Pexip Reverse Proxy and TURN Server

Deployment Guide

Software Version 6.1.2

Document Version 6.e

February 2021

]pexip[

Contents

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 of Pexip Infinity, 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:

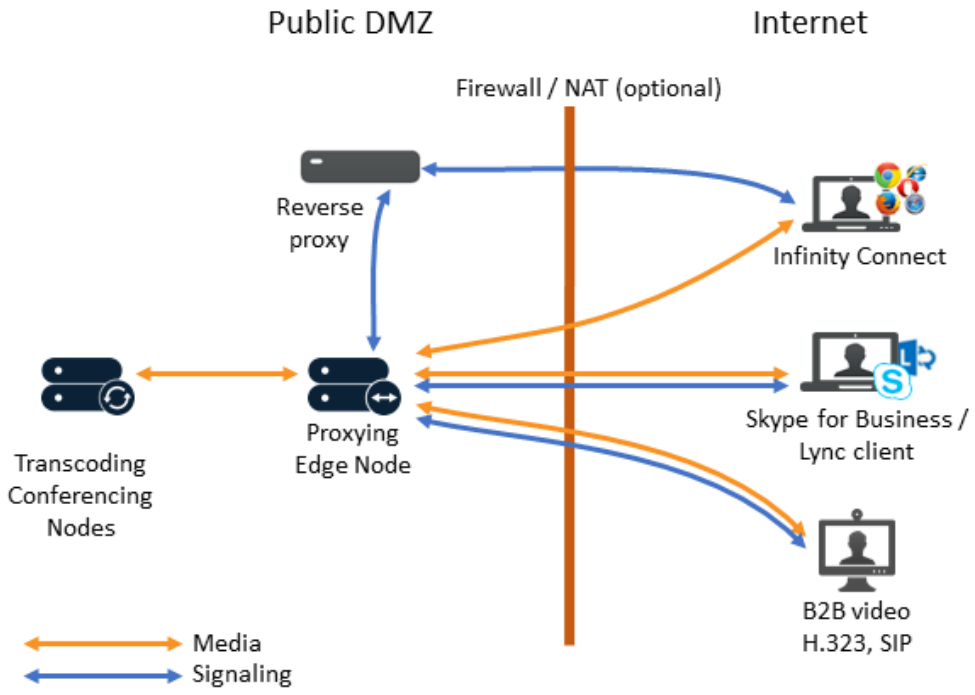

Deployments that do not use Proxying Edge Nodes

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

WebRTC clients (the Infinity Connect web app on the latest browsers, and the desktop and mobile clients) 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.

Note that Microsoft Edge browser version 44 and earlier (which is WebRTC-compatible) cannot use STUN and thus cannot send media to Pexip Infinity via a TURN server.

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 reverse proxy / TURN server 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.

As general good practice, we always recommend deploying the TURN server in a suitably secured network segment, such as a DMZ.

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 and mobile clients.

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.

When installing/enabling the TURN server you must specify the IP addresses of the Conferencing Nodes that will use the TURN server for media relay. This locks down the IP addresses that are allowed (safelisted) to communicate with the TURN server over UDP/3478.

For conferencing services, we recommend that you install your own SSL/TLS certificates on the reverse proxy for maximum security. If you are using VMR Scheduling for Exchange you must install your own 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 and load balance requests from Outlook clients.

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 (and they are not configured for static 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 its configured 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.
- As general good practice, we always recommend deploying the TURN server in a suitably secured network segment, such as a DMZ.

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 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 is referred to as such in this guide.

Example communication scenario

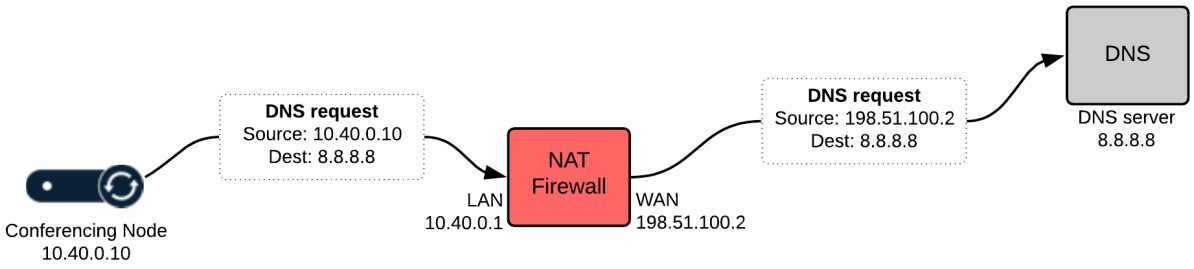

Using the above diagram as an example, the Conferencing Node has an IP address of 10.40.0.10 – this is a private/internal IP address which is not routable across public networks. When this Conferencing Node communicates with a host located on a public network (Internet), for instance a DNS server, traffic 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 appears 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 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 of Pexip Infinity, 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.

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 clients to access your Pexip Infinity services, and you may need to use a TURN server for Skype for Business / Lync clients. 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 reverse proxy / TURN server deployment: single NIC on public address

Here is 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 is deployment specific.

This example forms the basis of this guide. In this scenario:

- 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
  - Conferencing Nodes to send TURN packets to the TURN server on UDP port 3478
  - Conferencing Nodes to send STUN packets to the STUN server, typically on UDP port 3478, although stun.l.google.com uses port 19302.

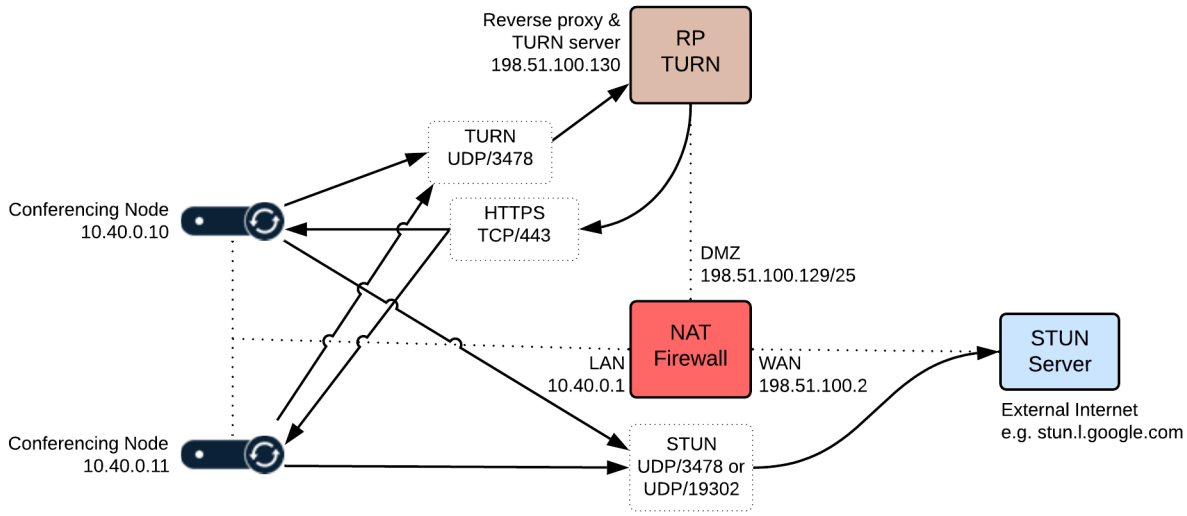

Example deployment used in this guide: single NIC on public address

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 is 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.

The Reverse Proxy and TURN Server appliance is also available as an Amazon Machine Image (AMI) on Amazon Web Services (AWS). When choosing your AMI, select **Community AMIs**, search for "Pexip" and select the Pexip Reverse Proxy.

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 6.1 of the Pexip RP/TURN OVA template from https://dl.pexip.com/rpturn/index.html (select the v6.1 directory) to the PC running the vSphere web client.

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).

See Reverse Proxy and TURN Server release notes for the latest information on new features and resolved issues.

Deploying the OVA template

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

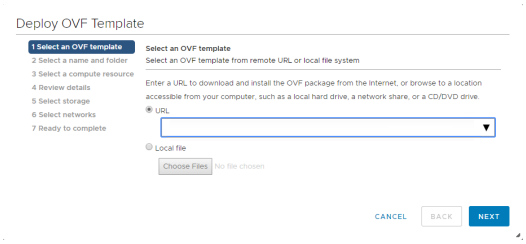

To deploy the OVA template:

1. Using the **vSphere web client**, go to **Hosts And Clusters**, click **File** and **Deploy OVF Template** (this option accepts OVA files).
2. 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.
3. 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 a console for the Reverse Proxy and TURN Server virtual machine.

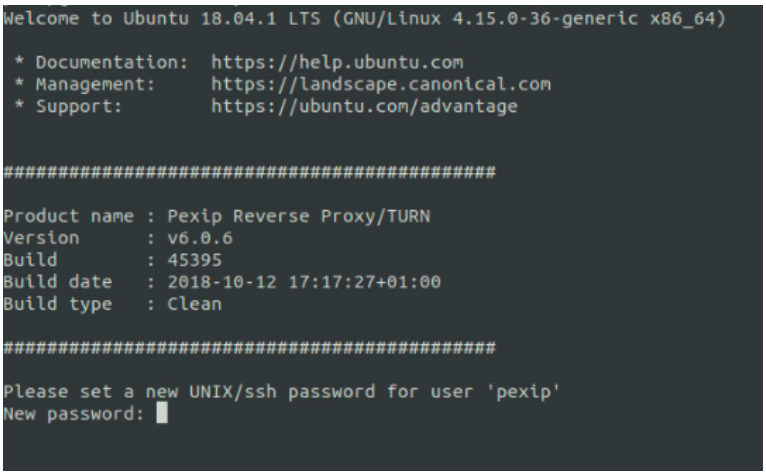

Initial login prompt

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. The password must:
   - have a minimum of 8 characters
   - satisfy at least 3 out of the following 4 conditions:
     - one lower case character
     - one upper case character
     - one special character
     - one digit.
3. After setting the new password, the install wizard starts and you log in again with the new password.

Running the installation wizard

The installation wizard is divided into several steps, which are explained below. Some steps are only presented if they are appropriate — for example, if dual network interfaces are detected, or if you choose to enable the TURN server.

The example configuration values described here are based on the Example reverse proxy / TURN server 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 is deployment specific.

The following table shows, for each step, the prompt text that is shown, an explanation of the step and some example input. If you subsequently rerun the installation wizard, the default values for the questions use the answers from the previous run (if they are still valid).

Step / Wizard prompt text	Example value	Description
1	**Network / NIC configuration**	
		Single NIC — these steps only apply if dual NICs are not detected, otherwise skip to dual NICs detected. In the example single NIC deployment scenario shown above, this step would be presented.

Step / Wizard prompt text		Example value	Description
1.1	IP Address?	198.51.100.130	The IP address of the appliance. Defaults to a value suggested by DHCP if available.
1.2	Netmask?	255.255.0.0	The network mask for the appliance. Defaults to a value suggested by DHCP if available.
1.3	Default gateway?	198.51.100.129	The IP address of the appliance's default gateway. Now skip to step 2 DNS servers.
Dual NICs detected			
1.4	<list of detected NICs> Do you want to configure Dual NIC?		This step only applies if 2 or more network interfaces are detected on the underlying virtual machine. In the example deployment scenario shown above, this step would not be presented. Values: yes/no Default: no
If Dual NICs are available but not selected (answer = no):			
1.5	Which network interface should be used?		Enter a NIC from the printed list (e.g. nic0) to be used as the (single) interface for the appliance. Then proceed as for a single NIC deployment above (step 1.1).
If Dual NICs are available and selected (answer = yes):			
1.6	Which is the internal-facing network interface?		Enter the name of the internal-facing network interface. Values: expects a NIC from the printed list (e.g. nic0).
1.7	Which is the external-facing network interface?		Enter the name of the external-facing network interface. It automatically fills out the answer if one NIC is left, otherwise it expects a NIC from the list which has not already been chosen (e.g. nic1).
In this case, steps 2 and 3 of the wizard now capture the interface addresses:			
2.1	IP Address for internal interface?		The IP address of the internal interface of the appliance (e.g. 10.44.0.5). Defaults to a value suggested by DHCP if available.
2.2	Subnet mask for internal interface?		The network mask for the internal interface (e.g. 255.255.0.0) Defaults to a value suggested by DHCP if available.

Step / Wizard prompt text			Example value	Description
	2.3	Add a custom network route for <internal NIC>?		Values: yes/no Default: no If "yes" you are prompted with: • "Network IP Address?" - enter a base IP address (e.g. 192.168.0.0). You can also supply an address in CIDR format (e.g. 192.168.0.0/16), in which case the following netmask question is skipped. • "Network Netmask?" - enter a netmask (e.g. 255.255.0.0). • "Via IP Address?" - enter an IP address that is accessible from the interface's IP/netmask (e.g. 10.44.0.1). • "Add another custom network route for <internal NIC>?" - enter "yes" to add another route (jumps back to "Network IP Address?"), or "no" to continue on to the next step. If you are rerunning the wizard, and you previously entered some custom routes, the previous routes are listed and you are asked "Use these default values?" instead. Reply "yes" (default) to reuse the previous routes, or "no" to change the routes via this step.
	3.1	IP Address for external interface?		The IP address of the external interface of the appliance (e.g. 198.51.100.130). Defaults to a value suggested by DHCP if available.
	3.2	Subnet mask for external interface?		The network mask for the external interface (e.g. 255.255.0.0) Defaults to a value suggested by DHCP if available.
	3.3	Default gateway for external interface?		The address of the appliance's default gateway (e.g. 198.51.100.129).
	After configuring dual NICs, the remaining step numbers are now as shown below + 2, e.g. DNS servers becomes step 4 and so on.			
2	**DNS servers**			
	2.1	<list of default DNS servers found> Use these default values?	yes	This step only applies if DNS servers were saved from a previous run or detected from DHCP. Values: yes/no Default: yes
	2.2	DNS server 1		This step only applies if not using the default DNS servers. Enter an IP address (e.g. 10.44.0.9) or an empty line. The question is repeated ("DNS server 2" etc) until an empty line is received. At least one DNS server must be configured. DNS is mainly used to resolve the hostnames of NTP servers.
3	**Hostname and Domain**			
	3.1	Hostname?	proxy	The hostname of the appliance. The hostname and domain (configured in the next step) must match the actual DNS name by which the appliance will be addressed.

Step / Wizard prompt text			Example value	Description
	3.2	Domain?	example.com	The domain suffix of the appliance. This means that if the host name in the previous step was configured as **proxy**, that the full FQDN of this appliance is **proxy.example.com**. Defaults to a value suggested by DHCP if available. If a custom SSL certificate is created for the 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.
4	**NTP servers**			
	4.1	<list of default NTP servers> Use these default values?	yes	We recommend that at least three NTP servers are used to ensure proper NTP time synchronization. The default list of NTP servers is 0.pexip.pool.ntp.org, 1.pexip.pool.ntp.org, 2.pexip.pool.ntp.org. Values: yes/no Default: yes
	4.2	NTP server 1		This step only applies if not using the default NTP servers. Enter an IP address (e.g. 10.44.0.5), an FQDN (e.g. 0.pexip.pool.ntp.org), or an empty line. The question is repeated ("NTP server 2" etc) until an empty line is received. At least one NTP server must be configured.
5	**Web reverse proxy**			
	5.1	Enable web reverse proxy?	yes	This step determines if the reverse proxy functionality is enabled. Values: yes/no Default: yes If "no", skip to step 6 TURN server.
	5.2	IP Address of signaling Conferencing Node 1?	10.40.0.10 [ENTER] 10.40.0.11 [ENTER] [ENTER]	Specify the Conferencing Nodes that will receive the proxied HTTP/HTTPS (signaling) requests. Enter the IP address of a Pexip Conferencing Node or an empty line when finished. We recommend that these nodes are configured as Transcoding Conferencing Nodes. The question is repeated ("Conferencing Node 2" etc) until an empty line is received. At least one Conferencing Node must be configured. If you are rerunning the wizard, and you previously entered some node addresses, the previous addresses are listed and you are asked "Use these default values?" instead. Reply "yes" (default) to reuse the previous addresses, or "no" to change the addresses via this step.

Step / Wizard prompt text		Example value	Description
5.3	<CSP information message> Enable Content Security Policy?	yes	Content-Security-Policy is an HTTP header that provides enhanced security against cross-site scripting attacks. Values: yes/no Default: yes Enter 'yes' to enable Content Security Policy. This is recommended if you are **not** using optional advanced features such as plugins 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, otherwise we recommend that you enable this option.
6	**TURN server**		
6.1	Enable TURN server?	yes	If you answered "no" to "Enable web reverse proxy?", this question is skipped and the TURN server is enabled by default. In our example deployment, this question is asked and the answer is "yes". Values: yes/no Default: yes
6.2	Do you want the TURN server to listen on port 443 instead of 3478?	no	This only applies if you answered "no" to "Enable web reverse proxy?". It allows you to configure the TURN server to listen on port 443 instead of 3478. In our example deployment the answer is "no". See Configuring the TURN server for TCP TURN relay for more information. Values: yes/no Default: no
6.3	Username?	pexip	The username and password are the credentials you must use when configuring Pexip Infinity with the access details for this TURN server. A username must be entered.
6.4	Password?	admin123	A password must be entered. If you subsequently rerun the installation wizard you must resupply this password.
6.5	IP Address of media Conferencing Node 1?		Specify the addresses of the Conferencing Nodes that will use this TURN server for media relay. ⓘ Only the addresses specified here will be allowed (safelisted) to communicate with the TURN server over UDP/3478. Enter the IP address of a Pexip Conferencing Node or an empty line when finished. The question is repeated ("Conferencing Node 2" etc) until an empty line is received. At least one Conferencing Node must be configured. This defaults to the list provided for the **IP Address of signaling Conferencing Nodes** step, if that question was asked. If you are rerunning the wizard, and you previously entered some node addresses, the previous addresses are listed and you are asked "Use these default values?" instead. Reply "yes" (default) to reuse the previous addresses, or "no" to change the addresses via this step. ⓘ If you later add more Conferencing Nodes to your platform, and those nodes need to use this TURN server, you will need to rerun the wizard and add the addresses of those new nodes.
7	**Management networks (IP tables rules for SSH access)**		

Step / Wizard prompt text			Example value	Description
	7.1	Add a management network?	yes	This allows the enterprise's management network to access this host over SSH. Values: yes/no Default: yes If you are rerunning the wizard, and you previously entered some management networks, the previous networks are listed and you are asked "Use these default values?" instead. Reply "yes" (default) to reuse the previous networks, or "no" to change the networks via this step. If "no" skip to step 8 Fail2ban. Note that the SSH service is disabled if no management networks are configured.
	7.2	IP Address for the management network 1?	10.0.50.0	Enter a base IP address. You can also supply an address in CIDR format (e.g. 10.0.50.0/24), in which case the following netmask question is skipped.
	7.3	Netmask for management network 1?	255.255.255.0	The network mask for the management network.
	7.4	Add another management network?	no	Enter "yes" to add another network (jumps back to "IP Address for the management network?"), or "no" to continue on to the next step.
8	**Fail2ban**			
	8.1	Enable Fail2ban?		Fail2ban is an intrusion prevention framework that can protect the reverse proxy from brute-force attacks on PIN-protected conferences. Values: yes/no Default: no See Enabling fail2ban for more information.
9	**SNMPv2c**			
	9.1	Enable SNMPv2c read only?		Enables SNMPv2c read only access. Values: yes/no Default: no If "no", skip to step 10 certificates.
	9.2	SNMP community?		The SNMP trap community name. Default: public
	9.3	SNMP system location?		A description of the appliance's location.
	9.4	SNMP system contact?		The contact details (for example, email address) of the person responsible for this particular appliance.
	9.5	SNMP system name?		A name for this appliance.
	9.6	SNMP system description?		A description for this appliance.

Step / Wizard prompt text			Example value	Description
10	**Certificates (these steps only apply if you are rerunning the installation wizard; on the first run, the certificates are regenerated without asking)**			
	10.1	Do you want to regenerate a new SSL certificate?		The SSL certificate is used when accessing conferencing services. Values: yes/no Default: yes See Replacing the default SSL certificate for instructions on how to upload your own TLS certificate. If you have uploaded your own SSL certificate and do not want it to be replaced, ensure that you answer "no" to this option.
	10.2	Do you want to regenerate a new SSH certificate?		The SSH certificate is used when connecting over SSH for maintenance tasks. Values: yes/no Default: yes

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 users can now access VMRs 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 for maximum security. If you are using VMR Scheduling for Exchange you must install your own 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 in the format:

   ```
   -----BEGIN CERTIFICATE-----
   <certificate>
   -----END CERTIFICATE-----
   ```

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

   ```
   -----BEGIN PRIVATE KEY-----
   <private key>
   -----END PRIVATE KEY-----
   ```

   Finally, there will normally be one or more intermediate CA certificates, where each intermediate has a section in the following format:

   ```
   -----BEGIN CERTIFICATE-----
   <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/macOS, 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, 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 systemctl restart nginx
```

Note that `sudo systemctl restart nginx` 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 systemctl restart nginx
```

If you rerun the installation wizard you are given the option `Do you want to regenerate a new SSL certificate?` Ensure that you answer "no" to this option if you want to preserve your own certificate.

Enabling fail2ban

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

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

This section describes how to enable fail2ban, modify the default configuration if required, and how to monitor the service.

Enabling fail2ban

You can enable or disable fail2ban when you run (or rerun) the installation wizard.

At the `Enable Fail2ban?` prompt in the installation wizard, answer "yes" to enable fail2ban, or "no" to disable it. The default is "no" when you first install the reverse proxy, and if you rerun the installation wizard it defaults to your previous setting.

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. Connect over SSH to the reverse proxy and log in as user **pexip**.
2. Run the following command to edit the fail2ban config file:
   ```
   sudo nano /etc/fail2ban/jail.local
   ```
3. 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

4. After editing and saving the file, run the following command to restart the fail2ban service:
   ```
   sudo systemctl restart fail2ban
   ```

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 does 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 browser version 44 and earlier (which is WebRTC-compatible) cannot 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 and mobile clients

The Infinity Connect desktop and mobile clients 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 desktop or mobile client attempts to place a call to 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 and mobile clients to initially send their HTTP requests to host **px01.vc.example.com** (our primary node) on TCP port 443. The clients will also try to use host **px02.vc.example.com** (our fallback node) if they cannot contact px01.

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

Infinity Connect desktop client

In this example, when a user attempts to place a call to **meet.alice@vc.example.com**, the client does **one** of the following:

- If the client is registered to Pexip Infinity and the global Route via registrar setting is enabled, the client will route all calls directly to the IP address of the Conferencing Node to which it is registered, regardless of the domain being dialed.

  For example, if the client is configured with a Registration Host of **registration.example.com**, then the client will perform an SRV lookup on **_pexapp._tcp.registration.example.com**.

  If the SRV lookup fails, or none of the returned hosts in the lookup can be contacted, the client will also attempt to connect directly to that domain, i.e. to **http://registration.example.com:443** (via DNS A-records for **registration.example.com**).
- 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 none of the returned hosts in the lookup can 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 are performed, and the client will report that it could not join the host domain.
- If a **serverAddress** has been configured, the client performs an SRV lookup on that domain, and attempts to contact the host(s) returned in that lookup.

  For example, if the **serverAddress** is **localserver.example.com** then the client performs an SRV lookup on **_pexapp._tcp.localserver.example.com**.

  If the SRV lookup fails, or none of the returned hosts in the lookup can be contacted, the client also attempts 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 are performed, and the client will report that it could not join the host domain.
- In all other cases, the client attempts 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 returns 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 next tries to contact **px02.vc.example.com**.

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

Infinity Connect mobile client

In this example, when a user attempts to place a call to **meet.alice@vc.example.com**, the client does **one** of the following:

- 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 none of the returned hosts in the lookup can 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 are performed, and the client will report that it could not join the host domain.
- In all other cases, the client attempts 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 returns 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 next tries to contact **px02.vc.example.com**.

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

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.

   The username and password credentials must match the values you specified in the Reverse Proxy and TURN Server installation wizard.

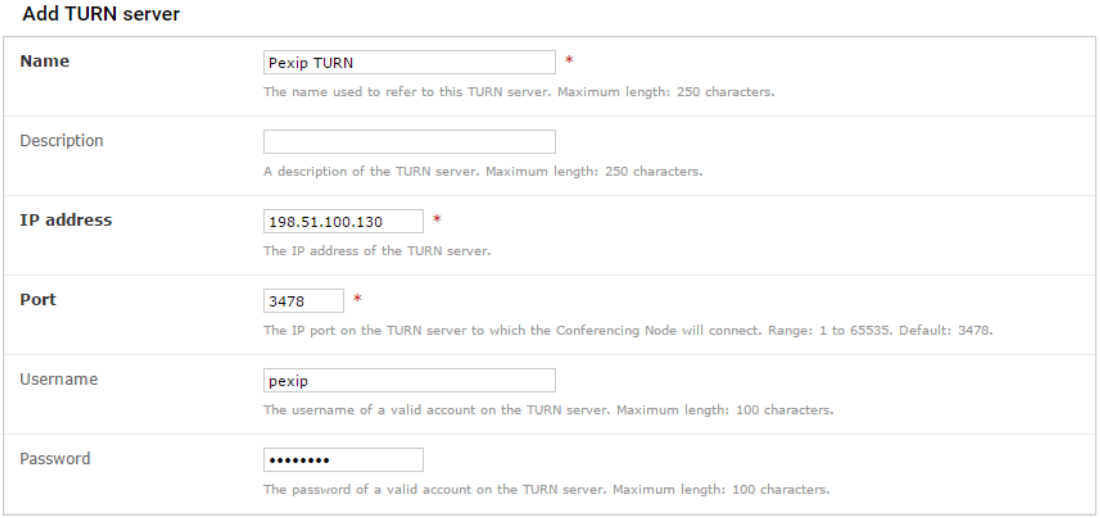

2. Go to **Platform > Locations**, and for each location, select the TURN server to be used for that location.

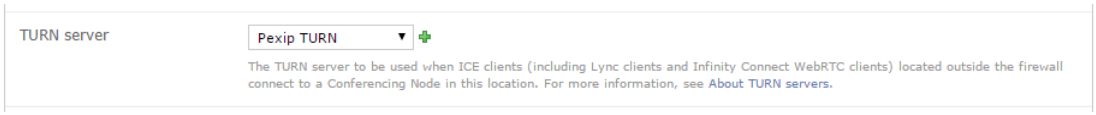

3. If you are using Pexip's TURN server appliance, as of version 6.1.0 you must ensure that it is configured with the IP addresses of the Conferencing Nodes that will use it for media relay. You do this during the installation wizard steps when deploying the TURN server.

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 used by 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 Infinity 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.

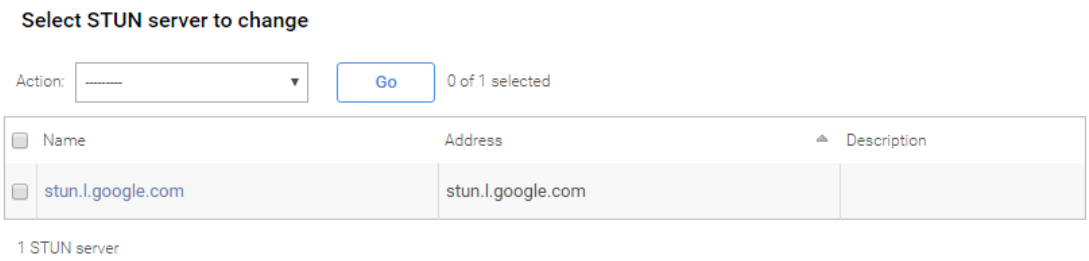

Associating STUN server addresses with Conferencing Nodes

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

1. Go to **Platform > 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.

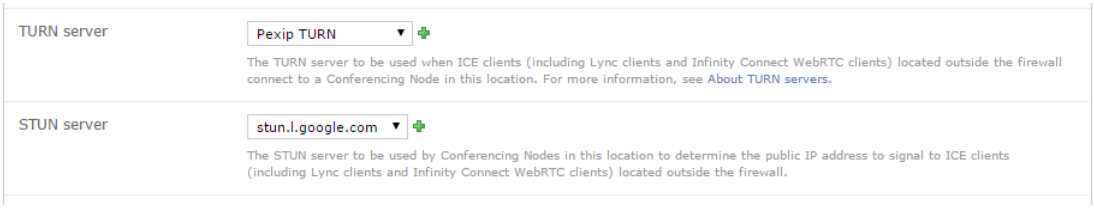

Associating STUN server addresses with gateway calls

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 **Services > Call Routing**.
2. Select the relevant rule.
3. Select a **STUN server** and select **Save**.

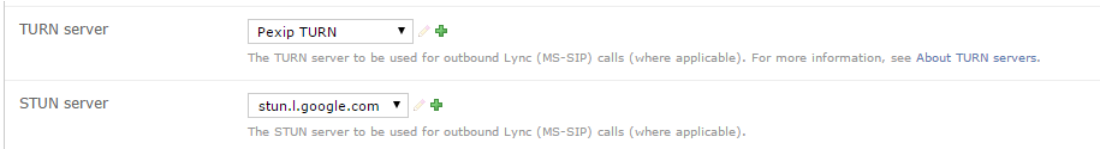

Configuring the STUN server addresses provided to Infinity Connect WebRTC clients

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 > 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**.

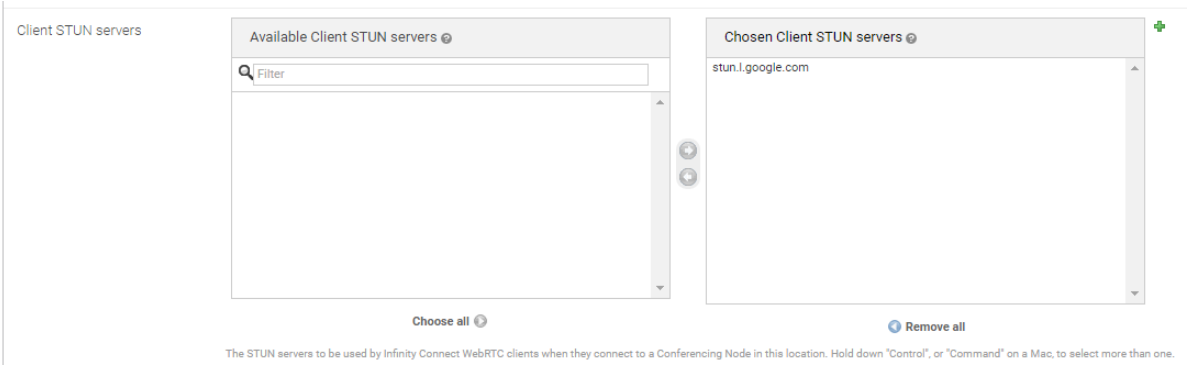

Appendix 1: Firewall ports

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 Infinity Connect clients in the public Internet:

Source address	Source port	Destination address	Dest. port	Protocol	Notes
<any> (Infinity Connect client)	<any>	Reverse proxy	80 / 443	TCP	HTTP/HTTPS
<any> (Infinity Connect client)	<any>	TURN server	3478	UDP	UDP TURN/STUN
<any> (Infinity Connect client)	<any>	TURN server	49152–65535	UDP	TURN relay media
<any> (Infinity Connect client)	<any>	TURN server	443	TCP	TURN relay media †
TURN server	49152-65535	<any>	<any>	UDP	RTP media
Reverse proxy / TURN server	<any>	DNS server	53	TCP/UDP	DNS
Reverse proxy / TURN server	<any>	NTP server	123	TCP	NTP

† Only applies if TURN over TCP/443 is enabled.

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:

Source address	Source port	Destination address	Dest. port	Protocol	Notes
Reverse proxy	<any>	Conferencing Nodes	443	TCP	HTTPS
Conferencing Nodes	40000–49999 **	TURN server	3478	UDP	UDP TURN/STUN
Conferencing Nodes	40000–49999 **	STUN server (if configured)	3478 / 19302	UDP	UDP TURN/STUN. Note that stun.l.google.com uses port 19302.
Management PC	<any>	Reverse proxy / TURN server	22	TCP	SSH
SNMP server	<any>	Reverse proxy / TURN server	161	UDP	SNMP ‡
Reverse proxy / TURN server	<any>	SNMP server	161	UDP	SNMP ‡

** Configurable via the Media port range start/end, and Signaling port range start/end options.

‡ Only applies if SNMP is enabled.

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 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 is deployment specific.

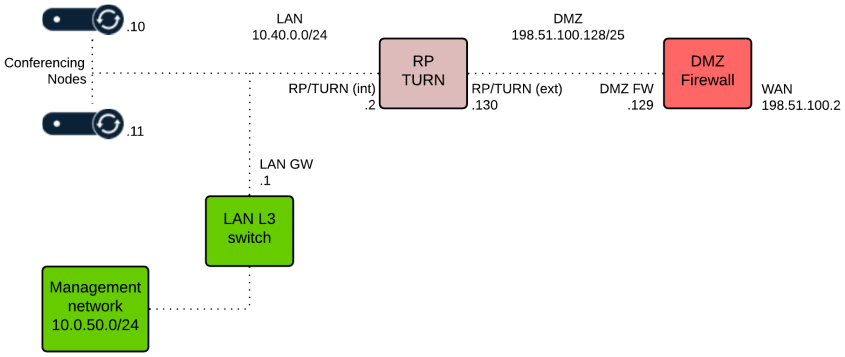

The reverse proxy/TURN server in this case has two network interfaces:

- 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. Complete the installation wizard steps for NIC configuration:

Prompt	Example input	Comments
<list of detected NICs> Do you want to configure Dual NIC?	yes	This option is only presented if multiple NICs are detected.
Which is the internal-facing network interface?	eth0	Enter the name of the NIC that will be the internal-facing interface.
Which is the external-facing network interface?	eth1	If there are just two NICs this is automatically filled with the name of the other interface.
IP Address for internal interface?	10.40.0.2	The IP address of the internal interface of the appliance.

Prompt	Example input	Comments
Subnet mask for internal interface?	255.255.255.0	The network mask for the internal interface.
Add a custom network route for <internal NIC>?	yes	In this example we want to configure a static route to the 10.0.50.0/24 management network via the LAN gateway at 10.40.0.1.
Network IP Address?	10.0.50.0/24	In this example the address is specified in CIDR format (thus you are not also prompted for a netmask).
Via IP Address?	10.40.0.1	This is the LAN gateway address
Add another custom network route for <internal NIC>?	no	In this example we only want to add one route.
IP Address for external interface?	198.51.100.130	The IP address of the external interface of the appliance.
Subnet mask for external interface?	255.255.255.128	The network mask for the external interface.
Default gateway for external interface?	198.51.100.129	The address of the appliance's default gateway.

3. Complete the remaining steps of the installation wizard as appropriate for your environment (see Running the installation wizard for a full list of the remaining installation wizard steps).
4. After the install wizard completes, the reverse proxy/TURN server will reboot with the new configuration.

If you want to confirm that the eth0 and eth1 interfaces correspond with the correct port group in VMware:

1. Log in to the reverse proxy/TURN server through SSH or VMware console as user 'pexip'.
2. Run the `ip addr` command.
   This shows the hardware MAC addresses of each interface, so that these can be matched against the virtual interfaces in VMware.

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 be 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:

Patching the operating system for the latest security bugs

We recommend that you keep the appliance's Operating System patched against the latest security bugs. The frequency with which you check for patches depends upon your local security policies but we recommend at least once per month, or whenever an important or critical CVE (Common Vulnerabilities and Exposures) has been resolved in Ubuntu.

Note that some updates may need some of the system services to be restarted or require a reboot of the VM, so we recommend performing these updates when a potential temporary loss of service is acceptable.

To install the latest security patches:

1. Take a VM snapshot of the appliance.
2. Log in to the reverse proxy/TURN server through SSH or VMware console as user 'pexip'.
3. Run the following command:
   ```
   sudo apt-get update && sudo apt-get dist-upgrade
   ```
4. Delete the snapshot after 1-2 days, when it has been confirmed that the patches have not had any detrimental effect.

Configuring the operating system to use a proxy for software upgrades

To configure the appliance's operating system to use a proxy for software upgrades:

1. Connect over SSH to the reverse proxy.
2. Create a proxy configuration file on the appliance with the following command:
   ```
   sudo nano /etc/apt/apt.conf.d/proxy.conf
   ```
3. Add the following contents via the text editor:
   ```
   Acquire {
     HTTP::proxy "http://user:password@proxy.server:port/";
     HTTPS::proxy "http://user:password@proxy.server:port/";
   }
   ```
   where the **user** and **password** credentials and the proxy's hostname/port are configured as appropriate, for example:
   ```
   Acquire {
     HTTP::proxy "http://proxyuser:Abcd$1234@proxy.example.com:8080/";
     HTTPS::proxy "http://proxyuser:Abcd$1234@proxy.example.com:8080/";
   }
   ```
4. Press Ctrl + O to save your changes.
5. Press Ctrl + X to exit nano.
6. To verify that your configuration file and proxy are working correctly, you can run the following command and check that no errors are reported:
   ```
   sudo apt-get update
   ```

Rerunning the install wizard

Many of the maintenance tasks to change the configuration of the Reverse Proxy and TURN Server involve rerunning the installation wizard. To do this:

1. Log in to the Reverse Proxy and TURN Server through SSH or VMware console as user 'pexip'.
2. At the command prompt, run the `installwizard` command.
3. At each step the default values use the answers from the previous run (if they are still valid).
   Press ENTER to accept the default value, or enter the value you want to use instead.
4. When all of the installation wizard steps have been completed, the appliance automatically reboots.

Adding or removing Conferencing Nodes from an existing reverse proxy or TURN server configuration

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

1. Log in to the reverse proxy/TURN server through SSH or VMware console as user 'pexip'.
2. At the command prompt, run the `installwizard` command.
3. Go through the installation steps accepting the default answers (to preserve the previous settings) until you get to step 5 **Web reverse proxy** (step 7 if you have dual NICs).
4. If you are using the reverse proxy functionality, answer "yes" when asked "Enable web reverse proxy?", and then:
   a. The existing signaling Conferencing Node addresses are listed and you are asked "Use these default values?".

      Reply "no".
   b. You are asked "IP Address of signaling Conferencing Node 1?"

      Re-enter all of the IP addresses of your new set of Conferencing Nodes i.e. from the existing list re-enter the addresses of the nodes you want to keep, add the addresses of any additional nodes, and do not re-enter the addresses of any nodes you want to remove from the list.

      Enter each address one at a time, pressing ENTER after each address, and then add an empty line when finished.
5. At step 6 **TURN server** (step 8 if you have dual NICs), if you are using the TURN server functionality answer "yes" when asked "Enable TURN server?", and then:
   a. The existing media Conferencing Node addresses are listed and you are asked "Use these default values?".

      Reply "no".
   b. You are asked "IP Address of media Conferencing Node 1?"

      Re-enter all of the IP addresses of your new set of Conferencing Nodes i.e. from the existing list re-enter the addresses of the nodes you want to keep, add the addresses of any additional nodes, and do not re-enter the addresses of any nodes you want to remove from the list.

      Enter each address one at a time, pressing ENTER after each address, and then add an empty line when finished.
6. Complete the remaining steps in the install wizard accepting the default answers (to preserve the previous settings).
7. When all of the installation wizard steps have been completed, the appliance automatically reboots.

Configuring the TURN server for TCP TURN relay

In normal circumstances, WebRTC clients connecting to a Pexip environment use HTTPS (TCP port 443) for call signaling, while sending RTP media (audio+video+presentation) via UDP (Chrome and Firefox also supports TCP media as a fallback). RTP media generally uses ephemeral highports — Conferencing Nodes use ports 40000-49999 by default and Pexip-provided TURN server uses UDP 49152-65535 for media.

These ephemeral highport ranges can cause a problem if the WebRTC client is behind a strict firewall setup. Strict firewalls tend to limit the outbound connections allowed from a client PC to well-known ports and/or protocols, for instance only allowing TCP/80 (HTTP), TCP/443 (HTTPS) and UDP/53 (DNS). In this case, for example, when the client attempts to send RTP media to a Pexip node on port 40120, these RTP packets will be blocked by the firewall, and thus the client cannot receive or send any media.

A workaround for this type of issue is to set up a TURN server specifically for such clients, where this TURN server listens for incoming TURN traffic on a well-known port (where this well-known port is likely to be allowed on the strict firewall). TCP port 443 is an example of such a well-known port, where it is likely, but not guaranteed, that the WebRTC client will be able to communicate with an arbitrary host located on the Internet over this port.

Note that sending TURN media over TCP 443 is not the same as encapsulating RTP inside of HTTP(S). In this case, we are simply using a well-known port (TCP 443) to send RTP media, where this well-known port in normal circumstances is used to send HTTPS traffic. This means that if the WebRTC client is located behind a web proxy (where all traffic between the client and the Internet is funneled via this web proxy), this approach is not likely to work since the web proxy will recognize that this RTP media is not actual HTTPS traffic and therefore drop the packets. This approach will therefore only work with "dumb" firewalls which only evaluate the port and network protocol for the traffic in question, and where the firewall disregards the application protocol that is being used across this network connection.

The following diagram shows a WebRTC client (running Chrome or Firefox), located behind a strict firewall, as well as a TURN server and Pexip Conferencing Nodes which are located on a public network. This strict firewall only allows a certain set of ports for outbound communication from Chrome client, and in this example, TCP/443 is one of the outbound destination ports that are allowed. When configured correctly, the client will first establish WebRTC signaling to the Conferencing Nodes (signaling is not shown in this diagram), and then receive TURN server provisioning information from the Conferencing Node which instructs the client to use STUN/TURN with this TURN server.

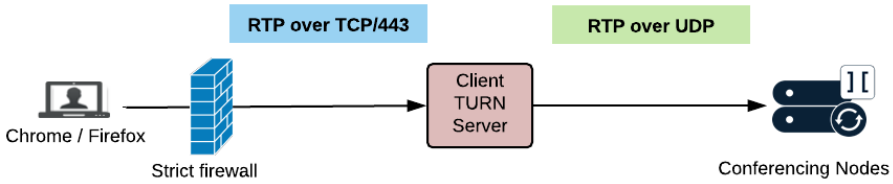

After the WebRTC client has established a relationship with the TURN server, the client can send RTP media to the TURN server over a single TCP port, and the TURN server will forward (relay) these RTP packets to the appropriate Conferencing Node over UDP as normal. Similarly, the Conferencing Node can send RTP media back to the WebRTC client via the same connection.

Note that:

- When using TCP/443 as the single TURN port, the client TURN server cannot coexist with a reverse proxy application as the reverse proxy also uses TCP/443. Therefore you should set up a dedicated VM instance for the client TURN role. The installation wizard does not allow you to enable both the reverse proxy and TURN over TCP/443 on the same appliance.
- In versions prior to 6.1.0 of the TURN server, clients may still use UDP media relays, if those media paths can be established (via ICE negotiation). As of version 6.1.0, clients are unable to use media relays over UDP/3478 as their addresses will not be in the safelist of allowed IP addresses for UDP media relay that is configured via the installation wizard.
- As of version 6.1.0 of the TURN server, clients can relay media only to the safelist of allowed IP addresses for UDP media relay (as communication between the TURN server and Conferencing Nodes always uses UDP/3478).
- As general good practice, we always recommend deploying the TURN server in a suitably secured network segment, such as a DMZ.
- In some deployments, the Conferencing Nodes are themselves behind dynamic NAT, which means that these Conferencing Nodes also require a TURN server to exchange RTP with external hosts. In this case we recommend that you deploy a separate TURN server (where TURN over TCP/443 is not enabled) for use by these Conferencing Nodes.

To deploy TCP TURN relay:

1. Deploy a TURN server for TURN over TCP/443:
   a. Start deploying the Pexip RP/TURN OVA template as usual.
   b. Follow the installation wizard and ensure that you make the following responses when asked:

Prompt	Response
Enable web reverse proxy?	no
Enable TURN server?	yes
Do you want the TURN server to listen on port 443 instead of 3478?	yes

2. Configure Pexip Infinity to provision Infinity Connect web app clients with instructions to use this TURN server:
   a. Use the Pexip branding portal to build your customized web app:
      i. Go to the Pexip branding portal (https://brandingportal.pexip.com).
      ii. Create a new **Customization** (or edit an existing customization).
      iii. On the **Application Settings** page, in the **Turn Servers** field, enter a JSON object that specifies the TURN server address and credentials, for example:

```
{ "url": "turn:turn.example.com:443?transport=tcp", "username": "user", "credential": "pass" }
```

      where you replace **turn.example.com** with the address of your client TURN server, and you replace **user** and **pass** with the correct credentials for your client TURN server. If the TURN server has dual NIC this should be the external-facing address.

Note that these credentials are not encrypted within the settings file. However, TCP relay is disabled, and the TURN server can only relay to safelisted addresses.

iv. When you have finished configuring your branding, go to the Dashboard, select the relevant Customizations (and any other elements if you are customizing other aspects of the web app) and **Build** your customization package.

This creates and downloads a **branding.zip** file containing your client customizations.

b. Upload the branding package to your Management Node:

i. Go to **Services > Web App Customization**.

ii. In the **Upload Web App branding** section, select **Choose File** and select the ZIP file containing your customizations.

iii. Select **Upload branding**.

The branding package will be uploaded. The upload process automatically detects which type of app branding is contained in the ZIP file and processes it accordingly.

Wait for the new branding to be replicated out to all Conferencing Nodes (typically after approximately one minute).

Now, when an Infinity Connect web app client connects to a Conferencing Node, it will be provisioned with the address of the client TURN server to use for its media routing.

Configuring NAT for the TURN server

To configure NAT for the TURN server, 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 systemctl restart coturn
```

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 between single NIC and dual NIC

To change the configuration of your appliance from single NIC to dual NIC (or vice versa) you should rerun the installation wizard:

1. Log in to the reverse proxy/TURN server through SSH or VMware console as user 'pexip'.
2. At the command prompt, run the `installwizard` command.
3. At step 1 of the install wizard it will detect how many NICs are available and present the appropriate options as described in Running the installation wizard.

   Answer the network interfaces questions and include any custom network routes as appropriate.
4. Complete the remaining steps in the install wizard accepting the default answers (to preserve the previous settings).
5. When all of the installation wizard steps have been completed, the appliance automatically reboots.

Changing the TURN server's access credentials

To set new access credentials for the TURN server, rerun the installation wizard:

1. Log in to the reverse proxy/TURN server through SSH or VMware console as user 'pexip'.
2. At the command prompt, run the `installwizard` command.
3. Go through the installation steps accepting the default answers (to preserve the previous settings) until you get to step 6 **TURN server** (step 8 if you have dual NICs).
4. Keep your default answers for the "Enable TURN server?" and "Do you want the TURN server to listen on port 443 instead of 3478?" prompts.
5. Specify your new TURN server credentials for the "Username?" and "Password?" prompts.

6. Complete the remaining steps in the install wizard accepting the default answers (to preserve the previous settings).
7. When all of the installation wizard steps have been completed, the appliance automatically reboots.

Note that the realm in use in your deployment is set to the value of Domain as specified in that installation wizard. You can also check this 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

To reset the appliance to its default state you must either redeploy the appliance VM from the original OVA file, or restore a snapshot taken after initially deploying the OVA template with its default values.

Reverse Proxy and TURN Server release notes

This topic provides a summary of the new features and fixed issues in each of the current and previous releases of the Pexip Reverse Proxy and TURN Server.

Version 6.1.2

Version 6.1.2 of the OVA template was released in February 2021 and includes the following changes:

- The following Microsoft and Office URLs have been safelisted on the reverse proxy:
  - **https://*.microsoft.com**
  - **https://*.office.com**

  This is to address CSP (Content-Security-Policy) issues when using the reverse proxy in conjunction with VMR Scheduling for Exchange.
- Addresses a security vulnerability (CVE-2020-26262) where Coturn allowed a malicious user to relay packets to the loopback interface.
- Any URLs in the format **https://<reverseproxy>/<alias>/<participant>** now get rewritten to the new recommended Pexip Infinity URL structure: **https://<conferencingnode>/webapp/?conference=<alias>&name=<participant>** when they are forwarded to a Conferencing Node.
- Any conference PINs are excluded from the nginx logs.

This release of the reverse proxy is only compatible with Pexip Infinity version 25 and later.

Version 6.1.1

Version 6.1.1 of the OVA template was released in July 2020 and addresses a security vulnerability (CVE-2020-4067) where Coturn does not initialize the STUN/TURN response buffer properly.

Note that if you are currently running version 6.1.0, this vulnerability would also be addressed by following Pexip's routine maintenance advice for patching the operating system for the latest security bugs.

Version 6.1.0

Version 6.1.0 of the OVA template was released in April 2020 and includes the following changes:

- There is a new step in the installation wizard to specify the IP addresses of the Conferencing Nodes that can use the TURN server for media relay. This addresses a security vulnerability (CVE-2020-11805) that allows an unauthenticated remote attacker to send arbitrary UDP traffic to destinations on the same network segment as the Pexip Reverse Proxy and TURN Server, by locking down the IP addresses that the TURN server is allowed (safelisted) to relay to over UDP. As general good practice, we always recommend deploying the TURN server in a suitably secured network segment, such as a DMZ.

Other improvements include:

- Resolved a minor file permission vulnerability for the default generated SSL certificate.
- Updated SSL ciphers, disabled TLS1.1 and enabled TLS1.3.
- Allows presentation images greater than 1 MB to be uploaded (up to a limit of 10 MB).

Version 6.0.10

Version 6.0.10 of the OVA template was released in May 2019 and includes the following changes:

- Resolved coTURN, TURN and STUN server security vulnerabilities CVE-2018-4056, CVE-2018-4058 and CVE-2018-4059.
- Improved logging:

  - Sensitive parameters such as client API conference tokens are now redacted in log output.
  - Better log rotation to prevent leaving nginx log files that are too large.
  - Ensure that log messages are written to the correct nginx log file.
- Prevent unwanted dialogs appearing when using `apt-get` to upgrade the `libpam-systemd` and `grub-pc` Ubuntu packages.
- Now also available as an Amazon Machine Image (AMI) on Amazon Web Services (AWS).

Version 6.0.7

Version 6.0.7 of the OVA template was released in November 2018 and includes the following changes:

- Enhanced installation wizard that contains options to:
  - select whether either the reverse proxy, the TURN server, or both applications are enabled
  - configure dual network interfaces
  - configure the TURN server for TCP relay via port 443
  - enable fail2ban
  - enable SNMPv2c
  - offer previous settings as the default when rerunning the wizard.
- Security improvements:
  - Updated underlying Ubuntu OS to ensure ongoing support for future security patches.
  - HSTS (HTTP Strict Transport Security) has been enabled. This means that if your deployment moves from using a valid TLS certificate to using an invalid certificate (e.g. you redeploy the reverse proxy, or your certificate expires or is invalidated for some reason) then certain web browsers will stop you from accessing the appliance via the web when using its DNS name, until you correct the certificate issue.
  - TLS v1.0 is disabled for HTTPS inbound connections.

There is no migration path from previous versions to version 6. You must uninstall the previous appliance and then perform a fresh install of version 6 using the same network and Conferencing Node addresses etc.