

---

Question

[Raymond Lawrence](#) · Oct 23, 2017

## Ensemble/Healthshare server sizing on Amazon AWS EC2 servers

I am designing the software architecture for an Ensemble/Healthshare production to be deployed on Amazon AWS EC2 servers (2 mirrored m4.large - 4 vCPUs / 16 GiB RAM running RedHat Linux 3.10.0-327.el7.x86\_64 and Healthshare for RHEL 64-bit 2016.2.1). It's a rather CPU-intensive production involving massive XSLT 2.0 transformations (massive both in terms of size and volume). I was wondering if anyone has experience configuring Ensemble productions on EC2 servers. My question or concern has to do with the following statement in the Ensemble documentation:

"InterSystems recommends that, as a maximum, you set the Actor Pool Size equal to the number of CPUs in your Ensemble server machine. You could set the number higher, but at any one time there are only as many jobs available as there are CPUs."

Given the presence of hyperthreading, does "number of CPUs" equate to number of vCPUs (4 in our case) or number of virtual cores (2)? I have read some articles in the Developer Community on Caché server sizing in virtualized environments suggesting that hyperthreading should not be taken into account, and for purposes of sizing, consider 1 virtual core to be equivalent to one CPU. However, that article was discussing VMWare, not EC2. I don't know if that should make any difference. I'd appreciate any information/suggestions from people with real-world experience working with production Ensemble/Healthshare healthcare applications deployed on EC2.

Thank you.

Ray Lawrence

Interoperability Team

QuintilesIMS

\*/

[#AWS](#) [#Cloud](#) [#Performance](#) [#Red Hat Enterprise Linux \(RHEL\)](#) [#Ensemble](#) [#HealthShare](#)

---

Source

URL: <https://community.intersystems.com/post/ensemblehealthshare-server-sizing-amazon-aws-ec2-servers>