

CPU Affinity...

April 28th, 2009 | 5 Comments |

I was just reading a discussion on the VMTN community on CPU affinity. The general opinion of the Experts is "Don't use CPU affinity". I fully agree with them, ESX is more than capable to handle the scheduling on it's own with just a limited overhead. And like Ken Cline also stresses it could harm performance because of NUMA load balancing for instance.

Something that's often overlooked though when one does CPU affinity is that people tend to give the VM vCPUs a 1:1 relationship with host cores. In other words a VM with two vCPUs will be pinned down to two cores on the host.

This does make sense doesn't it? No it actually doesn't. There's more to a VM than just it's vCPUs. I would like to refer to page 132 of the Resource Management Guide, aka HA-DRS Bible. In short, besides the vCPUs there are several VM associated threads that need to be scheduled as well. When affinity is set these threads, or worlds as VMware calls them, will be scheduled on the assigned cores. You can imagine that when you use the vCenter client to manage the client these threads (Video / Keyboard / Mouse / CD-Rom etc) will need to be scheduled on the same set of cores as the vCPUs need to be scheduled on... If you have a two vCPU VM and want to use CPU affinity pin it down to at least three cores! Before you start assigning cores to your VM also read the bullet points on page 133 why you shouldn't.

*The CPU affinity setting for a virtual machine applies not only to all of the virtual CPUs associated with the virtual machine, but also to all other threads (also known as "worlds") associated with the virtual machine. Such virtual machine threads perform processing required for emulating mouse, keyboard, screen, CD ROM and miscellaneous legacy devices.*

*In some cases, such as display intensive workloads, significant communication might occur between the virtual CPUs and these other virtual machine threads. Performance might degrade if the virtual machine's affinity setting prevents these additional threads from being scheduled concurrently with the virtual machine's virtual CPUs (for example, a uniprocessor virtual machine with affinity to a single CPU, or a two way SMP virtual machine with affinity to only two CPUs).*

*For the best performance, when manual affinity settings are used, VMware recommends that you include at least one additional physical CPU in the affinity setting in order to allow at least one of the virtual machine's threads to be scheduled at the same time as its virtual CPUs (for example, a uniprocessor virtual machine with affinity to at least two CPUs or a two way SMP virtual machine with affinity to at least three CPUs).*

**VMware I/O Performance**  
 VM to LUN optimization for your business critical apps  
[www.virtualinstruments.com](http://www.virtualinstruments.com)


Ads by Google

Posted in Server | Tags: ESX, performance

« SRM: Three updates, EMC / LSI / Hitachi SRA
 [Determining vClone growth rate »](#)

You can skip to the end and leave a response. Pinging is currently not allowed.

5 Responses to "CPU Affinity..."


**Virgil says:**  
 Wednesday, April 29, 2009 at 00:46



Agreed on general rule. The one environment I have seen affinity \*required\* for a supported configuration is Cisco's Unity deployment on V13. But they missed the point about +1 core as well.



[http://www.cisco.com/en/US/docs/voice\\_ip\\_comm/unity/virtualization\\_design/guide/cvvirtualdg010.html#wp61553](http://www.cisco.com/en/US/docs/voice_ip_comm/unity/virtualization_design/guide/cvvirtualdg010.html#wp61553)

**ABOUT YELLOW-BRICKS.COM**  
 Welcome to Yellow-Bricks.com a blog about virtualization, with a strong focus on VMware, HA and vCloud solutions.

Some facts:  
 999 Blog Articles  
 1950 Twitter Followers  
 5503 RSS Subscribers

SPONSORS


ADVERTISEMENTS

**Improve ESX Performance**  
 Quickly Identify Performance Issues Get VMware Capacity Analyzer Now.  
[www.VKernel.com/Performance](http://www.VKernel.com/Performance)

**Automate VMWare Processes**  
 Save Time, Reduce Errors & TCO. Learn More at our Resource Center.  
[UC4.com/VMWare](http://UC4.com/VMWare)

**Speed Up Virtualization**  
 Virtualization Slows With Frag. Try Dskeeper 2010 Free!  
[www.Dskeeper.com/VirtualServer](http://www.Dskeeper.com/VirtualServer)

RECOMMENDED


**vSphere 4.0 Quick Start Guide**  
 Shortcuts down the path of Virtualization  
 Buy it now just \$15.99.

ADVERTISEMENTS





**James S** says:

Wednesday, April 29, 2009 at 04:24

All I can say is that there are times when however unpalatable it is, CPU Affinity can be used to overcome issues and problems. For example...say you have a VM that is clearly CPU bound and is right on the edge of being able to run on your current host architecture...using CPU Affinity to keep a VM from executing on Core 0, where the SC currently is ALWAYS scheduled...could eke out a few small percents of performance gains for you. But yes...I also agree that Affinity should be avoided whenever possible.



**Joerg** says:

Wednesday, April 29, 2009 at 21:27

CPU Affinity doesn't work together with vMotion and Storage vMotion. It's also not possible to activate HA if one of the VMs have set a CPU affinity.

In a HA or DRS cluster you can't set CPU Affinity because the boxes are greyed out.

Regards  
Joerg



**Keith Chambers** says:

Monday, March 8, 2010 at 01:27

I wrote the Cisco Unity document referenced and did all the testing behind the solution. For 18 months it was my baby.

I arrived at the CPU Affinity requirement because it was the only way to maintain voice quality with other applications running side by side. This severely limited the solution but it was the reality of the beast.

I'm not questioning the technical merit behind the +1 but suggest that it's not a hard requirement. As VMware says, +1 gives you the best performance but that doesn't mean your app can't get by without it. During the Unity on VMware load testing I monitored voice quality closely and didn't experience hiccups (scheduling issues) when accessing the console without using +1.

I recommend people test their apps with end user perception in mind. Start with the least restrictive CPU protection strategy possible. Is end user perception acceptable? If not, take it up a notch and repeat. If you do have to use CPU Affinity, see if you can get away without the +1.



**Duncan Epping** says:

Monday, March 8, 2010 at 08:38

Keith, thanks for the comment but keep in mind that you were working on a unique situation where the host probably wasn't overcommitted on CPU. When it is you will start noticing that scheduling isn't as easy as expected especially not when you have a VM which has a lot of interaction going on and is actually using those helper worlds.

I agree that it is something that needs to be tested, but I would rather have an extra than running low with these kind of apps.

[www.axel.com](http://www.axel.com)

Ads by Google

## RECENT COMMENTS

- ▶ Aaron Delp on Standby NICs in an "IP-Hash" configuration
- ▶ Leif on Standby NICs in an "IP-Hash" configuration
- ▶ Greg on HACL
- ▶ Duncan Epping on Changing your dvPortgroup settings? (need your input!)
- ▶ Duncan Epping on VMworld San Francisco Run 2010

## TAGS

2.5 3.0.x 3.5 BC-DR Blogging Book  
Bugs cloud design dms ESX esxi ha  
Howto kb linux networking News patches  
performance powershell Scripting  
scripts security service console software srm  
Storage Tools training u2 update  
upgrade VCB vcdx vcenter vdi  
VirtualCenter vmtn VMware  
vmworld vSphere vstorage whitepaper  
Workstation

## DISCLAIMER

The views expressed anywhere on this site are strictly mine and not the opinions and views of VMware.

## Leave a Reply

Name (required)

Mail (will not be published) (required)

Website

Submit Comment

Copyright © **Yellow Bricks** - Building blocks for virtualization...

Powered by WordPress | Modern Style theme by FlexiThemes