The RedZoneCIO Podcast

RedZoneCIO Podcast Episode 1

Jason Kasch Interview

Bill - Intro: Hello and welcome, this is Bill Murphy. I want to welcome you to this podcast where I interview Jason Kasch, who is the Chief Information Officer at the Structural Group, which is based out of Maryland. Jason and I have known each other for many years. The company is focused in the construction industry, and has 30 plus or so offices around the country. Jason works with around 4,000 employees, 2,000 of which are mobile, which he'll get into this podcast. What we focus on today is talking about how he used file synching technologies to solve problems that he faced both at the CEO level all the way down to pulling off his ultimate strategies with disaster recovery, data security governance, and high availability. It was really interesting talking to him because he's one of the most innovative thinking CIOs that I know, brings a lot of creative juice to solving really interesting problems. He also sits on the board, and I've been in those meetings where I've seen the interaction at the board level with the CEO and the other executive members that sit on that board. He's also quite a heavyweight technologist himself and really pushes his staff to come up with creative solutions. There's quite a lot of ground that we cover and this is a really powerful set of solutions that he came up with to solve this particular problem. You can see the show notes, and links out to the specific products we talk about. I'll post the show notes on the blog site so you can refer back to them when we refer to vendors and books and interesting things that he's reading. I hope you enjoy, and welcome to this podcast.

Intro Music

Bill - 2:27:

I thought it would be good for us to talk a little bit about how your mind works in regards to putting together creative solutions to problems that are, I think in the industry would be considered pretty complex or that normal CIOs might lose their way thinking through the problems or maybe end up overspending or buying too many technologies to solve the problem that you essentially solved with one solution set. If you don't mind, if you could share a little bit about what challenge you had initially, and how that kind of flushed its way out as you went along.

Jason - 3:00:

Sure. Initially we had a couple of challenges that we were trying to solve and I didn't realize that they were all interlinked until we got a little bit further down the path. We had a couple of challenges; one was deployment of applications the other was access to data in our environment and a third really was how do you deploy devices faster and more efficiently. We headed down this really heavy path with a product call VMWare Horizon, which solves each of those problems kind of independently, but not well cohesively. We looked at this and we broke it down into 3 things. We said, the three things we're really trying to solve are: deployment of desktops, how fast can you turn around a desktop to an individual when they either lose it, or it has to be procured because it's brand new; the second problem was how do you give access to company supplied applications, things as mundane the office applications, Word, Excel, PowerPoint, Access, and as highly specialized as AutoCAD, or Mathcad, or some very specific engineering apps; and then the third problem was how do you give access to the data that these specialized applications need. We looked at it, we said really those are 3 independent problems we wanted to solve and we started to head down a path to solve each one of them. We started to look at products that could solve those, we started to hangout in some user forums and we got turned on to a couple of products. VMWare Horizon was one, there's a really cool product called Nasuni which solves some of these problems really interestingly, and this other product called DATAnywhere which solved the third of those, that problem, the one around data access, very elegantly.

Bill – 5:10:

Gotcha, Gotcha. Essentially, you wanted to accomplish 3 things: deploying application, access to data, and being able to push your applications and your data down to as many devices as you can and as efficiently as you can. How many devices do you have Jason?

Jason – 5:28:

We have probably, between smart devices, tablet devices, and work stations, I would say upwards of the 2,000 number.

Bill –5:40:

Okay, so it's not a small number of devices you have to manage. Now you actually engaged in a pilot right, of the VMWare Horizon? You didn't just look at the marketing literature you actually tested this out right?

Jason – 5:52:

We engaged VMWare directly through a 3rd party, a partner of theirs. We paid for an engineering proof of concept, for about $20,000. The proof of concept ran a good 3 months, we put together the type of use case scenarios we wanted to get solved and how we would see them solve it.

Bill – 6:20:

That's great. Is that a typical R&D budget that you have each year or is that a relatively new development for you as far as being able to not just look at the marketing brochure, download a piece of software and kick it around, but you went down a serious effort of discovery. Is that a normal budget for you for that kind of stuff?

Jason – 6:42

No, not generally. It came about because that product, Horizon, has a lot of hair on it. It's a really big product and it requires a lot of time and effort to get in your environment. VMWare is not very keen on giving their stuff away to try it out so it turned into us requiring to do a paid proof of concept, but as we've walked down this path I've found myself, in the last year to 18 months, doing more and more paid proof of concept based engineering arrangements much smaller in price and scope. I've found that this has been a recurring trend where up until about 18 months ago we really didn't do paid proof of concepts. We would ask either a vendor or a 3rd party integrator to come in and say here's our data, show us how you

would implement this. The process of actually putting it on your own hardware, putting it within your firewall, having your hands on it, be able to touch it, and feel it, and smell it, and see how it reacts, those have become paid engagements, and I think it's an area we're going to carve out of our budget to actually have moving forward. I think we got a lot more out of it, and we understood it a lot better than if we had just had an organization build it and then show us how it works.

Bill - 8:09:

You probably get a lot more of the heavyweights on the vendor side too I would imagine. They're probably lining up a little more seriously by doing that too. Did you find that the case?

Jason – 8:19:

I did, especially in the case of this VMWare. I mean the $20,000 or so that it cost for us to do the engineering POC, that product carried a half a million dollar year one price tag to it. It wasn't insignificant from a mid-market IT budget. We wanted to make sure we really got the product right and at the end of the day what we found was it didn't solve any of those 3 problems that we wanted to solve.

Bill – 8:49:

I want to keep moving down the other products because I think one of the things people need to be aware of is Citrix and VMWare; they're definitely application delivery technologies. I also found VMWare to be a challenge from an installation point of view, but I know what we're going to get into, in my opinion, turns the whole discussion on its head, but before we go on to the DATAnywhere let's talk about Nasuni. How did you run into Nasuni and how did you uncover them?

Jason – 9:23:

We ran into Nasuni as we were doing a bunch of web searching around this data access type of problem. Nasuni's head kind of popped up in all of the research we were doing and number 2 or number 3 player in that space. What they do very differently is take the data that you have in your environment and they replicate it. Hey use this term called consistently, I can't even remember the name of the term now, eventually consistent! You put a series of these devices in your environment and they talk to each other, and they eventually, over a series of days, weeks, months, however long it takes depending on the amount of data that you have, eventually they make the data consistent amongst all of those devices, and the they just move the delta changes between the data from site, to site, to site. The really interesting thing that

they do is then they leverage AWS, Amazon Web Services, and they have an instance of your data at AWS, but part of the challenge with their product is it doesn't use your inherent security model. It doesn't use AD; it doesn't use your security model, so you have to keep a duplicative method for keeping all of that information in sync. Then what really concerned me is once it's out at AWS, if we forget to remove an employee when we remove them out of AD, what happens to that? Do they still have access to that data? There are a lot of questions that they left and it also turned out to be a rather significant investment when you've got- we have 30 small to midsized offices, they have about, you know, anywhere from 10 to 40 people in it, so they're not overly huge, but we have a number of those scattered throughout the country. I think that product would probably be more relevant to someone that's got a couple of really large locations versus a bunch of small locations.

Bill – 11:33:

I think also part of your skill is you have a computer science background as well so you're able to really get in to the guts of when your team brings some of these challenges up you can really question them quite deeply on these products. Did you find that the case?

Jason – 11:49:

I do, and I found it helpful. Actually, the team that works for me is way deeper in technology understanding than I am at this point. I think the value I bring to the table is having a background in it, having done it in years past. A lot of these technologies aren't revolutionary. They come round, and around and around. We go from client server, to centralized, to distributed, to centralized, and that circles. That circle keeps going around and around and it's just, it's the names of the products and the way that they solve the problems are different, but the actual under the covers, it's either centralized or it's distributed. It's one or the other.

Bill – 12:35:

This is interesting. You have a real deep understanding of the technology, but then also, because I know from working with you, you're working right with the board at the top and you're actively engaged with them, understanding the business, where it's going. I know you actively sit with the CFO. You're very fluent with finance as well as the technology. I am particularly interested in understanding the challenge the CEO was having, and your approach to solving this problem. Can you share that a little bit?

Jason – 13:10:

The CEO had a concern that company files were residing in personal cloud systems, Drop Box specifically, and we didn't have a solution for that outside of the Microsoft platform.

Let me back up a little bit. We deploy and deliver computing devices centrally through IT. We don't have a formalized bring your own device policy or procedure, but what we do know is that people are definitely using them, and they're multiplying like rabbits in our environment, and we're forced to support that environment. Now, we should be supporting it because our job as IT professionals is to enable the business. I think that's one of the advantages when you ask, how am I able to speak fluently with them, it's stepping out of technology, for technology's sake and look at what those individuals do, and why they want to do it, and how they want to get it accomplished, and then try to match technology with it.

Our CEO is a typical CEO. He doesn't use a ton of customized applications that IT would generally support. He's not using AutoCAD every day. He's not using Mathcad every day. He's a civil engineer by trade, but he's also a CEO. He's very email heavy, he's very internet heavy, and really that lends itself to a tablet based device, and his particular tablet of habit is the IPad, which isn't a device that we have in our environment. We're a Microsoft application, and Microsoft software organization. We have general laptops, Surface devices, everything that's centralized around the windows platform, and he just doesn't use it. He had some concerns around sharing data because he does it today. How does our CEO share his data amongst his own devices when he's using a platform that we don't even support internally?

He does it through a 3rd party application, or what I would consider a consumer based application, with Drop Box. He came to us with some heavy concerns because he's an incredibly heavy user of Drop Box. He recognized that as he went around the company and asked about how he could share files with someone and how they shared files with him, it continually came back to this, well just put it on Drop Box, just put it on Drop Box. He and I had a very lengthy discussion on what that meant from a data governance perspective. Data governance, meaning as an IT word, he didn't say that. What he said was he had a concern about company files which were sitting in personal cloud systems.

Bill – 16:15:

Yeah, what's fascinating is how it's married from top to bottom. For our listeners, there's one other piece that I think people should know was sort of how I envisioned you being not only the hub for a 4,000 employee company and IT, you've got 30 locations and quite a mobile device challenge, but you're sort of sitting in the hub also of inputs coming at you, like how to solve this problem. You've got your CEO using Drop Box. You've got myself and my team recommending different systems and solutions for disaster recovery. You're looking at Nasuni, and you're looking at VMWare, and you're coming to decisions that I thought were very, very creative in the end. Maybe that would give us a good opportunity to talk about how you had started to pull this together, and how you ended up solving the problem with your CEO, your

disaster recovery, your high availability, your technical challenges that you wanted with file sync and such. Maybe you could talk a little bit about that.

Jason – 17:20:

Yeah, as you know we engaged RedZone to help us understand our disaster recovery, at the end of the day that's what we called it, but really it was about polling all of our senior executives about a) what they thought was important in the environment, and b) where they felt that IT was in this given point in time. Then what RedZone did was put a map together for us that said, here's the perception of where the organization thinks you are from an IT perspective, and here's the reality as you spoke with the IT individuals, and here's the gap analysis. Some things that stood out for us very readily were we had an outstanding backup methodology, but our restore methodology, in the event of a full disaster, was weeks before we could even think about restoring one single piece of data. Now, IT generally, backup and recovery, you would think that every IT organization does that well, and we thought that we did it well. What we didn't recognize is, like I said, is that there were some gaping holes. That's one.

The other was, we had a company policy, if you would, that said we don't keep any data local. As IT, we kind of turned a blind eye to what that really meant. The reality is, we issued laptops to people, and they have data on their local computers. Whether we backed that up, or didn't back that up, in the event that laptop got lost, stolen, run over by a forklift, that data is still relevant to that individual. Even though we had a policy that said we wouldn't store it local, the reality is they were either storing it local, or they were coming up with a creative way to back that up on their own, which then made it lose the visibility of that data from the corporate veil perspective. Those were a couple of items that were gaping holes for us when we said, what are some of the strategies that we have to employ. Then another was, we've got these 30 different locations, how do you make sure that if you lose a location for a disaster, or even for an epidemic; the disaster is probably the easier one to solve because you just restore a backup exactly the way it was, but what if there's an epidemic and everyone in the office is sick, but they can still work. If all pf the data still resides at their facility, how do you still get them access to their data so that they can work at their leisure? Because not everyone's coming to the office because maybe there's influenza and they don't want to get sick by coming in.

Those were the multitude of swirly, really loosely defined items that we were trying to work with in combination with this challenge of personal cloud data, and for me, as we started to talk through these exercises over a 2 or 3 month period, that 2 things started to stand out for me as distinct problems that could solve a multitude of these challenges. One was, how do you take your existing data where it lies and present it on any platform, anywhere, and can you have it fully redundant. That was the first piece, it was the data. Then the 2nd piece was, how do you give access to these applications, because the data in its native form isn't really good if you

don't have an application that can open that data. Let's take the case and point of an IPad. Let's say that you have that data stored as an Access file, but the IPad can't run Access. Well at the time it couldn't because Microsoft 365 hadn't deployed that on the I-devices. You didn't have it, so then how do you deliver that application to get access to that data? That's what led us down this original POC for Horizon because it seemed like they were they only one that could do it.

Bill – 21:37:

Gotcha, gotcha; now, you've looked at VMWare, you've looked at Nasuni, and now you're on to DATAnywhere, but if I remember correctly, your first conversations with Varonis didn't go very smoothly, or it didn't go quite where you were expecting them to go right?

Jason – 12:57:

It didn't. My hat's off to James Crifasi over at RedZone, because he kept hammering us about that product. He actually seemed to know more about the Varonis product than Varonis did themselves. We had their sales and engineering individuals on the call and we asked point blank questions. They said to us, that's not how it works. We got off the call, and James called us back and said, you know guys, you really need to peel the onion back on this Varonis product because based on what you're telling me you would want it to do, you can get it to do that. It's just not obvious to Varonis on the other side as to what you were trying to get solved.

Bill – 22:50:

Yeah I remember that was an interesting discussion. Probably because both of you guys are computer science guys, and probably confused them.

Jason – 23:00:

I think James and I both saw it as we- I mean you guys have implemented it so you know how the product does work, but as I started to hear what they said my mind starts thinking, it can absolutely do what we want it to do. There's a couple of components that it's missing, but if it could do those couple of components then this product absolutely will do exactly what we want.

Bill – 23:29:

What I find interesting, which I was thinking about because I know security isn't on the top of your rolodex every single day for a whole bunch of reasons; not the least of which is, you're a construction company; it's just not something that's on your radar so we never really talked much about security, but it's really interesting at the end of the day how you ended up having not only a high availability, disaster recovery, solving the CEO problem, but also having

data governance. Which is partially the issue you mentioned with the Nasuni having your data potentially in the amazon cloud and how do you deal with de-provisioning users at the core of your network, but you in essence have 30 sites in which you now can govern data access, which to me is absolutely fascinating. Maybe you could talk a little about that because I think you've uncovered something that a lot of people are wrestling with right now.

Jason – 24:26:

There are 2 keys to making the whole thing happen. For us it was to be able to have multiple replicas of all of our data at multiple locations. That was the first piece. That solves the disaster recovery piece. It has nothing to do with Varonis. We actually leveraged Microsoft's DFS, Distributed File System, to do a 2-way, full synchronization replication between all of our 30 locations and our corporate data center. That gave us 2 things: it gave us quick real-time access from the people who sit at that location, who are working on the data so they don't have to traverse the entire network to pick up a file and work on a file. Some of those files are massive. They're AutoCAD files, they're video files, and they're gigantic. Our customers are not going to wait for that file to get down so that they can make a change on an engineering document so that data's got to reside right at their fingertips on their local LAN, and every one of our LANs at all of our 30 plus locations, all have a file and print server so they can work local, they can work offline so if a circuit goes down it doesn't incapacitate their business. We also have manufacturing facilities in 4 locations and if the internet circuit goes down and you can't manufacture your product, then you're going to grind the company to a halt.

The first piece was, how do I take that data, and then the other challenge was, how do I get them out of band access to that data. We did that through replications using DFS back to our corporate headquarters. Now we have 2 copies of that data. If they happen to be sitting a corporate then they get a local copy of that data from our corporate file server. If they're sitting in let's say Dallas, then they get the Dallas version of that. If we have individuals who travel from location to location, because we a couple of locations that as business units they operate together, so a Kansas City office that works really close with a New Jersey office that works really close with a Pittsburgh office, we can set up a 3-way replication between those sites so that those sites have all of the same data at the same time, at their fingertips. That was the first piece of the puzzle that we solved that actually we stumbled on as we talked to Varonis, and started to talk about the way they implemented their tool. You spin up either a physical appliance or a virtual instance of this thing called Varonis DATAnywhere, it's a software product, and it attaches to your Active Directory so that it knows about your security rules and rights, and who can have access to what who can't. Then it web-enables, over a known port, all of that data, where it sits, without having to move it, and maintains your existing security structure. They've got an app that runs on the iOS device, on the Android device, in the metro

interface. You can get access to it through the web, through a URL, or through a syncing mechanism on a Windows desktop. Similar to how Drop Box worked. What their solution was, was to take that system and put it in 30 locations, and that gets really complex from a routing perspective, because how does an individual know whether that file system is in Dallas, and Denver, or Pittsburgh? Really, they could care less. What we said is, we can combine this DATAnywhere tool with Microsoft's DFS solution, and we can web-enable the consolidated DFS data that happens to sit in our corporate infrastructure, we can web-enable that. It's single instance. Whether you touch a file at our Dallas office, or you touch it at our corporate office, or you touch it from an IPad sitting at a conference, within seconds that file is now maintained consistently amongst all those data points. That solved disaster recovery, real-time availability, and bring your own device access because it's only using current AD credentials. We killed all 3 of those birds with one stone.

Bill – 29:25:

Right, right, that's fascinating. Very, Very important, what you found. I'm hoping that this is great information for everybody listening today. Couple rapid fire questions, what's your favorite non-IT, non-business book?

Jason – 9:45:

I have 2 teenage adults in my house as you know. Once they become teenagers, and they want to be adults, they don't want to talk to their parents very often. I read absolutely every book that they pick up, first and foremost, so I can have a conversation. You know, they're going sit quietly at the dinner table for half an hour, and then the minute you go to pick up something, now they want to talk; so I've read every popular, post-apocalyptic, teenage book out there: the Diversion series, *The Hunger Games*, *The Maze Runner*, all of them. They're exciting books, but my personal 2 favorites are Michael Crichton's *Prey*, and any of the *Jack Reacher* series by Lee Child.

Bill – 30:36:

Ah, interesting, that is great. That is a great idea on how to hang in there tough with your teenagers. I have one at the teenage years, but yeah that's a good point. Just read the darn books and then you have something to talk about right?

Jason – 30:51:

Yeah, because eventually they're going to want to talk about something and if you've already read ahead, so you're already past them, they can have that conversation, but then you can slide in, how was your math class today, you know and the dialog is already there.

Bill – 31:08:

That is great, I love it. What a great idea. Oh, perfect, perfect. Ok, what is your favorite IT book?

Jason – 31:20:

My all-time favorite IT book is *Business at the Speed of Thought* by Bill Gates, written in 1999. I think it's still germane today. I think if people were to go back and read it, he talks about a digital nervous system, which really in today's terminology, that's information and data governance. He was 15 years ahead of all of us.

Bill – 31:42:

For the aspiring CIO, someone that looks up to you and says, I want to be Jason when I get older, what's been the biggest thing that you would say to someone that you were mentoring?

Jason – 31:56:

There are probably 3 things. I think it would be to understand your company's data assets and their sprawl and what to do about it, because we forget that data is a huge asset for the organization that doesn't come and go like people do. The next would be to provide, or to figure out how to provide an always on architecture so people can work when they want to work, where they want to work, wherever that is in the world, and the third is to not forget about bandwidth. As data needs grow with movies, full feature high-definition training which we do in our company, large spread sheet files; bandwidth is going to get in the way eventually. Figure out how to get your hands around that so you can implement these things.

Bill – 32:52:

Great, that's great advice. This is great advice. I'm sure that the managers and wannabe CIOs are going to really look forward to getting those pieces of information Jason. Thanks so much for sharing your thoughts Jason; I look forward to catching up with you soon.