



# Sizing Service Engines

Avi Technical Reference (v17.1)

Copyright © 2018

# Sizing Service Engines

[view online](#)

Avi Networks publishes *minimum* and *recommended* resource requirements for new Avi Service Engines. These are somewhat abstract numbers though, as network and application traffic may vary dramatically. This guide provides some guidance on sizing. Avi recommends consulting with your local Avi sales engineer for more recommendation that is tailored to the exact requirements.

*Note for readers deploying into AWS: The [Installing Avi Vantage in Amazon Web Services](#) article provides additional sizing guidance unique to the AWS environment, including [this link](#) to Amazon's own guidance on various maxima.*

The below per-CPU-core data are obtained using Intel v4 CPUs.

	1 CPU core
Throughput	4 Gb/s
Connections/s	40k
SSL Throughput	1 Gb/s
SSL TPS (RSA2K)	~1,000
SSL TPS (ECC)	2500

Multiple performance vectors or features may have an impact on performance. For instance, to achieve 1 Gb/s of SSL throughput and 2000 TPS of SSL with EC certificates, Avi recommends two cores.

Avi Service Engines may be configured with as little as 1 vCPU core and 1 GB RAM, or up to 36 vCPU cores and 128 GB RAM. In write access mode, Service Engine resources for newly created SEs may be configured within the SE Group properties. When creating an Avi SE in read or no access modes, Avi SE resources are allocated via the hypervisor or by the size hardware used for bare metal servers.

## CPU

CPU scales very linearly as more cores are added. Avi does not use hyperthreading, as it has minimal improvement in performance. CPU is generally the bottleneck for SSL TPS, SSL throughput and HTTP compression. The default is 2 CPU cores, not reserved.

## Memory

Memory scales near linearly. It is used for concurrent connections and HTTP caching. Doubling the memory will double the ability of the Avi Service Engine to perform these tasks. The default is 2 GB memory, reserved within the hypervisor for VMware clouds. See also [SE Memory Consumption](#).

## PPS

For throughput-related metrics, the hypervisor is likely going to be the bottleneck. Limitations are different for each hypervisor and version. For instance, VMware's vSphere / ESX 5.x supports about 550k packets per second to a virtual

machine such as Avi's Service Engine. ESX version 6 doubles this number to around 1.1M PPS. Bare metal (no hypervisor) numbers depend greatly on the type of NIC used. Avi recommends Intel 82598 and 82599 10-Gb NICs for optimal performance.

## Disk

Avi Service Engines may store logs locally before they are sent to the Avi Controllers for indexing. Increasing the disk will increase the log retention on the SE. SSD is preferred over hard drives, as they can write the log data faster. This is not a requirement. The recommended minimum size for storage is 10 GB, which is also the default.

## Disk Capacity for Logs

Avi Vantage computes the disk capacity it can use for logs based on equations incorporating the following parameters: \* T, the SE's total disk capacity \* N, the number of CPU cores \* M, the main memory (RAM) of the SE \* p, `log_agent_max_storage_ignore_percent`, a parameter in SE runtime properties \* D, the capacity reserved for debug logs \* C (reserved for cores),  $2 * M * \text{MIN}$ , the minimum size allocated for logs \* L, the capacity reserved for client logs

Avi calculates D as follows:

$$D = p * T / 100$$

Initially, Avi sets L, the client log size, as follows:

$$L = T - D$$

Finally, Avi Vantage *adjusts* L and C based on the following conditions:

```
if L > C + MIN, then set L to (L - C)
  else if L > MIN, then set C to L - MIN and then set L to MIN
  else set C to 0
```