RunAs Radio

http://www.runasradio.com

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.

Text Transcript of Show #077
(Transcription services provided by PWOP Productions)

Steven Choy Measures Server Performance!
October 1, 2008

[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 #77, with guest Steven Choy, recorded Thursday, September 25, 2008. RunAs Radio is produced each week by PWOP Productions, providing professional media and podcasting services online at pwop.com.

Richard Campbell: Yes, you're listening to RunAs Radio. I'm your host, Richard Campbell, and with me as always my co-host, Greg Hughes.

Greg Hughes: Hey everybody. Richard, how are you again today?

Richard Campbell: You know, things are good, man. This is as much as we could do. I'm just glad we got a bunch of shows to record and this is going to be a really exciting one.

Greg Hughes: Yeah, keeping it busy and keeping it really interesting.

Richard Campbell: And gearing up for TechEd of course is in November. We still have some time to go. If you haven't heard already, .NET Rocks is running another sweepstakes contest; so that's at dotnetrocks.com, so pop over there and win yourself a free trip to TechEd in Barcelona. We'll be there.

Greg Hughes: Yeah; and you can choose this year or next year, is that right?

Richard Campbell: Right, although I hear next year is not in Barcelona.

Greg Hughes: Yeah, I've heard that too.

Richard Campbell: If you've got any questions, comments, show you'd like to see, you name it, fire us an email, info@runasradio.com. All right, let's get right to our guest here. Steven Choy is a Senior Premier Field Engineer with Microsoft Premier Field Engineering. In this role, he has provided support to government agencies and commercial clients. Among the projects he has worked on are: desktop deployment and migration, patch management and group policy. Currently, his work focuses on server virtualization and server health. Before joining Microsoft, Steven worked for PricewaterhouseCoopers, Lehman Brothers and T.J. Watson Research Center, which is an old branch of IBM if I remember correctly and good thing you're not working for Lehman Brothers anymore.

Stephen Choy: I know.

Greg Hughes: Hey Steven.

Stephen Choy: Hi. Thank you for having me.

Richard Campbell: You must have had some friends still over at Lehman.

Stephen Choy: Yeah, I have a couple of friends but at least they do have a job at Barkley.

Richard Campbell: Right. Well, I've always found talented people invariably land on your feet, you know, if you're good at what you do, somebody wants you.

Stephen Choy: Yup, that's true.

Richard Campbell: So I have not heard of the Microsoft Premier Field Engineering. This is not Microsoft Consulting Services.

Stephen Choy: No, but they are all under Microsoft Surface.

Richard Campbell: Okay.

Stephen Choy: We do more operation support and we also do some proactive engagement, a little bit different from MCS is they do more short-term and we do a longer term engagement stuff.

Richard Campbell: It sounds like it's much more IT oriented where I think consulting services is very dev oriented.

Stephen Choy: We have most of the time IT infrastructure. We also have some development engineers helping people for EDCs, helping people with their .NET coding and VBScript, that kind of stuff.

Richard Campbell: Oh, I get it, okay. That's cool and it must be fun to jump into all these different projects and get a chance to see what's going on.

Stephen Choy: Yup, you can see different kinds of environment and then you can compare from place to place. It's very exciting and they are challenging too.

Richard Campbell: We bumped into each other around server performance because I think it's one of those areas that people really don't know what to do when they really talk about how does an overloaded server look like and when I'm on stage talking about this, I say it's not like a sign pops out of the machine that says, "Help."

Greg Hughes: "I'm overloaded."

Stephen Choy: Yeah, yeah, typical. On my engagement, when I come in people will just call me, "Oh, my application is not running, or slow. What should we do?" and that's when we come in. On a server, we typically look on the disk, memory, processor, CPU, and network.

Richard Campbell: Right.

Stephen Choy: And see which one has high usage and then we can pinpoint what's the bottleneck and also we need to find out the role of the server because different roles of server, different resource will be more heavily used. For example, application server, we'll definitely look at the processor, memory, and CPU type. For a mail server, this is also important. We look at this processor and the memory too.

Richard Campbell: Right.

Greg Hughes: So processor, memory, CPU. Are there resources in a server environment that can impact performance?

Stephen Choy: Yeah, the network and the disk.

Greg Hughes: So IO and network.

Richard Campbell: It's really those four things, CPU, memory, disk, and network speed. For me, I've always tended towards the database side of things when doing a lot of performance tuning. You're so slave to the disk. It's the slowest piece of the machine.

Greg Hughes: IO, yeah.

Stephen Choy: Yup, and also the process too. A lot of times, my client will say you're running slow and have to run through and then I'll pinpoint, "You have too many bar handles. You have too many threads."

Richard Campbell: Right.

Stephen Choy: If you're miscoding a program, then you will of course from time to time suddenly your program will fall and say, "Oh, this doesn't work anymore," or it's so slow on the server.

Richard Campbell: That really gets into the scaling things where the app worked just fine in test but as we get it out into the wild, there is so much more demand on it that the thread count gets out of hand.

Stephen Choy: Yup, definitely.

Richard Campbell: So when it comes to figuring out performance like this, is it all about PerfMon? Is that the tool?

Stephen Choy: We are mainly using the PerfMon especially when we are talking about Windows 2008. We have an enhanced version of PerfMon. We call it VPN Performance Monitor and sometimes I call it PerfMon on steroids.

Richard Campbell: Nice.

Stephen Choy: Because it has a couple of enhancements. We put a lot of different tools combining one single UI launched from the Windows 2008 Performance Monitor. For example, they have a SPA, the Server Performance Advisor window currently available for download as a standalone package on Windows 2003 Server and that was integrated onto 2008 as well as Vista, they have the same.

Richard Campbell: There's a built right in, this Server Performance Advisor.

Stephen Choy: Yup and as you know, the Server Performance Advisor, they give you a comprehensive diagnostic report and very good on checking IIS and Active Directory.

Richard Campbell: Interesting.

Stephen Choy: And also they morph on it. They also have, if you use Server Performance Monitor with it, that is also incorporated into that tool too so you can create performance locks inside the tool. You define a template so it's much easier. You don't have to download individual package and actually, the Performance Monitor Vista will not work on Windows 2008.

Richard Campbell: Oh okay, so you have to use the ones that are built into 2008.

Stephen Choy: Yup, yup and I think one of the most useful features is they allow you to scale multiple content, automatic scale on vertical axis so you can easily see the graph. Before, our PSS engineer actually internally developed a tool to just doing that scaling all the performance context so that you can see the trend. People will say a picture is worth a thousand words.

Richard Campbell: Right.

Stephen Choy: So if you can see the graph, you can easily pinpoint what caused or you can see any memory leaked. That's one of the big part on Windows 2008.

Richard Campbell: The big thing for me with PerfMon is leaving it on for a while long enough to see those various little events that are sort of tipping points. When I've been fighting scaling websites, it's when did the garbage collection run and actually catching that in PerfMon was always tricky.

Stephen Choy: Yup, yup, and also on the new version, they collect a thousand data points now. In the previous version, they can only collect a hundred data points so the graph will display more accurately.

Richard Campbell: That's just a lot of data to have on a graph. It's got to be tough to read it.

Stephen Choy: Yup, yup, yup but on the new version, you have the timescale. You can easily move, expand it or contract it so it's helping you to look at the graph much easier.

Richard Campbell: Right. One of the complaints that I've heard routinely around PerfMon is that figuring out what to watch. It's like I feel like PerfMon is my own little internet. Everything I want to find is in there, I just can't find it.

Greg Hughes: Or you don't know which parts to use, right?

Stephen Choy: Yup, that's why we try to give out some special or typical content. For example, memory, you typically look at the page pool on CPU. You look at a certain date, process time, that kind of stuff, so that's one of my people under TechNet tried to give out some of the contents of you'll look for.

Richard Campbell: So, that means there are docs in TechNet to help you sort of pare through this but I guess this also gets into why Server Performance Advisor is so interesting is here is a tool that the people building the product have added in, you know, this is what good numbers look like.

Stephen Choy: Yup, yup and if you like Server Performance Advisor, I will recommend that you look for a tool called Performance Analysis of Logs and we call it PAL. It's actually written by one of colleagues, Clint Huffman in VBScript to help analyze performance kind of log and the tool comes with a lot of pre-defined specials.

Richard Campbell: This is on CodePlex? It's codeplex.com/pal for Performance Analysis of Logs.

Stephen Choy: Yup, that's the one and when you download the tool, it comes with other special files that contain most of the Microsoft folder like SQL Server, Biztalk, Active Directory, and IIS, and all those specials have been established during our Microsoft Y 2 Design Workshop and the tool has a GUI interface that you can modify, a special file to suit your client environment because some clients, some of the specials are too low to give out force positive or too high. At the end, they would create an HTML-based report file with all the graphic charts and how you'll exert when they offer their special and the PAL can run on XE, Vista, Windows 2003, Windows 2008, on both 32-bit and 64-bit.

Greg Hughes: Oh okay.

Stephen Choy: And one good thing is it's free and you can even -- the source code is licensed, as open license, so you can download and modify and incorporate and make changes on it if you want to.

Richard Campbell: Cool.

Greg Hughes: Great.

Richard Campbell: I even see there's a video here you can watch on how to use it.

Stephen Choy: Yup.

Greg Hughes: It's amazing how many resources there are on TechNet and other related sites out there that if you do just a little poking around, you can find an awful lot of great stuff.

Richard Campbell: I feel like I was supposed to know all this stuff and I didn't know about this. I need to run this on my servers.

Stephen Choy: Yeah, it's good.

Greg Hughes: What's amazing is can you ever really know everything that there is to know? That's what the resources are there for, right? Even Richard...

Richard Campbell: Dude, I've tried.

Stephen Choy: That's it too is just how you do apps, give you some data and how you save some time yourself looking at the tons of data from the PerfMon and then from then on you kind of pinpoint where that goes if it makes sense or not.

Richard Campbell: Yeah, this is great stuff, very interesting. So I see this raft of tools and so forth, maybe we should talk through some of the common issues you run into in field engineering on overloaded servers. What are people doing out there that's causing them grief?

Stephen Choy: Typical reasons. People misconfig with .ini. A lot of times they have a lot when 19 or 20 running out of a non-page and page pool memory. For example, on Exchange Server, you put

TCP switch on it and what that means, great, you get extra 1 gig memory for the user mode.

Richard Campbell: Right.

Stephen Choy: So that they can run much faster but there's no free lunch. Virtually, what you do is you take away 1 gig of kernel so everything is like based on PDE page or page pool is squeezed into 1 gig memory and constant creates a lot of problem.

Richard Campbell: I've run in exactly the same problem with SQL Server. It sounded like a good idea to use a 3-gig switch but you strangle the OS so much you actually end up slower.

Stephen Choy: Yup and sometimes people think with 3Gb, you put more memory will help. Actually, it will hurt because when you squeeze, the kernel space is less. You put more memory on it, it will have overhead on it. So, actually the system becomes more unstable than before.

Richard Campbell: And of course the correct answer to this is get to 64-bit and get there soon.

Stephen Choy: Yup because the 64-bit, the kernel, all those TCP -- you don't even need to use TCP by default, you can have even running on applications, you can 4 gig memory with yourself.

Richard Campbell: Right.

Stephen Choy: Also, other pool page memory and on page pool memory in the range of about gig memory so you shouldn't have any problem on it and also, typical people, the worse combination besides TCP is thinking, "Oh, I have more memory so maybe more than maybe 16 to 32 gig. Why do I put TCP and PAE on it?" PAE, we call it Physical Address Extension.

Richard Campbell: Right.

Stephen Choy: It's similar to old time expanded memory doing the mapping but with more memory on it, you are actually squeezing more mapping on the kernel space and then typically you'll get 2019 2020.

Richard Campbell: You really don't need to play with any of that once you're over the 64-bit realm.

Stephen Choy: No, no, you don't need that.

Richard Campbell: You leave it alone.

Stephen Choy: So that's why you should move to 64-bit passwords. Also, one common mistake is when people doing server consultation, so they are thinking, "Okay, I have a couple of servers that have 4-gig memory. Okay, when do we combine all of them into one server with 42-gig memory," and it should be fine but they forgot about the kernel space. It doesn't matter how many memory or physical memory you have, you're still limited on 32-bit access with 2-gig memory.

Richard Campbell: Sure. Server consolidation is an interesting angle on this and I've often found that mixing apps in one consolidated server having -- I would never put Exchange and SQL server on the same box. It just gives me chills.

Stephen Choy: No, no, definitely not.

Richard Campbell: It would be better off keeping it isolated but I guess virtualization is how we address this and now we have big boxes running isolated environments for each of our core apps.

Stephen Choy: Yup and that's why we have in Windows 2008 you can run on Hyper-V and do server consolidation.

Richard Campbell: Are you happy with Hyper-V? Do you think that's the way?

Stephen Choy: Yeah, I mean it's not perfect right now but I think it's very good compared to a market competitor, especially if you install your host as a core, Windows 2008 core, we don't have to order out GUI and you will make yours much quicker and use much less memory on it. Also, if you want to monitor a core, the only thing is because you can run Window with an ATI monitor, server monitor because they don't have a GUI interface on it, what you have to do is you run it remotely from a full version of Windows 2008 or they have ordered a log name. You can still run a log name locally on that box and they have to collect the log. You can bring it over to a full version of Windows 2008 analyzing it.

Richard Campbell: That is something I really wanted to dig into is how do the performance metrics change once we're in a virtualized environment? I mean don't some of those numbers become questionable? How will I really understand processor load once I'm in a virtual machine?

Greg Hughes: Shared environment, yeah.

Stephen Choy: Most of the performance techniques or conducts are similar except a couple of things that we add on the Hyper-V. We have the Hyper-V logical processor that will compare the overall resource utilization on other guest machine as well as the host machine and also the network because on Hyper-V, we also create a virtual network so we also have another content called Hyper-V Virtual Network to measure actually the network

performance on the virtual network. So, principally, all the other measurement techniques are the same, something like first you'll need to create a baseline. If you don't have a baseline, you don't know what you're going to compare.

Richard Campbell: Right.

Stephen Choy: And also the what load on the server also depends on your business cycle so typically you should have at least one baseline or width so you'll know all the up and down of what's going on in the server and up to that, you have something to base on it. Also, most of the specials are the same except if we're talking about Windows 2008 is the memory because on Windows 2008, we make change on the kernel space. It used to be kid's version. When you pull up all the PDE tools, the page pool have been set when the server pull up, but on Windows 2008, we have a dynamic memory manager that would dynamically allocate those components so hopefully with that we will allow in 2019 2020 running our systems page and on page pool memory and also on 64-bit, those are a couple of different scales, you're talking about the number. You can have much number on those system resources on it.

Richard Campbell: So, for the most part, these measurements come up the same way. I guess the challenge then -- doesn't it make sense to monitor the host machine on the virtualization side or is it all about the guest?

Stephen Choy: You need to monitor the host because the host is the host log, then your time will be definitely slow. So, you should monitor the host and also monitor individual virtual machines just as physical machines but then we add a couple of Hyper-V related content just to make sure that during the one-to-one mapping between physical and virtual, that it didn't get missing of the number of performance measure so we have something more than measure are similar except CPU and we get a couple in the network.

Richard Campbell: It seems to me that processors are not that big a problem when we're dealing with multiple virtuals because we have lots of processors and memory is pretty much fixed in its allocation so that's not that big a deal, but the real shared resources, the disk and the networking especially, I think networking is the one that's going to bite us where we got four or five VMs running in this machine and in essence they're all going across the same NIC.

Greg Hughes: Same interface.

Richard Campbell: Picking up the fact that you've saturated that NIC, anyone of those guests is nowhere near saturating the NIC, but the combination is now actually hurting the NIC.

Stephen Choy: True. So, you have a command, you have more than one NIC…

Richard Campbell: Right.

Stephen Choy: …on a host so that you won't get into the network bottleneck and also hopefully you use your switch so that you can stop your network chopping.

Richard Campbell: Well, I am seeing now more big virtual machines, big machines that are running lots of virtual servers where they are putting a dedicated NIC in almost per VM. It seems to be a strength of Hyper-V that Hyper-V actually allows the Virtual Machine to own a NIC.

Stephen Choy: Yeah, yeah.

Greg Hughes: Or dropping a fiber network connection and that's going to do super high bandwidth that you can actually split up a little bit if you don't need 100 gigabyte on each line.

Stephen Choy: True. I just want to mention because the purpose of using Hyper-V or virtualization is if you want to do some server consultation...

Richard Campbell: Right.

Stephen Choy: I went to a client once, and they want a couple of SQL Server or virtual machines and I told them, and also that SQL Server is heavily used, so I say it doesn't make sense why you want to pool on virtualized. They say, "Oh, there are so many directions." I'd say if your server is heavily used and you have thousands of transactions and people are hitting your server, it's better off as a standalone physical box.

Greg Hughes: Right.

Stephen Choy: You don't want to gain much because everything you serve the CPU and you serve the NIC, then virtually -- it's better off getting a physical box by itself.

Richard Campbell: Sure, I mean as soon as you're really straining the resources of the machine, it ought to be bare. You can't just virtualized -- well, we did a show on this, didn't we, Greg?

Greg Hughes: Yeah, we did.

Richard Campbell: Allan Hirt talking about virtualizing SQL Server. I think I started off with, "What are you, crazy?"

Greg Hughes: Exactly and in many cases, you are crazy I suppose, but there are cases where you can do it especially in like we said back then test environments or integration environments or what have you.

Richard Campbell: I'm interested in your opinion on this as well, Stephen. I also talked about web farms. As soon as I have multiple web servers, why would I virtualize those?

Stephen Choy: Yup. Well, if your web farm has a lot of topics, I think it's better off using a physical…

Richard Campbell: Yeah, be bare.

Stephen Choy: Because, otherwise, you just have caused more trouble.

Richard Campbell: I mean it just makes it harder to diagnose the problem if you're just adding this extra layer of complexity and for me by definition of web farm means bare. It's busy enough that I needed more than one web server so why would I virtualize that. Not that I'm down on virtualization. I'm virtualizing my brains out, but it's very true I hit the line on what should I virtualize and what shouldn't I.

Stephen Choy: Yup.

Greg Hughes: Yeah. There may be situations where if you have multiple web servers that are running in a virtual machine, but that's across multiple physical machines then you're basically, you know, you're trying to build your redundancy out that way and that's the reason for it as opposed to scale. I've seen that done before but I'm trying to think of other situations where you might think about virtualizing a web farm.

Richard Campbell: I don't know one and I hope I'm wrong. Maybe there's a listener out there who's like, "I know why you virtualize a web farm," I'd love to hear it.

Greg Hughes: Yeah, let us know. Just drop an email to info@runasradio.com. I'd be interested even just to find out what people are virtualizing and what their performance characteristics are on that.

Richard Campbell: For sure. So Steve, what kind of performance do we usually get moving from server 2003 to 2008? Is everything faster? Did stuff really change a lot? You mentioned the fact that 2008 manages memory differently, but are there other big things you've seen?

Stephen Choy: Other than that, all this performance specials are the same. Unless you move to 64-bit if you are still on 32-bit, it's the same as previous version.

Richard Campbell: Right.

Stephen Choy: The only thing we add is we have the view tool reliability and performance monitor that will help you to monitor the server and also I forgot to mention, the tool is called Reliability and Performance Monitor so they have a module called Reliability Monitor that they have a scale called the reliability index so if the scale is from 1 to 10, 10 is the most stable one and it was measured throughout the life of the computer, so when your machine is few, you will get 10 and then from then on, you're going to go down the hill hopefully not too much.

Richard Campbell: But it's a good idea that you set long-term viewpoint. Lots of people and certainly on the desktop side say, "My machine rots over time." I've often talked about it like software is fatty food and at some point your computer needs an angioplasty.

Stephen Choy: Yeah and also because they collect the data over a lifetime and also they amass the calculated index based on whether how often your server crash and the surface hangs and also including all the installation of any new application or any OS update so we'll say, "Hmm, my server was running fine last week and suddenly it doesn't work today." We can go back and see what had been done.

Richard Campbell: It helps you pinpoint when that happened.

Stephen Choy: Yup, yup.

Richard Campbell: I couldn't resist. While we're sitting here talking, I just popped open my reliability monitor on my big 64-bit Vista laptop.

Stephen Choy: What's your scale?

Richard Campbell: It's going up. I'm getting more reliable.

Stephen Choy: Good, good, good.

Richard Campbell: I think it's deluded.

Greg Hughes: The longer you pay attention to it, the more reliable it gets.

Richard Campbell: Yeah, something like that.

Greg Hughes: It's like a plant if you talk to it.

Richard Campbell: You know, it's interesting to think that yeah -- well, we are talking about Vista here, so I suspect that I'll find the reliability went up when SP1 went on and big driver patches too. That's the problem with being an early adopter here is we had bad drivers and as those drivers got better, reliability actually goes up but I think it's a great tool in terms of really understanding the stability of a machine over time to say, "Hey, there was clearly a breaking point here."

Greg Hughes: That's a good resource.

Stephen Choy: Yup.

Greg Hughes: Steve, I'm curious, for IT people out there who have not yet taken the deep dive into performance tuning or even performance monitoring and don't even really know where to start, maybe you have a top three or top couple of resources that you recommend people go to, to get a start on doing this.

Stephen Choy: There's a lot of information on TechNet, both TechNet Magazine and TechNet online. They have a lot of articles talking about -- I recently wrote an article on keeping your servers, how you watch content and you look at. It's not intended to be detailed. It's just written from the point of IT as a means so when they have someone calling them, they have problems with applications, those are the content they should check.

Greg Hughes: Got you.

Stephen Choy: And also, if anyone is interested on learning more techniques on server performance, we at Microsoft in support offers a workshop called Y design workshop.

Greg Hughes: Really?

Stephen Choy: A two-day workshop, yup, two-day workshop. They go for all the techniques, how to tune it and make it much easier to look at the graph and go to other major components, go deep dive on disk memory process, CPU, and network, what you should look for, and they also taught more details on Windows 2008. We cover everything so just in case you haven't moved on to Windows 2008, we still have the tools to tackle the issue.

Greg Hughes: That's great.

Richard Campbell: Where do these workshops take place?

Stephen Choy: Most of the workshops take place at a client's site. If you're a Microsoft client, try to talk to the TAM, technical account manager. They can arrange it either on-site or we offer the class on major training centers like Charlotte, Texas, or sometimes at Redmond.

Greg Hughes: Got you. So it's available to premiere support customers but is also available at your major training centers for people maybe that aren't on a premiere support contract.

Stephen Choy: I don't think they mean support contract. They can go over there when they offer the class or they can even come to the site, the customer, to make any particular information that you want to talk more or you want to remove some of the information you already know about.

Greg Hughes: Great. If you don't have technical account manager, is there a library source or a place where you can go to find out where those classes are being offered?

Stephen Choy: I can check it for you. I mean anyone who is interested, they can send me an email at schoy@microsoft.com. I can give them more information.

Greg Hughes: Sure.

Stephen Choy: And also, my co-worker Christopher Vang is working on his paper, a very detailed paper on Performance Monitor in Windows 2008 Best Practice. I don't know how many pages. He is happy working on it. Hopefully, he can publish it before end of this year. They will go through what's the difference between 2008 and 2003 and what other tools that they can use to monitor Windows 2008.

Richard Campbell: I hope he can publish it as a book, you know, that's one of those books I've always wanted to have is what are all these points of measure, what do they mean. If I'm looking for this problem, here are the things I should be watching and so on.

Stephen Choy: Yup, yup and he also goes through details on different server roles, on different server roles, what concept that you should pay more attention.

Richard Campbell: Right.

Stephen Choy: For example, web server, you should look at more cache, also the network and disk compared to the main controller because of different nodes so you should pay more attention on different content so when you have someone telling you, "Oh, the server is not working. Oh, it's busy. What should

I put on my performance content?" It's memory processor, it's the network, and you don't put cache because they see it doesn't do much on cache.

Richard Campbell: Right, the numbers of times I've just gone and grabbed a bunch of measurements off a PerfMon and they were all zero. Nothing will stop you there. You're looking in the wrong place. The machine is busy but that's not where it's busy.

Greg Hughes: How do you know if you're looking in the right place, that's the book that we need. I've always thought Mark Russinovich's case of presentations and articles that he writes where, you know, it's usually on a client operating system, there's something specific and he uses PerfMon and the sys internal tools and different stuff to chase it down and kind of drill down. So something like that to help on the server side would be terrific.

Stephen Choy: If the problem happened right now, I think the quick thing is you want to bring up the task manager so I want to point out that on the new 2008, the task manager, they have one additional pack called server. In the previous version, if you bring up task manager, you will look at the process pack as we see whole. We may see multiple with them, one of them driving all the resources and you don't know what is the root cause.

Greg Hughes: Yeah, t hey're all hidden in there.

Stephen Choy: But with the new one, we actually put a service pack. We can actually nail down, "Oh, that is the guy eating up all my resources."

Richard Campbell: Yeah, we're actually finally able to walk through the hierarchy to get to here's the bad boy and the other way too, you could be looking down at the service level and say what app is this and it will help to chase it up.

Stephen Choy: Yup, yup. So you don't need to use a process for it to go deep into what actually is the server using so I can incorporate into the top manager. Of course, for more compact and you build a monitor, you should be using the built in performance monitor that comes with Windows 2008.

Richard Campbell: Sure. Stephen Choy, we are out of time, just like I told you we'd end up. We go in here and boom, the time flies by.

Greg Hughes: Thanks a bunch.

Stephen Choy: It's a pleasure talking to you.

Richard Campbell: We really enjoyed talking to you. And we'll talk to you next week on RunAs Radio.