EPISODE 1412

[INTRODUCTION]

[00:00:00] JM: Rackspace is a multi-cloud solutions provider that has evolved beyond its cloud computing origins into a diverse set of services and support offerings. Customers work with Rackspace to adopt cloud application deployments, modern data analytics, and all the other opportunities offered by cloud computing. Much of this occurs through partnerships, where Rackspace provides teams of engineers to work with the customer.

Jeff DeVerter is the CTO of Products and Services at Rackspace. He's also the host of the podcast Cloud Talk, a show about trends in cloud and technologies. Jeff previously served as the CTO of Microsoft Technology at Rackspace. So our conversation began with a discussion of Microsoft-based migrations.

If you are interested in sponsoring Software Engineering Daily, we reach over 25,000 engineers per episode, and 250,000 engineers per month. If you are promoting a product, or looking to hire engineers, Software Engineering Daily can be a great place to get your message out to a large populace of experienced engineers. To find more information about sponsoring Software Engineering Daily, go to softwareengineering daily.com/sponsor. We'd love to hear from you.

[INTERVIEW]

[00:01:21] JM: Jeff, welcome to the show.

[00:01:22] JD: Hey, thanks so much. It's a pleasure to be here.

[00:01:24] JM: I want to start off with a fairly specific question. Let's say I have a large on-prem installation, and most of it is Microsoft infrastructure. What are the steps to moving that to the cloud?

[00:01:38] JD: Ooh! That's sounds like a call that I have five times a day. So it's a great question. So first of all, you need to destination. You mentioned it's Microsoft infrastructure.

Sure, makes us think Azure, but Azure isn't exclusive in this context. And so you need to know what your destination is. But I'd say even more than that. What are the goals of moving to the cloud? So helping to understand, is it as a capabilities play? Is it an availability play? Is it more tacking capabilities that exist in your cloud of choice that you really want to move to?

Once you sort of get that stuff baked in, it helps to direct to the right cloud, it helps us to think about what the right infrastructure might be once you get that way. But the other thing kind of before you jump into that migration or the tech bits as well is a realization that that move to the cloud isn't like previous migrations, where you might go from one data center to another. I mean, you're truly moving in to an environment that will evolve underneath your feet without you doing anything. And the capabilities there are really just beyond, I'd say, behind your imagination of what you can actually consume if your mindset has been totally, "Hey, I'm a Hyper VMware person running my Windows applications."

Once you get through all of that and you kind of get this mindset of constant innovation and constant upgrade, that whole ethos, is you got to know what you have. And that starts with an inventory. If we're thinking about a migration path that looks like, some might say, lift and shift, I prefer to say move and improve, where you're going to use – You want to get your feet wet in the cloud and not change too much at once. So if we're going to spin up some VMs, ultimately, to run a majority. Not necessarily all, but a majority of those services. So what are the networking requirements? What are the segmentation that you've got to deal with? How are we going to replicate that networking environment? Second to that, then, is what's the infrastructure look like that we're replicating? And generally speaking, you go one for one as you move out. You find that there's a lot of optimization opportunity once you're there. But do that once you're there through lessons learned and how much downtime you've got in machines or unused compute cycles.

And once you've got a full picture of what you need created, go in, start creating that. Of course, you have to couple of methods there to do it. You could either go in and through the interface and create that stuff. You could create a bunch of scripts that would automate that for moving as a Windows workload. We're going to move it out to – Say, we're going to Azure, because it's Microsofty, and you got Azure DevOps that you could use. But all of the usual, Ansible, Chef

and such can be used to deploy there as well. And then once the infrastructure there, it's what's your data migration plan, get the data in and recreate your applications, and of you go.

[00:04:27] JM: And is that significantly harder than moving a Linux-based installation?

[00:04:34] JD: Is it harder? No. All the same requirements still exist. The only thing that you might think about would be, in that move, does the fact that they are Windows workloads make you think that your only place – Just give you an artificial constraint of the fact that you need to go to Azure? And the answer is no, because Google and AWS will give you a first-class Microsoft server experience. But the challenge really becomes what does that – And I mentioned it earlier. What's your long-term roadmap for evolution? What are the services that you want to use to bring around your application as you go forward? Is your goal to get rid of running Active Directory servers when now you are saying you want to go and use Azure AD, so you'd move into Azure? So it's not necessarily harder for Windows or Linux. The two cautionary tales is don't think that just because it's Azure, or because it's Microsoft, it has to go to Azure. Second to that would be make sure you're thinking of the long-term roadmap and not just what's my first step into cloud?

[00:05:30] JM: Why is that long-term planning so important? Like what are some snafus you're going to get into if you don't prepare appropriately,

[00:05:37] JD: Well, ultimately could be signing yourself up for another migration. If your intent is, “Hey, I've got my Microsoft application. But my goal is to take in and ultimately mine that data. That's one of the reasons we want to move to the cloud.” We've haven’t even talked about why this customer is moving. But it's why they maybe they want to go. And maybe they have a BigQuery experience in their future. Well, if they make their artificial boundaries that says, “Well, because it's Microsoft, I've got to go out and move it to Azure. And now you got a second migration to go over to Google.” Or you are going to run two clouds at the same time, certainly doable. But just make sure you bear in mind the fact that you're signing up for a lot of fees on the networking side as the data leaves Azure and goes over to Google. It's generally free to get your data in, but it is not free on the egress side.

[00:06:30] JM: When people have these large installations from long-lived deployments, are there places where the pain shows up the most? Like if they're trying to move to the cloud, is there – Or is it mostly more of like institutional buy-in? Is that where the most friction lies?

[00:06:52] JD: You're challenged with long-term – I call them monolithic type applications. It's never just the one application. If it's been around the enterprise for any significant period of time, there are any number of other applications or services that rely on the data that exists inside of that one, or that this application relies on to function well.

And so consequently, I'll call it older application, you've got to look at all of the tangential applications that you've got to deal with as well, because you're going to break stuff when you move it. That becomes incredibly true when companies start to think about moving Oracle installations, SAP installations. Almost every other business in the enterprise with their own custom apps relies on data that comes out of those applications. So you got to think about the impact of moving your app when it's an older application.

The other piece to that is that if the app has been around for a long period of time, then the infrastructure that it's writing on, there's usually some other servers or services running on those VMs as well that the company has just forgotten about. Again, you mentioned institutional knowledge that the companies lose over time. That's usually what happens there, is they'll go and move it. And then a couple of businesses raise their hand and says, "Where did my app go? Where did my service go?"

[00:08:12] JM: How do companies prioritize this kind of migration among all the things that they could be doing? This is like a classic opportunity cost kind of situation where everybody always has new features that they need. And that's going to contrast with the effort to make some kind of large migration like this.

[00:08:34] JD: That's an interesting one. And it all goes back to how has leadership been sold on the need for this migration? For this move? Is it a cost savings move? Is it a feature-based move? And so I caution anyone to go to the cloud, because they're going to save money, because you aren't. You're not going to save money, because if you're picking up and moving a larger institutional enterprise application and moving it to the cloud, you're going to do a lift and

shift one for one, those machines cost more in the cloud on constant use basis than they do just going to buy-in your own compute.

Also, to that, is once a customer then starts to modernize and starts to save money and optimize it, then they're going to start to spend intentionally well. So it all goes back to how has leadership been sold on this move and how's it going to be graded, successful or not?

[00:09:27] JM: Can you speak more about that? Like do you have maybe a case study in mind?

[00:09:31] JD: Yeah, let me see if I can think of a customer I could tell you about. Oh, I got one. A Good one. So there is – I can mention them and by name. They're a very, very large, agri-food business and Midwestern sort of a company. They moved out into Azure a couple of years ago, and they were moving predominantly their development environments out there. And they took a model and a mindset that said, "Hey, I think we can save some money if we make this move. And also, if we do that, "Hey, it's a cloud. We get some more self-service capabilities. We don't always have to wait on the IT teams, the infrastructure teams."

And so off they went to the cloud, and all of a sudden, their bill starts coming in. And their bills are huge, because these developers are setting up very large environments. They're setting up – They're turning them on and they're leaving them on. It was very simple solution to solve. And so we came around and wrote an application, basically created some personas in their Active Directory. That said, based on the type of developer and the teams that they're on, they're allowed to do. And they work on these types of environments. We figured out what those environments look like. And created, through Azure DevOps, the ability to create those dynamically when they need them. Also, the ability outside of business hours that they automatically turned off and then turned on before they would show up for work.

Their Azure bill was cut literally in half the first month that went into production, and it was a significant bill. But then as we watched their spend over time, it started to go up. Well, it started to go up because they they've met their goals. They met their goals of availability. They met their goals of self service, and they met their goals of saving money. And then the leadership, the IT leadership, came back and said, "All right, you've proven that this works. Now let's expand this program out with keeping these boundaries in place, these guardrails in place of

who can do what, as well as when these machines would function and operate, these environments." Turning off on weekends. Turning off overnight.

And now they purposefully, with understanding of those costs, they were able to then expand that out to larger parts of the company. In the end, their bill was higher than it was when we first went in to fix it. But they were using so much more, and they were being so much more productive.

[00:11:39] JM: You've been in Rackspace for a while. And I'm sure you've seen firsthand perception of the cloud changing. What have been the biggest industry-wide perceptions that have changed?

[00:11:53] JD: It's a great question. So I think, first, I'd lead with the fact that the cloud isn't there to save you money. It's not going to save you money. I've said it a few times now. You're either going to spend poorly. You're going to fix that. You're going to get efficient with it. And you're going to – In fact, that's a big piece that we talk about a lot. And that's that FinOps, or financial operations for cloud-based environments, gets viewed a lot of times through the lens of it's there to save you money. It's not there to save you money. It's there to help make sure that every dollar that you spend inside of the cloud gets spent well, gets spent with purpose, gets spent to bring a very specific outcome. When you can do that, that's not a problem to have to spend the money to get it.

Other misconceptions that we saw, really very much in the beginning, is that we could take and treat cloud operations like we would treat regular data center operations. And everyone did it wrong from day one. They treated the cloud as if it were just more managed hosting or more servers running inside of a data center. And those workloads, first of all, they're running in an environment, as I mentioned before, that is constantly evolving underneath their feet. And so the role of operations goes away from, "How can I help you patch and backup and recover volumes or servers?" To, "How can I partner with the development teams to be able to help make sure that the infrastructure grows and modernizes through automation, through new capabilities as that cloud evolves underneath the feet? That they're moving in lockstep?"

[00:13:30] JM: And how has Rackspace's strategy changed among that change of perception?

[00:13:38] JD: Sure. So when I mentioned that everybody did cloud operations, in a sense, wrong from day one, we helped lead that charge. Meaning, we did it wrong. We treated it as regular managed hosting. Here's these cloud environments. I'm going to help make sure that they're patched and backed up. It's a cloud-based server. I shouldn't be patching and backing them up.

What I should be doing is working with the development teams to make sure that we have a dynamic ability to provision, or reprovision, or decommission based on demand, based on issues, based on need, as well as helping to make sure that, from a cloud operations model, that we're building machine learning models and using machine learning server apps to be able to make sure that we are dealing with those repeatable easy issues through software and letting humans do the hard work.

So we've gone away from models that are strict, in a sense, do you just pay us a percentage of infrastructure? And what we'll support for you is inside of this little box over here. And if it's not that, then we don't necessarily do it." Now we've moved to a do anything and a do with model that is designed to ensure that the infrastructure for the application is evolving with the application. And it's a model that the rest of the industry is copying, but we're – Flattery and all of that, which is nice. But that's been a big change that we have done around here from treating the cloud has managed hosting to treating the cloud as a dynamic environment that needs a dynamic support model that matches it.

[00:15:08] JM: So in that move from the focus on managed hosting to support, how is Rackspace had to change organizationally?

[00:15:19] JD: So it's a really cool questions. So when you think about our AI ops offerings, right now, Rackspace has one of the highest automation, service automation capabilities in the industry. We solve over 80 – We'll just say 80% of all system or user-generated tickets are solved systematically through software.

Well, we used to have Rackers who did all that work. They tended to be level one, level two type Rackers. What do we do with them? Well, what we do is we train them. We have a process and

a program in place that takes that takes traditional server engineers and turns them into cloud engineers through a 90-day program, paid program of 90 days of instruction. We run them through. And at the end, they literally get job offers into new teams in the cloud operations teams, as opposed to how to – Again, they're not patching servers. They're creating automation to help evolve the infrastructure over time.

[00:16:16] JM: And your role personally, like what have you had to change in your career as the company has made that shift?

[00:16:23] JD: Well, I've sort of a different sort of a view. Not necessarily view. But my career path has been a little different. So I started in 2008 at Rackspace. And I came over from a large financial services firm here in San Antonio. And I ran SharePoint, in a sense, for that large enterprise, 30,000 employees, utilizing that infrastructure to come help build out a SharePoint we call the dedicated single tenant SharePoint offering at Rackspace.

So my trajectory is a little different. I helped build out businesses inside of Rackspace, build them up and then move them on to other areas. Over time, I became the CTO of all the Microsofty things. I'm very much a Microsoft heritage, as you can imagine, starting in SharePoint. Actually left the company for about a year and 10 months and then came back as a CTO inside of the products division. And now Chief Technology Evangelist.

And I'll tell you, a big part of what I do is I spent us massive amount of my day helping the world understand that Rackspace technology is more than just servers in Iraq. It's about helping customers go through that whole transformative experience of moving to the cloud. Because again, the cloud isn't a geography change. It isn't just a data center change. It's an ethos change of a way of doing IT. 100% different.

So at Rackspace, now we have a full professional services that includes advisory services. Helping customers create that vision for the future, professional services, to go through that transformative change, and then the revamped managed services, or we call it Rackspace elastic engineering, to do the care, and feeding, and growth of that environment over time.

[00:17:55] JM: How does Rackspace interface with the variety of cloud providers that are out there? And how do you scale support, to being able to support the multifarious use cases that you can have with all these different cloud providers?

[00:18:11] JD: Yeah. So we have a pretty rich partner program here, and the hyperscale cloud providers, AWS, Azure and Google at the top of that food chain. Of course, our friends at VMware there as well. But the question is more around the cloud provider. So we manage them in a couple of different ways. We manage them as you would manage any partnership relationship. We have people who are assigned to that, and multiple people who are assigned to that. It's a very complex working business relationship. But the other piece of that then is the technical side.

So it starts with how do we make sure that everybody stays trained up? How do we take these service models and offerings that we have in one cloud and make them available as ubiquitously as possible? Not just across the other clouds, but across the other geographies. Rackspace is a global company with 7000 employees in 120 different countries and broken up into three primary regions of America's EMEA and then APJ with divisional leaders in each of those areas.

And so, for us, it's about ensuring that we've got everybody as trained as possible. I already mentioned the technical onboarding program, or TOPS we call it. But then there's also internal we call Rackspace University to help make sure that everybody is skilled and certified in the clouds that they have. And then consistency across delivery is just a management effort as well. So that's few of the ways of how we approach it. But it is a complex relationship, to say the least. And we love our hyperscale cloud providers. And we have regular conversations at the tops of each of those three companies. We drive a significant amount of revenue to each of them through the customers that we support inside of their clouds. So they have a vested interest in our happy partnerships.

[00:19:56] JM: These kinds of decisions around moving to the cloud and doing ongoing management of cloud resources and like, I guess, cost management, is this typically managed by C suite individuals? Or does it get pushed down? Does the decision making get pushed down to mid-level management or the engineering teams?

[00:20:23] JD: Well, you're certainly going to have your CTO, CIO roles who are going to be involved in it. They will be in communications and have conversations at that senior C suite level. They need to make sure that security is happy, that legal is happy. But there's also some – You mentioned reorganization inside of Rackspace. But any company that's going to go out and adopt cloud-based computing is adopting, again, that new ethos of way of working. I mean, if we're just looking at it for VMs, okay, fine. You probably don't to change very much.

But if we're moving into a DevOps, DevSecOps type of a model, we're integrating the entire company into that software development process. We are helping them to understand that agility comes through the entire company working together. So you have to have alignment at the higher levels. Now, were they involved in a day to day basis? No, not so much. Specific technologies that get chosen to be used inside for cloud of choice, whichever that one is.

Now we're starting to get down into mid-level leadership that is definitely taking their cues from those software developers, the operations team, of what is going to be best to integrate with their organization. And that integration needs to be thought of in a couple of different ways. Of course, we think of it initially as technical, but fits technically into our suite. Cultural, because culture fits into this as well. But also, on the business side. Which company fits from a business-wise into our organization well? All those things certainly play into it.

[00:21:48] JM: Have you seen any particular ways in which an ongoing management relationship of infrastructure between like outsource consultancy, or outsource consultancy, but I guess, like the partner technologists and the actual like in-house engineers at a company. How is that relationship best managed? Like do you typically have the partner developers, like the developers at Rackspace, for example, that are working with a company that's lifting and shifting? Do you just have them like acting essentially as employees of the company? Or is there a division, is there like a stark division of labor between the partners that are not endemic to the company?

[00:22:39] JD: So Jeff, our goal is to always be as employee-like to our customers as possible. That creates camaraderie. It creates closeness. We understand the business challenges better. We understand the business opportunities better. We understand the business outcomes better

in that type of a scenario, then we have the opportunity to really create some meaningful outcomes. And people enjoy their work better, which is important to think about as well.

Now, that is very much going to be driven, though, by the temperature that's being set by our customers. And I would encourage any company that's utilizing contractor-based workloads is to – They may be paid in a different way, but they need to be viewed as part of the team. It's not throw work over the wall. It's everybody creating solutions to a common outcome. Because we do have some customers, be it regulatory challenges, be it their own internal the way they do business, their business structure, that we just don't get as plugged in as we would prefer to. But when we do, when we are considered and brought in to be truly part of the team, then the right outcomes occur.

[00:23:51] JM: How has Rackspace dealt with the move to remote work during the pandemic? Has there been any significant issues since you guys actually have to manage a lot of physical infrastructure?

[00:24:04] JD: Great question. So when we think about the infrastructure that we manage in our own data centers, and we have 40 of them around the world, there's a lot of automation that plays into a lot of that. So the amount of staff required in data centers is not the most significant personnel staff that we have. Most of – When you think about supporting those types of workloads also, the majority of the teams that support workloads that run on that infrastructure, that can happen and always has happened remote even when they were in a Rackspace corporate facility. Rackspace corporate facilities that hold people don't hold servers. Data centers are always separate from those environments.

And so our move to working remote was really quite seamless. In fact, by our standards in the first six months of the move, productivity was equal to, if not a little higher. Now we settled back into the pace where we normally would be. But everybody was superconscious of, “Hey, we realize we're working from home. And we want to make sure that we're doing a good job.” There's a high work ethic that runs around Rackspace.

But for the rest of us, it was a pretty seamless move. We had started testing those systems that we had in place in late November, early December. So by the time January came around, we

were already running drills of 50%, 75%, in some cases, 100% of the employees working from home. So that when March did come around, and like everyone else, we pulled the plug and all went home. It was pretty seamless.

[00:25:35] JM: Can you tell me more about managing large scale server infrastructure? How do you, for example, do security monitoring and manage points of presence that are widely distributed enough? Give me an overview of how you successfully manage large fleets of physical infrastructure.

[00:25:59] JD: Sure. It's been the business. Was the start of the business back 23 years ago. So a collection of – In some cases, there's a little bit of homegrown applications running around. But there's also a lot of off-the-shelf stuff that runs as well. So Rackspace manages our own network operations center that looks after not only our own physical infrastructure to support customers, but also the customers infrastructure themselves. We've got two of these. One in the US, one in the UK. And that looks after the physical infrastructure from a networking point of view, also, from a physical point of view.

Then, depending on the size of the customer, if we think about – Of course, if we're talking about large deployments for individual customers, these are going to be larger ones, enterprise segment. And they've got suites of tools and services in place that look after all of the monitoring, that look after our asset management, that look after a patch status. All of that sort of stuff is managed inside of there.

And so in some cases, there're specific tooling for specific customers that's – But in other cases, it's just a repurposing of corporate tooling as well. So it's a very elaborate corporate set of infrastructure, of human capital, of software and physical infrastructure to help make sure that we're able to do that at enterprise-grade for some of our largest customers to our smallest customers.

[00:27:21] JM: What does a software update look like when you're trying to roll it out to multiple data centers?

[00:27:30] JD: Yeah. So we've got tooling to help make that happen. We have predefined patch days for – And that's really would be the software updates that we would do, would be patching to infrastructure, patching to servers. And we, as you can imagine, have tools to handle that. There are predefined days when customers do get upgraded on times of the month, as well as their own – Gives them the flexibility to choose different days or skip them altogether for that matter. So we handle all that centrally. There're remediation plans. There's notification that goes through. And every customer is assigned into an account management base that includes customer success managers, and engineers, and architects, and networking folks.

And the smaller customers, you're grouping more of those together in greater density with teams than the larger ones, which tend to be one to one. But everybody's notified when their customers are going to be upgraded. And everything gets kind of managed that way.

[00:28:29] JM: And can you tell me more about data center ops? Like what's the role of a data center operations person? And how has that changed over time?

[00:28:40] JD: I'd start with the way it's changed over time. Automation has really played a big key. If we think about somebody who's doing datacenter operations, of course, there's a leadership hierarchy that exists inside of there. But somebody who is, in a sense, I'll call them racking and stacking servers. They're just taking the architecture design that our solution engineers are creating. Once the gear arrives, or they pull it out of inventory, they are literally racking and stacking that gear. Cabling it in a way that meets the design spec. They are ensuring validation testing on all of that and setting it up for the automated tooling to then deploy operating systems in and begin the configuration process.

But as you can imagine, as more and more of that becomes more and more automated, the need for people physically in those data centers goes down. But we're not at the point of robotics where the robots are unboxing and building servers and putting them in racks. So that's still a human's job.

[00:29:34] JM: I want to come back to the discussion of moving to the cloud. If I move my large, Microsoft-based deployment to the cloud, what's next after that? Like once I get it into the cloud, what am I looking to do to leverage the cloud? Am I like changing my database deployments?

Am I doing things to slim down my application? Like containerizing? What's a typical usage pattern?

[00:30:06] JD: Well, in a large part, it's going to go back to the roadmap that we helped create for the customer. And so depending on what the needs are of the individual applications, and this is really sort of one of the first times we've talked about this on this call. And that is, while we can look at what infrastructure globally we want for our organizations, the path of architecture is going to be different and should be different for each application.

For instance, if we've got – This long-term application that's been running inside of the organization, we know it's only going to live for another couple of years. You know what? It doesn't make sense to go to great lengths to modernize this application – Even to put it into containers in some cases, it just doesn't make sense. So what makes sense for that app? Well, maybe there's a common identity infrastructure we can pull from. Maybe we can get its database to use a PaaS-based database and get some savings there. But at the end of the day, some reservations for the VMs themselves to get them at the best price and let it live out its life while we migrate the data elsewhere as part of an enterprise architecture strategy.

But for other applications that are going to be around for a while, well, it's important to think about what does success look like for them. And what is the technical aptitude of the firm themselves? I think Gartner's got the number that says that while most organizations want to adopt a cloud native type of an approach to cloud, 90% to 95% of them can't do it simply because they can't find the right people or their people are so consumed with keeping the lights on that they can't retrain them to go on that journey.

So understanding what the right path is into a technical environment that can be supported globally, can be coded to globally for the organization is super important. And then having a desired state, architectural state for enterprise applications, is important. That way, everybody has a north star that you're driving towards to usually manifest itself as three or four primary patterns inside of the organization. And then down the marching road you go. A lot of times containers is that next step. Once you move the environment over, you move into a containerized environment. Once you've containerized the environment, you've broken this

application down into – I wouldn't call them bite-sized chunks yet, but at least they're smaller than the monolithic app that it probably came from.

And so once it's in that environment, now we can start to strategically look at the application itself and look for areas where a serverless bit of infrastructure maybe can wrap around those containers and move us to more of a transactional model on our cloud consumption as well and save, in some cases, save some money, in some case, build resiliency, in some cases do both.

[00:32:58] JM: Have you seen there the companies that move to the cloud? Do they typically maintain some hybrid infrastructure? Do they maintain some servers actually running on-prem?

[00:33:11] JD: In my experience, most companies leave something behind. There is a mainframe somewhere. There is a Windows 2008 server that they just can't bring – They've lost the source code to and they just need to let it die. Usually, something stays behind in most occurrences. Now, bold companies will take it all and move it out. They will commit all in. And those that do it successfully, when they're doing it all in, we'll take a measured approach to it.

We all love some cloud native. We know the value that exists there. But making a transition from a VM-based environment holistically and completely out into cloud native in one step injects a ton of risk into the environment. So I talk about a roadmap. What is the roadmap toward a desired state? We'll work with even financial services firms that will spend a year leaving all of their data centers in just doing a lift and shift, doing it cautiously, doing it in a measured fashion and managing and maintaining it.

Once it's there, these individual applications now can undergo their own lifecycle of transformation. Because at that point, IT has sort of done its job. Because, Jeff, we haven't really talked about some of the other dynamics in making these decisions, because IT could come along and say, "Hey, we're leaving the data center. We're heading to Azure. We're heading to AWS. Buckle up. We'll tell you when your migration day is and when your outage window is." And we're just lifting and shifting.

Once we're there, it's up to you to upgrade your individual applications, utilizing the suite that exists there. So you can sometimes hit some speed bumps as far as adoption is concerned.

They've moved to Azure. They've moved to AWS or Google. And now goes and is incumbent on the individual lines of businesses inside of that organization to either do the work or pay for consultant to do that work.

So you could see some slowdowns there. But if the company is moving in lockstep, you asked the question before. How much is C suite involved in this? And when C suite is more involved, then those individual lines of businesses is working well with IT and is well-informed in the opportunities that exist now for their applications.

[00:35:18] JM: When you're engaging with a company, is it typically like – Do you have like one top level engagement with them? And then the engineering teams within Rackspace and the company engage beyond that? Or do you find yourself involved on an ongoing basis in these transformations?

[00:35:42] JD: Well, when we think about transformation, think about advisory services. What's our roadmap? I think about professional services. What's that big transformation work? And I think about managed services. How do we care for feed and grow that environment? Mature that environment over time?

So for companies like Rackspace, it's up to us to make sure that that advisory and professional services engagement goes amazingly so that they continue to engage with us in that care and feeding and maturing of the application inside of the cloud. Now, that being said – And I talk about companies that go whole hog out into the cloud. Not many companies go that route. Usually they're moving in a much more measured fashion. They're going to say, “Okay, we're taking division X. We're taking division Y. We're going after East data center. We're going after the UK Data Center.” And they'll move that bit over.

So for Rackspace, it's important for us to make sure that we are executing well regardless of the service contract. And usually, they show up in smaller segments. And in fact, we sort of encourage that as well, because it's easier to gauge success. And it's easier to adjust when conditions change, or the environment changes, or the business dynamic change. When you're in some monolithic long-term contract that says move it off from here to there and be done in

nine months, whatever the term is. Those are challenging, to say the least. Challenging for both Rackspace as well as the customer. It's just a lot to bite off.

[00:37:14] JM: What's the hardest part of moving an application? Is it the database itself? Or the actual application runtime? What do you find to be the most difficult part of moving?

[00:37:30] JD: I'll put my consultant head on and answer like a consultant. It depends. It depends on a lot of different factors. What's the size of the dataset? How much code is inside of the data? How large is it? Do we still have source code for this environment? Are we just lift and shifting to there? Which really transforming it. All of those factors play into it. But I'll go back to the hardest part of cloud transformation, is never the tech. Never the tech. Okay. So it means 99% of the time, it's not the tech. Sometimes the tech is hard. Mainframes are hard. Mainframes are hard for a lot of different reasons. But if we sort of set that to the side, the hardest part is organizational change. The hardest part is changing the way an organization operates and functions. Because to truly turn your application to be cloud native, your business needs to be cloud native. You need to as individuals adopt that new agile-based DevOps function, DevSec-minded world that says, "This is our new reality. This is our new ethos." And with that comes change. And getting companies to adopt that change successfully over time can be challenging if they don't realize it's more than just another migration.

[00:38:42] JM: That term cloud native, I typically think of that as associated with containers and Kubernetes. You've been industry long enough to also have witnessed the shift to virtualization. How have you seen the differences between the movement to virtualization versus the movement to containerization?

[00:39:02] JD: Virtualization was a lot easier. Virtualization was easier, because you could just say, "It's just another server." Remember all the stuff you had to do on your physical server? Now you do it on your virtual server. By the way, you can get greater density, and you'll get greater utilization. Oh, that's cool. I understand that. But when you start thinking about containerization, now that means we have to do – We're just moving runtimes or just moving data in some cases.

And so that's been a little harder for non-technical people on a day to day basis to wrap their brains around. And that creates inertia. Not understanding the change that's about to occur, always creates inertia. It's the thing that slows most companies down when they think about moving out to the cloud. They don't understand what that move is going to look like. Even moving into VMs? What is a GCP data center look like? Where's the address? Can I go visit it? Well, you can't, and you can't, and you can't.

But when companies work with a consultant whose primary business as helping them through that transformation and their services are geared towards different phases of the transformation, what starts to remove the unknown? And that removes friction. That removes inertia from projects.

And in those contexts, and this applies to containerization as well, starting small so that they can understand what that move looks like. Starting small helps them also understand what new infrastructure do I need to manage this stuff? Well, I need an orchestration system apparently for containers. What does that mean? How do I scale that? So those are some of the things that really slow that down. But the move to virtualization, so much easier than the move to containerization. Which is why not many people –And in the grand scheme of things, why containers aren't fully adopted yet.

[00:40:47] JM: Sorry. Sorry. Maybe it didn't understand exactly. So why was it so much harder? Why do you find containerization so much harder?

[00:40:55] JD: Because the concepts are a little different. Because we're now not replicating the entire server, but just replicating runtime inside of a server, but yet still needing orchestration and such. This confuses your business leader over the application. And that misunderstanding or lack of understanding is what creates the inertia.

[00:41:15] JM: Gotcha. To me, there's also a distinction. Because a lot of the virtualization stuff was closed source as opposed to the open source nature of the cloud native stuff. Does that make any difference for this age of transformation?

[00:41:31] JD: I would say not as much now. But I do think that it took a long time for enterprise to get their brain around the fact that open source is still valid to run an enterprise-grade inside of an organization. Really larger, more progressive companies adopted it pretty quickly. But open source created a mindset of how am I going to control security? How am I going to get the right version? Does that mean anybody can meddle in my software? Or is meddling in my software? So it created just another question mark. And question marks create inertia.

[00:42:09] JM: So given that you're the Chief Technologist for Microsoft technology at Rackspace, I take it, you have a pretty intimate perspective on, I guess, the reinvention of Microsoft from its past days. And Rackspace has had its own reinvention in its move from the focus on hosting to more focus on partnerships. Do you take any pages from the Microsoft book of internal organizational pivot?

[00:42:42] JD: Yeah. I think it's an interesting parallel. I haven't really thought about those two. Now, in fairness, my remit is now the all technology here. But absolutely, my first 10 years at Rackspace was very much Microsoft-focused. And their pivot has been nothing short of incredible, in my opinion. Watching a company go from being as closed as they were, as myopically focused as they were. And in a lot of cases, there a lot of rocks we can throw here. To being one that is so much more open. Their adoption of Linux is incredible. There's more Linux running in Azure than there is Microsoft workloads running out in their clouds, regardless of their location. A lot of their tools and technology are open-source-focused. They are one of the larger contributors to the open source world as well.

And when you think about what the heart of their transformation, was they became a customer-focused company. And they rallied around that mantra of how can we help – You could go back to their mission statement. What is it? To make every person in every country more productive or something, or achieve more something, to that effect. That's kind of a mission that you can rally a company around, and a set of tools, and servers, and services and all those sorts of things. And in a sense, that's what we've done at Rackspace, is we looked at our mission. And our mission here at Rackspace is we're a bunch of technologists who apply our technology to our customers to help them achieve more. To help them be more successful in what they do and achieve their own missions.

And so, for us, there's also a time component to that. Because of the fact it's not just, “Hey, I can upgrade you from VMware X to VMware Y, or Windows X to Windows Y.” It's we can help your company, your individuals and your technology evolve over time so that you can achieve whatever that business mission is that you've got as a customer.

And we've continuing to sharpen our pencil at that and what that looks like. And our expansion out into two professional services over the past three or four years is squarely focused on that. A majority of the people we're hiring right now are customer-facing professional services and advisor services individuals. And it's all about how do we help them through that transformation. Because again, we're not just moving data centers. We're transforming businesses through transforming applications.

[00:45:12] JM: Last point on the transformation of applications. So you wrote up a post about some predictions you had for 2022. And an area of that that I found quite relevant was talking about RPA and just more general data analytics and machine learning tools. What are some typical patterns for how a company that has been around for a long time starts to adopt these newer tools? Are there any case studies that you can share with how larger legacy companies have adopted these?

[00:45:51] JD: Case studies, not off the top of my head, other than cite Rackspace. We're a big consumer of RPA over here, robotic process automation. And it is literally saving us millions of dollars. We're literally funding other activities because we're able to do more – First of all, we're able to not have humans do half the things they used to do. And now we're able to apply those humans to work that's more meaningful. And that's ultimately the mission of RPA. And it's to take the mundane and scale it so that humans don't have to do it. They can work the exceptions. But now we can put them on work that actually has value. And that value can be anywhere else in the organization. That value could be the ones who are actually creating these bots to go out and do the thing.

I cite my little fictitious story of this fictitious person named Betty, whose job it was, was to transpose invoices from the fax inbox into the accounting system. Betty worked eight hours a day. She worked five days a week. She didn't work on the weekend. And she did not work overtime. And she was pretty finite. But she did know the business process for every aspect of

that thing. And Betty is an hourly employee. Betty does not have benefits. Betty does not have a long-term value in the company. But if Betty goes out and gets three months of training in automation anywhere, or fill in the blank, Microsoft Suits, or the others, and now she's creating those tools. Because the most important person in the business becomes the person who knows the flow of the business. If you know the flow of the business, then you're the person who can best understand what the right, in this case bot type of an environment, to get in and create to alleviate that manual boring job she has to do.

By the way, the person who does that says a starting salary of $65,000 in almost any company around the globe. Has benefits. Has long-term value and an amazing career, or beginning of could be a very rich programming career should they choose to go farther down that road. And that spawns easily into data. And as you look at those predictions, whether we're looking at machine tools, whether we're looking at machine learning tools, or whether we're looking at RPA or the other stuff, I think the opportunity that exists for enterprise in 2022 is how can the machines do more so that humans have a more engaging work to do that has greater value to your organization, and everyone continues to grow inside of your companies? Look, we're still in the middle of this great resignation. Everybody's flip flopping all over the place.

And I've read a really interesting article a couple of weeks ago from Harvard Business Review, and that was employers need to start recruiting their own employees again, and really helping them ensure that they see the value that they have in the organization, or help them find the right role that has that value.

So 2022 is about let the machines do the work, and then dig into the greater data insights, because that's where the magic is. When I talk about cloud transformation, I think of it in three swim lanes. So the first is infrastructure modernization, that lift and shift. Next one is going to be your app modernization or transformation. Truly getting into the cloud native tooling of containers and into serverless. And then the last I just call the innovation phase. And that's because we've moved so much of our data, of our applications, of our business processes out into a common cloud or two, that now we can start to play what if with our processes, with our data, and get so much greater insights. Because all of this data now exists in one cloud, let's get a pipeline that shoves it all off into a data lake and send some data scientist to go figure out what the correlations are between all of that stuff. Things that we could never do before, we can

now do, in some cases, for about 20 bucks, to see if it makes any sense in cloud spend. Where before, we'd spend 10s of 1000s to get infrastructure up to be able to play what if. This is the opportunity. Let's free up our people so that they can get in with that business knowledge and create better, greater insights that gives each of these businesses, or of our businesses, that greater competitive ability against their competitors.

[00:49:56] JM: Jeff, thank you so much for coming on the show. It's been a real pleasure talking to you.

[00:49:59] JD: Jeff, thanks. Man, it was great to be here.

[END]