FEDERAL TECH PODCAST

Protecting Federal Data

SUMMARY KEYWORDS
veeam, backup, data, storage, aws, ransomware, federal, cloud, people, ransomware attack, years, prem, public sector, workloads, absolutely, aws s3, replication, key, snapshots, called

Welcome to the federal tech podcast. My name is John Gilroy and I'll be your moderator. Well, here we are. We're downtown Washington DC at the AWS summit, I registered and I just went down the lobby and grabbed the first guy saw her room and said, Hey, let's talk about this stuff. And, and lo and behold happened to be Jeff Reichart, Vice President public sector and compliance strategy for Veeam Government Solutions. Jeff, how are you?

01:09
I'm doing great. It's a pleasure to be here, John.

01:12
No secret we we discussed this before, but it's always fun to grab people off the streets and find out what's going on. I love talking to different people at AWS and different conferences and townhouse because they really put the rubber to the road. Exactly what can your product do for the federal government, especially what convene do for the federal government? And how does it work with AWS? I thought we bring in Jeff. He's got a great background. I've heard his videos on YouTube. He's having mellifluous voice. I probably get fired from this podcast app, hear His voice. And so we're gonna examine what Veeam can do for my federal audience. Jeff, can you give us a quick nutshell, your background? And maybe a little bit about VGS, please?

01:46
Sure, absolutely. So I have been at Veeam on the product strategy team for three years now. But I've spent the last 25 years helping public sector and enterprise customers build backup and disaster recovery solutions, I tend to focus now a lot on compliance because the sort of lines are getting blurred between backup and security these days, especially with ransomware. And I really care about the mission with public sector customers. So I spent a lot of time working with with beam Government Solutions and our federal team.

02:15
About six months ago, I had one of your colleagues, Gil Vega. And I've done hundreds of interviews and guild really stands out. He was very sincere. And he focused on, you know, we want to develop strategic relationships with people and we want to be next to them and solve some of these difficult problems. And if you saw Gil, you'd trust him because he's a big strong guy. And he really is sincere about solving problems, the federal government. So from the perspective of someone working at department, interior, or NHS or so so how can we help them? And what's it got to do with AWS?

02:46
Well, that's a great question. Veeam is a leader in the enterprise backup space. And no matter whether your data is still on premises in the data center, or you've moved workloads up into AWS or other cloud platforms, it's still important that you have backups of that data for a variety of reasons. So we have strategic conversations with federal agencies, both civilian and DOD, all the time, about what their data requirements are, how their data is growing, what their mission critical apps are, and where that stuff is living so that we can make sure we protect it.

03:19
By the way, I think one of the people in your team is this guy, Earl Matthews. I saw his background is like this guy's been everywhere, got the T shirts, been in battles. And now he now you brought him on to help lead this endeavor. I mean, he's really an amazing guy. Maybe I shouldn't shoot for him the next time you absolutely

03:36
sure that Earl is incredible. He is a Harvard trained lawyer. No, he's amazing and has a long and storied career with the US Army as a general counsel, and also in the national security community. He was he was a an advisor to the last White House in that area. And he cares passionately about the mission for both civilian and federal agencies. And he's a great leader at beam Government Solutions.

04:01
I love it because he's just a regular guy from Philadelphia, you know, kind of tough background and fought his way up and just, he's, he's a wonderful guy to him. Maybe next time, we'll do that. So we'll talk a little about backup here. My wife teaches Latin and because of that, I get to use fancy word. And one of the words I'm going to use is nomenklatura nomenclature. So let's talk about nomenclature here. You know, there's there's backups, there's replication, there's a storage snapshot, I mean, so So just for the benefit of the novices out there maybe can differentiate between these three concepts.

04:31
Yeah, it's a great question. And it's important to differentiate between those three concepts. The key thing to remember about backup is that it is a periodic capture or protection of data that is going to a different kind of storage medium. So back in the day, it was all backups to tape. Now, of course, many people are doing their backups, for example, to s3 storage in the cloud with AWS. That is a different proposition than a storage snapshot. A storage snapshot doesn't get the day To off to a different medium, it stays on the storage array or it stays if your data is up in the cloud. If you have cloud native apps, it stays up in the cloud. And it's important to be aware of that because the two different capabilities backup and storage snapshots serve different purposes. Storage snapshots lets you back things up with a lot less impact. They let you protect things more frequently. But it's still key to back your data up to a second storage medium to make sure that if something goes wrong with the primary storage, or heaven forbid, with your cloud platform, that you're still protected, replication is actually taking that up a notch with replication, what we're doing is on a frequent basis, replicating either storage snapshots, or software based snapshots for something like VMware environments, or Hyper V

environments, or with Veeam, we can replicate natively in our software, with a recovery point objective that is with an ability to lose as little as two or three seconds of data. So it really, if you've got your really mission critical apps, that tends to be where replication really comes to the fore.

06:08

Before the interview, we were talking college basketball, believe it or not, and 20 years ago, Jeff, you were probably in college and 20 years ago, I was reading writing a weekly column for The Washington Post. And I'll never forget one time I wrote in the post, I said, if you don't backup, you will crack up. kind of unfair statements, good for photographers, good for audio engineers, and it's good for federal agencies as well. But it's just okay, I'll make a backup. The issue here is that there's malicious software out there that will enter a system and do nothing for six months. And then lo and behold, I find there's a problem. And I say, Well, I'm John Gilroy. I'm real smart. I'll go to my backup from yesterday. It's still there. And so, backup and crackup is really the easy part. I mean, the hard part is managing and have some kind of strategy behind this attack, isn't it?

06:55

Yeah, absolutely. There was a Senate report that was released yesterday that estimated that between state and local, federal and educational institutions, there were 2323 ransomware attacks last year alone. And that's a minimum, because the Senate report pointed out that very frequently, people don't tell and this information isn't being captured. It's really important. Now, Gartner actually had a white paper about this back in September about how to protect your backup systems from ransomware. Because John, exactly as you point out, when these adversaries get on the network, their dwell time, which is the length of time that they can be on there poking around escalating privileges, for federal agencies, in particular can be months and months. And one of the key things that they want to do by escalating those privileges, is delete your backups so that if they attack, you don't have any option but to pay a ransom.

07:46

Wow. Now, I would think that there must be technology out there that can examine the backup and make sure that is not infected. Is that right? Yeah, that's

07:55

absolutely the case. So job number one is to use, I'll throw out a term here a nomenclature term, I'll throw the word immutability, which I think probably most folks listening here have heard with respect to backup. That's just a fancy word. That means once the data is written, it can't be changed by anyone, even an administrator or someone with full admin credentials. That's really key, because it's most likely that your ransomware adversary is going to get admin credentials and try to delete. But John, as you pointed out, the other thing you have to be able to do is because you don't want to go back six months or seven months, when you restore your data, you're going to restore your data from more recently, but you need an automated way to clean it and sanitize it. Because to your point that you just made it will have malware on it. So you need a way to stage that data, the recent data, but still get the malware off of it so that it's secure. that that process is something that we have automated for years with our data labs capability in Veeam.

08:55

Let's let's go deeper into vocabulary here. I love acronyms and weird. Washington DC. This is a ton of acronyms. They shove it in this as you walk in and absolutely. So there's an acronym I think I've heard used several times is called S O br It's kind of light hearted, sober scale out backup repository repository. We got the interest of the nerds depository must be important.

09:17

Yeah, that's absolutely the case. So when Veeam began supporting cloud based storage years ago, so for example, AWS, s3 storage or other cloud vendors like like Azure Blob Storage, and there's other solutions out there, we realized that we needed a way to virtualize the pool of backup storage that our customers were using. And that's where the scale out backup repository had its origin. Basically, what that lets you do is have multiple tiers of backup storage. Some might be on premises for your workloads that are still on prem. Some might be in the cloud natively as a first tier of backup storage if your workloads are in the cloud, but you can scale that that initial landing zone what we call the performance tier with something that we call the capacity CTR, which is basically infinitely scalable cloud storage. And the wonderful thing about that, and we're here at the AWS public sector summit in DC, the wonderful thing about that with AWS specifically, is that their AWS s3, storage also includes the ability to add immutability. They have a use case, I'll just dig into a little bit because it's not public sector, but it works out great. AWS realized years ago, they needed a way for regulated industries to do things like post financial returns, where regulators could get to them, and guarantee the regulator's that they wouldn't be changed once they were posted. So no Enron shenanigans or anything like that going on. Amazon came up with a solution called s3, object lock in compliance mode. And it locks those that data so that even if the CEO is standing over the system administrator and saying, We've got to cook the books upload a different financial return, they can't do it, it can't be changed. It turns out that is perfect for ransomware. Because if the ransomware adversary has admin credentials, and tries to delete the backups, he still can't because of that compliance mode. And for stats, Veeam is on track this year to push about an exabyte of data up into public cloud storage. So it's it's a use case that's really resonated.

11:16

I've listened to your voice and a radio guy and he says, guy, this guy modulates real well, he should be on NPR, he has no problem with this. Now, when you go on NPR, I was an undercover 25 years, the engineers behind this class and because you count down three, two, and you get to hit it and go live, and it's really important. Three to one applies to your world, too, doesn't it? So what is this three to one thing in Veeam?

11:38

Well, you mentioned digital photography a few minutes ago, the three to one rule is not unique to Veeam. It was actually invented many years ago by a digital photographer, who realized he needed a way not to lose his life's work, if something went wrong with his storage system at home that he was storing all the data on. The three to one rule just says if you want your data to be safe, you need three copies of it, that's two backup copies and your primary copy. Those copies should be on at least two different types of media. So again, back to our backup conversation a couple of minutes ago, and one of them should be off site. That's the three to one rule. One of the nice things, for example, about AWS s3 storage is that it solves those use cases, it's a

different media, from your production storage, it's off site. And it's an easy and scalable way for you to get three copies. So we work a lot, we integrate deeply with AWS and with other cloud providers on that area.

12:32

Let's go back 20 years and you were in high school this time. You know, I used to talk to people about backup, and they would literally pull up a truck. And then little John Gilroy, you get his two wheeler and wheel out some tapes, they take it to a mountain. Absolutely. That's great. I had this idea. Okay, what if there's a problem, you gotta back up? So you got to contract, you know, like Manny, you get a truck and drive the mountain bring attack. How long does this take? I mean, this 1015 years ago, this must have been terrible. And so what about recovery time?

13:05

Yeah, that's a great question. One of the advantages of the revolutions in Nearline data storage for backups that have taken place, both on premises and in the cloud, particularly around object storage, which scales really, really well for backups is that your recovery times can be much, much faster than picking up the phone, calling Iron Mountain waiting for the for the tapes to come back on site, etc. And that is key. There's actually a stat from Sophos they do a big anti anti virus ransomware report every year and the 2022 report just came out. Sophos said that last year in 2021, the average downtime from a ransomware attack was a month. Right? Think about that, because you have to do the remediation that we just talked about. But you don't want any inefficiencies in your process making that any longer than it has to be. So having really good performance, secure backup storage, like what Veeam can do with AWS and with other partners is really key.

14:04

And I have read of not federal agencies, but there's a hospital in the state of Washington that was attacked, and they didn't pay the ransom. It took them that long, they had to buy all new hardware. And this is not a pleasant experience all new hardware and and then what happens if you roll it in with a broken arm and you have to go the next hospital, it's 40 miles over something's, it's a real serious problem. So the time for recovery seems to be an important question here. I have done many interviews with military folks and people with three letter agencies. I can't even say what they were. And its hardness and harden and all this kind of stuff. So I read about some kind of a hardened Linux repository that you folks have. So is this inside Fort Knox or something or a brick wall or what is this?

14:47

So you're currently at version 11 A Veeam. backup and recovery when version 11 came out, which happened in the spring of last year or February I think of last year. The one of the big innovation In version 11 of Veen backup and replication was the hardened Linux repository. It lets you take commodity hardware with a lot of disk inside of it. If you think for the technical folks out there think something like an HP Apollo server or a Cisco, UCS, 3260, s 3260. Those are platforms that worked really well for this, you put a commodity operating system on it, and by hardening and according to VMs recommendations, that is locking down access to it. And also taking advantage of a file system x Fs file system and Linux, what you can do is have for very low cost on prem, a very high performing layer of immutable backup storage that is ransomware approved backup storage.

The way that we thought of this is that it basically democratizes the process of giving folks immutable storage because they don't have to go with a more expensive solution they can do is commodity hardware and storage shelf, exactly the cheapest you can get. Yeah, right. So you gotta

15:57

go back to my wife, the Latin teacher, she doesn't think I'm very smart, and she's probably correct. I put things in boxes, okay, here's a baseball guy. There's a basketball guy. It's a swimmer category. And that's all he knows about. And this is a this and this guy only drives a decent car. And, and and I try to put veem into a box, things backup company. Okay, next. Now why? Well, well, well, well. Someone like you and Gil will correct me gently, real gently and just kills case. But I think listeners have been does backup Oh, well, I mean, backup today is is it's not your father's backup. It's a holistic approach. I mean, if you look at the NIST domains, it seems to touch on not just one or two domains, that seems to hit them all, doesn't it?

16:37

It absolutely does. We love the CIS, the NIST cybersecurity framework at Veeam. Gil is a big believer in it. It's it's the framework that we use internally to structure our security practice. And you're exactly right, John, you might think that the five functions of the NIST cybersecurity framework which are identify what you've got protected, detect an attack and then respond to the attack and recover from it, that veem would play very naturally in the Protect and recover areas. And we do. But we also have a lot of capabilities that actually map to other areas on the wheel, we've got automated disaster recovery orchestration, to try to speed up that one month of downtime if you get hit with a ransomware attack. We also have some data reuse capabilities. I mentioned the word data labs a couple of minutes ago. Basically veem offers two ways. And this is probably the best case for we can talk about how backup has evolved in recent years. The game offers two ways that users can access their backup data to do things with it one way is API driven, they can just mount the backups as though they were a disk on a Windows or Linux server and then do things like data classification or security forensics. But the other way, which frankly, has been around for longer is data labs, what data labs lets you do is take a server that you've backed up either physical or virtual, spin it up in a VMware or Hyper V environment, isolated from the network and do stuff to it. So I'll give you an example. There was a very destructive ransomware attack that happened in December of last year, it was a Cronos HR company, probably a lot of y'all have heard of this, a lot of public sector organizations, since we're here at AWS public sector Summit, were affected by this New York Transit Authority. The city of Cleveland, a bunch of public sector organizations had their HR and Payroll up in the Chronos application, they got hit by ransomware. The vulnerability that was used had been patched two years ago, but they had not patched it in that time. And that was the initial point of entry that was used to do that attack. One of the capabilities with sure backup is it lets you spin up your whole application suite. So Active Directory, my database server, my front end or web servers, and I can patch them in isolation, and find out at 10am on a weekday, if the patch is going to work out fine. And then when my when maintenance window comes during the weekend, I have confidence that I can do this without breaking my whole environment. So it really lets people keep up with their patch testing a lot better, using just backup data and no impact on production. And that, in turn, makes their whole environment more secure. And the interesting thing about that is it has nothing to do with recovering from downtime. It's just using that backup data for something else.

19:17

A creative use of a sandbox isn't it's a different twist on little sandbox. Exactly. What about sandboxes for years, I came in here in the train. And because they have gray hair, I'm sure people looked at me and said, Hey, that guy's the legacies are they said he's a geezer. What's to say used? But our federal listeners have legacy systems. And just because the state of the art works today, and we talked about this v 11. A and everything else. What about legacy systems? So you can walk a half mile from here, I'm sure walk into an agency, and they may have some systems that are long in the tooth. Right? And so how does Veeam work with those?

19:58

Fortunately, because we've been to During this for 13 years, we can support lots of things that are looked at as legacy like tape libraries that still have their place, especially for large datasets or for organizations that like that. But the key thing that we have is a licensing model and an operating model that let us deploy on prem or in the cloud without having to do a forklift upgrade or junk, your existing investment. We actually had IDC, which is an analyst firm International Data Corp, they do the most comprehensive market share analysis for all kinds of areas for hardware for software for lots of different it areas. And they did a study for us, they reached out to VM customers and asked, what was the effect of going to Veeam. In your case, what they found was that generally folks who use Veeam, because we work really, really hard, we spend a lot of development effort to be simple to use, that people who go to Veeam typically spend 50%, less time administering their backup solution than what they did before. And they also, on average, have over 80%, less downtime caused by data loss, because it's so easy to get their data back. This can have really transformative effects. It's nice to go to shows like this, and meet Veeam customers. Because when I talk with folks, I hear again and again and again. Yeah, we got hit by ransomware. Three years ago, we restored they didn't get the backups. And we were back up and running over a weekend, you know, and that's that's the kind of story that you really like to hear.

21:23

If you jump in a taxi, about two, three miles from here is the Kennedy Center. And national simply plays there. And the technology people have liberated a word from the Kennedy Center, and they stole the word orchestrate. Absolutely. And they say, Yes, we're going to orchestrate this cloud transition, we're gonna orchestrate that. And I see the word orchestrate used with immutable backup now that's gonna impress someone, hey orchestrate immutable backup. And so I guess this means everyone's on the same page. Now, when I think about an orchestra, that piano and go nowhere. It's there, right? But in this world, it's dynamic data has got to be able to go over. So you're, you're orchestrating immutable backups for the federal government. This is a high level of complexity, this isn't this is probably harder than playing the cello.

22:10

Right? It can be it can be, that's one of the reasons that we spend so much time developing around simplicity, you've got a lot of moving parts, let's face it, if I've got storage in the cloud, if I've got workloads in the cloud, if I have on prem, I need someone to have done the work for me, so that when my team goes to do that, they have less things to worry about less things to think about, we actually have a product called veem, disaster recovery, orchestrator that automates the process of failing over in VMware environments specifically. And one of the nice things about the way that works is it actually gives you reports on what your service level

agreements are, that you're meeting for the workloads that you're backing up. And you can do this exact data labs process that I talked about a minute ago, you can simulate failover, and get reports that you can hand to your regulators that you can hand to your security team and say, Yes, we tested it, you know, twice this month. And these were the results of the test. That's, that's really key. One of the things that we added actually around AWS and Azure and other cloud providers specifically, was the ability to orchestrate different tiers of backup storage, so that if you want to save certain data for a long time, certain records have a long retention period, instead of keeping everything in AWS s3, you can archive it off to Glacier or Azure archive, you can take advantage of that you can take advantage of glacier deep archive, and really just for pennies on the dollar, keep the data that you need to retain for a long time in a very cost effective way.

23:37
Gonna give you give you a kind of forward looking question here. And you can dodge it or answer it. I always wonder who's who is leading backup technology here? Is it is it banks and insurance companies, Jules Gill vague has experienced there is it intelligence community is it is at NIH. So who is reading and you had your feet in captivity? So who's leading who here?

24:01
That's That's a fantastic question. In my experience, regulated industries typically do as good a job or in some cases better, especially financial services, because they typically have a lot to spend on solutions and can afford to get replication and storage snapshots and backup on the platforms they're working on. But I've seen that financial services and pharma in particular, because they're heavily regulated industries tend to do a really, really good job with this. But on the federal side, some key things like multi factor authentication have been absolutely standard in federal organizations for a long time. I mean, you can't go and haven't been able to go for decades now onto military bases without everyone having using a CAC card to log into the system. So multi factor authentication, federal government is way ahead. And what I'm finding is that VMs federal customers and we have about 1400 of them completely get this conversation when we talk about the need for immutability when we talk about cloud and they are adopting more at a much more rapid pace than in the past, we actually do a data protection Trends report every year at Veen, we actually do the largest survey that's done by anyone, including analyst firms. And what we found in terms of the environments that have been deployed, are that overall, we surveyed this year about 3400. Folks, a globally all all verticals, all sectors, mostly large organizations. What we found is about half of workloads are already in the cloud 49% On average, and it's divided about 24 and 26%, between physical servers and virtualized servers on prem. What's interesting is that public sector Mac's maps to that almost exactly, there's very little difference. I think, in the past, we had this idea that maybe state and local or federal were slow adopters of cloud technologies. That's absolutely not the case anymore. If it was in the past, it's not any more because our data show that they have about half of their workloads up in the cloud, just like banks, and just like every other industry, and the other half divided between physical and virtual on prem.

26:05
Well, this has been a great conversation, Jeff, unfortunately, we're running out of time here. I'd like to thank our guest, Jeff Reichart, Vice President public sector and compliance strategy for Veeam Government Solutions. Thanks, Jeff.

26:17
Thank you, John.

26:20
Thanks for listening to the federal tech podcast. If you liked the federal tech podcast, please support us by giving us a rating and review on Apple podcasts.