



# **Pexip Infinity and Google Cloud Platform**

## **Deployment Guide**

**Software Version 32**

**Document Version 32.a**

**June 2023**

**]pexip[**

# Contents

<b>Introduction</b>	<b>4</b>
<b>Deployment guidelines</b>	<b>5</b>
Deployment options	5
Recommended instance types and call capacity guidelines	6
Security and SSH keys	6
Google Cloud VPN for private/hybrid cloud deployments	6
IP addressing	6
Assumptions and prerequisites	7
<b>Configuring your Google VPC network</b>	<b>8</b>
Google Cloud VPN for private / hybrid cloud deployments	8
Enabling communication between Pexip Infinity nodes	8
Inter-node communication requirements for multiple VPCs	9
Controlling access to the Management Node	9
Controlling access to Conferencing Nodes for installation/provisioning	10
Controlling access to Conferencing Nodes for conference participants	11
<b>Obtaining and preparing disk images for GCE Virtual Machines</b>	<b>13</b>
Obtaining the Pexip disk images	13
Uploading disk images to Google Cloud Storage	13
Preparing custom disk images	13
Prepare a Management Node image	13
Prepare a Conferencing Node image	14
<b>Deploying a Management Node in Google Cloud Platform</b>	<b>16</b>
<b>Initial platform configuration — GCP</b>	<b>20</b>
Accessing the Pexip Infinity Administrator interface	20
Configuring the Pexip Infinity platform	20
Next step	21
<b>Deploying a Conferencing Node in Google Cloud Platform</b>	<b>22</b>
Deploying the VM instance in GCP	22
Generating, downloading and deploying the configuration file	25
<b>Configuring dynamic bursting to the Google Cloud Platform (GCP)</b>	<b>28</b>
Configuring your system for dynamic bursting to GCP	28
Firewall addresses/ports required for access to the GCP APIs for cloud bursting	28
Setting up your bursting nodes in GCP and enabling bursting in Pexip Infinity	28
Configuring a GCP role, permissions and service account for controlling overflow nodes	29
Configuring the bursting threshold	30

Manually starting an overflow node .....	31
Converting between overflow and "always on" GCP Conferencing Nodes .....	31
<b>Managing Google Compute Engine VM instances .....</b>	<b>32</b>
Temporarily removing (stopping) a Conferencing Node instance .....	32
Reinstating (restarting) a stopped Conferencing Node instance .....	32
Permanently removing a Conferencing Node instance .....	32

# 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, or allow Pexip Infinity to handle this automatically for you via its dynamic bursting capabilities.

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.

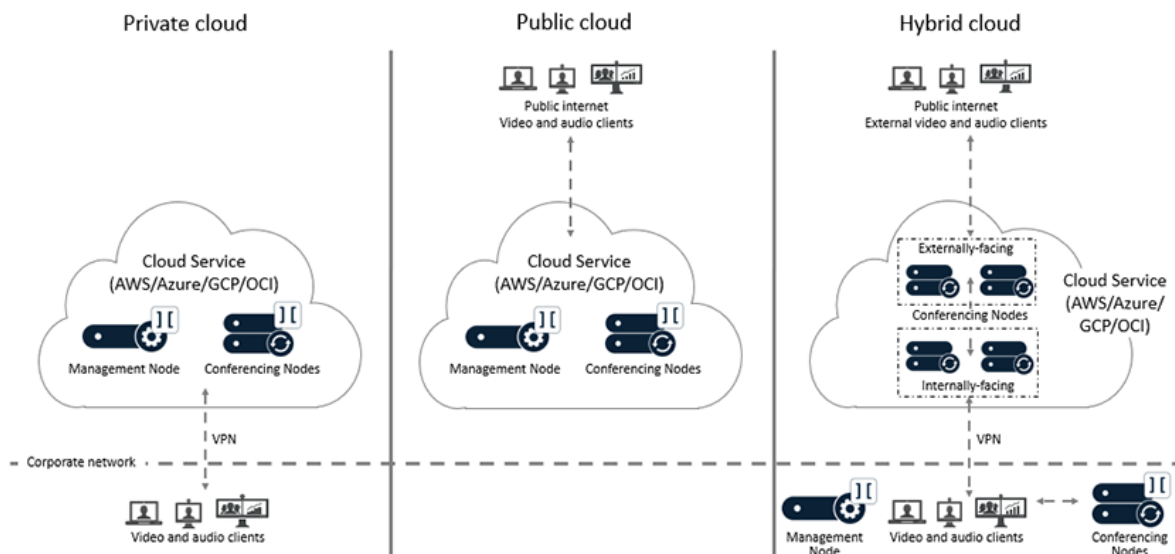
This flowchart provides an overview of the basic steps involved in deploying the Pexip Infinity platform on Azure:



## 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). You may also want to consider dynamic bursting, where the GCP-hosted Conferencing Nodes are only started up and used when you have reached capacity on your on-premises 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.

If available in your region, we recommend Ice Lake as the minimum CPU platform.

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

- **Management Node:** a machine type with 4 vCPUs such as n2-highcpu-4 (or larger such as n2-highcpu-8).
- **Transcoding Conferencing Nodes:** a machine type with 8 vCPUs 8 GB memory such as n2-highcpu-8 (or larger such as n2-highcpu-16 or n2-highcpu-32).
- **Proxying Edge Nodes:** a machine type with 4 vCPUs such as n2-highcpu-4 (or larger such as n2-highcpu-8).


This should provide capacity for approximately 16 HD / 35 SD / 300 audio-only calls per Transcoding Conferencing Node.

We recommend that you do **not** use AMD instance types.

For all available machine types see: [https://cloud.google.com/compute/pricing#predefined\\_machine\\_types](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"). You can also use other key types than rsa, such as ed25519. For more information about using and formatting SSH keys for GCP, see <https://cloud.google.com/compute/docs/instances/access-overview> and <https://cloud.google.com/source-repositories/docs/authentication#ssh>.

## 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 > 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

These 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.

- i** 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 **Other protocols and ports** from "icmp" to "icmp,esp".



### Protocols and ports ?

☐ Allow all

☒ Specified protocols and ports

☒ TCP

Ports

0-65535

E.g. 20, 50-60

☒ UDP

Ports

0-65535

E.g. all

☒ Other

Protocols \*

icmp,esp

Separate multiple protocols by commas, e.g. ah, sctp

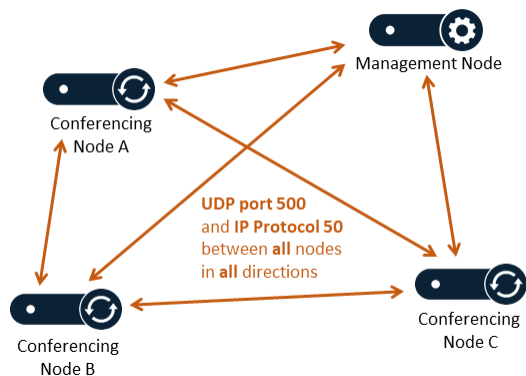
5. Select Save.

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.

## Inter-node communication requirements for multiple VPCs

In a basic deployment, your Pexip Infinity platform will be deployed within a single VPC.

In larger deployments you may choose to deploy your Conferencing Nodes across multiple VPCs — in which case there must be a directly routable path (no NAT) between all nodes that allows **UDP port 500 (IKE)**, and **IP Protocol 50 (IPsec ESP)** to pass between all nodes in both directions.



## 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 <i>Ingress</i> .
Action on match	Select <i>Allow</i> .
Targets	Select <i>Specified target tags</i> .
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 <a href="#">Deploying a Management Node in Google Cloud Platform</a> ).
Source filter	Select <i>IP ranges</i> .
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 <code>tcp:443</code>  Note that you may need to include <code>tcp:22</code> to allow SSH access if you intend to restrict or remove the <code>default-allow-ssh</code> rule.

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**.

- 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 <i>Ingress</i> .
Action on match	Select <i>Allow</i> .
Targets	Select <i>Specified target tags</i> .
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 <a href="#">Deploying a Conferencing Node in Google Cloud Platform</a> ).
Source filter	Select <i>IP ranges</i> .
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 <b>tcp:8443</b>

- 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:

- From the GCP project console, go to **VPC network > Firewall rules**.
- Select **Create firewall rule**.
- 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 <i>Ingress</i> .
Action on match	Select <i>Allow</i> .
Targets	Select <i>Specified target tags</i> .
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 <a href="#">Deploying a Conferencing Node in Google Cloud Platform</a> ).
Source filter	Select <i>IP ranges</i> .
Source IP ranges	Enter <b>0.0.0.0/0</b>  For a private deployment, the Source IP ranges should be restricted to the corporate intranet IP addresses.
Protocols and ports	Enter <b>tcp:80; tcp:443; tcp:1720; tcp:5060; tcp:5061; tcp:33000-39999; tcp:40000-49999; udp:1719; udp:33000-39999; udp:40000-49999</b>  Note that if you have enabled SIP UDP then <b>udp:5060</b> must also be included.

- Select **Create**.

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

[Ingress](#) [Egress](#)

<input type="checkbox"/> Name	Targets	Source filters	Protocols / ports	Action	Priority	Network
<input type="checkbox"/> default-allow-http	http-server	IP ranges: 0.0.0.0/0	tcp:80	Allow	1000	default
<input type="checkbox"/> default-allow-https	https-server	IP ranges: 0.0.0.0/0	tcp:443	Allow	1000	default
<input type="checkbox"/> pexip-allow-conferencing	pexip-conferencing	IP ranges: 0.0.0.0/0	tcp:80, tcp:443, 8 more ▼	Allow	1000	default
<input type="checkbox"/> pexip-allow-management	pexip-management	IP ranges: 81.143.209.108/32	tcp:443	Allow	1000	default
<input type="checkbox"/> pexip-allow-provisioning	pexip-provisioning	IP ranges: 81.143.209.108/32	tcp:8443	Allow	1000	default
<input type="checkbox"/> default-allow-icmp	Apply to all	IP ranges: 0.0.0.0/0	icmp	Allow	65534	default
<input type="checkbox"/> default-allow-internal	Apply to all	IP ranges: 10.128.0.0/9	tcp:0-65535, udp:0-65535, 2 more ▼	Allow	65534	default
<input type="checkbox"/> default-allow-rdp	Apply to all	IP ranges: 0.0.0.0/0	tcp:3389	Allow	65534	default
<input type="checkbox"/> 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\_v32\_GCP\_pxMgr\_<build>.tar.gz** for the Management Node.
- **Pexip\_Infinity\_v32\_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 **Cloud Storage > Browser**.
  - b. Select **Create Bucket**.
  - c. Enter a **Name** (for example, "pexip-v32"), 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>
  - d. Select **Create**.
2. Upload the Pexip images to the new bucket:
  - a. Select the new bucket e.g. pexip-v32.
  - 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:

Buckets / pexip-v20

<input type="checkbox"/> Name	Size	Type	Storage class	Last modified	Public access ⓘ	Encryption ⓘ	Retention expiry date ⓘ	Holds ⓘ
<input type="checkbox"/> Pexip_Infinity_v20_GCP_ConfNode_45400.tar.gz	517.77 MB	application/x-gzip	Regional	12/02/2019, 02:36:25 UTC	Not public	Google-managed key	–	None
<input type="checkbox"/> Pexip_Infinity_v20_GCP_pxMgr_45400.tar.gz	1.04 GB	application/x-gzip	Regional	12/02/2019, 02:38:43 UTC	Not public	Google-managed key	–	None

## 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-v32".
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-v32.
6. Select **Create**.

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

Compute Engine

← Create an image

Name ⓘ  
pexip-mgr-v20

Family (Optional) ⓘ

Description (Optional)

Labels ⓘ (Optional)  
[+ Add label](#)

Encryption  
Data is encrypted automatically. Select an encryption key management solution.  
☒ Google-managed key  
No configuration required  
☐ Customer-managed key  
Manage via Google Cloud Key Management Service  
☐ Customer-supplied key  
Manage outside Google Cloud

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. [Learn more](#)  
☒ pexip-v20/Pexip\_Infinity\_v20\_GCP\_pxMgr\_45400.tar.gz [Browse](#)

You will be billed for this image. [Compute Engine pricing](#) ⓘ

[Create](#) [Cancel](#)

## 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-v32".
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-v32.
6. Select **Create**.

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

Compute Engine

VM instances

Instance groups

Instance templates

Sole tenant nodes

Disks

Snapshots

Images

TPUs

Committed use discounts

Metadata

Health checks

Zones

Network endpoint groups

Operations

Security scans

Settings

Create an image

Name

pexip-node-v20

Family (Optional)

Description (Optional)

Labels (Optional)

+ Add label

Encryption

Data is encrypted automatically. Select an encryption key management solution.

☒ Google-managed key

No configuration required

☐ Customer-managed key

Manage via Google Cloud Key Management Service

☐ Customer-supplied key

Manage outside Google Cloud

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. [Learn more](#)

☒ pexip-v20/Pexip\_Infinity\_v20\_GCP\_ConfNode\_45400.tar.gz

Browse

You will be billed for this image. [Compute Engine pricing](#)

Create

Cancel

© 2023 Pexip AS

Version 32.a June 2023

Page 15 of 33

# 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 <b>4 vCPUs</b> (n1-standard-4).
Boot disk	Select the Management Node custom disk image: <ol style="list-style-type: none"> <li>a. Select <b>Change</b>.</li> <li>b. Select <b>Custom images</b>.</li> <li>c. Select your GCP project.</li> <li>d. Select the Management Node custom disk image, e.g. "pexip-mgr-v32".</li> <li>e. Select a <b>Boot disk type</b> of <i>SSD persistent disk</i>.</li> <li>f. Select <b>Select</b>.</li> </ol>
Identity and API access	For <b>Service account</b> , select <i>No service account</i> .
Networking:	Expand the <b>Advanced options</b> section and open the <b>Networking</b> section.
Network tags	Assign a <b>Network tag</b> to the instance, for example "pexip-management".  This is the tag that should be applied to your Management Node firewall rule (see <a href="#">Controlling access to the Management Node</a> ).
Networking:	You must decide whether or not to assign an external IP address to the instance.
External IP	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. <ol style="list-style-type: none"> <li>a. Expand the <b>Advanced options</b> section and open the <b>Networking</b> section.</li> <li>b. In the <b>Network interfaces</b> field, select the default interface to open the Network interface dialog.</li> <li>c. Select a <b>Subnetwork</b> 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).</li> <li>d. Select an appropriate <b>External IP</b>:               <ul style="list-style-type: none"> <li>■ <b>None</b>: no external IP address will be assigned. Use this where the instance does not need to have a publicly-accessible IP address.</li> <li>■ <b>Create IP address</b>: 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.</li> <li>■ <b>&lt;external address&gt;</b>: you can select a specific static external address if you have already created one in advance.</li> </ul> </li> </ol> <p>Do <b>not</b> select <i>Ephemeral</i> — if you stop and restart the instance a new address will be assigned.</p>



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. Open the **Security** section and then open the **Manage Access** section.
- b. Select **Add item** to add your own, existing SSH key. This produces a text box. Copy the contents of your public SSH key file and paste them into the text box. Modify the username element at the end of the key to "admin" or "admin@<domain>" if necessary.

See [Security and SSH keys](#) for more information.

---

**Compute Engine**

**VM instances**

Instance groups

Instance templates

Disks

Snapshots

Images

Committed use discounts

Metadata

Health checks

Zones

Operations

Quotas

Settings

**Create an instance**

**Name** ?

pexipmanager

**Zone** ?

europe-west1-d

**Machine type**

4 vCPUs 15 GB memory [Customize](#)

[Upgrade your account](#) to create instances with up to 64 cores

**Boot disk** ?

New 99 GB standard persistent disk

Image

pexip-mgr-v20 [Change](#)

**Identity and API access** ?

**Service account** ?

No service account

**Access scopes** ?

Select a service account to enable API access

**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-management

**Network interfaces** ?

default default (10.132.0.0/20)

[+ Add network interface](#)

**To create another network interface you need to have a new network first.**

[Less](#)

Your Free Trial credits, if available, will be used for this instance

[Create](#) [Cancel](#)

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:

IP address	The IP address must match the internal IP address allocated by GCE (the default should be correct).
Network mask	 The Network mask must be <b>255.255.255.255</b> .
Gateway	The Gateway must be the internal IP address of the gateway for the GCE region (the default should be correct).
Hostname Domain suffix	Enter your required Hostname and Domain suffix for the Management Node.
DNS servers	Configure one or more DNS servers. You must override the default values if it is a private deployment.
NTP servers	Configure one or more NTP servers. You must override the default values if it is a private deployment.
Web administration username Password	Set the Web administration username and password.
Enable incident reporting	Select whether or not to Enable incident reporting.
Send deployment and usage statistics to Pexip	Select whether or not 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 configure your Pexip Infinity platform licenses, VMRs, aliases, locations etc. as described in [Initial platform configuration — GCP](#) before you go on to add Conferencing Nodes.

## Initial platform configuration — GCP

After you have run the installation wizard, you must perform some preliminary configuration of the Pexip Infinity platform before you can deploy a Conferencing Node.

This section lists the configuration required, and provides a summary of each step with a link to further information.

All configuration should be done using the Pexip Infinity Administrator interface.

**i** **No changes** should be made to any Pexip VM via the terminal interface (other than as described when running the initial Pexip installation wizard) unless directed to do so by Pexip support. This includes (but is not limited to) changes to the time zone, changes to IP tables, configuration of Ethernet interfaces, or the installation of any third-party code/applications.

### Accessing the Pexip Infinity Administrator interface

The Pexip Infinity Administrator interface is hosted on the Management Node. To access this:

1. Open a web browser and type in the IP address or DNS name that you assigned to the Management Node using the installation wizard (you may need to wait a minute or so after installation is complete before you can access the Administrator interface).
2. Until you have uploaded appropriate TLS certificates to the Management Node, your browser may present you with a warning that the website's security certificate is not trusted. You should proceed, but upload appropriate TLS certificates to the Management Node (and Conferencing Nodes, when they have been created) as soon as possible.

The **Pexip Infinity Conferencing Platform** login page will appear.

3. Log in using the web administration username and password you set using the installation wizard.

You are now ready to begin configuring the Pexip Infinity platform and deploying Conferencing Nodes.

As a first step, we strongly recommend that you configure at least 2 additional NTP servers or NTP server pools to ensure that log entries from all nodes are properly synchronized.

It may take some time for any configuration changes to take effect across the Conferencing Nodes. In typical deployments, configuration replication is performed approximately once per minute. However, in very large deployments (more than 60 Conferencing Nodes), configuration replication intervals are extended, and it may take longer for configuration changes to be applied to all Conferencing Nodes (the administrator log shows when each node has been updated).

Brief details of how to perform the initial configuration are given below. For complete information on how to configure your Pexip Infinity solution, see the Pexip Infinity technical documentation website at [docs.pexip.com](https://docs.pexip.com).

### Configuring the Pexip Infinity platform

This table lists the Pexip Infinity platform configuration steps that are required before you can deploy Conferencing Nodes and make calls.

Configuration step	Purpose
<b>1. Enable DNS</b> (System > DNS Servers)	<p>Pexip Infinity uses DNS to resolve the hostnames of external system components including NTP servers, syslog servers, SNMP servers and web proxies. It is also used for call routing purposes — SIP proxies, gatekeepers, external call control and conferencing systems and so on. The address of at least one DNS server must be added to your system.</p> <p>You will already have configured at least one DNS server when running the install wizard, but you can now change it or add more DNS servers.</p>

Configuration step	Purpose
<b>2. Enable NTP</b> (System > NTP Servers)	<p>Pexip Infinity uses NTP servers to obtain accurate system time. This is necessary to ensure correct operation, including configuration replication and log timestamps.</p> <p>We strongly recommend that you configure at least three distinct NTP servers or NTP server pools on all your host servers and the Management Node itself. This ensures that log entries from all nodes are properly synchronized.</p> <p>You will already have configured at least one NTP server when running the install wizard, but you can now change it or add more NTP servers.</p>
<b>3. Add licenses</b> (Platform > Licenses)	<p>You must install a system license with sufficient concurrent call capacity for your environment before you can place calls to Pexip Infinity services.</p>
<b>4. Add a system location</b> (Platform > Locations)	<p>These are labels that allow you to group together Conferencing Nodes that are in the same datacenter. You must have at least one location configured before you can deploy a Conferencing Node.</p>
<b>5. Upload TLS certificates</b> (Certificates > TLS Certificates)	<p>You must install TLS certificates on the Management Node and — when you deploy them — each Conferencing Node. TLS certificates are used by these systems to verify their identity to clients connecting to them.</p> <p>All nodes are deployed with self-signed certificates, but we strongly recommend they are replaced with ones signed by either an external CA or a trusted internal CA.</p>
<b>6. Add Virtual Meeting Rooms</b> (Services > Virtual Meeting Rooms)	<p>Conferences take place in Virtual Meeting Rooms and Virtual Auditoriums. VMR configuration includes any PINs required to access the conference. You must deploy at least one Conferencing Node before you can call into a conference.</p>
<b>7. Add an alias for the Virtual Meeting Room</b> (done while adding the Virtual Meeting Room)	<p>A Virtual Meeting Room or Virtual Auditorium can have more than one alias. Conference participants can access a Virtual Meeting Room or Virtual Auditorium by dialing any one of its aliases.</p>

## Next step

You are now ready to deploy a Conferencing Node — see [Deploying a Conferencing Node in Google Cloud Platform](#) for more information.

# Deploying a Conferencing Node in Google Cloud Platform

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

Creating a new Conferencing Node is a two-step process:

1. Deploying a new VM instance in GCP.
2. Configuring the VM with the details of the specific Conferencing Node being deployed, using a file generated from the Pexip Infinity Management Node.

## Deploying the VM instance in GCP

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](#).

Note that if you have upgraded your Pexip Infinity software, you need a Conferencing Node disk image for the software version you are currently running.

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 <b>8 vCPUs (n1-highcpu-8)</b> .  We recommend selecting a minimum CPU platform. Select <b>Customize</b> and then select the most modern platform available that does not incur a surcharge, typically <b>Intel Broadwell or later</b> .  For more information see <a href="#">Recommended instance types and call capacity guidelines</a> .
Boot disk	Select the Conferencing Node custom disk image: <ol style="list-style-type: none"><li>a. Select <b>Change</b>.</li><li>b. Select <b>Custom images</b>.</li><li>c. Select the Conferencing Node custom disk image, e.g. "pexip-node-v32".</li><li>d. Select <b>Select</b>.</li></ol> We strongly recommend SSDs for Conferencing Nodes. General VM processes (such as snapshots and backups) and platform upgrades will be faster with SSDs.
Identity and API access	For <b>Service account</b> , select <b>No service account</b> .
Networking:	Expand the <b>Advanced options</b> section and open the <b>Networking</b> section.
Network tags	Assign <b>Network tags</b> to the instance, for example "pexip-provisioning pexip-conferencing".  These are the tags that should be applied to your Conferencing Node firewall rules (see <a href="#">Controlling access to Conferencing Nodes for installation/provisioning</a> and <a href="#">Controlling access to Conferencing Nodes for conference participants</a> ).

---

Networking:	You must decide whether or not to assign an external IP address to the instance.
External IP	<p>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.</p> <ol style="list-style-type: none"><li>Expand the <b>Advanced options</b> section and open the <b>Networking</b> section.</li><li>In the <b>Network interfaces</b> field, select the default interface to open the Network interface dialog.</li><li>Select a <b>Subnetwork</b> 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).</li><li>Select an appropriate <b>External IP</b>:<ul style="list-style-type: none"><li><b>None</b>: no external IP address will be assigned. Use this where the instance does not need to have a publicly-accessible IP address.</li><li><b>Create IP address</b>: select this option to create a static external address. You can enter a <b>Name</b> for the address and GCP will allocate a static IP address.</li><li><b>&lt;external address&gt;</b>: you can select a specific static external address if you have already created one in advance.</li></ul></li></ol> <p>Do <b>not</b> select <i>Ephemeral</i> — if you stop and restart the instance a new address will be assigned.</p>
SSH keys	<p>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.</p> <p>The username element of the SSH key must be "admin" or "admin@&lt;domain&gt;". To apply an instance-level key:</p> <ol style="list-style-type: none"><li>Open the <b>Security</b> section and then open the <b>Manage Access</b> section.</li><li>Select <b>Add item</b> to add your own, existing SSH key. This produces a text box. Copy the contents of your public SSH key file and paste them into the text box. Modify the username element at the end of the key to "admin" or "admin@&lt;domain&gt;" if necessary.</li></ol> <p>See <a href="#">Security and SSH keys</a> for more information.</p>

---

Compute Engine

VM instances

Instance groups

Instance templates

Disks

Snapshots

Images

Committed use discounts

Metadata

Health checks

Zones

Operations

Quotas

Settings

← Create an instance

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-v20

Change

Identity and API access ?

Service account ?

No service account

Access scopes ?

Select a service account to enable API access

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 as described below.

## Generating, downloading and deploying the configuration file

1. From the Pexip Infinity Administrator interface, go to **Platform > Conferencing Nodes** and select **Add Conferencing Node**.
2. You are now asked to provide the network configuration to be applied to the Conferencing Node, by completing the following fields:

Option	Description
Name	Enter the name to use when referring to this Conferencing Node in the Pexip Infinity Administrator interface.
Description	An optional field where you can provide more information about the Conferencing Node.
Role	<p>This determines the Conferencing Node's role:</p> <ul style="list-style-type: none"> <li>◦ <b>Proxying Edge Node:</b> a Proxying Edge Node handles all media and signaling connections with an endpoint or external device, but does not host any conferences — instead it forwards the media on to a Transcoding Conferencing Node for processing.</li> <li>◦ <b>Transcoding Conferencing Node:</b> a Transcoding Conferencing Node handles all the media processing, protocol interworking, mixing and so on that is required in hosting Pexip Infinity calls and conferences. When combined with Proxying Edge Nodes, a transcoding node typically only processes the media forwarded on to it by those proxying nodes and has no direct connection with endpoints or external devices. However, a transcoding node can still receive and process the signaling and media directly from an endpoint or external device if required.</li> </ul>
Hostname Domain	<p>Enter the hostname and domain to assign to this Conferencing Node. Each Conferencing Node and Management Node must have a unique hostname.</p> <p>The Hostname and Domain together make up the Conferencing Node's DNS name or FQDN. We recommend that you assign valid DNS names to all your Conferencing Nodes.</p>
IPv4 address	<p>Enter the IP address to assign to this Conferencing Node when it is created.</p> <p>This should be the GCE Internal IP address of the new VM instance.</p>
Network mask	<p>Enter the IP network mask to assign to this Conferencing Node.</p> <p>For GCP you should always enter <b>255.255.255.255</b></p> <p>Note that IPv4 address and Network mask apply to the eth0 interface.</p>
Gateway IPv4 address	<p>Enter the IP address of the default gateway to assign to this Conferencing Node.</p> <p>This is the default gateway address for the region in which the node is deployed. See <a href="https://cloud.google.com/compute/docs/vpc/#subnet-ranges">https://cloud.google.com/compute/docs/vpc/#subnet-ranges</a> for a table of regions and default gateway addresses.</p> <p>Note that the Gateway IPv4 address is not directly associated with a network interface, except that the address entered here lies in the subnet in which either eth0 or eth1 is configured to use. Thus, if the gateway address lies in the subnet in which eth0 lives, then the gateway will be assigned to eth0, and likewise for eth1.</p>
Secondary interface IPv4 address	Leave this option blank as dual network interfaces are not supported on Conferencing Nodes deployed in public cloud services.

Option	Description
Secondary interface network mask	<p>Leave this option blank as dual network interfaces are not supported on Conferencing Nodes deployed in public cloud services.</p> <p>Note that <b>Secondary interface IPv4 address</b> and <b>Secondary interface network mask</b> apply to the eth1 interface.</p>
System location	<p>Select the physical location of this Conferencing Node. A system location should not contain a mixture of proxying nodes and transcoding nodes.</p> <p>If the system location does not already exist, you can create a new one here by clicking  to the right of the field. This will open up a new window showing the <b>Add System Location</b> page.</p>
SIP TLS FQDN	A unique identity for this Conferencing Node, used in signaling SIP TLS Contact addresses.
TLS certificate	The TLS certificate to use on this node. This must be a certificate that contains the above <b>SIP TLS FQDN</b> . Each certificate is shown in the format <subject name> (<issuer>).
IPv6 address	The IPv6 address for this Conferencing Node. Each Conferencing Node must have a unique IPv6 address.
Gateway IPv6 address	<p>The IPv6 address of the default gateway.</p> <p>If this is left blank, the Conferencing Node listens for IPv6 Router Advertisements to obtain a gateway address.</p>
IPv4 static NAT address	<p>Configure the Conferencing Node's static NAT address, if you have assigned a public/external IP address to the instance.</p> <p>Enter the External IP address allocated by GCE for the VM instance.</p>
Static routes	From the list of <b>Available Static routes</b> , select the routes to assign to the node, and then use the right arrow to move the selected routes into the <b>Chosen Static routes</b> list.
Enable distributed database	This should usually be enabled (checked) for all Conferencing Nodes that are expected to be "always on", and disabled (unchecked) for nodes that are expected to only be powered on some of the time (e.g. nodes that are likely to only be operational during peak times).
Enable SSH	<p>Determines whether this node can be accessed over SSH.</p> <p><i>Use Global SSH setting:</i> SSH access to this node is determined by the global <b>Enable SSH</b> setting (<b>Platform &gt; Global Settings &gt; Connectivity &gt; Enable SSH</b>).</p> <p><i>Off:</i> this node cannot be accessed over SSH, regardless of the global <b>Enable SSH</b> setting.</p> <p><i>On:</i> this node can be accessed over SSH, regardless of the global <b>Enable SSH</b> setting.</p> <p>Default: <i>Use Global SSH setting</i>.</p>

3. Select **Save**.

4. You are now asked to complete the following fields:

Option	Description
Deployment type	Select <b>Generic (configuration-only)</b> .
SSH password	<p>Enter the password to use when logging in to this Conferencing Node's Linux operating system over SSH. The username is always <b>admin</b>.</p> <p>Logging in to the operating system is required when changing passwords or for diagnostic purposes only, and should generally be done under the guidance of your Pexip authorized support representative. In particular, do not change any configuration using SSH — all changes should be made using the Pexip Infinity Administrator interface.</p>

5. Select **Download**.

A message appears at the top of the page: "The Conferencing Node image will download shortly or click on the following link".

After a short while, a file with the name **pexip-`<hostname>`.`<domain>`.xml** is generated and downloaded.

Note that the generated file is only available for your current session so you should download it immediately.

6. Browse to **<https://<conferencing-node-ip>:8443/>** and use the form provided to upload the 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.

The Conferencing Node will apply the configuration and reboot. After rebooting, it will connect to the Management Node in the usual way.

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 becomes available, the Management Node reports 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 also 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.

# Configuring dynamic bursting to the Google Cloud Platform (GCP)

Pexip Infinity deployments can burst into the Google Cloud Platform (GCP) cloud when primary conferencing capabilities are reaching their capacity limits, thus providing additional temporary Conferencing Node resources.

This provides the ability to dynamically expand conferencing capacity whenever scheduled or unplanned usage requires it. The GCP cloud Conferencing Nodes instances are only started up when required and are automatically stopped again when capacity demand normalizes, ensuring that GCP costs are minimized.

For complete information about dynamic bursting, see [Dynamic bursting to a cloud service](#).

## Configuring your system for dynamic bursting to GCP

These instructions assume that you already have a working Pexip Infinity platform, including one or more primary (always on) Conferencing Nodes in one or more system locations. These existing Conferencing Nodes can be deployed using whichever platform or hypervisor you prefer.

### Firewall addresses/ports required for access to the GCP APIs for cloud bursting

Access to the GCP APIs for cloud bursting is only required from the Management Node.

The Management Node connects to [www.googleapis.com](https://www.googleapis.com) over HTTPS port 443.

## Setting up your bursting nodes in GCP and enabling bursting in Pexip Infinity


You must deploy in GCP the Conferencing Nodes that you want to use for dynamic bursting, and then configure the Pexip Infinity location containing those nodes as the overflow destination for the locations that contain your primary (always on) Conferencing Nodes:

1. In Pexip Infinity, configure a new system location for media overflow e.g. "GCP burst", that will contain your bursting Conferencing Nodes.  
(Note that system locations are not explicitly configured as "primary" or "overflow" locations. Pexip Infinity automatically detects the purpose of the location according to whether it contains Conferencing Nodes that may be used for dynamic bursting.)
2. In GCP, set up the service account, roles and permissions that the Pexip Infinity Management Node will use to log in to GCP to start and stop the node instances.  
See [Configuring a GCP role, permissions and service account for controlling overflow nodes](#) for more information.
3. Deploy in GCP the Conferencing Nodes that you want to use for dynamic bursting. Deploy these nodes in the same manner as you would for "always on" usage (see [Deploying a Conferencing Node in Google Cloud Platform](#)), except:
  - a. Apply to each cloud VM node instance to be used for conference bursting a label with a **Key** of `pexip-cloud` and an associated **Value** set to the **Tag value** that is shown in the **Cloud Bursting** section on the **Platform > Global Settings** page (the **Tag value** is the hostname of your Management Node).  
This tag indicates which VM nodes will be started and shut down dynamically by your Pexip system.
  - b. When adding the Conferencing Node within Pexip Infinity:
    - i. Assign the Conferencing Node to the overflow system location (e.g. "GCP burst").
    - ii. Disable (uncheck) the **Enable distributed database** setting (this setting should be disabled for any nodes that are not expected to always be available).
  - c. After the Conferencing Node has successfully deployed, manually stop the node instance on GCP.
4. In Pexip Infinity, go to **Platform > Global Settings > Cloud Bursting**, enable cloud bursting and then configure your bursting threshold, minimum lifetime and other appropriate settings for GCP:

Option	Description
--------	-------------

Enable bursting to the cloud	Select this option to instruct Pexip Infinity to monitor the system locations and start up / shut down overflow Conferencing Nodes hosted in your cloud service when in need of extra capacity.
Bursting threshold	<p>The bursting threshold controls when your dynamic overflow nodes in your cloud service are automatically started up so that they can provide additional conferencing capacity. When the number of additional HD calls that can still be hosted in a location reaches or drops below the threshold, it triggers Pexip Infinity into starting up an overflow node in the overflow location.</p> <p>See <a href="#">Configuring the bursting threshold</a> for more information.</p>
Tag name and Tag value	<p>These read-only fields indicate the tag name (always <code>pexip-cloud</code>) and associated tag value (the hostname of your Management Node) that you must assign to each of your cloud VM node instances that are to be used for dynamic bursting.</p> <p>In some circumstances the <b>Tag value</b> may auto populate as "unknown" instead of the hostname, in which case you should also use "unknown" on your cloud VM node instances.</p>
Minimum lifetime	An overflow cloud bursting node is automatically stopped when it becomes idle (no longer hosting any conferences). However, you can configure the <b>Minimum lifetime</b> for which the bursting node is kept powered on. By default this is set to 50 minutes, which means that a node is never stopped until it has been running for at least 50 minutes. If your service provider charges by the hour, it is more efficient to leave a node running for 50 minutes — even if it is never used — as that capacity can remain on immediate standby for no extra cost. If your service provider charges by the minute you may want to reduce the <b>Minimum lifetime</b> .
Cloud provider	Select <i>GCP</i> .
GCP private key	<p>The PEM-formatted private key for the Google Cloud Platform service account that the Pexip Infinity Management Node will use to log in to GCP to start and stop the node instances.</p> <p>If you have created and then downloaded a JSON private key file from the GCP service account, open the file in a plain text editor and copy everything between the double quotes following <code>"private_key"</code> : including the <code>-----BEGIN PRIVATE KEY-----</code> and <code>-----END PRIVATE KEY-----</code> header and footer, and then paste it into this field.</p> <p>After you have saved the settings, the private key will be stored in an encrypted format and subsequently displayed as asterisks.</p>
GCP service account ID	The <b>Service account ID</b> of the service account you have set up in GCP, for example, <code>bursting@example.com.iam.gserviceaccount.com</code> .
GCP project ID	The ID of the GCP project containing the bursting nodes.

- Go to **Platform > Locations** and configure the system locations that contain your "always on" Conferencing Nodes (the nodes/locations that initially receive calls) so that they will overflow to your new "GCP burst" location when necessary. When configuring your principal "always on" locations, you should normally set the **Primary overflow location** to point at the bursting location containing your overflow nodes, and the **Secondary overflow location** should normally only point at an always-on location.

-  Nodes in a bursting location are only automatically started up if that location is configured as a **Primary overflow location** of an always-on location that has reached its capacity threshold. This means that if a bursting location is configured as a **Secondary overflow location** of an always-on location, then those nodes can only be used as overflow nodes if they are already up and running (i.e. they have already been triggered into starting up by another location that is using them as its **Primary overflow location**, or you have used some other external process to start them up manually).

We recommend that you do not mix your "always on" Conferencing Nodes and your bursting nodes in the same system location.

## Configuring a GCP role, permissions and service account for controlling overflow nodes

Within GCP you must set up a service account to be used by Pexip Infinity to start up and shut down the Conferencing Node overflow instances.

First, you need to create a GCP role and its associated permissions:

- From the GCP project console, go to **IAM & Admin > Roles**.
- Select **+ Create Role**.

3. Configure a **Title** and **ID** e.g. "Bursting".
4. Select **+ Add Permissions**.
5. From the list of permissions, select **compute.instances.list**, **compute.instances.start**, **compute.instances.stop** and **compute.zoneOperations.list**.  
These permissions allow the service account to stop and start VM instances.
6. Select **Create**.

Next, you need to create a service account and apply the role to the account:

1. From the GCP project console, go to **IAM & Admin > Service Accounts**.
2. Select **+ Create Service Account**.
3. Configure a **Service account name** e.g. "bursting" and add a **Description**.  
The **Service account ID** is automatically created, based on the name you enter, and is what you need to configure in the **GCP service account ID** field in Pexip Infinity.
4. Select **Create**.
5. For **Service account permissions**, select the "Bursting" role you created above.
6. Select **Continue**.
7. On the next page you can optionally grant users access to the service account.
8. Select **+ Create Key** to create a private key for the account.
9. On the **Create key** page, select a key type of **JSON**.
10. Select **Create**.  
A JSON file that includes your private key is generated and saved to your local computer.  
Select **Close** to dismiss the message saying that the private key has been saved to your computer.
11. Select **Done** at the bottom of the **Create service account** page.

You can now use this service account and private key when configuring bursting in Pexip Infinity (**Platform > Global Settings > Cloud Bursting**).

## Configuring the bursting threshold

When enabling your platform for cloud bursting the most important decision you must make is the level at which to set the bursting threshold:

- The bursting threshold controls when your dynamic overflow nodes in your cloud service are automatically started up so that they can provide additional conferencing capacity. When the number of additional HD calls that can still be hosted in a location reaches or drops below the threshold, it triggers Pexip Infinity into starting up an overflow node in the overflow location.  
For example, setting the threshold to 5 means that when there are 5 or fewer HD connections still available in a location, an overflow node will be started up.
- When an overflow location reaches the bursting threshold i.e. the number of additional HD calls that can still be hosted on the Conferencing Nodes in the overflow location reaches the threshold, another overflow node in that location is started up, and so on.  
Note that the current number of free HD connections in the original location is ignored when deciding if the overflow location needs to overflow further — however, new calls will automatically use any available media resource that has become available within the original principal location.
- The bursting threshold is a global setting — it applies to every system location in your deployment.
- Note that it takes approximately 5 minutes for a dynamic node instance to start up and become available for conference hosting. If your principal deployment reaches full capacity, and the overflow nodes have not completed initiating, any incoming calls during this period will be rejected with "capacity exceeded" messages. You have to balance the need for having standby capacity started up in time to meet the expected demand, against starting up nodes too early and incurring extra unnecessary costs.

## Manually starting an overflow node

If you know that your system will need additional capacity at a specific time due to a predictable or scheduled spike in demand, but do not want to wait for the bursting threshold to be triggered before starting up the overflow nodes, you can manually start up any of your overflow nodes.

- i** Do not manually start an overflow node too early. If you manually start up a node more than the **Minimum lifetime** minutes before the node is needed, it will most probably get automatically stopped again before it is used.

You can start overflow nodes via the management API or via the Administrator interface:

- **Via the management API:** the `cloud_node` status resource can be used to list all of the available overflow nodes, the `cloud_monitored_location` and `cloud_overflow_location` resources retrieve the current load on the primary locations and any currently active overflow locations respectively, and the `start_cloudnode` resource can be used to manually start up any overflow node. This means that a third-party scheduling system, for example, could be configured to start up the overflow nodes via the management API approximately 10 minutes before a large conference is due to start.

For example, let's assume that you have:

- a regular spike in conferencing capacity demand at 9:00am every morning
- an even usage of about 20% of that spike level during the rest of the day
- a 30:70 ratio between your "always on" capacity and your overflow cloud capacity

we would recommend:

- configuring a low bursting threshold, such as 10-20% of your "always on" capacity (i.e. if your "always on" capacity is 80 HD calls, then set the bursting threshold to 12)
- getting your scheduling system to call the API to manually start up all of your overflow cloud nodes at 8:50am on weekdays.
- **Via the Pexip Infinity Administrator interface:** go to **Status > Cloud Bursting** and select **Start** for the required nodes (the **Start** option is in the final column of the **Cloud overflow nodes** table).

## Converting between overflow and "always on" GCP Conferencing Nodes

If you need to convert an existing "always on" GCP Conferencing Node into an overflow node:

1. In GCP:
  - a. Apply to the instance a label with a **Key** of `pexip-cloud` and an associated **Value** set to the **Tag value** that is shown in the **Cloud bursting** section of the **Platform > Global Settings** page.
  - b. Manually stop the node instance on GCP.
2. In Pexip Infinity:
  - a. Change the system location of the Conferencing Node to the overflow system location (e.g. "GCP burst").
  - b. Disable the node's **Enable distributed database** setting.
    - i** You should avoid frequent toggling of this setting. When changing this setting on multiple Conferencing Nodes, update one node at a time, waiting a few minutes before updating the next node.

If you need to convert an existing GCP overflow Conferencing Node into an "always on" node:

1. In GCP:
  - a. Remove the label with a **Key** of `pexip-cloud` from the GCP instance.
  - b. Manually start the node instance on GCP.
2. In Pexip Infinity:
  - a. Change the system location of the Conferencing Node to a location other than the overflow system location.
  - b. Enable the node's **Enable distributed database** setting.
    - i** You should avoid frequent toggling of this setting. When changing this setting on multiple Conferencing Nodes, update one node at a time, waiting a few minutes before updating the next node.

# Managing Google Compute Engine VM instances

This section describes the common maintenance tasks for [stopping](#), [restarting](#) and [permanently removing](#) Conferencing Node 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 > 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 > 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 > 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 > 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.