Pexip Reverse Proxy and TURN Server Deployment Guide

About this guide

This guide provides design and configuration guidelines for using a reverse proxy and a TURN server with Pexip Infinity.

Introduction

A reverse proxy and a TURN server are typically used in Pexip Infinity deployments where some clients cannot communicate directly with Pexip Conferencing Nodes, for example in on-premises deployments where the Pexip platform is located on an internal, enterprise LAN network while the clients are located in public networks on the Internet. In these cases a reverse proxy can be used to proxy the call signaling traffic between the externally-located client and the internal Conferencing Node. In addition, as the reverse proxy does not handle media, a TURN server acts as a media relay between the external client and the internal nodes.

Deployment recommendations

Since version 16, we recommend that you deploy Proxying Edge Nodes instead of a reverse proxy and TURN server if you want to allow externally-located clients to communicate with internally-located Conferencing Nodes. 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.

Note that you may still want to deploy a reverse proxy in front of your Proxying Edge Nodes if, for example, you want to:

- host customized Infinity Connect web app content
- use it as a load balancer for Pexip's VMR Scheduling for Exchange service, to proxy requests from Outlook clients to Conferencing Nodes.

The following diagram shows how the reverse proxy could be used in conjunction with Infinity Connect clients and Proxying Edge Nodes:

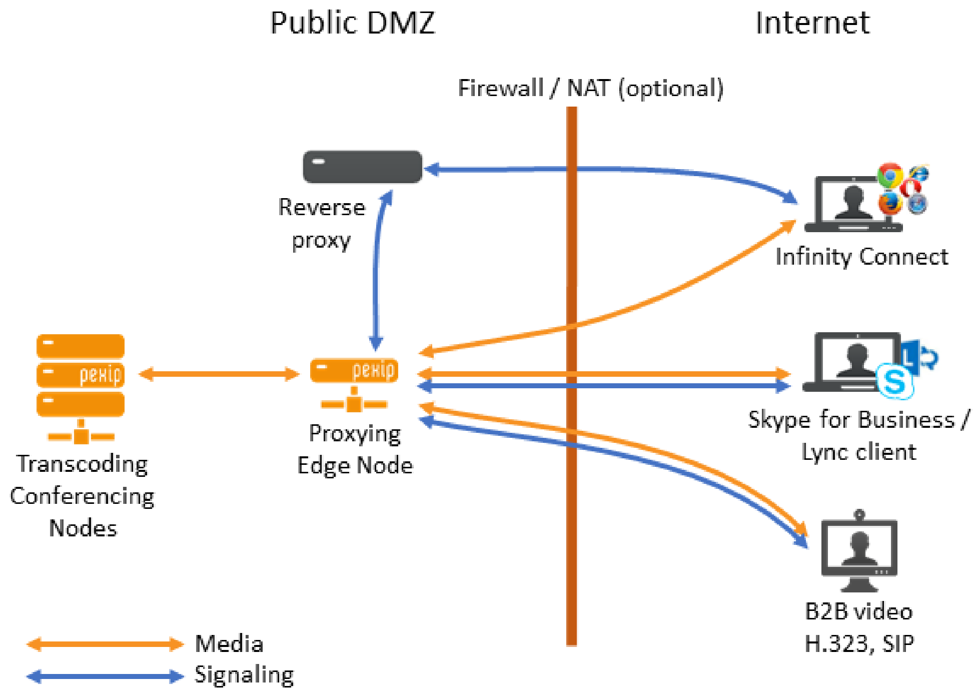

Deployments that do not use Proxying Edge Nodes

If you do not want to deploy Proxying Edge Nodes and thus want to route all signaling and media from external clients via a reverse proxy and a TURN server to your internal/on-premises nodes, then you should follow the rest of this Reverse Proxy and TURN Server guide and configure your on-premises nodes as Transcoding Conferencing Nodes.

Supported clients when using a reverse proxy/TURN

Pexip's Infinity Connect WebRTC clients (web app for Chrome, Firefox, Opera and Safari (version 11 onwards), the mobile clients for iOS and Android, and the desktop client) use ICE (Interactive Connectivity Establishment) to negotiate optimal media paths with Conferencing Nodes. Microsoft Skype for Business and Lync clients use a similar ICE mechanism, which means that Pexip can use TURN for all of these client types.

However, Infinity Connect clients on Internet Explorer and Safari (versions 6-10) browsers use the RTMP protocol, rather than WebRTC. While RTMP clients can connect to Conferencing Nodes via the reverse proxy, they cannot establish audio/video paths to Pexip Infinity via a TURN server. To establish audio/video media connectivity, RTMP clients need a direct TCP connection to a Conferencing Node.

Note that Microsoft Edge browsers (which are WebRTC-compatible) cannot currently use STUN and thus cannot send media to Pexip Infinity via a TURN server.

(An additional benefit of deploying Proxying Edge Nodes is that they provide connectivity for all browsers — WebRTC and RTMP-based.)

What's new in this release?

Version 5.0.5 of the OVA template was released in December 2017 and includes the following changes:

- Support for collecting usage statistics from the next-generation web app.

Version 5.0.3 of the OVA template was released in August 2017 and includes the following changes:

- The OVA template can now be deployed to VMware ESXi using vSphere web client 6.5.0.
- Content Security Policy no longer logs that it has blocked access to *.microsoft.com when used with Pexip's VMR Scheduling for Exchange add-in for Outlook clients.
- By default the TURN server now supports 15,000 sessions (previously 2,000).

Version 5 of the OVA template was released in April 2017 and includes the following changes and enhancements:

- The installation wizard contains a new option to enable Content Security Policy. This provides enhanced security against cross site scripting attacks.
- The fail2ban service no longer attempts to send mail (as there is no sendmail user).
- NTP synchronization now works correctly.
- TLS connections no longer use Triple DES.

Version 4 of the OVA template was released in November 2016 and offers improved scaling, OS security updates, and faster installation. It also includes the **fail2ban** service which provides protection against brute force attacks on PIN-protected conferences. Note that fail2ban is disabled by default. For more information, and instructions on how to enable fail2ban, see Enabling fail2ban.

Deployment options

Any type of HTTPS reverse proxy/load balancer or TURN server may be used with Pexip Infinity. However, this guide describes how to deploy these applications using the Reverse Proxy and TURN Server VMware appliance provided by Pexip.

This virtual VMware appliance is available as an OVA template which can be deployed on VMware ESXi 5 or later. The virtual appliance contains both reverse proxy and TURN applications.

The server hosting the reverse proxy requires a minimum of 2 vCPU, 2 GB RAM and 50 GB storage.

Depending on the network topology, the reverse proxy can be deployed with one or two network interfaces in various configurations:

- Single NIC, public address – see Example deployment: single NIC on public address
- Dual NIC, private and public addresses – see Appendix 2: Alternative dual NIC reverse proxy/TURN server deployment

In deployments with more than one Conferencing Node, the reverse proxy can load-balance HTTPS traffic between all Conferencing Nodes using a round-robin algorithm.

We recommend that the reverse proxy is configured with at least 3 Conferencing Nodes for resiliency as backend/upstream servers.

Prerequisites and requirements

Ensure that the following prerequisites are in place:

- The Pexip Infinity deployment (i.e. a Management Node and at least one Conferencing Node) must be configured and in a working state.
- Appropriate DNS SRV records must have been created in accordance with Using the reverse proxy with the Infinity Connect desktop client and Infinity Connect mobile client.

The reverse proxy and TURN applications require Pexip Infinity version 9 or later.

Security considerations

Infinity Connect clients (for conferencing services) and Outlook clients (for scheduling services) can only use encrypted HTTPS when communicating with Conferencing Nodes. The reverse proxy must therefore provide HTTPS interfaces through which the Infinity Connect and Outlook clients can communicate.

When configured correctly, the reverse proxy will allow HTTPS traffic to flow between the Infinity Connect and Outlook clients and the Conferencing Nodes only. Externally located clients will not be able to access other internal resources through the reverse proxy.

For conferencing services, we recommend that you install your own SSL/TLS certificates on the reverse proxy and TURN server for maximum security. If you are using VMR Scheduling for Exchange you must install your own valid certificates. For more information, see Replacing the default SSL certificate.

Version 4 of the reverse proxy introduced the **fail2ban** service which provides protection against brute force attacks on PIN-protected conferences. Note that fail2ban is disabled by default. For more information, and instructions on how to enable fail2ban, see Enabling fail2ban.

Design principles and guidelines

This section describes the design principles, guidelines and network requirements for a reverse proxy and for a TURN server when deployed with Pexip Infinity.

Reverse proxy application

In Pexip Infinity deployments, all Pexip Infinity Connect clients use HTTPS for the call signaling connections towards Conferencing Nodes.

The **reverse proxy application** is responsible for proxying HTTP/HTTPS requests from Infinity Connect WebRTC and desktop clients to one or more Conferencing Nodes. If you are using VMR Scheduling for Exchange, the reverse proxy application can also be used to proxy requests from Outlook clients, for the purposes of load balancing.

To proxy these requests, the reverse proxy application must be able to communicate with these externally-located clients as well as the Conferencing Nodes. This means that the reverse proxy must be able to reach any internally-located Conferencing Nodes either via a routed network or through NAT/port forwarding. The reverse proxy only needs to communicate with the Conferencing Nodes via HTTPS over TCP port 443 (when NAT/port forwarding is used to reach the Conferencing Nodes, the NATted port does not have to be 443, but the NAT/port forward must redirect to TCP/443 on the Conferencing Node).

TURN server application

As the reverse proxy does not handle media, the **TURN server application** enables external clients to exchange RTP/RTCP media (i.e. ensure audio/video connectivity) with the Conferencing Nodes.

A TURN server is a media relay/proxy that allows peers to exchange UDP or TCP media traffic whenever one or both parties are behind NAT. When Conferencing Nodes are deployed behind NAT, these nodes will instruct the WebRTC client to send its media packets to the TURN server, which will forward (relay) the packets to the Conferencing Nodes. Since this TURN server is normally located outside of the enterprise firewall, the Conferencing Node will constantly send media packets to this TURN server to "punch holes" in the firewall, allowing this TURN server to relay media packets back to the Conferencing Node, as the firewall will classify this as return traffic.

When using a TURN server with a Conferencing Node:

- Conferencing Nodes only use TURN over UDP (not TCP). However, Conferencing Nodes will perform ICE TCP negotiation.
- Conferencing Nodes always communicate with a TURN server over a single UDP port (default UDP/3478). UDP media is multiplexed from the Conferencing Node to that single port on the TURN server. The TURN server will reply back to the same port pair on the Conferencing Node. The TURN server never initiates a connection towards a Conferencing Node.

Another key responsibility of the TURN server is to act as a STUN server for the Conferencing Nodes – when a Conferencing Node is deployed behind a NAT (from the perspective of clients located on the Internet), the Conferencing Node uses STUN towards the TURN server to discover its public NAT address. The Conferencing Node sends a STUN request to the TURN server, which responds back to the Conferencing Node and tells it from which IP address it received the STUN request. Using this method, the Conferencing Node can discover its public NAT address, which is important in order for ICE to work between the Conferencing Node and clients using ICE (such as Skype for Business / Lync clients and Infinity Connect WebRTC clients). In relation to TURN and ICE, this public NAT address is also known as the server reflexive address or simply reflexive address, and will be referred to as such in this guide.

Example communication scenario

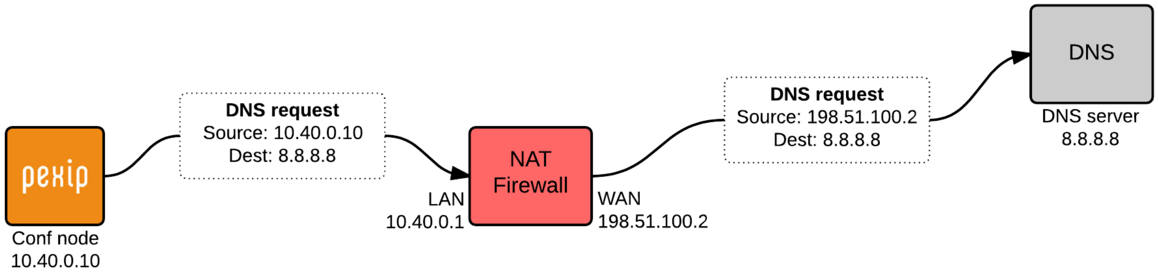

Using the above diagram as an example, the Conferencing Node has an IP address of 10.40.0.10 – this is a private/internal IP address which is not routable across public networks. When this Conferencing Node communicates with a host located on a public network (Internet), for instance a DNS server, traffic from this Conferencing Node passes through a NAT device (firewall/router), which will translate the source IP address for this traffic (10.40.0.10) to a public NAT address, in this case 198.51.100.2, before passing the traffic on to its destination. This means that when the DNS server receives the DNS request, the request will appear as coming from 198.51.100.2, which means that 198.51.100.2 is the **reflexive address** of the Conferencing Node.

For certain Skype for Business / Lync call scenarios to work correctly (notably RDP content sharing with external Skype for Business / Lync clients), it is essential that a Conferencing Node informs the remote Skype for Business / Lync client of this reflexive address. The Skype for Business / Lync client will in turn inform its Skype for Business / Lync Edge Server of this reflexive address so that the Edge Server will relay media packets from the Conferencing Node to the Skype for Business / Lync client.

In some deployment scenarios where the TURN server is not located outside of the enterprise firewall — and thus sees traffic from the Conferencing Nodes as coming from its private address e.g. 10.40.0.10 — a Conferencing Node will not be able to discover its reflexive address (its public NAT address, e.g. 198.51.100.2) by sending its STUN requests to the TURN server. In this case you may need to configure Pexip Infinity with the address of a separate STUN server, such as **stun.l.google.com**, so that each Conferencing Node can discover its reflexive address.

See When is a reverse proxy, TURN server or STUN server required? for summary guidelines of when each type of device is required in your deployment.

When is a reverse proxy, TURN server or STUN server required?

Since version 16, we recommend that you deploy Proxying Edge Nodes instead of a reverse proxy and TURN server if you want to allow externally-located clients to communicate with internally-located Conferencing Nodes.

However, if you do not want to deploy Proxying Edge Nodes, and all of your Conferencing Nodes are privately addressed, you will need to use a reverse proxy and a TURN server to allow external endpoints such as Infinity Connect and Skype for Business / Lync clients to access your Pexip Infinity services. A TURN server can also act as a STUN server, however, in some Pexip Infinity deployment scenarios where the TURN server is deployed inside your enterprise firewall, you may need to configure a separate, external STUN server.

When connecting to a privately-addressed Conferencing Node, Infinity Connect WebRTC clients that are behind a NAT may also use a STUN server to find out their public NAT address.

The following table shows when a reverse proxy, TURN server or STUN server needs to be deployed (if you are not using Proxying Edge Nodes). When used, they must be publicly accessible, and routable from your on-premises Conferencing Nodes.

External endpoint / client	Conferencing Node addresses	Reverse proxy	TURN server	STUN server (for Conferencing Nodes)	STUN server (for WebRTC clients behind NAT)
Infinity Connect WebRTC clients	Private (on-premises)	✓	✓	✓ (if the TURN server is inside the firewall)	✓
Skype for Business / Lync clients*	Private (on-premises)	-	✓ (only required if internal Conferencing Node cannot route to the public-facing interface of the SfB/Lync Edge server)	✓ (if the TURN server is inside the firewall)	✓
Any endpoint / client	Publicly reachable — either directly or via static NAT	-	-	-	-

* Also requires a Skype for Business / Lync Edge Server when Conferencing Nodes are privately addressed

Note that you may still want to deploy a reverse proxy in front of your Proxying Edge Nodes if, for example, you want to:

- host customized Infinity Connect web app content
- use it as a load balancer for Pexip's VMR Scheduling for Exchange service, to proxy requests from Outlook clients to Conferencing Nodes.

Example deployment: single NIC on public address

The diagram below depicts an example deployment where the reverse proxy and the TURN server application have been deployed with a single NIC on a public address.

The environment is split into two parts — an internal, private network segment and a DMZ network. The private network has two Pexip Infinity Conferencing Nodes, while the DMZ perimeter network contains the reverse proxy and TURN server.

Note that all IP addresses in this guide are examples only — actual IP addressing will be deployment specific.

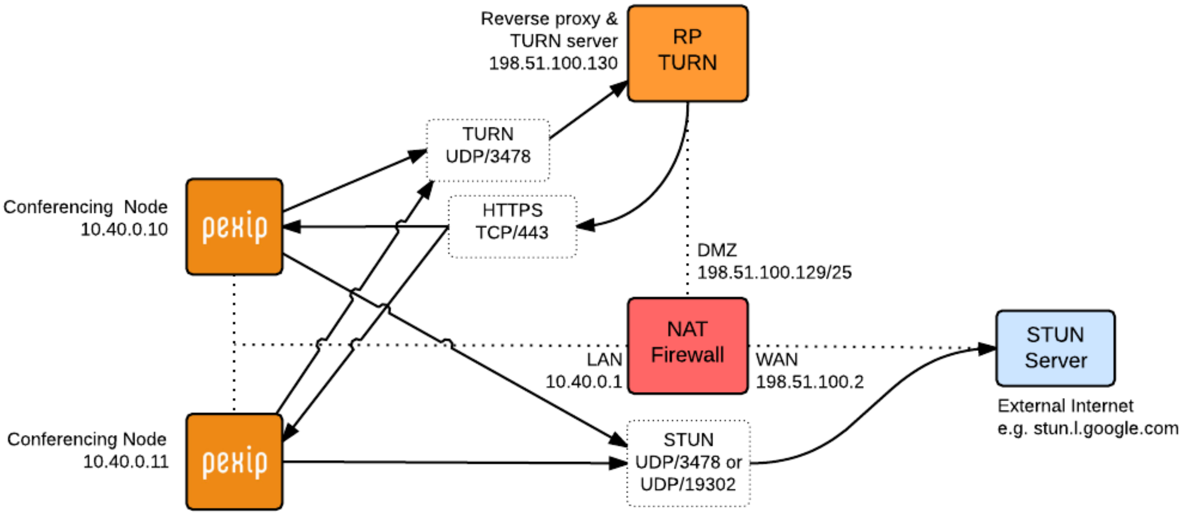

Example deployment used in this guide: single NIC on public address

This example forms the basis of this guide.

For this example deployment:

- Two Conferencing Nodes have been deployed in the LAN segment with IP addresses 10.40.0.10 and 10.40.0.11.
- The firewall in this scenario has three network interfaces:
  - LAN: 10.40.0.1/24
  - DMZ: 198.51.100.129/25
  - WAN: 198.51.100.2
- The DMZ network (198.51.100.129/25) can route network traffic to the LAN network (no NAT between LAN and DMZ).
- The reverse proxy and TURN server have been deployed in the DMZ subnet with IP address 198.51.100.130.
- As there is no NAT between the Conferencing Nodes in the LAN and the TURN server in the DMZ, the Conferencing Nodes have been configured to send their STUN requests to a STUN server in the public internet.
- The firewall has been configured to:
  - allow the reverse proxy to initiate HTTPS connections towards the Conferencing Node IP addresses
  - allow the Conferencing Nodes to send TURN packets to the TURN server on UDP port 3478
  - allow the Conferencing Nodes to send STUN packets to the STUN server, typically on UDP port 3478, although stun.l.google.com uses port 19302.

Deploying the reverse proxy and TURN server using an OVA template

Pexip provides a preconfigured Reverse Proxy and TURN Server appliance via an OVA template suitable for deployment on VMware ESXi. This OVA template is provided "as-is" and provides a reference installation which will be suitable for typical Pexip deployments where:

- Conferencing Nodes are deployed in internal, private networks.
- The reverse proxy and TURN server is deployed in a DMZ environment using one or two network interfaces.

Deployment steps

These steps involve:

- Downloading the OVA template
- Deploying the OVA template
- Setting the password for SSH/console access
- Running the installation wizard

Downloading the OVA template

Download the latest version 5 of the Pexip RP/TURN OVA template from https://dl.pexip.com/rpturn/index.html (select the v5 directory) to the PC running the vSphere client. Either the vSphere desktop client or web client can be used.

We recommend that you verify the OVA file integrity after downloading the OVA file by calculating the MD5 sum of the downloaded file (for instance using WinMD5 Free from www.winmd5.com) and comparing that with the respective MD5 sum found in file **readme.txt** (located in the same download location as the OVA images).

Deploying the OVA template

To deploy the OVA template:

- **vSphere Web Client**: go to **Hosts And Clusters**, right-click on a host server and choose **Deploy OVF Template...**.
- **vSphere Client for Windows**: go to **Hosts And Clusters**, click **File** and **Deploy OVF Template**.

We recommend 2 cores and 2 GB RAM for the host server VM.

During the OVA deployment, we recommend that you use the default options. Also make sure to assign the correct VMware network/port group for the network interface of the virtual machine.

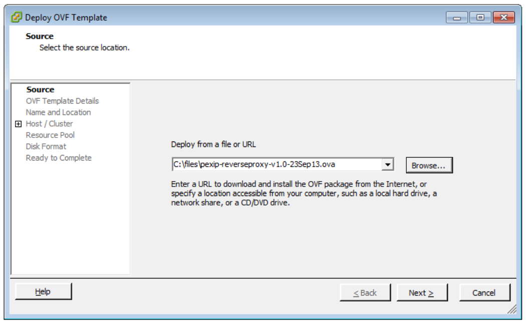

Deploy OVF Template

Source
Select the source location.

Deploy from a file or URL

C:\files\pexip-reverseproxy-v1.0-23Sep13.ova

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

vSphere desktop client

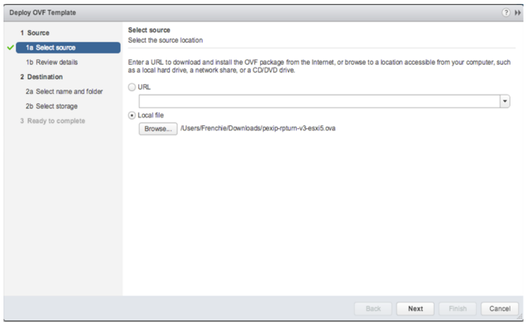

Select source

Select the source location

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

○ URL

● Local file

Browse... /Users/Frenchie/Downloads/pexip-rpturn-v3-esxi5.ova

vSphere web client

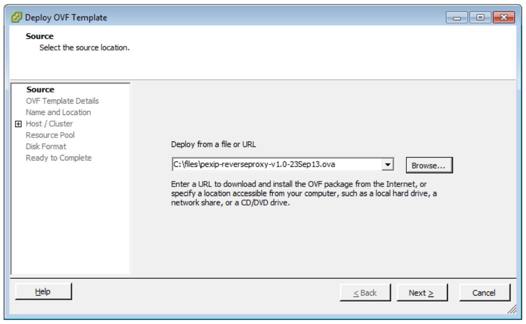

Source

Select the source location.

Deploy from a file or URL

C:\files\pexip-reverseproxy-v1.0-23Sep13.ova

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

vSphere desktop client

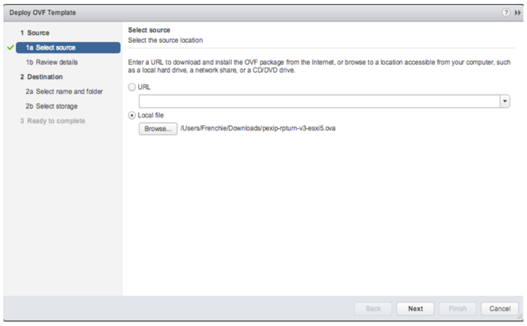

Select source

Select the source location

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

○ URL

● Local file

Browse... /Users/Frenchie/Downloads/pexip-rpturn-v3-esxi5.ova

vSphere web client

After the OVA template has been deployed, power on the newly-created virtual machine.

Setting the password for SSH/console access

After the virtual machine has powered on, open the VMware console for the Reverse Proxy and TURN Server virtual machine, right-click on the virtual machine in the vSphere client and choose **Open Console**.

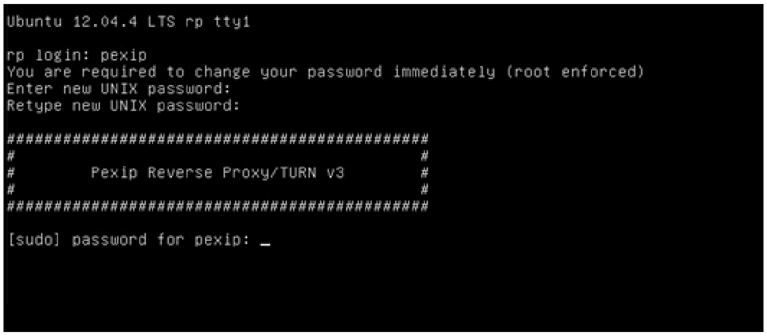

```
Ubuntu 12.04.4 LTS rp tty1

rp login: pexip
You are required to change your password immediately (root enforced)
Enter new UNIX password:
Retype new UNIX password:

#############################################
#                                           #
#        Pexip Reverse Proxy/TURN v3        #
#                                           #
#############################################

[sudo] password for pexip: _
```

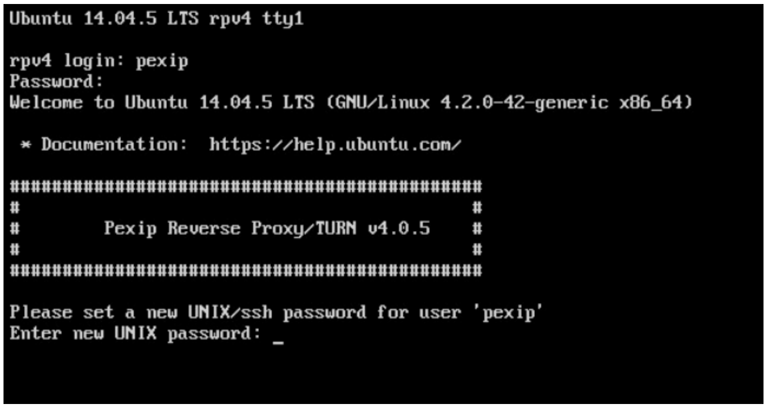

```
Ubuntu 14.04.5 LTS rpv4 tty1

rpv4 login: pexip
Password:
Welcome to Ubuntu 14.04.5 LTS (GNU/Linux 4.2.0-42-generic x86_64)

 * Documentation:  https://help.ubuntu.com/

#############################################
#                                           #
#        Pexip Reverse Proxy/TURN v4.0.5    #
#                                           #
#############################################

Please set a new UNIX/ssh password for user 'pexip'
Enter new UNIX password: _
```

Initial login prompt

Before you can start the install wizard, you must change the password. To do this:

1. Log in as user **pexip** with password **PEXIP** (these are case sensitive).
2. You are prompted to set a new account password. To do this you must enter the new password twice.
3. After setting the new password, the **Pexip Reverse Proxy/TURN** banner is shown, and you are prompted to enter the new password again.
   After this password has been entered, the install wizard will start.

Running the installation wizard

The installation wizard is divided into several steps, which are explained below. Some steps require a single line of input, while others allow multiple lines of input to be entered, one line at a time. Some steps also allow multiple entries on the same line.

The configuration example described here is based on the Example deployment: single NIC on public address, with the following additional assumptions:

- The reverse proxy and TURN server interface (198.51.100.130) resides in the same subnet as its default gateway (198.51.100.129).
- Router 10.40.0.1 is the next hop when accessing all internal hosts. The internal networks are defined by CIDR 10.0.0.0/8 (10.0.0.0-10.255.255.255).
- Hosts residing in the internal network 10.0.50.0/24 will access the reverse proxy and the TURN server over SSH.

Note that all IP addresses in this guide are examples only — actual IP addressing will be deployment specific.

The following table shows, for each step, the prompt text that is shown, the example text you should input, and an explanation of the step.

Step	Prompt text	Example input	Explanation for this example deployment
1	IP address	198.51.100.130 [ENTER]	198.51.100.130 is the IP address of the reverse proxy/TURN server interface. If the intention is to have dual NICs, specify the address of the internally-facing NIC (see Configuring a second NIC).
2	Subnet mask	255.255.255.0 [ENTER]	The reverse proxy/TURN server has a network mask of 255.255.255.0.
3	Default Gateway	198.51.100.129 [ENTER]	The reverse proxy/TURN server has a default gateway of 198.51.100.129.

Step	Prompt text	Example input	Explanation for this example deployment
4	Network address and subnet mask of network/host allowed to access this host over SSH	10.0.50.0 [SPACE] 255.255.255.0 [ENTER]	10.0.50.0 is the network address and 255.255.255.0 is the subnet mask of the enterprise's management network that is allowed to access this host over SSH.
5	Hostname	proxy [ENTER]	The reverse proxy has a hostname of **proxy**. The hostname and domain (configured in the next step) must match the actual DNS name by which the reverse proxy will be addressed (**proxy.example.com**).
6	Domain suffix	example.com [ENTER]	The reverse proxy has a domain name/suffix of **example.com**. This means, since the host name in the previous step was configured as **proxy**, that the full FQDN of this reverse proxy is **proxy.example.com**. If a custom SSL certificate is created for this reverse proxy, this FQDN needs to match the Subject Name and Subject Alternative Name of the SSL certificate. For more information, see Replacing the default SSL certificate.
7	DNS server(s), space separated	8.8.8.8 [SPACE] 8.8.4.4 [ENTER]	The reverse proxy is configured to use DNS servers **8.8.8.8** and **8.8.4.4**. DNS is in this case mainly used to resolve the hostnames of NTP servers.
8	NTP server(s), space separated	0.no.pool.ntp.org [SPACE] 1.no.pool.ntp.org [SPACE] 2.no.pool.ntp.org [ENTER]	The reverse proxy is configured to use three different pools of NTP servers, to ensure proper NTP time synchronization. We recommend that at least three NTP servers are used.
9	Enter the IP address of each conference node separately, followed by [ENTER]. Press [ENTER] on an empty line to finish entering conference nodes: Enter IP of next conference node or press [ENTER] if finished	10.40.0.10 [ENTER] 10.40.0.11 [ENTER] [ENTER]	Here, the IP addresses of both Pexip Conferencing Nodes are input, one at a time – **10.40.0.10** is the first Conferencing Node and **10.40.0.11** is the second Conferencing Node. We recommend that these nodes are configured as Transcoding Conferencing Nodes. (If you are deploying a TURN server only, input '169.254.0.1' followed by <ENTER>. 169.254.0.1 is a link-local IP address which is not used in production networks.)
10	Choose a space-separated username and password to enable the built in TURN server Enter 'disable' to disable the built in TURN Server Example: someusername [SPACE] somepassword [ENTER]	pexip [SPACE] admin123 [ENTER]	When deploying the TURN server application, you must enter some credentials, for example, entering: pexip [SPACE] admin123 [ENTER] would set the username to 'pexip' with a password of 'admin123'. These are the credentials you would use when configuring Pexip Infinity with the access details for this TURN server. (If you are deploying the reverse proxy only, input 'disable'.)
11	Content Security Policy	yes [ENTER]	Content-Security-Policy is an HTTP header that provides enhanced security against cross site scripting attacks. Enter 'yes' to enable Content Security Policy. This is recommended if you are **not** using optional advanced features such as plug-ins for Infinity Connect, externally-hosted branding, or externally-hosted pexrtc.js in your Pexip deployment. Enter 'no' for best compatibility with these optional advanced features.

When all of the installation wizard steps have been completed, the appliance will automatically reboot.

After the appliance has started up again it will be ready for use — Infinity Connect mobile client and Infinity Connect users will now be able to access Virtual Meeting Rooms from outside your network.

Replacing the default SSL certificate

For conferencing services, we recommend that you install your own SSL/TLS certificates on the reverse proxy and TURN server for maximum security. If you are using VMR Scheduling for Exchange you must install your own valid certificates.

To replace the built-in X.509 SSL certificate on the reverse proxy with a custom-created certificate:

1. Create a text file called **pexip.pem** which contains the following items in this specific order:
   - server certificate
   - server private key (which must be unencrypted)
   - one or more intermediate CA certificates (a server certificate will normally, but not always, have one or more intermediate CA certificates)

   Note that the contents MUST be in this specific order for the certificate to work properly.

   The first section with the server certificate should contain a single entry starting and ending with the following:

   ```
   ---- BEGIN CERTIFICATE ---- / ---- END CERTIFICATE ----
   ```

   The second section with the server private key should contain a single entry starting and ending with the following (although it may instead show '`BEGIN RSA PRIVATE KEY`'):

   ```
   ---- BEGIN PRIVATE KEY ---- / ---- END PRIVATE KEY ----
   ```

   Finally, there will normally be one or more intermediate CA certificates, where each intermediate has a section starting and ending with:

   ```
   ---- BEGIN CERTIFICATE ---- / ---- END CERTIFICATE ----
   ```

2. Using the SCP file transfer protocol, upload the **pexip.pem** file to the **/tmp** folder of the Reverse Proxy and TURN Server. This can be done using for instance WinSCP (www.winscp.net) or the 'scp' command-line utility for Linux/Mac OS X, using a command such as:

   ```
   scp pexip.pem pexip@198.51.100.130:/tmp
   ```

3. After the **pexip.pem** file has been transferred into the **/tmp** folder of the reverse proxy, connect over SSH to the reverse proxy, log in as user **pexip** and run the following commands, one at a time:

   ```
   sudo cp /etc/nginx/ssl/pexip.pem /etc/nginx/ssl/pexip.pem.backup
   ```

   ```
   sudo mv /tmp/pexip.pem /etc/nginx/ssl/pexip.pem
   ```

   ```
   sudo service nginx restart
   ```

   Note that `sudo service nginx restart` will restart the reverse proxy application and therefore interrupt the service briefly.

After these commands have been run, the reverse proxy should now be operational and using the new certificate.

If any problem occurs with the replaced certificate, the previous certificate can be restored using the following commands:

```
sudo cp /etc/nginx/ssl/pexip.pem.backup /etc/nginx/ssl/pexip.pem
```

```
sudo service nginx restart
```

Enabling fail2ban

Fail2ban is an intrusion prevention framework that can protect the reverse proxy from brute-force attacks on PIN-protected conferences. Fail2ban is disabled by default on the reverse proxy.

When enabled, fail2ban works by scanning the logs on the reverse proxy for repeated failed PIN entry attempts from the same IP address, and then blocks the source IP address that is responsible for that activity.

By default, it blocks access for 300 seconds if 10 PIN failures are logged from the same IP address within a 300 second window. The blocked IP address will be unable to connect to the reverse proxy for the duration of the ban. Note that each attempt to access a PIN-protected conference always creates one "failed PIN" log entry (even if the supplied PIN is correct), plus a second failure log entry if the supplied PIN is incorrect. Therefore in practical terms a user will be banned if they supply 5 incorrect PIN numbers (5 incorrect attempts x 2 log entries per attempt = 10) within the time window — see Limitations below for more information.

Benefits

It allows any source IP address a maximum of 5 incorrect PIN entries in a 5 minute (300s) window — and thus (on average) limits an attacker to about one PIN guess attempt per minute over the long term. If you have a six digit PIN you'd need an average of 500,000 PIN guess attempts to crack the average PIN (assuming the PIN was random) — which would take a single attacker, attacking from a single source IP, approximately 500000/60/24 = 347 days to crack (using a naive brute force attack). A four digit PIN would take just 5000/60/24 = 3.47 days to crack — so longer PINs really do provide a significant benefit.

Limitations

Due to the nature of the underlying protocol used to access PIN-protected conferences, every attempt to access a PIN-protected conference results in one "failed PIN" log entry being created — and then a further "failed PIN" log entry is recorded if an incorrect PIN is submitted. Thus, for example, two incorrectly submitted PINs for the same conference will result in four "failed PIN" log entries being created. An attempt to access a conference where an incorrect PIN is submitted before the correct PIN is supplied will result in three "failed PIN" log entries.

Therefore, to provide a balance between blocking intruders but allowing for normal use and genuine user errors, the default settings within the fail2ban service are configured to ban the source IP address if it detects 10 "failed PIN" entries in the log file (i.e. the intruder has submitted 5 incorrect PINs). Note that this also means, for example, that a source address would be blocked if it attempts to join 10 PIN-protected conferences within a 5 minute period, even if the correct PIN is always supplied. You can modify the default settings for the number of failures, ban duration and so on if required.

The fail2ban service won't deter a determined or well-resourced attacker who is prepared to conduct a brute force attack from multiple source IP addresses or ride out the ban duration repeatedly over a period of days, but it does make break-in harder.

Users behind NATs

All users that are behind the same NAT are seen as having the same source address. Therefore, if you have multiple users behind the same NAT who are accessing the reverse proxy, we recommend caution in enabling fail2ban as they could block themselves out even if they never enter a wrong PIN.

Alternative deployment scenario possibilities to cater for users behind a NAT include having a separate internal reverse proxy (with fail2ban enabled) for any internal users that are behind your own NAT and using split horizon DNS to ensure they are routed to the internal reverse proxy. Individual users who work from home or other remote location (from behind a remote NAT) will not normally be a problem as each user will typically be behind a different NAT (with a different IP address). However, a service-provider scenario where a publicly-deployed reverse proxy is providing conferencing facilities to a group of remote users who are all behind the same enterprise NAT will need to have fail2ban disabled.

Enabling, configuring and monitoring fail2ban

Fail2ban is disabled by default on the reverse proxy. This section describes how to enable fail2ban, modify the default configuration if required, and how to monitor the service.

Enabling fail2ban

To enable fail2ban on the reverse proxy, connect over SSH to the reverse proxy, log in as user **pexip** and run the following commands one at a time:

```
sudo update-rc.d fail2ban enable
```

```
sudo /etc/init.d/fail2ban restart
```

The response on the console to the `restart` command will be `* Restarting authentication failure monitor fail2ban` (note that this is not an error message).

If you want to disable fail2ban, use the commands:

```
sudo update-rc.d fail2ban disable
```

```
sudo /etc/init.d/fail2ban stop
```

Modifying the default fail2ban configuration

The default configuration of the fail2ban service is to block access for 300 seconds if it detects 10 "failed PIN" log entries from the same IP address within a 5 minute window.

You can modify each of these settings if required. To change the fail2ban configuration:

1. Run the following command to edit the fail2ban config file:
   ```
   sudo nano /etc/fail2ban/jail.local
   ```
2. You can modify the following ban-related settings:

Setting	Purpose	Default
bantime (DEFAULT section)	The number of seconds that a host address is banned.	300
findtime (DEFAULT section)	The time period in seconds that is monitored.	300
maxretry (you **must** modify the maxretry value specified in the **pexiprp** jail)	The number of failures that when reached will trigger the ban. A host is banned if it generates **maxretry** failures during the last **findtime** seconds.	10

3. After editing and saving the file, run the following command to restart the fail2ban service:
   ```
   sudo /etc/init.d/fail2ban restart
   ```

For more information on advanced configuration, see https://www.digitalocean.com/community/tutorials/how-to-protect-an-nginx-server-with-fail2ban-on-ubuntu-14-04.

Checking the status of the fail2ban service

You can check the status of the fail2ban service by running the following command:

```
sudo fail2ban-client status pexiprp
```

which gives the following status:

```
pexip@rp:/var/log$ sudo fail2ban-client status pexiprp
Status for the jail: pexiprp
|- filter
|  |- File list:    /var/log/nginx/pexapp.access.log
|  |- Currently failed:    0
|  `- Total failed:    0
`- action
   |- Currently banned:    0
   |  `- IP list:
   `- Total banned:    0
```

The `failed` and `banned` counts may be non-zero if the service has detected access failures.

If the service is not running you will see: `ERROR Unable to contact server. Is it running?`

Using the reverse proxy and TURN server with Infinity Connect and Skype for Business / Lync clients

Infinity Connect clients can connect directly to a Conferencing Node, but this will not provide a mechanism for balancing load between multiple Conferencing Nodes, or failing over in the event of a node failure. In addition, many customers may deploy Conferencing Nodes in a private network but would like to also provide access to external users using the Infinity Connect web app.

To resolve these issues, a reverse proxy in the DMZ can be used to forward the HTTPS traffic from the browser to the Conferencing Nodes, and a TURN server can be used to forward media from a private network to the public Internet.

It may be necessary to configure Pexip Infinity with the address of a STUN server, such as **stun.l.google.com**, so that each Conferencing Node can discover its reflexive address, which is essential for certain Skype for Business / Lync call scenarios to work correctly.

This section describes how to:

- connect to the reverse proxy from Infinity Connect clients and the Infinity Connect mobile client
- configure the Pexip Infinity platform with details of the TURN server
- configure the Pexip Infinity platform with details of a STUN server.

Using the reverse proxy and TURN server with the Infinity Connect web app

When the reverse proxy has been deployed, Infinity Connect users with WebRTC-compatible browsers can access conferences via **https://<reverse-proxy>/webapp/**, where **<reverse-proxy>** is the FQDN of the reverse proxy. This mechanism uses HTTPS for accessing the web pages and conference controls, and RTP/RTCP for the media streams (via a TURN server if necessary).

Note that Microsoft Edge browsers (which are WebRTC-compatible) cannot currently use STUN and thus cannot send media to Pexip Infinity via a TURN server.

(If the reverse proxy is not available, Infinity Connect web app users can connect via **https://<node>/**, where **<node>** is the IP address or URL of the Conferencing Node, providing the web app can reach the node's IP address directly.)

Using the reverse proxy with the Infinity Connect desktop client and Infinity Connect mobile client

The Infinity Connect desktop client and Infinity Connect mobile client work by sending HTTP GET and POST requests to a specific destination address to fetch information about a meeting (such as the participant list) and to send various commands (such as to mute or remove conference participants).

Clients discover the destination address for those HTTP requests through a custom DNS SRV lookup for **_pexapp._tcp.<domain>**. For instance, if the Infinity Connect desktop client or Infinity Connect mobile client has been configured with a meeting URI of **meet.alice@example.com**, it will perform a DNS SRV lookup for **_pexapp._tcp.example.com**.

Assume that the following **_pexapp._tcp.vc.example.com** DNS SRV records have been created:

```
_pexapp._tcp.vc.example.com. 86400 IN SRV 10 100 443 px01.vc.example.com.
_pexapp._tcp.vc.example.com. 86400 IN SRV 20 100 443 px02.vc.example.com.
```

These point to the DNS A-records **px01.vc.example.com**, port 443 (HTTPS), with a priority of 10 and a weight of 100, and **px02.vc.example.com**, port 443, with a relatively lower priority of 20 and a weight of 100.

- This tells the Infinity Connect desktop client to initially send its HTTP requests to host **px01.vc.example.com** (our primary node) on TCP port 443. The desktop client will also try to use host **px02.vc.example.com** (our fallback node) if it cannot contact px01.
- The Infinity Connect mobile client will send its HTTP requests either to **px01.vc.example.com** or to **px02.vc.example.com**, depending on the order of the returned SRV records. If it fails to contact the first host, it will not attempt to contact the second host address.

The connection logic in this example is explained in more detail below for each client.

Infinity Connect desktop client - next-generation version

ℹ If a next-generation Infinity Connect desktop client is registered to Pexip Infinity and the global Route via registrar setting is enabled, then the client will route all calls directly to the IP address of the Conferencing Node to which it is registered. In all other cases, the following example applies.

When a user attempts to access **meet.alice@vc.example.com**:

1. If the call is being placed via a preconfigured link that specifies a **host** domain, then the client will perform an SRV lookup on that domain, and attempt to contact one of the hosts returned in that lookup.

   For example, if the URL is **pexip://meet.alice@vc.example.com?host=localserver.example.com** then the client will perform an SRV lookup on **_pexapp._tcp.localserver.example.com**.

   If the SRV lookup fails, or the returned hosts in the lookup cannot be contacted, the client will also attempt to connect directly to that domain, i.e. to **http://localserver.example.com:443** (via DNS A-records for localserver.example.com).

   If that also fails, no further lookups will be performed, and the client will report that it could not join the host domain.
2. If the client still has not contacted a Conferencing Node, the client will attempt an SRV lookup on the domain portion of the address that was dialed, i.e. on **_pexapp._tcp.vc.example.com**.

   If the SRV lookup succeeds, it will return the records shown above, and the client will attempt to contact **px01.vc.example.com** (the record with the highest priority) on TCP port 443.

   If it cannot contact **px01.vc.example.com** it will next try to contact **px02.vc.example.com**.

   If it fails to contact either host, the client will also attempt to connect directly to the domain, i.e. to **http://vc.example.com:443** (via DNS A-records for vc.example.com).

   If that also fails, the desktop client will report that it has failed to contact a server.

Infinity Connect mobile client - next-generation version

In this example, when a user attempts to access **meet.alice@vc.example.com**, the client will attempt an SRV lookup on the domain portion of the address that was dialed, i.e. on **_pexapp._tcp.vc.example.com**.

If this SRV lookup succeeds, it will return the records shown above, and the client will attempt to contact the first host in the returned list on TCP port 443. Note that the addresses are returned in an arbitrary order and thus the first host may be either **px01.vc.example.com** or **px02.vc.example.com**.

If this SRV lookup fails, or it fails to contact the first host on the returned list, the client will attempt to connect to **http://vc.example.com:443** (via DNS A-records for vc.example.com).

Infinity Connect desktop client - legacy version

In this example, when a user attempts to access **meet.alice@vc.example.com**, the Infinity Connect desktop client will attempt an SRV lookup on **_pexapp._tcp.vc.example.com**:

- If the SRV lookup succeeds, it will return the records shown above, and the Infinity Connect desktop client will attempt to contact **px01.vc.example.com** (the record with the highest priority) on TCP port 443.

  If it cannot contact **px01.vc.example.com** it will next try to contact **px02.vc.example.com**.
- If it fails to contact either host, or the SRV lookup fails, and neither a **serverAddress**, **Connection server address** nor a **Registration server address** have been specified, the desktop client will report that it has failed to contact a server.
- If any of the **serverAddress**, **Connection server address** or a **Registration server address** have been specified, and are for a different domain to that of the dialed alias (vc.example.com in this case) the Infinity Connect desktop client will perform SRV lookups on those other domains, and attempt to contact the hosts returned in those lookups. For example, if the **Connection server address** is **localserver.example.com** then it will perform an SRV lookup on **_pexapp._tcp.localserver.example.com**.
- If each subsequent SRV lookup fails, or the returned hosts in those lookups cannot be contacted, the Infinity Connect desktop client will also attempt to connect directly to that domain, for example to **http://localserver.example.com:443** (via DNS A-records for localserver.example.com).

Infinity Connect mobile client - legacy version

In this example, when a user attempts to access meet.alice@vc.example.com, the Infinity Connect mobile client will attempt an SRV lookup on **_pexapp._tcp.vc.example.com**:

- If the SRV lookup succeeds, it will return the records shown above, and the Infinity Connect mobile client will attempt to contact the first host in the returned list on TCP port 443. Note that the addresses are returned in an arbitrary order and thus the first host may be either **px01.vc.example.com** or **px02.vc.example.com**.
- If the SRV lookup fails, or it fails to contact the first host on the returned list, the Infinity Connect mobile client will attempt to connect to **http://vc.example.com:443** (via DNS A-records for vc.example.com).

(Note that for the Android client, this example assumes that a **Connection server address** is not configured on the client. If a connection server address is specified, it would be used instead of the domain portion of the dialed conference address i.e. vc.example.com in this case.)

Note that the Infinity Connect mobile client will keep polling the reverse proxy periodically to update the participant list for a given virtual meeting room for as long as the application is active.

Configuring Pexip Infinity to use a TURN server

To relay media between the internal and external networks, a TURN server must be used. In addition to Pexip's TURN Server appliance, many other commercial TURN servers exist, including those on products such as a VCS Expressway, or those deployed using commercial or free software such as restund or rfc5766-turn-server.

The TURN server's details must be configured on the Pexip Infinity platform, and each location must nominate the TURN server that will be used automatically to forward media when required. To do this:

1. Go to **Call Control > TURN Servers**, and add details of the TURN server(s) to be used.

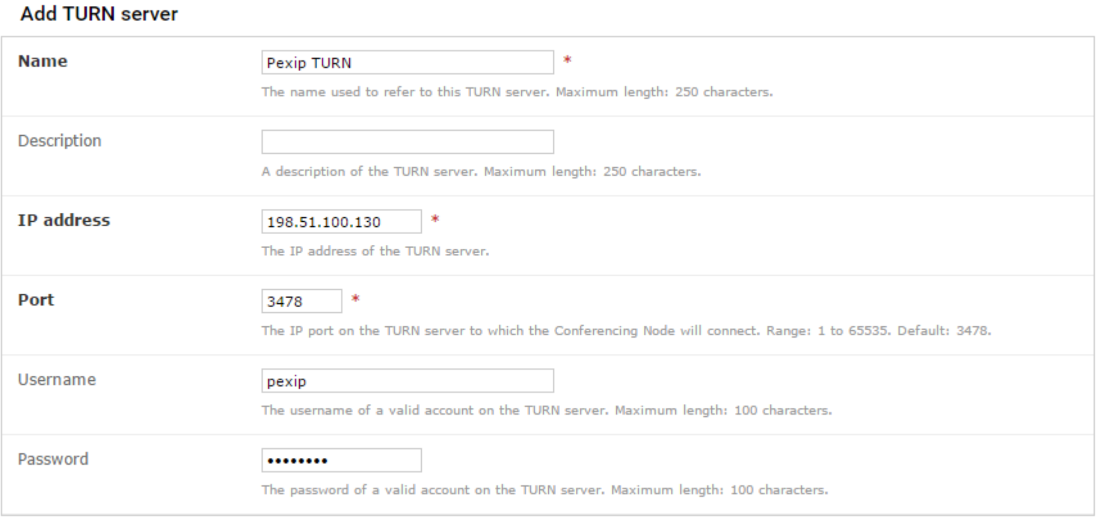

Add TURN server

Name	Pexip TURN * The name used to refer to this TURN server. Maximum length: 250 characters.
Description	 A description of the TURN server. Maximum length: 250 characters.
IP address	198.51.100.130 * The IP address of the TURN server.
Port	3478 * The IP port on the TURN server to which the Conferencing Node will connect. Range: 1 to 65535. Default: 3478.
Username	pexip The username of a valid account on the TURN server. Maximum length: 100 characters.
Password	•••••••• The password of a valid account on the TURN server. Maximum length: 100 characters.

2. Go to **Platform Configuration > Locations**, and for each location, select the TURN server to be used for that location.

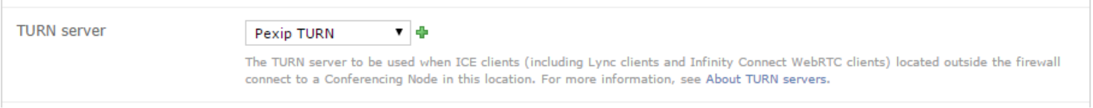

TURN server	Pexip TURN ▾ ✚ The TURN server to be used when ICE clients (including Lync clients and Infinity Connect WebRTC clients) located outside the firewall connect to a Conferencing Node in this location. For more information, see About TURN servers.

Configuring Pexip Infinity to use a STUN server

A STUN server allows clients, such as Conferencing Nodes or Infinity Connect WebRTC clients, to find out their public NAT address.

When a client is deployed behind a NAT, it can send a STUN request to the STUN server, which responds back to the client and tells it from which IP address it received the STUN request. Using this method, the client can discover its public NAT address, which is important in order for ICE to work between Conferencing Nodes and other ICE-enabled clients (for example, WebRTC and Skype for Business / Lync clients). In relation to ICE, this public NAT address is also known as the server reflexive address.

In Microsoft Skype for Business and Lync deployments it is essential that a Conferencing Node can discover its public NAT address.

Conferencing Nodes

If a Conferencing Node is deployed on a private network behind a NAT, its system location may already be configured with the details of a TURN server (such as the Pexip TURN server). Often, that TURN server can act as a STUN server and a separate STUN server is not normally required.

By default, Conferencing Nodes send their STUN requests to the TURN server, but if the TURN server is not located outside of the enterprise firewall then the Conferencing Node will not be able to discover its public NAT address. If this is the case in your deployment scenario, you must configure a separate STUN server — the Conferencing Node's STUN requests will then be sent to the STUN server, instead of the TURN server.

A STUN server is not required if:

- your Conferencing Nodes are publicly-addressable, either directly or via static NAT, or
- the STUN requests sent from the Conferencing Nodes to the TURN server will return the public NAT address of the Conferencing Node.

The STUN servers that you use with Pexip Infinity must be located outside of the enterprise firewall and must be routable from your Conferencing Nodes.

Infinity Connect WebRTC clients

When connecting to a privately-addressed Conferencing Node, Infinity Connect WebRTC clients that are behind a NAT may also use a STUN server to find out their public NAT address.

When an Infinity Connect WebRTC client connects to a Conferencing Node, the node will provision it with any **Client STUN server** addresses that are configured against that node's system location. The WebRTC client can then use those STUN servers to discover its public NAT address. This is typically only required if the WebRTC client is communicating via a TURN server.

For more information, see When is a reverse proxy, TURN server or STUN server required?.

How Conferencing Nodes decide which STUN server to use

The STUN server used by a Pexip Infinity Conferencing Node handling a call is determined as follows:

- **Conferences**: uses the STUN server associated with the location of the Conferencing Node that is handling the call signaling.
- **Point-to-point calls via the Pexip Distributed Gateway**: uses the STUN server associated with the Call Routing Rule that matched the call request. If there is no STUN server associated with the rule, then the STUN server associated with the location of the Conferencing Node that is handling the call signaling is used instead. Note that rules can optionally be configured on a per-location basis.

If a STUN server is not configured for a location or rule, but a TURN server is configured, the Conferencing Node will send STUN requests to that TURN server.

Nominating the STUN servers used by Pexip Infinity

Within Pexip Infinity you can configure the addresses of one or more STUN servers. You then associate those STUN servers with each System location (with separate configuration for the STUN server used by Conferencing Nodes in that location, and the STUN servers to offer to Infinity Connect clients connected to that Conferencing Node), and with each Call Routing Rule.

Configuring STUN server addresses

To add, edit or delete STUN server connection details, go to **Call Control > STUN Servers**. The options are:

Option	Description
Name	The name used to refer to this STUN server in the Pexip Infinity Administrator interface.
Description	An optional description of the STUN server.
Address	The IP address or FQDN of the STUN server. This should not be the same address as any of your configured TURN servers.
Port	The IP port on the STUN server to which the Conferencing Node will connect. Default: 3478.

Note that Pexip Infinity ships with one STUN server address already configured by default: **stun.l.google.com**. This STUN server uses port 19302 (rather than the common 3478) and can be assigned to system locations for use by Infinity Connect WebRTC clients.

You can use this STUN server or configure a different one.

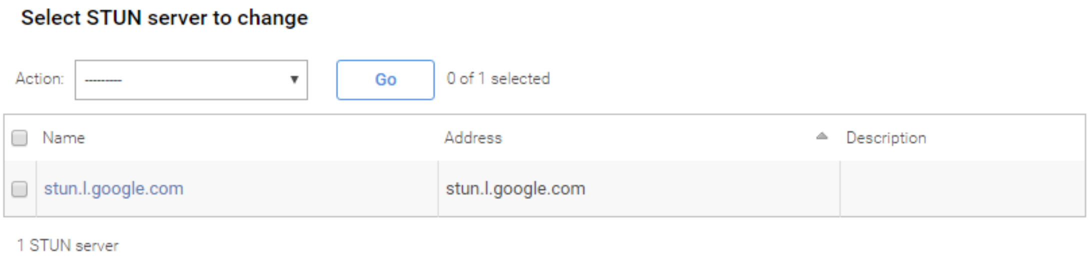

Select STUN server to change

□	Name	Address	Description
□	stun.l.google.com	stun.l.google.com	

1 STUN server

Associating STUN server addresses with Conferencing Nodes

To associate a STUN server address with a Conferencing Node, you must configure the node's system location:

1. Go to **Platform Configuration > Locations**.
2. Select the Conferencing Node's location.
3. Select a **STUN server** and select **Save**.

All Conferencing Nodes in that location will use the nominated STUN server for conference calls.

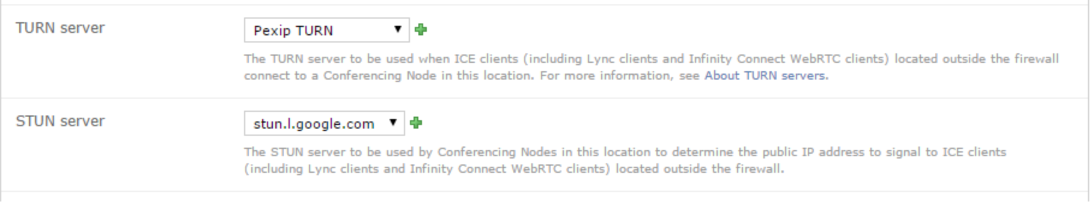

TURN server: Pexip TURN ▼

The TURN server to be used when ICE clients (including Lync clients and Infinity Connect WebRTC clients) located outside the firewall connect to a Conferencing Node in this location. For more information, see About TURN servers.

STUN server: stun.l.google.com ▼

The STUN server to be used by Conferencing Nodes in this location to determine the public IP address to signal to ICE clients (including Lync clients and Infinity Connect WebRTC clients) located outside the firewall.

Associating STUN server addresses with gateway calls

If a gateway call is being placed to an ICE-enabled client (such as Skype for Business / Lync clients and Infinity Connect WebRTC clients), the Conferencing Node placing the call will use the STUN server associated with the matching rule. (For gateway calls, the Conferencing Node does not use the STUN sever associated with its system location.)

To associate a STUN server address with a Call Routing Rule:

1. Go to **Service Configuration > Call Routing**.
2. Select the relevant rule.
3. Select a **STUN server** and select **Save**.

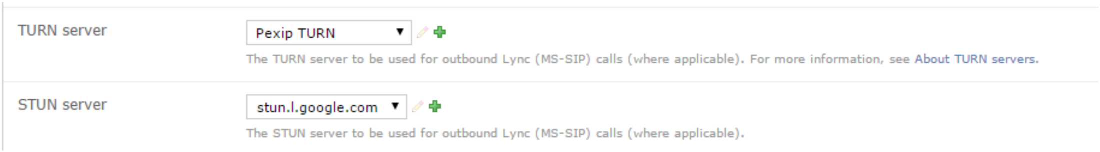

TURN server: Pexip TURN

The TURN server to be used for outbound Lync (MS-SIP) calls (where applicable). For more information, see About TURN servers.

STUN server: stun.l.google.com

The STUN server to be used for outbound Lync (MS-SIP) calls (where applicable).

Configuring the STUN server addresses provided to Infinity Connect WebRTC clients

To configure the specific STUN server addresses that are provisioned to Infinity Connect WebRTC clients, you must configure the system locations used by the Conferencing Nodes that the clients connect to:

1. Go to **Platform Configuration > Locations**.
2. Select the Conferencing Node's location.
3. Select one or more **Client STUN servers** and select **Save**.

When an Infinity Connect WebRTC client connects to a Conferencing Node in that location, the Conferencing Node will provide it with the addresses of the nominated STUN servers. These STUN servers are used by the client to discover its public NAT address.

If no **Client STUN servers** are configured for that node/location, the Infinity Connect client may still be able to communicate by using a TURN relay, if a TURN server is configured on the Conferencing Node, but this may cause delays in setting up media.

For clients on the same network as the Conferencing Nodes, where no NAT is present, users may find that WebRTC call setup time is improved by removing all **Client STUN servers**.

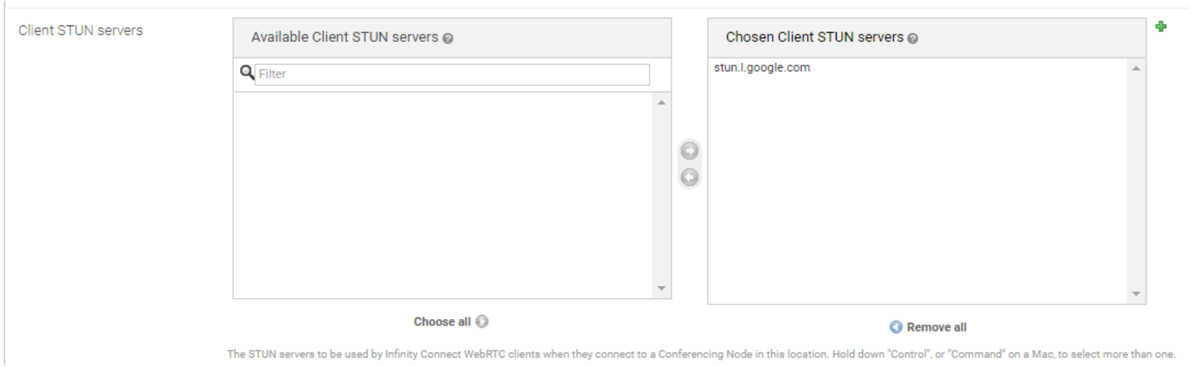

Chosen Client STUN servers: stun.l.google.com

The STUN servers to be used by Infinity Connect WebRTC clients when they connect to a Conferencing Node in this location. Hold down "Control", or "Command" on a Mac, to select more than one.

Appendix 1: Firewall ports

Traffic between the reverse proxy and TURN server and clients in the Internet

The following ports have to be allowed through any firewalls which carry traffic between the reverse proxy and TURN server in the DMZ and the Infinity Connect mobile client and Infinity Connect clients in the public Internet:

Purpose	Direction	Source IP	Protocol	Port	Destination IP
HTTP/HTTPS	Inbound	<any>	TCP	80 / 443	Reverse proxy
UDP TURN/STUN	Inbound	<any>	UDP	3478	TURN server
TURN relay media	Inbound	<any>	UDP	49152–65535	TURN server
RTP media	Outbound	TURN server	UDP	<any>	<any>
DNS	Outbound	Reverse proxy and TURN server	TCP/UDP	53	DNS server
NTP	Outbound	Reverse proxy and TURN server	TCP	123	NTP server

Traffic between the local network and the DMZ / Internet

The following ports have to be allowed through any firewalls which carry traffic between Conferencing Nodes and management stations in the local network and the reverse proxy and TURN server in the DMZ/internet:

Purpose	Direction	Source IP	Protocol	Port	Destination IP
HTTPS	Inbound	Reverse proxy	TCP	443	Conferencing Nodes
UDP TURN/STUN	Outbound	Conferencing Nodes	UDP	3478	TURN server
UDP TURN/STUN	Outbound	Conferencing Nodes	UDP	3478 / 19302	STUN server (if configured). Note that stun.l.google.com uses port 19302.
SSH	Outbound	Management PC	TCP	22	Reverse proxy and TURN server

Appendix 2: Alternative dual NIC reverse proxy/TURN server deployment

This section shows an alternative dual NIC network configuration for the reverse proxy and TURN server.

Dual NIC public/private address with routing to alternative VLAN for management

In this example, the environment is split into two parts: an internal, private network segment and a DMZ network. The private network has two Pexip Infinity Conferencing Nodes and the reverse proxy/TURN server internal interface, while the DMZ perimeter network contains the reverse proxy/TURN server external interface.

Note that all IP addresses in this guide are examples only — actual IP addressing will be deployment specific.

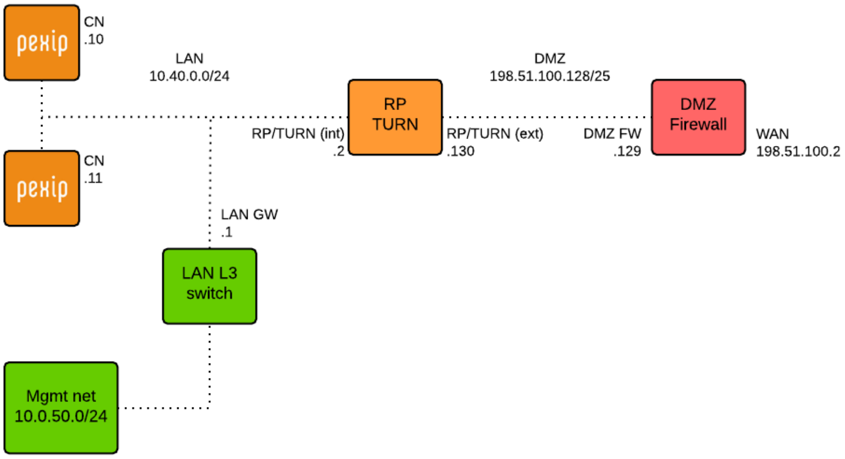

The reverse proxy/TURN server in this case has two network interfaces:

- An internal facing interface (IP address 10.40.0.2) which is connected to the internal LAN network.
- An external facing interface (IP address 198.51.100.130) which is connected to the DMZ network.

The internal interface of the reverse proxy/TURN server is configured on the same subnet as the Conferencing Nodes, while the external interface of the reverse proxy/TURN server is configured on the DMZ subnet. The Conferencing Nodes use 10.40.0.1 as their default gateway. This is also the gateway to the 10.0.50.0/24 management network, from where SSH connections to the internal interface of the reverse proxy/TURN server will originate.

There is no NAT between the outside and the DMZ network segment. The reverse proxy/TURN server uses the DMZ firewall 198.51.100.129 as its default gateway, and is also configured with a static route to the 10.0.50.0/24 management network via the LAN gateway at 10.40.0.1 so that SSH management traffic from this management network can function.

To deploy the reverse proxy/TURN server in this configuration:

1. Deploy the Reverse Proxy and TURN Server appliance OVA file and power on the VM instance.
2. In the install wizard, define the network configuration (steps 1 and 2 in the install wizard) for what will be the internal interface of the reverse proxy/TURN server (10.40.0.2 in this example).
3. For the default gateway (step 3), configure the LAN gateway (10.40.0.1).
4. For the trusted hosts/networks (step 4), configure 10.0.50.0 255.255.255.0.
5. After the install wizard completes, the reverse proxy/TURN server will reboot with the new configuration. When it is back up, connect over SSH to the reverse proxy from a host in the management network (or log in via the VMware console) as user 'pexip' and shut down the virtual machine using the following command:

   ```
   sudo shutdown -h now
   ```

   Input the password for user 'pexip' when prompted by sudo.

6. After the VM has been shut down, add one additional network interface to this VM instance and power it on again.
7. Log in to the reverse proxy/TURN server through SSH or VMware console and add the network configuration for the external network interface (198.51.100.130 in this example) and the static route configuration to /etc/network/interfaces, as described in Configuring a second NIC.

   In this example, the resulting **/etc/network/interfaces** file should contain the following (note that the existing gateway entry for interface eth0 is removed as the reverse proxy/TURN server should only have one default gateway defined):

   ```
   auto lo eth0 eth1
   iface lo inet loopback

   iface eth0 inet static
     address 10.40.0.2
     netmask 255.255.255.0

   iface eth1 inet static
     address 198.51.100.130
     netmask 255.255.255.128
     gateway 198.51.100.129
     dns-nameservers 8.8.8.8 8.8.4.4

   post-up route add -net 10.0.50.0 netmask 255.255.255.0 gw 10.40.0.1
   ```

   To ensure that the eth0 and eth1 interfaces correspond with the correct port group in VMware, running the `ifconfig` command on the reverse proxy/TURN server will show the hardware MAC addresses of each interface, so that these can be matched against the virtual interfaces in VMware.
8. Reboot the reverse proxy/TURN server VM to apply the new network settings.

STUN server configuration

If the reverse proxy/TURN server is deployed with a dual NIC public/private address, then Conferencing Nodes will typically not be able to discover their public NAT addresses as they will sending their STUN requests to the internal interface of the TURN server.

Therefore, you will need to configure Pexip Infinity with a separate STUN server (see Configuring Pexip Infinity to use a STUN server).

Appendix 3: Extra configuration and maintenance tasks

This section describes some additional configuration scenarios and maintenance tasks for the reverse proxy and TURN server:

Adding or removing Conferencing Nodes from an existing reverse proxy configuration

To add an additional Conferencing Node, or to remove a Conferencing Node from an existing reverse proxy configuration:

1. Run the following command to edit the Reverse Proxy and TURN Server config file:
   ```
   sudo nano /etc/nginx/sites-available/pexapp
   ```
2. To add a Conferencing Node:

   In section `upstream pexip` at the top of the config file, add additional `server` statements. In these additional entries, specify the IP address of each additional Conferencing Node; other parameters should be similar to the existing entries.
3. To remove a Conferencing Node:

   In section `upstream pexip` at the top of the config file, remove the `server` statements that specify the IP address of each Conferencing Node to be removed.
4. After editing the file, run the following command to reload the nginx configuration gracefully (not interrupting any existing sessions):
   ```
   sudo service nginx reload
   ```

Configuring NAT for the TURN server

To configure NAT, run the following command to edit the TURN server configuration:

```
sudo nano /etc/turnserver.conf
```

Anywhere in the config file, add the parameter `external-ip=1.2.3.4`, where `1.2.3.4` is the public NAT address of the TURN server.

After editing the file, run the following command to make the configuration change take effect:

```
sudo service rfc5766-turn-server restart
```

Note that this will interrupt any existing TURN sessions (e.g. any WebRTC or Skype for Business / Lync calls going via the TURN server), therefore these changes should be performed during a maintenance window that is outside of normal operating hours.

Configuring a second NIC

The first, existing NIC must be the internally-facing NIC. When a reverse proxy is deployed and the intention is to have dual NICs, the network configuration provided during the install wizard must be the internally-facing NIC, as SSH access will only be enabled for this initial NIC.

To configure a second NIC, run the following command to edit the network configuration file:

```
sudo nano /etc/network/interfaces
```

The existing interfaces file will be similar to this:

```
# >/etc/network/interfaces
# Written at 2014-04-23 10:30:34 UTC by InterfacesFileWriter
auto lo eth0
iface lo inet loopback

iface eth0 inet static
   address 10.40.0.2
   netmask 255.255.255.0
   gateway 198.51.100.129
   dns-nameservers 8.8.8.8 8.8.4.4
```

To add a second NIC:

1. Change the line `auto lo eth0` to `auto lo eth0 eth1`.
2. Create a new section called `iface eth1 inet static`, and create entries for `address`, `netmask`, `gateway` and `dns-nameservers`.
3. Remove the existing `gateway` parameter for eth0 (as eth0 will now be the internally-facing interface and thus should not have a default gateway; the default gateway should belong to the externally-facing eth1).

Important: static routes

If any static routes are needed in this scenario (which is usually the case unless the reverse proxy does not need to access any hosts outside of the internal eth0 subnet), these have to be defined at the bottom of the interfaces file, using syntax as follows:

```
post-up route add -net 10.0.50.0 netmask 255.255.255.0 gw 10.40.0.1
```

The above entry will create a static route for 10.0.50.0/24 via 10.40.0.1.

The resulting interfaces file will then look like this:

```
# >/etc/network/interfaces
# Written at 2014-04-23 10:30:34 UTC by InterfacesFileWriter
auto lo eth0 eth1
iface lo inet loopback

iface eth0 inet static
   address 10.40.0.2
   netmask 255.255.255.0
   dns-nameservers 8.8.8.8 8.8.4.4

iface eth1 inet static
   address 198.51.100.130
   netmask 255.255.255.128
   gateway 198.51.100.129
   dns-nameservers 8.8.8.8 8.8.4.4

post-up route add -net 10.0.50.0 netmask 255.255.255.0 gw 10.40.0.1
```

After the configuration has been changed and saved, the following command will make the changes take effect:

```
sudo service networking restart
```

Alternatively, you can reboot the Reverse Proxy and TURN Server appliance with the command `sudo reboot`.

TURN server

If the reverse proxy and TURN applications co-exist on the same virtual machine, you must also edit the TURN server configuration file to specify the external IP address.

Run the following command to edit the TURN server configuration file:

```
sudo nano /etc/turnserver.conf
```

Change the line that begins relay-ip= to specify the external IP address, e.g. relay-ip=198.51.100.130

After editing the file, run the following command to make the configuration change take effect:

```
sudo service rfc5766-turn-server restart
```

Note that this will interrupt any existing TURN sessions (e.g. any WebRTC or Skype for Business / Lync calls going via the TURN server), therefore these changes should be performed during a maintenance window that is outside of normal operating hours.

Changing the TURN server's access credentials

To set new access credentials for the TURN server, run the following command:

```
sudo turnadmin -k -a -b /etc/turnuserdb.conf -u <username> -r <realm> -p <password>
```

where `<username>` and `<password>` are the credentials to be applied, and `<realm>` is the domain of your deployment.

You can check the realm in use in your deployment by looking at your TURN server configuration file (by running `cat /etc/turnserver.conf`) and checking the `realm=` line.

If you change the credentials, remember to update the TURN server configuration on Pexip Infinity (**Call Control > TURN Servers**).

Restoring the Reverse Proxy and TURN Server to its default state

If you need to re-run the install wizard, you must first reset the appliance to its original default state. To do this you must either reinstall the application, or restore a snapshot taken after initially deploying the OVA template.