



Technical Report

NetApp Private Storage for Google Cloud Platform

Solution Architecture and Deployment Guide

Mark Beaupre and Alim Karim, NetApp
January 2018 | TR-4653

Abstract

This document describes the architecture of NetApp Private Storage for Google Cloud Platform (GCP).

TABLE OF CONTENTS

1	NetApp Private Storage for GCP Solution Overview	3
1.1	Assumptions	3
1.2	Use Case Overview	3
1.3	Technical Overview	3
2	NetApp Private Storage for GCP Solution Architecture	4
2.1	Solution Architecture Components	4
2.2	Solution Architecture Diagram	7
2.3	Solution Architecture Security Elements	8
3	NetApp Private Storage for GCP Deployment Overview	9
3.1	Planning: Preinstallation and Site Preparation	9
3.2	Install Equipment in the Equinix Data Center	12
3.3	Provision Google Cloud Dedicated Interconnect	13
3.4	Set Up the Customer Network Switch	19
3.5	Configure NetApp Storage	20
3.6	Test Connections and Protocol Access	20
3.7	Performance Testing Guidelines	24
	Where to Find Additional Information	25
	Version History	25

LIST OF TABLES

Table 1)	NetApp Private Storage IP address plan (sample)	11
----------	---	----

LIST OF FIGURES

Figure 1)	NPS for Google Cloud Dedicated Interconnect network architecture	5
Figure 2)	NetApp Private Storage for GCP solution architecture	8

1 NetApp Private Storage for GCP Solution Overview

This document describes the storage architecture of the NetApp Private Storage for the Google Cloud Platform (GCP) solution and provides procedures for deploying and testing the solution.

1.1 Assumptions

The document assumes that the reader has working knowledge of the following technologies:

- GCP services
- NetApp storage administration
- Network administration
- Windows and/or Linux administration

1.2 Use Case Overview

The NetApp Private Storage for GCP solution is a cloud-connected storage architecture that enables enterprises to build an agile cloud infrastructure that combines the scalability and flexibility of the GCP cloud with the control and performance of NetApp® storage.

NetApp storage is deployed at an Equinix colocation data center where the Google Cloud Interconnect (GCI) dedicated service is available, and the NetApp storage is connected to GCP compute resources through the GCP GCI network service.

Typical use cases for the NetApp Private Storage for GCP solution include:

- Development and testing
- Big data analytics
- Oracle, Microsoft SQL Server, and SAP primary workloads
- Disaster recovery (DR)
- Data with compliance requirements
- Data center migration and consolidation

For full information about NetApp Private Storage use cases, see [NVA-0009: NetApp Private Storage for Cloud](#).

1.3 Technical Overview

The NetApp Private Storage for GCP solution combines computing resources from GCP with NetApp storage deployed at a GCP Dedicated Interconnect location. Connectivity from the NetApp storage to the GCP cloud is made possible by the GCP Dedicated Interconnect network service.

At the time of the writing of this document (December 2017), customers who deploy the NetApp Private Storage for GCP solution at Equinix can use 10Gb/sec GCI connections provisioned manually by cross-connect. Up to eight 10Gb/sec GCI connections can be bonded to provide up to 80Gb/sec of bandwidth between the GCP cloud and NetApp storage.

In the GCP Dedicated Interconnect data center, the customer provides network equipment (switch or router) and NetApp storage systems. Virtual machines (VMs) in the GCP cloud connect to the NetApp storage through IP-based storage protocols (iSCSI, CIFS, or NFS). In addition, MPLS or point-to-point VPN network resources can be used to provide connectivity between GCP regions as well as connectivity to on-premises data centers.

2 NetApp Private Storage for GCP Solution Architecture

This section describes the components of the solution architecture and explains the data security elements that are included.

2.1 Solution Architecture Components

The solution architecture consists of the following components:

- Google Compute Engine
- Google Virtual Private Cloud (VPC)
- Google Cloud Dedicated Interconnect
- Equinix colocation data center (dedicated interconnect location)
- Border Gateway Protocol (BGP)
- Customer-provided layer 3 network equipment
- NetApp storage (AFF, FAS, FlexArray® storage virtualization software, and E-series)

Google Compute Engine

Google Compute Engine is a cloud service that provides resizable computing capacity in the cloud.

Locations

The Google Compute Engine service is available on a per-GCP region basis. Each GCP region is tied to a specific geographic location.

Virtual machines (VMs) used with the NetApp Private Storage for GCP solution can be deployed through the Google Cloud Platform Console web interface. Advanced GCP users can programmatically deploy VMs through APIs and scripts that use the `gcloud` command line tool.

For more information about locations where the Google Compute Engine service is available, see the [Google Compute Engine Product Information](#).

Machine Types

Google Compute Engine VMs have various machine types that support the customer's compute needs. Each machine type is a combination of CPU, memory, storage, and network bandwidth.

Machine types can be predefined or custom. Predefined machine types offer predefined VM configurations of CPU, memory, storage, and network bandwidth. Custom machine types allow customers to configure the VM to a workload.

For more information about custom machine types, see the [Custom Machine Type documentation](#).

Available Operating Systems

In addition to different machine types, Google Compute Engine VMs can run different operating systems, including Windows and Linux. For a list of operating systems available in GCP, refer to the [Google Compute Engine Images documentation](#).

For each operating system and application type, you can validate version compatibility with the NetApp client software and ONTAP® software version through the [NetApp Interoperability Matrix Tool](#).

Note: This site requires a NetApp Support account login.

Google Virtual Private Cloud

The Google VPC service provides isolated RFC 1918 IPv4 address ranges (10.0.0.0/8, 172.16.0.0/12, and 192.168.0.0/16) in the Compute Engine VMs that can be deployed.

A VPC can use a predefined IP address range, called auto-mode creation. Alternatively, a VPC can use a customized IP address range, subnets, and firewall rules; this is called VPC custom-mode creation.

A VPC can be connected to the customer's network, located in the Google Cloud Interconnect data center, or to on-premises customer networks through a point-to-point VPN.

In the GCP, a VPC is a global resource, but subnets are regional resources. VMs are deployed into VPC subnets.

For more information, see the [GCP VPC documentation](#).

Google Cloud Dedicated Interconnect

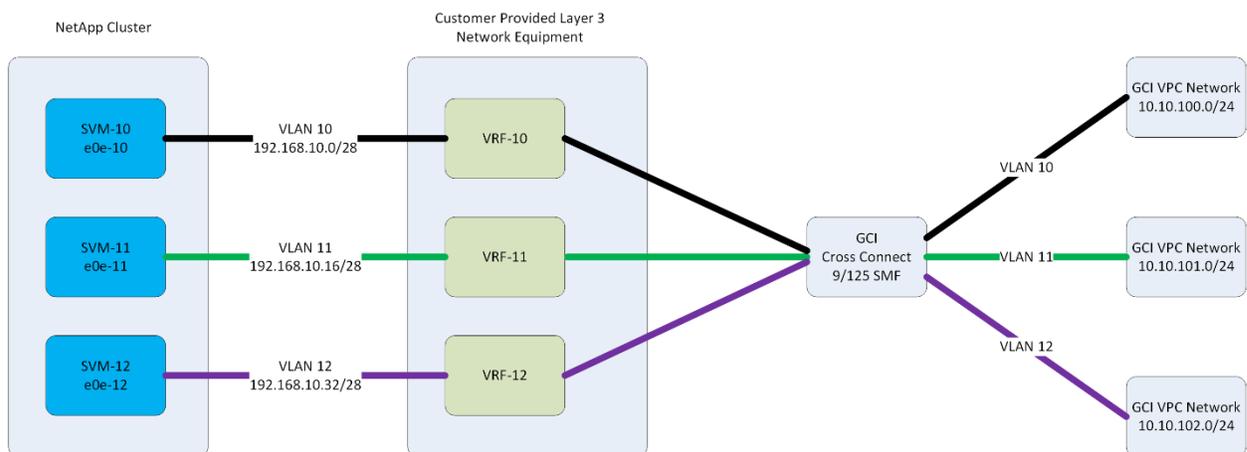
Google Cloud Dedicated Interconnect is used to connect a customer's infrastructure to the Google Cloud Platform.

In the NPS solution, Google Cloud Dedicated Interconnect is used to establish a dedicated network connection between the customer-provided network equipment in the Google Cloud Interconnect data center and the GCP VPC. Dedicated Interconnect supports the use of industry-standard 802.1Q virtual local area networks (VLANs) as well as 802.3ad link aggregation control protocol (LACP).

Multiple VLANs can be used to connect to multiple VPCs over the same physical Dedicated Interconnect network connection. The use of VLANs supports the segregation of network traffic from the storage, through the network equipment, and through to the VPC.

Figure 1 shows the high-level NPS for Google Cloud Dedicated Interconnect network architecture.

Figure 1) NPS for Google Cloud Dedicated Interconnect network architecture.



A single Dedicated Interconnect network connection is 10Gb Ethernet (10GbE). Customers can order up to eight Dedicated Interconnect network connections for a total of 80Gbps of bandwidth to one or more VPCs. When using multiple physical connections, the links are aggregated using LACP.

Dedicated Interconnect supports redundant network connections and can be combined with LACP.

For more information, see the [Google Cloud Dedicated Interconnect documentation](#).

Equinix Colocation Data Center (Dedicated Interconnect Location)

GCP Point of Presence

Google Cloud Dedicated Interconnect locations provide connectivity to the GCP through Dedicated Interconnect network connections. Equinix and other colocation providers have GCP Point of Presence in their data centers. This feature offers private connectivity to the GCP and does not go over the internet.

Note: Each Dedicated Interconnect location supports a subset of GCP regions.

For a list of locations and their supported regions, see the [Dedicated Interconnect Colocation Facility Locations documentation](#).

Some, but not all, Equinix data centers are close to the GCP region data centers; therefore, the latencies between Equinix and GCP can vary. NetApp recommends validating the latency of the network connectivity to the GCP before deploying workloads into the NetApp Private Storage for GCP solution.

Physical Security

Equinix data centers offer a secure, highly available environment for the customer-owned NetApp storage and network equipment for the NetApp Private Storage for GCP solution. Equinix provides a high degree of physical security.

Customers have the option of deploying their storage into dedicated secure cages or into secure cabinets in shared cages.

For more information about Equinix physical security, see the Equinix [Physical Security](#) webpage, or contact your Equinix account team.

Operational Security

Equinix data centers have a minimum N+1 power and cooling system redundancy. Many Equinix data centers have N+2 power and cooling system redundancy.

For more information about Equinix operational reliability, see the Equinix [Operational Reliability](#) webpage, or contact your Equinix account team.

Border Gateway Protocol

BGP is used to support network routing between the GCP VPC and the network in the GCI location data center over the Dedicated Interconnect network connection.

The network in the Equinix colocation data center is directly connected to the customer-provided layer 3 network equipment. The BGP configuration advertises local network routes to the VPC network over the GCI network connection. It also receives the BGP advertisements from the VPC network over the Dedicated Interconnect network connection.

Customer-Provided Layer 3 Network Equipment

The customer-provided network equipment is in the same GCI location data center as the NetApp storage. NetApp does not certify specific network equipment to be used in the solution; however, the network equipment must support the following features:

- BGP
- At least one 9/125 single-mode fiber (SMF) (10Gbps) port
- 1000BASE-T Ethernet ports
- 802.1Q VLAN tags
- 802.3ad LACP

The following features are optional:

- Virtual routing and forwarding (VRF)
- Redundant network switches
- Redundant 9/125 SMF (10Gbps) ports
- 40GbE ports (dependent on the NetApp storage)

Required Features

As previously noted, BGP is used to route network traffic between the local network in the Dedicated Interconnect data center and the GCP VPC network.

Dedicated Interconnect requires a minimum of one physical connection (9/125 SMF) from the customer-owned network equipment to the GCP.

1000BASE-T network ports on the switch provide network connectivity from the NetApp storage. Although these ports can be used for data, NetApp recommends using 1GbE ports for node management and out-of-band management.

802.1Q VLAN tags are used by GCI VLANs to segregate network traffic on the same physical network connection to different VPCs.

802.3ad LACP is used by GCI to aggregate physical network links to increase the amount of bandwidth between the NetApp storage and the Google Cloud Platform.

Optional Features

VRF is a way of segregating routing on a single layer 3 switch. Isolating routing is useful in multitenant environments to restrict network access between tenancies.

Redundant network switches protect against a loss of Dedicated Interconnect service caused by switch failure.

Note: For information about configuring redundant network switches, consult your network equipment vendor's documentation.

Redundant 9/125 SMF ports protect against a loss of Dedicated Interconnect service caused by a port or cable failure.

Some NetApp storage platforms support 40GbE connectivity natively. Using switches that can support 40Gb to 4 x 10Gbp connectivity (hydra cables) allows these platforms to be used with 10GbE switches. Alternatively, use network equipment that have 40GbE ports.

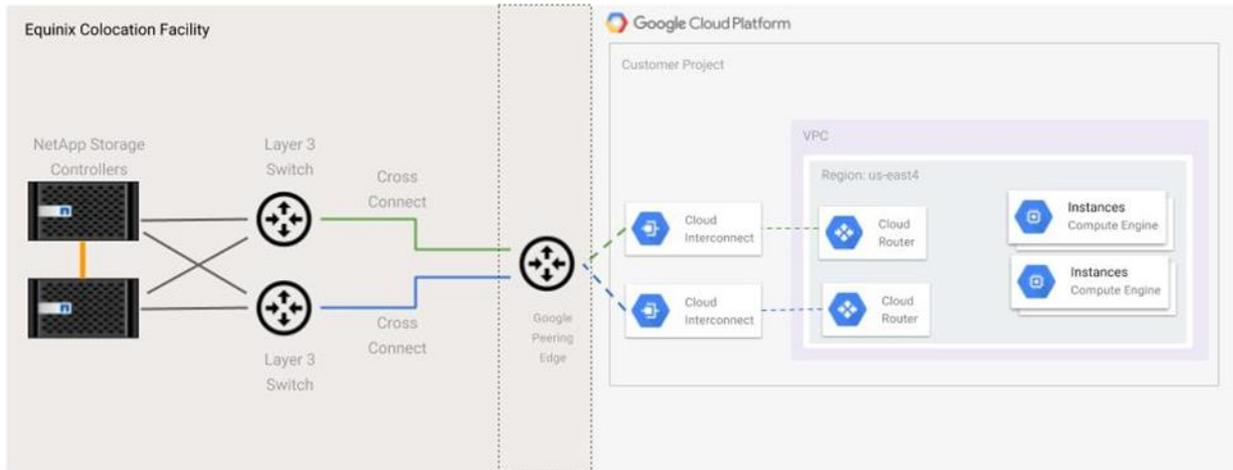
NetApp Storage

The NetApp storage platforms that can be used in the NetApp Private Storage solution are FAS and AFF, E-series, SolidFire® (Element® OS 9.0 or later), and StorageGRID® Webscale.

2.2 Solution Architecture Diagram

Figure 2 shows the architecture of the NetApp Private Storage for GCP solution.

Figure 2) NetApp Private Storage for GCP solution architecture.



2.3 Solution Architecture Security Elements

NetApp Private Storage for GCP allows customers to store their data on NetApp storage that they own or control, so that they can maintain compliance and enhance security of their persistent data.

The solution contains the following security-related elements:

- GCP VPC network
- Dedicated Interconnect network
- Physical security for the Equinix colocation data center
- NetApp storage encryption
- NetApp volume encryption

GCP Virtual Private Cloud Network

GCP VPC provides network isolation for the resources (VMs, services, and so forth) that are provisioned in it. Google Compute Engine VMs provisioned in a VPC can communicate with each other within the network. Resources external to the VPC do not have access to the resources in the VPC.

GCP VPCs can be accessed securely through a site-to-site VPN or through a GCI network connection at a GCI location data center such as Equinix.

Dedicated Interconnect Network

The Google Cloud Dedicated Interconnect network is a dedicated, private, secure network connection that does not traverse the Internet. GCP connects to NetApp storage in the Equinix colocation data center through physical cross-connects that are not shared with other customers.

Physical Security for Equinix Colocation Data Center

Equinix provides state-of-the-art physical security at all of its data centers where Google Cloud Dedicated Interconnect service is available. The data centers have security guards and security systems to provide video surveillance. The security systems have biometric hand scanners combined with mantrap interlocks to restrict access to authorized personnel only.

For more information about Equinix physical security, see the Equinix [Physical Security](#) webpage.

NetApp Storage Encryption

NetApp Storage Encryption is the NetApp implementation of full-disk encryption (FIPS 140-2 level 2, AES-256 encryption) that uses self-encrypting drives from leading vendors, allowing data on NetApp storage to be fully encrypted while maintaining storage efficiency and performance.

For more information, see [NetApp Storage Encryption](#).

NetApp Volume Encryption

NetApp Volume Encryption delivers software-based AES-256 encryption of individual volumes—each with its own unique key on any type of data on FAS and AFF storage systems.

For more information, see [NetApp Volume Encryption Datasheet](#).

3 NetApp Private Storage for GCP Deployment Overview

This section describes the standard deployment methodology for NetApp Private Storage for GCP. However, because no two customer environments are the same, NetApp has delivery partners who specialize in deploying NetApp Private Storage solutions. These partners are experienced and can help make your deployment a success. For information about NetApp Private Storage delivery partners, contact your NetApp account representative.

The high-level deployment workflow consists of the following phases and tasks:

1. Planning: Preinstallation and site preparation
2. Deployment:
 - a. Installing the equipment in the Equinix Data Center
 - b. Provisioning the Google Cloud Dedicated Interconnect
 - c. Setting up the GCP VPC Network
 - d. Setting up the Customer Network Switch
 - e. Configuring the NetApp storage
3. Validation: Testing connections and protocols

3.1 Planning: Preinstallation and Site Preparation

The following preinstallation and site preparation tasks take place during the planning phase of the NPS for GCP workflow.

To establish colocation power and space requirements, perform the following tasks, which are detailed in the following subsections:

1. Establish colocation power and space requirements.
2. Order space and power.
3. Order the necessary network, storage, and rack hardware.
4. Create an IP address plan.
5. Obtain Google and Equinix Customer Portal accounts.
6. Create an inbound shipment request through the Equinix Customer Portal.
7. Install Google Cloud Command Line Interface (`gcloud`) tools (optional).

In the final step of the planning phase, you validate that all preinstallation and site preparation tasks have been completed and the workflow is ready to move to the next phase.

Establish Colocation Power and Space Requirements

Use the [NetApp Hardware Universe](#) to determine the power and space requirements of the NetApp storage used with the NetApp Private Storage for GCP solution. Contact your NetApp account team for more information about the power and space requirements of the storage you want to deploy with NPS.

Use the technical specifications for the network equipment used by the NetApp Private Storage for GCP solution. Contact your network switch vendor about the power and space requirements of the network equipment you want to deploy with NPS.

Order Space and Power

There are two types of colocation space at Equinix: shared and dedicated. Shared space is a secure cage containing secure cabinets used by multiple customers. Customers are required to use Equinix racks in a shared space configuration.

Dedicated space is a secure cage that is assigned to a single customer. The smallest dedicated cage consists of five cabinets. Customers can either use their own Equinix racks or bring their own.

NetApp recommends that customers use redundant power connections connected to separate power distribution units (PDUs) so that the NetApp Private Storage solution can survive the loss of a single power connection.

The typical power connection configuration used with NetApp Private Storage is 208V/30A single-phase AC power. The number of volts can vary from region to region.

Contact your Equinix account team for information about available space and power options in the Equinix data center where you want to deploy NPS.

Order Necessary Network, Storage, and Rack Hardware

If more than six ports of power are required on a PDU, you need to purchase a third-party PDU, or order additional power connections from Equinix. Equinix sells PDUs that fit well with their cabinets.

The Equinix cabinets are standard 42U, 4-post racks.

Contact your NetApp account team to make sure that you order the correct rail kits for the cabinets you are using.

If you are using a secure cabinet in a shared cage, order a top-of-rack demarcation panel to connect the network equipment to GCP. The demarcation panel should be 24-port SC optical.

Create IP Address Plan

The creation of the IP address plan for NetApp Private Storage is a crucial step. Use the data in Table 1 directly in the configuration of the NPS network. As a reminder, the unit of tenancy is an SVM connected to a GCP VPC network through a GCI network connection.

Table 1 is a sample IP address plan.

Table 1) NetApp Private Storage IP address plan (sample).

Tenant	Tenant VLAN	NPS SVM Network	BGP Peering Network	BGP Authentication Key	BGP ASN	GCP VPC Network	GCP Region
	Assigned by Google		Assigned by Google	N/A			
	Assigned by Google		Assigned by Google	N/A			
	Assigned by Google		Assigned by Google	N/A			
	Assigned by Google		Assigned by Google	N/A			

Where:

- **Tenant.** The name or description of the NPS tenant.
- **Tenant VLAN.** The VLAN number that the NPS tenant uses to connect the NetApp storage assigned to them to the GCP VPC over a GCI network VLAN (for example, 100). Currently, the VLANs are assigned by Google.
- **NPS SVM network.** The network CIDR used by the NetApp SVM logical network interfaces. The network is typically a private network CIDR (such as 192.168.100.0/28).
- **BGP Peering network.** The BGP peering network is a /29 network. Google assigns a peering network to test the physical connection. After Google confirms that the network connection is working, they automatically send the peering network to be used.
Note: The lower IP address number of the peering network is assigned to the layer 3 interface on the network equipment in Equinix and the higher number is assigned to the GCP cloud router.
- **BGP authentication key.** At the time of this publication, Dedicated Interconnect does not use shared BGP keys.
- **BGP ASN.** The BGP Autonomous System Number (ASN) is the unique number assigned to the network equipment in Equinix. The ASN is a private ASN number. Private ASN numbers range from 64,512 to 65,535.
- **GCP VPC network.** The network CIDR of the GCP VPC (such as 10.10.100.0/24).
Note: You can allow auto-creation of the VPC network CIDR.
- **GCP region.** The GCP region that the VPC is created and connected through Dedicated Interconnect.

Obtain Google and Equinix Customer Portal Accounts

If you do not have an GCP account already set up, go to the [Google Accounts](#) webpage to create one.

Contact your Equinix account team to get your account set up in the [Equinix Customer Portal](#).

Create an Inbound Shipment Request Through Equinix

Equinix physical security procedures require an inbound shipping request for any shipments sent to an Equinix data center. The shipping addresses for the data center (also known as IBX) are located in the [Equinix Customer Portal](#).

In the inbound shipment request, you must provide the shipper, shipment tracking number, number of items in the shipment, weight of items in the shipment, and the date that the shipment is expected to arrive at the IBX.

When shipping equipment to the Equinix data center, use this format for the address: s

Name of cage/suite
c/o Equinix
Address of data center

For more information about Equinix shipping and receiving procedures for IBX, see the [Equinix Customer Portal](#) or contact your Equinix Client Services manager.

Install GCP Command Line Interface (gcloud) Tools

The [GCP Cloud SDK documentation](#) has easy-to-follow instructions for setting up and installing the GCP CLI (`gcloud`) Tools on Windows, Linux, or Mac.

Note: The `gcloud` CLI is available automatically in the Google Cloud Shell. If you are using Google Cloud Shell, you do not need to install `gcloud` manually. .

3.2 Install Equipment in the Equinix Data Center

You can begin to install the equipment in the data center after the preinstallation and site preparation phase is complete.

To set up the data center, perform the following tasks, which are detailed in the following subsections:

1. Set up security access to the Equinix data center and cage.
2. Make sure that all required materials (hardware, software, accessories, and so on) are available on site.
3. Install the NetApp storage in the rack.
4. Install the customer-provided network equipment in the rack.

Set Up Security Access to the Equinix Data Center and Cage

Use the [Equinix Customer Portal](#) to create a Security Access Request for the Equinix IBX where the NPS solution is being deployed. The security access registration process includes a biometric scan, PIN assignment, and security card assignment (depending on the IBX). You need to take government-issued identification to the IBX.

Note: The name on the security access request must be identical to the government-issued identification. If the names don't match, Equinix security will not process the request.

After the security access process is complete, you can visit the Equinix IBX without needing an Equinix work visit request.

Make Sure That All Required Materials Are Available On Site

You can take inventory of your shipment in person, or you can have the Equinix SmartHands technicians inventory the shipment. If you want to have the technicians inventory your shipment, use the [Equinix Customer Portal](#) to create a SmartHands request.

Install NetApp Storage in the Rack

If you are using an Equinix cabinet in a shared cage, you can install NetApp storage yourself, or you can have a NetApp partner install it.

If you are using a dedicated Equinix cage, the racks in the cage must be installed. Use the [Equinix Customer Portal](#) to create an Equinix SmartHands request to have the racks installed.

If you are having a NetApp partner install the storage, use the [Equinix Customer Portal](#) to create a work visit request for the partner engineers. The engineers must bring government-issued identification, and the names on the work visit request must match the government-issued identification.

Due to Equinix safety rules, the PDUs in the rack must be connected to Equinix power by an Equinix SmartHands technician. Use the [Equinix Customer Portal](#) to create a SmartHands request to connect the PDUs.

Install Customer-Provided Network Equipment in the Rack

The network equipment can be installed at the same time as the NetApp storage.

If the network equipment is installed at a different time, use the [Equinix Customer Portal](#) to create a work visit request for the partner engineers. The engineers must bring government-issued identification, and the names on the work visit request must match the government-issued identification.

3.3 Provision Google Cloud Dedicated Interconnect

To provision a Dedicated Interconnect network, perform the following tasks, which are detailed in the following subsections:

1. Order Google Cloud Dedicated Interconnect.
2. Order cross connects from Equinix.
3. Test the Dedicated Interconnect connections.
4. Create VPC.
5. Create VLAN attachments and establish BGP sessions.

Note: The Dedicated Interconnect provisioning process can be completed via the GCP Console web interface. Most steps in the provisioning process can be completed via the `gcloud` CLI in the Google Cloud Shell. This document uses `gcloud` steps where applicable.

Order Google Cloud Dedicated Interconnect

To order the Google Cloud Dedicated Interconnect, complete the following steps:

Note: A Google account and internet browser software installed on an internet-connected computer are required.

1. On the computer where the internet browser software is installed, go to <https://console.cloud.google.com> and click Activate Google Cloud Shell. The Google Cloud Shell opens in a new browser tab.

2. In the Google Cloud Shell, run the following command to create a new project:

```
gcloud projects create <<project-id>> --name="<<project-name>>"
```

Where:

- `<<project-id>>` (for example, `netapp-dev`) is the ID for the project you want to create. Project IDs must start with a lowercase letter and can contain lowercase ASCII letters, digits, or hyphens. Project IDs must be between 6 and 30 characters and must be unique.

- <<project-name>> (for example, NetApp-Dev) is the name of the project. If you do not specify a project name, the default value is project-id.
3. After you create the project, run the following commands to order GCI Dedicated Interconnect connections that will be used in a redundant configuration of two 10Gbps GCI connections:

```
gcloud compute interconnects create <<pri-gci-name>> \
--customer-name "<<customer-name>>" \
--description "<<gci-link-description>>" \
--interconnect-type IT_PRIVATE \
--link-type LINK_TYPE_ETHERNET_10G_LR \
--location <<gci-location>> \
--requested-link-count 1 \
--noc-contact-email <<customer-email>>

gcloud compute interconnects create <<sec-gci-name>> \
--customer-name "<<customer-name>>" \
--description "<<gci-link-description>>" \
--interconnect-type IT_PRIVATE \
--link-type LINK_TYPE_ETHERNET_10G_LR \
--location <<gci-location>> \
--requested-link-count 1 \
--noc-contact-email <<customer-email>>
```

Where:

- <<pri-gci-name>> (for example, nps-dev-pri) is the name of the first GCI Dedicated Interconnect connection.
- <<customer-name>> (for example, NetApp) is the name of the customer for the letter of authorization.
- <<gci-link-description>> (for example, iad-zone1-1-pri) is the description of the GCI Dedicated Interconnect connection.
- <<gci-location>> (for example, iad-zone1) is the name of the GCI Interconnect location. The list of names is on the [Colocation Facility Locations](#) webpage.
- <<customer-email>> is the e-mail address to which the letter of authorization should be sent.
- <<sec-gci-name>> (for example, nps-dev-sec) is the name of the second GCI Dedicated Interconnect connection.

Order Cross Connects from Equinix

To order cross connects from Equinix, complete the following steps:

1. Use the [Equinix Customer Portal](#) to create a cross-connect request to the GCP for a 10Gbps connection. Equinix uses the letter of authorizations sent by e-mail from Google to patch a cross connect from the GCP point-of-presence to the demarcation panel in the cage.

Note: If you have any questions about how to submit a cross-connect request, contact your Equinix client services manager.
2. For each Dedicated Interconnect cross connect, patch an SMF duplex cable from the demarcation panel where the cross connect is patched to the network equipment in the cage or cabinet.
3. After the cross connect is patched, schedule a network turn up by using the [Equinix Customer Portal](#) or through your Equinix client services manager.

Test Dedicated Interconnect Connections

Before you can use the Dedicated Interconnect connections, Google must verify that the cross connects are working. For each connection, Google sends IP configuration information that must be applied to the customer-provided network equipment. Google performs the following tests:

- Tests light levels by using the test IP addresses.

- After the first test passes, Google tests connectivity with the final IP addresses.

Here are the steps for each test.

1. For the first test, complete the following steps:
 - a. Log in to the primary customer-provided network equipment.
 - b. Create a port channel interface.
 - c. Configure the IP address on the port channel interface and assign the port patched into the primary Dedicated Connect cross connect to the port channel:

```
config t
int port-channel <<po-num>>
no switchport
ip address <<pri-link-addr>>/30
no shutdown

int eth <<pri-port>>
no switchport
port-channel <<po-num>> mode active
no shutdown

end
exit
```

2. Log in to the secondary customer-provided network equipment and complete the following steps:
 - a. Create a port channel interface.
 - b. Configure the IP address on the port channel interface and assign the port patched into the secondary Dedicated Connect cross connect to the port channel.

```
config t
int port-channel <<po-num>>
no switchport
ip address <<sec-link-addr>>/30
no shutdown

int eth <<sec-port>>
no switchport
port-channel <<po-num>> mode active
no shutdown

end
exit
```

Where:

- <<pri-port>> is the port on the primary customer-provided network equipment where the primary GCI Dedicated Interconnect connection is patched.
- <<pri-link-addr>> is the test IP address assigned to the primary Dedicated Interconnect connection.
- <<sec-port>> is the port on the secondary customer-provided network equipment where the secondary Dedicated Interconnect connection is patched.
- <<sec-link-addr>> is the test IP address assigned to the primary Dedicated Interconnect connection.

Note: If one or more of the connections fail the first test, refer to the [Google Cloud Dedicated Interconnect Troubleshooting Guide](#).

After the connections pass the first test, Google sends an e-mail with the production peering addresses for each connection.

- Remove the test IP addresses from the port-channel interfaces on the primary and secondary customer-provided network equipment.
- Configure the production peering IP address on the primary customer-provided network equipment by running the following commands:

```
config t
interface port-channel <<po-num>>
no switchport
no ip address
ip address <<pri-production-peering-IP>>/30
no shutdown
```

Where:

- <<pri-production-peering-IP>> is the production peering IP address provided by GCP for your primary interconnect.

- Configure the production peering IP address on the secondary customer-provided network equipment by running the following commands:

```
config t
interface port-channel <<po-num>>
no switchport
no ip address
ip address <<sec-production-peering-IP>>/30
no shutdown
```

Where:

- <<sec-production-peering IP>> is the production peering IP address provided by GCP for your secondary interconnect.

- After all the tests have passed, run the following command to view the status of the Dedicated Interconnect connections:

```
gcloud compute interconnects list
```

The output of the command should show the interconnects in an `OS_ACTIVE` for operational status:

NAME	LOCATION	OPERATIONAL_STATUS	ADMIN_ENABLED
nps-dev-pri	iad-zone1-1	OS_ACTIVE	True
nps-dev-sec	iad-zone2-1	OS_ACTIVE	True

Create VPC

There are two types of VPC networks: auto mode and custom mode. Auto-mode VPCs start with a single subnet in a specific region using a predefined address space.

For more information about GCP VPC, see the [GCP Virtual Private Cloud \(VPC\) Overview documentation](#).

For more information about creating a GCP VPC, see the [GCP Using VPC Networks documentation](#).

In the Google Cloud Shell, run the following command to create an auto-mode VPC:

```
gcloud compute networks create <<vpc-name>> --mode auto
```

Where:

- <<vpc-name>> is the name of the VPC (for example, `nps-vpc`). As a reminder, a VPC is a global resource and can contain subnets assigned to multiple regions.

Note: The IP address space and gateway for the `us-east4` subnet for auto-mode VPCs are: `10.150.0.0/20` and `10.150.0.1`. This subnet is created automatically.

Note: The complete list of auto-mode VPC network IP ranges can be found here: <https://cloud.google.com/compute/docs/vpc/#subnet-ranges>.

Create VLAN Attachments and Establish BGP Sessions

VLAN attachments on a Cloud Interconnect control which VPCs can connect to NetApp storage. Creating a VLAN attachment on a Cloud Interconnect automatically allocates a VLAN on the interconnect and provides BGP peering IP addresses. Each VLAN attachment is associated with a Cloud Router (which must exist before the VLAN attachment is created) in the region to which you want to establish connectivity.

For more information about VLAN attachments, see the [Google Cloud Dedicated Interconnect Creating VLAN Attachments documentation](#).

To create VLAN attachments and establish BGP sessions, complete the following steps:

1. Create cloud routers in the regions and VPCs to where you want to connect your NetApp storage. Run the following commands in the Google Cloud Shell:

```
gcloud compute routers create rtr-nps-pri --asn 64530 --network nps-vpc --region us-east4
gcloud compute routers create rtr-nps-sec --asn 64530 --network nps-vpc --region us-east4
```

2. After the Cloud Routers have been created, create VLAN attachments on both the primary and secondary interconnects:

```
gcloud compute interconnects attachments create att-nps-east4-pri --interconnect nps-dev-pri --
router rtr-nps-pri --region us-east4
gcloud compute interconnects attachments create att-nps-east4-sec --interconnect nps-dev-sec --
router rtr-nps-sec --region us-east4
```

3. Describe the VLAN attachments created in step 2 to retrieve the VLAN IDs and BGP peering addresses:

```
gcloud compute interconnects attachments describe att-nps-east4-pri --region us-east4
gcloud compute interconnects attachments describe att-nps-east4-sec --region us-east4
```

4. The output of these commands displays information similar to the following sample output:

Note: Make a note of the CloudRouterIPAddress, customerRouterIPAddress, and the 802.1q tag, because this information is required in subsequent steps.

```
cloudRouterIpAddress: 169.254.159.105/29
creationTimestamp: '2017-12-19T12:23:17.440-08:00'
customerRouterIpAddress: 169.254.159.106/29
id: '4474849221594631130'
interconnect: https://www.googleapis.com/compute/v1/projects/netapp-dev/global/interconnects/nps-
dev-pri
kind: compute#interconnectAttachment
name: att-nps-east4-pri
operationalStatus: OS_ACTIVE
privateInterconnectInfo:
  tag8021q: 1000
region: https://www.googleapis.com/compute/v1/projects/netapp-dev/regions/us-east4
router: https://www.googleapis.com/compute/v1/projects/netapp-dev/regions/us-east4/routers/rtr-
nps-pri
selfLink: https://www.googleapis.com/compute/v1/projects/netapp-dev/regions/us-
east4/interconnectAttachments/att-nps-east4-pri
```

5. From the VLAN attachment description, retrieve the link local IP address (cloudRouterIpAddress) that you need to assign to your Cloud Routers and attach to the assigned VLAN. In the Google Cloud Shell, run the following commands to create an interface on the Cloud Router associated with your primary interconnect:

```
gcloud compute routers add-interface rtr-nps-pri --region us-east-4 --ip-address
<<cloudRouterIpAddress pri>> --mask-length 29 --interface-name nps-east4-pri --interconnect-
attachment att-nps-east4-pri
```

Where:

- <<cloudRouterIpAddress pri>> is the IP address referenced as cloudRouterIpAddress in the primary interconnect VLAN attachment description.

6. Repeat step 5 to create an interface on the Cloud Router associated with your secondary interconnect:

```
gcloud compute routers add-interface rtr-nps-sec --region us-east-4 --ip-address <<cloudRouterIpAddress sec>> --mask-length 29 --interface-name nps-east4-sec --interconnect-attachment att-nps-east4-sec
```

Where:

- <<cloudRouterIpAddress sec>> is the IP address referenced as cloudRouterIpAddress in the secondary interconnect VLAN attachment description.

7. To complete the Cloud Router setup, use the customerRouterIpAddress that you retrieved from your VLAN attachment to assign the BGP peer for your Cloud Router. The customerRouterIpAddress is assigned to a VLAN interface on the customer-provided network equipment in section 3.4. Run the following command in the Google Cloud Shell:

```
gcloud compute routers add-bgp-peer rtr-nps-pri --interface nps-east4-pri --region us-east4 --peer-name bgp-nps-east4-pri --peer-ip-address <<customerRouterIpAddress pri>> --peer-asn <<peer-asn>>
```

Where:

- <<customerRouterIpAddress pri>> is the IP address referenced as customerRouterIpAddress in the primary interconnect VLAN attachment description.
- <<peer-asn>> is the ASN of the first customer-provided network equipment.

8. Repeat step 7 for the Cloud Router associated with the secondary interconnect:

```
gcloud compute routers add-bgp-peer rtr-nps-sec --interface nps-east4-sec --region us-east4 --peer-name bgp-nps-east4-sec --peer-ip-address <<customerRouterIpAddress sec>> --peer-asn <<peer-asn>>
```

Where:

- <<customerRouterIpAddress sec>> is the IP address referenced as customerRouterIpAddress in the secondary interconnect VLAN attachment description.
- <<peer-asn> is the ASN of the second customer-provided network equipment.

9. Verify the configuration of the Cloud Router by using the Cloud Shell:

```
gcloud compute routers describe rtr-nps-pri --region us-east4
```

The output should be similar to the following sample output:

```
bgp:
  asn: 64530
bgpPeers:
- interfaceName: nps-east4-pri
  ipAddress: 169.254.159.105
  name: bgp-nps-east4-pri
  peerAsn: 64514
  peerIpAddress: 169.254.159.106
creationTimestamp: '2017-12-19T12:06:58.386-08:00'
id: '271963318591475597'
interfaces:
- ipRange: 169.254.159.105/29
  linkedInterconnectAttachment: https://www.googleapis.com/compute/v1/projects/netapp-dev/regions/us-east4/interconnectAttachments/att-nps-east4-pri
  name: nps-east4-pri
kind: compute#router
name: rtr-nps-pri
network: https://www.googleapis.com/compute/v1/projects/netapp-dev/global/networks/nps-vpc
region: https://www.googleapis.com/compute/v1/projects/netapp-dev/regions/us-east4
selfLink: https://www.googleapis.com/compute/v1/projects/netapp-dev/regions/us-east4/routers/rtr-nps-pri
```

3.4 Set Up the Customer Network Switch

At this stage, the customer-provided network equipment can be configured. As previously noted, customers can provide any brand or model layer 3 network switch that meets the following requirements:

- Supports the BGP (licensed and enabled)
- Has at least one 9/125 SMF 1Gbps or 10Gbps port available
- Has 1000BASE-T Ethernet ports
- Supports 802.1Q VLAN tags
- Supports 802.3ad LACP

To set up the customer's network switch, complete the following steps:

1. Create and configure a VLAN interface.
2. Create and configure virtual routing and forwarding (VRF) instances.
3. Configure a new BGP peer session.

Note: See the switch manufacturer's documentation for specific configuration commands.

Sample Switch Commands

The following sample commands are for a Cisco Nexus switch running NX-OS:

```
configure terminal
vrf-context <<vrf-name>>

vlan <<pri-vlan-id>>

interface vlan <<pri-vlan-id>>
vrf member <<vrf-name>>
description "VLAN for att-nps-east4-pri"
ip address <<customerRouterIpAddress-pri>>/29
ip address <<local-subnet-gateway-address>>/<<cidr>> secondary
exit

router bgp <<asn>>
vrf <<vrf-name>>
address-family ipv4 unicast
network <<local-subnet>>/<<cidr>>
exit
neighbor <<cloudRouterIpAddress>> remote-as <<cloudRouterASN>>
address-family ipv4 unicast
exit
ebgp-multihop 4
exit
end

copy running-config startup-config
```

Where:

- <<vrf-name>> is the VRF name.
- <<pri-vlan-id>> is the VLAN ID from the VLAN attachment on the primary interconnect.
- <<customerRouterIpAddress-pri>> is the IP address referenced as customerRouterIpAddress in the primary interconnect VLAN attachment description.
- <<local-subnet-gateway-address>> is the gateway address of the local subnet used by the NetApp SVM.
- <<cidr>> is the CIDR number for the local SVM subnet.
- <<asn>> is the ASN of the local network.

Note: The value of the ASN of the local network must match the ASN supplied to the Cloud Router as <<peer asn>> (see step 7 in the [Create VLAN Attachments and Establish BGP Sessions](#) section).

- <<local-subnet>> is the local subnet network.
- <<cloudRouterIpAddress>> is the IP address referenced as cloudRouterIpAddress in the primary interconnect VLAN attachment description.
- <<cloudRouterASN>> is the ASN assigned to the primary Cloud Router.

Sample Switch Configuration

The following sample switch configuration is for a Cisco Nexus switch running NX-OS:

```
vrf context vrf-100

vlan 100
vrf member vrf-1000
interface vlan 100
  no shutdown
  vrf member vrf-100
  no ip redirects
  ip address 169.254.159.106/29
  ip address 192.168.100.1/28 secondary

router bgp 64514
vrf vrf-100
  address-family ipv4 unicast
  network 192.168.100.0/28
  neighbor 169.254.159.105 remote-as 64530
  address-family ipv4 unicast
```

3.5 Configure NetApp Storage

The configuration parameters of the NetApp storage are from the IP NetApp Private Storage for GCP address plan (see

Table 1).

To configure the NetApp storage, complete the following steps:

1. Create VLAN interface ports on cluster nodes.
2. Create a storage virtual machine (SVM) on the cluster.
3. Create logical interfaces (LIFs) on the SVM that uses the VLAN interface ports:
 - a. Management LIF
 - b. CIFS/NFS LIF
 - c. iSCSI LIFs

3.6 Test Connections and Protocol Access

To verify and test the GCI Dedicated Interconnect network connection and in the NPS/GCP environment, complete the following tasks, which are detailed in the following subsections:

1. Prepare GCP VM instance.
2. Test network connectivity.
3. Test iSCSI protocol connectivity.
4. Verify iSCSI LUN access.
5. Verify SMB protocol connectivity.
6. Verify NFS protocol connectivity.

7. Test AutoSupport.

Prepare GCP Virtual Machine Instance

A Google Compute Engine instance must be created in the VPC where your Cloud Routers exist.

If you know how to deploy a GCE instance, you can skip to step 2.

1. Provision a GCE instance in the VPC either by using the Google Cloud Console or by running the following command in the Google Cloud Shell:

```
gcloud compute instances create <<instance-name>> --zone <<zone>> --machine-type <<machine-type>> --network <<network>>
```

Where:

- <<instance-name>> is the name of your GCE instance.
- <<zone>> is the zone in which you want to create your instance, ideally located in the same region as your Dedicated Interconnects.
- <<machine-type>> is the machine type to be used for your GCE instance; defaults to n1-standard-1.
- <<network>> is the VPC network that contains your Cloud Routers.

The command output should be similar to the following example:

```
Created [https://www.googleapis.com/compute/v1/projects/netapp-dev/zones/us-east4-a/instances/test].
NAME      ZONE          MACHINE_TYPE  PREEMPTIBLE  INTERNAL_IP  EXTERNAL_IP  STATUS
test     us-east4-a    n1-standard-1          10.150.0.3   35.188.248.77  RUNNING
```

2. Log in to the GCE instance provisioned in step 1. Use an SSH client to connect to a Linux instance, or use an RDP client to connect to a Windows instance.
 - For more information about how to connect to a Linux instance, see the GCP [GCE documentation](#).
 - For more information about how to connect to a Windows instance, see the GCP [GCE documentation](#).

Test Network Connectivity

To test network connectivity, complete the following steps:

1. Use the ping utility on the GCP VM instance to verify network connectivity. On the VM, run the following command to ping the SVM network gateway on your switch in Equinix:

```
ping <<svm-gateway>>
```

Where:

- <<svm-gateway>> is the IP address of the layer 3 interface on your switch in Equinix (for example, 192.168.100.1)

The output of the command looks similar to this example:

```
Pinging 192.168.100.1 with 32 bytes of data:
Reply from 192.168.100.1: bytes=32 time=2ms TTL=251
Reply from 192.168.100.1: bytes=32 time=1ms TTL=251
Reply from 192.168.100.1: bytes=32 time=1ms TTL=251
Reply from 192.168.100.1: bytes=32 time=1ms TTL=251

Ping statistics for 192.168.100.1:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 1ms, Maximum = 2ms, Average = 1ms
```

Note: On the first ping attempt, there might be one or two dropped packets. However, after the first attempt, there should be no dropped packets.

Note: The output of the ping command varies based on the operating system used.

2. On the VM, run the following command to ping the NetApp SVM LIF:

```
ping <<svm-lif>>
```

Where:

- <<svm-lif>> is the IP address of the network interface on the NetApp SVM (for example, 192.168.100.2).

The output of the command looks similar to this example:

```
Pinging 192.168.100.2 with 32 bytes of data:
Reply from 192.168.100.2: bytes=32 time=2ms TTL=251
Reply from 192.168.101.2: bytes=32 time=1ms TTL=251
Reply from 192.168.100.2: bytes=32 time=1ms TTL=251
Reply from 192.168.100.2: bytes=32 time=1ms TTL=251

Ping statistics for 192.168.100.2:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 1ms, Maximum = 2ms, Average = 1ms
The output of the ping command will vary on the operating system used.
```

Test iSCSI Protocol Connectivity

To test iSCSI protocol connectivity, complete the following steps:

1. Use the iSCSI software initiator on your GCP VM instance to establish iSCSI sessions to the iSCSI LIFs created in section 3.5.
 - Note:** For information about how to establish an iSCSI session, see the documentation for the operating system of the GCP VM instance.
 - Note:** See the SAN Administration Guide on the [NetApp Support website](#) for the version of ONTAP that you are using on the NPS storage system.
2. The positive outcome for the test is that an iSCSI session has been successfully established from the iSCSI software initiator on the GCP VM instance to the iSCSI LIF on the NPS storage.

Verify iSCSI LUN Access

To verify iSCSI LUN access, complete the following steps:

1. From a local administration host, or from the GCP VM instance, create an aggregate, flexible volume, LUN, and igroup by using the ONTAP CLI or NetApp OnCommand® System Manager software.
 - Note:** The commands and/or workflows to create these storage primitives depend on the version of ONTAP used on the NPS storage system.
 - Note:** See the SAN Administration Guide on the [NetApp Support website](#) for the version of ONTAP that you are using on the NPS storage system.
2. After configuring the NetApp storage, use iSCSI tools on the GCP VM instance to discover the iSCSI LUN (for example, iscsiadm, Windows iSCSI control panel application, and so on).
 - Note:** For more information about how to discover the iSCSI LUN, see the documentation for the operating system of the GCP VM instance.
3. After the iSCSI LUN has been discovered by the GCP VM instance, create a file system on the LUN and mount the file system.

- Note:** For more information about how to discover the iSCSI LUN, see the documentation for the operating system of the GCP VM instance.
4. Use the CD utility on your GCP VM instance connected to the iSCSI LUN. Write a text file and save it to the iSCSI LUN.
Note: For more information about how to write and save a file, see the documentation for the operating system of the GCP VM instance.
 5. The positive outcome for this test is that you can access the LUN file system and write a file to it.

Verify SMB Protocol Connectivity

To verify SMB protocol connectivity, complete the following steps.

Note: To perform this test, you need a GCP VM instance running the Windows operating system deployed into the VPC that is connected to the GCI Dedicated Interconnect. If you do not have a Windows VM instance deployed, deploy one before you begin.

1. From a local administration host, or from the GCP VM instance, create a flexible volume and junction point on the NPS storage.
Note: See the File Access Management Guide for CIFS on the [NetApp Support website](#) for the version of ONTAP that you are using on the NPS storage system.
2. After creating the SMB share, use the GCP VM instance to access the share. Write a text file and save it to the SMB share.
3. The positive outcome for this test is that you can access the SMB share and write a file to it.

Verify NFS Protocol Connectivity

To verify NFS protocol connectivity, complete the following steps.

Note: To perform this test, you need a GCP VM instance running the Linux operating system deployed into the VPC that is connected to the GCI Dedicated Interconnect. If you do not have a Linux VM instance deployed, deploy one before you begin.

1. From a local administration host, or from the Linux VM instance, create a flexible volume and junction point on the NPS storage.
Note: See the File Access Management Guide for NFS on the [NetApp Support website](#) for the version of ONTAP that you are using on the NPS storage system.
2. After creating the NFS export, use the GCP VM instance to mount the export. Write a text file and save it to the NFS export.
3. The positive outcome for this test is that you can access the NFS export and write a file to it.

Test AutoSupport

For NetApp AutoSupport® to work, the NetApp storage must have access to the internet or to a mail host that has access to the internet. You can accomplish this in one of the following ways:

- Set up a mail host in the VPC that is connected to the storage.
- Set up a network connection to the internet in the colocation where the storage is located.
- Set up a network connection back on premises over a VPN or MPLS connection.

Note: See the System Administration Guide on the [NetApp Support website](#) for the version of ONTAP that you are using on the NPS storage system.

3.7 Performance Testing Guidelines

The concepts underlying performance testing with NetApp Private Storage for GCP are similar to those for performance testing in other environments. The following sections describe considerations to take into account when conducting performance testing in the NetApp Private Storage for GCP solution environment.

Goals of Performance Testing

Performance tests are used to validate the performance of the storage, network, and computing resources, given a specific workload that is an estimate of a real-world workload.

All architectures have limits to their performance. The goal of performance testing is not to see how much load you can put in the environment before things break. Instead, the goal is to follow an iterative, deliberate process that results in data that can be plotted and analyzed so that architects can anticipate performance based on a given workload (that is, performance curves).

NetApp Storage Considerations for Performance Testing

The considerations for sizing NetApp storage are the same in the NetApp Private Storage for GCP solution architecture as in typical deployments of NetApp storage. NetApp storage requires the following considerations:

- **Number and type of NetApp controllers.** Are the number and type of controllers used in the testing appropriate for the performance testing?
- **Number and type of disks in the aggregates.** Do the number and type of disks in the aggregate used in the testing have enough IOPS and storage capacity for the testing?
- **NetApp Flash Cache caching.** Are Flash Cache™ adapters installed in the storage controller nodes?
- **Cluster-node network connectivity.** What is the bandwidth of network connections, and how many connections are used to connect the storage to the network equipment in the Equinix colocation data center that is connected to the GCP?

Network Equipment Considerations for Performance Testing

The considerations for the network equipment in the NetApp Private Storage for GCP solution architecture are the same as those in typical network environments. The network equipment requires the following considerations:

- **Available CPU and memory.** Does the switch that is being used have enough resources to support the performance testing? Adding more workload to an oversubscribed network switch might contribute to invalid performance testing results.
- **Network ports used.** What is the bandwidth of GCI Dedicated Interconnect network connections (10Gbps), and what is the number of connections used to connect to the storage and to GCP? Is there enough bandwidth available to accommodate a performance test?

GCP Considerations for Performance Testing

It is important to understand how the components of the GCP can affect performance testing. The following considerations apply to the GCP:

- **GCI Dedicated Interconnect network connection.** Is there enough bandwidth available to accommodate performance testing? Contention for network bandwidth can affect performance testing results. Be sure that there is enough network bandwidth to support the testing.

- **GCP VM machine type.** Verify that you are using the proper instance type for performance testing. GCP throttles network throughput for smaller instance types and allocates more network bandwidth for larger instance types. Having the correct instance type is critical for a successful performance test.

Load-Generation and Monitoring Tools Used for Performance Testing

The load-generation and monitoring tools used for performance testing in the NetApp Private Storage for GCP solution architecture are the same as those used in typical NetApp storage environments. Consider the following guidelines:

- **Know which tool you want to use.** Each tool has advantages and disadvantages. Understanding the correct tool for your performance testing can provide more accurate test results.
- **Know your workload.** What kind of workload will you be testing? Understanding the I/O patterns of the workloads you are testing helps make it possible to configure the load generation tool correctly so that the testing can accurately model the performance.
- **Monitor the stack.** Implement monitoring for the computing, network, and storage resources so that bottlenecks can be identified. Collect performance data from each stack so that analysis can provide a more complete picture of how the NetApp Private Storage for GCP solution architecture is performing.

Where to Find Additional Information

To learn more about the information described in this document, refer to the following documents and/or websites:

- Google Cloud Platform Dedicated Interconnect Documentation
<https://cloud.google.com/interconnect/docs/>

Version History

Version	Date	Document Version History
Version 1.0	January 2018	Initial Release

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 2018 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.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

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.

TR-4653-0118