Pexip Infinity and Google Cloud Platform

Deployment Guide

Software Version 20

Document Version 20.a

October 2018

pexip

Contents

Introduction

The Google Compute Engine (GCE) service provides scalable computing capacity in the Google Cloud Platform (GCP). Using GCP eliminates your need to invest in hardware up front, so you can deploy Pexip Infinity even faster.

You can use GCP to launch as many or as few virtual servers as you need, and use those virtual servers to host a Pexip Infinity Management Node and as many Conferencing Nodes as required for your Pexip Infinity platform.

GCP enables you to scale up or down to handle changes in requirements or spikes in conferencing requirements. You can also use the GCP APIs and the Pexip Infinity management API to monitor usage and bring up / tear down Conferencing Nodes as required to meet conferencing demand.

Pexip publishes disk images for the Pexip Infinity Management Node and Conferencing Nodes. These images may be used to launch instances of each node type as required.

Deployment guidelines

This section summarizes the GCP deployment options and limitations, and provides guidance on our recommended GCP instance types, security groups and IP addressing options.

This flowchart provides an overview of the basic steps involved in deploying the Pexip Infinity platform on Azure:

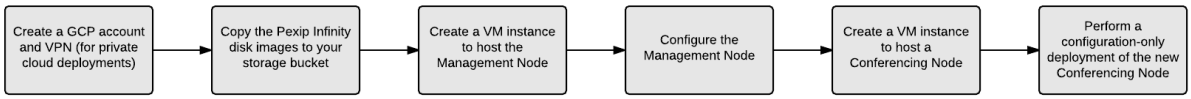

Deployment options

There are three main deployment options for your Pexip Infinity platform when using the Google Cloud Platform:

- **Private cloud**: all nodes are deployed within Google Cloud Platform. Private addressing is used for all nodes and connectivity is achieved by configuring a VPN tunnel between the corporate network and GCP. As all nodes are private, this is equivalent to an on-premises deployment which is only available to users internal to the organization.
- **Public cloud**: all nodes are deployed within GCP. All nodes have a private address but, in addition, public IP addresses are allocated to each node. The node's private addresses are only used for inter-node communications. Each node's public address is then configured on the relevant node as a static NAT address. Access to the nodes is permitted from the public internet, or a restricted subset of networks, as required. Any systems or endpoints that will send signaling and media traffic to those Pexip Infinity nodes must send that traffic to the public address of those nodes. If you have internal systems or endpoints communicating with those nodes, you must ensure that your local network allows such routing.
- **Hybrid cloud**: the Management Node, and optionally some Conferencing Nodes, are deployed in the corporate network. A VPN tunnel is created between the corporate network and GCP. Additional Conferencing Nodes are deployed in GCP and are managed from the on-premises Management Node. The GCP-hosted Conferencing Nodes can be either internally-facing, privately-addressed (private cloud) nodes; or externally-facing, publicly-addressed (public cloud) nodes; or a combination of private and public nodes (where the private nodes are in a different Pexip Infinity system location to the public nodes).

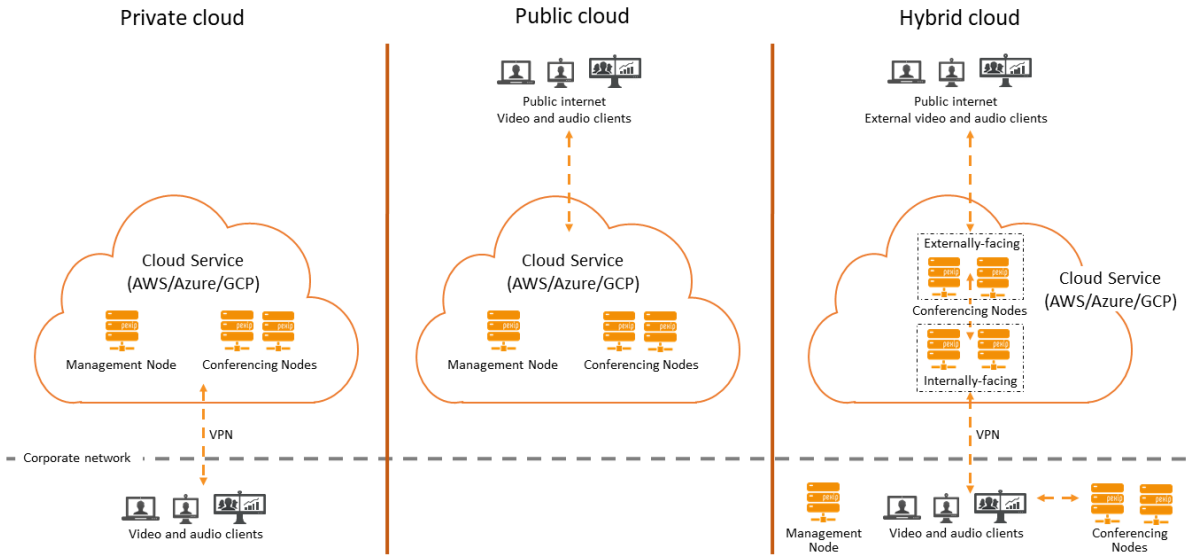

All of the Pexip nodes that you deploy in the cloud are completely dedicated to running the Pexip Infinity platform— you maintain full data ownership and control of those nodes.

Recommended instance types and call capacity guidelines

GCP instances come in many different sizes. In general, Pexip Infinity Conferencing Nodes should be considered compute intensive and Management Nodes reflect a more general-purpose workload. Our Server design recommendations also apply to cloud-based deployments.

For deployments of up to 20 Conferencing Nodes, we recommend using:

- **Management Node**: a machine type with 2 vCPUs 7.5GB memory (n1-standard-2) or larger.
- **Transcoding Conferencing Nodes**: a machine type with 8 vCPUs 7.2 GB memory (n1-highcpu-8) or larger.
- **Proxying Edge Nodes**: a machine type with 4 vCPUs (n1-highcpu-4).

This should provide capacity for approximately 11 HD / 27 SD / 140 audio-only calls per Transcoding Conferencing Node.

For all available machine types see: https://cloud.google.com/compute/pricing#predefined_machine_types.

Security and SSH keys

An SSH key must be applied to the VM instance that will host the Management Node (in order to complete the installation) and we also recommend applying SSH keys to your VM instances that will host your Conferencing Nodes. Keys can be applied project wide or for a particular VM instance. If a key is applied after the VM instance has been created then the instance must be rebooted for the key to take effect.

The username element of the SSH key must be "admin" or "admin@<domain>" i.e. the key takes the format:
`ssh-rsa [KEY_VALUE] admin` or
`ssh-rsa [KEY_VALUE] admin@vc.example.com` for example.

You can create key pairs with third-party tools such as PuTTYgen, or you can use an existing SSH key pair but you will need to format the public key to work in Compute Engine metadata (and ensure the username is modified to "admin"). For more information about using and formatting SSH keys for GCP, see https://cloud.google.com/compute/docs/instances/adding-removing-ssh-keys.

Google Cloud VPN for private/hybrid cloud deployments

For a private or hybrid cloud deployment, you must configure the Google Cloud virtual private network (VPN) to connect your on-premises network to the Google Cloud Platform.

For full information about how to configure the Google Cloud VPN, see https://cloud.google.com/compute/docs/vpn/overview.

IP addressing

All GCE VM instances are allocated a **Primary internal IP** (i.e. private) address. You can optionally also assign a static **External IP** (i.e. public) address to a GCE VM instance. You should assign a public address to all nodes in a public cloud deployment, and to any externally-facing nodes in a hybrid deployment, that you want to be accessible to conference participants located in the public internet.

Pexip Infinity nodes must always be configured with the private IP address associated with its instance, as it is used for all internal communication between nodes. To associate an instance's public IP address with the node, configure that public IP address as the node's **Static NAT address** (via **Platform Configuration > Conferencing Nodes**).

The private IP address should be used as the Conferencing Node address by users and systems connecting to conferences from the corporate network (via the Google Cloud VPN) in a private or hybrid cloud deployment. When an instance has been assigned an external IP address and that address is configured on the Conferencing Node as its **Static Nat address**, all conference participants **must** use that external address to access conferencing services on that node.

Assumptions and prerequisites

The deployment instructions assume that within GCP you have already:

- signed up to the Google Cloud Platform
- configured a Google Cloud VPN (for a private or hybrid cloud deployment)

For more information on setting up your Google Cloud Platform Virtual Machines, see https://cloud.google.com/compute/docs/instances/.

Configuring your Google VPC network

All Google Compute Engine (GCE) VM instances belong to a Google Virtual Private Cloud (VPC) network. You need to configure the VPC network to control access to the VM instances that will host your Pexip Infinity nodes in your Google Cloud Platform (GCP) deployment.

Google Cloud VPN for private / hybrid cloud deployments

For a private or hybrid cloud deployment, you must configure the Google Cloud virtual private network (VPN) to connect your on-premises network to the Google VPC network.

Google assigns a default range of private addresses to your VPC regions. You must ensure that the IP address ranges for the VPC regions in which you deploy your VM instances do not overlap with any subnets you use in your corporate network. If you do have overlapping subnets, you can create new subnets for each region in your Google VPC network, and then select that subnetwork when deploying your instance. See https://cloud.google.com/compute/docs/vpc/#subnet-ranges for information about the default VPC subnets per region.

For full information about how to configure the Google Cloud VPN, see https://cloud.google.com/compute/docs/vpn/overview.

A VPN is not required for public cloud deployments as you can access all of your nodes via their public IP addresses.

Enabling communication between Pexip Infinity nodes

To allow Pexip Infinity nodes to communicate, there must be a firewall rule in place to allow UDP and IPsec ESP protocol traffic between nodes. This applies to all deployment options (private, public and hybrid).

By default, the Google VPC network has a firewall rule called "default-allow-internal". This rule allows TCP, UDP and ICMP traffic between private addresses on the internal network, but it does not allow ESP traffic.

To modify this firewall rule to also allow ESP traffic:

1. From the GCP project console, go to **VPC network > Firewall rules**.
2. Select the **default-allow-internal** rule.
3. Select **Edit**.
4. Change **Protocols and ports** from "tcp:0-65535; udp:0-65535; icmp" to "tcp:0-65535; udp:0-65535; icmp; esp".

5. Select **Save**.

Note that this change adds ESP to the existing rule but does not remove or restrict any of the other default protocols and ports. This is because the **default-allow-internal rule** applies to all instances in your GCP project, and if you have something other than Pexip Infinity running (e.g. a reverse proxy, or something completely unrelated) then you probably want to allow UDP and TCP traffic to work.

Controlling access to the Management Node

We recommend that you lock down access to the Management Node to just the management stations that will administer your Pexip Infinity platform. This applies to all deployment options (private, public and hybrid), but is particularly important in public cloud deployments.

To create a new firewall rule to restrict access to the Management Node:

1. From the GCP project console, go to **VPC network > Firewall rules**.
2. Select **Create firewall rule**.
3. Complete the following fields (leave all other settings as default):

Name	Enter a name for the rule, for example "pexip-allow-management".
Direction of traffic	Select ***Ingress***.
Action on match	Select ***Allow***.
Targets	Select ***Specified target tags***.
Target tags	Enter a tag name, for example "pexip-management". You will use this name later when you create your Management Node VM instance to associate that instance with these firewall rules (see Deploying a Management Node in Google Cloud Platform).
Source filter	Select ***IP ranges***.
Source IP ranges	Enter the <IP address/subnet> of the management station/browsers that require access to the Management Node. Note that on a corporate network accessing a public cloud deployment, this should be the external public IP address of the corporate network and not the private address of the machine that is hosting the browser.
Protocols and ports	Enter **tcp:443**

4. Select **Create**.

Controlling access to Conferencing Nodes for installation/provisioning

We recommend that you lock down access to the provisioning interface on your Conferencing Nodes to just the management stations that will administer your Pexip Infinity platform. This applies to all deployment options (private, public and hybrid), but is particularly important in public and hybrid cloud deployments for nodes with an external IP address.

To create a new firewall rule to restrict access to the provisioning interface of a Conferencing Node:

1. From the GCP project console, go to **VPC network > Firewall rules**.
2. Select **Create firewall rule**.

3. Complete the following fields (leave all other settings as default):

Name	Enter a name for the rule, for example "pexip-allow-provisioning".
Direction of traffic	Select ***Ingress***.
Action on match	Select ***Allow***.
Targets	Select ***Specified target tags***.
Target tags	Enter a tag name, for example "pexip-provisioning". You will use this name later when you create your Conferencing Node VM instances to associate those instances with these firewall rules (see Deploying a Conferencing Node in Google Cloud Platform).
Source filter	Select ***IP ranges***.
Source IP ranges	Enter the <IP address/subnet> of the management station/browsers that require access to the Conferencing Nodes. Note that on a corporate network accessing a public cloud deployment, this should be the external public IP address of the corporate network and not the private address of the machine that is hosting the browser.
Protocols and ports	Enter **tcp:8443**

4. Select **Create**.

Controlling access to Conferencing Nodes for conference participants

A wider, more general access is typically required to the protocols and ports required to access conferences hosted on your Conferencing Nodes.

To create a new firewall rule to allow access to the conferencing-related ports and protocols of a Conferencing Node:

1. From the GCP project console, go to **VPC network > Firewall rules**.
2. Select **Create firewall rule**.
3. Complete the following fields (leave all other settings as default):

Name	Enter a name for the rule, for example "pexip-allow-conferencing".
Direction of traffic	Select ***Ingress***.
Action on match	Select ***Allow***.
Targets	Select ***Specified target tags***.
Target tags	Enter a tag name, for example "pexip-conferencing". You will use this name later when you create your Conferencing Node VM instances to associate those instances with these firewall rules (see Deploying a Conferencing Node in Google Cloud Platform).
Source filter	Select ***IP ranges***.
Source IP ranges	Enter **0.0.0.0/0** For a private deployment, the Source IP ranges should be restricted to the corporate intranet IP addresses.
Protocols and ports	Enter **tcp:80; tcp:443; tcp:1720; tcp:5060; tcp:5061; tcp:33000-39999; tcp:40000-49999; udp:1719; udp:33000-39999; udp:40000-49999** Note that if you have enabled SIP UDP then **udp:5060** must also be included.

4. Select **Create**.

After you have configured your firewall rules, your ingress rules will look similar to this:

Ingress Egress

Name	Targets	Source filters	Protocols / ports	Action	Priority	Network
default-allow-http	http-server	IP ranges: 0.0.0.0/0	tcp:80	Allow	1000	default
default-allow-https	https-server	IP ranges: 0.0.0.0/0	tcp:443	Allow	1000	default
pexip-allow-conferencing	pexip-conferencing	IP ranges: 0.0.0.0/0	tcp:80, tcp:443, 8 more	Allow	1000	default
pexip-allow-management	pexip-management	IP ranges: 81.143.209.108/32	tcp:443	Allow	1000	default
pexip-allow-provisioning	pexip-provisioning	IP ranges: 81.143.209.108/32	tcp:8443	Allow	1000	default
default-allow-icmp	Apply to all	IP ranges: 0.0.0.0/0	icmp	Allow	65534	default
default-allow-internal	Apply to all	IP ranges: 10.128.0.0/9	tcp:0-65535, udp:0-65535, 2 more	Allow	65534	default
default-allow-rdp	Apply to all	IP ranges: 0.0.0.0/0	tcp:3389	Allow	65534	default
default-allow-ssh	Apply to all	IP ranges: 0.0.0.0/0	tcp:22	Allow	65534	default

Obtaining and preparing disk images for GCE Virtual Machines

Pexip publishes Google Compute Engine (GCE) optimized disk images for the Management Node and for Conferencing Nodes.

Before you can use the published Pexip Infinity disk images, you must copy them to your storage bucket in the Google Cloud Platform (GCP). This guide refers to a disk image copied to your storage bucket as a **custom disk image**. All deployment operations use custom disk images.

Obtaining the Pexip disk images

To obtain your disk images, go to https://dl.pexip.com/infinity/index.html, select the appropriate directory for your software version, and then download the following files:

- **Pexip_Infinity_v20_GCP_pxMgr_<build>.tar.gz** for the Management Node.
- **Pexip_Infinity_v20_GCP_ConfNode_<build>.tar.gz** for a Conferencing Node.

Uploading disk images to Google Cloud Storage

The Pexip disk image packages must be uploaded to Google Cloud Storage.

1. Create a bucket to store the images:
   a. From the GCP project console, go to **Storage > Browser**.
   b. Enter a **Name** (for example, "pexip-v20"), and then select an appropriate **Storage class** and **Location** for your deployment. For more information about storage buckets, see https://cloud.google.com/storage/docs/creating-buckets
   c. Select **Create**.
2. Upload the Pexip images to the new bucket:
   a. Select the new bucket e.g. pexip-v20.
   b. Select **Upload Files**.
   c. In the dialog that appears, select the Management Node and Conferencing Node tar.gz files that you downloaded from Pexip.
   d. Select **Open**.

After you have uploaded your files, your bucket will look similar to this:

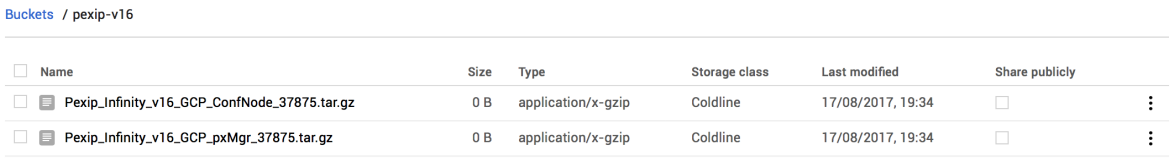

Preparing custom disk images

You must now prepare a custom disk image for the Management Node and for deploying a Conferencing Node.

Prepare a Management Node image

1. From the GCP project console, go to **Compute Engine > Images**.
2. Select **Create Image**.
3. Enter a **Name**, for example "pexip-mgr-v20".
4. Select a **Source** of ***Cloud Storage file***.
5. Select **Browse** and select the Management Node image package in your storage bucket e.g. pexip-v20.
6. Select **Create**.

You can now deploy the Management Node in Google Cloud Platform.

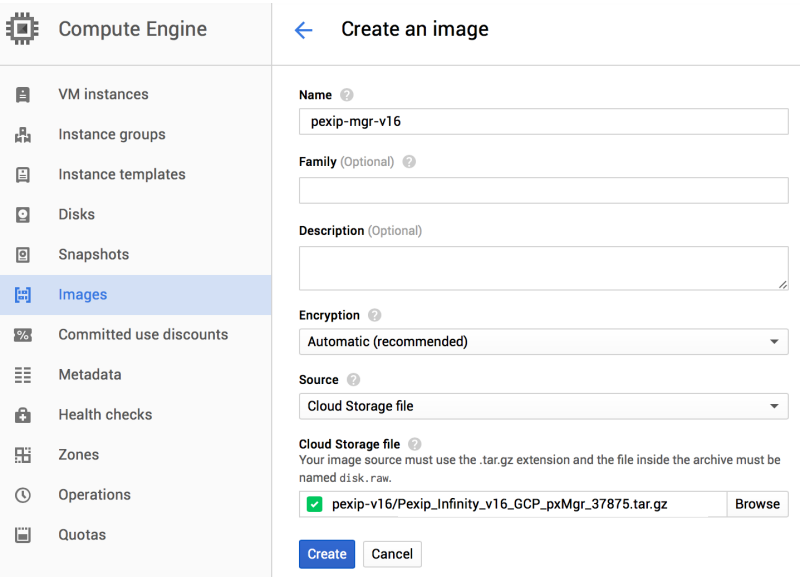

Prepare a Conferencing Node image

1. From the GCP project console, go to **Compute Engine > Images**.
2. Select **Create Image**.
3. Enter a **Name**, for example "pexip-node-v20".
4. Select a Source of ***Cloud Storage file***.
5. Select **Browse** and select the Conferencing Node image package in your storage bucket e.g. pexip-v20.
6. Select **Create**.

You can now deploy a Conferencing Node in Google Cloud Platform.

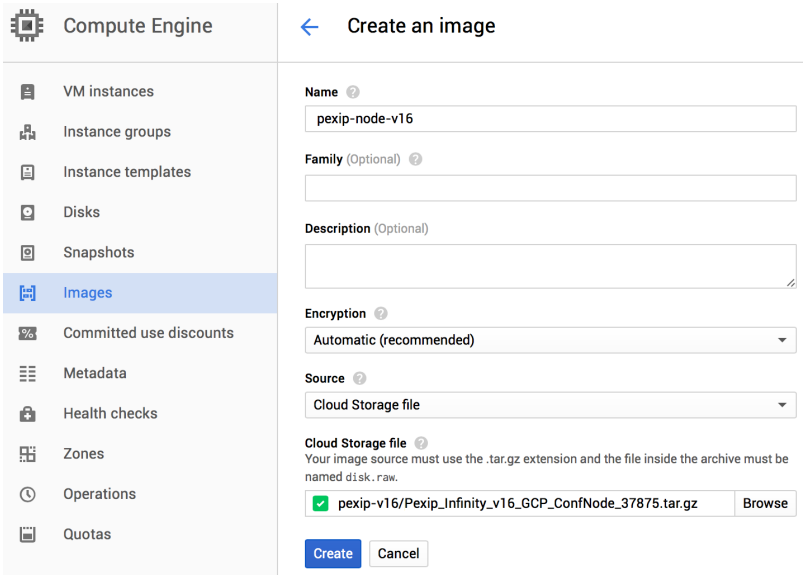
Compute Engine
VM instances
Instance groups
Instance templates
Disks
Snapshots
Images
Committed use discounts
Metadata
Health checks
Zones
Operations
Quotas
Create an image
Name
pexip-node-v16
Family (Optional)
Description (Optional)
Encryption
Automatic (recommended)
Source
Cloud Storage file
Cloud Storage file
Your image source must use the .tar.gz extension and the file inside the archive must be named disk.raw.
pexip-v16/Pexip_Infinity_v16_GCP_ConfNode_37875.tar.gz
Browse
Create
Cancel

Deploying a Management Node in Google Cloud Platform

As with all Pexip Infinity deployments, you must first deploy the Management Node before deploying any Conferencing Nodes. In a hybrid cloud deployment the Management Node may be deployed in the corporate network or in GCP. This section describes how to deploy the Management Node in GCP.

To deploy a Management Node on a Google Compute Engine VM:

1. Prepare a Management Node disk image. For more information on this, see Obtaining and preparing disk images for GCE Virtual Machines.
2. From the GCP project console, go to **Compute Engine > VM Instances**.
3. Select **Create Instance**.
4. Complete the following fields (leave all other settings as default):

Name	Enter a unique name for the instance, for example "pexipmanager".
Zone	Select an appropriate **Zone**. Typically you should choose a zone that is geographically close to the location from where it will be administered.
Machine type	Select ***2 vCPUs*** (n1-standard-2).
Boot disk	Select the Management Node custom disk image: a. Select **Change**. b. Select **Custom images**. c. Select the Management Node custom disk image, e.g. "pexip-mgr-v20". d. Select a **Boot disk type** of ***SSD persistent disk***. e. Select **Select**.
Networking: Network tags	Expand the **Management, disk, networking, SSH keys** section and select the **Networking** tab. Assign a **Network tag** to the instance, for example "pexip-management". This is the tag that should be applied to your Management Node firewall rule (see Controlling access to the Management Node).
Networking: External IP	You must decide whether or not to assign an external IP address to the instance. You must assign a static public/external IP address to the Management Node if you have a public cloud deployment. On a private or hybrid deployment you will typically access the Management Node over a VPC network, or host it on-premises. a. Expand the **Management, disk, networking, SSH keys** section and select the **Networking** tab. b. In the **Network interfaces** field, select the **default** interface to open the Network interface dialog. c. Select a **Subnetwork** if appropriate (e.g. if it is a private/hybrid deployment and you have created new subnets to avoid overlapping addresses in your corporate network). d. Select an appropriate **External IP**: ■ ***None***: no external IP address will be assigned. Use this where the node does not need to have a publicly-accessible IP address. ■ ***Create IP address***: select this option to create a static external address. You can enter a **Name** for the address and GCP will allocate a static IP address. ■ ***<external address>***: you can select a specific static external address if you have already created one in advance. Do **not** select ***Ephemeral*** — if you stop and restart the instance a new address will be assigned.

SSH keys	An SSH key must be applied to the Management Node instance if you are not already using a project-wide key for all of the instances in your project. The username element of the SSH key must be "admin" or "admin@<domain>". To apply an instance-level key: a. Select the **SSH Keys** tab and select **Show and edit**. b. Select **Add item**. This produces a text box. Copy the contents of your public SSH key file and paste them into the text box. c. Modify the username element to "admin" or "admin@<domain>" if necessary. d. Select **Save**. See Security and SSH keys for more information.

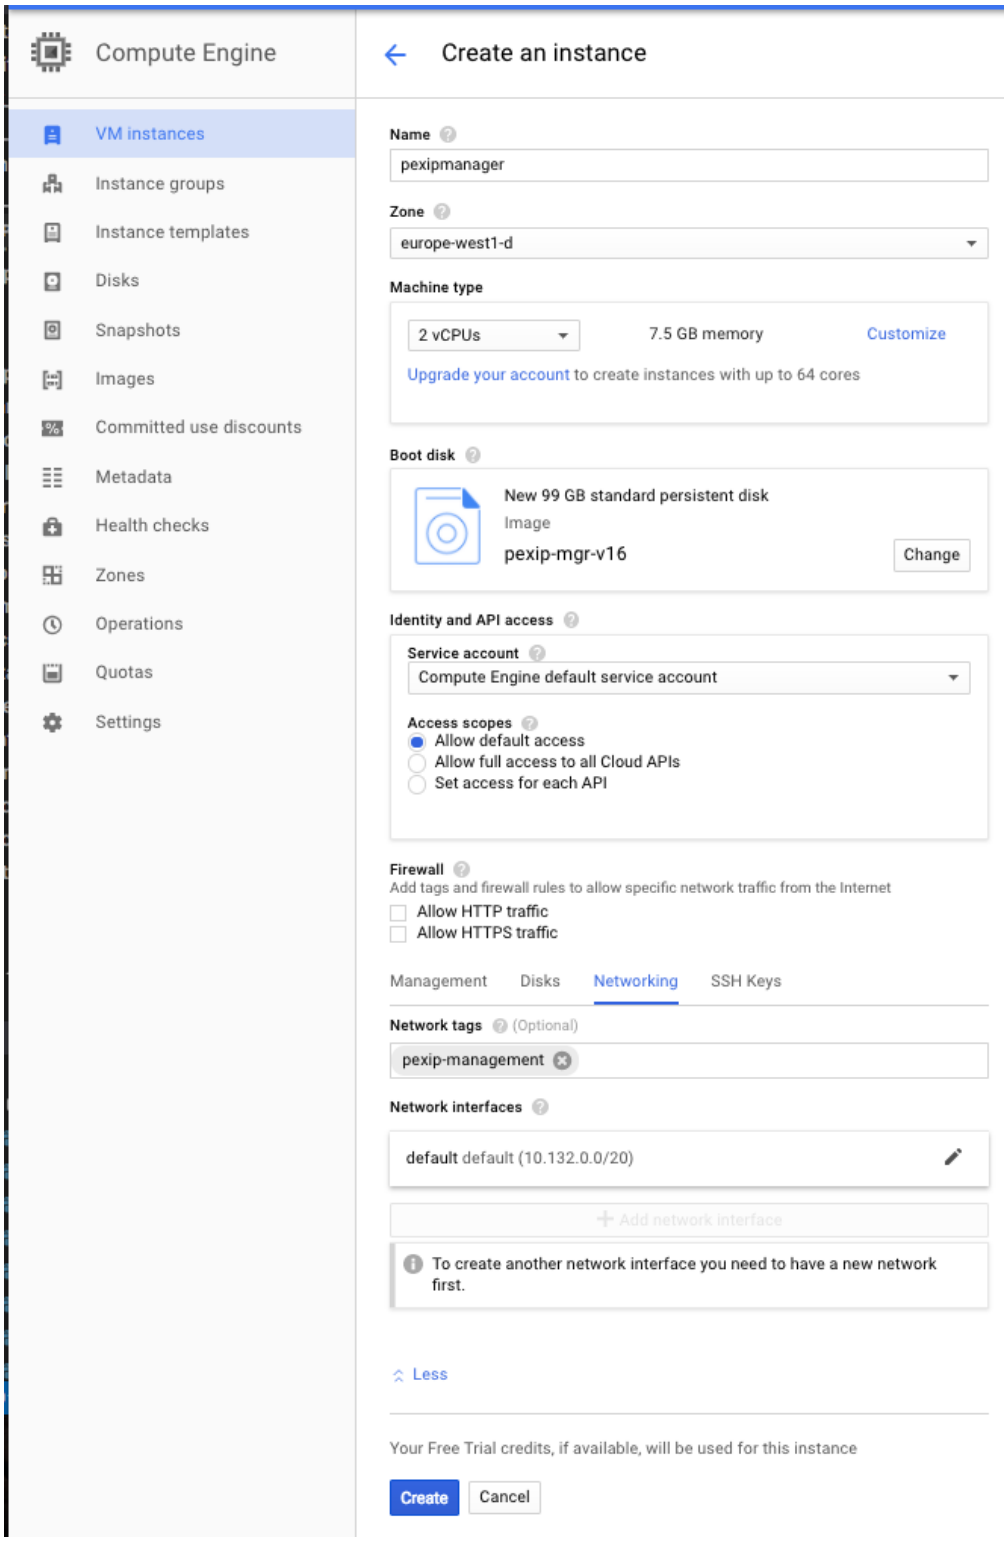

5. Select **Create** to create the instance.
6. You must now connect over SSH into the Management Node instance to complete the installation of Pexip Infinity:
   a. Use an SSH client to access the Management Node by its IP address, supplying your private key file as appropriate.
   b. At the "Enter new UNIX password:" prompt, enter your desired password, and then when prompted, enter the password again.

This will then log you out and terminate your SSH session.

7. Reconnect over SSH into the Management Node instance and continue the installation process:
   a. Log in again as ***admin***.
      You are presented with another login prompt:
      `[sudo] password for admin:`
   b. Enter the UNIX password you just created.
      The Pexip installation wizard will begin after a short delay.
   c. Complete the installation wizard to apply basic configuration to the Management Node:
      i. The **IP address** must match the internal IP address allocated by GCE (the default should be correct).
      ii. The **Network mask** must be **255.255.255.255**.
      iii. The **Gateway** must be the internal IP address of the gateway for the GCE region (the default should be correct).
      iv. Enter your required **Hostname** and **Domain suffix** for the Management Node.
      v. Configure one or more **DNS servers** and **NTP servers**. You must override the default values if it is a private deployment.
      vi. Set the **Web administration username** and **password**.
      vii. Select whether to **Enable incident reporting** and whether to **Send deployment and usage statistics to Pexip**.
      The DNS and NTP servers at the default addresses are only accessible if your instance has a public IP address. The installation wizard will fail if the NTP server address cannot be resolved and reached.
   After successfully completing the wizard, the SSH connection will be lost as the Management Node reboots.
8. After a few minutes you will be able to use the Pexip Infinity Administrator interface to access and configure the Management Node (remember to use https to connect to the node if you have only enabled access to **tcp:443** in your firewall rule, as shown in our example "pexip-allow-management" firewall rule).
   You can now configure your Pexip Infinity platform licenses, VMRs, aliases, locations etc, and add Conferencing Nodes.

Deploying a Conferencing Node in Google Cloud Platform

After deploying the Management Node you can deploy one or more Conferencing Nodes in GCP to provide conferencing capacity.

To deploy a Conferencing Node on a Google Compute Engine VM:

1. If you have not already done so, prepare a Conferencing Node disk image. For more information on this, see Obtaining and preparing disk images for GCE Virtual Machines.
2. From the GCP project console, go to **Compute Engine > VM Instances**.
3. Select **Create Instance**.
4. Complete the following fields (leave all other settings as default):

Name	Enter a unique name for the instance, for example "pexipnode-europe-1".
Zone	Select an appropriate **Zone**. Typically you should choose a zone that is geographically close to the location from where users will connect to it.
Machine type	Select ***8 vCPUs*** (n1-highcpu-8). We recommend selecting a minimum CPU platform. Select **Customize** and then select the most modern platform available that does not incur a surcharge, typically ***Intel Broadwell or later***.

Boot disk	Select the Conferencing Node custom disk image: a. Select **Change**. b. Select **Custom images**. c. Select the Conferencing Node custom disk image, e.g. "pexip-node-v20". d. Select **Select**. For Conferencing Nodes, SSDs are not a requirement, but general VM processes such as snapshots and backups will be faster with SSDs.
Networking: Network tags	Expand the **Management, disk, networking, SSH keys** section and select the **Networking** tab. Assign **Network tags** to the instance, for example "pexip-provisioning pexip-conferencing". These are the tags that should be applied to your Conferencing Node firewall rules (see Controlling access to Conferencing Nodes for installation/provisioning and Controlling access to Conferencing Nodes for conference participants).
Networking: External IP	You must decide whether or not to assign an external IP address to the instance. You must assign a static public/external IP address to the Conferencing Node if you want that node to be able to host conferences that are accessible from devices in the public internet. a. Expand the **Management, disk, networking, SSH keys** section and select the **Networking** tab. b. In the **Network interfaces** field, select the **default** interface to open the Network interface dialog. c. Select a **Subnetwork** if appropriate (e.g. if it is a private/hybrid deployment and you have created new subnets to avoid overlapping addresses in your corporate network). d. Select an appropriate **External IP**: ■ ***None***: no external IP address will be assigned. Use this where the node does not need to have a publicly-accessible IP address. ■ ***Create IP address***: select this option to create a static external address. You can enter a **Name** for the address and GCP will allocate a static IP address. ■ ***<external address>***: you can select a specific static external address if you have already created one in advance. Do **not** select ***Ephemeral*** — if you stop and restart the instance a new address will be assigned.
SSH keys	We recommend applying an SSH key to the Conferencing Node instance if you are not already using a project-wide key for all of the instances in your project. The username element of the SSH key must be "admin" or "admin@<domain>". To apply an instance-level key: a. Select the **SSH Keys** tab and select **Show and edit**. b. Select **Add item**. This produces a text box. Copy the contents of your public SSH key file and paste them into the text box. c. Modify the username element to "admin" or "admin@<domain>" if necessary. d. Select **Save**. See Security and SSH keys for more information.

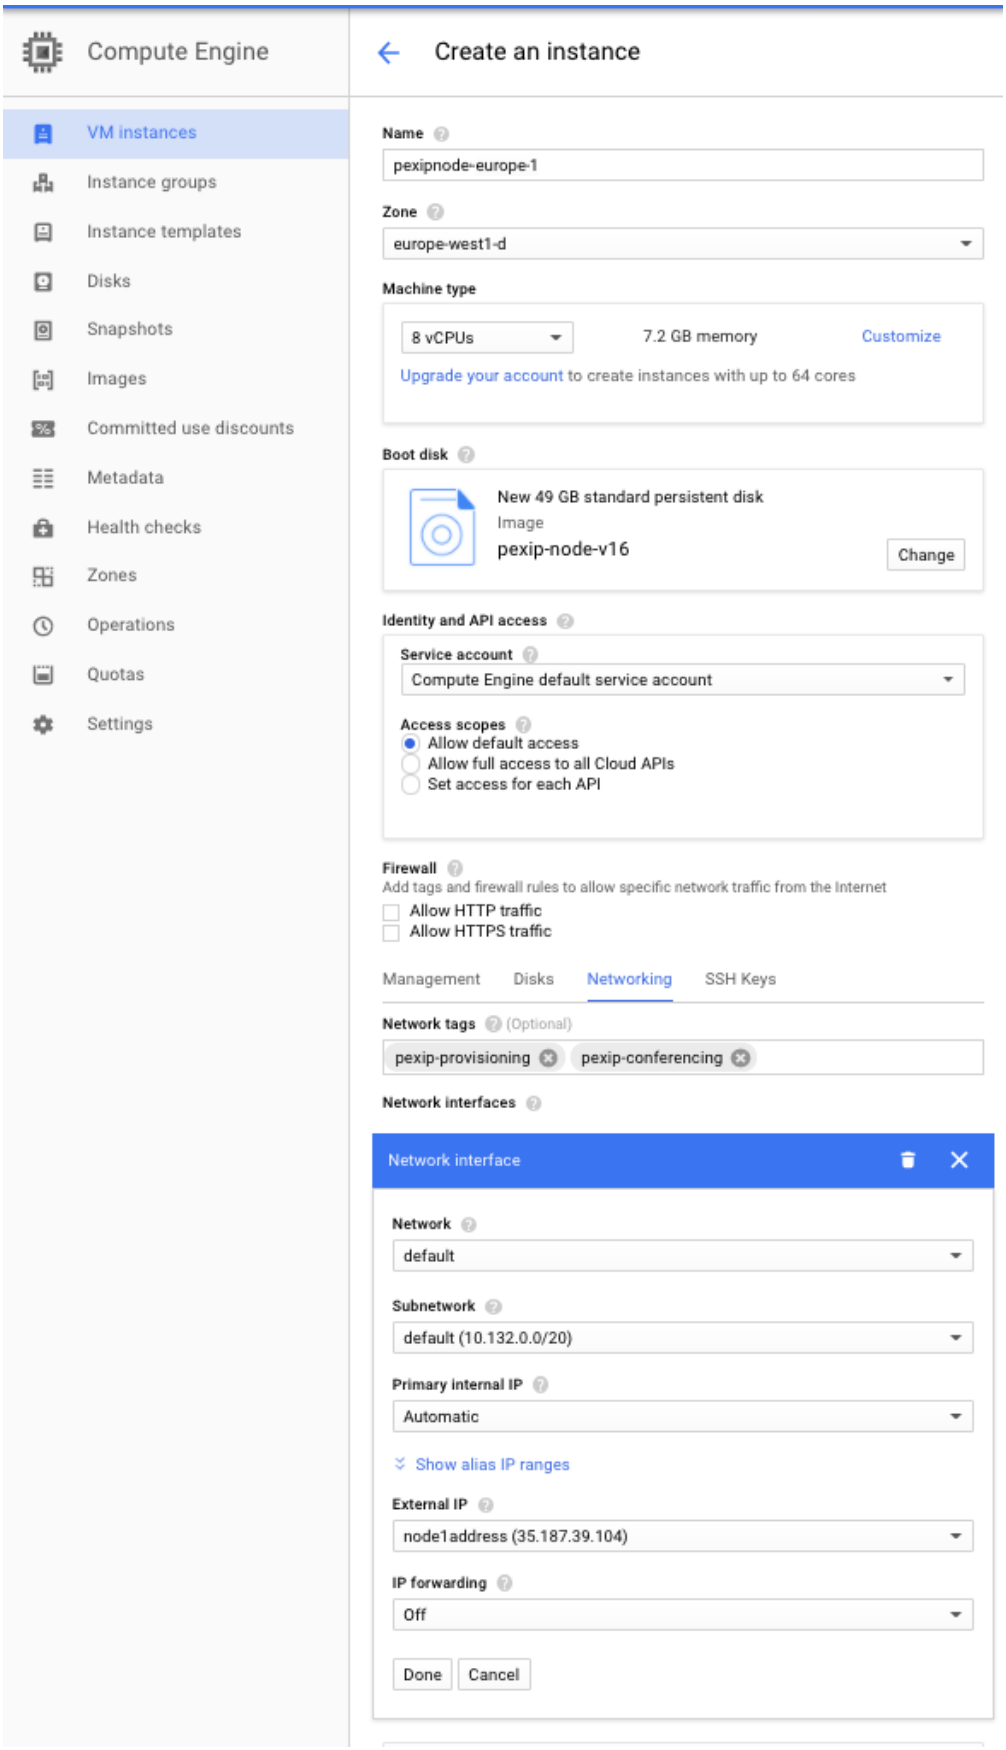
Compute Engine
Create an instance
VM instances
Instance groups
Instance templates
Disks
Snapshots
Images
Committed use discounts
Metadata
Health checks
Zones
Operations
Quotas
Settings
Name
pexipnode-europe-1
Zone
europe-west1-d
Machine type
8 vCPUs
7.2 GB memory
Customize
Upgrade your account to create instances with up to 64 cores
Boot disk
New 49 GB standard persistent disk
Image
pexip-node-v16
Change
Identity and API access
Service account
Compute Engine default service account
Access scopes
Allow default access
Allow full access to all Cloud APIs
Set access for each API
Firewall
Add tags and firewall rules to allow specific network traffic from the Internet
Allow HTTP traffic
Allow HTTPS traffic
Management
Disks
Networking
SSH Keys
Network tags (Optional)
pexip-provisioning
pexip-conferencing
Network interfaces
Network interface
Network
default
Subnetwork
default (10.132.0.0/20)
Primary internal IP
Automatic
Show alias IP ranges
External IP
node1address (35.187.39.104)
IP forwarding
Off
Done
Cancel

5. Select **Create** to create the instance.
6. On the **VM Instances** page, make a note of the "**Internal IP**" address, and the "**External IP**" address (if appropriate) that have been assigned to the new instance / Conferencing Node.
7. After the instance has booted, perform a configuration-only deployment on the Management Node to inform it of the new Conferencing Node:
   a. Log in to the Pexip Infinity Administrator interface on the Management Node.
   b. Go to **Platform Configuration > Conferencing Nodes**.
   c. Select **Add Conferencing Node**.
   d. For deployment type, choose ***Generic (configuration-only)*** and select **Next**.
   e. Enter the details of the new Conferencing Node, including:

IPv4 address	Enter the GCE **Internal IP** address of the new VM instance.
Network mask	Enter **255.255.255.255**
Gateway IP address	Enter the default gateway address for the region in which the node is deployed. See https://cloud.google.com/compute/docs/vpc/#subnet-ranges for a table of regions and default gateway addresses.
IPv4 static NAT address	Configure the Conferencing Node's static NAT address, if you have a assigned a public/external IP address to the instance. Enter the **External IP** address allocated by GCE for the VM instance.

   You must also specify other fields such as the **Name**, **Role**, **Hostname**, **Domain**, **System location** and assign a **TLS certificate**. For a full list of configuration fields, see https://docs.pexip.com/admin/deploy_vm_template.htm#deployment.
   f. Select **Finish**.
   g. Select **Download Conferencing Node Configuration** and save the XML configuration file.
   A zip file with the name **pexip-<hostname>.<domain>.xml** will be downloaded.
8. You must now upload the XML configuration file to the new Conferencing Node:
   a. Browse to **https://<conferencing-node-ip-address>:8443/** and use the form provided to upload the XML configuration file to the Conferencing Node VM.
   If you cannot access the Conferencing Node, check that you have allowed the appropriate source addresses in your ingress firewall rules for management traffic. In public deployments and where there is no virtual private network, you need to use the public address of the node.
      i. Select **Choose File** and select the XML configuration file.
      ii. Select **Upload**.
   b. The Conferencing Node will apply the configuration and then reboot. When it has rebooted, it will connect to the Management Node.
   You can close the browser window used to upload the file.

After deploying a new Conferencing Node, it takes approximately 5 minutes before the node is available for conference hosting and for its status to be updated on the Management Node. Until it is available, the Management Node will report the status of the Conferencing Node as having a last contacted and last updated date of "Never". "Connectivity lost between nodes" alarms relating to that node may appear temporarily.

When the node is up and running you can optionally remove the "pexip-provisioning" **Network tag** from the instance (or whichever tag you have associated with your provisioning firewall rule as described in Controlling access to Conferencing Nodes for installation/provisioning) as it is no longer required. Note, do not delete the firewall rule or remove the "pexip-conferencing" tag.

Managing Google Compute Engine VM instances

This section describes the common maintenance tasks for stopping, restarting and permanently removing Conferencing Node Google Compute Engine (GCE) VM instances on the Google Cloud Platform (GCP).

Temporarily removing (stopping) a Conferencing Node instance

At any time you can temporarily remove a Conferencing Node instance from your Pexip Infinity platform if, for example, you do not need all of your current conferencing capacity.

To temporarily remove a Conferencing Node instance:

1. Put the Conferencing Node into maintenance mode via the Pexip Infinity Administrator interface on the Management Node:
   a. Go to **Platform Configuration > Conferencing Nodes**.
   b. Select the Conferencing Node(s).
   c. From the **Action** menu at the top left of the screen, select **Enable maintenance mode** and then select **Go**.
   While maintenance mode is enabled, this Conferencing Node will not accept any new conference instances.
   d. Wait until any existing conferences on that Conferencing Node have finished. To check, go to **Status > Live View**.
2. Stop the Conferencing Node instance on GCP:
   a. From the GCP project console, select **Virtual Machines** to see the status of all of your instances.
   b. Select the instance you want to shut down.
   c. At the top of the VM instances page, select **Stop**.

Reinstating (restarting) a stopped Conferencing Node instance

You can reinstate a Conferencing Node instance that has already been installed but has been temporarily shut down.

To restart a Conferencing Node instance:

1. Restart the Conferencing Node instance on GCP:
   a. From the GCP project console, select **Virtual Machines** to see the status of all of your instances.
   b. Select the instance you want to restart.
   c. At the top right-hand of the page, select **Start** to restart the instance.
2. Take the Conferencing Node out of maintenance mode via the Pexip Infinity Administrator interface on the Management Node:
   a. Go to **Platform Configuration > Conferencing Nodes**.
   b. Select the Conferencing Node.
   c. Clear the **Enable maintenance mode** check box and select **Save**.

After reinstating a Conferencing Node, it takes approximately 5 minutes for the node to reboot and be available for conference hosting, and for its last contacted status to be updated on the Management Node.

Permanently removing a Conferencing Node instance

If you no longer need a Conferencing Node instance, you can permanently delete it from your Pexip Infinity platform.

To remove a Conferencing Node instance:

1. If you have not already done so, put the Conferencing Node into maintenance mode via the Pexip Infinity Administrator interface on the Management Node:
   a. Go to **Platform Configuration > Conferencing Nodes**.
   b. Select the Conferencing Node(s).
   c. From the **Action** menu at the top left of the screen, select **Enable maintenance mode** and then select **Go**.

While maintenance mode is enabled, this Conferencing Node will not accept any new conference instances.

d. Wait until any existing conferences on that Conferencing Node have finished. To check, go to **Status > Live View**.

2. Delete the Conferencing Node from the Management Node:
   a. Go to **Platform Configuration > Conferencing Nodes** and select the Conferencing Node.
   b. Select the check box next to the node you want to delete, and then from the **Action** drop-down menu, select **Delete selected Conferencing Nodes** and then select **Go**.
3. Delete the Conferencing Node instance on GCP:
   a. From the GCP project console, go to **Compute Engine > VM Instances**.
   b. Check the instance you want to permanently remove.
   c. Select **Delete** to remove the instance.