execute:

**\$ kubectl get pods -o wide**

Again, you may see the status as Pending or ContainerCreating. This is usually a normal situation:

```
1. bash
SJCMAC3024G8WL:AKS-k8s-north-south-inspection-master dspears$ kubectl get pods -o wide
NAME                                READY    STATUS    RESTARTS  AGE    IP          NODE
wordpress-8574f9c6f9-hm4cw          0/1     Pending  0          5s    <none>     <none>
wordpress-mysql-795bf5f54c-jqbgk    1/1     Running   0          10m   10.7.10.48  aks-default-56371607-0
```

```
1. bash
SJCMAC3024G8WL:AKS-k8s-north-south-inspection-master dspears$ kubectl get pods -o wide
NAME                                READY    STATUS    RESTARTS  AGE    IP          NODE
wordpress-8574f9c6f9-hm4cw          0/1     ContainerCreating  0          1m    <none>     aks-default-56371607-1
wordpress-mysql-795bf5f54c-jqbgk    1/1     Running   0          11m   10.7.10.48  aks-default-56371607-0
```



Again, verify that the pod gets to a running status.

```
1. bash
SJCMA3024G8WL:AKS-k8s-north-south-inspection-master dspears$ kubectl get pods -o wide
NAME                                READY   STATUS    RESTARTS   AGE   IP              NODE                                ...
wordpress-8574f9c6f9-hm4cw          1/1    Running   0           5m   10.7.10.31     aks-default-56371607-1             ...
wordpress-mysql-795bf5f54c-jqbgk    1/1    Running   0           15m  10.7.10.48     aks-default-56371607-0             ...
SJCMA3024G8WL:AKS-k8s-north-south-inspection-master dspears$
```

# Launch a two tiered Guestbook application

*In this activity, you will:*

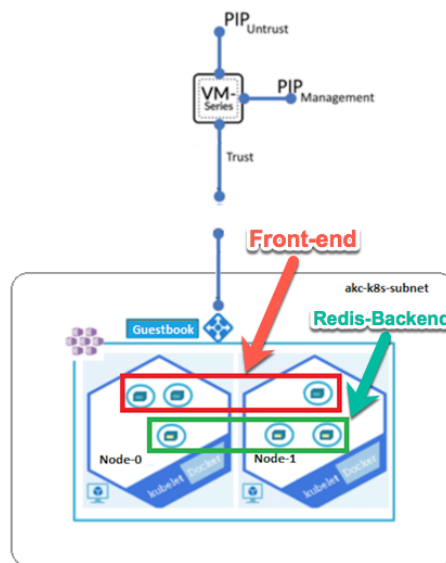
*Optionally: Explore the application's manifest file*

*Launch a two-tier WordPress application within your cluster*

## Task 1 – Guestbook Application Deployment YAML file

Guestbooks have been used by businesses for many years as a way to connect with customers and obtain contact information for future events and promotions. Today, businesses such as popular retail stores, 5-star hotels and even small family-owned B & B's are turning to iPad guestbook apps to help them gather information and enhance the customer's "in-biz" experience. Acquiring email addresses and a **social media** following is a crucial part of any marketing plan. With much of the population using computers on a daily basis, an email marketing plan is of the utmost importance. Using a guest book app in your store makes collecting email addresses a snap and offers enticing features with which the traditional paper and pen guestbook just can't compete. The guestbook application we will build and secure today could be used for Hotel website visits, shopping sites or any other business that wants to keep track of their customer and provide them with promotions or advertisements.

This lab will deploy the following simple Guestbook application on the cluster nodes created during the Terraform template deployment:



As you can see this is a two-tiered application with Pods that are dedicated to front-end web services and backend DB services.

If interested, the following section dives a bit deeper into the templates being used to create this application. This is a link to the application manifest. Optionally, click the link below and open it in a browser of your choice.

<https://github.com/PaloAltoNetworks/AKS-k8s-north-south-inspection/blob/master/guestbook-all-in-one.yaml>

The manifest file in this case we will deploy a 2-tier simple redis application with a fronted and backend tier. The backend tier will consist of a redis-master and redis-slave for db redundancy. Front-end Service:

```
90 apiVersion: v1
91 kind: Service
92 metadata:
93   name: frontend
94   annotations:
95     service.beta.kubernetes.io/azure-load-balancer-internal: "true"
96   labels:
97     app: guestbook
98     tier: frontend
99 spec:
100   # if your cluster supports it, uncomment the following to automatically create
101   # an external load-balanced IP for the frontend service.
102   type: LoadBalancer
103   ports:
104     - port: 80
105   selector:
106     app: guestbook
107     tier: frontend
108 ---
109 apiVersion: extensions/v1beta1
110 kind: Deployment
111 metadata:
112   name: frontend
113 spec:
114   replicas: 3
115   template:
116     metadata:
117       labels:
118         app: guestbook
119         tier: frontend
120     spec:
121       containers:
122         - name: php-redis
123           image: gcr.io/google-samples/gb-frontend:v4
124           resources:
125             requests:
126               cpu: 100m
127               memory: 100Mi
128           env:
129             - name: GET_HOSTS_FROM
130               value: dns
131             # If your cluster config does not include a dns service, then to
132             # instead access environment variables to find service host
133             # info, comment out the 'value: dns' line above, and uncomment the
134             # line below:
135             # value: env
136           ports:
137             - containerPort: 80
```

This tells the orchestrator that the service will be exposed via an internal load balancer

## Redis-backend-master :

```
1  apiVersion: v1
2  kind: Service
3  metadata:
4    name: redis-master
5    labels:
6      app: redis
7      tier: backend
8      role: master
9  spec:
10 #type: LoadBalancer
11 ports:
12 - port: 6379
13   targetPort: 6379
14 selector:
15   app: redis
16   tier: backend
17   role: master
18 ---
19 apiVersion: extensions/v1beta1
20 kind: Deployment
21 metadata:
22   name: redis-master
23 spec:
24   replicas: 1
25   template:
26     metadata:
27       labels:
28         app: redis
29         role: master
30         tier: backend
31     spec:
32       containers:
33         - name: master
34           image: gcr.io/google_containers/redis:e2e # or just image: redis
35           resources:
36             requests:
37               cpu: 100m
38               memory: 100Mi
39           ports:
40             - containerPort: 6379
```

## Redis-backend-slave:

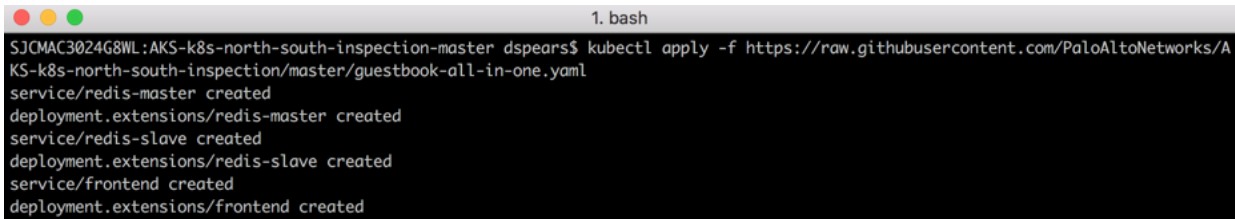
```
42  apiVersion: v1
43  kind: Service
44  metadata:
45    name: redis-slave
46    labels:
47      app: redis
48      tier: backend
49      role: slave
50 spec:
51 #type: LoadBalancer
52 ports:
53 - port: 6379
54 selector:
55   app: redis
56   tier: backend
57   role: slave
58 ---
59 apiVersion: extensions/v1beta1
60 kind: Deployment
61 metadata:
62   name: redis-slave
63 spec:
64   replicas: 2
65   template:
66     metadata:
67       labels:
68         app: redis
69         role: slave
70         tier: backend
71     spec:
72       containers:
73         - name: slave
74           image: gcr.io/google_samples/gb-redisslave:v1
75           resources:
76             requests:
77               cpu: 100m
78               memory: 100Mi
79           env:
80             - name: GET_HOSTS_FROM
81               value: dns
82             # If your cluster config does not include a dns service, then to
83             # instead access an environment variable to find the master
84             # service's host, comment out the 'value: dns' line above, and
85             # uncomment the line below:
86             # value: env
87           ports:
88             - containerPort: 6379
```

Even though there are two tiers in the application, only one (the frontend service) is exposed to the outside world via a load balancer. The annotation listed above tells GCP and Kubernetes that the load balancer would be of type: Internal.

## Task 2 – Launch the Application

Back in terminal shell type the following command to deploy the application pods:

```
$ kubectl apply -f https://raw.githubusercontent.com/PaloAltoNetworks/AKS-k8s-north-south-inspection/master/guestbook-all-in-one.yaml
```

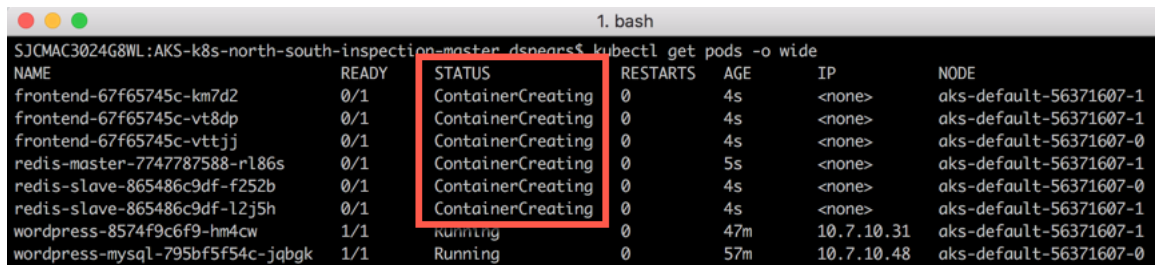


```
1. bash
SJCMAC3024G8WL:AKS-k8s-north-south-inspection-master dspears$ kubectl apply -f https://raw.githubusercontent.com/PaloAltoNetworks/AKS-k8s-north-south-inspection/master/guestbook-all-in-one.yaml
service/redis-master created
deployment.extensions/redis-master created
service/redis-slave created
deployment.extensions/redis-slave created
service/frontend created
deployment.extensions/frontend created
```

You should see the services and deployments being created. Next, validate the new pods in your cluster have been created. In your terminal execute:

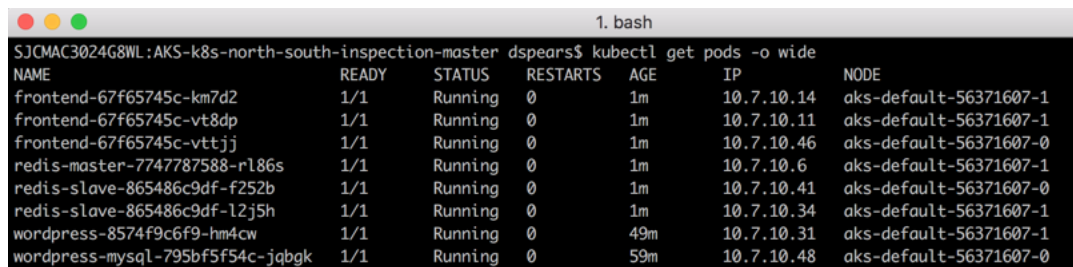
```
$ kubectl get pods -o wide
```

You may see the status as Pending or ContainerCreating. This is usually a normal situation:



```
1. bash
SJCMAC3024G8WL:AKS-k8s-north-south-inspection-master dspears$ kubectl get pods -o wide
NAME                                READY   STATUS             RESTARTS   AGE   IP              NODE
frontend-67f65745c-km7d2            0/1    ContainerCreating  0         4s   <none>         aks-default-56371607-1
frontend-67f65745c-vt8dp            0/1    ContainerCreating  0         4s   <none>         aks-default-56371607-1
frontend-67f65745c-vttjj            0/1    ContainerCreating  0         4s   <none>         aks-default-56371607-0
redis-master-7747787588-r186s      0/1    ContainerCreating  0         5s   <none>         aks-default-56371607-1
redis-slave-865486c9df-f252b       0/1    ContainerCreating  0         4s   <none>         aks-default-56371607-0
redis-slave-865486c9df-l2j5h       0/1    ContainerCreating  0         4s   <none>         aks-default-56371607-1
wordpress-8574f9c6f9-hm4cw         1/1    Running            0         47m   10.7.10.31    aks-default-56371607-1
wordpress-mysql-795bf5f54c-jqbgk   1/1    Running            0         57m   10.7.10.48    aks-default-56371607-0
```

By executing the **kubectl get pods -o wide** again, you start seeing that the Ready and Status of pods change as they start up. Verify that the pods gets to a running status.



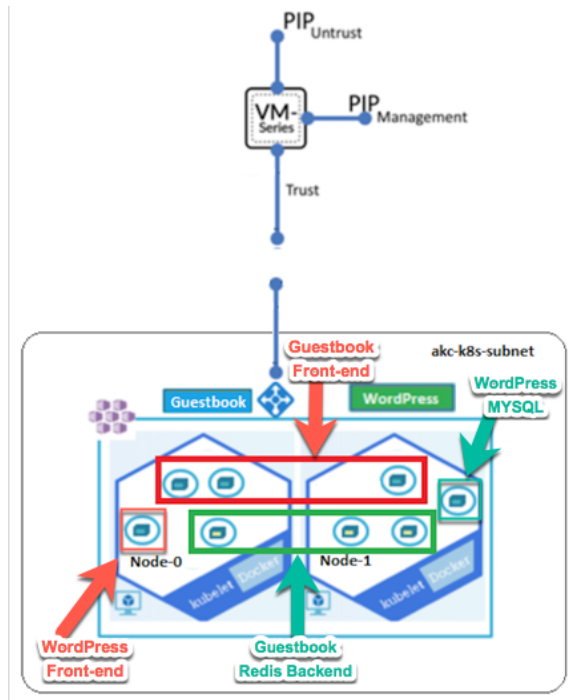
```
1. bash
SJCMAC3024G8WL:AKS-k8s-north-south-inspection-master dspears$ kubectl get pods -o wide
NAME                                READY   STATUS    RESTARTS   AGE   IP              NODE
frontend-67f65745c-km7d2            1/1    Running   0           1m   10.7.10.14     aks-default-56371607-1
frontend-67f65745c-vt8dp            1/1    Running   0           1m   10.7.10.11     aks-default-56371607-1
frontend-67f65745c-vttjj            1/1    Running   0           1m   10.7.10.46     aks-default-56371607-0
redis-master-7747787588-r186s      1/1    Running   0           1m   10.7.10.6      aks-default-56371607-1
redis-slave-865486c9df-f252b       1/1    Running   0           1m   10.7.10.41     aks-default-56371607-0
redis-slave-865486c9df-l2j5h       1/1    Running   0           1m   10.7.10.34     aks-default-56371607-1
wordpress-8574f9c6f9-hm4cw         1/1    Running   0           49m   10.7.10.31     aks-default-56371607-1
wordpress-mysql-795bf5f54c-jqbgk   1/1    Running   0           59m   10.7.10.48     aks-default-56371607-0
```

# Explore the newly deployed applications

In this activity, you will:

Explore aspects of the application deployments

The following diagram shows what has been instantiated:



Let's validate this by listing the new pods in your cluster. In your terminal window execute:

```
$ kubectl get pods -o wide
```

You should see the pods for both the WordPress and Guestbook Application:

```
SJCMAC3024G8WL:AKS-k8s-north-south-inspection-master dspears$ kubectl get pods -o wide
```

NAME	READY	STATUS	RESTARTS	AGE	IP	NODE
frontend-67f65745c-km7d2	1/1	Running	0	1m	10.7.10.14	aks-default-56371607-1
frontend-67f65745c-vt8dp	1/1	Running	0	1m	10.7.10.11	aks-default-56371607-1
frontend-67f65745c-vttjj	1/1	Running	0	1m	10.7.10.46	aks-default-56371607-0
redis-master-7747787588-r186s	1/1	Running	0	1m	10.7.10.6	aks-default-56371607-1
redis-slave-865486c9df-f252b	1/1	Running	0	1m	10.7.10.41	aks-default-56371607-0
redis-slave-865486c9df-12i5h	1/1	Running	0	1m	10.7.10.34	aks-default-56371607-1
wordpress-8574f9c6f9-hm4cw	1/1	Running	0	49m	10.7.10.31	aks-default-56371607-1
wordpress-mysql-795bf5f54c-jqbgk	1/1	Running	0	59m	10.7.10.48	aks-default-56371607-0

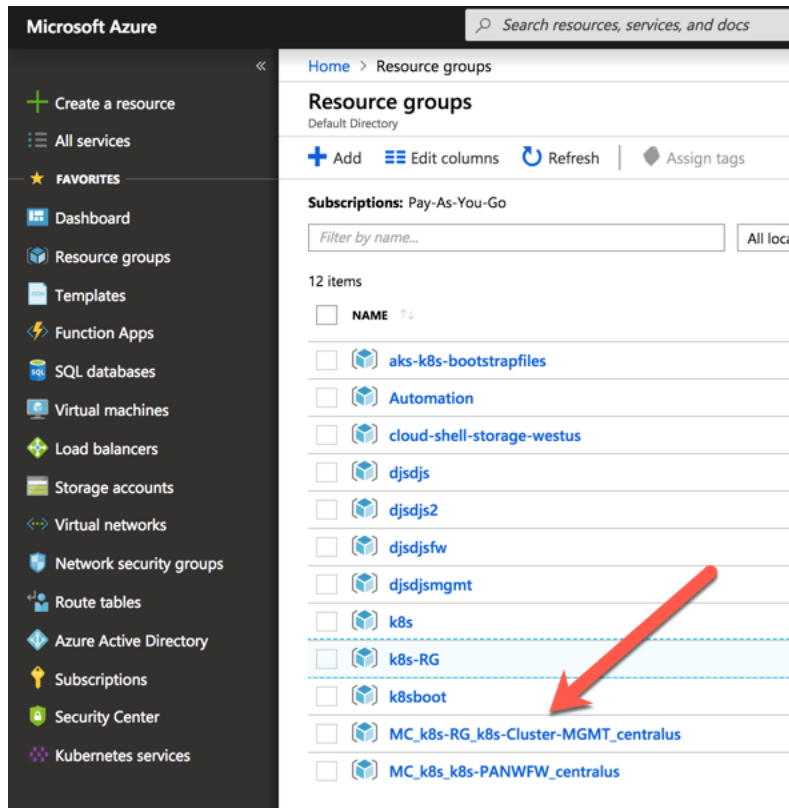
Next let's look at the load balancing service for the front-end pod. Execute the following command in the shell:

### \$ kubectl get svc

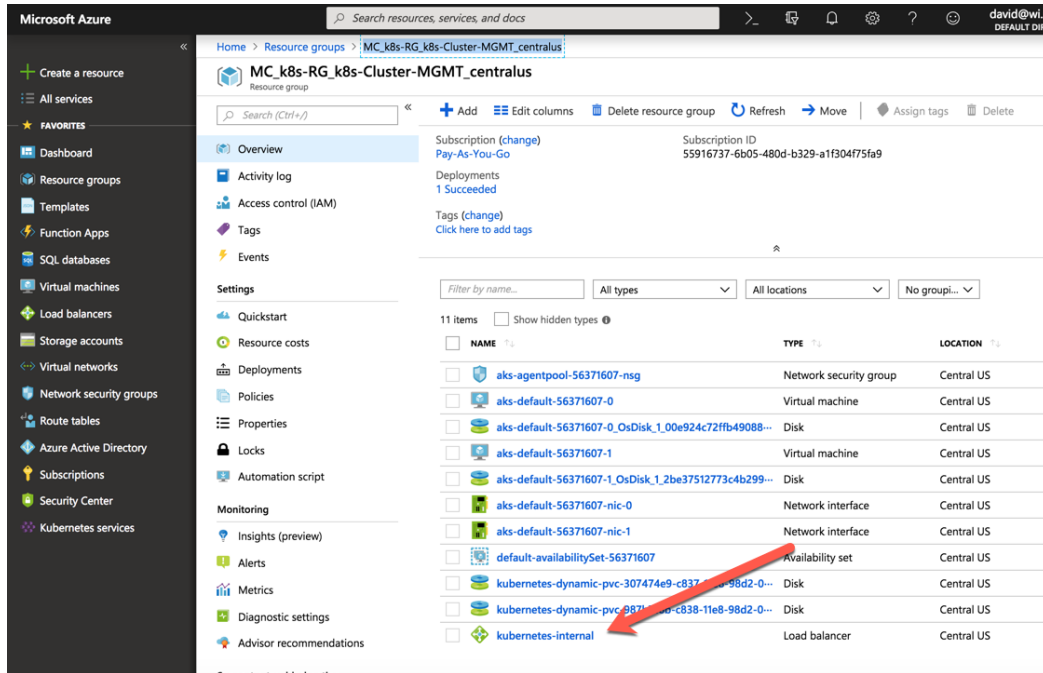
You can see there is a load balancer External IP for both the frontend Guestbook application and an External IP address for the WordPress server. Note that the IP address is in the 10.7.10.0/24 subnet. This is one of the subnets that was deployed in the Azure VNET during the Terraform execution.

```
SJCMAC3024G8WL:AKS-k8s-north-south-inspection-master dspears$ kubectl get svc
NAME                TYPE          CLUSTER_IP      EXTERNAL_IP      PORT(S)          AGE
frontend            LoadBalancer  10.21.151.190   10.7.10.67       80:32498/TCP     1h
kubernetes           ClusterIP     10.21.0.1       <none>           443/TCP          4h
redis-master         ClusterIP     10.21.51.117    <none>           6379/TCP         1h
redis-slave          ClusterIP     10.21.58.87     <none>           6379/TCP         1h
wordpress            LoadBalancer  10.21.165.47    10.7.10.66       80:30231/TCP     2h
wordpress-mysql      ClusterIP     none            <none>           3306/TCP         2h
SJCMAC3024G8WL:AKS-k8s-north-south-inspection-master dspears$
```

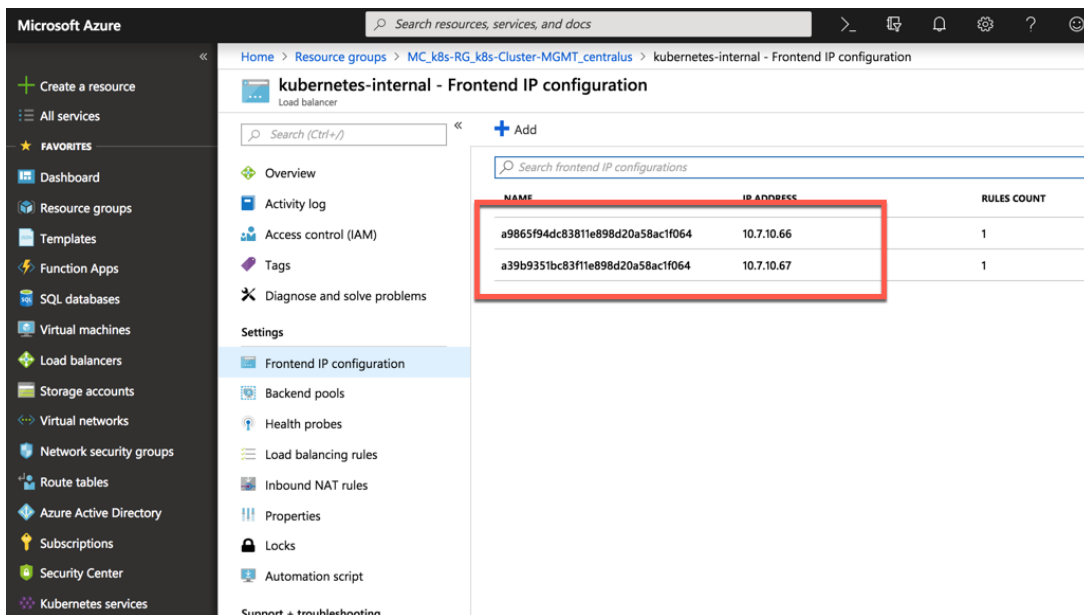
These load balancer IP addresses can be seen via the Azure Dashboard as well. Navigate to the Resource Groups and click on the “MC\_k8s-RG\_k8s-Cluster-MGMT\_centralus” Resource group. This group was created automatically for the k8s node resources.



Click on the Kubernetes-internal Load balancer:



Click on the Frontend IP configuration on the left Nav. The application load balancer IP ADDRESS are displayed:



# Securing Inbound Traffic

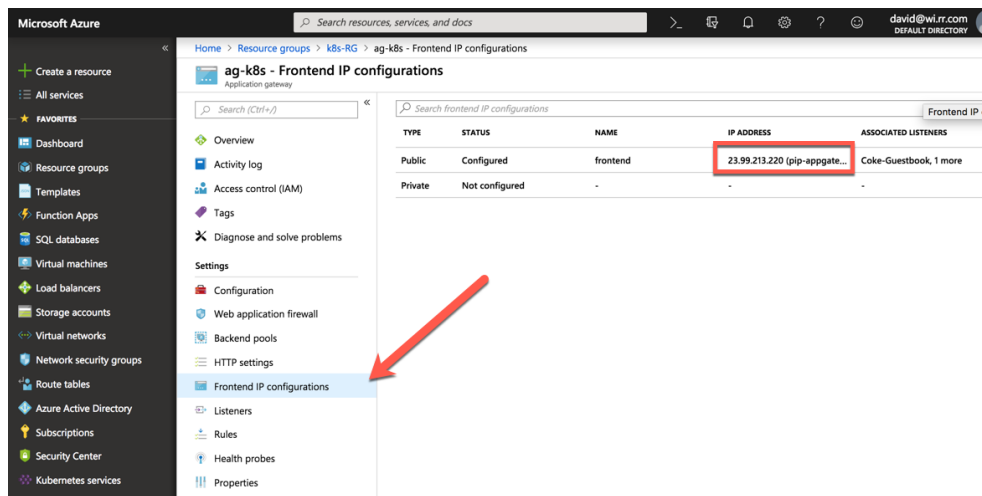
In this activity, you will:

Secure traffic that is inbound to your frontend services

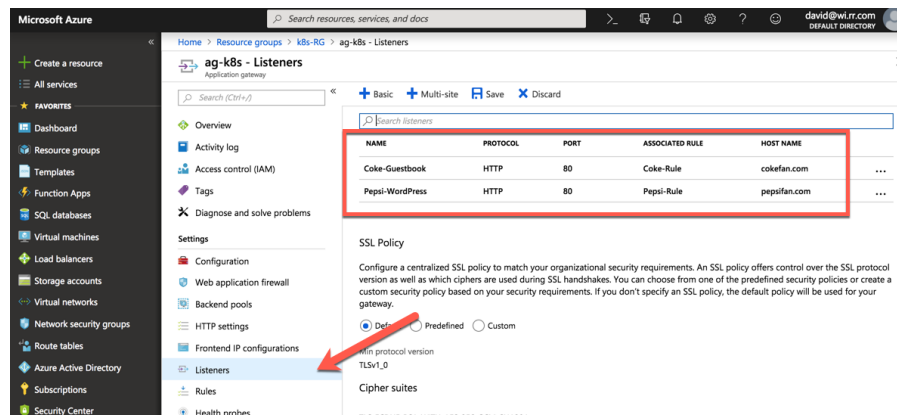
Validate that traffic is visible in the Firewall logs

## Task 1 – Azure Application Gateway IP Address

This Terraform deployment created an Azure Application gateway in front of the VM-Series firewall. As previously discussed, the Application Gateway is configured to do host header redirection. In order for this to function the frontend IP addresses must be identified and a few hosts entries need to be made on the testing machine. Open the Application Gateway Frontend IP configurations in the Resource groups > k8s-RG > ag-k8s blade:

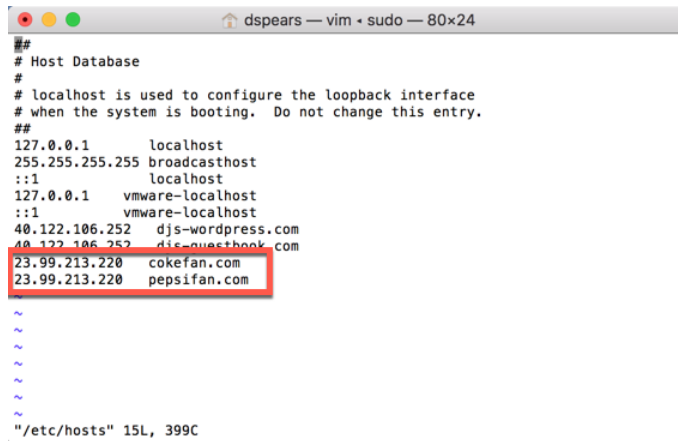


Copy this address as it will be needed to create a DNS entry in the local host file. Go to Application Gateway Listeners on the left Nav to see the DNS entries that the Application Gateway is configured to serve.





Open the local hosts file and create a pepsifan.com and cokefan.com entry. Each entry will have the IP address of the Application Gateway Frontend IP address:

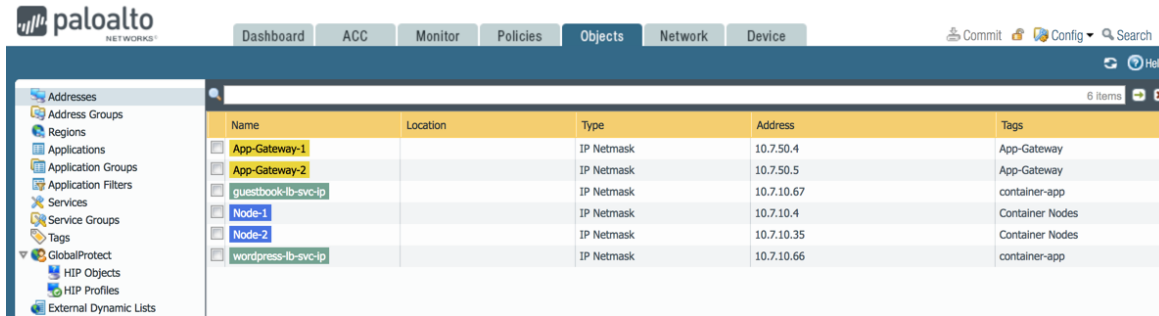


```
##
# Host Database
# localhost is used to configure the loopback interface
# when the system is booting. Do not change this entry.
##
127.0.0.1    localhost
255.255.255.255 broadcasthost
::1        localhost
127.0.0.1    vmware-localhost
::1        vmware-localhost
40.122.106.252 djs-wordpress.com
40.122.106.252 djs-guestbook.com
23.99.213.220 cokefan.com
23.99.213.220 pepsifan.com
~
~
~
~
~
~
~/etc/hosts" 15L, 399C
```

## Task 2 – Update the Firewall’s Address Objects

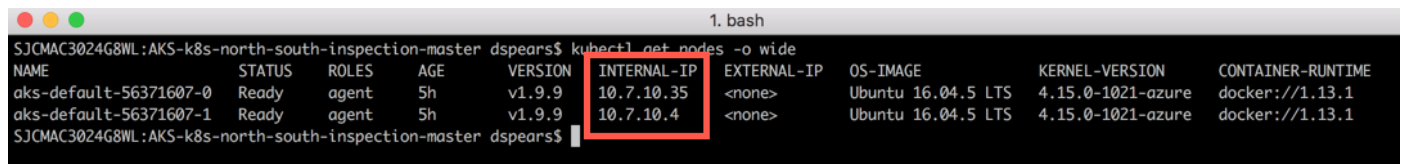
Open the VM-Series firewall. This design is not using any NATs for the inbound traffic flow. The bootstrapped configuration should have the correct addressing but this task will validate that.

Click the Objects Tab and navigate to “Addresses” on the left. The Addresses used in the policy are defined here:



Open the terminal window and check that the Address objects are correct. Execute the following command to verify the nodes:

```
$ kubectl get nodes -o wide
```

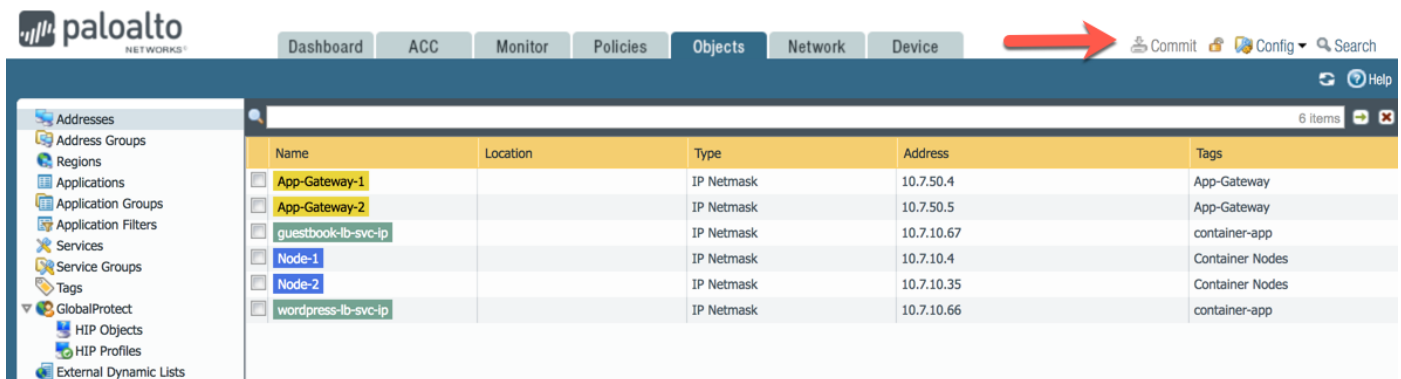


```
1. bash
SJCMA3024G8WL:AKS-k8s-north-south-inspection-master dspears$ kubectl get nodes -o wide
NAME                                STATUS    ROLES    AGE     VERSION    INTERNAL-IP    EXTERNAL-IP    OS-IMAGE           KERNEL-VERSION    CONTAINER-RUNTIME
aks-default-56371607-0              Ready    agent    5h     v1.9.9     10.7.10.35     <none>         Ubuntu 16.04.5 LTS 4.15.0-1021-azure docker://1.13.1
aks-default-56371607-1              Ready    agent    5h     v1.9.9     10.7.10.4      <none>         Ubuntu 16.04.5 LTS 4.15.0-1021-azure docker://1.13.1
SJCMA3024G8WL:AKS-k8s-north-south-inspection-master dspears$
```

Next enter the “**kubectl get svc**” command to verify the lb-svc-ip’s:

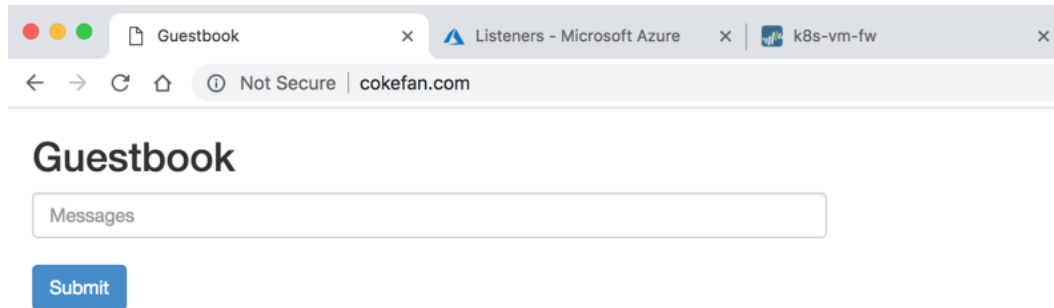
```
SJCMAC3024G8WL:AKS-k8s-north-south-inspection-master dspears$ kubectl get svc
NAME                TYPE          CLUSTER-IP    EXTERNAL-IP    PORT(S)          AGE
frontend            LoadBalancer  10.21.151.190  10.7.10.67     80:32498/TCP    2h
kubernetes           ClusterIP     10.21.0.1      <none>         443/TCP         5h
redis-master         ClusterIP     10.21.51.117   <none>         6379/TCP        2h
redis-slave          ClusterIP     10.21.58.87    <none>         6379/TCP        2h
wordpress            LoadBalancer  10.21.165.47   10.7.10.66     80:30231/TCP    3h
wordpress-mysql      ClusterIP     None           <none>         3306/TCP        3h
SJCMAC3024G8WL:AKS-k8s-north-south-inspection-master dspears$
```

If a change is needed, make the changes and click the commit link on the top right

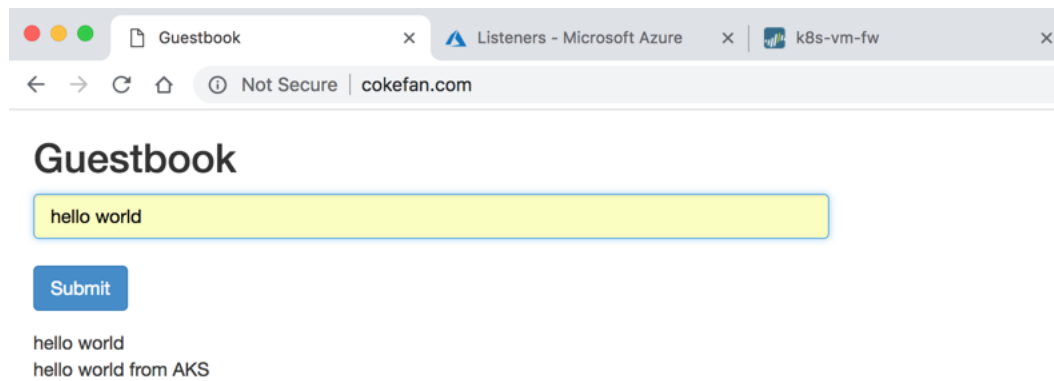


## Task 3 – Connect to the Guestbook Frontend

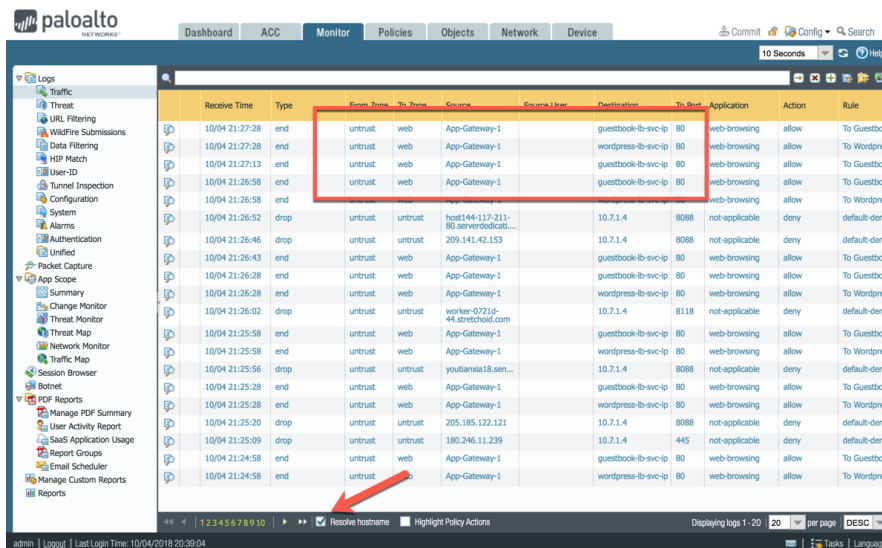
The VM-Series is now protecting your Kubernetes workload. In order to connect to the guestbook's frontend service, you will open a browser and navigate to the <http://cokefan.com> website:



Enter something in the Messages box and click submit. The messages should be echoed below:

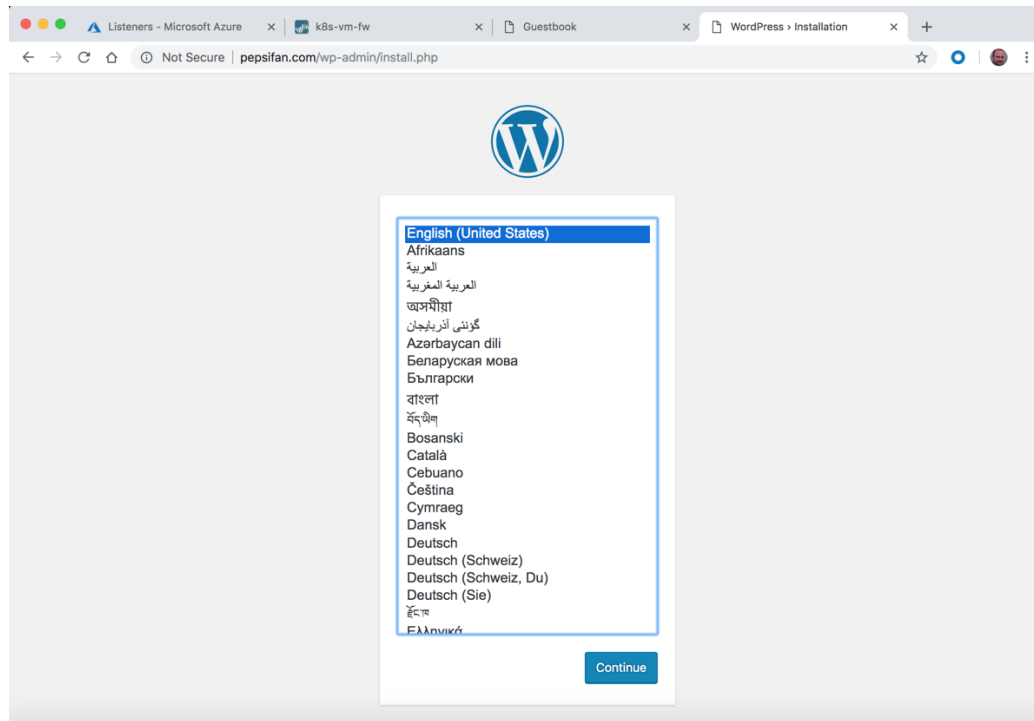


Open the VM-Series firewall monitor tab and validate that traffic is flowing through the firewall:



**ProTip:** Tick the Resolve hostname to make the logs more readable.

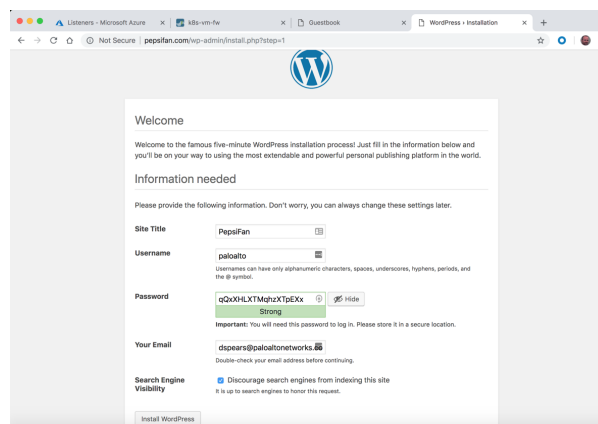
Now check that the <http://pepsifan.com> site works. Open a new tab and open the pepsifan.com site:



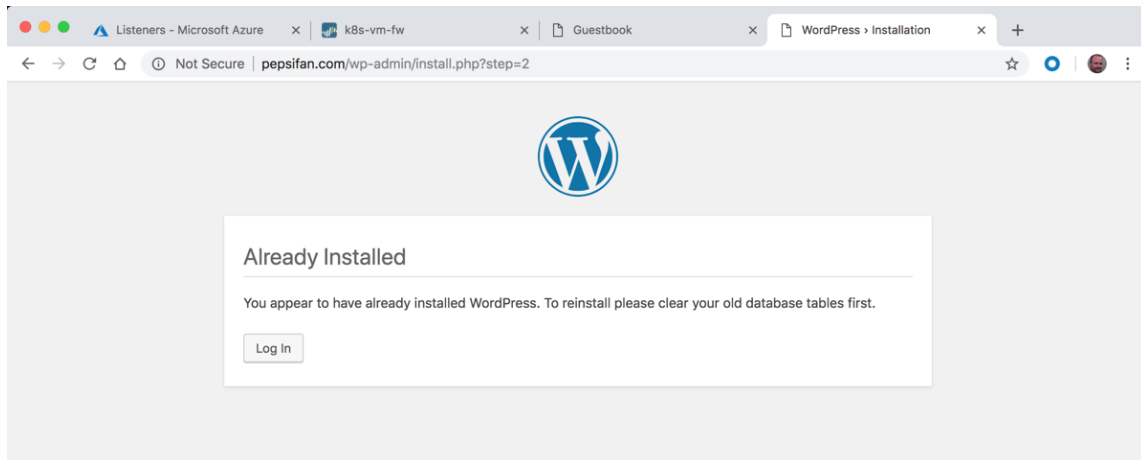
You may see the following error message. This is usually because the WordPress website takes a little time to get up and running. Click refresh a few times to get to the next step:



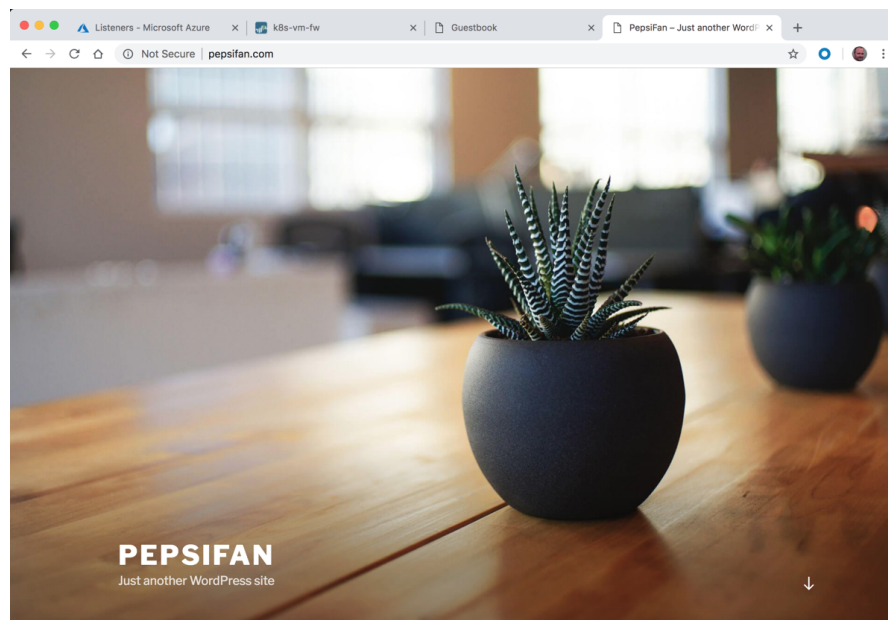
You should see the WordPress install page. Click Continue if you wish to go through the installation process:



After pressing Install WordPress, you might see a 502 error from the Application Gateway. If you press refresh a few times you should see the following:



Go to the root of the <http://pepsifan.com> site and you should now see the default theme:



Verify that the pepsifan.com traffic is running through the firewall:

The screenshot shows the Palo Alto Networks traffic logs interface. The left sidebar contains a navigation menu with categories like Traffic, Threat, URL Filtering, and more. The main area displays a table of traffic logs. The table has columns for Receive Time, Type, From Zone, To Zone, Source, Source User, Destination, To Port, Application, Action, and Rule. The logs show traffic from the 'untrust' zone to the 'web' zone, primarily from 'App-Gateway-1'. The application is identified as 'web-browsing' and the action is 'allow'. The rule applied is 'To Guestbook'.

Receive Time	Type	From Zone	To Zone	Source	Source User	Destination	To Port	Application	Action	Rule
10/04 21:45:58	end	untrust	web	App-Gateway-1		guestbook-fb-svc-ip	80	web-browsing	allow	To Guestbook
10/04 21:45:53	end	untrust	web	App-Gateway-1		wordpress-fb-svc-ip	80	web-browsing	allow	To Wordpress
10/04 21:45:28	end	untrust	web	App-Gateway-1		guestbook-fb-svc-ip	80	web-browsing	allow	To Guestbook
10/04 21:45:27	drop	untrust	untrust	5.188.206.248		10.71.4	3389	not-applicable	deny	default-deny
10/04 21:45:25	drop	untrust	untrust	205.185.113.156		10.71.4	8088	not-applicable	deny	default-deny
10/04 21:45:23	end	untrust	web	App-Gateway-1		wordpress-fb-svc-ip	80	web-browsing	allow	To Wordpress
10/04 21:45:08	drop	untrust	untrust	205.185.122.121		10.71.4	8088	not-applicable	deny	default-deny
10/04 21:44:58	end	untrust	web	App-Gateway-1		guestbook-fb-svc-ip	80	web-browsing	allow	To Guestbook
10/04 21:44:53	end	untrust	web	App-Gateway-1		wordpress-fb-svc-ip	80	web-browsing	allow	To Wordpress
10/04 21:44:28	end	untrust	web	App-Gateway-1		wordpress-fb-svc-ip	80	web-browsing	allow	To Guestbook
10/04 21:44:23	end	untrust	web	App-Gateway-1		wordpress-fb-svc-ip	80	web-browsing	allow	To Wordpress
10/04 21:44:19	drop	untrust	untrust	31.184.237.140		10.71.4	3491	not-applicable	deny	default-deny
10/04 21:43:58	end	untrust	web	App-Gateway-1		guestbook-fb-svc-ip	80	web-browsing	allow	To Guestbook
10/04 21:43:53	drop	untrust	untrust	host190-246-177-94.static.arabiacd...		10.71.4	8088	not-applicable	deny	default-deny
10/04 21:43:53	end	untrust	web	App-Gateway-1		wordpress-fb-svc-ip	80	web-browsing	allow	To Wordpress
10/04 21:43:52	drop	untrust	untrust	209.141.42.153		10.71.4	8088	not-applicable	deny	default-deny
10/04 21:43:28	end	untrust	web	App-Gateway-1		guestbook-fb-svc-ip	80	web-browsing	allow	To Guestbook
10/04 21:43:26	drop	untrust	untrust	5.188.206.14		10.71.4	11504	not-applicable	deny	default-deny
10/04 21:43:25	drop	untrust	untrust	180.247.43.104		10.71.4	8000	not-applicable	deny	default-deny
10/04 21:43:23	end	untrust	web	App-Gateway-1		wordpress-fb-svc-ip	80	web-browsing	allow	To Wordpress

# Securing Outbound Traffic

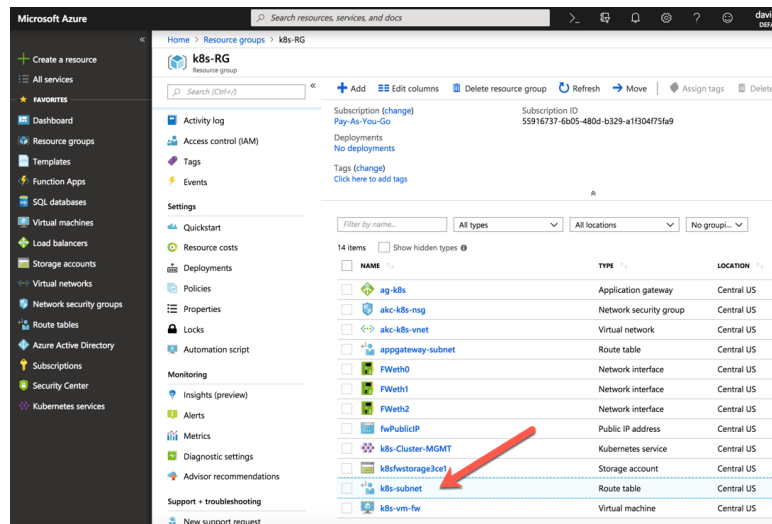
In this activity, you will:

Secure outbound traffic from the cluster nodes

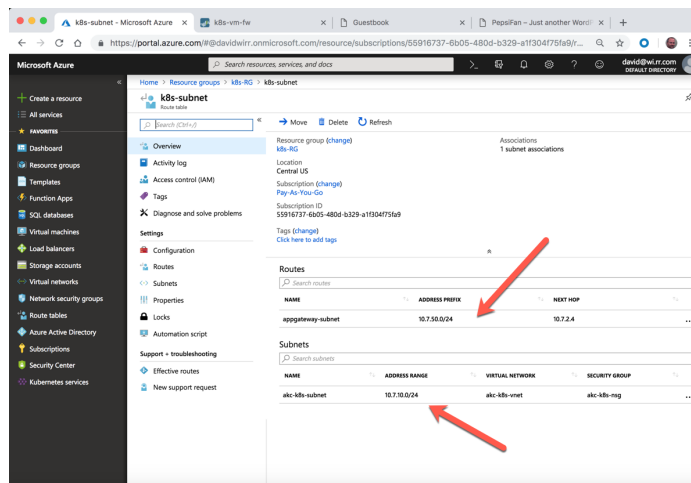
Validate traffic is in the Firewall logs

## Task 1 – Add Outbound Route

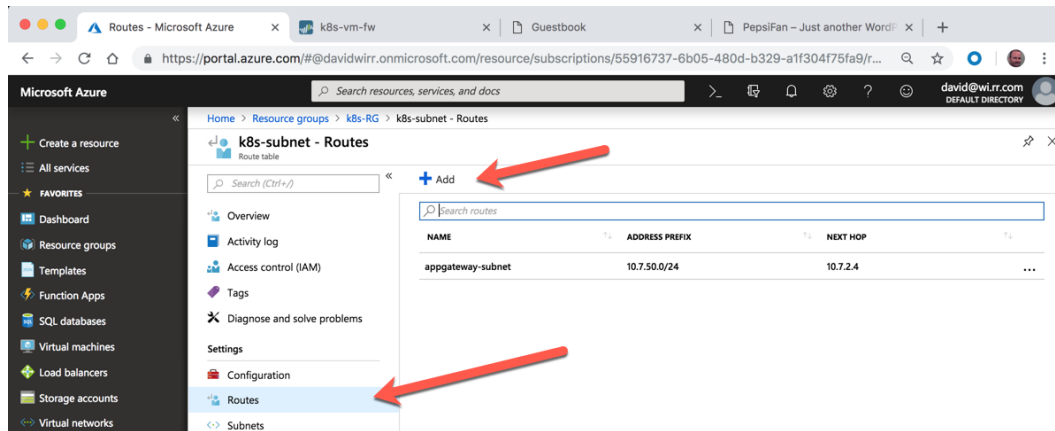
To secure any traffic that is originating from within the cluster we need to add a user defined route (UDR) to the routing table on the VNET subnet that the nodes are on. In this deployment that is the 10.7.10.0/24 subnet which is labeled akc-k8s-subnet. Navigate to the k8s-RG resource group and click on the k8s-subnet route tab:



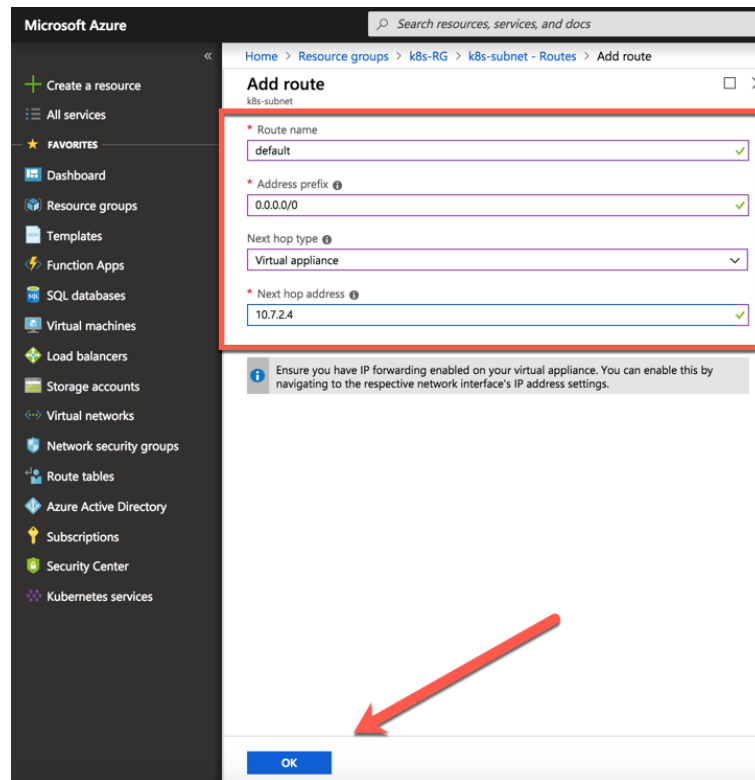
You can see a route to the app gateway subnet and that this is assigned to the 10.7.10.0/24 subnet:



Click on Routes on the left NAV and then click the “+Add” to add a new route:



Create a route with the following Parameters:

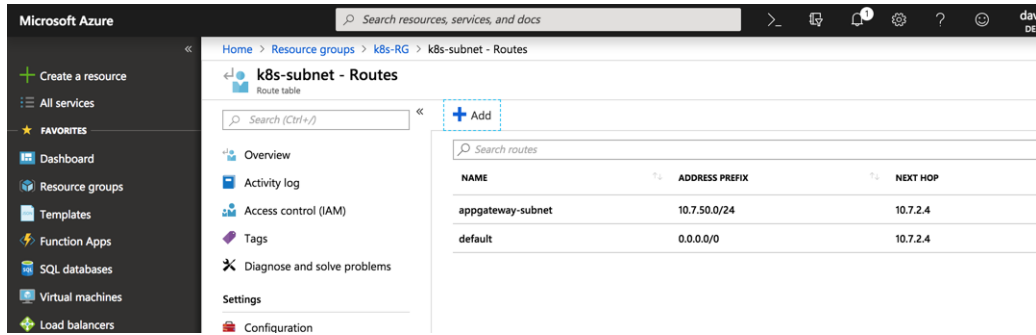


Route name: default  
Address prefix: 0.0.0.0/0  
Next hop type: Virtual Appliance  
Next hop address: 10.7.2.4

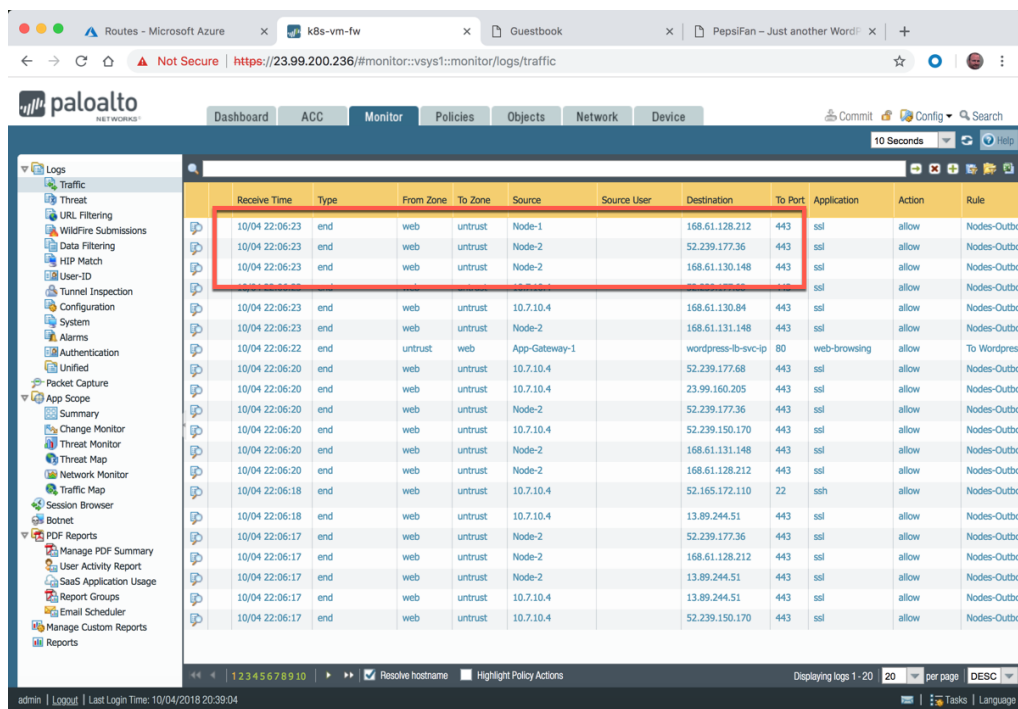
And then click Create



The new route should appear in the list:



Navigate back to the firewall monitor tab and you can now see outbound traffic as well from the cluster nodes.

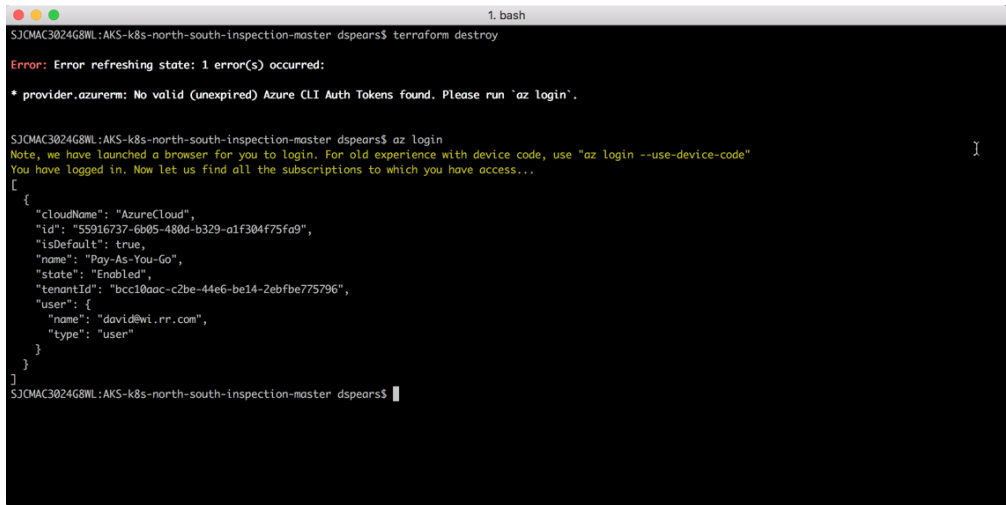


These source addresses are the instance addresses of the Kubernetes cluster node servers:

# Lab Termination

One advantage of Terraform is that it provides the ability to remove the deployment so it is not incurring ongoing cost but could be easily instantiated at a later time for testing and demonstrations. To destroy the lab, go to the terminal prompt and navigate to the directory that was used to deploy the environment and execute:

```
$ terraform destroy
```



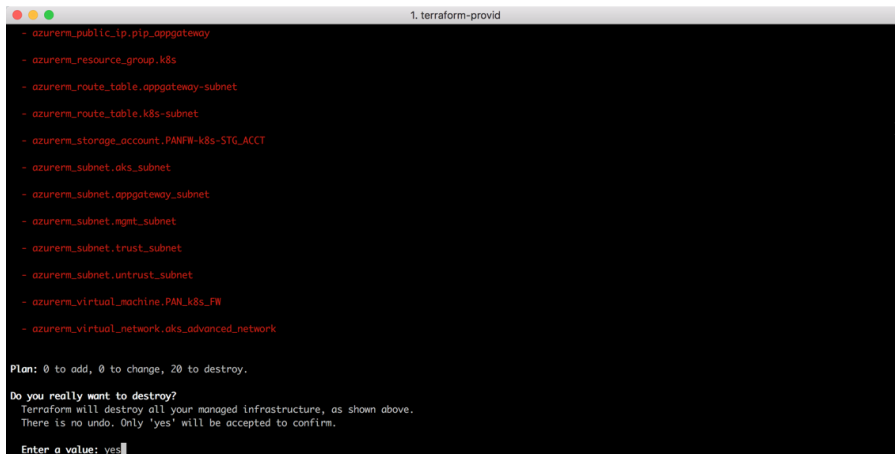
```
1. bash
SJM302468W:AKS-k8s-north-south-inspection-master dspears$ terraform destroy
Error: Error refreshing state: 1 error(s) occurred:

* provider.azurem: No valid (unexpired) Azure CLI Auth Tokens found. Please run 'az login'.

SJM302468W:AKS-k8s-north-south-inspection-master dspears$ az login
Note, we have launched a browser for you to login. For old experience with device code, use "az login --use-device-code"
You have logged in. Now let us find all the subscriptions to which you have access...
[
  {
    "cloudName": "AzureCloud",
    "id": "55916737-6b05-480d-b329-a1f304f75fa9",
    "isDefault": true,
    "name": "Pay-As-You-Go",
    "state": "Enabled",
    "tenantId": "bcc10aac-c2be-44e6-be14-2ebf775796",
    "user": {
      "name": "david@wi.rr.com",
      "type": "user"
    }
  }
]
SJM302468W:AKS-k8s-north-south-inspection-master dspears$
```

If an error message appears regarding the CLI Auth Tokens, run the **az login** command to get a new token.

Terraform will show the list of items that will be removed. Type **yes** at the prompt to start the process:



```
1. terraform-provid
- azurem_public_ip.pip_appgateway
- azurem_resource_group.k8s
- azurem_route_table.appgateway_subnet
- azurem_route_table.k8s_subnet
- azurem_storage_account.PANFW-k8s-STG_ACCT
- azurem_subnet.aks_subnet
- azurem_subnet.appgateway_subnet
- azurem_subnet.ngmt_subnet
- azurem_subnet.trust_subnet
- azurem_subnet.untrust_subnet
- azurem_virtual_machine.PAN_k8s_FW
- azurem_virtual_network.aks_advanced_network

Plan: 0 to add, 0 to change, 20 to destroy.

Do you really want to destroy?
Terraform will destroy all your managed infrastructure, as shown above.
There is no undo. Only 'yes' will be accepted to confirm.
Enter a value: yes
```

This should result in the complete removal of all the resources:

```
1. bash
azurerm_kubernetes_cluster.k8s: Still destroying... (ID: /subscriptions/55916737-6b05-480d-b329-...rvice/managedClusters/k8s-Cluster-MGMT, 9m20s elapsed)
azurerm_kubernetes_cluster.k8s: Still destroying... (ID: /subscriptions/55916737-6b05-480d-b329-...rvice/managedClusters/k8s-Cluster-MGMT, 9m30s elapsed)
azurerm_kubernetes_cluster.k8s: Still destroying... (ID: /subscriptions/55916737-6b05-480d-b329-...rvice/managedClusters/k8s-Cluster-MGMT, 9m40s elapsed)
azurerm_kubernetes_cluster.k8s: Still destroying... (ID: /subscriptions/55916737-6b05-480d-b329-...rvice/managedClusters/k8s-Cluster-MGMT, 9m50s elapsed)
azurerm_kubernetes_cluster.k8s: Still destroying... (ID: /subscriptions/55916737-6b05-480d-b329-...rvice/managedClusters/k8s-Cluster-MGMT, 10m0s elapsed)
azurerm_kubernetes_cluster.k8s: Still destroying... (ID: /subscriptions/55916737-6b05-480d-b329-...rvice/managedClusters/k8s-Cluster-MGMT, 10m10s elapsed)
azurerm_kubernetes_cluster.k8s: Still destroying... (ID: /subscriptions/55916737-6b05-480d-b329-...rvice/managedClusters/k8s-Cluster-MGMT, 10m20s elapsed)
azurerm_kubernetes_cluster.k8s: Still destroying... (ID: /subscriptions/55916737-6b05-480d-b329-...rvice/managedClusters/k8s-Cluster-MGMT, 10m30s elapsed)
azurerm_kubernetes_cluster.k8s: Still destroying... (ID: /subscriptions/55916737-6b05-480d-b329-...rvice/managedClusters/k8s-Cluster-MGMT, 10m40s elapsed)
azurerm_kubernetes_cluster.k8s: Still destroying... (ID: /subscriptions/55916737-6b05-480d-b329-...rvice/managedClusters/k8s-Cluster-MGMT, 10m50s elapsed)
azurerm_kubernetes_cluster.k8s: Still destroying... (ID: /subscriptions/55916737-6b05-480d-b329-...rvice/managedClusters/k8s-Cluster-MGMT, 11m0s elapsed)
azurerm_kubernetes_cluster.k8s: Still destroying... (ID: /subscriptions/55916737-6b05-480d-b329-...rvice/managedClusters/k8s-Cluster-MGMT, 11m10s elapsed)
azurerm_kubernetes_cluster.k8s: Destruction complete after 11m10s
azurerm_subnet.aks_subnet: Destroying... (ID: /subscriptions/55916737-6b05-480d-b329-...ks/aks-k8s-vnet/subnets/aks-k8s-subnet)
azurerm_subnet.aks_subnet: Destruction complete after 1s
azurerm_route_table.k8s-subnet: Destroying... (ID: /subscriptions/55916737-6b05-480d-b329-...rosoft.Network/routeTables/k8s-subnet)
azurerm_network_security_group.aks_advanced_network: Destroying... (ID: /subscriptions/55916737-6b05-480d-b329-...ork/networkSecurityGroups/aks-k8s-nsg)
azurerm_virtual_network.aks_advanced_network: Destroying... (ID: /subscriptions/55916737-6b05-480d-b329-...t.Network/virtualNetworks/aks-k8s-vnet)
azurerm_route_table.k8s-subnet: Destruction complete after 1s
azurerm_network_security_group.aks_advanced_network: Destruction complete after 1s
azurerm_virtual_network.aks_advanced_network: Still destroying... (ID: /subscriptions/55916737-6b05-480d-b329-...t.Network/virtualNetworks/aks-k8s-vnet, 10s elapsed)
azurerm_virtual_network.aks_advanced_network: Destruction complete after 11s
azurerm_resource_group.k8s: Destroying... (ID: /subscriptions/55916737-6b05-480d-b329-a1f304f75fa9/resourceGroups/k8s-RG)
azurerm_resource_group.k8s: Still destroying... (ID: /subscriptions/55916737-6b05-480d-b329-a1f304f75fa9/resourceGroups/k8s-RG, 10s elapsed)
azurerm_resource_group.k8s: Still destroying... (ID: /subscriptions/55916737-6b05-480d-b329-a1f304f75fa9/resourceGroups/k8s-RG, 20s elapsed)
azurerm_resource_group.k8s: Still destroying... (ID: /subscriptions/55916737-6b05-480d-b329-a1f304f75fa9/resourceGroups/k8s-RG, 30s elapsed)
azurerm_resource_group.k8s: Still destroying... (ID: /subscriptions/55916737-6b05-480d-b329-a1f304f75fa9/resourceGroups/k8s-RG, 40s elapsed)
azurerm_resource_group.k8s: Destruction complete after 48s

Destroy complete! Resources: 20 destroyed.
SJCMAC3024G8WL:AKS-k8s-north-south-inspection-master dspears$
```

This can be validated by executing the **terraform destroy** command one more time:

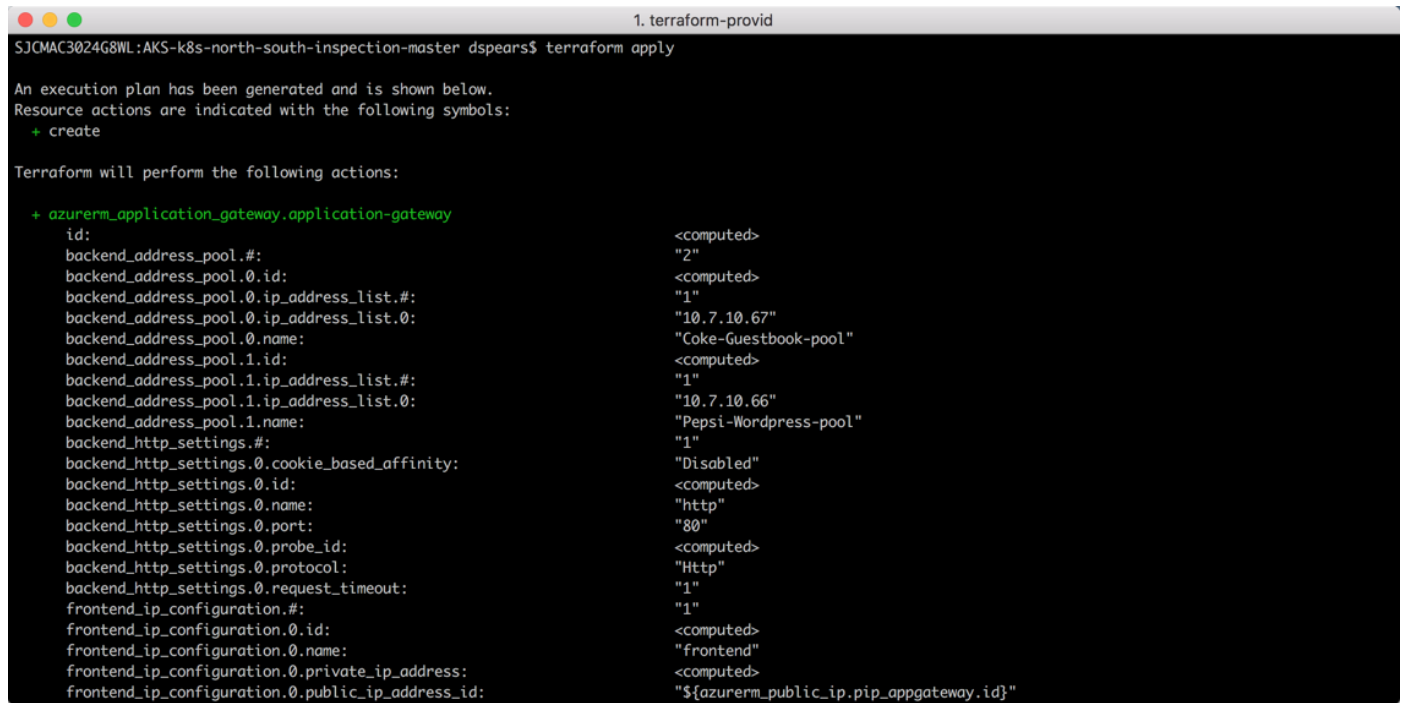
```
1. bash
azurerm_kubernetes_cluster.k8s: Still destroying... (ID: /subscriptions/55916737-6b05-480d-b329-...rvice/managedClusters/k8s-Cluster-MGMT, 10m50s elapsed)
azurerm_kubernetes_cluster.k8s: Still destroying... (ID: /subscriptions/55916737-6b05-480d-b329-...rvice/managedClusters/k8s-Cluster-MGMT, 11m0s elapsed)
azurerm_kubernetes_cluster.k8s: Still destroying... (ID: /subscriptions/55916737-6b05-480d-b329-...rvice/managedClusters/k8s-Cluster-MGMT, 11m10s elapsed)
azurerm_kubernetes_cluster.k8s: Destruction complete after 11m10s
azurerm_subnet.aks_subnet: Destroying... (ID: /subscriptions/55916737-6b05-480d-b329-...ks/aks-k8s-vnet/subnets/aks-k8s-subnet)
azurerm_subnet.aks_subnet: Destruction complete after 1s
azurerm_route_table.k8s-subnet: Destroying... (ID: /subscriptions/55916737-6b05-480d-b329-...rosoft.Network/routeTables/k8s-subnet)
azurerm_network_security_group.aks_advanced_network: Destroying... (ID: /subscriptions/55916737-6b05-480d-b329-...ork/networkSecurityGroups/aks-k8s-nsg)
azurerm_virtual_network.aks_advanced_network: Destroying... (ID: /subscriptions/55916737-6b05-480d-b329-...t.Network/virtualNetworks/aks-k8s-vnet)
azurerm_route_table.k8s-subnet: Destruction complete after 1s
azurerm_network_security_group.aks_advanced_network: Destruction complete after 1s
azurerm_virtual_network.aks_advanced_network: Still destroying... (ID: /subscriptions/55916737-6b05-480d-b329-...t.Network/virtualNetworks/aks-k8s-vnet, 10s elapsed)
azurerm_virtual_network.aks_advanced_network: Destruction complete after 11s
azurerm_resource_group.k8s: Destroying... (ID: /subscriptions/55916737-6b05-480d-b329-a1f304f75fa9/resourceGroups/k8s-RG)
azurerm_resource_group.k8s: Still destroying... (ID: /subscriptions/55916737-6b05-480d-b329-a1f304f75fa9/resourceGroups/k8s-RG, 10s elapsed)
azurerm_resource_group.k8s: Still destroying... (ID: /subscriptions/55916737-6b05-480d-b329-a1f304f75fa9/resourceGroups/k8s-RG, 20s elapsed)
azurerm_resource_group.k8s: Still destroying... (ID: /subscriptions/55916737-6b05-480d-b329-a1f304f75fa9/resourceGroups/k8s-RG, 30s elapsed)
azurerm_resource_group.k8s: Still destroying... (ID: /subscriptions/55916737-6b05-480d-b329-a1f304f75fa9/resourceGroups/k8s-RG, 40s elapsed)
azurerm_resource_group.k8s: Destruction complete after 48s

Destroy complete! Resources: 20 destroyed.
SJCMAC3024G8WL:AKS-k8s-north-south-inspection-master dspears$ terraform destroy
Do you really want to destroy?
Terraform will destroy all your managed infrastructure, as shown above.
There is no undo. Only 'yes' will be accepted to confirm.

Enter a value: yes

Destroy complete! Resources: 0 destroyed.
SJCMAC3024G8WL:AKS-k8s-north-south-inspection-master dspears$
```

At any point in the future, it is possible to come back to this directory and simply run the **terraform apply** command and quickly install the environment again:



```
1. terraform-provid
SJCMAC3024G8WL:AKS-k8s-north-south-inspection-master dspears$ terraform apply

An execution plan has been generated and is shown below.
Resource actions are indicated with the following symbols:
+ create

Terraform will perform the following actions:

+ azurem_application_gateway.application-gateway
  id: <computed>
  backend_address_pool.#: "2"
  backend_address_pool.0.id: <computed>
  backend_address_pool.0.ip_address_list.#: "1"
  backend_address_pool.0.ip_address_list.0: "10.7.10.67"
  backend_address_pool.0.name: "Coke-Guestbook-pool"
  backend_address_pool.1.id: <computed>
  backend_address_pool.1.ip_address_list.#: "1"
  backend_address_pool.1.ip_address_list.0: "10.7.10.66"
  backend_address_pool.1.name: "Pepsi-Wordpress-pool"
  backend_http_settings.#: "1"
  backend_http_settings.0.cookie_based_affinity: "Disabled"
  backend_http_settings.0.id: <computed>
  backend_http_settings.0.name: "http"
  backend_http_settings.0.port: "80"
  backend_http_settings.0.probe_id: <computed>
  backend_http_settings.0.protocol: "Http"
  backend_http_settings.0.request_timeout: "1"
  frontend_ip_configuration.#: "1"
  frontend_ip_configuration.0.id: <computed>
  frontend_ip_configuration.0.name: "frontend"
  frontend_ip_configuration.0.private_ip_address: <computed>
  frontend_ip_configuration.0.public_ip_address_id: "${azurerm_public_ip.pip_appgateway.id}"
```

**End of Activity**

## Conclusion

Congratulations! You have now successfully integrated the VM-Series firewall to gain visibility into North/South traffic for two container application hosted in a Kubernetes cluster.