



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 #0101*  
(Transcription services provided by [PWOP Productions](#))



**Alan Sugano Digs Into Storage Technologies!**  
**March 18, 2009**



[Music]

**Brandon Wenn:** From [runasradio.com](http://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 #101, with guest Alan Sugano, recorded Thursday, March 12, 2009. RunAs Radio is produced each week by PWOP Productions, providing professional media and podcasting services online at [pwop.com](http://pwop.com). You can follow the boys on Twitter at [twitter.com/runasradio](http://twitter.com/runasradio).

**Richard Campbell:** Thank you very much. This is Richard Campbell. You're listening to RunAs Radio and with me as always is my co-host Greg Hughes.

**Greg Hughes:** Hey, how is it going?

**Richard Campbell:** It's good. Hey, congratulations on our 100th show once again and that show came out really well.

**Greg Hughes:** Yeah and if I can be really honest it came out better than I thought it would but it was a lot of fun, wasn't it? I had the opportunity to go back and listen to it yesterday because I was getting these emails and tweets and different instant messages and stuff saying, hey, that was pretty cool so I figured I should go and find out aside from saying "uh" a whole bunch of time that actually I sounded pretty coherent. In most of them I thought I did.

**Richard Campbell:** You're a more interesting person than you realized, Mr. Hughes.

**Greg Hughes:** And I'm much more interesting now that I've done that, I suppose.

**Richard Campbell:** Okay, there you go.

**Greg Hughes:** Yeah, that was a lot of fun so thanks for the chance to do that. I had a good time, a good time talking with you.

**Richard Campbell:** Absolutely and now we're ready for the next hundred.

**Greg Hughes:** Yeah, it will only take us two more years.

**Richard Campbell:** Yeah at the pace we're going.

**Greg Hughes:** That's right.

**Richard Campbell:** So let's get started. All right, Greg, let's introduce our guest. Alan Sugano is the president of ADS Consulting Group founded in 1991, ADS specializes in networking, server virtualization, workstation virtualization, application virtualization,

security, exchange, custom programming, web development, SharePoint and SQL Server Development. Is that really specializing?

**Alan Sugano:** Not really.

**Richard Campbell:** That sounds like a full gambit.

**Alan Sugano:** Jack of all trades and master of none.

**Richard Campbell:** There you go, you do the whole thing. A regular speaker at various conferences including WinConnections. Hey, we're going to DevConnections.

**Alan Sugano:** Oh, you are. Maybe I'll see you guys there. I can hardly wait.

**Richard Campbell:** Absolutely.

**Alan Sugano:** I'm pretty sure DevConnections runs at the same time as WinConnections.

**Richard Campbell:** They might be a week apart actually. Maybe that's why we keep missing each other.

**Alan Sugano:** Yeah, I think you're actually right. Are you there next week or...?

**Richard Campbell:** The week after.

**Alan Sugano:** Yeah, okay, then they're one week apart. I can't remember the one show, the Spring show, they're apart, and then the Fall show, they're together.

**Richard Campbell:** Right because the venues are bigger. And also an MVP in Connected Systems?

**Alan Sugano:** Yes, yes. Well, actually it didn't get renewed for 2009. The magazine actually had made me hood an article on the 10 reasons not to upgrade the Vista and I think it might have upset some people at Microsoft so...

**Richard Campbell:** Uh-oh.

**Greg Hughes:** Yeah, that might do it.

**Alan Sugano:** I got "iixnay'd on the MVP eh."

**Richard Campbell:** That's too bad.

**Alan Sugano:** I got to call it like I see it.

**Richard Campbell:** Yeah.



**Alan Sugano:** Not like, you know. We have a lot of clients. I'd say some of them are in Vista but most of them are waiting for Windows 7.0 which actually looks very encouraging which is a good thing.

**Richard Campbell:** It is interesting to see, and we know perfectly well that Windows 7.0 is still the Vista at the core but it feels like they've just refocused themselves on things that I think matters a lot to the customers. For the first time ever, I feel with Vista and maybe we have the same experience. It's doing stuff I don't know about, like it's busy, I have this feeling that my computer is busy and I don't know why.

**Alan Sugano:** Yes, yes, you look at the hard drive and it flash like crazy and like what the heck is this thing doing.

**Greg Hughes:** Exactly. I'm running Windows 7.0 on the recording machine here on the beta and I've really enjoyed it so far. So I go on my security -- a little rant -- but from a security standpoint they've strengthened it quite a bit which is encouraging to see.

**Alan Sugano:** Yes. I think with Vista they might have been a little bit over too far one way to the point where it's little, well, at least it's very annoying but when they go to Windows 7.0 they really kind of address a lot of issues and the performance seems to be a little bit better. The upgrade cap is not quite as steep now that everybody has a chance to maybe upgrade to those from later software but, you know, when Vista first came out you had to upgrade almost everything to get it run properly with Vista.

**Greg Hughes:** Right.

**Richard Campbell:** Well, it does seem like it's just configuration changes now to get things right and the drivers are settling down so Win 7.0 is going to have the benefits of the pain we felt in Vista.

**Alan Sugano:** Yes, yes. I think Microsoft took a really good look at the issues that people were bringing up with Vista and it looks like they really addressed, not all of them but a good majority of them so it's good.

**Richard Campbell:** Absolutely. We were going to talk about storage technologies because apparently you spent a fair part of your time battling with this?

**Alan Sugano:** Yeah, in the context of everything we do so, you know, storage is something that you always have to be aware of and we try to keep on top of things and just be aware of what's coming out in the industry and make sure that we're

positioning our clients in the right place they need to be.

**Richard Campbell:** I think there are more choices now than ever for storage options.

**Alan Sugano:** That's 100% correct especially with this latest generation of hard drives, the SATA drives and SAS drives and there's this new thing, the hot new thing in the storage industry for SAN is this thing called Fibre Channel over Ethernet. It's got everybody pretty excited including myself primarily because the price points might be more attractive than traditional Fibre Channel SANs.

**Richard Campbell:** Well, when I think Fibre Channel hovering over Ethernet, I immediately think iSCSI. Isn't it the same thing?

**Alan Sugano:** No. Actually it's a little bit different. With iSCSI, basically what you're doing is you're taking, typically you're taking iSCSI calls and then there's an overhead process that basically converts the iSCSI request, the Ethernet request and that's typically done on a TCP Offload Engine or TOE card which typically reside on the Ethernet card itself. So you do have some overhead there and it's a little bit different in that once it's upon Ethernet that's when you can use iSCSI for like geographically distributed SAN or infrastructure like that, whereas Fibre Channel over Ethernet is a little bit different. It typically would use 10-gig Ethernet to do the transmission but it uses native 10-gig Ethernet as a physical transport layer to deliver the Fibre Channel load. So there's not really a conversion process because it's basically using Ethernet natively I'll say. So you don't have the overhead of converting iSCSI request to Ethernet.

**Richard Campbell:** So then you're also not trying to bang it through switches and things like that. I find most people set up iSCSI that way anyway, that there's literally just a cable running from the server into the iSCSI chassis. Is that the preferred sort of wiring Fibre Channel over Ethernet as well?

**Alan Sugano:** Well, in Fibre Channel over Ethernet it's a little bit different. It's not routable over a TCP/IP network. It doesn't have the TCP/IP portion. It's just Ethernet as the physical transport. The biggest advantage of it basically is its price point. You know there are several different flavors of fibre channels running from 2-gig, 4-gig, and 8-gig and the 10-gig stuff. The advantage of the 10-gig stuff is just coming out now, it's the price point. Everybody is hoping that it's going to be a lot cheaper and it will be faster than fibre channel traditionally and it depends on the vendor but a lot of times fibre channels still, I won't say insanely expensive, how about just mildly expensive, mildly insanely expensive. Usually, a lot of times on fibre channel switches you need, there's



an up-charge. Every time you bring up a port, you might have an 8-port switch but every time you make the port on the switch active, they charge you for it and it's a couple of grand pop there plus the redundant cards, blah, blah, blah so not everybody is happy with using 10-gig, the 10-gig Ethernet stuff, it's that there won't be a charge for that hopefully and for the per turn-out per charge and the actual equipment itself is going to be a little bit more reasonable and give better performance than the traditional fibre channels SAN stuff.

**Richard Campbell:** Although today, in 2009, 10-gigabyte Ethernet Nic's are still pretty darn pricey.

**Alan Sugano:** They are but it's all relative. You know, fibre channel, by the time you get the fibre channel card, it's very close to the -- it's probably, it depends on what vendor you go but they ride on the same price point but then you have the connectors and you have to do a redone in connection and I think the ace that's up the sleeves for Fibre Channel over Ethernet is that as it becomes more widely adapted, a lot of people are now deploying gig to the desktops so the next logical thing is 10-gig for the background. If the price is going to draw up like a rock, I mean when you look at the first gig switch that first came out, whether they're like a million dollars per switch or something like that, now you can get them for a couple of hundred bucks, no, maybe not that cheap but the price point have come down the price where they're very close to what a hundred megabyte Ethernet used to be even just a few years ago so that's what I'm hoping will happen with the 10-gig stuff and I was screaming even when iSCSI came out. I go why doesn't anybody have iSCSI over 10-gig yet because it's almost a no-brainer.

**Richard Campbell:** That seems to be on the radar as well, there will be 10-gig iSCSI.

**Alan Sugano:** Yes and again depending on what you -- if you look at it from a 10,000 feet level view theoretically, and again it depends on the hardware you're running and the configuration, 10 gigabytes iSCSI will probably be theoretically a little bit slower than Fibre Channel over Ethernet just because of the overhead of transporting this iSCSI request to Ethernet request. The advantage of iSCSI that it has is that the packets are routable.

**Richard Campbell:** Yeah. So the whole idea that I can't use switches to share out this stuff, a lot of infrastructure is available to me as soon as I go TCP/IP.

**Alan Sugano:** Exactly, exactly.

**Richard Campbell:** The normal configuration for Fibre Channel over Ethernet you're essentially talking

about a cable running from the server to the SAN or whatever the storage device that handles Fibre Channel.

**Alan Sugano:** Right, right. Well, you actually would run from the server, or the host, to the fibre channel switch and then the switch would go, then there would be another connection that switched the actual SAN. I mean that's very traditional and yet the only way to do it.

**Richard Campbell:** Right.

**Alan Sugano:** One thing though that iSCSI or any type of SAN, even the Fibre Channel over Ethernet, what we do recommend doing is unless you need to geographically stretch the SAN, we like to definitely run all that iSCSI stuff on a dedicated network separate from the normal traffic that the host see. That's something we almost insist on. We say if you're not going to do it this way, then use somebody else.

**Richard Campbell:** You don't bother.

**Greg Hughes:** Sure.

**Richard Campbell:** Is VLAN sufficient for that though?

**Alan Sugano:** Boy, that's a really good question and that brings a whole other ball of wax.

**Greg Hughes:** There's a logical separation of this, actually physically different.

**Alan Sugano:** In general, we actually, even for general traffic, we try to stay away from VLAN if we can and sometimes you just can't. I mean, I see where you're coming from. The main reason for doing a VLAN is to segment off the traffic but VLAN can still be relatively easily cracked with flood, packet flood...

**Greg Hughes:** Right.

**Alan Sugano:** Where if you flood the switch with a ton of traffic, the switch kind of gives up and basically turns your switch into a gigantic hub so it stretch broadcasting things. I'm not to say it have to be a conscious effort to do that but I think a lot of people do VLAN if they have a false sense of security that it's secure and we argue that if you can do it and your that paranoid about security, then spend an extra couple hundred bucks and put it on a physically separate switch on a physically separate network. If you want that network segmentation, then do it that way as it's possible. Now in some cases, it's not possible because you might have shared fibre links or something like that but in general we try to stay away from VLAN.



**Richard Campbell:** But the point being if you don't need to be routable for that traffic, don't bother, like it just seems like a headache. The main reason I even see iSCSI plugged into switches is because they're doing it for redundancy, for failover purposes.

**Alan Sugano:** Right.

**Richard Campbell:** That they have one port coming out of the machine and it can go to one or more storage devices and swap between them kind of seamlessly, but if you don't need to do that, don't. Stay out of the whole thing, just wire it directly.

**Alan Sugano:** Yeah and keep it on a physically separate network and a dedicated switch. The other thing that you can do is, and this would be more by accident, if you were VLAN off some of these are going to come into the switch and at 3:00 in the morning because they're really tired, they're going to make a booboo on the VLAN and blow up your entire storage sub-system. Not that that has ever happened to me but I've heard of it.

**Richard Campbell:** But it is a risk.

**Alan Sugano:** It is definitely a risk.

**Richard Campbell:** Have you got the sense that direct attached storage has sort of fallen out of favor?

**Alan Sugano:** You know, again it really depends on the price point and I talked to a couple of hardware vendors about this and I say, you know, you guys need to come up with like a low cost, whether it's iSCSI, Fibre Channel over Ethernet, or I don't think it's going to be fibre channel because it's going to be too expensive, but price point where people can get into the storage area and networker's SAN stuff at a decent price point. You know, even with iSCSI, I mean there are some vendors out there, oh, we have iSCSI sandbox for 10 grand, but by then you end up putting everything together and configuring it and being the TOE card, you're still at like probably 20 grand for something decent with a couple of terabyte storage.

**Greg Hughes:** A lot less that it used to be but still quite a chunk of change.

**Alan Sugano:** Yeah. So I think direct attached storage advantages of it are, you know, it's relatively cheap compared to a SAN and the performance is pretty good. You'll get actually, with the SAS rate, you'll get better performance on direct attached storage than you would on iSCSI in general just because the transport is nothing else that needed transport as faster. The iSCSI SAN typically runs it, assuming you're running gig Ethernet, it runs at one

gig per second and with SAS stuff you're looking at three theoretically although the trios can never keep up with that yet.

**Richard Campbell:** So for those who are confused, can we define out the difference between iSCSI SATA and SAS?

**Alan Sugano:** Yeah. Well, iSCSI, the old parallel iSCSI which would be kind of the Ultra 320, Ultra 160 iSCSI, blah, blah, blah, that is legacy stuff now and basically what it is is, well, Ultra 320 runs at 320 megabytes per second and it's a sheared bus and you can have, let's see, Ultra 320 I believe you can have 14 drives on the iSCSI bus...

**Greg Hughes:** Right.

**Alan Sugano:** And every drive is assign a unique ID, but the disadvantage of Ultra 320 iSCSI is that all over the traffic that goes along the bus, it's seen by all, every device on that bus.

**Richard Campbell:** Right.

**Greg Hughes:** Yeah.

**Alan Sugano:** So you can kind of think of it kind of like a hub so like a hub, you know, an Ethernet hub. When you broadcast traffic on a hub, everybody sees the traffic.

**Greg Hughes:** Sure.

**Alan Sugano:** It's essentially like a gigantic repeater. So the disadvantage of that, especially when it comes to things like rebuilding failed RAID drive and a RAID 5 Array, because it has to share the buss, you can't have simultaneous reads and writes because everybody has to wait for the buss to clear before they can send their information.

**Greg Hughes:** Yup, yup.

**Alan Sugano:** Contrast that with SATA or SAS, both of them basically has the same signaling technology, it's just that say there's a serial attached, SAS has serial attached to iSCSI and say it's serial ATA, typically SATA drives are used for workstations although you can use them in servers. SAS drives, which are roughly about 60% faster, are typically used for servers. Now, SAS and SATA, the biggest advantage there is that every device that's connected to the host adapter is on a separate channel, so it's more like a switch, and the advantage of that of course is that you can transmit data simultaneously across drives of the host adapter.



**Richard Campbell:** So it's as if every iSCSI drive was on its own channel entirely, they're totally isolated from each other.

**Alan Sugano:** Exactly. So where you really see things, the advantages of that, again like in the scenario of a RAID 5 Array where you have one drive that failed, you would see rebuild times significantly faster than you would Ultra 320.

**Greg Hughes:** So that brings up an interesting topic. Let's do our RAID for dummies for a few minutes if you don't mind. I know that in working for a lot of IT organizations, there are misconceptions, misunderstandings, and a lot of assumptions are made about different RAID configurations. I mean it can make a really big difference. From availability, from performance, from a data recovery and integrity standpoint, what RAID configuration do you use on a system? Would you mind just taking a couple of minutes in explaining what the different ones are and what each gives you?

**Alan Sugano:** Sure. There are actually quite a few different RAID levels. I'll just kind of go through the most common ones.

**Greg Hughes:** Sure.

**Alan Sugano:** At the very basic level, there's RAID zero which is basically no full tolerance. It's typically what you would see on a workstation or something like a JBOD which is short for a just a bunch of disks and the information has spread across all drives so if you have a RAID zero, 3-drive RAID array, let's keep it simple and you have 10-gig drives, then you would have a total usable space of 30-gig. So you would see one logical disk that looks like 30-gig drive although it would be striped across three 10-gig drives.

**Greg Hughes:** Right.

**Alan Sugano:** There's no-fault tolerance. The advantage to that is that it's pretty fast because you have three drives working for you, but the disadvantage is that if you lose any drive in that array, the whole entire enchilada would crash on you.

**Greg Hughes:** Correct, right.

**Alan Sugano:** And you would be very sad so it's not very good for servers and we don't recommend it, although sometimes we see it and if we do it's a red flag like you need to kind of not do that because if you have one drive failure, you're going to lose your server.

**Greg Hughes:** Sure.

**Alan Sugano:** So that's RAID zero. RAID 1, now it has changed a little bit with SAS and SATA drives but RAID1 is also every search was drive mirroring or drive-duplexing. There used to be a distinction between mirroring and duplexing. Mirroring was if you use the drive on the same channel, and a duplexing is this what if you need the drive on two separate channels. Today it doesn't really make that much difference because if you're using SAS or SATA, there are always undedicated channels there for the theoretically already duplex. It made more of difference when you're on Ultra 320, but basically what that is is you have two drives that essentially mirror the same information. So taking the example of the 10-gig drive, so if you have two 10-gig drives you have 10-gig of usable space and if one of the drives fail you could still continue to work on the remaining drive. So that's RAID 1. The best use that we use for RAID 1 drive, it's very, very good for sequentially written data so typically that's going to be one of two things, exchange logs or sequel logs.

**Richard Campbell:** Right.

**Alan Sugano:** And actually we also use it for loads of OS's for clustering hosts so a lot of times we load the OS on a RAID 1 Array. Typically, we try to avoid booting from the SAN when we like to load locally. It's a little bit more work to do that but the advantage of that is if you make a boo-boo on the boot partition on your SAN, you would take down everything simultaneously...

**Greg Hughes:** Yeah.

**Alan Sugano:** Rather than just if you make a boo-boo on one host, you would just mess up that one host but you could still boot in the remaining host.

**Richard Campbell:** Plus this fear of moving parts too, right, that I just get scared. As soon as you're booting a machine from something external to the machine, you're further and further away from recovery. It seems like there's a lot of moving parts to get a machine recover.

**Alan Sugano:** And I agree with you 100%. You know, when you're spending this much money for a SAN infrastructure and high availability, this is one of those you can't get a little pregnant arguments.

**Richard Campbell:** Yeah.

**Alan Sugano:** Again, it's a little bit more work upfront but kind of our philosophy is we'd rather do more work upfront to make our lives easier on the back-end rather than the other way around because whenever we stray from that, it always ends up -- it turns around and bites someone on the butt. It just



never fails. I don't know, I think Murphy is listening and realizes, ah, he tried to cut corners there so I guess what I mean is it's going to nail you, you're going to end up barbecued on the backend so we try to avoid that as much as we can.

**Richard Campbell:** RAID 10?

**Alan Sugano:** Actually RAID 5 is the next thing.

**Richard Campbell:** Okay.

**Alan Sugano:** The most common. RAID 5, we do a lot of RAID 5 and also there's another thing called RAID 6. Essentially what that is is you need at least three drives in the array and you have the overhead of  $N - 1$  so every drive has data and it has parity information that can be used to rebuild parity information on each scribe that can be use to build information if one of the drives fails. The advantage of that is you don't have the overhead quite so much of RAID 1 or even RAID 10. The disadvantage is that you can only lose one drive in the array before everything blows up. So the example there if you have three 10-gig drives, you would have 20 gigs of usable space for example and you could lose one drive in the array but if you lost more than one drive then the server would crash. RAID 6 is a kind of a different spin on RAID 5 where you basically have two sets of parity instead of one set of parity on the drive which allows you to lose two drives in the RAID array and still keep running. We typically don't use RAID 6 too much and the main reason why is that you have the overhead of writing two sets of parity in each drive. So it's not quite as fast as RAID 5. If the client is concern about that, what we typically would do is we would set up RAID 5 worth of one hot spare so you have the full tolerance of RAID 6 but you don't have the overhead of...

**Richard Campbell:** Double parity.

**Alan Sugano:** Writing two sets of parity. The disadvantage of that though is that you actually have to, if the drive fails, then you have to rebuild the information on this hot spare, but it's not quite as bad as it used to be with Ultra 320, with the new SAS and SATA drives, because the rebuild process work happens so much faster.

**Greg Hughes:** Right, right.

**Richard Campbell:** But I think there's another angle on this, which is also that when you're rebuilding in a SAS infrastructure, the existing drive still functions well. I found once you're in a rebuild process on iSCSI rig, your whole system is degraded because most of the bandwidth is consume in the rebuild.

**Alan Sugano:** Exactly and the main reason for that is the shared bus as oppose to the dedicated bus.

**Richard Campbell:** Yeah.

**Alan Sugano:** On a rebuild with Ultra 320 on a RAID 5 Array, you're going to hear about it. Trust me.

**Richard Campbell:** Oh yeah, everybody knows.

**Alan Sugano:** Everybody knows, yeah, yeah. Everybody knows the infrared doesn't have any clothes on. With SAS on a rebuild, they might know this. You know, it seems like things are running a little slow. Maybe it's just my imagination.

**Richard Campbell:** Right.

**Alan Sugano:** As oppose to getting phone calls off the hook.

**Richard Campbell:** Yeah, it's abundantly clear.

**Alan Sugano:** Yeah. So the last most common one that we would use is RAID 10 and that really is kind of like the Ferrari RAID Array. Basically what it is is it's called RAID 10 or RAID 1-0 because you have striped drives across one channel but then they're duplex across each other. So if I have a client that comes to me and says, hey Alan, I want you to build the fastest, normally is the most full tolerant, most expensive drive array, I'm going to say RAID 10.

**Richard Campbell:** Right.

**Alan Sugano:** The advantage of that of course is you have -- the disadvantages, you have 50% overhead so if you have let's say six drives in a RAID 10 Array, there are 10-gig apiece, you would only get half of that storage so you only have 300 gigs drive but you could lose three drives in the array and still keep running.

**Richard Campbell:** But it doesn't have to be the right three?

**Alan Sugano:** It has to be the right three. So it has to be one that channels or two channels, two drives on the -- well, actually, you know, it depends. Knock on wood, I haven't had more than one drive go out simultaneously in a RAID 10 Array so I guess it's something I could test one of these days but probably if I have some time. Theoretically, yeah.



**Richard Campbell:** Drives are very reliable now. As long as you have some kind of redundancy, you're going to survive.

**Alan Sugano:** Yes, yes. One of the things I should point out is that we do see some people that use SATA drives for servers which is fine but it doesn't have the performance of SAS and we always try to go SAS when we can.

**Richard Campbell:** Right.

**Alan Sugano:** The thing about SATA drives, you have to be careful if you use them on a server, it's you have to look at the duty cycle and the duty cycle in relation with the Mean Time Before Failure, about the MTBF, because a lot of times the SATA drives are rated out of maybe a 50% duty cycle because they're designed to be in a workstation. The duty cycle of the total running time is how much time of the drive is actually being used. So let's say a workstation what might be use us 50% of the time. On a server, the duty cycles are a lot higher so the duty cycle on a server might be 90%. I don't know where the drive is going to be using 90% of the time and sit idle 10% of this time. So we have seen some scenarios where people would put in say a SATA drive into a server and look at the Mean Time Before Failure is five years and they think they're good but the duty cycle is only 50% so then maybe after a year-and-a-half to three years the drives start failing and they're dropping like flies and they wonder what the heck is going on, and that's usually worse. The bottom line is if you do use SATA drives for servers, make sure that you look at the duty cycle and the Mean Time Before Failure. That's something that a lot of people overlook and then they wonder why their SATA drives keep going out, but in general we try to even stay away from SATA drives for servers and we like to go SAS just because.

**Richard Campbell:** Any thoughts on SSDs for the enterprise yet?

**Alan Sugano:** That is a really good question. Have you seen that sound from video?

**Richard Campbell:** I have seen this sound on video with 24 drives, that's awesome and I've been reading the stats on Intel's X25. I mean, I think the big concern that most people have around SSDs these days, at least in the enterprise level, is that they just don't have the lifespan.

**Alan Sugano:** Exactly.

**Greg Hughes:** How many reads and writes.

**Alan Sugano:** Yeah, they don't have a lifespan and they're still expensive but they do have

lower power consumption and they're a lot more shockproof than the traditional mechanical drives so I think the biggest advantage of those is right where they are is switches and laptops.

**Richard Campbell:** Right.

**Alan Sugano:** You know, a year from now I might be singing a different story saying like they are a lot more reliable than they used to be and they've come down in price. The other thing that we're seeing with the Solid State Drive is that the density isn't quite there yet in terms of the maximum size of the Solid State Drive you can get. Even today, it's not where it needs to be for service storage.

**Richard Campbell:** Especially at that price point. I mean, there are 256-gig drives out there but boy, oh, boy, are they costly.

**Greg Hughes:** Yeah, they're costly.

**Alan Sugano:** Yeah, I know, yeah, they're worth a million dollars apiece, but you know, I might be singing a different tune in a year from now and that would kind of be cool if they come down in price and they do prove to be reliable.

**Richard Campbell:** Yeah, I think that's the big whammy for me, it's am I really going to get two or three years out of this before it fails.

**Alan Sugano:** Right and you don't want to be the guy that sticks out their neck to go, hey, I think we should go Solid State and have it blow up and then have a gun on your face.

**Richard Campbell:** Oh yeah. Well, that's the IT mantra. Change is good, you go first.

**Alan Sugano:** Yeah, that's right. Go ahead, take a jump off the cliff, let me know how the landing was and just get back to me on...

**Richard Campbell:** Yeah, get back to me and we'll talk.

**Alan Sugano:** Yeah, hopefully if you can still talk, or as Mark Minasi likes to say there might be an RGE, a Resume Generating Event.

**Richard Campbell:** Yes. So when it comes to SANs, I just find that folks, it's almost like they read the brochure and they just believe, because I've certainly found in my test and working in companies that are struggling that SANs means feature is not performance, it's reliability and folks forget that.



**Alan Sugano:** Yes, so in some scenarios, especially in iSCSI SANs, you'll actually see direct attached storage app perform iSCSI.

**Richard Campbell:** Sure, especially by 10's.

**Alan Sugano:** Oh yeah, definitely. Really, the main reason for a storage area being SAN is for high availability so it typically would be combining with some type of cluster solutions, Microsoft Cluster or a VMware cluster or something like that, but really the advantage of the SAN is you could -- basically what direct attached storage is only the physical server that's connected to it can see the storage. What the SAN advantage is that two or more nodes can actually see the storage so it's basically shared storage but to make that happen you have some significant overhead to do that and that's why SAN is, you know, as a Fibre Channel SAN, all in, you'll probably, even a low-end one today, you're still looking at probably about 30, maybe minimum \$30,000 to \$50,000 to do something halfway decent. You know, not to say that I'm sure there's going to be a SAN where they can go, oh, we can do fibre channels for 10 grand but you know it's like...

**Greg Hughes:** Sure. Well, I can remember spending hundreds of thousands of dollars for, what was it, a couple of terabyte storage like nine years ago or eight years ago. It's a different story now.

**Alan Sugano:** Yeah, it's a different story. The other reason why I think there is a need, and it's coming but still relatively slow, in a decent, reliable, good entry SAN is for virtualization because when you get a host that's running 5 or 6 or 10 virtual servers on it, that goes down. It's a big deal. It's like losing 10 servers all at once.

**Greg Hughes:** Sure.

**Alan Sugano:** So that's where the high availability would come in for a sturdier network, it's when you hand that scenario you could fail over to some waiting hardware so that you won't actually go down, but then you need a SAN to do that unfortunately....

**Greg Hughes:** With the SAN, I guess you could do geographic distribution or redundancy of the SAN and the whole concept of being able to sync and split to do back-ups and take all your data offline and get a real snapshot and do that kind of thing, it's all stuff that I think it's fairly unique to a SAN but those are all those extra overhead that goes along with those capabilities to give you that advantage.

**Alan Sugano:** Yeah, yeah and you also have the overhead of even learning the SAN operating because the SAN itself typically has a built-in

operating system. The SAN interfaces are getting a lot easier than they used to be. I mean, when the first stuff came out it was so cryptic and complex, we just let the SAN vendor set up and burn the LUNs and stuff like that.

**Greg Hughes:** Right.

**Alan Sugano:** But now, you could almost figure it out yourself. You might need a little bit of help but at least you have a chance to do it, whereas before it's like, okay, I know enough not to touch this so I'm just going to let the SAN vendor set it up for me.

**Richard Campbell:** But I also get to sense that people abuse SAN. Once they get one, they believe this is the great, master storage and they can add too many servers to it and then they complain about performance problems because in the end there is so much competition for spindles and in some ways some vendors love that because the answer is buy more.

**Greg Hughes:** Yeah...

**Alan Sugano:** Yeah, it's like, oh, we can't fix this, it would just be a million dollars to fix it.

**Richard Campbell:** Yeah, just as easily solved.

**Alan Sugano:** And for your request, then SAN world domination or whatever it is, but yeah, you're right which brings up another point, you know, it's very critical when you're sitting your SAN and how the LANs are burned or how the LANs are configured because typically one LAN will logically correspond to an array on the SAN. What we typically see is we'll set up, again it depends on the SAN, but we'll set up like some high performance LANs and then some medium performance LANs and maybe use in some blower kind of near-line storage LANs depending on what the client is doing, and the high performance stuff would be set up with either RAID 1 or RAID 10 using SAS drives. The medium performance LANs would be RAID 5, RAID 6, or RAID 5 with the hot spare using SAS, and on the low-end stuff it would probably be RAID 5 with SATA and typically a lot of vendors delight in mix and match drive SATA versus SAS on the same SAN now which is kind of nice. Then you can kind of plan, okay, well, because you don't need Ferrari performance typically for every single thing you do but you could probably be going to need it for your SQL Server or your SharePoint Server, probably Exchange, and maybe a really high loaded file server, but for like mail archiving that would probably be low performance. If you have a bunch of pictures that need to be archive, that could be put on the SATA lower performance LAN. So it really does require quite a bit of planning, and then if



you're going to be rolling out a Microsoft cluster, there are certain LANs that have to be burned before the core driver and things like that that you need to plan for when the SAN is going in.

**Richard Campbell:** Yeah, I do get to sense that people often -- like you said it's \$50,000 just to get started with the SAN, and I think really it's more than that because if you have at least two heads and a couple of extra drives for each head, you're fooling yourself in terms of SAN performance.

**Alan Sugano:** Yes.

**Richard Campbell:** I think people jump away from direct attached storage too quickly or simpleralized iSCSI solutions.

**Alan Sugano:** Yes, yes, I agree. I mean, usually when we get a shelf for a SAN, we'll fill it up with tries and question why do we need to fill it up, it's like, again it's one of those things you can't get a little pregnant.

**Greg Hughes:** Right.

**Alan Sugano:** It's really kind of neat to fill up the shelf because then it affects how you burn your LAN, and if you only fill it up halfway then you end up having two LANs that are not big enough to really do anything that you needed to do, blah, blah, blah, so I think you'll end up pulling the LAN, rebuilding it, restoring the data, and that cause you an extra couple of days of work to do it the right way. So whenever we get a shelf on a -- a shelf would be a logical chassis which typically holds 14 drives to 15 drives. We usually fill up the entire enchilada with drives or just done with it and we found that although...

**Greg Hughes:** Kind of better if it's built there.

**Alan Sugano:** It's a little more expensive upfront. We never regret doing it on the backend.

**Greg Hughes:** Right.

**Alan Sugano:** We never -- that never happened.

**Greg Hughes:** You're buying a SAN for flexibility of storage and be able to create things on the fly so, you know, shortchanging yourself for a few bucks for a couple of extra drives could be pretty unfortunate.

**Alan Sugano:** Again, it's that you can't get a little pregnant argument there. It doesn't make sense not to fill up the shelf with all the drives just because you're going to regret it.

**Greg Hughes:** So the other head cost, if you will, to SAN is having one more people around it that know, as you mentioned before, how to do all that work, how to make it shine, how to make scream, and even just how to configure it depending on the vender of the SAN so there's overhead associated on the people side as well.

**Alan Sugano:** Yeah and then, you know, back-up and then it's not just the SAN itself, it's all the stuff supported around it, so yeah, for back-up typically. People snap the SAN to JBUD for back-up, and then from JBUD they'll spin it to Tape so you have the overhead of getting or purchasing a JBUD for your fan and then some SAN then just trajects through the SAN and that of course is very expensive, and then whenever you mention SAN with any type of storage, you know, it added zero or a couple of zeroes because software is more expensive, the hardware is more expensive, the drives are more expensive, it's just more expensive.

**Greg Hughes:** And you can't get real benefits out of it that can make the investment worthwhile but it's never -- it's like buying a boat or an airplane, it's going to cost you more than that front piece so you need to plan for that.

**Alan Sugano:** Yeah, maybe not quite as bad but close, well, maybe a small boat for it's just a small hole in the ocean that you have to fill up, as oppose to a big one.

**Greg Hughes:** Right.

**Richard Campbell:** Yeah, there's a hole in your datacenter you have to keep shoveling money into.

**Alan Sugano:** Yeah, yeah. One of the things also that, just SAN and clustering in general, that we've found is Microsoft clustering with LAN, you can have potential owners. So with Microsoft Cluster, you can have two or more hosts, the potential owner of the LAN, but you can only have one host that's active on the LAN at any given time. So in other words, the active host has to release the LAN before the back-up host can take over the LAN. On VMware clustering, it's a little bit different. You can have two or more hosts and this is really the way the file system is written. You can have two or more hosts have simultaneous access to the LAN, they can have simultaneous access to a file on the LAN. Where that really comes into play is if they're using clustering in virtualization even with Hyper-V, that's why in Microsoft they have -- I don't know if you're familiar with the quick migration with Hyper-V...

**Richard Campbell:** Yeah.



**Alan Sugano:** Where you actually you can move one machine over from one host to another host but you actually have to shut down the host and then bring it back up on the new host. You have to shut down the virtual server guest and then bring it up on the new host, and the reason why that is is because you have to actually do a failover on the LAN in order to get that to work. So if you're using Hyper-V in a virtualization environment with clustering, the other issue there too is that you have to failover all of the contents of the LAN when you do a failover scenario. So the disadvantage of that is that if you have five machines on a single LAN, you have to failover all five virtual guests to a different host simultaneously. I call the entire enchilada failover because with Microsoft Cluster, you can't have simultaneous access to the LAN for more than one host. In contrast to that, with VMware it's a little more flexible because since you can have simultaneous access with the LAN you can do a granular failover. So you can have, if you have one host go down and it was running 10 guests, you can have five of those guests failover to one host, the five failover to another host and that's what makes the technology like the motion which basically is the hot migration of virtual server from one host to another in real time and what will allow that to happen is the way they wrote their file system which allows simultaneous access to the LAN. It's subtle but important difference in feature set, and supposedly Microsoft is coming out with -- the next release of Hyper-V is going to include hot migration and I could see why they didn't do it with this release, it's because if you look at the way the clustering works, you can't have simultaneous access to the LAN. The first time we set up a Microsoft Cluster, I thought I'd actually not set it up correctly until I really took a look and said, oh, I see, you can have possible owners but you can only have an active owner at any given time.

**Richard Campbell:** Right. Well Alan, I think we've run out of time. Any final words, call outs, hi moms?

**Greg Hughes:** Hi, mom.

**Alan Sugano:** No, not too much. I'm actually going to be speaking on WinConnections coming up this week in Orlando, Florida so if anybody gets the chance to go out there and they hear their show, come up and certainly say hi.

**Richard Campbell:** You bet.

**Alan Sugano:** I also have an article coming out in Windows IT Pro on "Exchange Recovery Soup to Nuts" so that will probably be in another couple of months, hopefully here, because we're finishing it up now.

**Richard Campbell:** Awesome.

**Alan Sugano:** I appreciate you guys having me on the show and hopefully the listeners have got some -- now, they're thoroughly confused, they're out of storage.

**Richard Campbell:** Yes.

**Greg Hughes:** Better educated.

**Richard Campbell:** How we properly worried you about your storage.

**Alan Sugano:** Yeah.

**Greg Hughes:** Thanks for dropping by.

**Richard Campbell:** Yeah, thanks for coming on the show, Alan. We'll talk to you again.

**Alan Sugano:** Okay, great.

**Richard Campbell:** And we'll talk to you next week on RunAs Radio.