



<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 #081
(Transcription services provided by [PWOP Productions](#))



Shane Creamer Goes Deep on Performance Monitor!
October 29, 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 #81, with guest Shane Creamer, recorded Thursday, October 16, 2008. RunAs Radio is produced each week by PWOP Productions, providing professional media and podcasting services online at pwop.com.

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

Greg Hughes: That's me.

Richard Campbell: Sir, how are you?

Greg Hughes: I'm doing pretty well. How about yourself?

Richard Campbell: Great and odds are when people listen to this show, we are not home.

Greg Hughes: Or at least you're not home.

Richard Campbell: Well, yeah, might be PDC, might be DevConnections, might be TechEd. You know, this is the season of much travel.

Greg Hughes: That's right. It's confusing enough just thinking about where we have to travel, when do we start adding in time shift that shows and trying to map those up in your brain. It gets really confusing, doesn't it?

Richard Campbell: Well, what I'm excited about is when we're together, of course we get the opportunity to do a bunch of shows in person...

Greg Hughes: Right.

Richard Campbell: Sometimes with an audience, sometimes in private but I like those face-to-face shows. They're a lot of fun to do.

Greg Hughes: Yeah. You know, it's a ton of fun doing the shows the way we usually do it which is I have studio space in my home office, you have your studio space and we're talking to somebody...

Richard Campbell: Somewhere.

Greg Hughes: Typically over the phone, that kind of thing, or over the net, their face-to-face. It's really cool to be able to sit there and have that type of interaction too and then share that with the people that listen to our show.

Richard Campbell: And the whole life feel of the buzz of the conference is just -- those shows are interesting too. I'm looking forward to the interviews we're going to get.

Greg Hughes: Well, we always worry about the background noise, don't we? But it ends up -- I actually kind of like it a little bit...

Richard Campbell: Absolutely.

Greg Hughes: It feels a little different and the conversations seem to always turn to be really, really interesting.

Richard Campbell: Absolutely. It's going to be fun. I think we're going to get a lot of shows. Between Barcelona and DevConnections, lots and lots of interesting people to talk to.

Greg Hughes: Yup and of course in Barcelona, we've got Speaker Idol going on and a bunch of other stuff and so anybody who is listening, if you'll be at the TechEd IT week or the IT forum week in Barcelona at the EMEA conference, then be sure to stop by and say hi, we'll be glad to see you.

Richard Campbell: Yeah. We're right in the middle of the venter floor at the community stage. All right, let's introduce our guest. Shane Creamer started working with Windows in the early 1990's. He's worked at Microsoft since 1995 and currently works in the Microsoft Premier Field Engineering as a Performance and Debug engineer. He's filled past roles as a customer advocate program manager and was an extended part of the Operating System team that helped ship Windows 2000, way back when. Shane's passion has been serving as a customer advocate to help the Microsoft community self-identify and self-solve many of the typical stability and performance issues that occur in the murky, gray area known as Performance Analysis.

Greg Hughes: Very murky and gray huh.

Richard Campbell: Since a "Dummies Guide to Using PerfMon Effectively" does not exist, Shane has been creating the Vital Signs workshop which takes premier customers on a three-day experience where they learn the proper use of the Performance Analysis tools like Performance Monitor and Server Performance Advisor, and then 25 counters that are on the must know list of counters needed to properly address performance and stability issues regardless of what role the server has or what process it is running. This workshop has been delivered to over 250 customers worldwide. Shane's hobbies are home theaters, exercising, health, doing charitable work, helping parents of Autistic children on the ASD spectrum, his son is one of them on the spectrum,



and he has trained as an emotional counselor, that's kind of come in handy for performance problems...

Greg Hughes: Oh cool.

Richard Campbell: Helping people and couples rebuild their lives when he isn't rebuilding server performance. He currently lives in Eastern Washington and loves spending time with his family and watching his son become more functionally normal each month, thanks to the latest ASD spectrum research. Shane Creamer can be reached at shanec@microsoft.com. Welcome Shane Creamer.

Greg Hughes: Hi, Shane.

Richard Campbell: Nice. That's quite an intro actually.

Greg Hughes: It's a cool intro.

Richard Campbell: Yeah, some challenging parts of your life, no two ways about it. Server performance is the easy part I mentioned.

Shane Creamer: You know I like helping people and it's one of the reasons why I left the product group and decided to come back over two disports out of the house because one of the things you can learn to be is a teaching search engine...

Richard Campbell: Absolutely.

Greg Hughes: Right.

Shane Creamer: And I just love it.

Richard Campbell: And you know an interesting part, I think a lot of folks don't understand the sophistication level of the support available from Microsoft.

Greg Hughes: That's terrific stuff.

Richard Campbell: Its not just phone calls anymore.

Shane Creamer: No sir.

Greg Hughes: I know I've leverage Microsoft premiere support in some pretty hairy situations and also in some terrific learning situations so learned at what you do whether it's going some place to get that information or if it's on the phone or if it's premiere support coming to our location to present it. It has always been -- and a lot of people don't even realize that it's available but there is a lot of great content and a lot of very, very smart people with terrific knowledge to tap into.

Shane Creamer: You know, I couldn't agree more. One of the things that I tell my customers every time I engage them on, and it's the same with you, and that is there are brilliant people on both sides of the phone and the ideas are simply, you know, instead of looking across the table at each other, just kind of look down the road together and say, "You know, let's go where the data leads us. You got extreme knowledge of the environment, I have some good tools, let's put it to use and let's see where we can go for the fix"

Greg Hughes: Right, exactly. So I guess maybe a quick question before we jump into our topic for today. If people are not very aware, business people are not very aware of what's available to them in their premiere support area or other areas, who is the right person for them to contact at Microsoft? Who is their point person to get more information and to find out what they can take advantage of?

Shane Creamer: There are two layers to that answer. The first one is going to sound really good. The second one is going to be like I'm going to have to find out. If a customer has some kind of an engagement scenario already like they have a professional level support, they have a premiere agreement, then they'll have whenever they join that contract with Microsoft and Support Relationship, they're made aware of key contact that they can use and leverage. If we're talking about a customer out there like, you know, I have a certain number of machines in my environment and I have a no support relationship with Microsoft, I'd be interested in getting one. You know honestly to my embarrassment, I don't know but you know something, I'm willing to find out because that's a good question. What if I have a company and my business has been growing pretty well and suddenly I realize I could use a support relationship. To much of my embarrassment, I don't know.

Greg Hughes: If I have a support relationship, it's probably my technical account manager or somebody like that that I will call to do that kind of thing; or maybe if not, then maybe the local office or the local sales office is a good place to start but yeah, that would be in fact I think, Richard, it would be interesting just to have a conversation about resources available to customers. That might be a great show to do someday.

Richard Campbell: Absolutely. We only have half an hour today and...

Greg Hughes: Not today, yeah.

Richard Campbell: I knew we want to drill into this pretty deeply. This is a workshop; you've been doing



a workshop around Performance Analysis in PerfMon and so forth?

Shane Creamer: That's correct. It started off around 2004. One of the things that I intend to find is, in fact this might surprise you, but one of the things we do at Microsoft is every time we reach closure, in other words we talk with a customer and we say have we achieve root cause, have we got you a solution whether it's a bug fix or a work around that you agree that this is issue is resolved, this fire is put out; and when they say yes, we can actually code the issue. So we have basically kind of a call entry much like, you know, kind of like your file folder structure in Windows. We might have a folder and then folder one, folder two, folder three underneath and so for example what if the issue was the OS and under the OS we say it was a kernel pool. It's not just a kernel pool; it's a kernel pool page. Under kernel pool page, there was a third party driver that leak it basically bled the pool dry so when you run hundreds of thousands of curls through that, you can have a very accurate picture, a highly accurate picture of what customers are escalating to us.

Richard Campbell: Right.

Shane Creamer: Now obviously, many, many intelligent people like you can stamp out the day-to-day issues so what we tend to see are the nasty of the nastiest that are escalated thus and it should be, and what we tend to find is what if I told you that 50% of all support calls as of October 2007 that come into Microsoft are actually storage issues, they're issues with their disk subsystem.

Richard Campbell: Wow.

Shane Creamer: Which means in a sense, we are -- and this doesn't bother us at all that you know, 50% of our calls basically where we're acting as frontline support you can almost assert four storage vendors and that is okay because customers say their servers are slow, they want to talk with Microsoft because it's the logical place to start especially we got a good engineer and so one of the things that I do for the customers that I engage with is as a performance and debugging chair, I say if you're servers is slow, let's talk a look and so PerfMon is one of the tools that, unfortunately, you know, it's kind of there in the '90s and I think as we get into the 21st century, a lot of people including myself, as your experience grows, you prefer to differ tools and so sometimes PerfMon kind of get sit on the shelf as kind of this tool is kind of there, kind of effective, kind of not and what we really wanted to do is make sure that customers know. In the hands of somebody who is improperly trained, this book has some amazing details without having to crawl inside of a process and create a log that's gigabytes in size.

Richard Campbell: Right.

Greg Hughes: So operative words, their properly trained. It's kind of Greek until you really know what you're looking at.

Shane Creamer: Agreed.

Richard Campbell: And I find a lot of folks that are very frustrated with PerfMon that they don't know it's like your own little internet, everything you want to find is in there, it just don't know where it is.

Shane Creamer: And to my embarrassment, when you love your company, I don't mean to just drink the Kool-Aid to drink it cool, but I mean when you find that you're a part of a company that tries really good things especially in the engineering group, the dev teams do great stuff but sometimes it's either that they miss with the right features and one of the things that we tend to find is usability is really tough to get in front of developers insisting on shipping a product and all companies need to ship a product so I don't want to cast aspersions this version but for example PerfMon, where is the dummies guide to PerfMon?

Richard Campbell: Yeah.

Greg Hughes: Yeah.

Shane Creamer: Where is this -- if you look, ironically enough if you step back and think and look at all the PerfMon articles out there and even look at the rest of the documentation we create, there is a built-in assumption that you already know what you're doing.

Richard Campbell: Yes.

Shane Creamer: And so one of the things that I love to do is take very intelligent people that have deep skills in SQL and Exchange and IIS and certificate services and active directory and take them. We take the first day, the first day is just basically introducing them to the tune how to use them right and to watch these people go from junior woodshop. The Jedi path to unlearn is to Jedi masters in a couple of weeks is awesome. I love getting the emails that says, "Shane, we have this issue that's plaguing for six months. Thanks to your workshop, here's my analysis, here's the screenshot, what do you think?" I just rubber stamp and I go I couldn't have done a better job myself.

Greg Hughes: So one of the things that PerfMon is it's very, very visual. So part of what we're doing today is you are going to be as we're talking sharing a screen with us which we will then



make available to our viewers to also be able to view, to take a look at as you go through it.

Shane Creamer: Absolutely.

Greg Hughes: So there are no dummies, yellow dummy guides if you will trademark -- it's probably a trademark of some company that I should say right now, but your class that you've taught to some of your premiere customers and other customers really is sort of that what does it take to really make it work for you. Correct?

Shane Creamer: It is a heroic attempt to fill that space until that book becomes available, yes.

Greg Hughes: The reviews I've read have been terrific so we're looking forward to jumping into it.

Shane Creamer: Okay. So at this point, would you like me to go ahead and take your listeners on a little journey about good PerfMon techniques?

Richard Campbell: Absolutely.

Shane Creamer: Okay. So for the listeners on the radio, what we have in front of us here is PerfMon and we're using the -- there are two quest to PerfMon. There's the Windows 2000, Windows XP Server 2003 version of PerfMon and then there's the Windows Server 2008 and Vista version of PerfMon which is what I have in the screen today and one of the benefits about it if you have ever been in PerfMon is: on the bottom of the screen you've got the time stamp built in, you can hover your mouse over something and get some real-time information. So for example like if I kind of tooltip this a little bit, notice right there on the screen I can see it right now, that value is four. There is no more guessing. One of the things, if anybody has used PerfMon, it's both an art and a science.

Richard Campbell: Sure.

Shane Creamer: And what we look for is relationship so right off the get-go we tend to tell our -- some of the real focus on the top potholes that we see people are using PerfMon, good intelligent people using PerfMon having a go at it without much training and the make and so that's my goal today to make sure that you steer clear of that if possible. So anytime you look at the bottom of the graph where I have my mouse right now, there are values, there's minimum and there's maximum values and there's an average...

Richard Campbell: Right.

Shane Creamer: And one of the things we tell our customers right off the get-go is you focus on the minimums and the maximums. If you have a three-day capture where you've got some line at the very bottom of the graph and it's really doing nothing, say it just transfers for a second and we're really doing nothing on a disk and then for 5 or 10 minutes, we do three to four minutes of heavy, heavy disk transfer to the point where we see massive latency, it doesn't do any good to core an average there because what we care about is that 5 to 10 minutes where, you know, like in Exchange server, servicing 5,000 mail boxes which were just simply unresponsive.

Richard Campbell: Yeah, we just paralyze them, it will load. In the end it's always peak moments that matter.

Greg Hughes: Yup.

Shane Creamer: Exactly. So the first thing that we always try to tell our students is focus on the peaks in the valleys. The averages will take care of themselves as most of people can kind of eyeball. I mean if you look at your checkbook and you put your check and then you saw or imagine it dwindling, you can kind of in your mind kind of imagine the graph. So we're pretty good as human beings in coming up with our own mental averages. It's the peaks and the throes really where the action is. The second part of that is to give peaks and throes contexts. You know, we always want to make sure we know what time the issue occurred. So what I have here is this is an active directory domain controller that was escalated to Microsoft as a very high priority. Now this customer is intelligent but they had really no PerfMon experience so most of us can go to Task Manager and see what our CPU utilization is but how many people are really willing to stare at it for hours to days. So obviously we need to record that so that we have a running log of what is occurring. What you see here is that the very time they reported this problem around 8:40 and they are a part of this problem around 8:50, for a couple of minutes the server went offline. So what I'm going to do here in PerfMon 2.0 and I like to call which is, you know, Vista and Search 2008 PerfMon. You can take a counter here, right click it and choose scale and scripting it automatically scales for you which is huge time savings because the second problem is that if you're going to look at a number in the graph, having it in the middle of the graph is important. If it's too squished down, it looks benign and if it's so high that as a graph of 0-100, it's slammed against the top, you really can't see what's going on. So here we have this, so for the listeners if they're only listening to the audio portion only, what you have here is at the very time that the helpdesk staff of coders and issues, we can see the memory goes from about 424 megs all the way down to zero so...



Richard Campbell: That's really bad. That's really bad.

Greg Hughes: That's the available memory, right.

Shane Creamer: And so not only that, but what are the things we can see here because we have the graph appropriately. There is this nasty little downward trend here and that's kind of indicative. When you have unremitting allocation and no de-allocation of releasing of resources, we start to -- those who have been in the PerfMon, they will start to sniff a leak line.

Richard Campbell: Everyone looking at the line of available memory that's gradually going down, that can't be good.

Greg Hughes: Yeah.

Shane Creamer: That can't be good. Now the next thing that's a perfect lead in is the next we want to look at. We always tell our PerfMon users or PerfMon students to keep an eye on duration. Now notice this is between 8:26 a.m. in the graph now and 9:02. This is a 30-minute capture so one could assert that this could simply be, as the peak user load comes on, naturally resources are going to decline because active directory or Exchange or any product is going to use allocate more resources to satisfy the needs of the user so if this is a two-day capture; we might see it come back up. So the idea is that we let duration be our friend to say do I have like a two-day capture here where I can see the rise and fall of what the server really looks like or am I just looking at just enough of the graph where I could theoretically be head faked into thinking that something that it isn't. So we let duration be our friend.

Richard Campbell: Right.

Greg Hughes: Right, right.

Richard Campbell: When you talk about taking a recording for two days, am I just saving this to disk? What is the correct way to say this out; this is a file to be use as a SQL server? Does any of that matter?

Shane Creamer: You know, it's an excellent question. So what we're talking about when we talk about saving something out to disk is you fire a PerfMon and you start to capture and obviously in PerfMon you can create something called a template where kind of like ordering food in a restaurant, you look at the menu and you select counters that you want to record and by recording those counters, you can choose then where you want to save it. I'm obviously saying this for those who are kind of like

looking at PerfMon the first time really. You can choose objects. So just to give an example, here on the screen, I have lots of objects here that were recorded on the server: memory, logical disk. So we can choose objects and then we can have the server record those at interval. Typically, we use five minutes and if you use the right counters, we don't graph big objects like the thread object and so with this developer objects. You can make a PerfMon sample like every five minutes only graphs 64 to 12pk of data. So if you're graphing every five minutes, you can have a capture for a couple of weeks and be less than 50 megs.

Richard Campbell: Nice.

Greg Hughes: Sure.

Richard Campbell: What about the observer effect. How much impact does PerfMon have on the machine that you're observing?

Shane Creamer: Excellent question. The Heisenberg Certainty Principle, you know.

Richard Campbell: Right.

Shane Creamer: By looking at it, do we induce more load? The good news is that -- well, there is good and bad news. If PerfMon has, like you said it's like the internet. I have everything I want there if you're a developer and you're looking at your application thread, you're going to use the thread object but the thread object, if you got 244,000 threads running on an active 64-bit server, just running the thread object means you're capturing data on all 244,000 threads so you just turn PerfMon into SQL server.

Richard Campbell: Yeah.

Shane Creamer: But on the other hand, one of the things that we teach in terms of about 25 counters that are like dog counters that developers uses, that we use at Microsoft, they are like the key counters, it's kind of like building a house, it's your foundation, you can add those counters to it but you never want to leave this behind.

Richard Campbell: Right.

Shane Creamer: And with those 25 counters, you're talking about less than one percent hit to a CPU at a time of sample and about 64 to 128k of data per sample intervals. So you can argue that a PC400 from the 20th century with 500-megs of RAM running some web page could easily handle that and if a server that open handles it, obviously our modern 20th century, you know 2003, 2008 hardworking can handle that without a hitch properly tuned.



Richard Campbell: I've heard the mention of these 25 great counters.

Greg Hughes: And 10 great ones for IT and 25 for devs, is that what I heard too?

Shane Creamer: You know, it's 25 that we use for everybody so whether you're in development or you're an IT professional managing a server environment, there are some counters that we have found that are must haves and what I've got here in front of us on the screen is kind of an excerpt from the vital signs workshop and what these are are these are the counters and the numbers. One of the things that I'm going to do here is that I want to make that PDF, that PDF is publicly available. In other words, any customer is allowed to have this document so you guys have the PDF here, welcome to share it with this presentation because we want everybody, we want to create cart carrying members of the PerfMon army.

Richard Campbell: You bet.

Shane Creamer: And so to do that, we need to equip our soldiers.

Richard Campbell: Yes. So we will definitely provide a link to this PDF file so that everybody can get a good look at this.

Shane Creamer: Because it got the threshold and it's got the counters and that's all we need. We want to know what the counters are and what are some truly crisp thresholds that we can leverage so that we know if something crosses this, then the environment is mainly inducing latency on top of a critical application.

Richard Campbell: Yeah and I think that's one of the challenges for folks that are going to PerfMon new, it's just not okay I found a number that seems to be meaningful. It moves in relation to my problem, what's good and what's bad.

Shane Creamer: Exactly.

Richard Campbell: So when you're talking about anything under 50% dial time on the physical disk is a concern. Now we start watching that.

Shane Creamer: Right and percent basically is our favorite version of average or current Disk Queue length.

Richard Campbell: Rather than looking because you can't look at the Disk Queue length directly, right.

Shane Creamer: Well, you can and you can't. One of the things that we also talk about the

workshop is that most of these PerfMon counters were created in the '90s when everything was direct attached to the server.

Richard Campbell: Right.

Shane Creamer: Dave Cutler and his team did a fantastic -- I mean it blows my mind to realize that the counters that I use today were created back in NT3.1 when these counters were around.

Richard Campbell: Right.

Shane Creamer: So that's pretty thoughtful foresight, I mean, you know, when he created VAX and DMS, when we hired him, we didn't mess around, we got the right guy.

Richard Campbell: You got the guy.

Shane Creamer: Yeah. So these counters are still useful today but some of them we had to retire. For example, percent disk time. Percent disk time is a value where you're supposed to know the exclusive number of dedicated hard disks so if you got a value of 666, if you don't know how many exclusively dedicated hard drives you have to divide that so you get a number between 0-100, you can't do the math.

Richard Campbell: Right.

Shane Creamer: So the other thing is that why I really had this burden for our larger Microsoft community out there is to, you know, here are the counters that are useful, there are other counters that you can kind of land on top of this like a tapestry but this is the ones that we use all the time. So to count the king of disk as far as we're concern is average disk second read or write. That's very much like a SAN person, you know, taking a look at the back of the SAN frame which is what is the raw latency to and from the disk.

Richard Campbell: Right and it doesn't matter whether they are direct attached storage going through some kind of a rate controller or iSCSI or some LUN-based SAN, that number is still meaningful.

Shane Creamer: That's right. If Wile E Coyote falls off the cliff, it doesn't matter if he falls off by a foot or a hundred feet. It has the good news that most of these are precipices. For example, if I was to -- so for those of you that are on the radio show that might be tempted, so in the video, what I'm doing is I'm scrolling down to the memory object which is really kernel memory and one of the things we tend to find here is free system paste table entries. If you go underneath 5,000 paste table entries, that's when a server will start to go really dumb, it will go



Shane Creamer Goes Deep on Performance Monitor!

October 29, 2008

unresponsive, it will throw some really weird errors, it's like if your checking account runs out of ledger in your checks and you keep writing checks, you lose track of who you've given your resources to and the OS has the same problem.

Richard Campbell: Right. So once you're running low, again I got to think a great gauge no matter what you're running of just is this machine sick.

Shane Creamer: That's right and that's why we coin a very corny phrase vital signs because when you walk to a corner's office and you're not feeling well, the first thing they want to check are the vitals.

Richard Campbell: Yeah.

Shane Creamer: Because that will give them a huge clue.

Richard Campbell: What about generally free memory, its available megabytes.

Shane Creamer: You know, available megabytes, what we have here is, so most people understands if you got 2 gigs for your server, you're doing pretty good.

Richard Campbell: Right.

Shane Creamer: The question starts to come, what happens if my server has a 100 megs or 250 megs and that's really where we have two, really we have two positions. One is it is theoretically possible to run the Vista on a 1998 era of P2-400 laptop with 256 or 512 megs of RAM. You might not enjoy the experience but it will functionally come up for you and we tend to find out the same thing. Most of Microsoft white papers really without saying it say how much memory is necessary to keep the server alive and out of trouble. So you might say things like at least 10 megabytes is free. Of course we all know, let's say we have an Exchange server or a SQL server, a big iron server, even an Oracle server, I mean something in the Microsoft chronology and its back up time and we're not going to back up a 36 terabyte database. Ten megs isn't good a cutter for buffering that and that's the difference between backing up in two hours and backing up in 20.

Richard Campbell: Yeah.

Shane Creamer: So the idea there is that is why our ratio is that hey if you got 20, if you got less than 10% of free RAM available, that can be a warning, and that apply even to a 64-gig server. Let say you have a 64-gig server and it's running SQL and SQL has taken 60 gigs of it because SQL can pin its memory, that leaves 4 gigs for the OS. So here comes a 36 terabyte database so we're going to back up the OS,

we'll use every byte available as kind of buffering the file of activities so you can get the maximum true put out of your disk so in that case, you really want to watch your server under six gigs free, even 6-gig is a big number, store big data on it and the OS will take and go wanting all of it.

Richard Campbell: Yeah. That makes sense to me. It's not just the actual number but the percentage of overall memory.

Shane Creamer: Absolutely. So what we try to say is that most of our white papers out there are good quality papers and what they really want to communicate is what's necessary to keep the patient alive, not to make a drugstore and I think most of us don't want to just keep the server alive, we really want the server to run very fast, as fast as...

Greg Hughes: Right.

Shane Creamer: And so here's kind of a performance setting so that's where we really talk about having like a gig or two or more free at the time where you can do a big eye operation and then of course here is the minimum bar which is what's necessary to keep the server alive. If I put you underneath the water and I give you a straw, depending on what it describes whether or not you will be able to breathe or solely asphyxiate.

Richard Campbell: Yes and you're still alive but you are going down. Let's keep moving along here. Beyond memory means there's a lot of different values here. We talked a little bit about disk, talk little bit about memory, what about processor? How important is processor in this?

Shane Creamer: You know, processor, an interesting thing occurred. Now, this is where we're going to run on just a little bit geeky but I figured geeks like your show.

Richard Campbell: Yeah, I think so.

Shane Creamer: I'm a geek and I looked at some of the authors and some of the people you've interviewed and that's like fantastic so, you know, if they can geek, I can geek so I'm spinning my propeller hat as we speak. So what we have here at processors is in Server 2003, an interesting change has made at the kernel level and this is where once again Landy Wang and Dave Cutler and a lot of the guys at the base OS team, you know, some of the coolest stuff we do and never gets advertised, I mean where is the advertisement: hey, guess what we have done in kernel, I mean that doesn't quite play to the average store front sign but for those of us in the IT community, some of these changes are huge and they should have like IT geek signs. Guess what, in



kernel, we now make processor time less worrisome than it was before. Hey, that's the feature I'm interested in...

Richard Campbell: Absolutely.

Shane Creamer: So in Windows 2000 NT 4.0, one of the things we always wanted to see was if you had say, 4 processors, you want the all four processors running at the same level.

Richard Campbell: Right.

Shane Creamer: You wanted to see the load evenly dispersed.

Richard Campbell: Even though we pretty much knew that couldn't happen, that's really what we wanted but you know it's just given work clothes, run on given processors and sometimes some of them are rougher than others. You know what frustrates you is I've got four cores, I've got eight cores, two of them are pin and the rest are sleeping.

Shane Creamer: And why is that?

Richard Campbell: Yeah, why is that?

Shane Creamer: And here comes the answer. So when Windows 2000 NT 4.0, 3.5, and 3.5.1, it was very important that what the system, the trade manager basically did is it looked at it and said okay. You know I want to make sure that all threads are getting evenly dispersed across all the processors because we thought at that time that was a good thing. Well, the base at West team which is like your deep, deep guys, you know the guys who work at the house, these are the guys that breathe assembly for a living. it's kind of like the metrics, you know they always look at things in code raw absolutely. These are the same people.

Richard Campbell: Right.

Shane Creamer: So they are like half human, half cyborg. You have to look in their eyes whether or not you're talking to a computer at that moment or you're talking to a person.

Richard Campbell: Give them time to come out of it.

Shane Creamer: They're the best of the best. I need a human response here, don't give me an assembly language now, and what they did is they realized they got really close in the hardware. They work with the guys at AMD and Intel and they realized that a thread is like the atomic coin of the realm, that's the smallest unit of work that DOS will do, it's to operate on thread, the same had some instructions

we call function calls, some registers punches them at the CPU and says run mate. What we found is that if you move a thread from CPU to another, there is a hidden secret on hardware called the messy bus and it's kind of messy...

Richard Campbell: Nice.

Shane Creamer: But what happens is that say you have two processors. If you move that thread from processor 1 to processor 2, it has to stop both processors because you have to copy over the L1 cache, you have to copy over the L2 cache, so the more processors you have, the worse it is to context switch. So if you've got 16 processors, all 16 processors have to stop if CPU 1 wants to transfer its thread to CPU 7 and so that's kind of the hidden secret. So the more CPUs to have, the more horrendous the overhead becomes. So we realized wow, another one moving in the multiple cores, this is going to kill customers and this is going to kill productivity, I mean your CPUs are busy but yet they are not doing as much as they could. So what they decided to do is that the OS determines if this thread really likes to run here on the CPU. Unless the CPU is 80% or 90% busy, I'm not going to move it because in the multi-core environment, it's actually better to have like three to four CPUs busy and three or four CPUs not because it saves on ring transitions...

Richard Campbell: Right.

Shane Creamer: And it saves all these prime and so there are some hidden reasons that accident makes a good decision where it's actually better for you to have your CPUs running at different rates.

Greg Hughes: Very interesting.

Richard Campbell: I remember reading about this and being excited about this so I guess I qualify as a geek too.

Shane Creamer: Awesome. You guys are obviously, when I read your descriptions, you guys are geeks deluxe so you know I'm an aspiring geek so knowing that my information is welcome and I'm in a good company makes me feel good, makes me feel like I'm on the right path.

Greg Hughes: The PerfMon ninjas are aspiring geeks. That's cool.

Richard Campbell: Nice.

Shane Creamer: That's right. So obviously, percent processor becomes important but it's not so important now that especially in the era of multi-cores. So if we see a couple of CPUs that are rather busy, let's not forget about driver affinity. You know, one of



Shane Creamer Goes Deep on Performance Monitor!

October 29, 2008

the things that have been around for years is HP, Intel, and Emulex, the people that makes these controllers, HBA controller and for those who are like what's an HBA, there goes Microsoft not to finding an acronym again.

Richard Campbell: Yeah.

Shane Creamer: It's this SCSI controller for the sand.

Richard Campbell: Right.

Shane Creamer: It kind of have ATM network card and has a SCSI controller and what they do is they write their drivers to take advantage of the last two CPUs, the last one of two CPUs because they know the OS wants to run most of the application threads on the first couple of CPUs. So percent processors are important to understand the relationship but as long as we don't see a CPU tapped out at 90% or more, that's important. The other part I like to bring in there is that there are two more counters that we're going to be adding to this document which is percent privilege time which represents kernel mode and percent user time which represents user modes. So what does that have to do with anything? Good question. What it means is what if I have an application like SQL or Exchange or IIS and I see my CPU at 50% and I see 80% of it is user mode and 20% of it is kernel. What happens if I see my CPU at a 100% and 99% of it is kernel? What that tells an IT professional that something has gone wrong, something has been lacking in kernel because that's not how my server usually look and so we can look for patterns of behavior.

Greg Hughes: Very cool.

Richard Campbell: Absolutely.

Greg Hughes: One thing that -- not to get off topic, but this comes to mind that and I'm thinking of this as we're talking right now -- first of all this is terrific and great information and it's very cool that you can basically, tell me if I'm wrong, save a template which you can then reuse in performance monitor to do the same type of analysis over and over again, right?

Shane Creamer: Absolutely. In fact, one of the things we teach in the workshop that, you know, like I said if you guys like this kind of information, we can make this kind of like a little weekly show like we'll talk about disk at some point, tips on disk, best tips on processor and user mode and what not, but yes, we can actually talk about performance fundamentals which is how to take a bunch of counters and save them out as a template and then script it out where you can punch that in to a hundred of the servers and

not only that, how about creating a black box cache that is always running even if the server...

Greg Hughes: There you go.

Shane Creamer: Blue screens or it has to get rebooted because it's security patch Tuesday, how about always having birth Monday because one of the things we all know is how many times have servers had an issue and we have to say do we have any data, no, well, then we have to set up a capture. How about always having something running 250, maybe 500 megs in size, it's circular like an event log so that if the wire ever trips, we've got data right there. Those are things we love to teach your users out there how to do.

Greg Hughes: That's interesting because that's kind of where I was going, it was how can we make this a little more modular, a little more automatic and certainly the templating is a great way but that you hit the nail right on the head, that the nail allows thinking off.

Shane Creamer: Absolutely and like I said, sure this is a three-day workshop but there are pieces here I'd like everybody to know.

Richard Campbell: Absolutely and maybe we need to do more shows.

Shane Creamer: Yeah.

Richard Campbell: More shows, we need to do a PerfMon series. It's just such a huge topic. You know, we stepped on this or stepped into this with Stephen Choy a month ago or so and...

Shane Creamer: Yeah, he is a great guy, great guy.

Richard Campbell: Also a part of the previous support team and just sort of realized this is such a huge topic so I'd like to get some feedback from the listeners. Do we need to dive deeper into this? I think we've given them some key reads, made a little video here to see a bit of what's going on, the PDF file, the sort of key measures in it and we haven't even gotten through all of this network interface, there are some other important bits of what you got to measure on every machine. It's got to be a challenge to see how can we break this down into individual elements. I think we could just sit and spend half an hour on disk access, what to measure and why, what good, what does it look like embed?

Greg Hughes: You're right though at the beginning of the show, you know trying to develop a PerfMon army. This is all information that people can use now but also could have been using for the past



Shane Creamer Goes Deep on Performance Monitor!

October 29, 2008

many, many years and it's a pretty small percentage of folks out there that really know how to use PerfMon.

Shane Creamer: And it shouldn't have to be that way.

Richard Campbell: Yeah. I think it's a great idea and interesting to teach. Maybe you could just go over the fundamentals of so I want to leave PerfMon running on my server all the time but not kill myself. I now have this PDF file on my hands so I know what I should be measuring. It's always frustrating to me that we have a failure and that's when we go instrument and wait for the next failure. So I want to be able to have that instrumentation ahead of time. What's the trick to have it get start on boot and always be running?

Shane Creamer: Well, let me bring up a little something here from the workshop. So if I'm come and bring up, I go to a student, I go to Module 3 PerfMon techniques. Now what I'm doing here is I'm bringing up a PDF from the training. So this will be in the video for those of you but otherwise we'll do our very best to make sure we describe this as a word picture for the folks that are listening. So what we essentially do, it involves two things. One is using a logman. So a logman is kind of the PerfMon for command lines.

Richard Campbell: Right.

Shane Creamer: So if you want to do something on PerfMon, you know, command line version, you're going to use what's called a logman. What we really do here is you use logman and there's kind of a -- I mean, I'll kind of rattle it off but obviously those of you on the radio either have to kind of type this along with me or naturally you want to pick up the video but for example you use logman.exe and then what, where, we're going to use the create option. What would you like to create? I'd like to create a counter capture. What would you like to call it? We'll call it black box, much like the airport, you know the airplanes flying in the air, we have a black box.

Richard Campbell: Right.

Shane Creamer: What would you like to use? We'll use the dash V because we want to append like the month-month, day-day, hour-hour, minute-minute so that's an mm-dd-hh-mm and then what is really cool is you can take all the counters that you want to use, put them on a text ball. I happen to call it counters.text and to use the best cfswitch, it will pull those counters and then you can use a Dash S which is what's my sample interval and then what's the format. The format is I'd like to be binary circular so it doesn't run away as my hard drive.

Richard Campbell: What would you like the output to be?

Shane Creamer: I like the output to be C:\black box_percent computer name and what's so sleek about that is what if I have a hundred servers and two weeks from next Tuesday server 77 and server 25 have these issues. No problem, I copy them to my computer but I've got two black box captures. How would I distinguish them? Well, they'll have the computer name at end of it like server 77 and it will have the month-month, day-day, hour-hour, minute-minute time stamp so I know which capture. If I've got 20 captures in there, I know which one I need to look at.

Richard Campbell: Nice.

Shane Creamer: And not but not the least is to use the Dash Mac switch and we set that for 250 megabytes and so that creates, that add something you can script out to hundreds of machines, have it go punch it in, next thing you know you go into PerfMon and there's the template and then you simply use a logman start script, go out to your machines with the Dash-s, Dash-s wack, wack server one, some I call wack wack server 2, maybe you want to get a little fancy, use like percent server name percent and create a separate text file on the four next loop, you go punch that one and there are lots of techniques, use a logman start and then you can start all these captures out there or to the black box. What you do is you use basically a start-up script and your start-up script has something maybe called like black box start.net and it has one line in there. Logman-start black box so whether or not you wanted to play via group policy to all your machines where you're handy, but what if you're just trying this for the first time, you don't want to ask your group policy guys to do this just check and it's kind of like your private test, then just modify the local computer policy and then wow, shazam! Every time the server reboots, as soon as the performance logs and alert service comes back up, the start up script comes up afterwards, it starts up the black box capture. We've got customers like Intel, DHL, Disney, Wells Fargo, Lockheed Martin, Boeing, big customers that just love this workshop and some of them had it 7 or 8 times because they want all the right teachings. When I'm teaching and they get it, they start reducing their issues, they will tire of reactive issues and the time to closures starts going down. There are teams like hey...

Richard Campbell: I want that. That's awesome.

Shane Creamer: One of those things, ultimately I think we're going to have to do is I'm really hoping eventually a tech net guy who writes books for a living picks this up and go "Shane, how'd you and I like to



partner together and get this book out to the rest of the world." I was like, "I'd love to." That would be my dream, that people could download or buy a book that takes them through chapter by chapter and I think it's awesome the Dummies Guide format, I mean I've read them. It's got the fishing because my son is sick so he wants to go fishing and it's like I don't fish much so I've got the Dummies Guide to Fishing. Great book.

Richard Campbell: Right.

Shane Creamer: I'd love something like that for our customers.

Richard Campbell: Yeah, it sounds like the book that needs to exist. We got to figure that out. Shane, we're out of time. We're long but it's such good stuff. Thank you so much for coming along on the show.

Greg Hughes: Great info.

Shane Creamer: It was my pleasure and again if the listener response is, hey, we'd love to have some of these topics, I'm here, I mean they're my customers. Whether I know them or not, I want to help and so I'd be glad to help any way I can.

Richard Campbell: I think it feels like a series. Maybe some of them we get an insert once a month just to get people up to speed or maybe a completely separate track to go. We have to work it out.

Shane Creamer: Richard, Greg, it was a pleasure and a privilege to talk to you.

Greg Hughes: Likewise, likewise, very good.

Richard Campbell: And we'll see you next week on RunAs Radio.