



RUNAS RADIO



<http://www.runasradio.com>



Richard
Campbell

RunAs Radio is a weekly Internet Audio Talk Show for IT Professionals working with Microsoft products. The full range of IT topics is covered from a Microsoft-centric viewpoint.



Greg
Hughes

Text Transcript of Show #148
(Transcription services provided by [PWOP Productions](#))



Clint Huffman Does Hyper-V Performance Tuning!
February 17, 2010



[Music]

Brandon Wenn: From runasradio.com, you're listening to RunAs Radio, the Internet audio talk show for IT professionals with Richard Campbell and Greg Hughes. This is Brandon Wenn, announcing show #148, with guest Clint Huffman, recorded Tuesday, January 12, 2010. RunAs Radio is produced each week by PWOP Productions, providing professional media and podcasting services online at pwop.com. You can follow the boys on Twitter at twitter.com/runasradio,

Richard Campbell: Thank you, Brandon. This is Richard Campbell. You're listening to RunAs Radio. With me as always, my co-host, Greg Hughes.

Greg Hughes: That is me. Hello.

Richard Campbell: Here we are again and, you know, this is a shooting fish in the barrel kind of show because our friend, Clint Huffman, has come back.

Greg Hughes: Right, yeah. So we always get something good out of Clint, don't we?

Richard Campbell: Well, he's the guy who really tip us off on all the cool stuff going on the PerfMon and PAL and...

Greg Hughes: Exactly.

Richard Campbell: Clint, I think you've been here too much. You don't qualify for a bio. We're just going to start talking to you.

Greg Hughes: Yeah. We're just going to jump on you.

Clint Huffman: Yeah. I'm like a regular, just give me the usual.

Richard Campbell: Just the usual. And now it's the Clint Huffman section.

Greg Hughes: Ladies and gentleman, Clint Huffman. Take it away, Clint.

Richard Campbell: What have you been working on, Clint?

Clint Huffman: Okay. Yeah, in the past few months I've been working on some interesting stuff. Well, it's actually more than a few months, it's in the past couple of years but I've been working on Hyper-V performance issues and I think I've got enough that I've kind of reach a critical mass with this information and it's actually taken me about a point where I can say, yeah, these are good things to do and these are relatively bad things to do and things to watch out for in Hyper-V and Hyper-V being kind of a buzz thing right now, or at least more specifically virtualization I'll

Clint Huffman Does Hyper-V Performance Tuning!

February 17, 2010

say that, that is becoming more of a buzzword. I figured it might be a good topic to talk about for the show.

Richard Campbell: I've got to presume there's like a PAL instrument pack for Hyper-V now.

Clint Huffman: Well, yeah, there was. Well, there is actually a PAL v1.3.5 that's out right now, the latest 1x version. Yes, there is a Hyper-V threshold file in there that looks at the Hyper-V, Hypervisor Counters. It needs a lot more work. I've learned a lot since then so in PAL 2.0 I hope to give it a nice overhaul.

Richard Campbell: Cool. Yeah.

Greg Hughes: Where do we see performance issues when you're virtualizing user Hyper-V? Is it in great big installations, single machine type of situations, everything in between? Where do you focus most of your energy?

Clint Huffman: Yes. That's the answer to that question, Greg. But mostly what I've been using it for is, one, here at home I use it for reproducing problems. I've got my own virtual BizTalk lab here that I used Hyper-V in everything like that so it's nice there. But mainly where I've been focusing as far as job-wise is really Enterprise Services, just large BizTalk installations that just need to scale out. Web servers are really popular as far as virtualizing now. So I've been basically doing with the small personal level and the high enterprise level and I'm sure people are doing a little bit between there. The Hyper-V was originally intended for kind of a backward compatibility. You know, if you've got an old NT 4.0 box or something that's really old that gets to either migrate and still considered critical, then you want to consider the hardware consolidation and put all that Hyper-V and things of that nature. So now with the original on set, I haven't really heard of anybody really doing that yet but maybe that's just what I do that takes away from that.

Greg Hughes: So what are the common problems that you see or that you're dealing with from a performance perspective or whatever kind of problems you tend to run up against?

Clint Huffman: Well, I deal with here lately it's been mainly focused on Hyper-V performance. There might be other problems that people had as far as compatibility maybe and such, but my area has been performance related and I've got a couple of tips and things I want to talk about on the show here about what to watch out for. Now, granted most of these tips and tricks that I'm about to tell you are really from a guy named Tony Voellm who was previously the dev lead for Hyper-V performance on the product team. He's still with the product team. He's with the Windows product team now but he deals with more



battery life and performance tuning of Windows itself but he's still kind of the de facto source for Hyper-V in performance issues.

Richard Campbell: So where do we start?

Clint Huffman: Okay. So really the thing is making sure that you've done all the best practices for Hyper-V, but primarily speaking here make sure that you're aware that you're going to have a performance hit by going to a virtualized environment. Anytime you add an extra layer of operating system like this you're always going to incur a little bit of cost.

Greg Hughes: Sure.

Clint Huffman: The cost really varies. Most people see about 10% to 20% degradation of performance depending on what kind of application you're doing. As far as with virtualization, I mean there's really no like two need per se that you could do but there are a lot of best practices that you can follow in Hyper-V environment to make things working better. One of the main things to know about this is that you're going to clock in the -- the virtual guest is going to be skewed slightly. What that means is something like percent processor time is going to look a little -- it's going to look a little higher at the virtual guest but on the host machine the processor time that you see from the Hypervisor Counters is going to appear to be a little bit less. So if you're 100% CPU in the virtual guest, it may look like 90% CPU as far as real processor time at the physical level from that virtual machine. It's a little bit askew there as far as that.

Richard Campbell: And you said the Hyper-V counter is here because Task Manager lies to you in the Hyper-V situation.

Clint Huffman: That is correct. Where did you hear that?

Richard Campbell: Well, I'm running enough Hyper-V now to know. Well, the saving grace has been SCVMM which is the System Center Virtual Machine Manager. It gives you better numbers. I'm looking at the System Center Virtual Machine Manager numbers and going "How come Task Manager doesn't show this?"

Clint Huffman: I thought you were reading my blog there, Richard.

Richard Campbell: Oh, I read your blog too, sir.

Clint Huffman: Yeah, so I actually have a blog entry. I do now use Task Manager in a Hyper-V environment or something along those lines, and yeah, you're right. Looking at Task Manager on the room partition, which is really the host machine, is not going to give you the true processor overhead. You

actually have to look at the Hyper-V Hypervisor Counters or in this case the SCVMM to actually see the true numbers of what's actually going over there. As a matter of fact, if you go to percent processor time on the host machine that will also lie. It only gives you the processor time of that machine. The reason it's like that is like the room partition of the host is actually considered like another virtual machine running outside of the Hypervisor and that's why you see it kind of standalone on its processor time.

Richard Campbell: Right. So you're literally just seeing the processor time that the host machine consumes. You don't see the guest machines at all without going into PerfMon and going specifically. What's the counter? Hyper-V...

Clint Huffman: It's Hyper-V Hypervisor Logical Processor % Total Runtime.

Greg Hughes: And that will give you the accurate numbers.

Clint Huffman: Yeah, that will be the physical processor as far as what they're actually using off the box and yeah, you could have your physical processors at 100% CPU and you look at the Task Manager on your room partition and you're like "Oh look, it's only consuming 5% across all the processors," while in reality it's just overwhelmed.

Richard Campbell: Yeah and what's sneaky about that is it can be completely concealing. One of the things I've struggled with is understanding when I should add another processor to a guest OS which I can do now.

Clint Huffman: Yeah.

Richard Campbell: I can actually say give them two logical processors. So I'm just trying to figure out when does that make sense to do. It has been fairly challenging to instrument.

Clint Huffman: Yeah. The way it works, a lot of people assume that the virtual processors are like affinities to a physical processor and unfortunately that's not the case. I'm not saying unfortunately but that's just not how it's architected. The virtual processors are really just thread issuers where the virtual guest will simply get a lot of threads that need to have attention to the processors and then those threads are then like load balance across the physical processors. So for example, let's say you have four physical processors on the physical box and you have a virtual machine that has one virtual processor. If that virtual processor goes 100% CPU you'll see the physical processor distribute that load and you'll see about 25% CPU across that, across all four processors.



Richard Campbell: And it's always happening here as Hyper-V is sharing the threads across these different processors. This is presuming there's a bunch of threads running which is almost certainly there are.

Clint Huffman: Yes, exactly and that's absolutely correct. You know, any one thread can only run on one processor at a time.

Richard Campbell: Right. So even more so what does it actually mean to add another logical processor to a Hyper-V instance?

Clint Huffman: Well, in that case you would because if you want that virtual machine to use more CPU you have to give it another processor because it's only consuming up to one virtual processor that it could handle and then distributing that. At this point you're only using 25% of the total amount of CPU possible on the virtual box.

Richard Campbell: Given that we have four cores.

Clint Huffman: Given four cores precisely, yes. If you add another virtual processor now or if you're still on a four physical processor box, that will all go up to now 50% because now you have two virtual processors consuming 100% CPU and dividing that amongst four processors that now ends up being...

Richard Campbell: So I started out thinking so this is just sort of an arbitrary value of how fast one logical processor can run. But is Hyper-V actually smart enough to map that, the power of that logical processor to one of the physical processor performance-wise so that it does work out perfectly given four cores, 100% of a logical equals 25% of the four cores.

Clint Huffman: Yeah. Are you leading me there, Richard? Do you already know the answer here?

Richard Campbell: No, I'm just trying to...

Greg Hughes: He noted my question too actually.

Richard Campbell: I'm trying to understand, my friend, honestly.

Clint Huffman: I know, I know, I know. Yeah, so I think you know more than you let on sometimes. Yes, so there is a way that there's this thing Enlightenments that where if the virtual guest happens to be "un-light," that means there's more Hyper-V aware or where are the Hypervisor, and what I can do is it can do what they call Hypercalls straight to the Hypervisor highly bypassing the room partition or the host in this case. So yes, it can make it much more efficient but there's a cost to that and the cost to

make that happen is really you have to have the guest on Windows Server 2008.

Richard Campbell: Okay. It's such a great name, Enlightenments. I'm cognizant of a higher power.

Greg Hughes: Going back to our earlier conversation, so the only way really to be enlightened is to use Hyper-V and the only way to be enlightened with Hyper-V is to basically give it and go all and go Windows Server 2008.

Clint Huffman: Yeah, exactly. I've been on several engagements so far where they guess it was Windows Server 2003 and they're like "Can we get any better efficiency out of this," and I'm like, "Well, not unless we go to Windows Server 2008 to make it more efficient, get better response time, etc, from the virtual machine" assuming that there's no other bottleneck because that's one thing that I go on site for first to say "Okay, is the virtual environment is that the virtual environment the actual bottleneck in the entire solution?"

Richard Campbell: Right.

Greg Hughes: Right.

Clint Huffman: Because if you have a physical web server that is using it back in SQL box and you have a virtual web server that's using the exact same SQL box, well, if SQL Server is the bottleneck then the two servers are going to perform the same way. You know, you're not going to see much difference and so a lot of that effort really comes down to it's not really a Hyper-V or virtualization process, it's really about identifying where your bottleneck is. So that's the first thing I do when on site.

Richard Campbell: Yeah. I do a great conference session where I actually scale up a website from one server to two and get no performance gain at all because the database is completely buried. That is the linchpin. What resource do you run out of most often in a Hyper-V situation?

Clint Huffman: Well, that's always going to be dependent on your application.

Richard Campbell: Right.

Clint Huffman: So I mean depending on your needs, on what you're doing. So just like any application, it really comes down to what resources does it use.

Richard Campbell: Yeah.

Clint Huffman: Is that resource pretty much consumed. Here lately though my past couple of issues have really been processors.



Richard Campbell: Really.

Clint Huffman: And I'm really quite amazed by that.

Richard Campbell: Yeah, because the processors don't like the lowest man on the totem pole. For me, I find that if it's SQL Server we're out of drive performance. If it's IIS we're typically running low on memory. Like their sort of the things that eats when it gets really busy are fairly consistent and processors never seem to be the one.

Clint Huffman: I know and whenever I teach Windows Performance Analysis I always teach processor last because it's always the guy who holds the bag. If you had a really bad disk performance problem, then it's going to eat up memory to compensate for that, the system cache consuming, you know, doing kind of three rides and holding all that in the memory which now becomes a memory problem and then the processor then has to kind of compensate for all of these by working harder to manage all the memory in the Disk I/O. So yeah, it's actually pretty rare to see processors being bottlenecks. This particular application is really optimized I guess for threading handling. Most of these applications I've been dealing with, they spawned up new processors, really short-lived processors. They only last maybe a few seconds and then do their work and then die off. So by changing that to a thread model they should get better performance. I typically see this of course from UNIX boxes where the UNIX is really good at spawning processors but not so much threads, while on Windows it's more of the inverse of threads versus due process.

Richard Campbell: Right. The overhead of constructing a process is high enough that you're not getting a good performance benefit for doing it that way.

Clint Huffman: Right, right. All right, and of course it makes sense to me that we would now lean -- because processors have had an easy time of it, they've always been so much faster in everything else in the box, we're starting to lean more on processor demanding type of things so virtualization is inherently processor centric.

Richard Campbell: Yeah. It's just that that's where all the work ends up going and throwing that extra whammy and it's bloody hard to instrument.

Clint Huffman: Yes, sometimes. Oh, we do have Xperf which definitely helps us out a lot in identifying where we're called.

Richard Campbell: All right, let's talk about Xperf. What is it and why do we care?

Clint Huffman Does Hyper-V Performance Tuning!

February 17, 2010

Clint Huffman: Yeah. I want to do some of those. In fact, I can't...

Richard Campbell: Yeah.

Clint Huffman: The Microsoft Xperf comes with the Microsoft Windows Server 2008 Performance Toolkit and it's a toolkit that has been in development over 10 years by the Windows product team and what it does is basically collects ETW information or event tracing for Windows. It collects this data information which is a lot of events going on in Windows and then basically puts it together in kind of a PerfMon-style format to say here, I'm giving all these events, here is the relevant information that we can provide to you with that, and you can actually see like for example from the platter like the disk platter, you can actually see the head moving back and forth doing like a whole bunch of sequential writes all over the place. So it's an amazing technology.

Richard Campbell: So is this a replacement for PerfMon, you think?

Clint Huffman: No, no. The Xperf is in addition to. Think of Xperf as the microscope for really digging into a very, very specific problem as far as re-gathering a massive amount of data in a very quick way. PerfMon would be kind of the microscope I guess or kind of going back a little bit and saying, okay, you can gather Performance Counters for hours, days, or weeks or months.

Richard Campbell: Right.

Clint Huffman: But you can't gather ETW data for hours, weeks, or months. It's really design for short term gatherings. With that said, you can config your Xperf to do what's called a revolving buffer where you basically say like 100 megabytes buffer in a RAM is yours and completely keep overwriting that data until finally you have a problem and now you can stop Xperf and then look at the data and see what it has to say.

Richard Campbell: Right. It becomes just sort of black box recorder. It only has the last 20 minutes of the machine but that's the important 20 minutes.

Greg Hughes: Yeah.

Clint Huffman: Right. When we were dog footing Windows 7.0 and Vista internally at Microsoft, we would have these things running constantly. They had a little applet that we run doing this and as it was running all the time and then all we had to do is like if you have a problem or when we had a problem, you would hit say "Stop this connection, I was opening Outlook or doing something in Outlook," type that in, hit submit and then that ETL file once its given trace log file gets set up to the product group or whomever it has a really need for it, Outlook team maybe, and



then they would analyze it, compile. Yeah, I think most people in the company are having this problem, a performance problem, and that's really when Windows 7.0 kind of came about. It was all that work we did in Vista.

Richard Campbell: That's awesome. Okay, so Xperf drills into particular detailed aspects of the machine so you can sort of get a grasp when you really got serious problems there. That really helps in the Hyper-V context too?

Clint Huffman: It does, yes. As a matter of fact, most of the Hyper-V performance analysis that we do is using Xperf and you want to do it primarily on the other room partition of the Server 2008 box to see what's going on there. Yeah, great stuff. I actually know what it does. I can actually see a VHD that was fragmented just from the disk analysis. I got "Look, your VHD file is fragmenting there," just because I can see the head kind of swiveling back and forth trying to do this sequential read.

Richard Campbell: Interesting, yeah.

Greg Hughes: Cool.

Richard Campbell: And that's another common complain that I've heard with Hyper-V, it's that you reach a point where -- well, you explain it, Clint. What is this defragmentation thing?

Clint Huffman: Oh, okay. Defragmentation is where basically pages of memory, either on the disk or in memory, it gets fragmented. Kind of think of like a little bit of speculating where let's say you have a completely nice sheet of paper and then you're like, okay, I've allocated 100% of my memory or 100% of disk space on this paper and then you only start saying, okay, now I want to freeze some memory and you start poking holes in a piece of paper. You've got all these holes in this piece of paper, you know, little finger-sized holes and you've got them all over the place to the point where maybe you have more than 30% free space and disk memory because of all these little holes you poked into it but now you've got this big block, this 4 x 4 block that you now need to put into this memory space on this paper through a hole of some kind, but yeah, you don't have any holes big enough or at least this contiguous holes that are big enough to put that block in and now you're able to call a heap fragment and at that point you can't put the block in there. I'm not talking memory fragmentation but same thing can happen with disk where if it is not contiguous then it doesn't perform well.

Richard Campbell: And how do you recover from that?

Clint Huffman: Defrag. You run the defrag tool in Windows or third party and it just basically puts

Clint Huffman Does Hyper-V Performance Tuning!

February 17, 2010

everything together in a nice contiguous block so that when a head reads off the disk it is written all in one rotation basic layer however long it takes.

Richard Campbell: So we should be running a defragger on the host OS to defrag the VHD files?

Clint Huffman: I didn't say that.

Richard Campbell: Okay. I know.

Clint Huffman: I'm just telling you what defrag does in the said stuff. I would say, yes, the disk response time on the roots is having some problems meaning it works great in 15 milliseconds on average for doing reads and writes and if the VHD file seems to be the one hit the most during this timeframe and if you've done a defrag analysis and the VHD file appears to be most fragmented, yeah, you probably wanted to defrag at that point. The reason I'm a little hesitant to say doing defrag is because if you're using SAN, it is common to use SAN with the Hyper-V environment and in the Enterprise environment, sometimes the SAN gets worse performance if defragged. So I would say always consult with your SAN vendor or hardware vendor you work with to make sure that defragging is the right thing for you.

Greg Hughes: Got you.

Richard Campbell: Do you run this defrag task within the guest or in the host?

Clint Huffman: Oh, good point. I would guess if you're seeing bad performance in both, sure, try them both. Do you know something about this that I don't, Richard?

Richard Campbell: No, no. I'm just curious because when you talk about the disk from a guest perspective, it's not like you really have control over the media per se. The host sort of decides where things are actually going. It's virtualizing all that to the guest so I would think the host is the only thing that could control it. But by that token if I'm running a defrag on the host, can I actually leave the guest operating while I do that?

Clint Huffman: I have never tried that.

Richard Campbell: It's an interesting...

Clint Huffman: I don't know.

Richard Campbell: Because the correct answer I think most of the time with the SAN is you can't run a defragger, it won't help.

Clint Huffman: Yeah.



Greg Hughes: Because the SAN is optimizing for you.

Richard Campbell: Yeah, it writes the stuff out of the disk in its own way but yeah, this is interesting, an area that's fascinating to me in just the consequences of this and how you actually undo it intelligently.

Greg Hughes: Can you defrag an active VHD file while it's running in Hyper-V on the host? I would guess you probably can't.

Clint Huffman: Yeah. It will probably be locked out or locked by the operating system or something. I don't know, I've never tried it. More importantly though if you're having disk problems in a guest machine, say you virtualized SQL Server that requires massive amounts of Disk I/O, you really need to be using what's called Pass-through Disks and the way Pass-through really works is really you're giving the LAN itself a physical run represented to the operating system.

Greg Hughes: Right.

Clint Huffman: The way you do it is you go to disk management on the room partition, the host machine, and you actually right click on the physical LAN, a physical disk and stay offline.

Richard Campbell: Right.

Clint Huffman: And so the operating system, the room partition no longer sees that disk, or it's not online at least. Then you go to the Hyper-V manager and now you go to settings on the virtual machine, now you say assign that LAN, that physical disk to the virtual machine now and that's the first to give you a lot better performance. We did it for SQL Server and it actually made a huge difference in performance for us.

Greg Hughes: So a physical disk on a host assigned to a virtual machine because you basically allocated it just for that.

Clint Huffman: Right and I suggest that for my last customer. Unfortunately they're like, well, we can't really or they don't want to put another physical disk attached to the server to do that.

Richard Campbell: Right.

Clint Huffman: So they stop with VHDs in that case. So yeah, that would give you much better efficiency because now you go straight to the Hypervisor versus going through the -- assuming it's a light stop too that it won't go through the room partition or host anymore.

Richard Campbell: Yeah and avoid some virtualization stepping. It owns the resource. Is this

the same thing with NIC now, right? That you literally assign the NIC to a certain VM and it owns it. It's no longer virtualizing the networking.

Clint Huffman: The way it works is that you actually assign the NIC, the physical NIC to a virtual network in Hyper-V.

Richard Campbell: Okay.

Clint Huffman: And then you assign the virtual machines to that network. So yeah, you could do a one-to-one ratio there.

Richard Campbell: Right, but essentially it still has that layer of abstraction. It's still sharing it.

Clint Huffman: Yes.

Richard Campbell: So you can assign multiple virtual machines to a given NIC even though you've assigned the NIC into the pool. I find the Hyper-V networking more complicated. It does take some time to figure out what am I doing with this? Is the host using it as well? Is the sharing within the machines in the Hyper-V setting? There are some interesting stuff there.

Clint Huffman: Yeah. We actually have an example in VMware. Yes, exit environment where this media system or media server that a customer have had all four virtual machines going to the same physical NIC.

Richard Campbell: Right.

Clint Huffman: And they're streaming audio and video and stuff like that and so you actually see the output queue length on the network adaptor just go way up to like 60, I mean 60 packets constantly waiting to go on the wire.

Richard Campbell: Wow. Were they watching that from the host?

Clint Huffman: From the guest perspective.

Richard Campbell: Okay, but yeah, that's a busy NIC, no two ways about it.

Clint Huffman: Yes, it was.

Richard Campbell: Well, I'm far more used -- back in the 10-based T days it wasn't that surprising when you backlog in NIC but I think we've just gotten sort of lazy with gigabyte NICs now that they've got so much headroom we just never fill them up. I mean when did we ever look at network performance and not see like 4% use. You know, it's never busy.



Clint Huffman: Yeah, but as far as constructing and deconstructing TCP/IP sessions, now that's what gets you.

Richard Campbell: Right.

Clint Huffman: The latency of creating and storing those connections is the big killer nowadays.

Richard Campbell: Yeah. Now you can buy substantially more expensive NICs where they offload that work on to the NIC and it's really fast at it.

Clint Huffman: Yeah and Hyper-V can take advantage of that somewhat. There's this thing called virtual memory queues that can help take advantage, the Hyper-V can take advantage of. There's actually a whitepaper on it where you can use chimney offload and virtual memory queues, either one or the other, to really take advantage of the hardware and the NIC, otherwise, it can't take advantage of that. It will slow things down a little bit in performance. Now the virtual guest actually had a 10-gigabyte connection, a virtual 10-gigabyte connection, and if you actually do a file transfer from guest to guest on the same physical box, they actually do a memory map versus through a network course basically and so it can actually be pretty fast in a lot of cases. But going off the box it can actually take care of the hardware and actually be pretty fast if you take advantage of that. Along those lines you want to use what's called synthetic drivers or synthetic devices versus -- what's the other thing?

Richard Campbell: Physical?

Clint Huffman: Emulated, emulated.

Richard Campbell: Oh okay, yeah.

Clint Huffman: So what I mean by that is wherever you're using network adaptor, I think you actually run into this too, Richard, where you want to use -- when you query a new network adaptor you want to use, that's called a synthetic one. But if you go to the legacy network adaptor, that's considered an emulated one.

Richard Campbell: Right.

Clint Huffman: But that's not as fast as far as performance.

Richard Campbell: I ran into a piece of hardware where the synthetic one wouldn't work but the emulated one did.

Clint Huffman: Yeah, simple backwards compatibility.

Richard Campbell: Right, yeah. I also run into the interesting problem where Hyper-V does optimize the processor and I have an AMD Server and an Intel

server and that meant that portability didn't work automatically where I couldn't just use SCVMM, point to a VM and say put that on the other machine and go "I'll take a different processor, sorry."

Clint Huffman: Yeah. The other caveat -- yeah, go ahead.

Richard Campbell: There's actually a switch to say put this in compatibility mode so it wouldn't be portable between processors but you lost a big chunk of performance when you do that.

Clint Huffman: Okay. Yeah, along those lines there's also a technology called SLAT, Second Level Address Translation, that's in the latest processors and AMD calls this feature Rapid Virtualization Indexing. It's on their Phenom and Opteron processors.

Richard Campbell: Oh yeah.

Clint Huffman: So that's supposed to help with making it a lot faster. An Intel versions of it is on the column-based processors and Xeon 5500 and above, 5500 series. The Intel Core i7 desktops also have the SLAT features.

Richard Campbell: Right.

Clint Huffman: So giving them an extra boost and processor efficiency. I have not seen the performance of these processors directly yet, but the technology is what I hear is out there.

Richard Campbell: Funny I've got one of each. I've got Opterons, some got some Xeon 5550s and I just put an i7 machine together as a new workstation so I guess I'm going to experience all of these.

Clint Huffman: And you're running Windows Server 2008 as a guest?

Richard Campbell: On some of the machines, yes.

Clint Huffman: Okay because the Server 2008 box should begin in better performance being "unlicensed."

Richard Campbell: Yeah.

Clint Huffman: Hyper-V aware of like...

Richard Campbell: Taking advantage of it, yeah. It's going to be interesting to trial that. I can spend all my day just testing this stuff. I have so much fun with it.

Clint Huffman: Yeah.



Richard Campbell: Well, Clint, I think we're about out of time. Any final words, places folks should be looking about Hyper-V performance?

Clint Huffman: So Tony Voellm's blog. Tony Voellm is T and then Voellm. I want to make sure he gets credit for a lot of his work because he's got a fantastic blog up there about how do you do Hyper-V performance. Also Ewan Fairweather and I did a lot of work on BizTalk and Hyper-V performance so just search for our names. It's up on MSDN about doing BizTalk and Hyper-V and we talk a lot about doing more everyday Hyper-V things for general operating systems performance as well so that's a good one. There's more information out there for performance tuning Hyper-V. I only kind of touch on the major ones but I plan on doing a blog specifically on Hyper-V performance pretty soon and my blog is blogs.technet.com/clinth, and for a final word, "Hi mom."

Greg Hughes: Thanks, Clint.

Richard Campbell: Thanks, Clint, for coming on the show again and we'll talk to you next week on RunAs Radio.