



Deployment Steps

NetApp HCI Solutions

NetApp
August 29, 2024

This PDF was generated from https://docs.netapp.com/us-en/hci-solutions/anthos_task_configure_management_switches.html on August 29, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- Deployment Steps 1
 - 1. Configure management switches 1
 - 2. Configure Data Switches 5
 - 3. Deploy NetApp HCI with the NetApp Deployment Engine 12
 - 4. Configure the vCenter Server 21
 - 5. Deploy and Configure the F5 Big-IP Virtual Edition Appliance 25
 - Complete Anthos prerequisites 28
 - 7. Deploy the Anthos admin workstation 34
 - 8. Deploy the admin cluster 37
 - 9. Deploy Additional User Clusters: NetApp HCI with Anthos 47
 - 10. Enable access to the cluster with the GKE console 50
 - 11. Install and configure NetApp Trident storage provisioner 53

Deployment Steps

This section provides detailed protocols for implementing the NetApp HCI solution for Anthos.

This deployment is divided into the following high-level tasks:

1. [Configure management switches](#)
2. [Configure data switches](#)
3. [Deploy NetApp HCI with the NetApp Deployment Engine](#)
4. [Configure the vCenter Server](#)
5. [Deploy and configure the F5 Big-IP Virtual Edition Appliance](#)
6. [Complete Anthos prerequisites](#)
7. [Deploy the Anthos admin workstation](#)
8. [Deploy the admin cluster](#)
9. [Deploy user clusters](#)
10. [Enable access to cluster with the GKE console](#)
11. [Install and configure NetApp Trident storage provisioner](#)

Next: [Configure management switches](#).

1. Configure management switches

Cisco Nexus 3048 switches are used in this deployment procedure to provide 1Gbps connectivity for in- and out-of-band management of the compute and storage nodes. These steps begin after the switches have been racked, powered, and put through the initial setup process. To configure the switches to provide management connectivity to the infrastructure, complete the following steps:

Enable advanced features for Cisco Nexus

Run the following commands on each Cisco Nexus 3048 switch to configure advanced features:

1. Enter configuration mode.

```
Switch-01# configure terminal
```

2. Enable VLAN functionality.

```
Switch-01(config)# feature interface-vlan
```

3. Enable LACP.

```
Switch-01(config)# feature lacp
```

4. Enable virtual port channels (vPCs).

```
Switch-01(config)# feature vpc
```


5. Set the global port-channel load-balancing configuration.

```
Switch-01(config)# port-channel load-balance src-dst ip-l4port
```

6. Perform the global spanning-tree configuration.

```
Switch-01(config)# spanning-tree port type network default  
Switch-01(config)# spanning-tree port type edge bpduguard default
```

Configure ports on the switch for in-band management

1. Run the following commands to create VLANs for management purposes.

```
Switch-01(config)# vlan 2  
Switch-01(config-vlan)# Name Native_VLAN  
Switch-01(config-vlan)# vlan 16  
Switch-01(config-vlan)# Name OOB_Network  
Switch-01(config-vlan)# vlan 3480  
Switch-01(config-vlan)# Name MGMT_Network  
Switch-01(config-vlan)# exit
```

2. Configure the ports ETH1/29-32 as VLAN trunk ports that connect to management interfaces on each HCI storage node.


```
Switch-01(config)# int eth 1/29
Switch-01(config-if)# description HCI-STG-01 PortA
Switch-01(config-if)# switchport mode trunk
Switch-01(config-if)# switchport trunk native vlan 2
Switch-01(config-if)# switchport trunk allowed vlan 3480
Switch-01(config-if)# spanning tree port type edge trunk
Switch-01(config-if)# int eth 1/30
Switch-01(config-if)# description HCI-STG-02 PortA
Switch-01(config-if)# switchport mode trunk
Switch-01(config-if)# switchport trunk native vlan 2
Switch-01(config-if)# switchport trunk allowed vlan 3480
Switch-01(config-if)# spanning tree port type edge trunk
Switch-01(config-if)# int eth 1/31
Switch-01(config-if)# description HCI-STG-03 PortA
Switch-01(config-if)# switchport mode trunk
Switch-01(config-if)# switchport trunk native vlan 2
Switch-01(config-if)# switchport trunk allowed vlan 3480
Switch-01(config-if)# spanning tree port type edge trunk
Switch-01(config-if)# int eth 1/32
Switch-01(config-if)# description HCI-STG-04 PortA
Switch-01(config-if)# switchport mode trunk
Switch-01(config-if)# switchport trunk native vlan 2
Switch-01(config-if)# switchport trunk allowed vlan 3480
Switch-01(config-if)# spanning tree port type edge trunk
Switch-01(config-if)# exit
```

Configure ports on the switch for out-of-band management

1. Run the following commands to configure the ports for cabling the IPMI interfaces on each HCI node.


```

Switch-01(config)# int eth 1/13
Switch-01(config-if)# description HCI-CMP-01 IPMI
Switch-01(config-if)# switchport mode access
Switch-01(config-if)# switchport access vlan 16
Switch-01(config-if)# spanning-tree port type edge
Switch-01(config-if)# int eth 1/14
Switch-01(config-if)# description HCI-STG-01 IPMI
Switch-01(config-if)# switchport mode access
Switch-01(config-if)# switchport access vlan 16
Switch-01(config-if)# spanning-tree port type edge
Switch-01(config-if)# int eth 1/15
Switch-01(config-if)# description HCI-STG-03 IPMI
Switch-01(config-if)# switchport mode access
Switch-01(config-if)# switchport access vlan 16
Switch-01(config-if)# spanning-tree port type edge
Switch-01(config-if)# exit

```



In the validated configuration, we cabled odd-node IPMI interfaces to Switch-01, and even-node IPMI interfaces to Switch-02.

Create a vPC domain to ensure fault tolerance

1. Activate the ports used for the vPC peer-link between the two switches.

```

Switch-01(config)# int eth 1/1
Switch-01(config-if)# description vPC peer-link Switch-02 1/1
Switch-01(config-if)# int eth 1/2
Switch-01(config-if)# description vPC peer-link Switch-02 1/2
Switch-01(config-if)# exit

```

2. Perform the vPC global configuration.


```

Switch-01(config)# vpc domain 1
Switch-01(config-vpc-domain)# role priority 10
Switch-01(config-vpc-domain)# peer-keepalive destination <switch-02_mgmt_address> source <switch-01_mgmt_address> vrf management
Switch-01(config-vpc-domain)# peer-gateway
Switch-01(config-vpc-domain)# auto recovery
Switch-01(config-vpc-domain)# ip arp synchronize
Switch-01(config-vpc-domain)# int eth 1/1-2
Switch-01(config-vpc-domain)# channel-group 10 mode active
Switch-01(config-vpc-domain)# int Po10
Switch-01(config-if)# description vPC peer-link
Switch-01(config-if)# switchport mode trunk
Switch-01(config-if)# switchport trunk native vlan 2
Switch-01(config-if)# switchport trunk allowed vlan 16,3480
Switch-01(config-if)# spanning-tree port type network
Switch-01(config-if)# vpc peer-link
Switch-01(config-if)# exit

```

[Next: Configure Data Switches](#)

2. Configure Data Switches

Mellanox SN2010 switches provide 25Gbps connectivity for the data plane of the compute and storage nodes. To configure the switches to provide data connectivity to the infrastructure, complete the following steps:

Create MLAG cluster to provide fault tolerance

1. Run the following commands on each Mellanox SN210 switch for general configuration:
 - a. Enter configuration mode.

```

Switch-01 enable
Switch-01 configure terminal

```

- b. Enable the LACP required for the Inter-Peer Link (IPL).

```

Switch-01 (config) # lacp

```

- c. Enable the Link Layer Discovery Protocol (LLDP).

```

Switch-01 (config) # lldp

```

- d. Enable IP routing.


```
Switch-01 (config) # ip routing
```

- e. Enable the MLAG protocol.

```
Switch-01 (config) # protocol mlag
```

- f. Enable global QoS.

```
Switch-01 (config) # dcb priority-flow-control enable force
```

2. For MLAG to function, the switches must be made peers to each other through an IPL. This should consist of two or more physical links for redundancy. The MTU for the IPL is set for jumbo frames (9216), and all VLANs are enabled by default. Run the following commands on each switch in the domain:

- a. Create port channel 10 for the IPL.

```
Switch-01 (config) # interface port-channel 10
Switch-01 (config interface port-channel 10) # description IPL
Switch-01 (config interface port-channel 10) # exit
```

- b. Add interfaces ETH 1/20 and 1/22 to the port channel.

```
Switch-01 (config) # interface ethernet 1/20 channel-group 10 mode
active
Switch-01 (config) # interface ethernet 1/20 description ISL-SWB_01
Switch-01 (config) # interface ethernet 1/22 channel-group 10 mode
active
Switch-01 (config) # interface ethernet 1/22 description ISL-SWB_02
```

- c. Create a VLAN outside of the standard range dedicated to IPL traffic.

```
Switch-01 (config) # vlan 4000
Switch-01 (config vlan 4000) # name IPL VLAN
Switch-01 (config vlan 4000) # exit
```

- d. Define the port channel as the IPL.

```
Switch-01 (config) # interface port-channel 10 ipl 1
Switch-01 (config) # interface port-channel 10 dcb priority-flow-
control mode on force
```


- e. Set an IP for each IPL member (non-routable; it is not advertised outside of the switch).

```
Switch-01 (config) # interface vlan 4000
Switch-01 (config vlan 4000) # ip address 10.0.0.1 255.255.255.0
Switch-01 (config vlan 4000) # ipl 1 peer-address 10.0.0.2
Switch-01 (config vlan 4000) # exit
```

3. Create a unique MLAG domain name for the two switches and assign an MLAG virtual IP (VIP). This IP is used for keep-alive heartbeat messages between the two switches. Run these commands on each switch in the domain:

- a. Create the MLAG domain and set the IP address and subnet.

```
Switch-01 (config) # mlag-vip MLAG-VIP-DOM ip a.b.c.d /24 force
```

- b. Create a virtual MAC address for the system MLAG.

```
Switch-01 (config) # mlag system-mac AA:BB:CC:DD:EE:FF
```

- c. Configure the MLAG domain so that it is active globally.

```
Switch-01 (config) # no mlag shutdown
```



The IP used for the MLAG VIP must be in the same subnet as the switch management network (mgmt0).



The MAC address used can be any unicast MAC address and must be set to the same value on both switches in the MLAG domain.

Configure ports to connect to storage and compute hosts

1. Create each of the VLANs needed to support the services for NetApp HCI. Run these commands on each switch in the domain:

- a. Create VLANs.

```
Switch-01 (config) # vlan 1172
Switch-01 (config vlan 1172) exit
Switch-01 (config) # vlan 3480-3482
Switch-01 (config vlan 3480-3482) exit
```

- b. Create names for each VLAN for easier accounting.


```
Switch-01 (config) # vlan 1172 name "VM_Network"
Switch-01 (config) # vlan 3480 name "MGMT_Network"
Switch-01 (config) # vlan 3481 name "Storage_Network"
Switch-01 (config) # vlan 3482 name "vMotion_Network"
+
```

2. Create hybrid VLAN ports on ports ETH1/9-10 so that you can tag the appropriate VLANs for the NetApp HCI compute nodes.

a. Select the ports you want to work with.

```
Switch-01 (config) # interface ethernet 1/9-1/10
```

b. Set the MTU for each port.

```
Switch-01 (config interface ethernet 1/9-1/10) # mtu 9216 force
```

c. Modify spanning-tree settings for each port.

```
Switch-01 (config interface ethernet 1/9-1/10) # spanning-tree
bpdudfilter enable
Switch-01 (config interface ethernet 1/9-1/10) # spanning-tree port
type edge
Switch-01 (config interface ethernet 1/9-1/10) # spanning-tree
bpduguard enable
```

d. Set the switchport mode to hybrid.

```
Switch-01 (config interface ethernet 1/9-1/10 ) # switchport mode
hybrid
Switch-01 (config interface ethernet 1/9-1/10 ) # exit
```

e. Create descriptions for each port being modified.

```
Switch-01 (config) # interface ethernet 1/9 description HCI-CMP-01
PortD
Switch-01 (config) # interface ethernet 1/10 description HCI-CMP-02
PortD
```

f. Tag the appropriate VLANs for the NetApp HCI environment.


```
Switch-01 (config) # interface ethernet 1/9 switchport hybrid
allowed-vlan add 1172
Switch-01 (config) # interface ethernet 1/9 switchport hybrid
allowed-vlan add 3480-3482
Switch-01 (config) # interface ethernet 1/10 switchport hybrid
allowed-vlan add 1172
Switch-01 (config) # interface ethernet 1/10 switchport hybrid
allowed-vlan add 3480-3482
```

3. Create MLAG interfaces and hybrid VLAN ports on ports ETH1/5-8 so that you can distribute connectivity between the switches and tag the appropriate VLANs for the NetApp HCI storage nodes.

- a. Select the ports that you want to work with.

```
Switch-01 (config) # interface ethernet 1/5-1/8
```

- b. Set the MTU for each port.

```
Switch-01 (config interface ethernet 1/5-1/8) # mtu 9216 force
```

- c. Modify spanning tree settings for each port.

```
Switch-01 (config interface ethernet 1/5-1/8) # spanning-tree
bpdufilter enable
Switch-01 (config interface ethernet 1/5-1/8) # spanning-tree port
type edge
Switch-01 (config interface ethernet 1/5-1/8) # spanning-tree
bpduguard enable
```

- d. Set the switchport mode to hybrid.

```
Switch-01 (config interface ethernet 1/5-1/8 ) # switchport mode
hybrid
Switch-01 (config interface ethernet 1/5-1/8 ) # exit
```

- e. Create descriptions for each port being modified.


```
Switch-01 (config) # interface ethernet 1/5 description HCI-STG-01
PortD
Switch-01 (config) # interface ethernet 1/6 description HCI-STG-02
PortD
Switch-01 (config) # interface ethernet 1/7 description HCI-STG-03
PortD
Switch-01 (config) # interface ethernet 1/8 description HCI-STG-04
PortD
```

f. Create and configure the MLAG port channels.

```
Switch-01 (config) # interface mlag-port-channel 115-118
Switch-01 (config interface mlag-port-channel 115-118) # exit
Switch-01 (config) # interface mlag-port-channel 115-118 no shutdown
Switch-01 (config) # interface mlag-port-channel 115-118 mtu 9216
force
Switch-01 (config) # interface mlag-port-channel 115-118 lacp-
individual enable force
Switch-01 (config) # interface ethernet 1/5-1/8 lacp port-priority 10
Switch-01 (config) # interface ethernet 1/5-1/8 lacp rate fast
Switch-01 (config) # interface ethernet 1/5 mlag-channel-group 115
mode active
Switch-01 (config) # interface ethernet 1/6 mlag-channel-group 116
mode active
Switch-01 (config) # interface ethernet 1/7 mlag-channel-group 117
mode active
Switch-01 (config) # interface ethernet 1/8 mlag-channel-group 118
mode active
```

g. Tag the appropriate VLANs for the storage environment.


```
Switch-01 (config) # interface mlag-port-channel 115-118 switchport
mode hybrid
Switch-01 (config) # interface mlag-port-channel 115 switchport
hybrid allowed-vlan add 1172
Switch-01 (config) # interface mlag-port-channel 116 switchport
hybrid allowed-vlan add 1172
Switch-01 (config) # interface mlag-port-channel 117 switchport
hybrid allowed-vlan add 1172
Switch-01 (config) # interface mlag-port-channel 118 switchport
hybrid allowed-vlan add 1172
Switch-01 (config) # interface mlag-port-channel 115 switchport
hybrid allowed-vlan add 3481
Switch-01 (config) # interface mlag-port-channel 116 switchport
hybrid allowed-vlan add 3481
Switch-01 (config) # interface mlag-port-channel 117 switchport
hybrid allowed-vlan add 3481
Switch-01 (config) # interface mlag-port-channel 118 switchport
hybrid allowed-vlan add 3481
```



The configurations in this section must also be run on the second switch in the MLAG domain. NetApp recommends that the descriptions for each port are updated to reflect the device ports that are cabled and configured on the other switch.

Create uplink ports for the switches

1. Create an MLAG interface to provide uplinks to both Mellanox SN2010 switches from the core network.

```
Switch-01 (config) # interface mlag port-channel 101
Switch-01 (config interface mlag port-channel) # description Uplink
CORE-SWITCH port PORT
Switch-01 (config interface mlag port-channel) # exit
```

2. Configure the MLAG members.

```
Switch-01 (config) # interface ethernet 1/18 description Uplink to CORE-
SWITCH port PORT
Switch-01 (config) # interface ethernet 1/18 speed 10000 force
Switch-01 (config) # interface mlag-port-channel 101 mtu 9216 force
Switch-01 (config) # interface ethernet 1/18 mlag-channel-group 101 mode
active
```

3. Set the switchport mode to hybrid and allow all VLANs from the core uplink switches.


```
Switch-01 (config) # interface mlag-port-channel switchport mode hybrid
Switch-01 (config) # interface mlag-port-channel switchport hybrid
allowed-vlan all
```

4. Verify that the MLAG interface is up.

```
Switch-01 (config) # interface mlag-port-channel 101 no shutdown
Switch-01 (config) # exit
```


[Next: Deploy NetApp HCI with the NetApp Deployment Engine](#)

3. Deploy NetApp HCI with the NetApp Deployment Engine

NDE delivers a simple and streamlined deployment experience for the NetApp HCI solution. A detailed guide to using NDE 1.6 to deploy your NetApp HCI system can be found [here](#).

These steps begin after the nodes have been racked, and cabled, and the IPMI port has been configured on each node using the console. To Deploy the NetApp HCI solution using NDE, complete the following steps:

1. Access the out-of-band management console for one of the storage nodes in the cluster and log in with the default credentials ADMIN/ADMIN.

The image shows a NetApp login interface. At the top left is the NetApp logo. The main heading is "Please Login". Below this, there are two input fields: "Username" with the text "ADMIN" entered, and "Password" with five dots (masked). Below the password field is a "Login" button.

2. Click the Remote Console Preview image in the center of the screen to download a JNLP file launched by Java Web Start, which launches an interactive console to the system.

3. With the virtual console launched, a user can log in to the HCI storage node using the ADMIN/ADMIN username and password combination.
4. The Bond1G interface must have an IP, a netmask, and a gateway set statically; its VLAN set to 3480; and DNS servers defined for the environment.

```
Bond10G
Method          : static
Link Speed      : 50000
IPv4 Address     :
IPv4 Subnet Mask :
---->
IPv4 Gateway Address :
---->
MTU             : 9000
---->
Bond Mode        : LACP [ActivePassive, ALB, LACP]
---->
LACP Rate        : Fast [Fast, Slow]
---->
Status           : UpAndRunning [Down, Up, UpAndRunning]
---->
Virtual Network Tag :
---->
Routes           : Number of routes: 0.
---->
```



Select an IP that is within the subnet you intend to use for in-band management but not an IP you would like to use in production. NDE reconfigures the node with a production IP after initial access.



This task must only be performed on the first storage node. Afterward, the other nodes in the infrastructure are discovered by the Automatic Private IP Address (APIPA) addresses assigned to each storage interface when left unconfigured.

5. The Bond 10G interface must have its MTU setting changed to enable jumbo frames and its bond mode changed to LACP.


```

Bond10G

Method          : static

Link Speed      : 50000

IPv4 Address     :

IPv4 Subnet Mask :
---->

IPv4 Gateway Address :
---->

MTU             : 9000
---->

Bond Mode       : LACP   [ActivePassive, ALB, LACP]
---->

LACP Rate       : Fast   [Fast, Slow]
---->

Status          : UpAndRunning   [Down, Up, UpAndRunning]
---->

Virtual Network Tag :
---->

Routes          : Number of routes: 0.
---->

```



Configure each of the four storage nodes in the NetApp HCI solution this way. The NDE process is then able to discover all the nodes in the solution and configure them. You do not need to modify the Bond10g interfaces on the two compute nodes.

6. After completion, open a web browser and visit the IP address you configured for the management port to start NetApp HCI configuration with NDE.
7. On the Welcome to NetApp HCI page, click the Get Started button.
8. Check each associated box on the Prerequisites page and click Continue.
9. The next page presents End User Licenses for NetApp HCI and VMware vSphere. If you accept the terms, click I Accept at the end of each agreement and then click Continue.
10. Click Configure a New vSphere Deployment, select vSphere 6.5U2, and enter the Fully Qualified Domain Name (FQDN) of your vCenter Server. Then click Continue.

vSphere Configuration

You may elect to configure a new vSphere deployment or to join an existing vSphere deployment.

- ☒ Configure a new vSphere deployment
 - ☐ Configure Using vSphere Version 6.7 Update 1
 - ☒ Configure Using vSphere Version 6.5 Update 2
- ☐ Join and extend an existing vSphere deployment

If you have set up a DNS record for your new vCenter server, then configure your server using its fully qualified domain name and DNS server IP address:

- ☒ Configure Using a Fully Qualified Domain Name  *Best Practice!*

vCenter Server Fully Qualified Domain Name

anthos-vc.cie.netapp.com 

Note: The domain name must resolve to an unused IP address.

DNS Server IP Address

10.61.184.251 

If you have not set up a DNS record for your new vCenter server, you may configure using an IP address that we define:

- ☐ Configure Using an IP Address 

Note: Once defined, the IP address cannot be changed.

[Back](#)

[Continue](#)

11. NDE asks for the credentials to be used in the environment. This is used for VMware vSphere, the NetApp Element storage cluster, and the NetApp Mnode, which provides management functionality for the cluster. When you are finished, click Continue.

Credentials

Define the user name and password that will be used for the storage cluster, vCenter, and the management node.

User Name

Password

Password must contain:

- ✓ At least 8 characters
- ✓ No more than 20 characters
- ✓ 1 uppercase letter that is not the first character ?
- ✓ 1 lowercase letter
- ✓ 1 of the following special characters: !@\$
- ✓ Allowed characters: A-Z a-z 0-9 !@\$
- ✓ 1 number that is not the last character ?

Re-enter Password

Back

Continue

12. NDE then prompts for the network topology used to cable the NetApp HCI environment. The validated solution in this document has been deployed using the two-cable option for the compute nodes, and the four-cable option for the storage nodes. Click Continue.

Network Topology


Select a compute node topology and a storage node topology appropriate for your hardware installation.

Compute Node Topology

☐ 6 Cable Option

The 6 cable option provides dedicated ports for management (2 x 1/10 GbE), virtual machines (2 x 10/25 GbE) and storage (2 x 10/25 GbE).


☐ Use vSphere Distributed Switch? [?](#)



(H300E,H410C,H500E,H700E)

☒ 2 Cable Option

The 2 cable option provides shared management with ports for virtual machines and storage (2 x 10/25 GbE). The 2 cable option uses vSphere Distributed Switch. [?](#)




(H300E,H410C,H500E,H700E)

Storage Node Topology

☒ 4 Cable Option

The 4 cable option provides dedicated ports for management (2 x 1/10 GbE) and storage (2 x 10/25 GbE).



(H300S,H410S,H500S,H700S)

[Back](#)

[Continue](#)

13. The next page presented by NDE is the inventory of the environment as discovered by the APIPA addressed on the storage network. The storage node that is currently running NDE is already selected with a green check mark. Select the corresponding boxes to add additional nodes to the NetApp HCI environment. Click Continue.

Inventory

Verify the available nodes and select **at least 2 compute nodes and 4 storage nodes** to include in your installation.

Refresh Inventory

Compute Nodes

Filter								
<input checked="" type="checkbox"/>	Serial Number ↕	Chassis Serial Number / Slot ↕	Node Type ↕	Software Version ↕	Physical CPU Cores ↕	Memory ↕	1 GbE Ports ↕	10 GbE Ports ↕
<input checked="" type="checkbox"/>	HM17CS002729	002170990158 / B	H410C	1.6	8	384 GB	0 of 2 detected	2 of 4 detected
<input checked="" type="checkbox"/>	HM181S002024	002170990158 / A	H410C	1.6	8	384 GB	0 of 2 detected	2 of 4 detected
1 - 2 of 2 results					20			

2 compute nodes selected

Storage Nodes

Filter								
<input checked="" type="checkbox"/>	Serial Number ↕	Chassis Serial Number / Slot ↕	Node Type ↕	Raw Capacity ↕	Element Version ↕	Drive Count ↕	1 GbE Ports ↕	10 GbE Ports ↕
<input checked="" type="checkbox"/>	221814003506	221814003436 / C	H500S	5.76 TB	11.3.1.5	6 of 6 detected	2 of 2 detected	2 of 2 detected
<input checked="" type="checkbox"/>	221818004613	221814003436 / D	H500S	5.76 TB	11.3.1.5	6 of 6 detected	2 of 2 detected	2 of 2 detected
<input checked="" type="checkbox"/>	221826005865	002170990158 / C	H500S	5.76 TB	11.3.1.5	6 of 6 detected	2 of 2 detected	2 of 2 detected
<input checked="" type="checkbox"/>	221826005866	002170990158 / D	H500S	5.76 TB	11.3.1.5	6 of 6 detected	2 of 2 detected	2 of 2 detected
1 - 4 of 4 results					20			

4 storage nodes selected

Back

Continue



If there are any nodes missing from the inventory screen, wait a few minutes and click Refresh Inventory. If the node still fails to appear, additional investigation of environment networking might be required.

14. You must next configure the permanent network settings for the NetApp HCI deployment. The first page configures infrastructure services (DNS and NTP), vCenter networking, and Mnode networking.

Network Settings

Provide the network settings that will be used for your installation.

Live network validation is: **On** ?

Infrastructure Services

DNS Server IP Address 1

10.61.184.251



DNS Server IP Address 2 (Optional)

10.61.184.252



NTP Server Address 1 ?

10.61.184.251



NTP Server Address 2 (Optional)

10.61.184.252



To save time, launch the easy form to enter fewer network settings. >



vCenter Networking

VLAN ID	Subnet ?	Default Gateway	FQDN	IP Address
3480	172.21.224.0/24	172.21.224.1	anthos-vc.cie.netapp.com	172.21.224.10

Management Node Networking ?

Management Network		iSCSI Network
VLAN ID	VLAN ID	
3480	3481	
Subnet ?	Subnet ?	
172.21.224.0/24	172.21.225.0/24	
Default Gateway		
172.21.224.1		
Hostname	Management IP Address	Storage (iSCSI) IP Address
anthos-mnode	172.21.224.50	172.21.225.50

15. The next page allows you to configure each node in the environment. For the compute nodes, it allows you to configure the host name, management network, vMotion network, and storage network. For the storage nodes, name the storage cluster and configure the management and storage networks being used for each node. Click Continue.

Compute Node Networking

		Management Network	vMotion Network	iSCSI A Network	iSCSI B Network
VLAN ID		3480 ✓	3482 ✓	3481 ✓	3481 ✓
Subnet ?		172.21.224.0/24 ✓	172.21.226.0/24 ✓	172.21.225.0/24 ✓	172.21.225.0/24 ✓
Default Gateway		172.21.224.1 ✓	Default Gateway (Optional)	Default Gateway (Optional)	Default Gateway (Optional)

Serial Number	Hostname	Management IP Address	vMotion IP Address	iSCSI A - IP Address	iSCSI B - IP Address
HM17CS002729	Anthos-ESXi-01 ✓	172.21.224.11 ✓	172.21.226.11 ✓	172.21.225.11 ✓	172.21.225.111 ✓
HM18IS002024	Anthos-ESXi-02 ✓	172.21.224.12 ✓	172.21.226.12 ✓	172.21.225.12 ✓	172.21.225.112 ✓

Storage Node Networking

Storage Cluster Name

Anthos-Store ✓

Note: The storage cluster name cannot be changed after deployment.

		Management Network	iSCSI Network
VLAN ID		3480 ✓	3481 ✓
Subnet ?		172.21.224.0/24 ✓	172.21.225.0/24 ✓
Default Gateway		172.21.224.1 ✓	Default Gateway (Optional)
Management Virtual IP (MVIP) ?		172.21.224.20 ✓	Storage Virtual IP (SVIP) ?
			172.21.225.20 ✓

Serial Number	Hostname	Management IP Address	Storage (iSCSI) IP Address
221814003506	Anthos-Store-01 ✓	172.21.224.21 ✓	172.21.225.21 ✓
221818004613	Anthos-Store-02 ✓	172.21.224.22 ✓	172.21.225.22 ✓
221826005865	Anthos-Store-03 ✓	172.21.224.23 ✓	172.21.225.23 ✓
221826005866	Anthos-Store-04 ✓	172.21.224.24 ✓	172.21.225.24 ✓

Back

Live network validation is: On ?


Continue

- On the next page, review all the settings that have been defined for the environment by expanding each section, and, if necessary, click Edit to make corrections. There is also a check box on this page that enables or disables the Mnode from sending real-time health and diagnostics information to NetApp Active IQ. If all the information is correct, click Start Deployment.



If you want to enable Active IQ, verify that your management network can reach the internet. If NDE is unable to reach Active IQ, the deployment can fail.

- A summary page appears along with a progress bar for each component of the NetApp HCI solution, as well as the overall solution. When complete, you are presented with an option to launch the vSphere client and begin working with your environment.



Next: [Configure the vCenter Server](#)

4. Configure the vCenter Server

NDE deploys the solution with vCenter server and integrates the solution with the Element cluster by provisioning the Mnode VM and installing the NetApp Element Plug-in for vCenter.




Note that NDE deploys vSphere 6.7U1. You can upgrade the Virtual Appliance and individual ESXi hosts by following the instructions from VMware [here](#).

After deployment, you must make a few modifications to the environment, including the creation of additional vDS portgroups, datastores, and resource groups for the deployment of the Anthos on VMware solution.

Complete the following steps to configure your vCenter Server:

1. Log into the VMware vCenter server using the [Administrator@vsphere.local](#) account and the password chosen for the admin user during NDE configuration.



2. Right-click `NetApp-HCI-Cluster-01` created by NDE and select the option to create a new resource pool. Name this pool `Infrastructure-Resource-Pool` and accept the defaults by clicking OK. This resource pool is used in a later configuration step.

Name	<u>Infrastructure Resource</u>		
▼ CPU			
Shares	Normal ▼	4000	
Reservation	0 ▼	MHz ▼	Max reservation: 54,128 MHz
Reservation Type	<input checked="" type="checkbox"/> Expandable		
Limit	Unlimited ▼	MHz ▼	Max limit: 58,128 MHz
▼ Memory			
Shares	Normal ▼	163840	
Reservation	0 ▼	MB ▼	Max reservation: 751,064 MB
Reservation Type	<input checked="" type="checkbox"/> Expandable		
Limit	Unlimited ▼	MB ▼	Max limit: 756,820 MB

CANCEL

OK



The reservations in this resource pool can be modified based on the resources available in the environment. NetApp HCI is deployed as an all-in-one solution. Therefore, NetApp recommends reserving the resources necessary to provide availability for the infrastructure services by placing them into this resource pool and adjusting the resources appropriately. Infrastructure services include vCenter Server, NetApp Mnode, and F5 Big-IP Load Balancer.

3. Repeat this step to create another resource pool for VMs deployed by Anthos. Name this pool Anthos-Resource-Pool, and click the OK button to accept the default values. Adjust the resource availability based on the specific environment in which you are deploying the solution. This resource pool is used in a later deployment step.
4. To configure Element volumes to be used as vSphere datastores, click the dropdown menu and select NetApp Element Management from the list.
5. A Getting Started screen appears with details about your Element cluster.



6. Click Management, and the vSphere client presents a list of datastores. Click Create Datastore to create one datastore to host VMs and another to host ISOs for future guest installs.
7. Next click the Network menu item in the left panel. This displays a screen with information about the vDS deployed by NDE.
8. Several virtual port groups are defined by the initial configuration. NetApp recommends leaving these alone to support the infrastructure, and additional port groups should be created for user-deployed virtual guests. Right-click the NetApp HCI VDS 01 vDS in the left panel, and then select Distributed Port Group followed by the New Distributed Port Group option from the expanded menu.
9. Create a new distributed port group called `Management_Network`. Then click Next.
10. On the next screen, select the VLAN type as VLAN, and set the VLAN ID to 3480 for management purposes. Click Next, and, after reviewing the options on the summary page, click Next again to complete the creation of the distributed port group.
11. Repeat these steps to create distributed port groups for the `VM_Network` (VLAN 1172) as well as any other networks that might be used in the NetApp HCI environment.



Additional networks can be defined to segment any additional deployed VMs. Examples of this use could be for a dedicated HA network for additional F5 Big-IP appliances if provisioned. Such configurations are in addition to the environment deployed in this validated solution and are considered out of scope for this NVA document.

[Next: Deploy and Configure the F5 Big-IP Virtual Edition Appliance](#)

5. Deploy and Configure the F5 Big-IP Virtual Edition Appliance

Anthos enables native integration with F5 Big-IP load balancers to expose services from each pod to the world.

This solution makes use of the virtual appliance deployed in VMware vSphere as deployed by NDE. Networking for the F5 Big-IP virtual appliance can be configured in a two-armed or three-armed configuration based on your network environment. The deployment in this document is based on the two-armed configuration. Additional details for configuring the virtual appliance for use with Anthos can be found [here](#).

To deploy the F5 Big-IP Virtual Edition appliance, complete the following steps:

1. Download the virtual application Open Virtual Appliance (OVA) file from F5 [here](#).



To download the appliance, a user must register with F5. They provide a 30-day demo license for the Big-IP Virtual Edition Load Balancer. NetApp recommends a permanent 10Gbps license for the production deployment of an appliance.

2. Right-click the infrastructure resource pool and select Deploy OVF Template. A wizard launches that allows you to select the OVA file that you just downloaded in Step 1. Click Next.

Deploy OVF Template

1 Select an OVF template

2 Select a name and folder

3 Select a compute resource

4 Review details

5 Select storage

6 Ready to complete

Select an OVF template

Select an OVF template from remote URL or local file system

Enter a URL to download and install the OVF package from the Internet, or browse to a location accessible from your computer, such as a local hard drive, a network share, or a CD/DVD drive.

☐ URL

☒ Local file

BIGIP-15.0.1-0.....ALL-vmware.ova

CANCEL

BACK

NEXT

3. Click Next to continue through each step and accept the default values for each screen presented until you reach the storage selection screen. Select the VM_Datastore that was created earlier, and then click Next.
4. The next screen presented by the wizard allows you to customize the virtual networks for use in the environment. Select VM_Network for the External field and select Management_Network for the Management field. Internal and HA are used for advanced configurations for the F5 Big-IP appliance and are not configured. These parameters can be left alone, or they can be configured to connect to non-infrastructure, distributed port groups. Click Next.
5. Review the summary screen for the appliance, and, if all the information is correct, click Finish to start the deployment.
6. After the virtual appliance is deployed, right-click it and power it up. It should receive a DHCP address on the management network. The appliance is Linux-based, and it has VMware Tools deployed, so that you can view the DHCP address it receives in the vSphere client.
7. Open a web browser and connect to the appliance at the IP address from the previous step. The default login is admin/admin, and, after the first login, the appliance immediately prompts you to change the admin password. It then returns you to a screen where you must log in with the new credentials.

f5 BIG-IP Configuration Utility
F5 Networks, Inc.

Hostname
bigip1

IP Address
172.21.224.20

Username
admin

Password

Log in

Welcome to the BIG-IP Configuration Utility.
Log in with your username and password using the fields on the left.

(c) Copyright 1996-2019, F5 Networks, Inc., Seattle, Washington. All rights reserved.
[F5 Networks, Inc. Legal Notices](#)

8. The first screen prompts the you to complete the Setup Utility. Begin the utility by clicking Next.

9. The next screen prompts you for activation of the appliance license. Click Activate to begin. When prompted on the next page, paste either the 30-day evaluation license key you received when you registered for the download or the permanent license you acquired when you purchased the appliance. Click Next.



For the device to perform activation, the network defined on the management interface must be able to reach the internet.

10. On the next screen, the End User License Agreement (EULA) is presented. If the terms in the license are acceptable, click Accept.
11. The next screen counts the elapsed time as it verifies the configuration changes that have been made so far. Click Continue to resume with the initial configuration.
12. The Configuration Change window closes, and the Setup Utility displays the Resource Provisioning menu. This window lists the features that are currently licensed and the current resource allocations for the virtual appliance and each running service.
13. Clicking the Platform menu option on the left enables additional modification of the platform. Modifications include setting the management IP address configured with DHCP, setting the host name and the time zone the appliance is installed in, and securing the appliance from SSH accessibility.
14. Next click the Network menu, which enables you to configure standard networking features. Click Next to begin the Standard Network Configuration wizard.
15. The first page of the wizard configures redundancy; leave the defaults and click Next. The next page enables you to configure an internal interface on the load balancer. Interface 1.1 maps to the vmnic labeled Internal in the OVF deployment wizard.



The fields in this page for Self IP Address, Netmask, and Floating IP address can be filled with a non-routable IP address for use as a placeholder. They can also be filled with an internal network that has been configured as a distributed port group for virtual guests if you are deploying the three-armed configuration. They must be completed to continue with the wizard.

16. The next page enables you to configure an external network that is used to map services to the pods deployed in Kubernetes. Select a static IP from the VM_Network range, the appropriate subnet mask, and a floating IP from that same range. Interface 1.2 maps to the vmnic labeled External in the OVF deployment wizard.
17. On the next page, you can configure an internal-HA network if you are deploying multiple virtual appliances

in the environment. To proceed, you must fill the Self-IP Address and the Netmask fields, and you must select interface 1.3 as the VLAN Interface, which maps to the HA network defined by the OVF template wizard.

18. The next page enables you to configure the NTP servers. Then click Next to continue to the DNS setup. The DNS servers and domain search list should already be populated by the DHCP server. Click Next to accept the defaults and continue.
19. For the remainder of the wizard, click Next to continue through the advanced peering setup, the configuration of which is beyond the scope of this document. Then click Finish to exit the wizard.
20. Create individual partitions for the Anthos admin cluster and each user cluster deployed in the environment. Click System in the menu on the left, navigate to Users, and click Partition List.
21. The displayed screen only shows the current common partition. Click Create on the right to create the first additional partition and name it Anthos-Admin. Then click Repeat, name the partition Anthos-Cluster1, and click the Repeat button again to name the next partition Anthos-Cluster2. Finally click Finished to complete the wizard. The Partition list screen returns with all the partitions now listed.

Next: [Complete Anthos prerequisites.](#)

Complete Anthos prerequisites

Now that the physical environment is set up, you can begin Anthos deployment. This starts with several prerequisites that you must meet to deploy the solution and access it afterward. Each of these steps are discussed in depth in the Anthos [GKE On-Prem Guide](#).

To prepare your environment for the deployment of Anthos on VMware, complete the following steps:

1. Create a Google Cloud project following the instructions available [here](#).



Your organization might already have a project in place intended for this purpose. Check with your cloud administration team to see if a project exists and is already configured for access to Anthos on VMware. All projects intended for use with Anthos must be whitelisted by Google. This includes the primary user account, additional team members, and the access service account created in a later step.

2. Create a deployment workstation from which to manage the installation of Anthos on VMware. The deployment workstation can be Linux, MacOS, or Windows. For the purposes of this validated deployment, Red Hat Enterprise Linux 7 was used.



This workstation can be hosted either internal or external to the NetApp HCI deployment. The only requirement is that it must be able to successfully communicate with the deployed VMware vCenter Server and the internet to function correctly.

3. Install [Google Cloud SDK](#) for interactions with Google Cloud. It can be downloaded as an archive of binaries for manual install or installed by either the apt-get (Ubuntu/Debian) or yum (RHEL) package managers.


```

[user@rhel7 ~]$ sudo yum install google-cloud-sdk
Failed to set locale, defaulting to C
Loaded plugins: langpacks, product-id, search-disabled-repos,
subscription-manager
Resolving Dependencies
--> Running transaction check
---> Package google-cloud-sdk.noarch 0:270.0.0-1 will be installed
--> Finished Dependency Resolution

Dependencies Resolved

=====
=====
Package                               Arch          Version          Repository
Size
=====
=====
Installing:
google-cloud-sdk                      noarch         270.0.0-1        google-cloud-
sdk                                  36 M
Transaction Summary
=====
=====
Install 1 Package

Total download size: 36 M
Installed size: 174 M
Is this ok [y/d/N]: y
Downloading packages:
6d81c821884ae40244c746f6044fc1bcd801143a0d9c8da06767036b8d090a24-google-
cloud-sdk-270.0.0-1.noar | 36 MB  00:00:00
Running transaction check
Running transaction test
Transaction test succeeded
Running transaction
  Installing : google-cloud-sdk-270.0.0-1.noarch
1/1
  Verifying   : google-cloud-sdk-270.0.0-1.noarch
1/1

Installed:
  google-cloud-sdk.noarch 0:270.0.0-1

Complete!

```




The gcloud binary must be at least version 265.0.0. You can update a manual install with a gcloud components update. However, if SDK was installed by a package manager, future updates must also be performed using that same package manager.

4. With the workstation configured, log in to Google Cloud with your credentials. To do so, enter the login command from the deployment workstation and retrieve a link that can be copied and pasted into a browser to allow interactive sign-in to Google services. After you have logged in, the web page presents a code that you can copy and paste back into the deployment workstation when prompted.

```
[user@rhel7 ~]$ gcloud auth login
```

Go to the following link in your browser:

```
https://accounts.google.com/o/oauth2/auth?code_challenge=7oPNSySHr_Sd2ZZ4K83koIeGTLVcdbjc8omr6zCbAI&prompt=select_account&code_challenge_method=S256&access_type=offline&redirect_uri=urn%3Aietf%3Awg%3Aoauth%3A2.0%3Aoob&response_type=code&client_id=32655940559.apps.googleusercontent.com&scope=https%3A%3F%2Fwww.googleapis.com%2Fauth%2Fuserinfo.email+https%3A%2F%2Fwww.googleapis.com%2Fauth%2Fcloud-platform+https%3A%6F%2Fwww.googleapis.com%2Fauth%2Fappengine.admin+https%3A%2F%2Fwww.googleapis.com%2Fauth%2Fcompute+https%3A%2F%2Fwww.googleapis.com%2Fauth%2Faccounts.reauth
```

Enter verification code: 6/swGAh52VVgB-

TRS5LVrSvP79ZdDlb9V6ObYUGqoY67a3zp9NPciIKsM

You are now logged in as [user@netapp.com].

Your current project is [anthos-dev]. You can change this setting by running:

```
$ gcloud config set project PROJECT_ID
```

5. Enable several APIs so that your environment can communicate with Google Cloud. The pods deployed in your clusters must be able to access <https://www.googleapis.com> and <https://gkeconnect.googleapis.com> to function as expected. Therefore, the VM_Network that the worker nodes are attached to must have internet access. To enable the necessary APIs, run the following command from the deployment workstation:

```
[user@rhel7 ~]$ gcloud services enable --project anthos-dev \  
cloudresourcemanager.googleapis.com \  
container.googleapis.com \  
gkeconnect.googleapis.com \  
gkehub.googleapis.com \  
serviceusage.googleapis.com \  
stackdriver.googleapis.com \  
monitoring.googleapis.com \  
logging.googleapis.com
```


6. Create a working directory called `anthos-install`, and change into that directory.

```
[user@rhel7 ~]$ mkdir anthos-install && cd anthos-install
[user@rhel7 anthos-install]$
```

7. Before you can install Anthos on VMware, you must create four service accounts, each with a specific purpose in interacting with Google Cloud. The following table lists the accounts and their purposes.

Account Name	Purpose
component-access-sa	Used to download the Anthos binaries from Cloud Storage.
connect-register-sa	Used to register Anthos clusters to the Google Cloud console.
connect-agent-sa	Used to maintain the connection between user clusters and the Google Cloud.
logging-monitoring-sa	Used to write logging and monitoring data to Stackdriver.



Each account is assigned an email address that references your approved Google Cloud project name. The following examples all list the project `Anthos-Dev`, which was used during the NetApp validation. Make sure to substitute your appropriate project name in syntax examples where necessary.


```

[user@rhel7 anthos-install]$ gcloud iam service-accounts create
component-access-sa \
    --display-name "Component Access Service Account" \
    --project anthos-dev
[user@rhel7 anthos-install]$ gcloud iam service-accounts keys create
component-access-key.json \
    --iam-account component-access-sa@anthos-dev.iam.gserviceaccount.com

[user@rhel7 anthos-install]$ gcloud iam service-accounts create connect-
register-sa \
    --project anthos-dev
[user@rhel7 anthos-install]$ gcloud iam service-accounts keys create
connect-register-key.json \
    --iam-account connect-register-sa@anthos-dev.iam.gserviceaccount.com

[user@rhel7 anthos-install]$ gcloud iam service-accounts create connect-
agent-sa \
    --project anthos-dev
[user@rhel7 anthos-install]$ gcloud iam service-accounts keys create
connect-agent-key.json \
    --iam-account connect-agent-sa@anthos-dev.iam.gserviceaccount.com

[user@rhel7 anthos-install]$ gcloud iam service-accounts create logging-
monitoring-sa \
    --project anthos-dev
[user@rhel7 anthos-install]$ gcloud iam service-accounts keys create
logging-monitoring-key.json \
    --iam-account logging-monitoring-sa@anthos-
dev.iam.gserviceaccount.com

```

8. The final step needed to prepare your environment to deploy Anthos is to limit certain privileges to your service accounts. You need the associated email address for each service account listed in Step 7.
 - a. Using the component-access-sa account, assign the roles for `serviceusage.serviceUsageViewer`, `iam.serviceAccountCreator`, and `iam.roleViewer`.


```
[user@rhel7 anthos-install]$ gcloud projects add-iam-policy-binding
anthos-dev\
  --member "serviceAccount:component-access-sa@anthos-
dev.iam.gserviceaccount.com" \
  --role "roles/serviceusage.serviceUsageViewer"
[user@rhel7 anthos-install]$ gcloud projects add-iam-policy-binding
anthos-dev\
  --member "serviceAccount:component-access-sa@anthos-
dev.iam.gserviceaccount.com" \
  --role "roles/iam.serviceAccountCreator"
[user@rhel7 anthos-install]$ gcloud projects add-iam-policy-binding
anthos-dev\
  --member "serviceAccount:component-access-sa@anthos-
dev.iam.gserviceaccount.com" \
  --role "roles/iam.roleViewer"
```

- b. Using the connect-register-sa service account, assign the role for gkehub.admin.

```
[user@rhel7 anthos-install]$ gcloud projects add-iam-policy-binding
anthos-dev \
  --member "serviceAccount:connect-register-sa@anthos-
dev.iam.gserviceaccount.com " \
  --role "roles/gkehub.admin"
```

- c. Using the connect-agent-sa account, assign the role for gkehub.connect.

```
[user@rhel7 anthos-install]$ gcloud projects add-iam-policy-binding
anthos-dev \
  --member "serviceAccount:connect-agent-sa@anthos-
dev.iam.gserviceaccount.com" \
  --role "roles/gkehub.connect"
```

- d. With the logging-monitoring-sa service account, assign the roles for stackdriver.resourceMetadata.writer, logging.logWriter, monitoring.metricWriter, and monitoring.dashboardEditor.


```
[user@rhel7 anthos-install]$ gcloud projects add-iam-policy-binding
anthos-dev \
  --member "serviceAccount:logging-monitoring-sa@anthos-
dev.iam.gserviceaccount.com" \
  --role "roles/stackdriver.resourceMetadata.writer"
[user@rhel7 anthos-install]$ gcloud projects add-iam-policy-binding
anthos-dev\
  --member "serviceAccount:logging-monitoring-sa@anthos-
dev.iam.gserviceaccount.com" \
  --role "roles/logging.logWriter"
[user@rhel7 anthos-install]$ gcloud projects add-iam-policy-binding
anthos-dev\
  --member "serviceAccount:logging-monitoring-sa@anthos-
dev.iam.gserviceaccount.com" \
  --role "roles/monitoring.metricWriter"
[user@rhel7 anthos-install]$ gcloud projects add-iam-policy-binding
anthos-dev\
  --member "serviceAccount:logging-monitoring-sa@anthos-
dev.iam.gserviceaccount.com" \
  --role "roles/monitoring.dashboardEditor"
```

9. Download the vCenter certificate for the VMWare CA; this is used later to authenticate to the vCenter during installation.

```
[user@rhel7 anthos-install]$ true | openssl s_client -connect anthos-
vc.cie.netapp.com:443 -showcerts 2>/dev/null | sed -ne '/-BEGIN/,/-
END/p' > vcenter.pem
```

Next: [Deploy the Anthos admin workstation](#)

7. Deploy the Anthos admin workstation

The admin workstation is a vSphere VM deployed within your NetApp HCI environment that is preinstalled with all the tools necessary to administer the Anthos on VMware solution. Follow the instructions in this section to deploy the Anthos admin workstation.

To deploy the Anthos admin workstation, complete the following steps:

1. Download the gkeadm binary into your working directory

```
[user@rhel7 anthos-install]$ gsutil cp gs://gke-on-prem-
release/gkeadm/1.6.1-gke.1/linux/gkeadm ./
[user@rhel7 anthos-install]$ chmod +x gkeadm
```


2. Use the gkeadm tool to create an admin workstation configuration file.

```
[user@rhel7 anthos-install]$ ./gkeadm create config
```

3. Two files are created: credential.yaml and admin-ws-config.yaml. Fill out each of these files.
 - a. credential.yaml contains your username and passwords for your VMware vCenter server.

```
kind: CredentialFile
items:
- name: vCenter
  username: "administrator@vsphere.local"
  password: "vSphereAdminPassword"
```

- b. admin-ws-config.yaml contains other information about your vSphere environment as well as the physical and networking options for the admin-workstation VM.

```
gcp:
  # Path of the whitelisted service account's JSON key file
  whitelistedServiceAccountKeyPath: "/home/anthos-install/service-
keys/access-key.json"
  # Specify which vCenter resources to use
  vCenter:
    # The credentials and address GKE On-Prem should use to connect to
    vCenter
    credentials:
      address: "anthos-vc.cie.netapp.com"
      datacenter: "NetApp-HCI-Datacenter-01"
      datastore: "VM_Datastore"
      cluster: "NetApp-HCI-Cluster-01"
      network: "VM_Network"
      resourcePool: "Anthos-Resource-Pool"
  # Provide the path to vCenter CA certificate pub key for SSL
  verification
    caCertPath: "/home/anthos-install/vcenter.pem"
  # The URL of the proxy for the jump host
  proxyUrl: ""
  adminWorkstation:
    name: gke-admin-ws-200915-151421
    cpus: 4
    memoryMB: 8192
  #The boot disk size of the admin workstation in GB. It is recommended
  to use a disk with at least 50 GB to host images decompressed from
  the bundle.
  diskGB: 50
```



```

# Name for the persistent disk to be mounted to the home directory
(ending in
.vmdk).
# Any directory in the supplied path must be created before
deployment.
  dataDiskName: gke-on-prem-admin-workstation-data-disk/gke-admin-ws-
200915-151421-data-disk.vmdk
# The size of the data disk in MB.
  dataDiskMB: 512
  network:
# The IP allocation mode: 'dhcp' or 'static'
  ipAllocationMode: "dhcp"
# # The host config in static IP mode. Do not include if using DHCP
# hostConfig:
#   # The IPv4 static IP address for the admin workstation
#   ip: ""
#   # The IP address of the default gateway of the subnet in
which the admin workstation
#   # is to be created
#   gateway: ""
#   # The subnet mask of the network where you want to create
your admin workstation
#   netmask: ""
#   # The list of DNS nameservers to be used by the admin
workstation
#   dns:
#   - ""
# The URL of the proxy for the admin workstation
proxyUrl: ""
ntpServer: ntp.ubuntu.com

```

4. Create the admin workstation.


```
[user@rhel7 anthos-install]$ ./gkeadm create admin-workstation
The output will be verbose as the workstation is created. In the end you
will be prompted with the IP address to login to the workstation if you
chose DHCP.
...
Getting ... service account...
...
*****
Admin workstation is ready to use.

Admin workstation information saved to /usr/local/google/home/me/my-
admin-workstation
This file is required for future upgrades
SSH into the admin workstation with the following command:
ssh -i /home/user/.ssh/gke-admin-workstation ubuntu@10.63.172.10
*****
```

Next: [Deploy the admin and the first user cluster](#)

8. Deploy the admin cluster

All Kubernetes clusters deployed as a part of the Anthos solution are deployed from the Anthos admin workstation that you just created. A user logs into the admin workstation using SSH, the public key created in a previous step, and the IP address provided at the end of the VM deployment. An admin cluster controls all actions in an Anthos environment. The admin cluster must be deployed first, and then individual user clusters can be deployed for specific workload needs.



There are specific procedures for deploying clusters that use static IP addresses [here](#), and procedures for environments with DHCP can be found [here](#). In this guide, we use the second set of instructions for ease of deployment.

To deploy the admin cluster, complete the following steps:

1. Log into your admin-workstation using the SSH command prompted at the end of the deployment. After successful authentication, you can list the files in the home directory, which are used to create the admin cluster and additional clusters later on. The directory also includes the copied vCenter cert and the access key for Anthos that was created in earlier steps.


```
[user@rhel7 anthos-install]$ ssh -i ~/.ssh/gke-admin-workstation
ubuntu@10.63.172.10

Welcome to Ubuntu 18.04.5 LTS (GNU/Linux 5.4.0-1001-gkeop x86_64)

 * Documentation:  https://help.ubuntu.com
 * Management:    https://landscape.canonical.com
 * Support:       https://ubuntu.com/advantage
Last login: Fri Jan 29 15:46:35 2021 from 10.249.129.216

ubuntu@gke-admin-200915-151421:~$ ls
admin-cluster.yaml
user-cluster.yaml
vcenter.pem
component-access-key.json
```

2. Use scp to copy the remaining keys for your Anthos account over from the workstation you deployed the admin-workstation from.

```
ubuntu@gke-admin-200915-151421:~$ scp user@rhel7:~/anthos-
install/connect-register-key.json ./
ubuntu@gke-admin-200915-151421:~$ scp user@rhel7:~/anthos-
install/connect-agent-key.json ./
ubuntu@gke-admin-200915-151421:~$ scp user@rhel7:~/anthos-
install/logging-monitoring-key.json ./
```

3. Edit the admin-cluster.yaml file so that it is specific to the deployed environment. The file is very large, so we will address it by sections.
 - a. Most of the information is already filled in by default based on the configuration used to deploy the admin-workstation by gkeadm. This first section confirms the information for the version of Anthos being deployed and the vCenter instance it is deployed on. It also allows you to define a local data disk (VMDK) for Kubernetes object data.


```

apiVersion: v1
kind: AdminCluster
# (Required) Absolute path to a GKE bundle on disk
bundlePath: /var/lib/gke/bundles/gke-onprem-vmware-1.6.0-gke.7-
full.tgz
# (Required) vCenter configuration
vCenter:
  address: anthos-vc.cie.netapp.com
  datacenter: NetApp-HCI-Datacenter-01
  cluster: NetApp-HCI-Cluster-01
  resourcePool: Anthos-Resource-Pool
  datastore: VM_Datastore
  # Provide the path to vCenter CA certificate pub key for SSL
  verification
  caCertPath: "/home/ubuntu/vcenter.pem"
  # The credentials to connect to vCenter
  credentials:
    username: administrator@vsphere.local
    password: "vSphereAdminPassword"
  # Provide the name for the persistent disk to be used by the
  deployment (ending
  # in .vmdk). Any directory in the supplied path must be created
  before deployment
  dataDisk: "admin-cluster-disk.vmdk"

```

- b. Fill out the networking section next, and select whether you are using static or DHCP mode. If you are using static addresses, you must create an IP-block file based on the instructions linked to above, and add it to the config file.



If static IPs are used in a deployment, the items under the host configuration are global. This includes static IPs for clusters or those used for SeeSaw load balancers, which are configured later.


```

# (Required) Network configuration
network:
# (Required) Hostconfig for static addresses on Seesaw LB's
hostConfig:
  dnsServers:
    - "10.61.184.251"
    - "10.61.184.252"
  ntpServers:
    - "0.pool.ntp.org"
    - "1.pool.ntp.org"
    - "2.pool.ntp.org"
  searchDomainsForDNS:
    - "cie.netapp.com"
ipMode:
  # (Required) Define what IP mode to use ("dhcp" or "static")
  type: dhcp
  # # (Required when using "static" mode) The absolute or relative
  # path to the yaml file
  # # to use for static IP allocation
  # ipBlockFilePath: ""
  # (Required) The Kubernetes service CIDR range for the cluster.
  # Must not overlap
  # with the pod CIDR range
  serviceCIDR: 10.96.232.0/24
  # (Required) The Kubernetes pod CIDR range for the cluster. Must
  # not overlap with
  # the service CIDR range
  podCIDR: 192.168.0.0/16
vCenter:
  # vSphere network name
  networkName: VM_Network

```

- c. Fill out the load balancer section next. This can vary depending on the type of load balancer being deployed.

Seesaw example:

```

loadBalancer:
  # (Required) The VIPs to use for load balancing
  vips:
    # Used to connect to the Kubernetes API
    controlPlaneVIP: "10.63.172.155"
    # # (Optional) Used for admin cluster addons (needed for multi
    # cluster features). Must
    # # be the same across clusters

```



```

# # addonsVIP: "10.63.172.153"
# (Required) Which load balancer to use "F5BigIP" "Seesaw" or
"ManualLB". Uncomment
# the corresponding field below to provide the detailed spec
kind: Seesaw
# # (Required when using "ManualLB" kind) Specify pre-defined
nodeports
# manualLB:
# # NodePort for ingress service's http (only needed for user
cluster)
# ingressHTTPTNodePort: 0
# # NodePort for ingress service's https (only needed for user
cluster)
# ingressHTTPSNodePort: 0
# # NodePort for control plane service
# controlPlaneNodePort: 30968
# # NodePort for addon service (only needed for admin cluster)
# addonsNodePort: 31405
# # (Required when using "F5BigIP" kind) Specify the already-
existing partition and
# # credentials
# f5BigIP:
# address:
# credentials:
# username:
# password:
# partition:
# # # (Optional) Specify a pool name if using SNAT
# # snatPoolName: ""
# (Required when using "Seesaw" kind) Specify the Seesaw configs
seesaw:
# (Required) The absolute or relative path to the yaml file to use
for IP allocation
# for LB VMs. Must contain one or two IPs.
ipBlockFilePath: "admin-seesaw-block.yaml"
# (Required) The Virtual Router Identifier of VRRP for the Seesaw
group. Must
# be between 1-255 and unique in a VLAN.
vrid: 100
# (Required) The IP announced by the master of Seesaw group
masterIP: "10.63.172.151"
# (Required) The number CPUs per machine
cpus: 1
# (Required) Memory size in MB per machine
memoryMB: 2048
# (Optional) Network that the LB interface of Seesaw runs in

```



```
(default: cluster
#   network)
vCenter:
#   vSphere network name
networkName: VM_Network
#   (Optional) Run two LB VMs to achieve high availability
(default: false)
enableHA: false
```

- d. For a SeeSaw load balancer, you must create an additional external file to supply the static IP information for the load balancer. Create the file `admin-seesaw-block.yaml`, which was referenced in this configuration section.

```
blocks:
- netmask: "255.255.255.0"
  gateway: "10.63.172.1"
  ips:
  - ip: "10.63.172.152"
    hostname: "admin-seesaw-vm"
```

F5 BigIP Example:

```
# (Required) Load balancer configuration
loadBalancer:
# (Required) The VIPs to use for load balancing
vips:
# Used to connect to the Kubernetes API
controlPlaneVIP: "10.63.172.155"
# # (Optional) Used for admin cluster addons (needed for multi
cluster features). Must
# # be the same across clusters
# # addonsVIP: "10.63.172.153"
# (Required) Which load balancer to use "F5BigIP" "Seesaw" or
"ManualLB". Uncomment
# the corresponding field below to provide the detailed spec
kind: F5BigIP
# # (Required when using "ManualLB" kind) Specify pre-defined
nodeports
# manualLB:
# # NodePort for ingress service's http (only needed for user
cluster)
#   ingressHTTPTNodePort: 0
# # NodePort for ingress service's https (only needed for user
cluster)
#   ingressHTTPSNodePort: 0
```



```

# # NodePort for control plane service
# controlPlaneNodePort: 30968
# # NodePort for addon service (only needed for admin cluster)
# addonsNodePort: 31405
# # (Required when using "F5BigIP" kind) Specify the already-
existing partition and
# # credentials
f5BigIP:
  address: "172.21.224.21"
  credentials:
    username: "admin"
    password: "admin-password"
  partition: "Admin-Cluster"
# # # (Optional) Specify a pool name if using SNAT
# # snatPoolName: ""
# (Required when using "Seesaw" kind) Specify the Seesaw configs
# seesaw:
# (Required) The absolute or relative path to the yaml file to
use for IP allocation
# for LB VMs. Must contain one or two IPs.
# ipBlockFilePath: ""
# (Required) The Virtual Router Identifier of VRRP for the Seesaw
group. Must
# be between 1-255 and unique in a VLAN.
# vrid: 0
# (Required) The IP announced by the master of Seesaw group
# masterIP: ""
# (Required) The number CPUs per machine
# cpus: 4
# (Required) Memory size in MB per machine
# memoryMB: 8192
# (Optional) Network that the LB interface of Seesaw runs in
(default: cluster
# network)
# vCenter:
# vSphere network name
# networkName: VM_Network
# (Optional) Run two LB VMs to achieve high availability
(default: false)
# enableHA: false

```

- e. The last section of the admin config file contains additional options that can be tuned to fit the specific deployment environment. These include enabling anti-affinity groups if Anthos is being deployed on less than three ESXi servers. You can also configure proxies, private docker registries, and the connections to Stackdriver and Google Cloud for auditing.


```

antiAffinityGroups:
  # Set to false to disable DRS rule creation
  enabled: false
# (Optional) Specify the proxy configuration
proxy:
  # The URL of the proxy
  url: ""
  # The domains and IP addresses excluded from proxying
  noProxy: ""
# # (Optional) Use a private Docker registry to host GKE images
# privateRegistry:
#   # Do not include the scheme with your registry address
#   address: ""
#   credentials:
#     username: ""
#     password: ""
#   # The absolute or relative path to the CA certificate for this
#   registry
#   caCertPath: ""
# (Required): The absolute or relative path to the GCP service
# account key for pulling
# GKE images
gcrKeyPath: "/home/ubuntu/component-access-key.json"
# (Optional) Specify which GCP project to connect your logs and
# metrics to
stackdriver:
  projectID: "anthos-dev"
  # A GCP region where you would like to store logs and metrics for
  # this cluster.
  clusterLocation: "us-east1"
  enableVPC: false
  # The absolute or relative path to the key file for a GCP service
  # account used to
  # send logs and metrics from the cluster
  serviceAccountKeyPath: "/home/ubuntu/logging-monitoring-key.json"
# # (Optional) Configure kubernetes apiserver audit logging
# cloudAuditLogging:
#   projectid: ""
#   # A GCP region where you would like to store audit logs for this
#   # cluster.
#   clusterlocation: ""
#   # The absolute or relative path to the key file for a GCP service
#   # account used to
#   # send audit logs from the cluster
#   serviceaccountkeypath: ""

```




The deployment detailed in this document is a minimum configuration for validation that requires the disabling of anti-affinity rules. NetApp recommends leaving this option set to true in production deployments.



By default, Anthos on VMware uses a pre-existing, Google-owned container image registry that requires no additional setup. If you choose to use a private Docker registry for deployment, then you must configure that registry separately based on instructions found [here](#). This step is beyond the scope of this deployment guide.

4. When edits to the `admin-cluster.yaml` file are complete, be sure to check for proper syntax and spacing.

```
ubuntu@gke-admin-200915-151421:~$ gkectl check-config --config admin-  
cluster.yaml
```

5. After the configuration check has passed and any identified issues have been remedied, you can then stage the deployment of the cluster. Since we have already checked the validation of the config file, we can skip those steps by passing the `--skip-validation-all` flag.

```
ubuntu@gke-admin-200915-151421:~$ gkectl prepare --config admin-  
cluster.yaml --skip-validation-all
```

6. If you are using a SeeSaw load balancer, you must create one before deploying the cluster itself (otherwise skip this step).

```
ubuntu@gke-admin-200915-151421:~$ gkectl create loadbalancer --config  
admin-cluster.yaml
```

7. You can now stand up the admin cluster. This is done with the `gkectl create admin` command, which can use the `--skip-validation-all` flag to speed up deployment.

```
ubuntu@gke-admin-200915-151421:~$ gkectl create admin --config admin-  
cluster.yaml --skip-validation-all
```

8. When the cluster is deployed, it creates the kubeconfig file in the local directory. This file can be used to check the status of the cluster using `kubectl` or run diagnostics with `gkectl`.


```

ubuntu@gke-admin-ws-200915-151421:~ $ kubectl get nodes --kubeconfig
kubeconfig
NAME                                STATUS    ROLES    AGE
VERSION
gke-admin-master-gkvm              Ready    master   5m
v1.18.6-gke.6600
gke-admin-node-84b77ff5c7-6zg59    Ready    <none>    5m
v1.18.6-gke.6600
gke-admin-node-84b77ff5c7-8jdmz    Ready    <none>    5m
v1.18.6-gke.6600
ubuntu@gke-admin-ws-200915-151421:~$ gkectl diagnose cluster --
kubeconfig kubeconfig
Diagnosing admin cluster "gke-admin-gkvm"...- Validation Category:
Admin Cluster VCenter
Checking Credentials...SUCCESS
Checking Version...SUCCESS
Checking Datacenter...SUCCESS
Checking Datastore...SUCCESS
Checking Resource pool...SUCCESS
Checking Folder...SUCCESS
Checking Network...SUCCESS- Validation Category: Admin Cluster
Checking cluster object...SUCCESS
Checking machine deployment...SUCCESS
Checking machineset...SUCCESS
Checking machine objects...SUCCESS
Checking kube-system pods...SUCCESS
Checking storage...SUCCESS
Checking resource...System pods on UserMaster cpu resource request
report: total 1754m nodeCount 2 min 877m max 877m avg 877m tracked
amount in bundle 4000m
System pods on AdminNode cpu resource request report: total 2769m
nodeCount 2 min 1252m max 1517m avg 1384m tracked amount in bundle 4000m
System pods on AdminMaster cpu resource request report: total 923m
nodeCount 1 min 923m max 923m avg 923m tracked amount in bundle 4000m
System pods on UserMaster memory resource request report: total
4524461824 nodeCount 2 min 2262230912 max 2262230912 avg 2262230912
tracked amount in bundle 8192Mi
System pods on AdminNode memory resource request report: total 6876Mi
nodeCount 2 min 2174Mi max 4702Mi avg 3438Mi tracked amount in bundle
16384Mi
System pods on AdminMaster memory resource request report: total 465Mi
nodeCount 1 min 465Mi max 465Mi avg 465Mi tracked amount in bundle
16384Mi
SUCCESS
Cluster is healthy.

```


9. Deploy Additional User Clusters: NetApp HCI with Anthos

With Anthos, organizations can scale their environments to incorporate multiple user clusters and segregate workloads between teams. A single admin cluster can support up to five user clusters, and each user cluster can support up to twenty-five nodes.

To add additional user clusters to your deployment, complete the following steps:

1. Copy the `config.yaml` file to a new file named `anthos-cluster02-config.yaml`.

```
ubuntu@Anthos-Admin-Workstation:~$ cp config.yaml anthos-cluster02-  
config.yaml
```

2. Make the following edits to the newly created file:

1. Comment out the sections that refer to the existing admin cluster with (#).
2. When you get to the `usercluster` section, update the following fields:
 1. Update the partition name under the `bigip` section.
 2. Update the `controlplanvip` and `ingressvip` values under the `vip` section.
 3. Update the `clustername` value.

```
usercluster:  
  # In-Cluster vCenter configuration  
  vcenter:  
    # If specified it overwrites the network field in global  
vcenter configuration  
    network: ""  
    # # The absolute or relative path to the yaml file to use for  
static IP allocation.  
    # # Do not include if using DHCP  
    # ipblockfilepath: ""  
    # # Specify pre-defined nodeports if using "manual" load  
balancer mode  
    # manullbspec:  
    #   ingresshttpnodeport: 30243  
    #   ingresshttpsnodeport: 30879  
    #   controlplanenodeport: 30562  
    #   addonsnodeport: 0  
    # Specify the already-existing partition and credentials to use  
with F5  
    bigip:  
      # To re-use credentials across clusters we recommend using
```



```

YAML node anchors.
# See https://yaml.org/spec/1.2/spec.html#id2785586
credentials:
  address: "172.21.224.22"
  username: "admin"
  password: "NetApp!23"
partition: "Anthos-Cluster02-Part"
# # Optionally specify a pool name if using SNAT
# snatpoolname: ""
# The VIPs to use for load balancing
vips:
  # Used to connect to the Kubernetes API
  controlplanevip: "10.63.172.108"
  # Shared by all services for ingress traffic
  ingressvip: "10.63.172.109"
  # # Used for admin cluster addons (needed for multi cluster
features). Must be the same
  # # across clusters
  # addonsvip: ""
# A unique name for this cluster
clustername: "anthos-cluster02"
# User cluster master nodes must have either 1 or 3 replicas
masternode:
  cpus: 4
  memorymb: 8192
  # How many machines of this type to deploy
  replicas: 1
# The number of worker nodes to deploy and their size. Min. 2
replicas
workernode:
  cpus: 4
  memorymb: 8192
  # How many machines of this type to deploy
  replicas: 3
# The Kubernetes service CIDR range for the cluster
serviceiprange: 10.96.0.0/12
# The Kubernetes pod CIDR range for the cluster
podiprange: 192.168.0.0/16

```

3. Run the following command to check the config file again to verify that there are no syntax errors. Because you have removed the admin section, you must reference the `kubeconfig` file for the admin cluster named `kubeconfig` (found in the working directory).


```
ubuntu@Anthos-Admin-Workstation:~$ gkectl check-config --config anthos-
cluster02-config.yaml --kubeconfig kubeconfig
- Validation Category: Config Check
  - [SUCCESS] Config

- Validation Category: Docker Registry
  - [SUCCESS] gcr.io/gke-on-prem-release access

- Validation Category: vCenter
  - [SUCCESS] Credentials
  - [SUCCESS] Datacenter
  - [SUCCESS] Datastore
  - [FAILURE] Data Disk: vCenter data disk already exists
  - [SUCCESS] Resource Pool
  - [SUCCESS] Network

- Validation Category: F5 BIG-IP
  - [SUCCESS] Credentials
  - [SUCCESS] Partition

- Validation Category: Network Configuration
  - [SUCCESS] CIDR, VIP and static IP (availability and overlapping)

- Validation Category: VIPs
  - [SUCCESS] ping (availability)

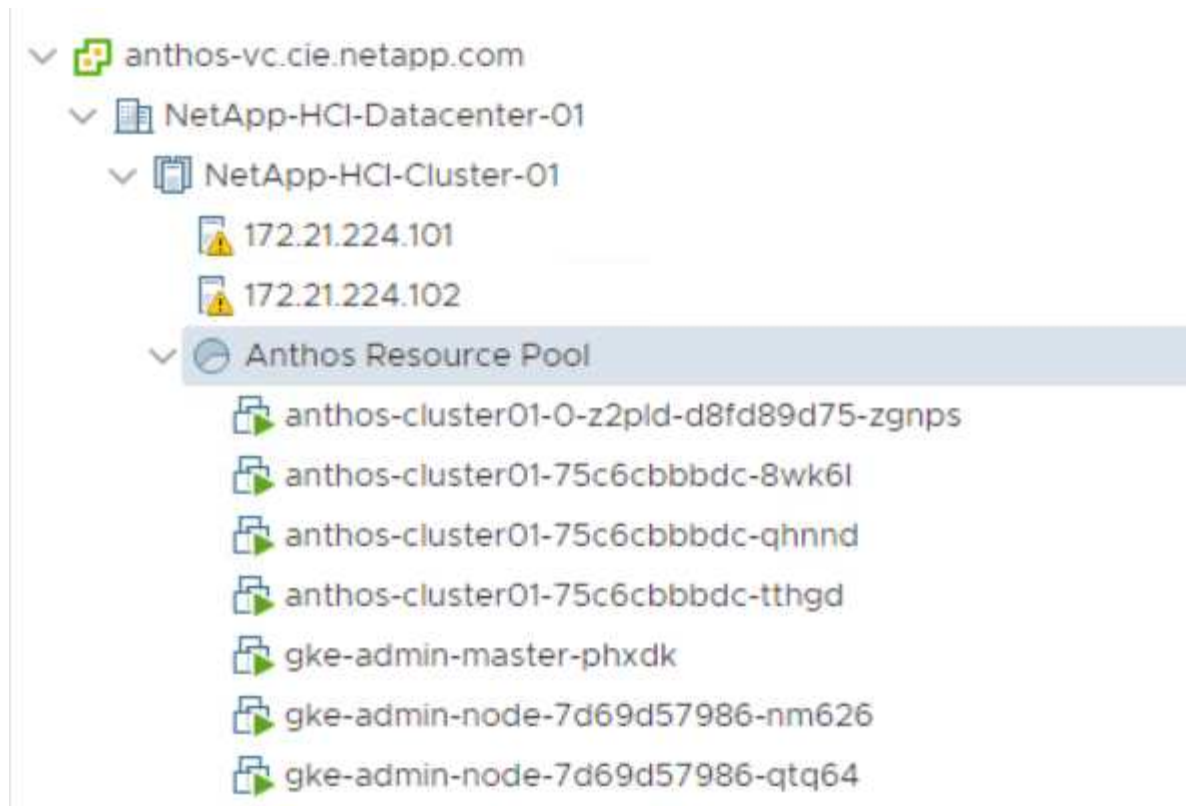
- Validation Category: Node IPs
  - [SUCCESS] ping (availability)

Some validations FAILED or SKIPPED. Check report above.
```

4. If all the checks succeed as expected, you can deploy this new user cluster in a manner very similar to the first cluster creation, referencing the `kubeconfig` file from the admin cluster.

```
ubuntu@Anthos-Admin-Workstation:~$ gkectl create cluster --config
anthos-cluster02-config.yaml --kubeconfig kubeconfig
```

5. As with the previous deployment, the process runs for several minutes and can be monitored on screen and in vCenter by watching the resource pool as the VMs populate. When complete, you should be able to see the new user cluster (four nodes).



6. You can access and execute commands against the deployed user cluster using the `kubectl` command line tool and the `kubeconfig` file generated by the process (stored in the working directory).

```
ubuntu@Anthos-Admin-Workstation:~$ kubectl get nodes --kubeconfig
anthos-cluster02-kubeconfig
NAME                                STATUS    ROLES    AGE      VERSION
anthos-cluster02-84744f5bd8-8rqk6  Ready    <none>   9m16s    v1.13.7-
gke.20
anthos-cluster02-84744f5bd8-f1786  Ready    <none>   9m28s    v1.13.7-
gke.20
anthos-cluster02-84744f5bd8-fnsmp  Ready    <none>   9m21s    v1.13.7-
gke.20
```

10. Enable access to the cluster with the GKE console

After clusters are deployed and registered with Google Cloud, they must be logged into with the Google Cloud console to be managed and to receive additional cluster details. The official procedure to gain access to Anthos user clusters after they are deployed is detailed [here](#).



The project and the specific user must be whitelisted to access on-premises clusters in the Google Cloud console and use Anthos on VMware services. If you are unable to see the clusters after they are deployed, you might need to open a support ticket with Google.

The non-whitelisted view looks like this:

The following figures provides a view of clusters.

To enable access to your user clusters using the GKE console, complete the following steps:

1. Create a `node-reader.yaml` file that allows you to access the cluster.

```
kind: clusterrole
apiVersion: rbac.authorization.k8s.io/v1
metadata:
  name: node-reader
rules:
- apiGroups: [""]
  resources: ["nodes"]
  verbs: ["get", "list", "watch"]
```

2. Apply this file to the cluster that you want to log into with the `kubectl` command.

```
ubuntu@Anthos-Admin-Workstation:~$ kubectl apply -f node-reader.yaml
--kubeconfig anthos-cluster01-kubeconfig
clusterrole.rbac.authorization.k8s.io/node-reader created
```

3. Create a Kubernetes service account (KSA) that you can use to log in. Name this account after the user that uses this account to log into the cluster.

```
ubuntu@Anthos-Admin-Workstation:~$ kubectl create serviceaccount netapp-
user --kubeconfig anthos-cluster01-kubeconfig
serviceaccount/netapp-user created
```

4. Create cluster role-binding resources to bind both the view and newly created node-reader roles to the newly created KSA.

```
ubuntu@Anthos-Admin-Workstation:~$ kubectl create clusterrolebinding
netapp-user-view --clusterrole view --serviceaccount default:netapp-user
--kubeconfig anthos-cluster01-kubeconfig
clusterrolebinding.rbac.authorization.k8s.io/netapp-user-view created
ubuntu@Anthos-Admin-Workstation:~$ kubectl create clusterrolebinding
netapp-user-node-reader --clusterrole node-reader -
--serviceaccount default:netapp-user --kubeconfig anthos-cluster01-
kubeconfig
clusterrolebinding.rbac.authorization.k8s.io/netapp-user-node-reader
created
```


5. If you need to extend permissions further, you can grant the KSA user a role with cluster admin permissions in a similar manner.

```
ubuntu@Anthos-Admin-Workstation:~$ kubectl create clusterrolebinding
netapp-user-admin --clusterrole cluster-admin --serviceaccount
default:netapp-user --kubeconfig anthos-cluster01-kubeconfig
clusterrolebinding.rbac.authorization.k8s.io/netapp-user-admin created
```

6. With the KSA account created and assigned with correct permissions, you can create a bearer token to allow access with the GKE Console. To do so, set a system variable for the secret name, and pass that variable through a `kubectl` command to generate the token.

```
ubuntu@Anthos-Admin-Workstation:~$ SECRET_NAME=$(kubectl get serviceaccount netapp-user --kubeconfig anthos-cluster01-kubeconfig -o jsonpath='{$.secrets[0].name}')  
ubuntu@Anthos-Admin-Workstation:~$ kubectl get secret ${SECRET_NAME} --kubeconfig anthos-cluster01-kubeconfig -o jsonpath='{$.data.token}' | base64 -d  
eyJhbGciOiJSUzI1NiIsImtpZCI6IiJ9.eyJpc3MiOiJrdWJlcm5ldGVzL3NlcnZpY2VhY2NvdW50Iiwia3ViZXJuZXRlcy5pby9zZXJ2aWNlYWNBjB3VudC9uYWwlc3BhY2UiOiJkZWZhdmx0Iiwia3ViZXJuZXRlcy5pby9zZXJ2aWNlYWNBjB3VudC9zZWNYZXQubmFtZSI6Im5ldGFwcC1lc2VyLXRVa2VuLWJxd3piIiwia3ViZXJuZXRlcy5pby9zZXJ2aWNlYWNBjB3VudC9zZXJ2aWNlLWFfjY291bnQubmFtZSI6Im5ldGFwcC1lc2VyIiwia3ViZXJuZXRlcy5pby9zZXJ2aWNlYWNBjB3VudC9zZXJ2aWNlLWFfjY291bnQudWlkIjoibmFtZS0xMWWVhLWEzMGUtNmFiZmRlYyJyZWNBbmIiwic3ViIjoic3lzdGVtOnNlcnZpY2VhY2NvdW50OmRlZmF1bHQ6bmV0YXBBLXVzZXIifQ.YrHn4kYlb3gwxVKCLyo7p6Jl1f7mwwIgZqNw9eTvIkt4PfyR4IJHxQwawnJ4T6RljIFcbVSQwvWi1yGuTJ98lADdcwtFXHoEfMcOa6SIn4OMVw1d5BGloaesN8150VCK3xES2DHAmLexFBqhVBgckZ0E4fZDvn4EhYvtFVpKlRbSyAE-DHD59P1bIgPdIoikREgbOddKdMn6XTVsuiP4V4tVKhkctcdRNRAuw6cFDY1fPol3BFHr2aNBIe6lFLkUqvQN-  
9nmD63JGdHL4hfXu6PPDxc9By6LgOW0nyaH4__gexy4uIa6lfNLKV2SKe4_gAN41ffOCKe4Tq8sa6zMo-8g
```

7. With this token, you can visit the [Google Cloud Console](#) and log in to the cluster by clicking the login button and pasting in the token.

Log in to cluster

Choose the method you want to use for authentication to the cluster

☒ Token

`›xc9By6LgOW0nyaH4__gexy4ula61fNLKV2SKe4_gAN41ffOCKe4Tq8sa6zMo-8g`

☐ Basic authentication

☐ Authenticate with Identity Provider configured for the cluster

[CLOSE](#) [LOGIN](#)

1. After login is complete, you see a green check mark next to the cluster name, and information is displayed about the physical environment. Clicking the cluster name displays more verbose information.

Next: [Install and Configure NetApp Trident Storage Provisioner.](#)

11. Install and configure NetApp Trident storage provisioner

Trident is a storage orchestrator for containers. With Trident, microservices and containerized applications can take advantage of enterprise-class storage services provided by the full NetApp portfolio of storage systems for persistent storage mounts. Depending on an application's requirements, Trident dynamically provisions storage for ONTAP-based products such as NetApp AFF and FAS systems and Element storage systems like NetApp SolidFire and NetApp HCI.

To install Trident on the deployed user cluster and provision a persistent volume, complete the following steps:



The following instructions are screen-capped from a Trident 21.01 install, but the same steps to manually deploy the Trident Operator also apply to the current 21.04 release.

1. Download the installation archive to the admin workstation and extract the contents. The current version of Trident is 21.04, which can be downloaded [here](#).

```
ubuntu@gke-admin-ws-200915-151421:~$ wget
https://github.com/NetApp/trident/releases/download/v21.01.0/trident-
installer-21.01.0.tar.gz
--2021-02-17 12:40:42--
https://github.com/NetApp/trident/releases/download/v21.01.0/trident-
installer-21.01.0.tar.gz
Resolving github.com (github.com)... 140.82.121.4
Connecting to github.com (github.com)|140.82.121.4|:443... connected.
HTTP request sent, awaiting response... 302 Found
Location: https://github-
releases.githubusercontent.com/77179634/0a63b600-6273-11eb-98df-
3d542851f6ff?X-Amz-Algorithm=AWS4-HMAC-SHA256&X-Amz-
```



```

Credential=AKIAIWNJYAX4CSVEH53A%2F20210217%2Fus-east-
1%2Fs3%2Faws4_request&X-Amz-Date=20210217T173945Z&X-Amz-Expires=300&X-
Amz-
Signature=58f26bcac7eeee64673a84d46696490acec357b97a651af42653f973b778ee
88&X-Amz-
SignedHeaders=host&actor_id=0&key_id=0&repo_id=77179634&response-
content-disposition=attachment%3B%20filename%3Dtrident-installer-
21.01.0.tar.gz&response-content-type=application%2Foctet-stream
[following]
--2021-02-17 12:40:43--  https://github-
releases.githubusercontent.com/77179634/0a63b600-6273-11eb-98df-
3d542851f6ff?X-Amz-Algorithm=AWS4-HMAC-SHA256&X-Amz-
Credential=AKIAIWNJYAX4CSVEH53A%2F20210217%2Fus-east-
1%2Fs3%2Faws4_request&X-Amz-Date=20210217T173945Z&X-Amz-Expires=300&X-
Amz-
Signature=58f26bcac7eeee64673a84d46696490acec357b97a651af42653f973b778ee
88&X-Amz-
SignedHeaders=host&actor_id=0&key_id=0&repo_id=77179634&response-
content-disposition=attachment%3B%20filename%3Dtrident-installer-
21.01.0.tar.gz&response-content-type=application%2Foctet-stream
Resolving github-releases.githubusercontent.com (github-
releases.githubusercontent.com)... 185.199.111.154, 185.199.108.154,
185.199.109.154, ...
Connecting to github-releases.githubusercontent.com (github-
releases.githubusercontent.com)|185.199.111.154|:443... connected.
HTTP request sent, awaiting response... 200 OK
Length: 38527217 (37M) [application/octet-stream]
Saving to: `trident-installer-21.01.0.tar.gz'

100%[=====
=====>] 38,527,217  84.9MB/s
in 0.4s

2021-02-17 12:40:44 (84.9 MB/s) - `trident-installer-21.01.0.tar.gz'
saved [38527217/38527217]

```

2. Extract the Trident install from the downloaded bundle.

```

ubuntu@gke-admin-ws-200915-151421:~$ tar -xf trident-installer-
21.01.0.tar.gz
ubuntu@gke-admin-ws-200915-151421:~$ cd trident-installer

```

3. First set the location of the user cluster's kubeconfig file as an environment variable so that you don't have to reference it, because Trident has no option to pass this file.


```
ubuntu@gke-admin-ws-200915-151421:~/trident-installer$ export
KUBECONFIG=~/.anthos-cluster01-kubeconfig
```

4. The `trident-installer` directory contains manifests for defining all the required resources. Using the appropriate manifests, create the `TridentOrchestrator` custom resource definition.

```
ubuntu@gke-admin-ws-200915-151421:~/trident-installer$ kubectl create -f
deploy/crds/trident.netapp.io_tridentorchestrators_crd_post1.16.yaml
customresourcedefinition.apiextensions.k8s.io/tridentorchestrators.tride
nt.netapp.io created
```

5. If a `Trident` namespace does not exist, create one in your cluster using the provided manifest.

```
ubuntu@gke-admin-ws-200915-151421:~/trident-installer$ kubectl apply -f
deploy/namespace.yaml
namespace/trident created
```

6. Create the resources required for the `Trident` operator deployment, such as a `ServiceAccount` for the operator, a `ClusterRole` and `ClusterRoleBinding` to the `ServiceAccount`, a dedicated `PodSecurityPolicy`, or the operator itself.

```
ubuntu@gke-admin-ws-200915-151421:~/trident-installer$ kubectl create -f
deploy/bundle.yaml
serviceaccount/trident-operator created
clusterrole.rbac.authorization.k8s.io/trident-operator created
clusterrolebinding.rbac.authorization.k8s.io/trident-operator created
deployment.apps/trident-operator created
podsecuritypolicy.policy/tridentoperatorpods created
```

7. You can check the status of the operator after it's deployed with the following commands:

```
ubuntu@gke-admin-ws-200915-151421:~/trident-installer$ kubectl get
deployment -n trident
NAME                READY    UP-TO-DATE    AVAILABLE    AGE
trident-operator    1/1      1              1             54s
ubuntu@gke-admin-ws-200915-151421:~/trident-installer$ kubectl get pods
-n trident
NAME                                READY    STATUS    RESTARTS    AGE
trident-operator-5c8bbf6754-h957z    1/1      Running    0            68s
```

8. With the operator deployed, we can now use it to install `Trident`. This requires creating a `TridentOrchestrator`.


```

ubuntu@gke-admin-ws-200915-151421:~/trident-installer$ kubectl create -f
deploy/crds/tridentorchestrator_cr.yaml
tridentorchestrator.trident.netapp.io/trident created
ubuntu@gke-admin-ws-200915-151421:~/trident-installer$ kubectl describe
torc trident
Name:          trident
Namespace:
Labels:        <none>
Annotations:   <none>
API Version:   trident.netapp.io/v1
Kind:          TridentOrchestrator
Metadata:
  Creation Timestamp:  2021-02-17T18:25:43Z
  Generation:         1
  Managed Fields:
    API Version:      trident.netapp.io/v1
    Fields Type:      FieldsV1
    fieldsV1:
      f:spec:
        .:
        f:debug:
        f:namespace:
  Manager:           kubectl
  Operation:         Update
  Time:              2021-02-17T18:25:43Z
  API Version:      trident.netapp.io/v1
  Fields Type:      FieldsV1
  fieldsV1:
    f:status:
      .:
      f:currentInstallationParams:
        .:
        f:IPv6:
        f:autosupportHostname:
        f:autosupportImage:
        f:autosupportProxy:
        f:autosupportSerialNumber:
        f:debug:
        f:enableNodePrep:
        f:imagePullSecrets:
        f:imageRegistry:
        f:k8sTimeout:
        f:kubeletDir:
        f:logFormat:
        f:silenceAutosupport:

```



```

      f:tridentImage:
      f:message:
      f:namespace:
      f:status:
      f:version:
    Manager:      trident-operator
    Operation:    Update
    Time:         2021-02-17T18:25:43Z
    Resource Version: 14836643
    Self Link:    /apis/trident.netapp.io/v1/tridentorchestrators/trident
    UID:         0e5f2c3b-6ca2-4b85-8453-0382e1426160
  Spec:
    Debug:      true
    Namespace:  trident
  Status:
    Current Installation Params:
      IPv6:
      Autosupport Hostname:
      Autosupport Image:
      Autosupport Proxy:
      Autosupport Serial Number:
      Debug:
      Enable Node Prep:
      Image Pull Secrets:      <nil>
      Image Registry:
      k8sTimeout:
      Kubelet Dir:
      Log Format:
      Silence Autosupport:
      Trident Image:
    Message:      Installing Trident
    Namespace:    trident
    Status:       Installing
    Version:
  Events:
    Type      Reason      Age      From      Message
    ----      -
    Normal    Installing  23s      trident-operator.netapp.io  Installing
    Trident
    Normal    Installed  15s      trident-operator.netapp.io  Trident
    installed

```

9. You can verify that Trident is successfully installed by checking the pods that are running in the namespace or by using the tridentctl binary to check the installed version.


```
ubuntu@gke-admin-ws-200915-151421:~/trident-installer$ kubectl get pod
-n trident
NAME                                READY   STATUS    RESTARTS   AGE
trident-csi-2cp7x                   2/2     Running   0           4m16s
trident-csi-2xr5h                   2/2     Running   0           4m16s
trident-csi-bnwvh                   2/2     Running   0           4m16s
trident-csi-d6cfc6bb-lxm2p          6/6     Running   0           4m16s
trident-operator-5c8bbf6754-h957z   1/1     Running   0           8m55s

ubuntu@gke-admin-ws-200915-151421:~/trident-installer$ ./tridentctl -n
trident version
+-----+
| SERVER VERSION | CLIENT VERSION |
+-----+
| 21.01.1        | 21.01.1        |
+-----+
```

10. The next step in enabling Trident integration with the NetApp HCI solution and Anthos is to create a backend that enables communication with the storage system. NetApp has been validated for several different protocols through the Anthos-ready partner storage validation program. This allows NetApp Trident to provide support in Anthos environments for NFS through our ONTAP platforms and iSCSI from both the ONTAP and Element storage used in NetApp HCI.



A NetApp HCI platform deploys with NetApp Element storage by default. In this guide we configure a backend for this system specifically. In addition to this, a customer can choose to connect to a remote ONTAP storage system or deploy an ONTAP Select software-defined storage system as a virtual appliance in VMware vSphere to provide additional NFS and iSCSI services. The configuration of each of these additional storage backends is beyond the scope of this guide.

11. There are sample backend files available in the downloaded installation archive in the `sample-input` folder. Copy `backend-solidfire.json` to your working directory and edit it to provide information detailing the storage system environment. For Element-based iSCSI connections, copy and edit the `backend-solidfire.json` file.

```
ubuntu@gke-admin-ws-200915-151421:~/trident-installer$ cp sample-
input/backend-solidfire.json ./
ubuntu@gke-admin-ws-200915-151421:~/trident-installer$ $ vi backend-
solidfire.json
```

- a. Edit the user, password, and MVIP value on the EndPoint line.
- b. Edit the SVIP value.


```
{
  "version": 1,
  "storageDriverName": "solidfire-san",
  "Endpoint": "https://trident:password@172.21.224.150/json-
rpc/8.0",
  "SVIP": "10.63.172.100:3260",
  "TenantName": "trident",
  "Types": [{"Type": "Bronze", "Qos": {"minIOPS": 1000, "maxIOPS":
2000, "burstIOPS": 4000}},
            {"Type": "Silver", "Qos": {"minIOPS": 4000, "maxIOPS":
6000, "burstIOPS": 8000}},
            {"Type": "Gold", "Qos": {"minIOPS": 6000, "maxIOPS":
8000, "burstIOPS": 10000}}]
}
```

12. With this backend file in place, run the following command to create your first backend.

```
ubuntu@gke-admin-ws-200915-151421:~/trident-installer$ ./tridentctl -n
trident create backend -f backend.json
+-----+-----+
+-----+-----+-----+-----+
|      NAME              | STORAGE DRIVER |              UUID
| STATE  | VOLUMES |
+-----+-----+
+-----+-----+-----+-----+
| solidfire-backend | solidfire-san | a5f9e159-c8f4-4340-a13a-
c615fef0f433 | online |      0 |
+-----+-----+
+-----+-----+-----+-----+
```

13. With the backend created, you must next create a storage class. Just as with the backend, there is a sample storage class file that can be edited for the environment available in the sample-inputs folder. Copy it to the working directory and make necessary edits to reflect the backend created.

```
ubuntu@gke-admin-ws-200915-151421:~/trident-installer$ cp sample-
input/storage-class-csi.yaml.template ./storage-class-basic.yaml
ubuntu@gke-admin-ws-200915-151421:~/trident-installer$ vi storage-class-
basic.yaml
```

14. The only edit that must be made to this file is to define the `backendType` value to the name of the storage driver from the newly created backend. Also note the `name-field` value that must be referenced in a later step.


```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: basic-csi
provisioner: csi.trident.netapp.io
parameters:
  backendType: "solidfire-san"
```

15. Run the `kubectl` command to create the storage class.

```
ubuntu@gke-admin-ws-200915-151421:~/trident-installer$ kubectl create -f
sample-input/storage-class-basic.yaml
```

16. With the storage class created, you must then create the first persistent volume claim (PVC). There is a sample `pvc-basic.yaml` file that can be used to perform this action located in `sample-inputs` as well. The only edit that must be made to this file is ensuring that the `storageClassName` field matches the one just created.

```
ubuntu@gke-admin-ws-200915-151421:~/trident-installer$ vi sample-
input/pvc-basic.yaml
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
  name: basic
spec:
  accessModes:
    - ReadWriteOnce
  resources:
    requests:
      storage: 1Gi
  storageClassName: basic-csi
```

17. Create the PVC by issuing the `kubectl` command. Creation can take some time depending on the size of the backing volume being created, so you can watch the process as it completes.


```
ubuntu@gke-admin-ws-200915-151421:~/trident-installer$ kubectl create -f
sample-input/pvc-basic.yaml
```

```
ubuntu@gke-admin-ws-200915-151421:~/trident-installer$ kubectl get pvc
--watch
```

NAME	STATUS	VOLUME	CAPACITY
ACCESS MODES	STORAGECLASS	AGE	
basic	Pending		
basic	1s		
basic	Pending	pvc-2azg0d2c-b13e-12e6-8d5f-5342040d22bf	0
basic	5s		
basic	Bound	pvc-2azg0d2c-b13e-12e6-8d5f-5342040d22bf	1Gi
RWO	basic	7s	

[Next: Reference videos.](#)

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.