

## **INTERVIEW WITH RICH VINING**

- Eric Green: Hello and welcome to a new podcast series from IBM software that explores the challenges IT managers and business professionals are facing today. I'm Eric Green and I'll be talking with a range of experts to discover new perspectives, approaches and examples that can help meet these challenges and introduce you to the capabilities of smarter software from IBM. So let's get started. So I have the pleasure today of talking with Rich Vining, with Tivoli Storage Software at IBM. Rich, thanks so much for joining us today.
- Rich Vining: Oh my pleasure, Eric.
- Eric Green: So I wanted to start with asking, what are the key problems or challenges today for storage management?
- Rich Vining: Well we're seeing a couple of major trends. One has been going on for quite a while, which is that data is continuing to grow. More and more systems are becoming interconnecting, more things are becoming online, companies, organizations, governments, they're using data for more important things, and this data is just continuing to grow out of control. Also, people are making more and more copies of their data to make sure they can recover it, which is also increasing the amount of storage that's required.
- The second thing that we're seeing also is an increase in complexity in the IT environment. As new applications come on and new delivery models come on and new opportunities for business come, we're seeing, you know, a proliferation of different kinds of systems, different kind of operating systems, different kinds of applications, certainly seeing a distribution of these resources into different locations, different offices around the world. And we're seeing increasing requirements around some of these applications whereby, you know, you can't lose any data, you can't afford any downtime. Whereas other applications don't. So you're seeing a proliferation of different policies that have to be applied across all these different things. So that's making it very, very difficult for storage administrators and people who run storage as their job to manage all of these needs as they're growing.
- Eric Green: And what about dealing with data restore and data protection?
- Rich Vining: Well, yeah. So the more of these things that happen, the more these different systems come online or these different requirements

come across, they definitely complicate the data protection and recovery aspect of the job. You know, there's lots of different things that can go wrong, for example. Right? Somebody could accidentally delete a file. Or you could have a virus attack, where you don't find your database is corrupted for two weeks. You could have a server or disc crash, you know, hardware failure. And then you could have the normal disaster, you know, the fire, the flood, the nuclear meltdown, whatever it may be. All of these things have different requirements for how you restore the data following one of these events.

So it's incumbent upon the backup administrator or the storage administrator to find all the different ways that they need to apply recovery technologies, and as a result, a lot of them are implementing point solutions to handle all the different things that have to go wrong. And that in itself adds to the complexity, because now you have many different interfaces where you have to go. You have to get training for all of these different products. You have to have the right person with the right login to be able to find the right data at the right time to put it in the right place and not break anything when you're doing a restore. And all of that happens, of course, in an emergency. So all of this stuff is kind of boiling to a head and really become a huge issue for storage and backup administrators.

Eric Green: That's really interesting. So continuing from there, what do you see as changing with storage management and data protection?

Rich Vining: Well one of the biggest things that's changing, and it's adding into this complexity theme, is the proliferation or the increased deployment of virtual machines, virtual servers. For a long time, virtual servers were used for application testing, for quality control testing, for application development, those back office kinds of things. Now they're being more and more deployed into production server environments, application environments. You're seeing people running SAP and Oracle and DBII and all these critical applications on virtual machines. And that's great for the business, it's great for the IT environment in that it helps to reduce the number of physical servers they have to manage, and by it, it reduces the energy that's needed. It also improves application development time and deployment time. You don't have to go out and buy a server and take all the time it takes to get it shipped and installed and configured. Now you can spin up a new virtual machine in a matter of minutes.

So it's all great from that side – from the business. But on the storage side, it's actually making the world more difficult. Because just because you're reducing the number of physical servers you're using, you're not reducing the amount of data that you're creating. In fact, you're probably creating more because of all the extra servers that you're bringing online. And also the ability to spin up, spin down and move these virtual machines at very easy will creates another challenge for the storage administrator. What happens to the storage as it's – you know, as this virtual machine moves, what happens to the data? If you decide to decommission a virtual machine, what happens to that data?

So again, it's adding more complexity but that's the biggest thing that's changing right now because companies are scrambling to deploy these virtual technologies because they do offer such cost savings.

Eric Green:

So just as an interesting sidebar on that, I mean, it seems like it's quite the conundrum. It's always been a problem as code writers and the like continue to write these vast amounts of code and you've got these vast amounts of data whether they're applications or it's actually storing patient data and things like this. And I guess the fact that storage in general, I mean storage management is one thing, but storage itself has gotten easier, faster and cheaper, that it's kind of further led to the laziness or the increase in records and increase in applications and the like. I mean is that something that you're seeing also?

Rich Vining:

Absolutely. Yeah. So yeah, the storage itself is not the problem. Right? The capacity of the drives continues to double every couple of years. And as a result, the cost per terabyte or cost per gigabyte continues to go down pretty dramatically. So buying more storage isn't a problem in and of itself, but the management of that storage, the powering, the cooling, the floor space for that storage is becoming an issue. And certainly managing all that data that's being stored in those boxes is definitely an issue. So companies are really looking for smarter ways of managing their data. Even though, you know, maybe the overall cost of acquiring more storage isn't going up, companies are seeing that for every dollar they spend on storage capacity, they're going to spend ten to twelve dollars managing it, whether that be software, maintenance, what have you. So they know that they can't keep doing what they're doing because, you know, their data may be growing 40% a year but their budgets are only growing 5% a year. So there's

this huge gap that they have to address. They are looking for smarter solutions for managing their data and getting the data growth under control so they don't have to keep buying more storage.

Eric Green: And that's a perfect transition right there. So how in fact is storage management getting smart or more intelligent?

Rich Vining: Well, certainly the idea of data movement and migration, placing data on the right tier of storage is very important in this whole thing, and then moving it to the right tier based on where it is in its life cycle, to the point of actually expiring it or deleting it when it's done so that you're not taking up precious storage resources for data that you no longer need. That's part of it. But also, you know, a more holistic approach to data reduction. You know, eliminating a lot of the duplicate that you may be creating. If you're doing full backups every week, for example, that full backup is going to consist of probably about 80 to 90% the same data that you backed up the week before. And if you're doing that every week for say three months, you're keeping three months worth of backups, you've got, you know, tons of, maybe 12, 15 copies of the same data sitting on those backup tapes or those backup discs. And that's crazy. So take a smarter approach to that. Don't do full backups. You know, use an incremental only approach to backups, which would certainly reduce your storage footprint dramatically.

Data de-duplication is a big thing in the storage market to help, again, reduce duplicate data so that you don't have to store it. And there's a lot of technologies around that where de-duplication can be deployed in lots of places throughout the environment. Data compression. IBM just did an acquisition recently of a company that does real-time compression of network storage, and that's really effective to reduce the storage footprint. So lots of different ideas and technologies can be brought to bear to help reduce data. But then you also have to address the complexity of the IT environment and all the different policies and processes and different requirements that you have for data protection and recovery.

And what we've been working on is this concept called unified recovery management, where you can take all this stuff, all these different locations, systems, requirements, all the things that can go wrong, and deploy the tools that you need to deal with all of those things, but manage them all from one place. So one admin with

training on the entire system that knows what to do when something goes wrong and do the right thing in terms of doing the restore that's needed. You know, do you need to restore a file? Do you need to restore a disc volume? Do you need to restore an office? That can all be done now from one admin console.

Eric Green: Excellent. Very interesting stuff. And so I think the next logical question or step here for our audience is, I mean, you see this all the time with different types of customers and people that you're consulting with. Can you give our audience perhaps an example of a modern day storage management – how modern day storage management has improved efficiency, costs or whatever for an organization?

Rich Vining: Sure. So one great example is, we've got a customer, it's a major cabinet level department in the U.S. government. They've got over 400 locations with about 2,600 servers across those locations. And given the nature of their business, they were doing only full backups on the weekends. So they weren't doing backups during the week. They would run a full backup to tape on the weekends. That presented some challenges, obviously, because if you're only doing a backup on a Saturday, you know, and something crashes on Friday, you've lost an entire week's worth of data. And, you know, maybe the government can get away with that, but certainly a normal business cannot. Right? If you've lost the entire week's worth of data, your orders for the week, or all the development you've done on a product for a week. It would be disastrous.

So they needed a smarter solution for doing backups. And they implemented one of our products that helps them to do backups to disc on a regular basis. You know, in the background without shutting anything down. Because again, you don't want to do anything that's going to detract from your business. You know, you can't take your system down to run a backup if that system's critical. So we have solutions that allow you do snapshots or continuous data protection in the background that don't disrupt the operations. So what this customer is doing now is they're backing up their systems six or eight times a day depending on the system.

So now what they've done is they've reduced their amount of data at risk from a whole week's worth to just three or four hours worth. So that's a huge improvement. But what they're also seeing now is vastly improved recovery times. So it used to take them 48, 50 hours to do a restore of their Exchange server, their e-mail server. So you know, if they lost that server and they had to restore it

would take them basically two days. And now it takes them 6 to 10 minutes. So an instant restore of that e-mail system. Same thing with file systems. If it would take them 16 to 18 hours to do a restore, now it takes them 6 to 10 minutes. So pretty impressive, you know, restore times, which result in, you know, it's called downtime. Right? It's reducing the amount of time that whatever the application is not available to your customers. So very important, you know, things that we're bringing into market that allow customers to A) reduce the amount of data at risk, and B) reduce the amount of time that they're down when something does go wrong. So that's one example.

Eric Green: That's great! Thank you very much. And so continuing down that path, how else do you see IBM as continuing to drive innovation in this storage management space?

Rich Vining: Well, we're continuing to develop the products around a number of areas. Certainly the unified recovery management area and the data reduction area that I talked about we have a wide range of technologies and solutions that we've brought to market already. We're now investing pretty heavily around the virtual environment area. We just brought a new product out called Tivoli Storage Manager for Virtual Environments which really provides advanced types of backups for VM without impacting the virtual machines themselves. And that's always been a problem with virtual machines is that if you try to do your standard type in-guest backup like you would do with a physical server, but you're doing that across, you know, 15 or 18 machines within the same physical machine, that thing would come to a grinding halt and you wouldn't meet your backup windows.

So this new product basically does the backups without impacting the virtual machines at all. It does it without adding any new hardware for doing that, and it also provides some really great flexible recovery options. From the single backup that you're doing, you can do a file level restore, you can do a disc volume restore, or you can do an image, you know, a virtual image restore. And using that technology, you can move images around, you can, you know, use your backup application as your image management solution if you want to do that. So some new things in that area we're working on, we're continuing to invest in our virtual disc and our virtual server environments.

Eric Green: So we're running a bit short on time, but I think our audience would really benefit with any last pointers you might have. Perhaps some bulleted best practices for storage management?

Rich Vining: Sure. And again, we'll get to the two main areas where customers are having problems. One is data growth. And I think the first important step with data growth is to figure out what you have. You know, and we have some sophisticated tools, we've got some services that are available to help you do an assessment of what you have, where your data is, who owns it, how old it is, how long it's been since it's been touched, and make some intelligent decisions about what do to with that data. Do you keep it? Do you archive it? Do you put it into some secondary tier storage? Those sort of things. And then craft the policies around data life cycle management that will help that stay in place over the long period of time.

The second thing is, you know, stop duplicating data. You know, if you're using a backup system that's doing full backups, stop it, because you're just creating a lot of extra data. And there is a fix for that. There is the de-duplication appliances that the other vendors would be happy to sell you. IBM will too, we've got a de-duplication system. But, you know, it's not necessary. Those systems are really geared toward reducing all of the data that you're creating from full backups, so stop doing that.

The next thing is, yeah, do deploy de-duplication where it makes sense in your environment. Deploy compression, deploy archiving and HSM to move the data to the right tiers of storage. Lots of things that can be done to help you get control of your storage environment and stop the growth or at least slow the pace of the growth. The other thing too is get control of the management of all of this backup stuff. Many customers have more than one backup vendor. You know, they've either had one from a different department, or they've done a merger and acquisition, or just somebody came on board that had a favorite vendor and they brought that in. The easiest thing you can do simplify your environment is to get rid of anything more than one. Use one solution to manage your entire backup infrastructure across the entire environment, across all the different needs that you have and look to what we call unified recovery management to do that.

Eric Green: Well Rich, I want to thank you very much for joining us today.

Rich Vining: Well you're welcome Eric, it was my pleasure.

Eric Green: Thanks for listening. Please do visit [IBM.com/software](http://IBM.com/software) to connect with our experts, continue the conversation, and to learn more about smarter software from IBM. Let's build a smarter planet.