

SYSTRAN Pure Neural® Server 9.9

Hardware and Software Requirements

The following requirements apply to SYSTRAN Pure Neural® Server 9.9 (SPNS9).

1 Server software requirements

1.1 Server Supported Platforms

SPNS9 is supported on a “fresh” OS minimal install (BareMetal or virtualized environments):

- Linux RedHat Enterprise 7 64-bit or CentOS 7 64-bit (versions ≥ 7.3)
- Linux RedHat Enterprise 8 64-bit.

1.2 Server Software Dependencies Pre-requisites:

These required 3rd party components are provided and come along the installation kit:

- RabbitMQ required version $\geq 3.9.x$, < 3.10
- Erlang version: supported by RabbitMQ version¹
- MongoDB required version = 4.2.x
- Redis required version =6.0.x for RHEL/Centos 7 and =5.0.x for RHEL 8
- Node JS required version $\geq 14.x$, < 15
- Docker engine version ≥ 17.12
- (As an option) Elasticsearch required version $\geq 7.17.0-1$
- (As an option) Java version supported by Elasticsearch version². Recommended: OpenJDK 11

2 Server sizing

2.1 Configuration examples:

These sizing examples are based on “Collaboration usage”:

	Minimal configurations <small>2nd gen. PNMT (L): global throughput ~200 char/sec</small>		Recommended configurations <small>2nd gen. PNMT (L): global throughput ~800 char/sec³</small>	
	CPU	RAM / Disk	CPU	RAM / Disk
Small system: 1 bidirectional language	16 cores	16 GB RAM 40 GB disk	32 cores	32 GB RAM 60 GB disk
Medium sized system: 5 bidirectional languages		24 GB RAM 80 GB disk		64 GB RAM 150 GB disk
Large system: 10 bidirectional languages		48 GB RAM 150 GB disk		128 GB RAM 300 GB disk

Notes:

- Global throughput corresponds to the cumulated throughput of the system when languages are used in parallel.
- Throughput has been calculated by translating in parallel sets of raw text of 5 000 characters from Wikipedia.

¹ <https://www.rabbitmq.com/which-erlang.html>

² https://www.elastic.co/fr/support/matrix#matrix_jvm

³ For 1 bidirectional language system, global throughput is ~1500 char/sec

2.2 SYSTRAN Server Detailed Requirements

A Language Pair (LP) is composed of two translation engines, one per direction. SPNSg supports Rule Based (RBMT), Statistical (SMT/SPE) and Neural (PNMT) translation engines. PNMT engines are recommended.

2.2.1 Processor

According to the translation engine technology used, performance on CPU has been measured as follows⁴:

- 2nd gen. PNMT: **~200 char/sec** per instance of 8 logical cores AVX2⁵ for size L engines.
(2nd gen. PNMT ~750 char/sec per instance of 8 logical cores AVX2⁶ for size M engines)

Recommended: 1 instance per engine per server.

Other components requirements are from 4 to 16 cores according to usage (number of requests and users).

2.2.2 GPU

PNMT engines can be deployed on GPU servers to provide high performance. GPU deployments require a multiple computing node backend license. You must ensure that the official Nvidia driver is installed and running on your server prior to the installation. GPU performance is ~5 000 char/sec according to GPU generation for M Engines.

2.2.3 RAM

For each instance of a translation engine, the requirements are the following:

- 2nd gen. PNMT engine: **~3GB** per instance for size L engines
(2nd gen. PNMT engine: ~2.5GB per instance for size M engines)

The Filter component, which embeds OCR technology requires 1 Gb per instance. 2 instances are recommended.

Other components including MongoDB, RabbitMQ, Elasticsearch, Redis, SYSTRAN services require between 10 to 20 Gb according to the translation workload.

In case of a virtualized environment, we do not recommend memory overcommitting and NUMA⁷.

2.2.4 Hard Disk

We recommend SSD and a system with extension capabilities. Installation is under /opt and DB data is stored in /var. Provide at least 80% of disk capacity to /opt and /var. Swap size minimum is 10GB.

3 Clients

For accessing web Translation Tools and the User Interface, a latest generation browser with JavaScript enabled is recommended: Up-to-date version of Edge, Chrome, Firefox or Safari.

Client application software requirements:

- Windows 7 (min Windows 7 SP1), Windows 8/8.1 or Windows 10
- Skype/Lync plugin: Skype for business 2015 32-bit / Skype for business 2016 32-bit / Lync 2010, 2013 32-bit
- Legacy Office plugin⁸: 32-bit or 64-bit versions of Microsoft® Office® 2007, 2010, 2013, 2016

Office 365 plugins are supported from SYSTRAN Pure Neural Server 9.7 and require oauth2 feature to be configured.

⁴ Measured on Intel(R) Xeon(R) CPU E5-2620 v4 @ 2.10GHz translating sets of raw text of 5 000 characters. May vary depending on the language.

⁵ <https://software.intel.com/en-us/node/185211>

⁶ <https://software.intel.com/en-us/node/185211>

⁷ <https://docs.mongodb.com/manual/administration/production-notes/>

⁸ whether used with a license or an Office 365 subscription