RT: Doug Hazelman is the Vice President of Technical Marketing at CloudBerry Lab in Florida. They're the provider of cloud backup software for built-for-cloud storage for companies such as Amazon, Microsoft and Google and the number one cross-platform cloud backup software.

As well as providing backup for corporate and ordinary users, CloudBerry also provide a managed backup service for IT solution providers. Doug is a veteran of the IT consulting space, having spent almost 10 years as Chief Evangelist at Veem Software. His role at CloudBerry sees him building a community of enthusiasts and customers.

Let's talk about the role of MSPs in backup. I'm a former MSP owner myself, and traditionally backup to me has been where an SMB (small and medium businesses) client might have had a tape drive onsite which they sometimes changed.

We've moved away from that now, and most of the MSPs I know use the cloud for backup, and I've heard the term 'backup as a service', so perhaps we can start by looking at how this works in the new world of MSP (managed service provider) offerings.

DH: Any MSP worth their salt is going to want to offer backup to their customers. It's a critical

component of any infrastructure, and there are multiple ways to go about it. I look at the history of backup software, and tape is still alive today, but what's happened over the years is it moved to disc.

What we're seeing now is backup to cloud, and one of the big reasons why is because if something happens to your customer's site, where everything including their backup is, you can't help them.

You need to get that backup offsite, and the cloud and cloud storage is the best way to move the data. It's much easier than you going to collect tapes and having to store them somewhere. The data is safe in the cloud in event of a disaster, and it means you can restore it and get your customer back online much quicker.

RT: I'd like to talk about the disaster recovery aspect of things, but I'd also like to look at how you're educating MSPs on the opportunities for their business with backup. There will still be lots of clients out there using tapes and so on, so what does cloud backup look like in practical terms, and how do you at CloudBerry educate MSPs on this?

DH: Cloud backup takes many forms – some providers have you backup your data to their cloud, done through an appliance at the customer's site which transfers to their data centre, which is quite common.

At CloudBerry, we've taken a different approach, because we understand that not one size fits all. From a cloud storage perspective, there are a lot of options, and they mean different price points and therefore different margins for MSPs.

What we allow MSPs to do is what we call 'bring your own cloud storage' and let them choose the storage that works best for them and their customers. That may mean they offer a range of different cloud storage providers to customers, depending on whether they're price-conscious or concerned about disaster recovery, and giving them a choice.

The other thing is, because you're bringing your own cloud, we only facilitate the backup from the computer or server to the cloud; we don't actually store the data – it goes directly from the client site into the cloud and it doesn't pass through our servers.

This gives an extra level of protection, because you as the service provider are setting all the security parameters for your customers so they can take control of protecting their data.

RT: This is something that really attracted me to CloudBerry when I first became aware of your solutions – you don't force any of the MSPs to use a set solution, a specific application or cloud service, because you cover lots of different ones.

That allows the MSPs to build up a 'best of breed service', so what does that look like in practical terms? What are some of the MSPs you see using as platforms and configurations?

DH: The most common platform is definitely Amazon S3, because it's the biggest and best,

but we also have a lot of customers who use Azure or Google. Some of the newer options are Backblaze B2 as well as Wasabi, and they've got some very aggressive pricing. That's the layout as it is today, and although we support over 30 different cloud options, those are the top five.

RT: I want to rewind a little bit and see where we've come from. Traditionally, we might have seen a lot of file-based backup within MSP businesses, and most of those I speak to today employ image-based backup. How do you see the role of cloud in that area?

DH: From an image perspective, if you're doing a backup to the cloud, many of the providers offer cloud compute alongside storage. Platforms such as Amazon have S3 for storage and EC2 for compute, which means you can take the image and spin it up as a virtual machine instance.

We currently provide that for Amazon and Azure and will provide it in the future for Google as well. For those cloud servers that provide compute, an image can now be stored up there and up and running in case of disaster.

That's what we see on the server side, and on the workstation side, where data is still being stored, we see a lot of file-based backup. When you're talking about client PCs and laptops, you typically have an image for those so you can rebuild them very quickly.

Doing a file backup to the cloud is much more efficient and uses less bandwidth than trying to upload the entire or parts of the image. That's typical of what our solution providers do –image backup for their servers and then storing them locally as well as in the cloud (we call this 'hybrid mode'). For the workstations and PCs, they do file backup which goes into the cloud.

RT: This comes back to what I would call the real-world situation, where one size doesn't fit all – clients have different budgets and needs, so I think this is a very cool approach. What is the most common scenario you see out there? Is it file-based for workstations and image-based for servers?

DH: That's probably the most common scenario, although even on servers they don't necessarily care about the OS (operating system), they just care about the data, so some customers use file-based for that and don't worry about the image.

I talked to a customer recently who are unique as they're not a managed service provider but an application provider, which they offer to a particular subset of customers. The application has data, so they provide as an additional service to back that up so they can recover it for the customer in the event of disaster.

For them, it's a very homogenous environment, because they're only backing up the data for the application rather than the entire PC, but they have hundreds of customers using this service. I think this is a unique service in how you can look at different ways of doing things depending on what business you're in.

RT: As the saying goes: “The backup is only as good as the data it recovers or restores” – maybe I made that up! When we talk about restoring files, I know as a former MSP myself it could be a real chore. How do you make life easier for MSPs when it comes to the restoration of files?

DH: We provide a SaaS (Software as a Service) based application, so if a customer contacts you to ask for a missing file, you can go right into the web interface and initiate the restore for that file, whether it’s from a file or an image backup.

Additionally, if you choose, you can enable a customer portal, so the customer can go in and recover their own files. Most of our MSPs don’t do that, as they’d rather provide the recovery service themselves and they don’t have to manage as many logins for the customers.

We try to make it as easy as possible to recover those individual file objects and get them to the customer when they need it and get them back online.

RT: Some customers will want the ability to restore the files themselves and others won’t, so we’re coming back to this idea that one size really doesn’t fit all.

DH: When I talk about CloudBerry I talk about flexibility, and the fact that we provide a platform that’s also easy to use. Sometimes that’s great, and sometimes people need to be told what to do, but I think the flexibility is key because every business is different and needs different offerings.

RT: I like to put myself in the shoes of an MSP business owner listening to this interview, and I know that people will be impressed by what you offer, but they’ll be wondering how to make sure that the end users and the internal MSP staff can only see the data they’re meant to see.

DH: From an MSP perspective, they can go in and access the data on an as-needed basis as they set it up, and you can assign accounts within the MSP. For the customer, if you choose to set them up individually they can log into a portal and only see the data that’s backed up using the credentials for that account. They’re walled off from seeing everyone else’s data.

RT: What about auditing? Quite often with backups it’s the weakest point, because people could backup an entire client site but the security wouldn’t be there. Do you have any audit trail for who’s accessing files?

DH: Everything is logged, so whether the MSP goes in through the SaaS or the customer goes through the portal, we can log what they do. Information is available via reports to show who’s had access, when and what they did.

RT: You’ve been in the industry for a long time. The MSPs you work with today, what size are they and what do they do?

DH: The majority are small businesses, with one to five members of staff and they cover a small geographic area. We do have some larger companies serving wider areas, but many of our current customers are SMEs with a small customer base of their own.

Almost all of them are servicing small business customers – the MSPs aren't supporting five clients with 25,000 users each; they're more likely to have 10 clients with around 15 users each. That's where our sweet spot has been, although there are some bigger MSPs, or companies who are using the product in unique ways.

RT: What about integration? The type of MSPs you've described tend to like the idea of 'best of breed' and get their own tools in. What sort of integration do you have with the common tools on the market, such as PSA (professional service automation) and RMM (remote monitoring and management)?

DH: For RMM, we have integration with ConnectWise Automate (aka LabTech), Kaseya and Ninja and these are available within the portal to connect them together. Interestingly, Ninja have done their own integration, so if you're a Ninja customer you can go in and enable CloudBerry to deploy backup.

RT: I'm hearing great things about Ninja. What's been your experience of working with them?

DH: They've been great to work with. The back and forth on the integration has been really good. It's only recently launched, but they're getting ready to roll it out in a big way. If you're a Ninja customer, ask them to tell you more about the CloudBerry integrations.

RT: And which other tools do you integrate with?
DH: For PSA, we have support for Autotask and ConnectWise. We haven't added a whole lot more, because as a lot of our customers are so small they don't use PSA and we haven't needed it. They'll use RMM before choosing a PSA, so we do what our customers want.

RT: Would I be right in saying that you've got API (application programming interface) access to hook into other tools?

DH: We do.

RT: That makes more sense than trying to hook into all the different tools, as the MSPs can use the API integration.

DH: There's an API integration that Ninja use to integrate with us, but also from the backup client perspective we have a very robust command line interface for that. Even if we're not integrated with your RMM, you can still use it to deploy the CloudBerry client using the command line toolset. You can even batch mode it and so on, so you can be creative in how you deploy that backup client to your customers.

RT: I'd be interested on your opinion on this – I see the role of backup as something that's changed significantly over the past few months, whereas previously long-term backup would go back six months or a year and protect data.

I don't think that's still the case, because with the rise of ransomware, you might be backing up data that's already been compromised but not yet activated. How are you and CloudBerry adapting to this new world of backup?

DH: Backup is one of the best protections against ransomware, provided you keep your backups air-gapped, which the cloud is great for. What we've done at CloudBerry is to add ransomware protection.

This is essentially a set of heuristics that view the data as it's being backed up, to look at what has changed from the last backup – the percentage of changes and what they were.

That way, if those files have been modified, we set up an alert to tell the client that they might have been hit by ransomware. This is in no way a replacement for your firewalls and other ransomware protection you should have in place, it's just another checkbox to find something that might modify your files on the next backup.

We want to alert the customer so they can check everything out to protect them from problems further down the road. It's just an extra checkbox we've put in, because as you said, a lot of times the ransomware could sit dormant for weeks and then be activated. The customer doesn't know when they were infected, which is a problem, because they also won't know which backups to recover from. It's not easy.

RT: Talking of end users, this is a question that came to me quite recently, from a member of my team who was confused by the idea of accessing backup through a web browser. I was intrigued by the CloudBerry Explorer app that you provide, which you mention on your website. Tell me a little bit more about that?

DH: That was actually one of the first things we released before backup. Essentially, CloudBerry Explorer allows you to connect to your cloud storage buckets and present it in a Windows Explorer-type interface. You're using Windows Explorer to access your files, drive letters and so on, and CloudBerry Explorer works in the same way.

You install it and connect it to your cloud storage and then you see it in a familiar file/folder interface where you can copy, paste, drag and drop the files. It's not meant for Dropbox and those types of things, it's meant for cloud storage bucket and blob access.

RT: And am I right in thinking there's a Freeware version of that?

DH: There is. If you still love the mapped drive, we also have CloudBerry Drive, which essentially maps a drive letter to the cloud storage bucket.

RT: That's cool, because there will be clients behind the scenes who don't care where their data is stored, they just care about doing things the way they always have in terms of drive letter, so I can see why that would be a great option.

When it comes to the cloud storage itself, most of the cloud providers I've spoken to or looked at suggest that their backups are absolutely secure, and I think for the most part they are. I know that CloudBerry Labs provide additional encryption options for some of the backups – why do you recommend that MSPs and clients use additional encryption in this way?

DH: The more you can protect the data, the better. The last thing you want is to be sending data to the cloud unencrypted, because someone can see the data in-stream. You need to encrypt it as you're sending it, because even if there's an https connection, that doesn't necessarily mean it's secure.

Another way to secure the data even more is through what we call file name encryption, so as we back the data up we encrypt the names. If someone gets access in one way, they won't be able to see file names, passwords or .txt that's been backed up from a customer's PC. Different cloud suppliers also offer in-cloud site encryption as an additional layer.

This means you're managing multiple, different passwords for all these encryption settings, but I think at the end of the day, that type of pain is definitely worth it when you're talking about protecting your customers' data.

RT: The more sophisticated MSPs already have systems in place for capturing passwords in PSA tools such as through using IT Glue.

I want to talk about another application of cloud storage, one that I don't think has been massively adopted by MSPs but which I think is a huge opportunity, and that is archiving to the cloud.

I have to admit that I don't fully understand the difference between what is referred to as hot, cool and cold tiers in the cloud. Help me understand what these mean, and why MSPs should be looking at long-term data storage.

DH: All these different tiers are also different service levels, which is why 'hot' is going to be best, not only for putting things in the cloud but also for retrieving it very quickly – it's hot and ready to go.

If you look at cool storage, that means you can put the data there quickly, but when you need to retrieve it it may take a little longer, because it's on a lower tier of storage (spinning discs and those types of things).

Then you've got the cold storage, such as Amazon Glacier, where you can put the data there as fast as you want, but to retrieve it may take up to 24 to 48 hours, because it's stored on some media that Amazon isn't forthcoming about.

Each one has a different price point, and costs vary not only for the storage but for the retrieval requests. A request from hot will be more than from cold, and so on and so forth.

The pricing helps to look at the archive capabilities when a customer says: “I want this data in hot for 30 days and then after that, move it off to archive.” We can automatically do that with Amazon today, where we can age the data into different tiers. We’ll have that capability for Azure at some point in the future.

The customer needs two weeks of fast retention and the rest of it they want to archive for a year, so you can age that data into the other tiers, store the archive and meet the regulatory requirements for those.

The other thing is also the regions, and that’s extremely important when we think about things like regulatory requirements, such as GDPR. Depending on where you’re based, there are different requirements for data storage.

If you’re in the United States, you might not want your data in a particular region or country, so you can choose this. Most of the cloud service has region support, so you can decide where the data stays, including across multiple data centres.

There’s a level of flexibility, and you get that with CloudBerry because we don’t own the cloud or have data centres in every country, but our suppliers probably do. That way you can choose what fits within your region, requirements and regulations.

RT: Going back to what we talking about around the hot, cool and cold tiers of storage, I remember in my MSP days I’d say to the client, ‘we’re going to back up your data, can you tell us what’s important?’ and they’d say ‘all of it’. When they got the bill, then they’d want to try to reduce the amount.

Most people listening to this will know it’s nigh on impossible to get the client to separate the data, but with hot, cool and cold you can move away the data that’s not been accessed recently. The client gets to keep it and it’s always there as a fail-safe, but it’s not costing them the earth to store it.

DH: Exactly.

RT: What’s next for CloudBerry Labs? What do you see changing in the world of backup over the next few months and years, and what are you doing to help MSPs with those changes?

DH: I think for us, we have an extremely good platform, so it’s about education. We need to get our name out there and let people who we are, what we do and how we can help them grow their business.

I think you're going to see a lot more push to the cloud, because from an economic perspective, it's not as expensive as people think to utilise it. When you look at the data that's put in cloud storage and the security, it's not dear.

In terms of CloudBerry, we continue to make improvements. We've got development staff working hard, not just on our portal for the SaaS-based application in the cloud, but also for all the capabilities we have within our backup.

We've got a lot of new features rolling out, including for Mac and Linux as well as Windows, and these should be available soon. Our aim is to keep pushing, to take our customer feedback and add that in to help us make better products.

RT: I've heard that you're really good at listening to feedback as a company, and I see the interactions you have with people on social media as well.

If anyone listening wants to reach out to you personally, Doug, or to find out more about CloudBerry, where can they find you online?

DH: I'm on social media as VMDoug – on Twitter and Instagram as @VMDoug. You can also visit www.cloudberrylab.com to find out more about us and visit our forum. I'm also active on Spiceworks and Reddit – search for VMDoug.