



Pexip Infinity KVM Installation Guide

Introduction

This installation guide describes the minimum steps required to deploy and test a simple Pexip Infinity platform in a KVM environment.

- ❗ Full information on configuring and using Pexip Infinity is available:
 - on the [Pexip Infinity technical documentation website](#) (from where all documentation can also be downloaded in PDF format)
 - as online help, by clicking the **Help** link in the top right corner of the Pexip Infinity Administrator interface (available after the Management Node has been deployed).
- ❗ You must ensure you have completed all necessary platform-based [Planning and prerequisites](#) prior to installation.

Configuring KVM for Pexip Infinity

This section describes the KVM configuration required before you [install the Management Node](#) or [install a Conferencing Node](#).

Prerequisites

This guide assumes that you have a KVM environment already installed.

Pexip Infinity requires your KVM environment to include Linux kernel 3.10.0 or later, and QEMU 1.5.0 or later. This means the following distributions: Debian 8, RHEL 7, SLES 12, or Ubuntu 14.04 (or later, where appropriate).

You must ensure that:

- the KVM environment may be managed through **libvirt**
- the **virt-manager**, **virsh** and **virt-install** applications are available for use on the client system that is managing the KVM environment; to do this, run the command:

```
sudo apt-get install virt-manager libvirt-bin
```

The **libvirt** command line tools are used to import the disk image and create the VM. Note that we use the libvirt command line tools to perform the import as they provide greater control than Virtual Machine Manager. For more information, see <http://libvirt.org/drvqemu.html>.

Synchronizing time

- ⓘ All host servers **must** be synchronized with accurate time before you install the Pexip Infinity Management Node or Conferencing Nodes on them.
- NTP **must** also be enabled on the Management Node VM before you deploy any Conferencing Nodes (this is done during installation of the Management Node).
- 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 will ensure that log entries from all nodes are properly synchronized.
- Pexip Infinity Management Node and Conferencing Node VMs use the UTC timezone, and all logs are in UTC. Do not attempt to change the timezone on these systems.

Install the NTP package on your host server using the appropriate package management tool for your Linux distribution.

Installing the Management Node

Before installing the Management Node we recommend that you review the [KVM configuration prerequisites](#) and the .

Installation of the Management Node is a two-step process:

1. **Deploying the VM template:** this creates a new unconfigured instance of a generic Management Node VM. The process for deploying the template in a KVM hypervisor environment is described below.
2. **Running the installation wizard:** after deploying the Management Node template, the [installation wizard](#) allows you to enter the basic configuration details for your Management Node VM.

Deploying the Management Node template

To deploy a new instance of a Pexip Infinity Management Node on to a KVM host, you must:

1. [Download](#) the Management Node OVA image file and convert it for use with KVM.
2. Create a new volume on your KVM server and [upload](#) the disk image.
3. Create the Management Node [virtual machine](#).
Note that we use the libvirt command line tools to perform the import as they provide greater control than Virtual Machine Manager.
4. Enable the virtual machine for [automatic startup](#).

These steps are described in detail below.

Download the Pexip .ova image and convert it for use with KVM

You must download the Management Node OVA image file and then, to use it with KVM, you must convert it from VMDK to raw format.

From within your KVM environment:

1. Download the **Pexip Infinity OVA** file (Pexip_Infinity_v18_generic_pxMgr_<build>.ova) from www.pexip.com/software-download.
2. Unpack the .ova image, using the command:

```
tar xf Pexip_Infinity_v18_generic_pxMgr_<build>.ova
```

This unpacks a set of files including `pexipcumgr-disk01.vmdk`.
3. If required, you can verify the `pexipcumgr.mf` manifest file, using the command:

```
shasum -c pexipcumgr.mf
```
4. Convert the disk image from VMDK to raw, using the command:

```
qemu-img convert -O raw pexipcumgr-disk01.vmdk pexipcumgr-disk01.raw
```

(This conversion process can take several seconds.)

Create a new volume and upload the disk image

Next, you create a new volume on your KVM server and upload the converted disk image. From within your KVM environment:

1. Use **virsh** to create a new volume on your KVM server:

```
virsh --connect qemu://<hostname>/system vol-create-as <poolname> <volume_name> 99G --format raw
```

where:

`<hostname>` is the hostname of your KVM server. Note that you can omit the `<hostname>` if you are running **virsh** commands on the local server i.e. you can use `virsh --connect qemu:///system`.

`<poolname>` is the name of the storage pool in which to create the volume; typically you would use **default**. (To determine the storage pools available on the target system, use `virsh --connect qemu://<hostname>/system pool-list`.)

`<volume_name>` is the name of your new volume.

99G is the virtual size of the volume; always use 99G for a Management Node.

For example:

```
virsh --connect qemu://host1.example.com/system vol-create-as default pexip-manager 99G --format raw
```

This example creates a volume named **pexip-manager** of size 99 GB and format raw in the storage pool named **default**.

2. Upload the converted disk image to the newly created volume:

```
virsh --connect qemu://<hostname>/system vol-upload <volume_name> pexipmcumgr-disk01.raw --pool <poolname>
```

For example:

```
virsh --connect qemu://host1.example.com/system vol-upload pexip-manager pexipmcumgr-disk01.raw --pool default
```

This example uploads the **pexipmcumgr-disk01.raw** image to the newly created volume, **pexip-manager**, in the storage pool named **default**.

Create the virtual machine

After the disk image has been uploaded, you can create the virtual machine to use it.

Note that we use the libvirt command line tools to perform the import as they provide greater control than Virtual Machine Manager.

1. Identify the filesystem path of the newly uploaded disk image:

```
virsh --connect qemu://<hostname>/system vol-path <volume_name> --pool <poolname>
```

For example:

```
virsh --connect qemu://host1.example.com/system vol-path pexip-manager --pool default
```

This prints out the absolute path to the disk image file, for example:

```
/var/lib/libvirt/images/pexip-manager
```

This path is used in the **disk path** parameter in the next step.

2. Use the **virt-install** command line tool to create the virtual machine:

```
virt-install \
  --import \
  --hvm \
  --name=<vm_name> \
  --arch=x86_64 \
  --vcpus=2 \
  --ram=4096 \
  --cpu host \
  --os-type=linux \
  --connect=qemu://<hostname>/system \
  --virt-type kvm \
  --disk path=<image_file_path>,bus=virtio,format=raw,cache=none,io=native \
  --network bridge=br0,model=virtio \
  --memballoon virtio \
  --graphics vnc,listen=0.0.0.0,password=<password>
```

This creates a new VM (KVM domain) from the converted disk image.

The command options are described below (items in **bold** may be changed as necessary):

Option	Description
<code>--import</code>	Build guest domain around pre-installed disk image; do not attempt to install a new OS.
<code>--hvm</code>	Create a fully virtualized (i.e. not paravirtualized) VM.
<code>--name=<vm_name></code>	Name of the new VM, where <code><vm_name></code> is, for example, pexip-management-vm .
<code>--arch=x86_64</code>	CPU architecture of new VM (must be <code>x86_64</code>).

Option	Description
<code>--vcpus=2</code>	Number of CPUs allocated to new VM. Typically, this is 2 for the Management Node.
<code>--ram=4096</code>	Memory allocated to new VM (in megabytes).
<code>--cpu host</code>	Expose all host CPU capabilities to new VM (CPUID).
<code>--os-type=linux</code>	The guest OS is Linux.
<code>--connect=qemu://<hostname>/system</code>	Connect to KVM on the target system, where <hostname> is the hostname of your KVM server.
<code>--virt-type kvm</code>	Use KVM to host the new VM.
<code>--disk path=<image_file_path>, bus=virtio, format=raw, cache=none, io=native</code>	<ul style="list-style-type: none"> Define the location of the disk image file, where <image_file_path> is as determined in the previous step, for example <code>/var/lib/libvirt/images/pexip-manager</code>. Expose it to the guest on the virtio paravirtualized bus (as opposed to IDE/SCSI). Define the image file as being in raw format. Instruct the host system not to cache the disk contents in memory. Use the native IO backend to access the disk device.
<code>--network bridge=br0, model=virtio</code>	<ul style="list-style-type: none"> Create a network interface connected to the br0 bridge interface on the host. Expose it to the guest as a virtio paravirtualized NIC.
<code>--memballoon virtio</code>	Expose the virtio memory balloon to the guest.
<code>--graphics vnc, listen=0.0.0.0, password=<password></code>	Expose the graphical console over VNC, listening on 0.0.0.0 (i.e. all addresses on the target system) and with an access password of <password> .

You may receive a warning "Unable to connect to graphical console: virt-viewer not installed"; if so, this message can be safely ignored.

After the VM has been created, it may be managed using the Virtual Machine Manager desktop interface (virt-manager application) or via the command line interface (virsh).

The new node should start automatically. If it does not you can use the Virtual Machine Manager to start the node, or the CLI command:

```
virsh --connect qemu://<hostname>/system start <vm_name>
```


Note that you can list existing VMs by using `virsh --connect qemu://<hostname>/system list`

Enabling automatic startup

After deploying a new Management Node in KVM, you should enable automatic startup of that virtual machine (VM). In KVM, automatic startup is disabled by default for every new VM. This means that if the host server is powered down for any reason, when it restarts the VM will not restart and must be started manually.

You can only enable automatic startup after the Management Node has been deployed.


To enable automatic startup using Virtual Machine Manager:

1. Connect to the Virtual Machine Manager (virt-manager) that is managing the node's VM.
2. Select the node's VM and then, from the toolbar, select the **Show the virtual machine console and details** icon . A new window for that VM is opened.
3. If necessary, select **View > Details** to display the VM information.
4. From the sidebar menu, select **Boot Options**.
5. Select the **Start virtual machine on host boot up** check box.
6. Select **Apply**.

Running the installation wizard

To run the installation wizard, which configures the Management Node, you must open a console window on the Management Node VM.

Opening a console window

1. Connect to Virtual Machine Manager (virt-manager) that is managing the Management Node's VM.
2. Ensure that the Management Node VM is Running.
3. Select the new Management Node VM and then, from the toolbar, select the **Show the virtual machine console and details** icon .

A new window for that VM is opened.

4. If necessary, select **View > Console** to display the console.

Running the installation wizard

1. At the prompt, enter the username *admin*.

The display will read:

```
You are required to change your password immediately (root enforced)
Enter new UNIX password:
```

2. Create a password for the Management Node operating system by typing the password, pressing enter, retyping the password, and pressing enter again.
3. Ensure you record the password in a secure location. After you have finished running the installation wizard you will not need the password again unless you need to access the Management Node using SSH.

You are presented with another login prompt:

```
Running Pexip installation wizard...
[sudo] password for admin:
```

4. Log in again with the password you just created.

The Pexip installation wizard will begin.
5. Follow the prompts to set the following configuration for the Management Node.

If you press enter, the default value will be applied:

Setting	Default value	Multiple entries allowed?	Can be changed via Pexip Infinity Administrator interface?
IP address	192.168.0.100	No	No ‡
Network mask	255.255.255.0	No	No ‡
Gateway	192.168.0.1	No	No ‡
Hostname	<no default>	No	No ‡
Domain suffix	<no default>	No	No ‡
DNS servers	8.8.8.8	Yes, if separated by a space	Yes
NTP servers †	Two of the following: <ul style="list-style-type: none"> ○ 0.pexip.pool.ntp.org ○ 1.pexip.pool.ntp.org ○ 2.pexip.pool.ntp.org ○ 3.pexip.pool.ntp.org 	Yes, if separated by a space	Yes
Web administration username	admin	No	No ‡

Setting	Default value	Multiple entries allowed?	Can be changed via Pexip Infinity Administrator interface?
Web administration password	<no default>	No	Yes
Enable incident reporting (yes/no)	<no default>		Yes
Send deployment and usage statistics to Pexip (yes/no)	<no default>		Yes

† The NTP server must be accessible by the Management Node at the time the startup wizard is run. Installation will fail if the Management Node is unable to synchronize its time with an NTP server.

‡ After they have been configured, do not attempt to change these settings by any other means. To change these settings on server-based deployments, you must re-run the installation wizard.

The installation will begin and the Management Node will restart using the values you have configured.

Initial platform configuration

After you have run the installation wizard, you must perform some preliminary configuration of the Pexip Infinity platform before you can then 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, and configuration of Ethernet interfaces.

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 service 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.

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
1. Enable DNS (System Configuration > DNS Servers)	At least one DNS server must be added to your system. DNS is required when using hostnames instead of IP addresses when setting up NTP servers, syslog servers and VM managers. It is also used for call routing that depends on FQDNs. 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.

Configuration step	Purpose
<p>2. Enable NTP</p> <p>(System Configuration > NTP Servers)</p>	<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 3 distinct NTP servers or NTP server pools to ensure 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>
<p>3. Add licenses</p> <p>(Platform Configuration > Licenses)</p>	<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>
<p>4. Add a system location</p> <p>(Platform Configuration > Locations)</p>	<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>
<p>5. Upload TLS certificates</p> <p>(Platform Configuration > TLS Certificates)</p>	<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>
<p>6. Add Virtual Meeting Rooms</p> <p>(Service Configuration > Virtual Meeting Rooms)</p>	<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>
<p>7. Add an alias for the Virtual Meeting Room</p> <p>(done while adding the Virtual Meeting Room)</p>	<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>

Manually deploying a Conferencing Node on a KVM host

To manually deploy a new Conferencing Node onto a KVM host, you must:

1. Use the Pexip Infinity Administrator interface to [generate](#) and download the .ova image.
2. [Convert](#) the .ova image for use with KVM.
3. Create a new volume on your KVM server and [upload](#) the disk image.
4. Create the Conferencing Node [virtual machine](#).
Note that we use the libvirt command line tools to perform the import as they provide greater control than Virtual Machine Manager.
5. Enable the virtual machine for [automatic startup](#).

These steps are described in detail below.

Generate and download the .ova image

Before you start, ensure that you are currently using the same machine that you will subsequently want to use to upload the generated file on to your host server.


Use the Pexip Infinity Administrator interface to generate and download the .ova image:

1. Go to **Platform Configuration > Conferencing Nodes** and select **Add Conferencing Node**.
2. From the **Deployment type** field, select **Manual (KVM)**.
3. Select **Next**.
4. You are now asked to provide information regarding the CPUs and memory of the Conferencing Node VM, by completing the following fields:

Option	Description
Number of virtual CPUs to assign	Enter the number of virtual CPUs to assign to the Conferencing Node. We recommend no more than one virtual CPU per physical core, unless you are making use of CPUs that support hyperthreading.
System memory (in megabytes) to assign	Enter the amount of RAM (in megabytes) to assign to the Conferencing Node. The number entered must be a multiple of 4. We recommend 1024 MB (1 GB) RAM for each virtual CPU.

5. Select **Next**.
6. 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 that will be used to refer to this Conferencing Node in the Pexip Infinity Administrator interface.
Description	An optional field where you can provide more information about the Conferencing Node.

Option	Description
Role	<p>This determines the Conferencing Node's role:</p> <ul style="list-style-type: none"> ○ Proxying Edge Node: 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. ○ Transcoding Conferencing Node: 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.
Hostname Domain	<p>Enter the hostname and domain to be assigned 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	Enter the IP address to be assigned to this Conferencing Node when it is created.
Network mask	Enter the IP network mask to be assigned to this Conferencing Node.
Gateway IPv4 address	Enter the IP address of the default gateway to be assigned to this Conferencing Node.
Secondary interface IPv4 address	The optional secondary interface IPv4 address for this Conferencing Node. If configured, this interface is used for signaling and media communications to clients, and the primary interface is used for communication with the Management Node and other Conferencing Nodes.
Secondary interface network mask	The optional secondary interface network mask for this Conferencing Node.
System location	<p>Select the physical location of this Conferencing Node. A system location should not contain a mixture of Proxying Edge Nodes and Transcoding Conferencing 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 Add System Location 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 SIP TLS FQDN . Each certificate is shown in the format <code><subject name> (<issuer>)</code> .
IPv6 address	The IPv6 address for this Conferencing Node. Each Conferencing Node must have a unique IPv6 address.
Gateway IPv6 address	The IPv6 address of the default gateway.
IPv4 static NAT address	The public IPv4 address used by this Conferencing Node when it is located behind a NAT device. Note that if you are using NAT, you must also configure your NAT device to route the Conferencing Node's IPv4 static NAT address to its IPv4 address .
Static routes	From the list of Available Static routes , select the routes to assign to the node, and then use the right arrow to move the selected routes into the Chosen Static routes 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).

Option	Description
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 Enable SSH setting (Platform Configuration > Global Settings > Enable SSH).</p> <p><i>Off:</i> this node cannot be accessed over SSH, regardless of the global Enable SSH setting.</p> <p><i>On:</i> this node can be accessed over SSH, regardless of the global Enable SSH setting.</p> <p>Default: <i>Use Global SSH setting.</i></p>
SSH password	<p>Enter the password to be used when logging in to this Conferencing Node's Linux operating system over SSH. The username will always be <i>admin</i>.</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>

7. Select **Finish**.

You are taken to the **Manually Deploy Conferencing Node** page.

8. Select **Download Conferencing Node**.

This downloads a zip file with the name **pexip-`<hostname>`.`<domain>`.ova**.

Convert the .ova image for use with KVM

To use the Conferencing Node OVA image file with KVM, you must convert it from VMDK to raw format:

1. Copy the downloaded OVA file (named **pexip-`<hostname>`.`<domain>`.ova**) to the server running KVM.
2. Unpack the .ova image, using the command:

```
tar xf pexip-<hostname>.<domain>.ova
```

This unpacks a set of files including **pexip-disk01.vmdk**.

3. Convert the disk image from VMDK to raw, using the command:

```
qemu-img convert -O raw pexip-disk01.vmdk pexip-disk01.raw
```

(This conversion process can take several seconds.)

Create a new volume and upload the disk image

Next, you create a new volume on your KVM server and upload the converted disk image. From within your KVM environment:

1. Use **virsh** to create a new volume on your KVM server:

```
virsh --connect qemu://<hostname>/system vol-create-as <poolname> <volume_name> 49G --format raw
```

where:

`<hostname>` is the hostname of your KVM server. Note that you can omit the `<hostname>` if you are running **virsh** commands on the local server i.e. you can use `virsh --connect qemu:///system`.

`<poolname>` is the name of the storage pool in which to create the volume; typically you would use **default**. (To determine the storage pools available on the target system, use `virsh --connect qemu://<hostname>/system pool-list`.)

`<volume_name>` is the name of your new volume.

49G is the virtual size of the volume; always use 49G for a Conferencing Node.

For example:

```
virsh --connect qemu://host1.example.com/system vol-create-as default pexip-conf-01 49G --format raw
```

This example creates a volume named **pexip-conf-01** of size 49 GB and format raw in the storage pool named **default**.

2. Upload the converted disk image to the newly created volume:

```
virsh --connect qemu://<hostname>/system vol-upload <volume_name> pexip-disk01.raw --pool <poolname>
```

For example:

```
virsh --connect qemu://host1.example.com/system vol-upload pexip-conf-01 pexip-disk01.raw --pool default
```

This example uploads the **pexip-disk01.raw** image to the newly created volume, **pexip-conf-01**, in the storage pool named **default**.

Create the virtual machine

After the disk image has been uploaded, you can create the virtual machine to use it.

Note that we use the libvirt command line tools to perform the import as they provide greater control than Virtual Machine Manager.

1. Identify the filesystem path of the newly uploaded disk image:

```
virsh --connect qemu://<hostname>/system vol-path <volume_name> --pool <poolname>
```

For example:

```
virsh --connect qemu://host1.example.com/system vol-path pexip-conf-01 --pool default
```

This prints out the absolute path to the disk image file, for example:

```
/var/lib/libvirt/images/pexip-conf-01
```

This path is used in the **disk path** parameter in the next step.

2. Use the **virt-install** command line tool to create the virtual machine:

```
virt-install \
  --import \
  --hvm \
  --name=<vm_name> \
  --arch=x86_64 \
  --vcpus=4 \
  --ram=4096 \
  --cpu host \
  --os-type=linux \
  --connect=qemu://<hostname>/system \
  --virt-type kvm \
  --disk path=<image_file_path>,bus=virtio,format=raw,cache=none,io=native \
  --network bridge=br0,model=virtio \
  --memballoon virtio \
  --graphics vnc,listen=0.0.0.0,password=<password>
```

This creates a new VM (KVM domain) from the converted disk image.

The command options are described below (items in **bold** may be changed as necessary):

Option	Description
--import	Build guest domain around pre-installed disk image; do not attempt to install a new OS.
--hvm	Create a fully virtualized (i.e. not paravirtualized) VM.
--name=<vm_name>	Name of the new VM, where <vm_name> is, for example, pexip-conf01-vm .
--arch=x86_64	CPU architecture of new VM (must be x84_64).
--vcpus=4	Number of CPUs allocated to new VM. By default, this is 4 for the Conferencing Node.
--ram=4096	Memory allocated to new VM (in megabytes).
--cpu host	Expose all host CPU capabilities to new VM (CPUID).
--os-type=linux	The guest OS is Linux.

Option	Description
-- connect= qemu://<hostname>/system	Connect to KVM on the target system, where <hostname> is the hostname of your KVM server.
--virt-type kvm	Use KVM to host the new VM.
--disk path=<image_file_path>, bus=virtio ,format=raw,cache=none,io=native	<ul style="list-style-type: none"> Define the location of the disk image file, where <image_file_path> is as determined in the previous step, for example <code>/var/lib/libvirt/images/pexip-conf-01</code>. Expose it to the guest on the virtio paravirtualized bus (as opposed to IDE/SCSI). Define the image file as being in raw format. Instruct the host system not to cache the disk contents in memory. Use the native IO backend to access the disk device.
--network bridge=br0,model=virtio	<ul style="list-style-type: none"> Create a network interface connected to the br0 bridge interface on the host. Expose it to the guest as a virtio paravirtualized NIC.
--memballoon virtio	Expose the virtio memory balloon to the guest.
--graphics vnc,listen= 0.0.0.0 , password=<password>	Expose the graphical console over VNC, listening on 0.0.0.0 (i.e. all addresses on the target system) and with an access password of <password>.

You may receive a warning "Unable to connect to graphical console: virt-viewer not installed"; if so, this message can be safely ignored.

After the VM has been created, it may be managed using the Virtual Machine Manager desktop interface (virt-manager application) or via the command line interface (virsh).

The new node should start automatically. If it does not you can use the Virtual Machine Manager to start the node, or the CLI command:

```
virsh --connect qemu://<hostname>/system start <vm_name>
```

Note that you can list existing VMs by using `virsh --connect qemu://<hostname>/system list`


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".)

Enabling automatic startup

After deploying a new Conferencing Node in KVM, you should enable automatic startup of that virtual machine (VM). In KVM, automatic startup is disabled by default for every new VM. This means that if the host server is powered down for any reason, when it restarts the VM will not restart and must be started manually.

You can only enable automatic startup after the Conferencing Node has been deployed.

To enable automatic startup using Virtual Machine Manager:

1. Connect to the Virtual Machine Manager (virt-manager) that is managing the node's VM.
2. Select the node's VM and then, from the toolbar, select the **Show the virtual machine console and details** icon . A new window for that VM is opened.
3. If necessary, select **View > Details** to display the VM information.
4. From the sidebar menu, select **Boot Options**.
5. Select the **Start virtual machine on host boot up** check box.
6. Select **Apply**.

Next steps after initial installation

Making a test call

When you have deployed a Conferencing Node and configured a Virtual Meeting Room and an alias, you can make a test call to check that your system is working.

An easy way to do this is by using the Infinity Connect web app to dial the alias of one of the Virtual Meeting Rooms you've already created, as follows:

1. Open a browser (we recommend Chrome) and type in the IP address (or FQDN, if you've set it up already) of one of the Conferencing Nodes.
2. In the **Person or conference to call** field, enter the alias of the VMR you want to use for testing.
3. Select **Connect** and check your microphone, camera and speakers are working as expected.
4. Select **Start**.
You will be connected to the VMR.
5. From another machine, join the conference in the same way.

The two participants should be able to see and hear each other, and share content.

See [About the Infinity Connect web app](#) for more information.

Further configuration

You are now ready to continue [configuring the Pexip Infinity platform](#) and [services](#) and deploying more [Conferencing Nodes](#).

Specifically, you should now do the following:

- [Assigning hostnames and FQDNs](#)
- [Enabling SNMP on Conferencing Nodes](#)

At some point you may also want to:

- [integrate the Pexip Infinity platform with your call control system](#)
- [configure the Pexip Distributed Gateway](#)
- [register devices directly to the Pexip Infinity platform](#)
- [customize the user experience](#)

Integrating with a call control system

To integrate Pexip Infinity with your call control system, you must configure a trunk or neighbor zone towards each of the Conferencing Nodes.

For further information about how to configure your specific call management system to work with Pexip Infinity, see the following documentation:

- [Pexip Infinity and Microsoft Skype for Business / Lync Deployment Guide](#)
- [Pexip Infinity and Cisco VCS Deployment Guide](#)
- [Pexip Infinity and Cisco Unified Communications Manager Deployment Guide](#)
- [Pexip Infinity and Polycom DMA Deployment Guide](#)

Configuring the Pexip Distributed Gateway

The Pexip Distributed Gateway service enables endpoints to make calls to other devices or systems that use different protocols and media formats.

For more information, see [About the Pexip Distributed Gateway service](#).

Registering devices directly to the Pexip Infinity platform

SIP and H.323 endpoints, and some Infinity Connect clients can register directly to Pexip Infinity Conferencing Nodes. This allows Pexip Infinity to route outbound calls to those registered devices without having to go via a SIP proxy or H.323 gatekeeper, or rely on DNS.

For more information, see [Registering devices to Pexip Infinity](#).

Customizing the user experience

You can easily apply your own corporate branding to the Pexip Infinity platform, and produce a personalized user experience for all of your Pexip Infinity services.

For more information, see [Customizing the user experience](#).

Pexip Infinity installation checklist

Use this checklist to identify the key tasks involved in preparing for and deploying the Pexip Infinity platform. Also, there is a configuration [datasheet](#) below to help you gather the key network and configuration information required.

Prior to installation

1. Download the appropriate Pexip Infinity Management Node installation file from www.pexip.com/software-download.
2. Ensure that you have appropriate host servers (see [Server design guidelines](#)).
3. Assign network IP addresses and host names for the Management Node and Conferencing Nodes.
4. Create DNS records for your Management Node administration.
5. Create DNS records to allow endpoints/clients to discover your Pexip Infinity Conferencing Nodes (see [DNS record examples](#)).
6. Generate or request certificates (Base64-encoded X.509 PEM format) for the Management Node and Conferencing Nodes (see guidelines at [Certificate creation and requirements](#)).

Hypervisor / host servers

1. Note the CPU model number and the number of cores per socket on the host server to be used with the Conferencing Nodes, as this determines the maximum number of vCPUs to assign for the Conferencing Nodes.
2. Prior to deploying the Management Node or a Conferencing Node, ensure that all host servers are synchronized to NTP servers.
3. Upload the OVA file (or ZIP for Hyper-V) of the Management Node and run the setup wizard from the hypervisor console.

Pexip Infinity Administrator interface

1. Configure basic Management Node settings after installation (licenses, VM managers, any additional DNS or NTP servers).
2. Add a system location.
3. Deploy Conferencing Nodes to the location (and use your hypervisor management tools for manual deployments).
4. Configure the SIP TLS FQDN on the Conferencing Nodes.
5. Verify your node's DNS records. (You can use the tool at <http://dns.pexip.com> to lookup and check SRV records for a domain.)
6. Replace the self-signed server certificates on the Management Node and Conferencing Nodes with your own certificates that have been signed by either an external CA or a trusted internal CA (see [Managing TLS certificates](#)).
7. Upload any required chain of intermediate CA certificates to the Management Node.
You can use a tool such as <https://www.sslshopper.com/ssl-checker.html> to verify certificates and the chain of trust (specify port 5061 i.e. use the format <domain>:5061 for the server hostname to ensure that SIP TLS connections are checked).
8. Configure your VMRs and aliases.
9. Configure the Pexip Distributed Gateway (via Call Routing Rules), if required.

Hypervisor maintenance

1. Enable automatic startup on every VM.
2. Backup your Management Node VM, and optionally, your Conferencing Node VMs.

Pexip Infinity configuration datasheet

Use this datasheet to help you gather the key network and configuration information required for your deployment.

Management Node (installation wizard)

Management Node IP address:

Network mask:

Gateway IP address:

Management Node hostname:

Management Node domain:

DNS server 1:

DNS server 2:

NTP server 1:

NTP server 2:

Management Node (configuration)

VM name:

System location 1 name:

License entitlement key:

VM manager IP/hostname:
(VMware automatic deployments only)

Conferencing Nodes

CPU cores per socket on host server:
(to determine the size of each node)

Conferencing Node 1 name / VM name:

Conferencing Node 1 IP address:

Conferencing Node 1 hostname and domain:

Conferencing Node 2 name / VM name:

Conferencing Node 2 IP address:

Conferencing Node 2 hostname and domain:

For complete information on how to configure your Pexip Infinity solution, see the Pexip Infinity technical documentation website at docs.pexip.com.