EPISODE 1218

[INTRODUCTION]

[00:00:00] JM: AWS offers over 200 services as part of its infrastructure platform and that number continues to grow. Organizing all of these services and tracking the costs that they incur can be a significant challenge often requiring teams of AWS certified sysadmins working together to get a handle on an enterprise scale system. Vantage provides an alternative streamlined AWS console that makes it easier to manage AWS services and track associated costs. Users link their AWS account to Vantage and it automatically profiles all their services and aggregates the information into a dashboard. Users can customize how their advantage console appears and allows users to break down service usage by region.

Ben Schaechter is the co-founder of Vantage. Before founding Vantage he was a senior product manager at AWS and DigitalOcean. Ben joins the show today to talk about how Vantage helps streamline the AWS experience and why teams of all sizes can benefit from a better user experience on the AWS platform.

Ben, welcome to the show.

[00:00:58] BS: Hey, thank you. Good to be here.

[00:01:00] JM: You work on Vantage, which is an alternative AWS console. For people who have not worked deeply with AWS, could you explain what the role of the console is?

[00:01:10] BS: Yeah, for sure. So the console is used for a few different use cases typically for new people who are getting up and running with AWS. They'll use it to explore services. And for existing customers or people who might have infrastructure already, it's typically used for observability and debugging use cases. So they want to verify that some resources are in a state that they think they are or that the resources are there altogether, but those are typically kind of the main use cases that people use the console for.

[00:01:43] JM: You've worked in the past at AWS. What did you learn about customer use cases when you were at AWS?

[00:01:52] BS: So at AWS I was primarily focused on container services and I wasn't working on the AWS console, but I think the biggest learning for me from there was actually an organizational observation, which is there's a whole bunch of different service teams and each service team decides to staff their console resources a little bit differently or at all. And so I think the opportunity that I saw for Vantage was really having a centralized experience across all of the AWS services whereas each individual service at AWS has its own unique console experience, varying degrees of investment, and as a result kind of different experiences for each individual service console.

[00:02:40] JM: Tell me a little bit more about the problems with the AWS console that you've seen over time.

[00:02:46] BS: Yeah. So there're a few different problems that we typically hear from our customers. So first and foremost, for certain customers it's just very difficult for them to see what resources they have. So when you log into the AWS console, there're hundreds of different services and it's opaque to you which of those services you're using. And so first and foremost Vantage does a really good job of just showing you what resources and what services you're using on top of AWS. The second thing that people get a lot of value from Vantage is that we've done a lot of work on figuring out all of the various pricing permutations that can happen with resources on AWS. And the AWS console doesn't show very much or at all pricing information on a per resource basis. It requires you to go into the cost explorer service, provision a whole bunch of reports yourselves. Whereas out of the box, Vantage will go ahead and show you, “Hey, here's what you're paying for these EC2 instances or these container surfaces and so on and so forth without you having to do any configuration whatsoever.”

[00:03:59] JM: Now, building a solution to a UX problem that is endemic to a platform like AWS, that's a tricky business to build because you're basically building on somebody else's

platform. So can you tell me about how you got the idea for building an alternative AWS console?

[00:04:24] BS: Yeah, for sure. So the idea actually has been germinating for a while. And I'll explain the story of how we arrived here. But I've worked at another provider named DigitalOcean, and from my time there what I saw is that people love DigitalOcean because it's really easy to get up and going and provision some services. And after a certain period of time some of the customers would just reach a level of sophistication and have to migrate off. So the developer experience was great, but the underlying services didn't scale tremendously well. And from my time at AWS I saw the inverse problem where the underlying services were really fantastic, but the developer experience was always very difficult. So if you're a large organization that needs to scale, you can integrate all these services, but there's a lot of time being spent just learning the services, how they work together, how to configure things.

And so the idea started really simply is just what if we took the developer experience of DigitalOcean and applied it to the large public cloud providers? AWS being the first of which. And in terms of business value, I think that there's enough features that Vantage can provide either now or in the future that provide a lot of utility to customers. There're two things that I think are not going to change for developers going forward. One, people are always going to want easier ways to interact with their infrastructure, which I think Vantage can provide over the AWS console. And two, I think they're always going to want to simplify their understanding of their costs and where they're coming from. And I think Vantage has kind of a unique product offering that allows it to really provide a companion to the AWS console. It may not do everything under the sun that the AWS console provides, but for the information that's lacking there, people are willing to sign up and pay for with Vantage.

[00:06:22] JM: Tell me more about the specific problems that you are trying to solve with an alternative AWS console.

[00:06:31] BS: Yeah. So there're a few different problems that we see. I think first of all, we're just trying to provide a more seamless uniform developer experience across AWS services and

resources. So what this means is instead of you having to learn 200 plus different consoles on AWS, Vantage has a single one with uniformity across all AWS services and resource types. So the value that's being provided here is that we're really able to get you to the information that you're looking for more quickly. The second one is primarily around cost transparency. So making it very easy for you to understand where your costs are coming from. Making it very easy for you to see how things are trending over time and then eventually having features for you to actually mitigate against these costs. So I think that customers are going to find value kind of across the board in terms of just a more seamless experience that saves them time and eventually be able to save them money as well through some of the cost transparency features that we have.

[00:07:42] JM: Interesting. Like how would you be able to save a customer money?

[00:07:47] BS: Yeah, it's a good question. So there's a few different ways. So first of all, there's a very, very basic use case, which is just show me what I'm paying for. Aand surprisingly just by making it easy for people to see the resources that they have, kind of unlike the more individual and SMB basis, people just really don't know what resources they have or they might have resources in a region that they weren't using that they can go and clean up. Over time, Vantage can provide right-sizing recommendations or even helping you identify orphaned resources.

For example, if you're using infrastructure as code to spin things up and down, not always are they perfectly cleaned up. And Vantage can set conditions for these resources being present and alert you to them so you can go and clean them up. Another example is we're able to profile resources on a per resource basis. So a simple example is we might be able to see that you have an EC2 or RDS instance that you're paying on demand for that might be a good target for savings plans or reserved instances. But we're able to profile resources at an individual basis so we can say, "Hey, this EC2 instance or this RDS instance might be a good candidate based off of these metrics that we've looked at from CloudWatch. Or you know you've been paying for it on-demand and you don't currently have any reserved instances or savings plans. So there's things along those lines that we're taking a look at, and I think that's

a very basic example, but there's a long list of things that we could do to provide people with value taking a look at egress and request metrics from S3, looking at right sizing container services and so on and so forth.

[00:09:30] JM: This kicks off an interesting entire area of conversation, the cost optimization, cloud cost optimization branch of startups. There are so many cloud cost optimization startups. Is there room for the number of different cost optimization solutions? I mean, it seems like I guess I don't really know what the customer base is like. Are people looking for solutions to monitor their software and give them price recommendations?

[00:10:01] BS: It's a fair question. Interestingly, Vantage did not intend to be a tool that focused on pricing as much as it does today. What's been interesting is from our early access and initial launch, the vast majority of customer feedback is around cost transparency and cost optimization. So what we've seen is really the pull from the market and working backwards from the customer in terms of building these features. I think as we move forward we'll market ourselves less as an alternative to the AWS console and more as a tool that can focus on cost transparency. So I think the initial feedback from users has shown that there's definitely a need for it.

I think the other unique thing that Vantage is doing is that it's not even necessarily cost optimization, and I think if we market ourselves as cost optimization we wouldn't stack up really well against some of the other competitors in the market, but we've done a really good job at something that we refer to as cost transparency. It may be nuanced, but we're just purely making it easy for you to see what resources you have and what you're paying for, and over time we might be able to have some cost optimization features, but there's clearly a need for it from the uptick that we've seen and just the feedback that we've received. I think time will tell, but it's been a surprise to see kind of the feedback in that realm of what people are requesting.

[00:11:32] JM: So let's say I plug into vantage and I have my AWS infrastructure scanned and I can now use Vantage as my AWS console. Do you have the entire set of AWS console functionality built into Vantage? Because that's a whole lot of functionality.

[00:11:53] BS: It is. It is. So the way that Vantage actually works today is anyone can sign up. You confirm your email and you authenticate with Vantage through something called a cross-account IAM role, and that cross-account IAM role is actually defaulted to read-only permissions. So by default you are not able to take any right actions or management actions and it focuses primarily on observability of your infrastructure over time, and we have a few of these right actions that are available in Vantage today. Let's say you wanted to complete a right action or an update action on AWS using Vantage, and let's use the example of Route 53 because we have support for that right now. Vantage will show you the actions that are available to you. So in this example might be create a hosted zone or update a record and it will be aware of what permissions it needs to get the job done and be able to request those permissions accordingly. There're some specific Route 53 IAM permissions that you would just grant to Vantage and it would get the job done.

I don't think that Vantage is trying to be feature parody one-to-one with AWS. I think AWS has to provide a large number of management actions through the console just by the fact that that's their business and they need to provide that. I think where vantage can provide a lot of value is actually some basic write functionality on top of certain services, but also offer kind of like what I would describe as like pro functionality or pro tools on top of AWS. So one very basic example might be, “Hey, I want to assign a particular IAM user a certain permission with a TTL,” so that they have access to something for a certain amount of time and then that permission is kind of tossed out. And that's something that Vantage can do really well because it's coordinating multiple different actions. It's persisting the state of what needs to happen with those actions going forward. And so I think that's more the realm of where Vantage will go and we're not trying to be a one-to-one feature parody with the AWS console, and primarily because we think people are pretty well-baked into infrastructure as code or other automation tools for interacting with their infrastructure. And in some regards we don't believe that you should be doing a lot of right actions through a console.

[00:14:23] JM: What are some other best practices that you try to encourage developers to be ushered into through the design advantage?

[00:14:33] BS: Yeah. In terms of best practices, I mean, I think some things that Vantage does particularly well is we create essentially like a relational model of all of your resources behind the scenes. So you can see related AWS resources very easily. So let's say you have an EC2 instance and want to see related EBS volumes and networking primitives. We group things together very seamlessly and kind of by default on Vantage. The other thing that I think is more interesting is Vantage has this concept or primitive called a view. And a view you can think of as essentially a set of conditions that matches a subset of your AWS infrastructure. And so our customers are using Vantage to group resources together kind of for different applications or microservices or environments and this allows them to answer questions like, “What is my staging environment cost versus my production environment? What is my search service costing me versus my caching service?” And so Vantage takes a different approach, whereas on AWS you're kind of grouping things by service or by region or by account. Vantage does a really good job of actually logically grouping things based off of what matters to your infrastructure directly. And so that's a lot of the utility that people are getting out of Vantage right now, and with those groupings you get additional features like cost transparency reports as they exist right now, and in the future there's other avenues that we're looking at in terms of security recommendations or other things along those lines.

[00:16:19] JM: If you were trying to convince somebody who had been using the AWS console for a long time to gravitate towards Vantage, what would you say to convince them?

[00:16:30] BS: In general, there's not too much that we have to say. The one liner that we have on our website and what we kind of say publicly is that we're an AWS companion or console that's focused on developer experience and cost transparency. One truism that I’ve kind of seen across the board is that most AWS customers are not particularly satisfied with the console from a usability experience and they're not particularly thrilled about being able to understand what resources they have and what they're paying for.

So even just saying it as simply as that tends to have people at least get up and onboarded with Vantage. And then from there what we've seen is we get a whole slew of different use cases or feature requests from people that we're actively building in. But I would say right now I think in some regards we have a more seamless experience just for interacting or observing your infrastructure and over time it will get better and better.

[00:17:34] JM: So you've mentioned working at two cloud providers. You've worked at both AWS and DigitalOcean. What have you learned about the difficulties of designing a cloud provider? And do you think that those kinds of difficulties will bubble up eventually? I mean, as Vantage grows, I imagine the surface area of the product will grow and you're going to have to grapple with the same kinds of issues that these companies have grappled with at scale.

[00:18:02] BS: Yeah, it's a good question. I mean, I think – So from my time at DigitalOcean, I think when I was there they had a fairly small set of products that people would leverage. So an equivalent to EC2, an equivalent to EBS, an equivalent to S3, and it was really easy for people to get up and going with basic applications. And at AWS, it was the opposite where you had everything under the sun. But DigitalOcean's feature set has grown tremendously, and I think there's a consistent problem across all cloud providers that working at these two has taught me, which is all of them have the same bunches of underlying primitives for you to build your applications. The thing that's particularly challenging is just learning how they all inter-operate together. How do you use IAM with all the other AWS services or on DigitalOcean? How do you actually provision a database and make it accessible to your droplets or EC2 equivalents? And so I think one thing that Vantage can do going forward, and there's a whole deployment story that we haven't talked about and actually tying different services together that I think is like a big, big area of value for customers and something that we could do particularly well having experience in public cloud.

I mean, one very basic example is like I think what Heroku has done and has been doing for years is really, really awesome, and they've simplified down the user experience but still allow customers to scale. And I think there's a world where a Heroku-like experience could exist, but

it's ultimately backed by resources in your AWS account. And so I think that there's a chance that we may get into that story of deploying and managing infrastructure. It's likely not something on the near term, but if we were to walk down that path, then we would definitely be using kind of best practices learned from these public cloud providers to simplify things down and give you a more seamless way of managing your applications.

[00:20:18] JM: Are there any challenges to implementing or onboarding a customer onto Vantage if the footprint of their infrastructure is really large? Like if I have a huge deployment and a huge AWS deployment, are there any roadblocks that I'm going to hit if I'm trying to onboard with Vantage?

[00:20:41] BS: Yeah, it's a good question. So I'll tell you that we're still learning ourselves. We have customers in the hundreds of thousands of dollars of spend per month and we've talked with a few in the millions and tens of millions per month and we haven't had anything break at this point in terms of reaching people at scale. We've been pretty intentional about how we're actually interacting with getting information about your account or accounts on AWS and have done it in a way that we believe can scale. That being said, we haven't had like a customer in the millions or tens of millions of dollars of spend actually connect.

Some of the things that we're seeing in terms of early challenges are working with AWS organizations and sub-accounts. I can say that I think we're doing a pretty good job of handling it right now and that experience will get better over time. But in general, the vast majority of people are able to get up and going have all of their resources imported correctly in the events that we don't have kind of support for one of their AWS services. They typically get in touch with. I think recently we had a customer who mentioned that we didn't have support for Elasticsearch. We added support for it in about three or four days. So we can move pretty quickly in terms of additional AWS service support. I mean, I think it's like we'll learn as our customers push us as with everything else. So no major concerns at this point in time, but I think we're still waiting for the ultimate test of one of these really, really large customers.

[00:22:23] JM: When you're building Vantage, the biggest problems that you encounter are they engineering problems or are they design problems?

[00:22:31] BS: Yeah, that's a good question. I would say it's a mix of both. I mean, I can tell you our designer probably thinks that there's considerable challenges on the design front for just the surface area like you mentioned earlier of AWS services and resources and making them all kind of fit together nicely. And at the end of the day it's a whole bunch of data and random attributes and things along those lines that are specific to each individual resource. So I think there's like a whole bunch of work that's just a lot of difficulty on making it easy for people to manage their infrastructure.

In terms of engineering, I don't think that Vantage has had a particularly difficult time in terms of engineering challenges. There was a lot of early work done on our syncing engine with AWS to make sure that we can sync accounts very quickly in a scalable manner, and that was a lot of upfront investment, but the majority of our engineering efforts are more so around feature development and product development that I don't think is too difficult relative to other companies. But I would say our sweet sauce in terms of a differentiated offering is mainly around making sure that all of that information is there and a really seamless and intuitive developer experience. And that's probably going to be some of the greater challenges that we have going forward.

So to kind of distill that down, I think design is probably more of a challenge for us going forward then engineering. And just to pull on that thread a little bit more, design we look at is not only visual design, but what's the user experience in terms of different flows around the Vantage console? What is the user experience around our documentation? What is the user experience around our support? And looking at things in the future, what is the user experience around potential CLI offering or cost savings? So there's a lot of thinking that's going into that going forward.

[00:24:39] JM: Can you just tell me more about the data flow? Like I've got Vantage, it's a big layer of monitoring and dashboarding and console over AWS. How does data from AWS get pulled and propagated into the Vantage view? Because you're sucking in a lot of data.

[00:25:01] BS: Yeah. So like I mentioned earlier, we interact with accounts through something called a cross-account IAM role, and we provide a cloud formation stack to create that role. And the reason I mention this is that it basically lists a large number of permissions for the APIs that were we're calling on behalf of your accounts. And if you take a look at those, it's basically a whole bunch of list and describe actions. And while AWS is busy managing not only these APIs, but also the underlying services, we can operate with a little bit more swiftness, because ultimately we're just looking at the metadata about all of your infrastructure. And so our syncing engine that we've built for Vantage will kind of periodically take a look at the various resources that are in your accounts by using these list and describe actions and do a full diff of have you scaled things up? Have you scaled things down? Have you updated configuration? And we do that across all AWS services and all AWS regions in about 15 to 20 seconds. So we spent a long, long time in terms of making that a very quick experience.

In terms of the volume of data, it's growing and it's not nothing, but it's not a huge engineering challenge in terms of collecting that data, indexing and making sure it's very, very quick. And we feel that all of this was necessary just to provide a really fast and intuitive experience on Vantage. So I would encourage people, if you are used to an experience on AWS of clicking around and trying to get the information and you do similar things on Vantage, chances are it's going to perform incredibly quickly on Vantage relative to AWS just because of how much time and effort we've put into this syncing engine and how we store that data.

So that's essentially what we're doing behind the scenes. Customers also have the ability, if they would like, to issue like a hard refresh on Vantage. So in the navigation bar of the Vantage console there is a refresh button, and that will issue a sync across all regions and all services at that point in time. And again, it takes like 15 to 20 seconds. So if you do make an immediate change you need it picked up by Vantage, you can quickly get that picked up. But anyway, that's a high-level overview of what we're doing behind the scenes.

[00:27:28] JM: So Vantage is pulling data from AWS? It's not like AWS is pushing data to Vantage?

[00:27:34] BS: Correct. Yup, that's correct. We are periodically pulling for information.

[00:27:40] JM: Gotcha. Now, I think we should revisit, what exactly are you trying to accomplish other than the improvements in the UX? So obviously there's the cost improvement, but you could – I mean, if you're trying to build a layer of usability over the entire AWS console, you could get into logging and monitoring and metrics. You could basically get into everything that AWS does if you wanted to. You could almost be like a Heroku style situation going on if you wanted to. I guess I'm just curious a little bit more about what the vision is and how you're honing that.

[00:28:20] BS: Yeah, it's a good question. I think that some of this we are still actively exploring from feedback from our customers. Everything we do at Vantage is entertained and governed by what people request and give us feedback on. A large part of the feedback that we've received is around pricing. And so I think that's one vertical of where we're going to focus in the near term. The way that I look at things going forward is there's likely three to four other verticals that we could explore and we would do so each intentionally, but I think all of them are kind of governed around these tenants of saving the developer time and saving the developer money. And I think that there's a few different ways of looking at this.

So let's take a vertical of security. I think security can actually hold up to both of these tenants, and this might be things like Vantage sees that you have security groups that are too open or S3 buckets that are publicly available. One, I think, Vantage is your companion that can just be monitoring these things in the background so you don't need to worry about it and alert you to them. So it's saving you time. And it's potentially saving you money in the events that something is going wrong, right? Either through losing customer trust, or having an exploit, or something along those lines. And I think there's a whole bunch of other verticals that we could walk down that would also fit these two kind of tenants or principles. So deployments, like I

mentioned, is a whole area where I think we could save people a ton of time and time being spent you know researching different services, researching how they actually interact with each other to deploy an application. I have no doubt in my mind that we can save engineering time there and likely save people money as well. So have you misconfigured things? Do you need to worry about service auto scaling that we can handle on your behalf? So I think that's a whole other vertical that we can walk down.

In terms of where Vantage can go, we're super, super early in building the product in the company and I think we're maybe just shy of being six months in and we're really happy with where the product is at and the feature set is at. But I think we'll have a lot more clarity on what makes sense as next steps in another six months or in another year. But right now I think we're just so firmly focused on pricing and cost transparency and developer experience on navigating your resources that that will likely keep us busy for another six months at least.

[00:30:58] JM: Can you go a little bit deeper into the pricing stuff? Like how do you have the insights to give recommendations on pricing?

[00:31:07] BS: Yeah. So what Vantage has done behind the scenes is we've basically built our own pricing database based off of all the various price permutations that can happen on AWS, of which I think we've got a couple million different price permutations. So taking a look at all EC2 instance types across all regions, across various storage costs. There's kind of a whole number of different verticals and life cycles. So is this on demand? Is it spot? Is it reserved? Is it savings plan? So when you have a resource that's important to Vantage, Vantage is able to take a look at what that resource is and apply a certain rate to it just for getting a price.

By nature of importing the individual resource we're able to look at things that other providers don't necessarily have imported. So we're able to look at things like CloudWatch metrics. Like are you using CPU memory and is there egress coming from this instance? And I think that data allows us to make really informed recommendations at one point. For example, let's say we see a customer with an EC2 instance that is a compute instance, but there's no CPU or network on it. Vantage can identify that and make a recommendation to you that other

providers can't necessarily do. They might be able to say, “Your EC2 costs have increased.” But what Vantage can do is, “Hey, your EC2 costs have increased. Here's the line item cost of each one of them. And oh, by the way, we see that there's no activity on these certain instances. Do you want to “right size” them or get rid of them or do you want to set conditions to right size or get rid of them?” And I think that's a whole area that can really be interesting for customers that Vantage can do very, very well. And so, yeah, I would say that's kind of how I think about that going forward in terms of what we can do to provide customers with value.

[00:33:13] JM: Do you have any perspective on where the biggest cost savings lie? Where are people losing a lot of money?

[00:33:22] BS: Yeah, it varies from customer to customer, organization to organization. There's basic things that people I think either don't know, don't think about or honestly just have difficulty doing themselves. For example, reserved instances and savings plans are a whole huge area that what we see is that kind of larger organizations who have higher bills kind of take a look at these things, but one thing that's always surprised me is that I think these things can provide value to the long tail of individuals and SMBs that may not be taking advantage of them. So even if you're someone with a single EC2 instance, you could purchase a reserved instance and save 30% typically at least if you're willing and able to pay up front. And so reserved instances and savings plans I think are some easy wins. After that, it's kind of looking at each individual AWS service and figuring out what are the top three cost centers of that service. So with S3, there's storage, but there's also request costs that people may not be aware of. So AWS actually charges you per thousand requests on S3 and they make it hard for you to actually even see what number of requests you're getting on that S3 bucket. And oh, by the way, they have different pricing for different request types, and this is on top of the actual egress traffic that they're charging you.

So one example is Vantage can profile your S3 buckets show you with high-fidelity what request metrics you're getting, what the egress costs are and give you recommendations. But anyway, this is essentially an anecdote that we can apply and we'll be applying to every paid service that we support in terms of –With EC2, the top thing might be reserved instances,

savings plans and right sizing. With S3 it might be fronting that with a CDN or just changing the way that you interact with accessing those resources. And with other services, there're kind of more bespoke ways for you to save money. So that's kind of how we think about it. And I think as we walk down this path of feature development we'll be taking a look at each individual service, but typically we can save people a lot with the top 10 services on AWS and go from there.

[00:35:45] JM: Tell me about your go to market strategy? How have you found users for Vantage?

[00:35:51] BS: Yeah. So we launched on January 12th. We have not done a ton actually on purposeful marketing or go to market. The launch day for us involved having TechCrunch cover our launch, which was really great. And we also did a show Hacker News post that was extremely well received, which I'm really thankful for. And we have some other ideas around go to market and channels for ways that people can find at Vantage that are largely going to be based around pricing transparency, and we'll share a bit more on that in a few weeks. But almost everything that we've had in terms of people finding us or customers signing up has been word of mouth at this point. We haven't spent any time or money on advertising and we continue to have a steady stream of people finding us. So from speaking with customers or people getting interested in Vantage I always ask them like, "Hey, how did you hear about us?" And typically it's, "Someone told me about it, or I saw it on Twitter, or someone was discussing this in our company Slack." And so we've got some organic momentum. Over time, I think there's a whole content play as well and this is something that I saw that DigitalOcean did particularly well where content about infrastructure if done well can really build up organic SEO in terms of helping people accomplish problems. They'll learn about Vantage. Some subset of those users will actually sign up for the service. And so we have some ideas around that stuff going forward, but that's kind of where things are at from a go to market perspective right now.

[00:37:32] JM: Do you see it as necessary to expand into other cloud providers or do you think just focusing on AWS is enough?

[00:37:42] BS: It's an interesting question. So we started with AWS primarily because we believe that they're the market leader and it's a large enough market to tackle. We have had some customers contact us about Google and Azure and DigitalOcean. I think we're going to dip our toes into this soon with showing prices across different cloud providers, which I think is going to be really interesting. So for example, when you're in the AWS console, all you see are AWS services, but there's kind of an idea where it might be the best thing by the customer to actually market them other services. And this might be other public cloud providers. It might be things like Cloudflare. It might be things like Datadog. And I think that's like a really interesting way for us to expand our reach and do the right thing by our customer.

I mean, one very basic example is if we allowed you to integrate with other providers like Cloudflare and we can see S3 bucket egress happening, why can't we facilitate an integration between those in a click of a button to front all of your S3 or some of your S3 with Cloudflare and instantly save you a ton of money and potentially time just with setting up that integration. And so I think you know there's thousands of startups that are building businesses that may have competitive products with AWS that are actually better products than what AWS provides themselves. So we don't even necessarily relegate things to just other public cloud providers, but we think of, “Hey, there might be a subset of people who compete with AWS on EC2.” We might think GCP, Azure, DigitalOcean. But on the 200 plus other services, there's a whole bunch of other providers that could be really interesting. And so I think that's one area where you're going to see a lot of exploration and support for in the Vantage console going forward.

[00:39:45] JM: Do you have any other perspectives on how cloud will change in the next five to ten years?

[00:39:52] BS: I will say that I would expect there to be probably more primitives launched than less. I mean, I think from what I’ve seen, and it was the same at DigitalOcean as it was at AWS, there're so many unique things to so many different organizations. And I think AWS gets in some regards a bad rap around the complexity that they have with certain products. Like I don't even know how many EC2 instance types there are. But the reason those exist is because customers have really unique needs, and I think this is only going to change over time.

I expect over the next five years there to be no slowdown of people converting from on-prem to cloud, and with that there will be more use cases that provide more primitives.

I think there's going to be a lot of companies that are trying to do things like what Vantage does and like what other kind of products in the market do, which is try to simplify the interaction between these primitives. And so you see this happening in kind of certain arenas. Like I think that the rise of success for companies like Versailles or Netlify is really just simplifying down a few different AWS experiences and providing a good developer experience on top of that. And so, I mean, I think that's the extent to which I've thought about it. I would say there's probably going to be more, not less products. And I think that there's an opportunity for people to kind of simplify the experience across these products going forward, but that's really the extent that I've thought deeply about what's for cloud in the next five to ten years.

[00:41:36] JM: If you aren't working on Vantage, do you have any other ideas for businesses in the infrastructure space?

[00:41:45] BS: Ooh! This is a good question, because actually something recently popped up that I think is just such a product that would be – I mean, we would pay for this at Vantage, and I've seen people on Twitter talk about this kind of stuff. And I think the idea is around providing a product that allows people to build CLIs, and I'm happy to explain this more if it doesn't make sense. But I think it's really hard for people to create CLIs, command line interface, for interacting with their infrastructure. And I think there's a pretty sizable business that could be built around scaffolding out a CLI for startups or businesses or whatever, whoever wants one, and then having a translate, like an API proxy layer to whatever their backend is for actually serving that CLI. And I think this would be a great business because you could charge a subscription just for the CLI and then charge for the number of requests that the CLI is getting. I don't know, I think that's one area. Like if I wasn't working on Vantage, I'd probably be trying to do something there because there's so many startups that are getting into dev tools and all of them are trying to build individual CLIs and some of them can do it, others can't, and then there's a whole bunch of other startups that just don't have the time or expertise to do

something along those lines. So yeah, I would happily invest in someone doing that or become a customer of someone doing that, or if not working on Vantage, that'd be an area I'd explore. But beyond that, I think I would still be working in cloud somewhere to be honest. I mean, I think there's just so much opportunity and the market is growing so quickly, but I don't have any other particularly interesting ideas beyond that one.

[00:43:29] JM: Okay. Well, that seems like a good place to close off. Ben, thank you so much for coming on the show. It's been a real pleasure talking to you.

[00:43:34] BS: Yeah, thank you so much, Jeff. It was nice to meet you and I appreciate the time.

[END]