EPISODE 1277

[INTRODUCTION]

[00:00:00] CQ: Hello, I'm Corey Quinn, Chief Cloud Economist at The Duckbill Group. I also host two podcasts myself; The AWS Morning Brief, and Screaming in the Cloud. But because that's not enough to keep me going, I write the Last Week in AWS newsletter as well. I've been taking over hosting duties for Software Engineering Daily for this week and taking you on a tour of the cloud.

The fifth of those five episodes, or final one, today focuses on a bit of a strange cloud outlier, for lack of a better term, Oracle Cloud. Here today to suffer my slings and arrows is Oracle's group vice president, cloud engineering, Cloud@Customer, Salman Paracha.

[00:00:45] JM: A few announcements before we get started. One, if you like Clubhouse, subscribe to the Club for Software Daily on Clubhouse. It's just Software Daily. And we'll be doing some interesting Clubhouse sessions within the next few weeks. And two, if you're looking for a job, we are hiring a variety of roles. We're looking for a social media manager. We're looking for a graphic designer. And we're looking for writers. If you are interested in contributing content to Software Engineering Daily, or even if you're a podcaster, and you're curious about how to get involved, we are looking for people with interesting backgrounds who can contribute to Software Engineering Daily. Again, mostly we're looking for social media help and design help. But if you're a writer or a podcaster, we'd also love to hear from you. You can send me an email with your resume, jeff@softwareengineeringdaily.com. That's jeff@softwareengineeringdaily.com.

[INTERVIEW]

[00:01:48] CQ: Salman, thank you for joining me.

[00:01:49] SP: Hey, Corey. Thanks for having me. I appreciate you bringing me on the show.

[00:01:53] CQ: Oh, of course. So you are Oracle Cloud's group vice president, cloud engineering, Cloud@Customer. Did someone get bonused by every time they slipped the word the cloud into your job title?

[00:01:49] SP: I think it was a nice coincidence that all came together. Part of the naming philosophy at Oracle also comes from a bit of past behavior if you like, because some folks have come to work from a cloud provider that does a really good job of naming. So I think that carries over here as well to some degree.

[00:02:20] CQ: I'm relatively unconvinced that any provider in the world, when it comes to cloud, has done a good job of naming services. You can make a possible exception for DigitalOcean. But by and large, every other provider tends to sort of befit for lack of a better term.

[00:02:35] SP: Well, it's a long name, but the goal is to make sure it's descriptive enough so people can remember what we're trying to go pitch to them when we talk about the third category.

[00:02:43] CQ: So let's talk a little bit about your own trajectory. You were personally at AWS for, I want to say, seven years, something like that?

[00:02:52] SP: Yeah, close to eight years. Would have been eight years if I stayed another three and a half more months. Yeah, eight years,

[00:02:59] CQ: Which feels like forever for those of us who've never stayed at a job longer than two years ourselves until I started this place. But you were there doing a lot of things. You were promoted from a principal PM role, and eventually became, if I'm not mistaken, the general manager, or I guess highest person who is directly responsible for a contained service for the serverless application repository.

[00:03:19] SP: That's correct. I started my journey actually in amazon.com, where I lead a product called Amazon Coins. That's actually how I got into Amazon. And Amazon Coins was this rewards currency for our Kindle Fire ecosystem. So customers spend some, they get some

coins, rewards currency. And then they can spend those coins back on apps, games or in game in-app purchases.

I always had engineering roots. I have a computer science degree from the University of Florida. And so I just want to get back into more engineering roots. And so I joined AWS in 2014. I looked over workspaces, or end user computing product, and from workspaces. That's where I became a principal PM, at workspaces. And then I moved into a single threaded leader role to launch services, which is the application repository, and also nurtured as a GM for the serverless application models. I owned engineering and product for both of those services. One is an open source tooling framework, if you would, and the other is a service. But that's what I looked over before I joined Oracle back in 2020, last year.

[00:04:20] CQ: Yeah, that's hard to believe, but it's been a little over a year that you've been there. And in that time, you went from a VP of product management specifically for dedicated region Cloud@Customer, which we'll get into, and were promoted to group vice president with a bunch of the word cloud recurring after that statement. So, congratulations. Promotions at most different companies tend to be very different things. And even at a place that's a rapidly growing as Oracle Cloud, that's no small feat.

[00:04:49] SP: Okay, thanks for that. I think this additional responsibility is a lot of support that I get from the current set of people who brought me in. So I appreciate them leaning in. I think I've got a whole bunch to do to prove myself out in this role, but how can you take it on and be part of this journey.

[00:05:04] CQ: Back when you first joined Oracle, you and I had had a brief conversation, because I've heard this from a bunch of different folks who I knew in the AWS ecosystem, sometimes customers, sometimes AWS employees, and you all sort of come and approach me with this sort of similar hangdog look about you, like I've taken a new job over at Oracle. Like you expect me to somehow hit you with a stick or something. And the first time someone did that, it's, "Oh, my God! Why would you work at Oracle? Oh, my God! They're paying you in yachts?" And the answer was basically, "Yeah, pretty much." But the more people I see starting to make the transition into working at Oracle Cloud, the clearer it became that there was something else going on here, something that fundamentally went well beyond what we were

seeing as far as just money explaining it. What was it that drew you to decide, “You know what? Oracle, one of the most hated names in tech in some ways, I'm going to go there and help build out their cloud.”

[00:06:01] SP: Yeah, it's a lot to unpack on that one. But I think that when I looked at any job, at that time, think about what I’ll gain personally in terms of skill set, I was looking at what Oracle is trying to go solve for and what I'll be able to do for this organization. So there're a lot of my past AWS folks that have worked with directly at AWS who are now at Oracle. So that obviously helps in, “Hey, what are they trying to go to?” They are trying to solve for this via fourth cloud provider. Why would they want to do that, or sort of secret sauce there that they can latch on to. And so my conversations with a lot of the engineering counterparts that I worked with on duty basis did lead me to believe that the core fundamentals of what they have built, what they call their second generation cloud. We can get into naming. Because they actually had a first generation cloud offering, which was appliance-based, which didn't quite really meet the market, what customers wanted.

[00:06:54] CQ: No. It really needed the quotation marks around cloud in some respects. So the gen two cloud, which is often talked about by your cofounder, Larry Ellison, on Earnings Calls. And I always hear Gen two is if it's the Linux distribution, which, “Oh, that doesn't sound great. Oh, nevermind.”

[00:07:09] SP: Yeah. So Gen two was an actual cloud offering that built from the ground up that had a lot of the core principles of what cloud was about, which was instantly scale, pay for what you use, just-in-time provisioning, etc., the benefits of the cloud that as customer saw it. the thing that Oracle knew about its customer base was that a lot of their customers were using legacy software, and they continue to have some legacy software. And the customers they serve have a very different journey to the cloud. So a lot of customers have changed the way they think about cloud journey and have adopted certain principles and services, for sure. But Oracle really wanted to sort of say, “If I were to get customers to using –” Of course, Oracle products, and Oracle serves not just database. It has a whole bunch of software products in our global business unit space and others, and get them to sort of seamlessly move to cloud. What would that looked like?

And I think that's where that story got a bit interesting. For example, there's a whole bunch of investments that Oracle has made an L2, or layer two, networking to virtualize it. That's tough engineering work, because usually what happens in a cloud environment is a layer three, the IP packets are virtualized so that they can have an overlay network and you can send traffic from one instance to another instance. And it all feels as if you're in a physical network, but you're really not.

So as I enumerate a few things, so they really looked deeply at, “Hey, what if we were to do layer two? What if we were to start with off-the-box virtualization for increased performance of networking?” And networking became a very core focal point. For example, Oracle was the first ones to offer RDMA or a 100 gigabit per second network throughput between instances at scale in the cloud, because customers were doing all these tough workloads and didn't want to have to change their workloads to move or take the advantages of the cloud, the price performance, the data ingress/egress challenges, etc.

So I can elaborate on these things in more detail. But the fundamentals of building a cloud offering that feels very much like on-premises software with little change and enabling customers to lift and shift, I think has merit. Technical merits for sure. And I think customers are just getting familiar with what Oracle Cloud can offer to them. Because, frankly, Oracle, in a second generation offering, has been relatively late to what was already established in the marketplace.

[00:09:29] CQ: Right. And historically, it's been mostly a story that I've seen in public around, “Oh, if you have existing Oracle workloads, this is the best place to cloudify them,” for lack of a better term. That said, I'm starting to see breakthrough success stories were folks who do not have, at this day, disclosed on-premises Oracle environment are starting to seriously consider Oracle Cloud for different workloads. The thing that caught my mind, you allude to networking, but let's be very direct here. You are 10 times less expensive for outbound data transfer than AWS, which is Similar to the other tier-one providers. So for anything that's heavily data transfer-based, that immediately becomes an object of significant interest and concern.

Now, everyone loves to say, “Well, once you're at significant scale, of course, no one pays retail.” Sure. But that also applies to you folks, as well. And when you're sitting here trying to

sketch out a startup that might be heavy on data transfer, you're going to use the publicly posted pricing just to get back in the envelope calculations. And when your orders of magnitude away from being a viable answer, a lot of the expensive retail pricing doesn't seem to make a whole heck of a lot of sense. So folks will never have those conversations with a number of cloud providers just because it doesn't look externally to be viable.

[00:10:44] SP: Yeah, there's some truth to that for sure. One of the philosophies, if you would, apart from the engineering principles, and I chatted a lot of the principal engineers on AWS who now come in and have built out Oracle Cloud, on top of the engineering principles was a simplicity option. So for example, pricing simplicity, Oracle Cloud will have same instance price across all regions. So now you can budget for it, but you have to think about what are the instance costs in one region or the other? And now we do have about tens of these regions. So it's not trivial. And so simplicity of pricing and this whole notion of having a single bucket of credits that you purchase and you can deploy it against any services was also novel out for its time.

The simplicity of pricing, transparency of pricing, and what you can easily use to construct your budget against I think it's an important additional element of the Oracle Cloud story, which I think customers are slowly but surely seeing and taking advantage of. Clearly, the customers that we have on our website that we talked about that have used Oracle's egress to its advantage is Zoom, and others in that same bucket like 8x8. But just the general principles of making pricing simple using same instance pricing across all regions for all types of services, it's a big deal. And I think customers are really appreciating that.

[00:12:01] CQ: They seem to be. My question for you is are you seeing an improved uptake of Oracle Cloud from customers that do not have pre-existing deep-seated decades-long relationships with Oracle proper?

[00:12:16] SP: For sure. So I see a few parts of the Oracle Cloud business. So I can speak to those more convincingly. I think if you look at all the networking and **[inaudible 00:12:26]** we've done. So a lot of our customers are not thinking about how to really run VMware workloads on Oracle Cloud, apart because those VMware workloads can't live without L2. VMware application is we're ready for physical networks, ones particularly with clusters of compute nodes that share

same broadcast domains and use features for on-premises that are supported in L3 virtual network stack.

A few examples of that, like assignment of MAC's and IPS without a proceeding API call, low latency reassignment of MAC and IPS for high availability and live migration, multiplexing by MAC address, etc., etc., and VLAN support. So we are seeing a lot of customers really interested in running their VMware stack on Oracle Cloud, because it looks very much like bare metal that they were running on-premises. But taking advantage of the scaling function of the cloud, taking advantage of the price performance function of the cloud and being able to really be more agile. So that's one category.

The other category related to networking as well is our HPC fabric. Because HPC requires nodes and clusters to be able to speak to each other, whether that'd be computational fluid dynamics, whether that'd be visual rendering, whether that'd be Monte Carlo simulation, they need the network to be really performant at all times at every time. And so HPC and the use cases there is Mazda, Nissan, and a whole bunch of research that's happening even in MIT that takes advantage of both our instances and our network stack to achieve certain scale points. I think we're seeing a lot of customers use it for that. Simplicity of pricing, of course, plays into all that, but HPC, our legacy infrastructure, such as customers on-premises and VMware installations are really attracting a lot of our customers. The other category, which I'm not super close to, but others in Oracle are, is our ISV category, which are looking for simplicity of deployment and significance.

[00:14:21] CQ: Those are independent software vendors, correct? So it's always good to disambiguate these things. So sometimes if I have to think for a second, I'm probably not the only one. Please continue.

[00:14:29] SP: No, no. Sure. So independent software vendors as a whole category of customers that are continuing to build an Oracle so that they can take advantage of the price performance that Oracle offers by providing a superior service to their customers and take advantage of cloud. And in that space, we have a whole bunch of customers and have already publicly talked about like **[inaudible 00:14:48]** and Zoom and Altair and others that are in that bucket. This workforce recently that we added to that as well. And so these are non-Oracle

workloads. They have nothing to do with Oracle technology. Some of them may have been running Oracle databases in the past, maybe a portion of it. But by and far large, their purpose in life is to develop a new modern stack or deliver some value. And Oracle has been able to support these types of workloads.

[00:15:16] CQ: Let's dive into the HPC story that you just mentioned a minute ago, because one of the things that I found when I talked to folks who are using HPC, or high-performance computing, is outside of some very specific use cases generally found in academia or research, a lot of HPC customers don't think of themselves in the context of HPC. On some level, this harkens back to the problem of, “Oh, this is for your data lake,” and you talk to folks, and “Oh, we don't have a data lake.” And you sort of stare at them and, “No, you have about 30 petabytes of data sitting in your object storage. What do you think that is?” “Oh, you mean all the logging nonsense, or all the assets, or all the insert fill in the blank here.” It becomes something that customers don't self-identify with. So what is HPC for the vast majority of people out there who might not know what that is, first of, and secondly, might not realize it applies to them?

[00:16:07] SP: Yeah. Fair enough. So HPC is this category of compute technologies where you're really thinking about putting clusters of computers close to each other, so they can achieve as unit of work. So in manufacturing, this could be computer-aided engineering, computational fluid dynamics. Figuring out if your car is really going to – If it crashes, is it really going to sustain body injury as an example? Or is it going to be able to – If it falls off a bridge and goes in the water, what's the rate at which water will flow into the car? A lot of interesting simulation work in manufacturing, but it needs a lot of computational power to run these types of simulations. And they all need to be done in concert with each other trying to achieve a unit of work.

Similarly, in finance, it's like trading platforms and risk modeling exercises. And media is visual effects rendering, which requires like rendering farms. And you need a whole bunch of those. And they need to do coordinated work and complete that coordinated work in a reasonable amount of time. Some of that work is not showing up in sort of deep learning as well. It’s like coordinated works. It's not like the old classical high-performance computing, what has been labeled. We’ll be seeing that technology transfer into that type of workload as well. And research, and medical research, drug discovery, genomics, climate change. All of that requires a

whole bunch of computer servers, we put together, do coordinated work, run algorithms, and spit out a result and do that quickly and reasonably well. Last but not least, life sciences, analyzing the physical movements of atoms or molecules to perform genomic sequencing.

So these are all the various types of industries that use high-performance computing to achieve a result. Hope that helps the audience what HPC really means. And for that, you need bare metal instances, low latency cluster networks using 100 gigabits per second RDMA technology and high-performance storage to achieve that outcome. So I hope that it gives you a bit of color what HPC is.

[00:18:09] CQ: No. It absolutely does. I've spoken to a couple of large hedge fund style computational-heavy workloads. And what they wound up discovering was, for what they were doing, they already built out a data center. They couldn't handle interrupts for their application stack super well. And they were debating. Do we build out a second data center? Do we build out a cloud environment? And the answer became pretty clear, this was a couple of years ago, that there was no viable path economically to put this in the cloud unless they got some truly astronomical discounting, which is always on a fixed term basis. So there's no guarantee that it'll be renewed at those generous levels. So for their use case, and given their sensitivity around how proprietary all this stuff was, building a second data center economically made sense. I don't normally come up with that assessment. But I'm here to help customers come up with the right answer. Not try and push in any particular direction.

[00:18:57] SP: Yeah, I think there's some truth to that as well, right? Customers, eventually, if you don’t have the right price performance economics across compute, storage, networking, and achieve the rate of performance they want the unit of work to complete, they will choose to go business as usual. And so when we looked at this problem, we said we have to solve this across the stack and have that stock be readily available to have the work as complete, to be completed. And then the data that has to flow out of it also has to be super cheap. So I think that combination does make sense. Mazda, Nissan, Siemens, they're all using Oracle Cloud today amongst the many others for HPC style workloads. And that's one category that I read as we talked about. And then the other category we touched a little bit on was VMware.

But the thing I didn't mention, and I think it's worth a few minutes of the chat for sure, is what do you do for all your on-premises workloads that are sitting on blade servers today, or commodity servers, or a specialized servers today?

[00:19:57] CQ: Generally, I will wind up pulling the fire alarm that winds up dumping all the foam out of the **[inaudible 00:20:01]** servers unworkable **[inaudible 00:20:02]** written off at a total loss by insurance and start over somewhere else. But that's why they don't let me in the data centers anymore.

[00:20:08] SP: That's probably not a bad reason if you were going to do all of that at a customer that monitors a billion dollars' worth of transactions every day. But I think the other thing that we were looking at solving a tough problem was customers training in a single rack of a blade server, for example, for a single rack of a cloud server does not necessarily get them along their way of transforming the data center. So data center transformation is happening across the industry, whether customers are shrinking the data centers, moving it to colos, moving at all to a public cloud is going and happening in one of these various types of shapes.

And if you look at this market space and you say, "Well, how can we help 80% of that workload get to the cloud quickly?" And part, to be honest, this is where Oracle actually has a lot of business, our exadata platform, the combination of hardware and software today that our exadata provides us, provides some very significant performance gains for our customers. And we were looking at all that customer base and looking at this problem of 80% of workload. This is the other tough engineering problem that I think Oracle has solved for, which is how do you get regions to scale down to a very smaller size so you can start small and then linearly scale them out? I think the notion of always having multiple availability zones to be part of a region deployment has been challenged in the market. Like even in the market you'll see local zones, or smaller footprints, and the engineering effort to get all your services into a package format, if you would, that can help customers really transform and change your data center strategy is the area that I'm focused on.

Frankly, the reason is, you asked me earlier, why did I join? Because I saw some of the engineering behind what it means to shrink data centers without having to forego capabilities like high-availability, and durability, etc. And that's another area where I see a lot of customers'

interest in. They want to make the right choice long-term for their on-premise as a state. And Oracle has, I think, a set of products that can enable them to do that. And so from an engineering perspective, there's a whole bunch of efforts made on thinking about power distribution, and placement, anti-affinity, affinity mapping of compute instances. And a whole bunch of that I can go into more detail of.

[00:22:24] CQ: One of my problems has always been that while I understand the engineering reasons behind doing a lot of these things, the painful part is that I have to sit here and figure out on some visceral level what it's going to take to wind up spreading this thing across multiple availability zones from day one. I'm incurring data transfer charges between those availability zones in that model. And it's still within the same disaster radius for a whole bunch of disasters that my insurer cares about, that my compliance folks care about. If I have to have multiple copies of this thing running, I'd sooner put it multi-region, not multi-AZ. That's always been my starting philosophy. And I understand there are workloads where this makes zero sense and can't be done. I am sympathetic to that. But what you're saying does resonate with my own experiences.

[00:23:07] SP: I think it's class of workloads, agreed. Like there's some class of workloads you probably want to have multi region so that you have full DR and full back-up until your compliance and insurance adaptations are met. Totally got it. And then there's high availability, which is not necessarily same thing as DR or a back-up for which there are these multiple zones. The argument that I was making was that customers on-premises data today is served in a single type of zone infrastructure today, and they rack and stack a whole bunch of hardware, software, and continue to meet the latency performance characteristics to serve their customers. And when they think about breaking that thing down into moving into the cloud, the challenge is what portion do you move to the cloud, the public cloud, without fundamentally breaking the properties of the system, right? Their dependencies, their multiple downstream and upstream dependencies are all put together into a software package. Now, we can argue that's like, "Why did you do that?" But those are the tools of the trade back 10, 20 years from now. There weren't services. There wasn't the notion of cloud.

And so one of the things to think, as we were thinking about this problem, as I mentioned over dedicated regions, was how do you bring the entire cloud experience to customers so that they

can easily move and without having to worry about the networking costs, networking performance between these two different islands in their data center, modernize it all and then have an escape hatch to wherever they want to take it once it all has been modernized? I see that as being an emerging trend where customers want to move to the cloud economics and cloud agility. But so much of their data – And being abode, I have heard so much of regulation, so much of latency prevents them from doing that. And I think what we have done there on our own engineering perspective where we shipped proper region into a smaller footprint and bring all the services along with it. I think it's a pretty interesting engineering problem that Oracle solved.

[00:25:05] CQ: I would agree with you. I think that there's also, I guess, a criticism that I would be remiss if I didn't bring up, which is Charles Fitzgerald, who runs Platformonomics, has been doing his whole series for multiple years now on Follow the CAPEx, specifically parsing through a bunch of the earnings calls for all the cloud providers you care to name and several we don't, and figuring out, “Okay, who's serious about this versus who's not?” based upon a somewhat opinionated analysis of their capital expenditures that they're incurring as a function of cloud investment? And his working thesis is that if you're throwing serious money into cloud, you're serious. But if you're not making heavy capital investments, then you're clearly not. And I have an opinion on that. But I'm curious to hear what your point on that is, given that Oracle is not spending anywhere near the multiple tens of billions on capital expenditure in a quarter that AWS, Azure and GCP are.

[00:26:02] SP: So I think, I don't know the exact specific on the financials of what we spend and we don't.

[00:26:06] CQ: And I would not ask you to disclose them should you know them internally, to be very clear. But it's hard to easily hide tens of billions of dollars of spend in quarters and not wind up getting led out of the building in handcuffs right around the third time you've done it. The SEC does not play around.

[00:26:20] SP: Fair enough. Okay. So I'll say I think I know at least one portion convincingly that I can talk about on this front, which is we wanted to get close to customers where they were and be able to do that in a way that we can earn their business and then scale linearly. So Oracle

has a large enterprise customer base. We're very fortunate to have that. We've served them for many years. And we wanted to get close to their workloads first. It's not a surprise. I think anyone who knows Oracle Cloud, we will say that Oracle workloads running on Oracle Cloud provides the best price performance economics than anywhere else.

And so when we think about capital expenditure, there's a bucket of capital expenditure for all types of cloud consumers. So all of the investments in the US and Europe that we're building in multiple data centers and quickly growing the data center footprint is trying to manage the heterogeneity of our cloud. But in some areas, for example, in Oman, when we did a dedicated region, we're like, "Okay, we have to spend a certain amount of capex to get very specific workloads." We're in a different cycle of our business, right? And as a result of that, we have to choose where exactly we're going to double down on our investment. So we really have focused ourselves on fixing an engineering problem that says we can start small and linearly scale out. And as soon as we see the sign that we need to linearly scale out, we do that very rapidly.

And so the flip side is, "Oh, actually, you should have the exact same architecture, the exact same capital expenditure to give customers the price performance economics that you need over time." But we're already doing that. So the question is, what are we going to gain from that additional capital expenditure with the exception of, "Hey, we have this complete illusion of infinite capacity." Are we actually serving, let's say, 3000 customers in a locality, and we know exactly the workloads we want, and we know what the linear scale-up model looks like.

So five years ago, when they thought about region design, they took very deliberate choices. And this is what I think makes the capex story interesting. Because once you have ability to sort of shrinkwrap to a certain amount, of course, at some bare basic capacity you need. And if you're seeing the demand pickup, can you incrementally do more? I think that, for a fact, Oracle can do. When we start growing and scaling our business, we continue to add incremental capacity versus having to spend tens of billions of dollars upfront. And there's some use cases. For example, as a result, like we don't have a full edge story yet, right? And so I'm sure there's a lot of capex on that front as well. And so we're trying to make sure we're doing a judicious investment cycle as customers. And to be honest, the first customer trust that we are going after is the Oracle customer base, the Oracle enterprise customer. We have been fortunate that the basic primitives that we have put in place for networking and compute pricing and simplicity has

led to additional wins. And that's just the fruit of good engineering, like the HPC and the VMware workload, the legacy migration workload and whatnot. Hopefully that gives you some color to how we think about this investment cycle.

[00:29:23] CQ: One area that is near and dear to you that I find fascinating that no other provider can touch is the idea of Cloud@Customer @Cloud, @Cloud, @Cloud. Sorry, I got stuck in that loop again, or dedicated regions, or whatever it is you want to call it. But fundamentally, you roll out a full suite of every service Oracle Cloud offers on-premises at a customer. Talk to me about that, because that is unheard of. We have Azure Stack, which does some subset. We have AWS Outpost, which can only be described as begrudging. We have GCP doing whatever it's doing with Anthos. It's such a boring name. I have to look it up every time, and it's not clear at all where it starts or where it stops. But Oracle Cloud remains the only cloud service that I've seen that effectively takes over a customer data center. Tell me about that.

[00:30:08] SP: Yeah, so technically, we don't take over a customer's data center, right. What we do is we say, “Give us a little bit of space, a little bit of power. And we'll land an entire region on your premises.”

[00:30:18] CQ: When you say a little bit of space and a little bit of power, let's begin there, what are those requirements?

[00:30:23] SP: About 2000 square feet of contiguous block less space, and about a half a megawatt of power.

[00:30:29] CQ: That's not small by my spare room standards here. But for a data center, that is easily doable for any data center that isn't completely bursting at the gills already.

[00:30:37] SP: So what we found when we did this initial search to target markets, about 10% to 15%, of basically what the space would look like. And that's where we were focusing our efforts on. There's some room for us to grow into the region. So this is why we asked for a little bit more space and a little bit more power. We're also doing some engineering work to further to reduce our ask here. But that's the shape of what we're asking customers.

So the idea here is – And frankly, this journey started three, five years ago. That said, if we were to go land to the whole capex conversation kind of ties nicely into this. Like if we have to go land in Tokyo in the next 30 days, what does it take for us to land a region in Tokyo 30 to 60 days? I mean, that's the type of problem we have to go solve for. Or if we have to go land in Brazil or with land and go some other remote location, where we have, of course, think about logistics and a whole bunch of other legal stuff. But from an engineering perspective, can we start with a footprint, support 70 of our services on day one, both control plane and data plane, with a certain amount of capacity that can linearly scale? Can we do that? And that led to a whole bunch of interesting engineering choices on upstream power and downstream power.

But fundamentally, what it resulted in was we could actually take that footprint and stamp out more regions and starting small and then linearly scaled them as demand grows. That entire concept is being brought to the customers. Now, over time they've gotten – Oracle has got an even smarter about what it means to do the right deployment strategies. How do we make sure we do fairness algorithms for our services teams? What does it mean to not have service teams request above and beyond capacity they don't need based on signals we get from customers. So there's a whole bunch of engineering work that led to can we do scaling out of regions at anywhere in the globe at 30 to 60 days and start small and to building it? That whole engineering is coming to customers' on-premises?

And so why do customers need that engineering on-premises? What do they need a region on-premises? So I also look at our exadata Cloud@Customer business, which is just a special purpose database offering, which runs the Oracle database on a high-performance software and hardware stack. So we have that offering as well. But what we heard from customers was, "What do we do about our middleware that we have to modernize? What do we do about the application that calls into the database? All of that has to be modernized. All that needs a space to live in. So that eventually, if we are no longer dealing with data regulation or no longer dealing with a latency problem, we have modernized, then we eventually have an escape hatch to where we need to go to.

But this product really is in the making for three and a half, four years. Now we've gotten to a point we can say, "Okay, now we can ship it into your data center. We've figured out how to do

that part too." But it brings with it the entire gambit of Oracle 70 plus services, including our developer services, our core compute network storage services, our AI, data science services, etc. That's how we do it fundamentally.

[00:33:29] CQ: One of the things that I get flack for sometimes is when I talk about Oracle Cloud, everyone gets up in arms and yells at me. They call me a sellout. They call me a shill, etc., which I understand that. I get accused of that all the time for my coverage of AWS. I get accused of being too mean towards AWS, as well as being a sellout to AWS in response to the exact same thing. So apparently, the Internet doesn't know what it wants, and that's fine. My position has always been that you can buy my attention. You cannot buy my opinion. And I will say that Oracle Cloud is technically excellent. I've spun things up there. It worked as expected. The only onboarding problem I had a few years ago when I started down this path was there was some challenge with getting out of fraud control, because I was doing this mid-air. So it was through some aggregator in Chicago, I think it was, from an IP perspective. I'm claiming to be somewhere else. And it probably looked fairly suspicious. I can't really fault Oracle Cloud for that. But since then, they did $1 temp-off that they then were refunded. And that was the last time they charged me for anything. You're always free tier is always free, as it says on the tin, and should be. There are no surprise bills. I read through, because this is Oracle. Let's not kid ourselves here. I read through the terms and conditions carefully. You aren't asserting ownership of anything I build. It's very reasonable, very standard. There's nothing lurking there that I was able to see that made this a terrible idea. And I still get periodic email updates on the instance I left running telling me that there's been a maintenance event and you folks are. Here's what happened. So rather than leaving it to me to play guess and check.

Further, a lot of the things that you did with the Iron.io folks acquisition to build out your serverless functions, it was in many ways superior to what AWS Lambda had been doing. There was an awful lot of really neat stuff going in there. The single big drawback that I argue is kept people away is the Oracle at the beginning of all of that. I mean, you can only put so many – Going through the word cloud in so many times. But that still winds up having the Oracle word coming first. I'm not asking you necessarily to dunk on your employer that seems career-limiting in most places. But how do you square that circle?

[00:35:37] SP: I've been cut from the cloud that focus is just on your customers. So I continue to just focus on how we're going to serve our customers. And if a company like Oracle is making a very deliberate choice in serving their customers differently and they want people who have done this part of the business before and want to both change culturally inwards and outwards, then I think that's the change that we're impacting. So we've got to be humble and sort of have the humility in front of customers saying, “We have a platform. We have business practices, and we have a support structure that is here to serve you, and serve you as adequately as you can serve in the past by any other cloud provider. And just give us the right direction with that business.” So yes, there is probably some of that over-high part, because the business model has been different in the past, I suppose. It's just that I don't know that past. All I know is what we're trying to build towards the future.

And customers, when they look at the service that they can get from Oracle, usually give us a second take. Frankly, that's what I'm looking for, which is I need the second take for me to describe why we can provide value to you. And do that in clear transparent terms repeatedly. So apart from just technology, if you look at what we've done recently with our cloud lift services, we're paying for customers’ migrations. Meaning my team goes in and tries to migrate customers’ workloads on our dime and behalf, because we believe strongly that the platform will be where you want to be. So that's our cost we incurred. We have cloud net services as part of our program we’ve launched two, three months ago where we said to customers, “Bring your workloads.”

So I think earning trust is a very long term thing. You just got to chip away at it and prove out case by case on how we can serve them well. I don't have a magic bullet for it. But I think customers are taking notice and giving us even more of their business to us. And it's about keeping their trust now and making them successful towards deployment, which is a good portion of my focus and the entire organization's focus at this point.

[00:37:33] CQ: The idea that effectively every other provider has in the space, “Oh, we’ll charge you at a discounted rate to help you migrate in.” I don't mean to be unkind here, but I've never understood the thinking behind that strategy. Because once someone starts spending money on your cloud service, they are going to continue to spend money on your cloud service. Spoiler, I've been down this road with a bunch of different customers, and I've seen how this goes. You

should be paying customers to let you migrate things on. Covering the cost is one of the least difficult steps that companies could do. But somehow you're the only ones I've seen doing this in the wild. I further do want to point out that I've spoken to a variety of Oracle Cloud customers.

Now, historically, most of these have been big e-enterprises, very large companies, the blue chips that you would expect. But I'm starting to see you move into the startup space a bit more. And one thing I have not heard through any of those conversations is a complaint about Oracle Cloud, either a technical shortcoming or a business practice complaint. People will complain about Oracle, the overall company, without any provocation. And let's be honest, a lot of this is very well deserved. But Oracle Cloud only gets grief that is reflected.

What I start seeing that I think is, I guess, a real sign of maturity and growth in the platform, is there’re starting to be stories around Oracle Cloud that don't feature the word databases or, heaven forbid, autonomous databases, which is Oracle's database division’s whole big thing these days. It starts to stand on its own legs and its own merits rather than depending upon its giant install base in existing enterprises as its primary means for marketing and selling. Are you seeing an uptick in the startup world of Oracle Cloud these days?

[00:39:15] SP: I'm not super close to our startup ecosystem, to be really honest. So substantiate any claim that I make here. I do know that a lot of startups that are born in the cloud that use Oracle Cloud have shared the same sentiment. So to the extent I can give you a trend line or some sort of an indication that startups are now a big area of on focus for Oracle, or they're coming to our platform organically. I just don't have the data points on that one. But I do know ISVs, whether born in the cloud, new or ones that are migrating, they are. So that’s a good proxy, I think, because probably everyone is trying to develop a software stack and they need a place to put it both from that price performance perspective from our business practices as well. But I think the overall sentiment is changing, and not only because we have autonomous database, or not only because we have an Oracle database product associated with the cloud. I think that is true. And that's a fact. Although I've come to really appreciate the word autonomous. I don't know what changed my thinking on that one. But I'll leave that for a different debate. But I think you're right, in part because we've been very deliberate about it, which is Oracle Cloud is more than just database. Oracle Cloud is about cloud technologies. Things that you come to expect from other cloud providers if you want to develop new software

or bring over existing software. And the proof really right now that I see in sort of big bold letters is our independent software vendors and ISVs that are coming onto the platform rapidly.

[00:40:41] CQ: I would like to point out that I didn't give you a list of questions I was going to be asking you before we did this, the best guidance I gave you. Let's hope none of us gets fired by the end of this. And just now is a perfect example of the sort of thing I'm talking about. You could have said, “Oh, yeah. We're seeing a whole bunch of startup stuff that I've heard of clients taking stuff thinking of adopting our service. I can't talk about any offhand, of course. But, oh, yeah, it's great.” This ties into, in many ways, an echo of what I've seen at your previous employer. I have never yet caught someone at AWS or at Oracle Cloud lying to me, even when right now you could have come out and said, “Oh, yeah, we're seeing whatever it is will sound good.” You're intrinsically honest. And I appreciate that about not just you, but it seems to be a cultural ethos I'm seeing throughout my conversations with folks working with Oracle Cloud. This is the sort of thing that makes me tell people, “Yeah, I know. It sounds like I'm having a joke at your expense.” Just suspend disbelief and have a conversation with them. And so far, it seems to be going relatively well. One day, I'm hoping I won't have to wind up giving that disclaimer first. But that's going to take a few more years, I suspect.

[00:41:45] SP: Yeah, I don't know the timeline on that. I'm in a smaller bubble, because I'm in the Oracle Cloud infrastructure piece. And I know customers – And in fact, I think to be honest, our SaaS teams, the Oracle fusion products, are doing really well with customers as well. And we're doing a whole bunch of work with our SaaS platform teams who run on Oracle Cloud and make it a canonical platform through which to deliver the SaaS products for improved latency, improved user experience and whatnot. I don't think even if I was in AWS, I would say that we are ever going to be happy without the experience that we offer to customers. We just have to continually in this environment, and especially the infrastructure, I just have to be discontent and strive for what we really can do better.

So part of my reason not to give a glowing scorecard to Oracle or to anybody is just general discontent with where we are and what we can do to be better. And there's so much we can do to be even better for our customers. So I think if I **[inaudible 00:42:39]** if I was back in AWS, I would say that about – If you ask me like, “How are startups adopting AWS?” I would have probably said I wish we could do better. But nonetheless, I think customers are taking notice.

Simplicity, focus on really tough engineering problems, and transparency in how we do business with them both at the presale side, the for-sale side and support. I think the combination of this, I might hope, is that customers will continue to give us more of their business.

[00:43:09] CQ: I expect that they will. I think that you've done a lot of the right things. I think that what you're doing is heading in the right direction. I also want to thank you for taking the time to sit here and suffer my slings and arrows. If people want to learn more about Oracle Cloud and make up their minds for themselves, where should they go to start?

[00:43:23] SP: They should go to Oracle.com. On the page, they should get an easy link to get to our Oracle Cloud infrastructure and get started. Deploy our first application of VM. Setup a virtual cloud network, whatever have you. So go to Oracle.com, on the main page, and variable you'll see Oracle Cloud infrastructure. Just click on it and it will take you to some canned parts on what problem you want to go solve with cloud infrastructure and hopefully takes you getting to where you need to be pretty quickly.

[00:43:49] CQ: Thank you very much for taking the time. I appreciate it.

[00:43:52] SP: You're most welcome. Thanks for having me, Corey. I appreciate the chat.

[00:43:55] CQ: Salman Paracha, Oracle's group vice president, cloud engineering, Cloud@Customer. This concludes my guest hosting of Software Engineering Daily, but the tour of the cloud isn't over quite yet. Subscribe to Screaming in the Cloud on your podcast platform of choice to hear the final episode of the cloud tour saga. And of course, don't forget to go to lastweekinaws.com to subscribe the Last Week in AWS newsletter. And follow me on Twitter @quinnypig for all my snarky takes.

[END]