



Pexip Infinity

Hardware Resource Allocation Guide

Software Version 33

Document Version 33.a

October 2023

] pexip[

Contents

Introduction	3
Hardware resource allocation rules	4
Resource allocation examples	6
Virtual Meeting Rooms	6
Single non-distributed VMR	6
Single distributed VMR (standards-based, WebRTC and SfB endpoints)	7
Distributed VMR — 3 nodes	7
Gateway calls to Microsoft Teams	8
Single non-distributed gateway call (standards-based endpoint to Teams)	8
Single gateway call via Proxying Edge Node (standards-based endpoint to Teams)	8
Multiple distributed gateway calls via Proxying Edge Node (mixed endpoints to same Teams meeting)	8
Gateway calls to Google Meet	9
Single non-distributed gateway call (standards-based endpoint to Google Meet)	9
Single distributed gateway call (standards-based endpoint to Google Meet)	10
Multiple distributed gateway calls (mixed endpoints to Google Meet)	10

Introduction

This guide describes how hardware resources are allocated in a Pexip Infinity deployment. It includes the following sections:

- [Hardware resource allocation rules](#): how hardware resources are allocated and consumed in a Pexip Infinity deployment.
- [Resource allocation examples](#): how resources are used in a number of specific scenarios, such as distributed and non-distributed conferences, gateway calls, calls using standards-based endpoints, Skype for Business / Lync clients, or a combination of clients.

Hardware resource allocation rules

A number of different types of connections to Transcoding Conferencing Nodes are required for a conference to take place, or for a gateway call to be made.

A connection can be a call or presentation from an endpoint to a Virtual Meeting Room or Virtual Auditorium, a backplane between Transcoding Conferencing Nodes, or a call into or out of the Infinity Gateway. In this context, a **connection** is analogous to a **port**. Note that a Skype for Business client may require two connections (one for the video call, and one for presentation content).

When a connection is proxied via a Proxying Edge Node, the proxying node also consumes connection resources in order to forward the media streams on to a Transcoding Conferencing Node. A transcoding node always consumes the same amount of connection resources regardless of whether it has a direct connection to an endpoint, or it is receiving the media streams via a proxying node.

Each connection to a Transcoding Conferencing Node requires a certain amount of resource, or capacity.

In general, when compared to a single high definition **HD 720p** call:

- a **Full HD 1080p** call uses twice the resource
- an **SD** standard definition call uses half the resource
- an **audio-only** call uses one sixteenth of the resource.

However, note that:

- A WebRTC call using the **VP8** codec uses the same amount of resource as H.264, and the **VP9** codec uses around 25% more resource, so VP9 at 720p uses the equivalent of 1.25 HD resources, and VP9 at 1080p uses the equivalent of 2.5 HD resources. Note that within the same conference some participants may use VP9 (if they are connected to a Conferencing Node using the AVX2 or later instruction set) while other participants may use VP8 (if they are connected to a Conferencing Node on older hardware).
- Conferences or gateway calls that use the **Adaptive Composition** layout, and Teams conferences that use the Teams-like layout, consume additional Conferencing Node resources. The actual amount of additional resource depends on many factors, but as a guide, it uses an additional 1 HD of resource per conference, for up to 3 other video participants, plus approximately another 0.5 HD for each additional (4th, 5th etc.) video participant that is on stage. This is regardless of the call quality / resolution of the conference itself and each individual participant's connection (codec, bandwidth and so on).
- H.323 audio-only calls are treated the same as video calls for resource usage purposes.
- Connections to a Media Playback Service use 1.2 times as much resource as a connection to a VMR.
- When transferring a participant, the transferee can temporarily take two sets of resources while in the process of being transferred. This should not normally last more than a few seconds.

If you want to limit video calls to specific resolutions (and limit the transcoding node resources that are reserved for calls), you should use the **Maximum call quality** setting.

The following rules determine how hardware resources are allocated and consumed by conference and gateway calls in a Pexip Infinity deployment. They are measured in terms of HD resources, and assume that the **Maximum call quality** of the conference is set to HD (more or less resources are used for Full HD and SD respectively, as described above):

- **Standards-based endpoints** (SIP or H.323) participants:
 - Each participant uses 1 HD resource.
 - Extra resources are consumed when sending main video and presentation content to the endpoint, than when compared to only sending main video. The amount of additional resource depends on many factors, including the resolution and frame rate of both sets of content, and the available bandwidth (plus the general maximum settings for call quality, bandwidth and presentation bandwidth ratio). Thus, if for example, a Conferencing Node is using 1 HD of resource to send just main video, and it then starts also sending presentation content, it may typically use a total of 1 HD of resource for presentation plus 0.5 HD for main video. However, in other circumstances (in particular, the bandwidth configured for the call), it could use 1 HD for content plus 1 HD for video, or potentially just 0.5 HD for content and 0.5 HD for video.
 - When sending two video streams to a dual-screen endpoint, the call bandwidth is always split 50-50 between video and presentation content (and cannot be changed), and it uses twice the resource as when only sending main video.
 - No additional resource is consumed when receiving presentation content from the endpoint.
- **WebRTC VP8/VP9** participants:

- VP8 participants use 1 HD resource and VP9 participants use 1.25 HD resources for main video.
- Extra resources are consumed when sending main video and presentation content to the WebRTC endpoint, in the same manner as described above for standards-based endpoints.
- No additional resource is consumed when receiving presentation content from the endpoint.
- **Skype for Business / Lync participants:**
 - Each participant uses 1 HD resource.
 - They each require 1 additional HD resource when sending or receiving presentation.
- If an endpoint is restricted due to bandwidth limitations to a lower resolution than the conference's maximum call quality, the transcoding node will use the appropriate lower level of resource for that endpoint's connection. For example, if it is limited to an SD connection it will use 0.5 HD resources.
- Each conference instance on each Transcoding Conferencing Node reserves a backplane connection at a resource level corresponding to the conference's **Maximum call quality** setting, to allow the conference to become geographically distributed if required. The exceptions to this are:
 - Deployments with a single Conferencing Node. In such cases, no backplanes will ever be required, so capacity is not reserved.
 - Conferences that are audio-only (in other words, where the conference has its **Conference capabilities** set to **Audio-only**). In such cases, capacity equivalent to one audio connection is reserved for the backplane.
- Only one backplane connection is used for each conference on each Transcoding Conferencing Node, regardless of the number of other transcoding nodes that are involved in the conference. Note that on the Administrator interface, backplanes are displayed as unidirectional, so a single link between two nodes is shown with two backplanes, one in each direction. Likewise, a conference between three nodes is shown with six backplanes. However, for resource allocation purposes each node would require just one backplane connection.
- Pexip Infinity always tries to optimize **gateway calls**:
 - A gateway call does not reserve resource for a backplane, but will use one if required (for example, if the participants are connected via different Transcoding Conferencing Nodes).
 - For a gateway call to **Google Meet**, the connection to Google Meet always uses 1 HD resource (it uses VP8) for main video. The resources required for the VTC leg of the connection depend upon the type of endpoint and the **Maximum call quality** setting. If the VTC endpoint starts to present content then an extra 1 HD resource is used for the connection from Pexip Infinity to Google Meet. However, no additional resources are required on the Google Meet leg if presentation content is sent from Google Meet, but **0.5 HD of additional resource** would typically be required for each endpoint receiving presentation.
 - For a gateway call to a **Microsoft Teams** meeting, the connection to Teams uses 1.5 HD of resource if **Maximum call quality** is SD or HD, otherwise it uses 1.5 Full HD resources. The resources required for the VTC leg of the connection depend upon the **Maximum call quality** setting. If any participant presents content, additional resources (typically **0.5 HD**) would be required, either on the Teams backplane (when an endpoint presents) or on the node handling the endpoint's media connection (when a Teams client presents). The exact amount of resource used depends on the codec, resolution and frame rate of the presentation stream.
- If an API participant is the first participant to join a conference, it will reserve a backplane for the conference.

Proxying Edge Node resource requirements

When a connection is proxied via a Proxying Edge Node, the proxying node also consumes connection resources in order to forward the media streams on to a Transcoding Conferencing Node.

A proxying node uses approximately the equivalent of 3 audio-only resources to proxy a video call (of any resolution), and 1 audio-only resource to proxy an audio call.

Resource allocation examples

The examples below are designed to give you an idea of how the [Hardware resource allocation rules](#) apply as the capabilities of the Pexip Infinity platform are extended in the various scenarios:

- [Virtual Meeting Rooms](#)
- [Gateway calls to Microsoft Teams](#)
- [Gateway calls to Google Meet](#)

In each of these cases, the hardware capacity requirements are described in terms of the HD **resources** required for a connection. A connection can be a call or presentation from an endpoint to a Virtual Meeting Room or Virtual Auditorium, a backplane between Transcoding Conferencing Nodes, or a call into or out of the Infinity Gateway. In this context, a **connection** is analogous to a **port**. Note that a Skype for Business client may require two connections (one for the video call, and one for presentation content).

All of these examples are based around using the default **Maximum call quality of HD**. A **Full HD** connection typically uses twice the resource of a HD connection and a standard definition SD connection typically uses half the resource. WebRTC connections are assumed to be using the VP9 codec. VP9 calls consume around 1.25 times the resource for calls for main video (i.e. 1.25 HD resources for 720p, and 2.5 HD for 1080p).

All of the examples assume that the endpoints are connecting directly to Transcoding Conferencing Nodes. However, the Teams examples also show the effect of routing those calls via a Proxying Edge Node to the Teams Connector. (When a connection is proxied via a Proxying Edge Node, the proxying node also consumes connection resources in order to forward the media streams on to a Transcoding Conferencing Node. A transcoding node always consumes the same amount of connection resources regardless of whether it has a direct connection to an endpoint, or it is receiving the media streams via a proxying node.)

Virtual Meeting Rooms

Single non-distributed VMR


- There are multiple Conferencing Nodes in a single system location, but the conference is hosted on a single node.
- There are three endpoints (two standards-based and one WebRTC VP9*) connecting directly to the node.
- Each endpoint uses HD video and audio.

Pexip Infinity **requires 4.25 HD resources**. 3.25 HD resources are required by the endpoints (the two standards-based endpoints each use 1 HD resource and the WebRTC VP9 client uses 1.25 HD), and 1 HD resource is reserved by the backplane.

- A presentation is then sent from a standards-based endpoint (see diagram).

Pexip Infinity then typically uses an **extra 0.5 HD resource** for each standards-based or WebRTC endpoint that is receiving the presentation, meaning the **total resources used on Node 1 is now 5.25 HD**.

The conference uses 3 concurrent call licenses.



Single distributed VMR (standards-based, WebRTC and SfB endpoints)

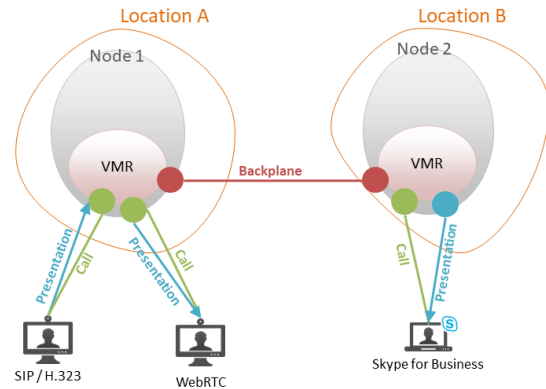
- There are two Conferencing Nodes in two different system locations.
- One standards-based endpoint and one WebRTC VP9* endpoint connect to the VMR in the Location A node.
- One SfB endpoint connects to the VMR in the Location B node.
- Each endpoint uses HD video and audio.

Pexip Infinity **requires 3.25 HD resources** on Node 1 (2.25 for the connected participants, and 1 for the backplane), **and 2 HD resources** on Node 2, (1 for the connected participant, and 1 for the backplane).

- A presentation is then sent from the standards-based endpoint (see diagram).

Pexip Infinity uses an **extra 0.5 HD resource** on Node 1 for the WebRTC client to receive presentation, meaning the **total resources used on Node 1 is now 3.75 HD**. And it **requires 1 additional HD resource** on Node 2 for the SfB client to receive presentation, so the **total resources used on Node 2 is now 3 HD**.

The conference uses 3 concurrent call licenses.



Distributed VMR — 3 nodes


- There are three Conferencing Nodes in three different system locations.
- One WebRTC VP9* endpoint connects to the VMR in Location A.
- One SfB endpoint connects to the VMR in Location B.
- One standards-based endpoint connects to the VMR in Location C.
- Each endpoint uses HD video and audio.

Pexip Infinity **requires 2.25 HD resources** on node 1 (1.25 for the WebRTC participant, and 1 for the backplane) and **2 HD resources** on nodes 2 and 3 (1 for the connected participants, and 1 for the backplane).

- The WebRTC participant then shares their screen (see diagram).

Pexip Infinity **requires no additional resources** on Node 1 for the WebRTC endpoint to send presentation, but does require **1 additional HD resource** on Node 2 for the SfB client to receive presentation, and **0.5 HD additional HD resource** on Node 3 for the endpoint to receive presentation, so the **total resources used on Node 1 is still 2.25 HD, Node 2 is now 3 HD and Node 3 is now 2.5 HD**.

The conference uses 3 concurrent call licenses.




Gateway calls to Microsoft Teams

Single non-distributed gateway call (standards-based endpoint to Teams)

- We have one standards-based endpoint connecting via a single transcoding node to a Microsoft Teams meeting via the Teams Connector.
- Two Teams clients are also connected to the meeting.
- Each endpoint uses HD video and audio.

Pexip Infinity **requires 2.5 HD resources**. The standards-based endpoint uses 1 HD resource, and the backplane connection to the Teams Connector uses 1.5 HD resources.



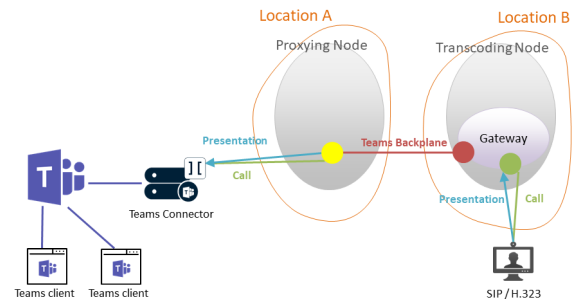
- A presentation is then sent from the standards-based endpoint (see diagram).

Pexip Infinity would typically use an **extra 0.5 HD resource** on the Teams backplane connection for the Microsoft Teams meeting to receive presentation. If one of the Teams clients sent presentation content it would instead typically use an **extra 0.5 HD resource** for the standards-based endpoint to receive presentation.

The call uses 2 concurrent call licenses in both cases.

Single gateway call via Proxying Edge Node (standards-based endpoint to Teams)

- There is a proxying node in location A and a transcoding node in location B.
- A gateway call is placed from a standards-based endpoint connected to the transcoding node to a Microsoft Teams meeting. The call is routed via the proxying node and the Teams Connector.
- Two Teams clients are also connected to the meeting.
- Each endpoint uses HD video and audio.



Pexip Infinity **requires 0.2 HD resources** on the proxying node, and **2.5 HD resources** on the transcoding node (1 for the standards-based endpoint and 1.5 for the backplane connection).

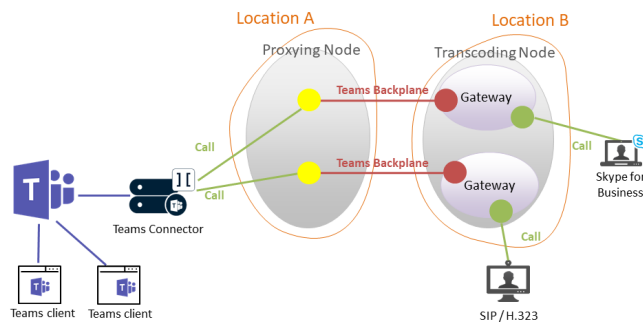
- A presentation is then sent from the standards-based endpoint (see diagram).

Pexip Infinity would typically use an **extra 0.5 HD resource** on the Teams backplane connection for the Microsoft Teams meeting to receive presentation. If one of the Teams clients sent presentation content it would instead typically use an **extra 0.5 HD resource** for the standards-based endpoint to receive presentation.

The call uses 2 concurrent call licenses in both cases.

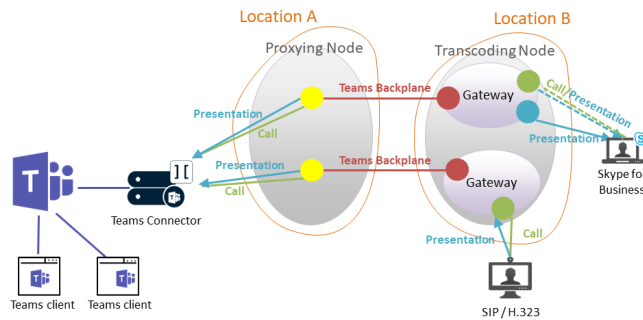
Multiple distributed gateway calls via Proxying Edge Node (mixed endpoints to same Teams meeting)

- There is a proxying node in location A and a transcoding node in location B.
- Two gateway calls are placed from endpoints connected to the transcoding node to a Microsoft Teams meeting. One endpoint is a standards-based endpoint and the other is a Skype for Business client. Both calls are routed via the proxying node and the Teams Connector.
- Two Teams clients are also connected to the Microsoft Teams meeting.
- Each endpoint uses HD video and audio.



Pexip Infinity **requires 0.4 HD resources** on the proxying node and **5 HD resources** on the transcoding node (the standards-based endpoint requires 1 for its connection to the gateway, and 1.5 for the backplane; the Skype for Business client requires 1 for its connection to the gateway, and 1.5 for the backplane) (see diagram above).

- A presentation is then sent from the standards-based endpoint (see diagram below).



Pexip Infinity would typically use an **extra 0.5 HD resource** on the Teams backplane connection for the Microsoft Teams meeting to receive presentation. The Skype for Business client **requires 1 additional HD resource** on the transcoding node to receive presentation.

* WebRTC connections are assumed to be using the VP9 codec. VP9 calls consume around 1.25 times the resource for calls for main video (i.e. 1.25 HD resources for 720p, and 2.5 HD for 1080p).

Gateway calls to Google Meet

Single non-distributed gateway call (standards-based endpoint to Google Meet)


- We have one standards-based endpoint connecting via a single node in a single location to a Google Meet meeting.
- Two Google Meet clients are also connected to the meeting.
- Each endpoint uses HD video and audio.

Pexip Infinity **requires 2 HD resources**. The standards-based endpoint uses 1 HD resource. The connection to Google Meet uses 1 HD resource (it uses VP8).

- A presentation is then sent from the standards-based endpoint (see diagram).

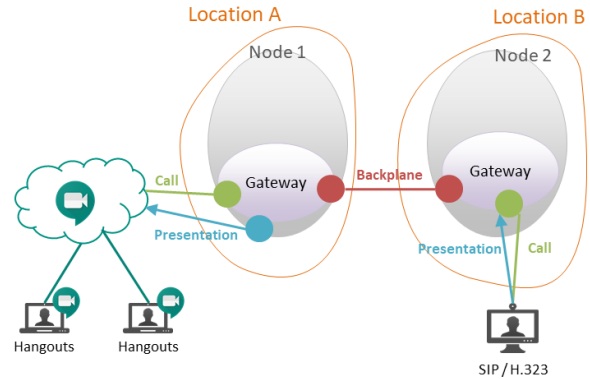
Pexip Infinity **requires an additional 1 HD resource** for Google Meet to receive presentation. However, no additional resources are required on the Google Meet leg if presentation content is sent from Google Meet, but **0.5 HD of additional resource** would typically be required for each endpoint receiving presentation.

The call uses 2 concurrent call licenses in both cases.



Single distributed gateway call (standards-based endpoint to Google Meet)

- There are two Conferencing Nodes in two different system locations.
- A gateway call is placed from a standards-based endpoint connected via node 2 to a Google Meet meeting connected via node 1.
- Two Google Meet clients are also connected to the meeting.
- Each endpoint uses HD video and audio.



Pexip Infinity **requires 2 HD resources** on node 1 (1 for the connection to Google Meet and 1 for the backplane), and **2 HD resources** on node 2 (1 for the standards-based endpoint and 1 for the backplane), so **4 HD resources in total**.


- A presentation is then sent from the standards-based endpoint (see diagram).

Pexip Infinity **requires an additional 1 HD resource** for Google Meet to receive presentation. However, no additional resources are required on the Google Meet leg if presentation content is sent from Google Meet, but **0.5 HD of additional resource** would typically be required for each endpoint receiving presentation.

The call uses 2 concurrent call licenses in both cases.

Multiple distributed gateway calls (mixed endpoints to Google Meet)

- There are two Conferencing Nodes in two different system locations.
- Two gateway calls are placed to the same Google Meet meeting connected via node 1: one from a standards-based endpoint connected via node 2, and one from a Skype for Business client also connected via node 2.
- Two Google Meet clients are also connected to the meeting.
- Each endpoint uses HD video and audio.



Pexip Infinity **requires 4 HD resources** on node 1 (each of the 2 gateway calls requires 1 HD for the connection to Google Meet and 1 for the backplane). It also requires **4 HD resources** on node 2 (the standards-based endpoint requires 1 for its connection to the gateway, and 1 for the gateway's backplane; the Skype for Business client requires 1 for its connection to the gateway, and 1 for the gateway's backplane), so **8 HD resources in total** (see diagram above).

- A presentation is then sent from the standards-based endpoint (see diagram below).

Pexip Infinity **requires an additional 1 HD resource** on node 1 to send the presentation to Google Meet, but does not require additional resources on the other connection that is receiving the presentation from Google Meet. The Skype for Business client **requires 1 additional HD resource** on node 2 to receive presentation, so **10 HD resources in total** across the two nodes.

