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 #138
(Transcription services provided by PWOP Productions)

Sumeet Bansal and Chris Featherstone Go Solid State Driving!
December 09, 2009

[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 #138, with guests Sumeet Bansal and Chris Featherstone, recorded Monday, November 30, 2009. 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. As always, my friend, things are good. It's the fall, you know the weather is bad, a good time to stay inside and listen to a podcast.

Greg Hughes: Yeah, that works pretty good.

Richard Campbell: Yeah.

Greg Hughes: Although we've certainly had our share of very, very wet, cool weather here.

Richard Campbell: Yeah. You'd think it's wintertime or something.

Greg Hughes: Or something, yeah.

Richard Campbell: All right. Hey, I've got a really exciting show coming up here and I know you're looking forward to it as well.

Greg Hughes: Yeah.

Richard Campbell: So let's get straight to our guest here. We've got two on the line here. Sumeet Bansal -and Chris Featherstone both from Fusion-io. Smith is the principal solutions architect and Smith is responsible for designing architectural solutions that use the Fusion-io product as core technology to deliver unprecedented levels of scalability. Performance is simplicity. Smith brings more than 11 years in IT and database infrastructure experience to the position. Chris is the product marketing and senior PM at Fusion-io as well. He is heavily involved in the product marketing team focused on next generation software products offered by Fusion-io. Guys, we'll talk further into that, but the topic today is the Solid State Storage revolution because things are going nuts.

Greg Hughes: There are an awful lots of questions I know I have around Solid State Storage.

Chris Featherstone: With a lot of people we've talked to, they seem to be a matter of when and a matter of if.

Richard Campbell: Well, and these days it's just seems like everybody is buying a laptop with an SSD in it and so I feel like the workstations are done. They're using SSDs, they work. The real debate here is are SSDs ready for the enterprise.

Sumeet Bansal: Absolutely, absolutely. No question about that. We are getting a lot of interest, a lot of clients and you know how it start is a company calls us, they say "Look, we're having severe performance problems. Can you help?" And we say, "All right, why don't you try it out?" We send them the product. They try and they just drop by six inches. You know, they're amazed at what this technology is doing for them and we can't keep up, we can't keep up with how much demand there is for something like this.

Greg Hughes: So let's talk about what we're talking about I guess. What is Solid State Storage? For the uninitiated and uninformed, what is the technology?

Sumeet Bansal: Right. In very simplest terms it is storage, so it is persistent storage but the core difference is that there are no moving caps. There are no spinning classes and there are no moving heads. So given that, the idea, the whole idea behind this kind of storage is to provide a very high performance both in terms of reads and writes. If you look at a spinning disk, because of the way it works, because of the pressure and the moving head design, it is limited. It has limitations which are critical and they're mostly governed by Newton's Law. The platter can only spin so fast because if it spins faster than that then it will have quality issues. Having an SSD-based storage forward comes some of those issues. So the whole idea is for performance and there are many ways in which companies are offering this kind of a product primarily there are products which work on the SAS, SATA, or Fiber Channel Interface, and then there are companies that are providing SSDs that work on the PCIE interface and obviously these interfaces they determine how fast or what kind of low latency this kind of technology provides.

Richard Campbell: Right. You've mentioned that Newton's Law and some of the limitations on drives moving parts they can break, one of the things that I know I hear over and over again about SSD technology is you can only do so many reads and writes, or more specifically writes before, it just won't work anymore. Is that true or is that false or is it somewhere in between?

Sumeet Bansal: Yes. So those kinds of limitations are true but that is true as applied to the native components of an SSD which is the nan-Flash itself.

Greg Hughes: Okay.

Sumeet Bansal: So nan-Flash by itself, it's a commodity piece and one can only write so much to it. However, if you look at SSD products in general, they are more than just nan-Flash. The idea behind it is to collect a lot of these nan-Flash chips together on a piece of hardware and to put infrastructure around it to put controllers, to put software, to circumvent some of those issues. So the idea is if you know that nan-Flash or seven pixels on nan-Flash can only be written for so much, what is done is writes are distributed. So the idea is to make sure that when data is written to a bank of this nan-Flash, it's not written consistently to the same location and the rights are distributed in such an even way that you get a much longer life. Just to give an example just by means of implementing the right kind of algorithm and infrastructure. It is possible, and I'll just give you an example of one of the Fusion-io drive, it's our 160 gigabytes SSD product. This is to speak in terms of its endurance. If one writes 2.5 terabytes of data everyday. It is expected to last for 19 years.

Richard Campbell: Well, you're not going to get 19 years out of an average hard drive anyway.

Greg Hughes: I'd be lucky if I get three years out of a hard drive.

Sumeet Bansal: Exactly.

Chris Featherstone: Exactly.

Sumeet Bansal: So the idea is that yes, nan-Flash it being a commodity, it has its flaws and that's what makes the value proposition go strong. It's because we take this thing, we take this core piece which has flaws but then we build on it. You know, we build on it by other hard drive components and software components and ultimately we deliver the product which really has a very high business value proposition.

Greg Hughes: You know, I had a notebook with an early SSD drive in it that is really kind of a cool idea but it was so slow, and what I'm hearing you say is that what you build around this technology is probably what might control whether it performs well or it doesn't perform well. It seemed like at least early on there is a trade off. You either got size or you speed.

Sumeet Bansal: Once again, those are really design issues or the implementation issues and we have really gotten around all of them so we really give you the best of both worlds and then give you high density. So in other words, we can provide a single card that has 640 gigabytes of storage capacity on it and on top of that, you know, we can provide fairly high performance. For example, a 640-gigabyte card that we have has the ability to move more than 1.2 gigabytes, that is gigabytes not gigabits. It can move more than 1.2 gigabytes per second. That's the kind of software that we're talking about. In the beginning, I talked about implementation and specifically the interface implementation. You see, what happens today is most of this SSD technology that you see in laptops or even for server market, it's based on the more traditional SAS, SATA, or Fiber Channel Interfaces.

Richard Campbell: Right.

Sumeet Bansal: The problem with these interfaces is that it increases latency because there is a lot of work that needs to be done to translate these protocols and our product, Fusion-io, drives the work on PCIE interface. The benefit of this is that it is very close to the CPU that has very high bandwidth on the PCIE slot in a server and it doesn't have to go to any of those translations and as a result we provide microsecond level of latency. All of our products, they have anywhere between 40-microsecond to 70 microsecond latencies.

Richard Campbell: Obviously, in the workstation environment we're seeing the state of the interface being the way they go just because everybody has already got one, and in a traditional hard drive the performance just hasn't been that big of a deal of the latency there. I think most people get the benefits of transfer rate more than anything with SSDs. They're just 10 times faster than regular drives?

Greg Hughes: Are they really 10 times faster?

Sumeet Bansal: It depends on the SSDs again.

Richard Campbell: Yeah.

Sumeet Bansal: And again it depends on whether the SSDs are for full workstations or for the servers.

Greg Hughes: if you look at just the commodity hardware, what's the variation of just the commodity parts of it in terms of speed and size and capacity, and is there really a big variety there or is all sorts of close to the same?

Chris Featherstone: Well, you know, I think nan-Flash and nan-Flash, there's SLC and MLC based and then I think the key is when you get into -- so obviously that's lots of the commodity that's sitting on the board themselves. The key is then when you start to look at what interface does the sit behind, if they have a physical controller or software controller which we have, that's when you get into some of the key latency as well as true put and performance issues because we're really not sitting behind anything because we're sitting directly in the PCIE bus. Just by virtue of that one thing, because we're not behind again another physical controller like a SATA or SAS,

or even a Fiber Channel Controller, we can increase the speed that much more.

Richard Campbell: And this gets back to this whole thing of regular drives have a few milliseconds of latencies so that latency of the interface isn't a big thing when you're dealing with the drive, but once you get down to this NAN speed, the latency of the interface matters.

Chris Featherstone: Sure.

Richard Campbell: But still, I think that seek times are so good on these things. That just seems to me like the interface is not that big a deal, but this is transfer rate that most people notice for performance on these drives.

Chris Featherstone: So what happens is again it all depends on the application.

Richard Campbell: Right.

Sumeet Bansal: And different kinds of applications have different kinds of requirements. If you take any storage medium and if you try to define its performance profile, you can define this performance profile based on a few factors. One is IOPS input off its ability for input, output operations per second. One is true put, how much they think it can move per second, and the third one is latency. All play a role depending on what kind of an environment it is. So for example, if you take an enterprise database environment and let's say it happens to be an ecommerce environment, it's driving your website, then you can kind of presume that it's an OLTP or Online Transaction Processing base environment. So there by definition, the kind of transactions that you expect are small, small transactions, small really drive transactions but in large quantity, in large volume. When that happens, then the factors or the metrics of this performance profile that we created earlier, the two factors that matter most are the latency and the I/Os.

Richard Campbell: Okay.

Sumeet Bansal: Latency is important because I mean what is latency? It's the time taken for data movements from your storage medium to the processor and especially when you're talking about an improvement of say millisecond scale to microsecond scale. That shows, that comes up especially when a transaction count is high.

Richard Campbell: Sure. When I look at the typical IT guy who's already running an infrastructure of spinning media and lots of it, I got to think he wants to be able to just swap out, to say can I take these drives out and install them as say in an external chassis and just put SSDs in there instead. I'm thinking particularly stuff like clustered servers around an iSCSI array. It seems to me these things just drop in and are going to work. I guess the biggest concern here is the performance we're giving out at best comes from RAID within the SSD and then we're layering on this additional RAID infrastructure. Are we just beating these drives unnecessarily?

Sumeet Bansal: No, not necessarily and I think that model is valid. The issue, however, is it has to realize what are SSDs giving you. If you put SSD in a chassis, then you could very easily end up in a situation where your bottleneck becomes the connectivity between your chassis and the server because if you say put less than 10, if you put 10 SSD drives in a chassis, it is possible that you're violating them. You might be getting true put and I/Os so high that you can't actually make use of it because you can't move that much data through the wire to your server.

Richard Campbell: That's a good problem to have though, to actually be transferring with those kinds of arrays.

Sumeet Bansal: Well, absolutely and again there are many companies that are working on solving just that. I mean, one simple concept is data compression. If you know that you have, you can move certain amount of data. There are companies out there today they are trying to figure out how to compress and decompress data on the fly on both ends.

Greg Hughes: Are there certain no-brainer just obvious applications where using this technology yields great, big, huge results that you see every single time?

Sumeet Bansal: Absolutely, databases. I mean I just can't tell you enough how easy it is to improve the performance of your system by three times, four times, five times and to do it is literally so easy. All you have to do is you open up your box and if you're using SAS database as it is, then obviously you most likely already have the drive-based to accept them. If you're using PCIE-based SSDs, then you need to have PCIE slot.

Richard Campbell: Right.

Sumeet Bansal: It's really as simple as that. You pop up in your 10 box, you put this disk in and you just pinpoint your databases to put all of the data files and log files on this media and immediately you can realize there's huge improvement. I mean, can you imagine that cost, like you go to any business and if they have an objective of improving their performance to double, they want to double the performance, it will be usually expensive to do that say using traditional means today which is adding more spindles and adding more trace to their SAN infrastructure.

Greg Hughes: That really takes me to my other question which is from a cost perspective, what's the delta, what's the difference in the cost if I have a terabyte of hard drives spinning base storage versus Solid Stage Storage if you think of it in industry terms?

Sumeet Bansal: If the only matter that you will look at is dollars, then yes. The hard drives or the spinning media will build in much fewer dollars than say something like SSD.

Greg Hughes: If I'm looking at it in terms of dollars per true put unit or dollars per the other components you're talking about, the speed, I understand that that's going to be a different figure. Isn't that right?

Sumeet Bansal: Exactly and when it comes to dollars but true put or dollars per performance, then that's where SSD shines.

Richard Campbell: Sure.

Sumeet Bansal: In fact, if you go to the TPC website, you can actually find, if you look at the top 10 by price performance, you will actually see that the top that hovers are actually tests done by some server vendor and some SSD vendors.

Chris Featherstone: I think one thing too to keep in mind especially with cost, since we're looking at it, obviously SSDs could be perceived as expensive if it's just a price per gigabyte type of scenario, whereas like we've been talking with price per performance and/or overall TCO especially when you look at the amount of operating cost, you can reduce in terms of cooling as well as in terms of the number or electricity that's involved in running SSD spindle.

Greg Hughes: What is the difference in power consumption? I hadn't thought about that as a primary factor, but what does it mean in terms of cooling in terms of power?

Chris Featherstone: So if you look at the price, you almost go price per performance per kilowatt, it would be kind of the natural measurement.

Greg Hughes: Sure.

Chris Featherstone: And we're almost 2300 times less than a standard hard drive. I mean, you know, in terms of the amount of kilowatts that we can run over let's say, and you know if you're going to deliver 100,000 IOPS on a 15,000 RPM drive, it's going to be about 2,000 more in terms of kilowatts per year than one of our drives per se, one of our SSD drives.

Greg Hughes: So obviously from a cooling perspective, it's going to be a big difference as well.

Chris Featherstone: Absolutely.

Greg Hughes: I mean, I know. I deal with datacenters where you can't get more power from the Power Grid and then you give them the physical location to put more stuff in the datacenter between a cooling perspective and how much power is available respective. There's empty floor space.

Chris Featherstone: Uh-hmm, correct.

Greg Hughes: So from what I'm hearing you say is that this is one option to try to find a way to, a) increase performance, and b) be able to get better density in your datacenter.

Chris Featherstone: Absolutely. I mean, IT administrators are going to see the biggest paying for the buck because we know that they're feeling a ton of pressure to reduce cost, reduce server footprints while increase performance and how do you do that. You know, I'm sure a lot of those guys their heads swell. Technology is like they use with some of these storage are way to the back and actually become a reality very quickly because of what they provide.

Richard Campbell: I was going down the path here that the SSDs don't have the capacity of spinning media so you end up with more of them, but that is not true anymore.

Chris Featherstone: Correct.

Richard Campbell: When I think about 15K drives, I think the biggest I can get is about a 320 and that's still a 2-1/2 or 3-1/2 inch form factor and I can get a 2-1/2 inch 512-giga SSD. You know, the big thing that I still don't get the feel of here is that when I put drives into an enterprise infrastructure I buy enterprise class drives with five-year warranties, 100% duty cycle and I've seen the difference between the drive by buy that's of that caliber compared to a workstation drive because I know an enterprise class drive spins and works for full-bore 24 hours a day. It's in an array and it's being hammered, and I don't see that distinction in the SSD world. It just seems to me like they've been doing the same things with -- they've been making that distinction with SSDs.

Sumeet Bansal: So with SSDs, there are just different kinds of things that one needs to be aware of. So I spoke about longevity. There's a huge difference between say how much an OCZ drive is expected to last versus how much an entire drive is expected to last.

Richard Campbell: Right.

Sumeet Bansal: Versus how much a Fusion-io drive is expected to last. So longevity is very

important. What are the mechanisms implemented to make sure that the data is safe on your SSD.

Richard Campbell: Right.

Sumeet Bansal: There's a huge difference between how different vendors are implementing protections. What happens when the data suddenly power up outage? Do your in-flight transactions get persisted to your media? And again that is that there are different technologies out there and there's a huge difference in how much effort some things are making to make sure the data doesn't get corrupted, or stays protected. So it's tricky. A lot of these issues have not been -- they're not really apparent when people talk about spinning drives or spinning media because people have just gotten used to them.

Richard Campbell: Yeah.

Sumeet Bansal: I mean, that has been there for so long now but this is a different kind of storage and so there are many, many points of differentiation which are not apparent on surface because what is the first thing that people do when they get their hands on SSD? They install it, and then there are some basic disk benchmarking.

Richard Campbell: Right.

Sumeet Bansal: How fast is it reading? How fast is it writing? But how many people actually go to the trouble of putting the disk scan running a load and then say plug pulling the cord out?

Richard Campbell: And I remember doing this in the '80s with databases. You know, we're trying to prove the fact that the moment it says the transaction is complete, even if you yank the power plug out, when the machine comes back up that transaction is there. We've got bitten by this with RAID controllers that had caches on them where they needed to have battery back-up and so forth so if you did lose power, it will be able to recover properly from it.

Chris Featherstone: Thank goodness that NaN is nonvolatile.

Richard Campbell: Yeah and so if it's written, it's written and that's all there is to that.

Chris Featherstone: I think one of the other key things too, I mean it's just like spinning with mechanical drives. When you take a mechanical drive and then obviously per vendors are going to put their secret sauce if you will on top of that to make it their key differentiator and their value add and I think that's one thing to keep in mind especially with the way that since NaN is a commodity out there, everybody can get it, everybody can utilize it, that's going to be the key differentiator. It's what we provide over an X competitor to provide that value out to the IT administrators, four things that some of you is talking about like in case of data loss or power loss or things like that, that's going to be the absolute key especially going back to the original point of how am I going to have those types of guarantees that I can get today with mechanical disks.

Richard Campbell: I think the other aspect of familiarity here is that IT folks are used to the failure modes of spinning hard drives. We've got smart, we've got these technologies that let's us know, hey, you start to lose sectors on this drive, this drive is heading towards a failure, get it out now. What is SSD failure looks like, how does it die?

Chris Featherstone: One of the things that we do for those types of cases is that we will basically ask the drive. Bwe talk about the endurance and the wear leveling and things...

Richard Campbell: Sure.

Chris Featherstone: We know we can set those thresholds to actually then throw up let's say for an SNMP trap or WMI event to say, hey, this drive is getting close to its end of life, go ahead and replace it. So we have actually as part of the software on the hardware itself that will give the ability to look and peer, see the ware, have that information and those statistics and be able to bubble them up too in let's say a monitoring environment.

Richard Campbell: Okay but you don't go down the smart route which should be make it more like a regular hard drive. You're using SNMP, like you said different ways to communicate out the fact that I know where my ware is on this drive.

Chris Featherstone: Uh-hmm. That's actually the lower level. Obviously when we put our technology in the systems, they come up as a block device.

Richard Campbell: Right.

Chris Featherstone: So you know, again whatever file system or whatever operating system, we're looking at as a block device. But as we get in the lower levels of the firmware and the drivers, we can give that information up, involve a better information up.

Richard Campbell: Yeah and you just get the sense that people can get comfortable with. I know when this drive is going to fail because every drive fails, and can anticipate and move things along like that. I recently saw OCZ announcing a one terabyte SSD so size thing seems to be about done.

Chris Featherstone: Yeah.

Sumeet Bansal: So on that point, one thing that I do want to point out is size is important and we do

believe that high density is important, but again consumers should take the time to dissect the product that they're getting into...

Richard Campbell: Yes.

Sumeet Bansal: If they are getting one terabyte of storage, they should know what is inside that. How are they accomplishing the one terabyte storage capacity? For example, they are vendors out there, and no names, and essentially what they're doing is they're taking say four or five distinct SSD modules, and you know what? This maybe actually SATA or SAS module, but they're packaging them together on one physical enclosure and they're putting their own controller, their own RAM, their own even a processor, some small microprocessor and they're covering it up in a package and then they're sort of building a PCIE plug on it. So it looks like PCIE-based storage but it really isn't. It's very deceptive, but more than that it can be very unreliable. If you are simply taking four commodities as it is and physically packaging them in one enclosure, then you're actually introducing a component which has a very high chance of failure.

Richard Campbell: Yeah, I understand and I think this is part of the challenge that folks are having with us, is this in general there's been a few different approaches to this. Some are less expensive, some are more expensive and it's hard to know the ones you can really count on. In the early days, and when I mean early days I mean like last year, it was hard to find an SSD with a warranty longer than a year, but now we're starting to see three-year and five-year warranties on these SSDs so folks are starting to believe enough to warrant their product for long durations but priced accordingly too. The expensive SSDs seem to be the ones with the long warranties and the higher reliabilities.

Chris Featherstone: There is parity there absolutely.

Richard Campbell: Yeah, it seems to be pretty consistent there. We certainly, I think, have beaten our own on latency affair bit here. The fact when we go to the simpler and closer interfaces like PCIE, we get really low latency. But what do you see that is real in transfer rates? What is an actual physical transfer rate that people can believe in? I know I've gotten hard drives in ideal conditions up to the hundred megabytes per second range, but in reality 40 and 50 is real. So what's real in an SSD?

Sumeet Bansal: Are you talking about like is it generalized, talking about workstation SSDs or server SSDs?

Richard Campbell: I guess I'm interested in the difference there, like why would they be different in terms of transfer rate? What are the ranges that they tend to fall into?

Sumeet Bansal: I will start with the enterprise SSDs.

Richard Campbell: Okay.

Sumeet Bansal: And it's very easy to get rates of over 700 megabytes per second and what more is one can actually get a rate higher than that if you take multiple SSDs and interface them together and it's real, it's very real, and then obviously when you drop down everything and you go to the workstation kind of SSD, again data rates are from 200 to 300 megabytes per second, they're very real, very tangible.

Richard Campbell: And do find that that speed decays overtime, that as the drive ages the transfer rate declines?

Sumeet Bansal: It does but the rate of decline also depends on the vendor and specific IP that the vendor has developed around it. So for some instances, the decay or the rate of decay is more than 50% and so as time goes on, as the drive stops filling up or maybe the drive is 50% full – there are some products out there where they drop in performance by more than 50% and then again it's something like that's where we have put a lot of power and we have the right algorithms and software magic in place to ensure that performance is more or less consistent.

Richard Campbell: Let's talk a little bit about Fusion-io because I ran across an article in hardware about you doing exactly this. When you took an array of Intel X-25, which I think are the benchmark SSD in the say form factor out there, and compare them against one of your Fusion-io products. Can you tell me a little bit about that?

Sumeet Bansal: Right. It's a good tent and so the thing to consider here is, and I think the question goes back to density again, if you have two options to get the same density by getting say one CAD and it doesn't have to be Fusion CAD but it can be in any SSD CAD.

Richard Campbell: Right.

Sumeet Bansal: What if say you can get the same capacity by reading multiple counts, which is the better option? By design, the one CAD option is a better choice because you have less components in your architecture, therefore you have less points of failure. Also you have to consider that when you are reading multiple SSDs, do you need a physical rate controller to do that? Because if you do, then you have to know that you will have a bottleneck because they're physical rate controller itself. It doesn't matter how fast your SSDs are, but your controller will and I believe that the bottleneck there is 300 megabytes per second for the critical controllers.

Richard Campbell: Right.

Sumeet Bansal: So you're already bottle-necking yourself. You are kind of using the technology to try to fulfill the promise of high performance but you don't have a lot of room to grow.

Richard Campbell: Well, and 300 megabytes per second was awesome when you're using spinning media because you never hit that limit, but now we're finally dealing with drives that are doing 600, 700, 800 megabytes per second and so the rate of the infrastructure is now the limiting factor.

Sumeet Bansal: Correct, correct and as time goes on I'm sure that rate controllers will improve themselves in design and they will try to remove the bottleneck, but for now if one can implement SSDs without having to use a physical rate controller, then that's when you're truly realizing the benefit of SSDs.

Richard Campbell: And this really gets back to I think one of the points you've made early on here which is that this is a totally different way of handling storage and so the infrastructure first spinning media is inappropriate, it's obsolete essentially.

Sumeet Bansal: I think so, I think so. I'll just give you an example. So wine.com is one of our premiere customers and they used to have NetApp SAN, that's what they were using to run their primary database systems.

Richard Campbell: Right.

Sumeet Bansal: It has been completely eliminated by Fusion-io. The last time I checked, their SANs where still sitting empty. Probably they're using it to store MP3s if anything.

Richard Campbell: That's some expensive MP3 store.

Sumeet Bansal: Businesses have been for far too long used to – they're just being used to using the SAN because they haven't anything better. When they wanted performance, all right, add more trace, add more spindle.

Richard Campbell: Right.

Sumeet Bansal: That was the mantra, but no more because now there's a much more scalable and much more easy to manage solution.

Chris Featherstone: One thing to keep in mind as well, especially with this storage, because there are really two areas where we're talking using this for storage and then using this for upfront local attached storage 4.0 application or things like that, but when you get into the storage you really relegate those different tiers from your general purpose. Now you kind of push tier 3 back to straight archival and really almost eliminate two or three because now you have Solid State disk cap in front, you can then move your tiers up and now we basically should be storage area archive instead of search area network or something like that.

Richard Campbell: You know, you're almost reflecting the CPU architecture where the RAM wasn't fast enough so we had this the layers of cache coupled closely to the CPU for that really, really high performance. I mean, the honest truth here, looking at this story in hard hardware was it was 160 gig I/O drive they're using and it was unbelievably smoking fast, but if I had 10 terabytes of total archive stories I need, that's an awful lot of I/O drives. I'm going to run out of PCIE bus at some point so keep what's actually – use it like a nonvolatile cache that the stuff that I'm working on right now is in that fast, fast, fast, fast, fast environment and other stuff is being pushed back to the SAN.

Chris Featherstone: Right, we believe that it will fundamentally change the way the applications are built, the way that storage is looked at because it really is a new memory tier for storage, for applications to utilize because of the blazing performance, because of reduction in cost because of all those key elements of providing a really, really great hardware for a lot of the infrastructure out there.

Richard Campbell: Smith, Chris, these are fascinating technologies. I really appreciate your time in coming to talk to us a bit about it.

Chris Featherstone: It's our pleasure and looking forward to some of the great things we can do with these technologies. It's exciting and I think the most part is just like when we went from core about COM and Java and things like that, it's just, to a managed environment, is just really, really refreshing technology.

Richard Campbell: Cool. All right, and we'll talk to you next week on RunAs Radio.