



# Pexip Infinity

## Server Design Guide

## Introduction

This document describes the recommended specifications and deployment for servers hosting the Pexip Infinity platform. It starts with a [Summary of recommendations](#), which are supplemented by further details and explanations in the following Appendices:

- [Appendix 1: Detailed hardware requirements](#) provides a more detailed breakdown of the minimum and recommended hardware requirements for servers hosting the Management Node and Conferencing Nodes respectively.
- [Appendix 2: Host server components](#) provides a diagram and descriptions of the different components of the host server (e.g. Processor, cores, RAM) when used by Conferencing Nodes.
- [Appendix 3: Achieving high density deployments with NUMA](#) provides details on NUMA architecture and how this impacts server architecture and overall performance of Conferencing Nodes.

## Summary of recommendations

This section summarizes the recommended specifications and deployment guidelines for servers hosting the Pexip Infinity platform.

### Management Node

#### Host server specifications

- 2 cores\* (most modern processors will suffice)
- 4 GB RAM\*
- 100 GB storage
- The Pexip Infinity VMs run on bare-metal hypervisors so no OS should be installed.

\* Sufficient for deployments of up to 30 Conferencing Nodes.

### Conferencing Node

#### Host server specifications

- We recommend Intel Xeon E5-2600/4600 series (Sandy Bridge-EP architecture) processor or later
- 2.7 GHz (or faster) clock speed
- 8, 10 or 12 physical cores per socket
- 1 GB RAM for each core that is allocated to the Conferencing Node
- 4 memory modules per processor socket, with all memory channels populated
- 50 GB storage
- VMware ESXi 5.1 or later
- The Pexip Infinity VMs run on bare-metal hypervisors so no OS should be installed.

## General deployment recommendations

- Prefer processors with high core count (8 cores or more per CPU).
- Prefer processors with a high clock speed (2.7 GHz and higher).
- Prefer a smaller number of large Conferencing Nodes (e.g. 4 x 8-core nodes), rather than large number of small Conferencing Nodes (e.g. 8 x 4-core nodes).
- A single Conferencing Node must not be assigned more vCPU than the amount of physical cores on each processor socket.
- For each physical CPU core:
  - configure 1 vCPU
  - assign at least 1 GB RAMFor example, on an E5-2680v2 CPU with 10 cores per CPU, assign 10 vCPU and 10 GB of RAM.
- Populate memory equally across all NUMA nodes/sockets on a single host server. All memory channels (typically 4 per CPU for E5-2600) must be populated.
- Do not over-commit either RAM or CPU resources on hardware hosts. In other words, the Conferencing Node and Management Node each must have dedicated access to their own RAM and CPU cores. Pexip Conferencing Nodes use real-time media, which needs dedicated capacity.
- Ensure all BIOS settings pertaining to power saving should be set to maximize performance rather than preserve energy. (Setting these to an energy-preserving mode may impact transcoding capacity.) The actual settings will depend on the hardware vendor; some examples are given below:

### Typical HP settings

- HP Power Profile: Maximum Performance
- Power Regulator modes: HP Static High Performance mode
- Energy/Performance Bias: Maximum Performance
- Memory Power Savings Mode: Maximum Performance

### Typical Dell settings

- System Profile: Performance Optimized

## Appendix 1: Detailed hardware requirements

The following table lists the recommended hardware requirements for the Management Node and Conferencing Node host servers.

	Management Node	Conferencing Node
Server manufacturer	Any	Any
Processor make* (see also <a href="#">Performance considerations</a> )	Any	Intel Xeon E5-2600/4600 series (Sandy Bridge-EP architecture) or later
Processor instruction set	Any	AVX2
Processor architecture	64-bit	64-bit
Processor speed	2.0 GHz	2.7 GHz or faster
No. of physical cores*	2**	8 cores per socket
Processor cache	no minimum	20 MB or greater (2.5 MB L3 cache per core)
Total RAM*	4 GB**	1 GB per core
RAM makeup	Any	All channels must be populated with a DIMM, see <a href="#">Memory configuration</a> below. (For example, E5-24xx supports 3 DIMMs per socket, E5-26xx supports 4 DIMMs per socket.)
Hardware allocation	The host server must not be over-committed in terms of either RAM or CPU. In other words, the Management Node and Conferencing Nodes each must have dedicated access to their own RAM and CPU cores.	
Storage space required	100 GB	50 GB
GPU	No specific hardware cards or GPUs are required.	
Network	Gigabit Ethernet connectivity from the host server.	
Operating System	The Pexip Infinity VMs run on bare-metal hypervisors so no OS should be installed.	
Hypervisor (see also <a href="#">Performance considerations</a> )	<ul style="list-style-type: none"> <li>VMware ESXi 4.1 or later</li> <li>Microsoft Hyper-V 2012 or later</li> </ul>	<ul style="list-style-type: none"> <li>VMware ESXi 5.1 or later</li> </ul>

\* This does not include the processor and RAM requirements of the Hypervisor.

\*\* Sufficient for deployments of up to 30 Conferencing Nodes. For larger deployments, you will need to increase the amount of RAM and number of cores. For more information, consult your Pexip support representative.

## Capacity

The number of calls, or ports, that can be achieved per server in a Pexip Infinity deployment will depend on the specifications of the particular server and the bandwidth of each call.

As a general indication of capacity:

- When deployed on our recommended hardware (Intel Sandy Bridge, 8 cores, 2.7 GHz), Pexip Infinity can connect up to two High Definition 720p30 calls per CPU core. This is based on 1.1 GHz per HD call plus 20% headroom. Capacity for higher speeds can be linearly calculated based on these figures.
- The same recommended hardware can connect a higher number of lower-resolution calls per CPU core. For example, up to 20 audio-only AAC-LD calls at 64 kbps.
- Deployments on servers that are older, have slower processors, or have fewer CPUs, will have a lower overall capacity.

## Performance considerations

The type of processors and Hypervisors used in your deployment will impact the levels of performance you can achieve. Some known performance considerations are described below.

### AMD processors

We have observed during internal testing that use of AMD processors results in a reduction of capacity (measured by ports per core) of around 40% when compared to an identically configured Intel platform. This is because current AMD processors do not execute advanced instruction sets at the same speed as Intel processors.

AMD processors older than 2012 may not perform sufficiently and are not recommended for use with the Pexip Infinity platform.

### Microsoft Hyper-V

We have observed during internal testing that use of Microsoft's Hyper-V hypervisor may result in a reduction of performance of up to 50% (as measured in quantity of ports per physical core allocated to the Conferencing Node) when compared to VMware ESXi 5.1. This is due to slower pass through execution of advanced processor instruction sets.

### VMware ESXi 4.1

We have observed during internal testing that use of VMware's ESXi 4.1 hypervisor may result in a reduction of performance of approximately 20% (as measured in number of ports per physical core allocated to the Conferencing Node) when compared to VMware ESXi 5.x. This is due to slower pass through execution of advanced processor instruction sets.

## Memory configuration

Memory must be distributed on the different memory channels (i.e. 4 channels per socket on the E5-2600).


There must be an equal amount of memory per socket, and all sockets must have all memory channels populated (you do not need to populate all slots in a channel, one DIMM per channel is sufficient). Do not, for example, use two large DIMMs rather than four lower-capacity DIMMs - using only two per socket will result in half the memory bandwidth, since the memory interface is designed to read up from four DIMMs at the same time in parallel.

### Example

E5-2600 dual socket system:

- Each socket has 4 channels
- All 4 channels must be populated with a DIMM
- Both sockets must have the same configuration

Therefore for a dual socket E5-2600 you need 8 identical memory DIMMs.



## Appendix 2: Host server components

The diagrams below outline the physical components of a host server, and shows how the components of a Pexip Infinity Conferencing Node virtual machine relate to these.

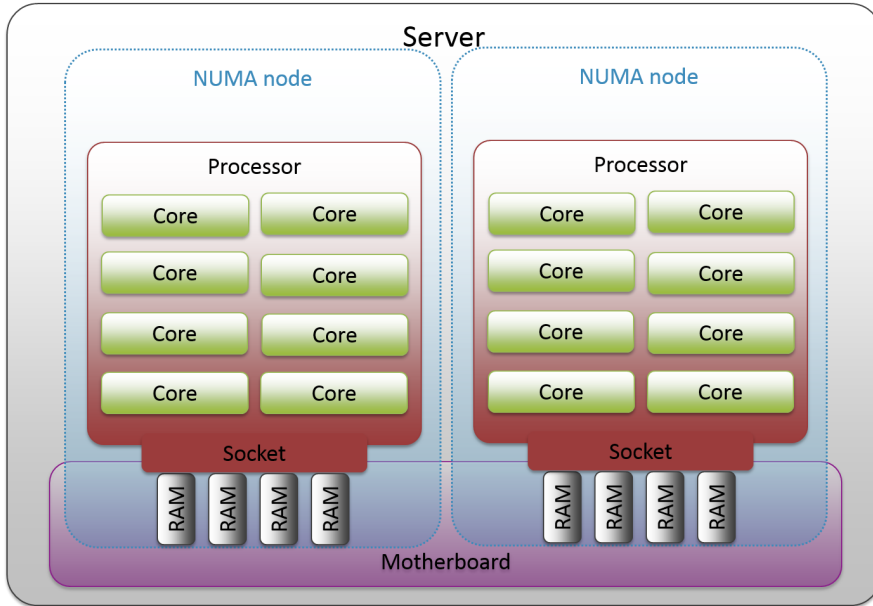


Diagram showing the physical components of a host server

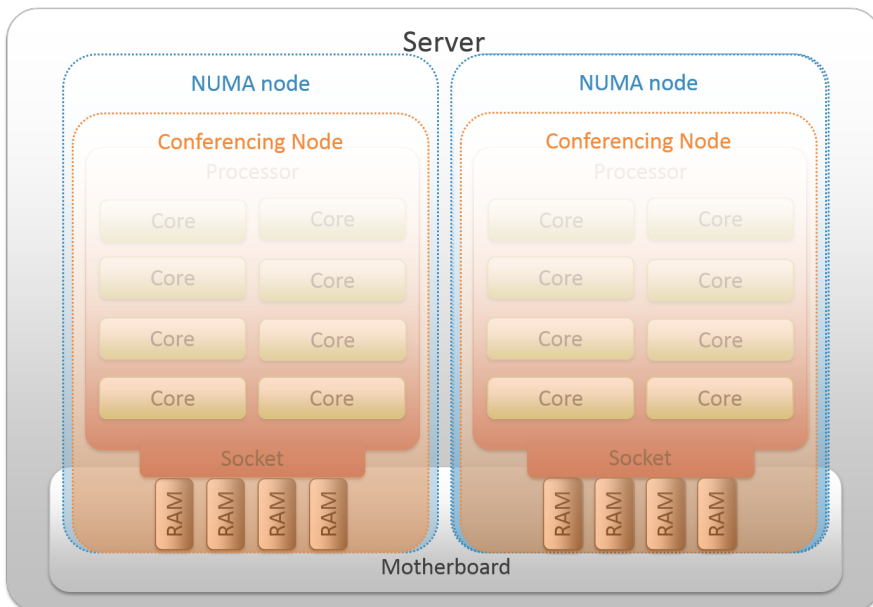


Diagram showing how the Pexip Infinity Conferencing Node VM relates to the host server

The table below provides descriptions for the terms used in the diagrams above and elsewhere in this guide, in the context of a Pexip Infinity deployment.

Term	Description
Processor	The hardware within a computer that carries out the basic computing functions. It can consist of multiple cores.
Core	One single physical processing unit. Intel Xeon E5 typically has 8 cores (10, 12 or more in newer versions)
Socket	The socket on the motherboard where one processor is installed.
RAM	Also referred to as "memory modules". The hardware that stores data which is accessed by the processor core while executing programs.
Virtual CPU (vCPU)	<p>The VM's understanding of how many CPU cores it requires. Each vCPU appears as a single CPU core to the guest operating system.</p> <p>When configuring a Conferencing Node, you are asked to enter the number of virtual CPUs to assign to it. We do not recommend that you assign more virtual CPUs than there are physical cores on a single processor on the host server. For example, if the host server has 2 processors each with 8 physical cores, we recommend that you assign no more than 8 virtual CPUs.</p>
NUMA node	The combination of a processor (consisting of one or more cores) and its associated memory.

## Appendix 3: Achieving high density deployments with NUMA

### Introduction

There are many factors that can affect the performance of Virtual Machines (VMs) running on host hardware. One of these is how the VM interacts with NUMA.

This section provides an overview of NUMA and how it applies to Pexip Infinity Conferencing Nodes. It gives a summary of our recommendations and suggested best practices for maximizing performance, and an explanation of the reasoning behind these.

### About NUMA

NUMA stands for non-uniform memory access. It is an architecture that divides the computer into a number of nodes, each containing one or more processor cores and associated memory. A core can access its local memory faster than it can access the rest of the memory on that machine. In other words, it can access memory allocated to its own NUMA node faster than it can access memory allocated to another NUMA node on the same machine.

### Conferencing Nodes and NUMA nodes

We strongly recommend that a Pexip Infinity Conferencing Node is deployed on a single NUMA node in order to avoid the loss of performance incurred when a core accesses memory outside its own node.

In practice, with modern servers, each socket represents a NUMA node. We therefore recommend that:

- one Pexip Infinity Conferencing Node is deployed per socket of the host server, and
- the number of vCPUs that the Conferencing Node is configured to use is the same as or less than the number of cores available in that socket.

It is possible to deploy smaller Conferencing Nodes over **fewer** cores than are available in a single socket, although this will reduce capacity.

Deploying a Conferencing Node over **more** cores than provided by a single socket will cause loss of performance, as and when remote memory is accessed. This must be taken into account when moving Conferencing Node VMs between host servers with different hardware configuration: if an existing VM is moved to a socket that contains fewer cores than the VM is configured to use, the VM will end up spanning two sockets and therefore NUMA nodes, thus impacting performance.

To prevent this occurring, ensure that either:

- you deploy Conferencing Nodes only on servers with large number of cores per processor
- the number of cores used by each Conferencing Node is the same as (or less than) the number of cores available on each NUMA node of even your smallest hosts.

### VMware and NUMA

As well as the physical restrictions discussed above, the hypervisor can also impose restrictions. VMware provides virtual NUMA nodes on VMs that are configured with more than 8 CPUs. This default value can be altered by setting **numa.vcpu.min** in the VM's configuration file.



## Summary of deployment recommendations

We are constantly optimizing our use of the hardware and expect that some of this advice will change in later releases of our product. However our current recommendations are:

- Prefer processors with a high core count.
- Prefer a smaller number of large Conferencing Nodes rather than a larger number of smaller Conferencing Nodes.
- Deploy one Conferencing Node per NUMA node (i.e. per socket)
- Configure one vCPU per core on that NUMA node.
- Populate memory equally across all NUMA nodes on a single host server.
- Do not over-commit resources on hardware hosts.