

## Transcript of Prasad Sankaran & Larry Socher, Accenture | Accenture Cloud Innovation Day 2019, 9/12/19

YouTube link: [https://www.youtube.com/watch?v=\\_fTrTYQKDoo](https://www.youtube.com/watch?v=_fTrTYQKDoo)



(intro music)

>> Hey, welcome back everybody, Jeff Frick here from theCUBE. We are high atop San Francisco in the Accenture Innovation Hub. It's in the middle of the Salesforce Tower. It's a beautiful facility. I think it had the grand opening about six months ago, we were here for the grand opening. Very cool space, they got maker studios, they've got all kinds of crazy stuff going on. But we're here today to talk about cloud and this continuing evolution about cloud in the enterprise and hybrid cloud and multi cloud and public cloud and private cloud and we're really excited to have a couple guys who are really helping customers make this journey, 'cause it's really tough to do by yourself. CIOs are super busy, they're worried about security and all kinds of other things. So Accenture's often the trusted partner. We've got two of the leaders from Accenture joining us today, Prasad Sankaran, he's the Senior Managing Director of Intelligent Cloud Infrastructure for Accenture, welcome. And Larry Socher, the Global Managing Director of Intelligent Cloud and Infrastructure offering from Accenture, gentleman welcome.

>> Hiya Jeff.

>> I love it, "Intelligent Cloud." What is an intelligent cloud all about? You've got it in your title, it must mean something pretty significant.

>> Yeah, I think Jeff, first of all, thank you for having us, but yeah, absolutely, everything's around, becoming more intelligent around using more automation and the work we deliver to our clients and cloud as you know is the platform to which all of our clients are moving, and so it's all about bringing the intelligence not only into infrastructure, but also into cloud generally and it's all driven by software.

>> Right, it's just funny to think of where we are in this journey. We talked a little before we turned the cameras on, and Larry, you made an interesting comment when I said, you know, "When did this cloud "for the enterprise start?" And you took it back to SaaS-based applications, which

>> Yeah we're sitting in the Salesforce building

>> That's true. it's the tallest building in San Francisco, right?

>> You know, everyone's got a lot of focus on AWS's rise, et cetera, but the real start was really getting into SaaS. I mean, I remember we used to do a lot of Siebel deployments for CRM and we started to pivot to Salesforce, we were moving from Remedy into ServiceNow, and we went through on-premise collaboration email to 0365, so we've actually been at it for quite awhile, in that, you know particularly in the SaaS world, and it's only more recently that we're starting to see, the kind of the push to the you know, public paths, and it's starting to cloud native development, but this journey started, you know, what was that, seven, eight years ago that we really started to see some scale around it.

>> And tell me if you agree, really what the Salesforces of the world and the ServiceNows of the world and Ops 365 kind of broke down some of those initial barriers, which were all really about security and security, security, security, it's always to here. Where now, security is actually probably an attribute. That the cloud can bring.

>> Absolutely. In fact, I mean those barriers took years to break down. I still saw clients where they were forcing Salesforce or ServiceNow to put you know, instances on param, and I think you know, they finally woke up to you know, these guys invest a ton in their security organizations, you know, there's a little bit of that needle in the haystack you know, if you breach a data set you know what you're getting after, but when you're up into Salesforce it's a lot harder, so you know, I think the security problems have certainly gone away. We still have some compliance regulatory things, data sovereignty, but I think security, not that it's solved by any means, you know, it's always going to be an ongoing problem, but I think they're getting more comfortable with their data being up in the public domain.

>> Right. >> Well, not public.

>> I think it also, helped them with their progress towards getting cloud native. So you know, you pick certain applications which were obviously hosted by Salesforce and other companies, and you did some level of custom development around it, and now I think that's paved the way for more complex applications and different workloads now going into the public cloud and the private cloud, but that's the next part of the journey.

>> Right. So let's back up a half a step, because as you said, a bunch of stuff then went into public cloud, right? Everyone's putting in AWS and Google, IBM's got a public cloud, there was a lot more, there's not quite so many as there used to be. But then we ran into a whole new host of issues, right? Which has kind of opened up this hybrid cloud, this multi-cloud world, which is, you just can't put everything into a public cloud. There's certain attributes you need to think about from the application point of view before you decide where you deploy that. So I'm just curious if you can share what you guys do with clients. How should they think about applications, how should they think about what to deploy where?

>> Yeah, I think, I'll start, and then Larry

>> Sure.

>> Has a lot of experience in this area. I think, you know, we have to obviously start from an application-centric perspective. You've got to take a look at, you know where your applications have to live, what are some of the data implications around the applications? What do you have by way of regulatory and compliance issues? What do you have to do as far as performance? Because certain applications have to be in a high-performance environment certain other applications don't, so I think a lot of these factors will then drive where these applications need to reside. And then what we're seeing in today's world is really a complex situation where you have a lot of legacy, but you also have private as well as public cloud. So you approach it from an application perspective.

>> Yeah, I mean if you really take a look at our, I mean, look at Accenture's clients. I mean we're typically focused on the upper end of the market, Global 2000, so very, and our clients typically have application portfolios ranging from 500 to 20,000 applications. And really, if you think about the purpose of cloud or even infrastructure for that, they're there to serve the applications. No one cares if your cloud or infrastructure's not performing if the apps aren't, so we start off with an application modernization approach, and ultimately looking, you know, with our tech advisory guys coming in, our intelligent engineering services who do the cloud native and app mod work, our platforms guys who do, you know, everything from Salesforce through

SAP. They should drive a strategy on how those applications are going to evolve, whether it's 500 or 20,000 and determine, hey, and usually using something like the six Rs methodology, you know, am I going to retire this? Am I going to retain it? You know, am I going to replace it with SaaS? Am I going to refactor and format? And it's ultimately that strategy that's really going to dictate a multi and hybrid cloud story. So it's based on the applications, data gravity issues, where they're going to reside, the requirements around regulatory, the requirements for performance, et cetera, that will then dictate the cloud strategy. So I'm, you know, not a big fan of going in there and just doing a, you know, multi-hybrid cloud strategy without a really good upfront application portfolio approach.

>> Right.

>> How are we going to modernize that?

>> And how do you segment? That's a lot of applications, and you know, how, the old thing, how do you know that one byte at a time, right? How do you help them prioritize where they should be focusing on?

>> Yeah, so typically what we do is work with our clients to do a full application portfolio analysis, and then we're able to then segment the applications based on, you know, importance to the business, and some of the factors that both of us mentioned. And once we have that then we come up with an approach where certain sets of applications you move to SaaS, certain other applications you move to PaaS, so you know, you're basically doing the refactoring and the modernization. And then certain others, you know, you can just, you know, lift and shift, so it's really a combination of both modernization as well as migration. It's a combination of that, but to do that, you have to initially look at the entire set of applications and come up with that approach.

>> I'm just curious where within that application assessment where is cost savings, where is, "This is just old" and where is opportunities to innovate faster, 'cause as we know, a lot of talk early days is cost savings, but what the real advantage is, is execution speed, if you can get it tuned in.

>> You could go back three or four years and we had, there was a lot of CIO discussions around cost savings. I really have seen our clients shift. Cost never goes away, obviously, right? But there's a lot greater emphasis now on business agility, you know? How do I innovate faster, get new capabilities to market faster, to change my customer experience? So it's really, IT's really trying to step up and you know, enable the business to compete in the marketplace. So we've seen a huge shift in emphasis or focus at least, starting with, "How do I get better business agility, "how do I leverage the cloud and cloud native development "to get there?" You know, service levels, we're starting to see an increase on, "Hey, these applications need to work," et cetera. So obviously cost still remains a factor, but we're seeing a much more, you know, much more emphasis on agility, you know, enabling the business, you know, giving the right service levels, the right experience to the users or the customers, big pivot there.

>> Okay, and let's just get the definitions out, because you know, a lot of conversation about public cloud's easy, private cloud's easy, but hybrid cloud and multi cloud, and confusion about what those are, how do you guys define 'em? How do you help your customers think about the definitions?

>> Yeah, I think that's a really good point. So what we're starting to see is there were a lot of different definitions out there, but I think as I talk with more clients and our partners, I think we're all starting to come to a, you know, same kind of definition. On multi cloud, it's really about using more than one cloud. But hybrid I think is a very important concept because hybrid is really all

about the placement of the workload, or where the application is going to run. And then again, of course to all of these points that we talked about, data gravity, performance and other things, other factors, but it's really all about where do you place the specific workload?

>> Yeah, and if you look at that, think about public, I mean, it's obviously gives us the innovation of the public providers, you look at how fast Amazon comes out with new versions of Lambda et cetera, so it's the innovations there, obviously agility can spin up environments really quickly, which is, you know one of the big benefits of consumption economic models, so there's a number of drivers that are pushing in the direction of public. You know, on the private side, there's still quite a few benefits that don't get talked about as much. So you know, if you're lookin' at performance, you know, the public world, you know, although they're scaling up larger t-shirts et cetera, they're still trying to do that for a large array of applications. On the private side, you can really tailor something to very high performance characteristics, whether it's 32-, 64- you know terabyte HANA, I mean you can get much more focused precision environment for business critical workloads like that. You know, Oracle, Oracle Rack, you know. The dupe clusters if you think about fraud analysis. So that's a big part of it. Related to that is the data gravity that Prasad just mentioned, you know. If I've got a 64-terabyte HANA database, you know, sitting in my private cloud, it may not be that convenient to go and put, you know get that data shared up in RedShift or in Google TensorFlows, so there's some data gravity on networks just aren't there, the latency of moving that stuff around is a big issue. And then a lot of people have investments in their data centers, I mean, the other piece that's interesting is legacy. You know, you know as we start to look at a world, there's a lot, there's a ton of code still living, you know, whether it's Unix systems, it's IBM mainframes, there's a lot of business value there, and a lot of times the business cases aren't necessarily there to replace them.

>> Right.

>> And in a world of digital decoupling, where I can start to use microservices, we're seeing a lot of trends, we work with one hotel to take their reservation system, you know, wrap it in microservices, we then did a, you know, open shift couch-based front-end, and now when you go against, you know, when you go and browse in properties, you're looking at rates, you're actually going into a distributed database cache, you know, and using the latest, you know, cloud-native technologies that can be dropped every two weeks, or every three or four days for a mobile application. And it's only when it goes, you know, when the transaction goes back to reserve the room that it goes back there. So we're seeing a lot of power with digital decoupling, but we still need to take advantage, you know, we've got these legacy applications. So the data centers are really, we're trying to evolve them and really just, you know, how do we learn everything from the world of public and start to bring those similar type efficiencies to the world of private? And really, what we're seeing is this emerging approach where I can, you know, start to take advantage of the innovation cycles, the Lambdas, the RedShifts, the Azure functions of the public world, but then maybe keep some of my more business-critical regulated workloads, you know? That's the other side of the private side, you know?

>> Right.

>> If I've got GXP compliance, if I've got HIPAA data that I need to worry about, GDPR, this is all, you know, the whole set of regulatory requirements. Now, over time, we do anticipate the public guys will get much better and more compliant. In fact, they made great headway already. But there's still a number of our clients, they're still, you know, not 100% comfortable from a compliance perspective.

>> Right. They got to meet Theresa Carlson, she'll change 'em. (laughs) Who runs AWS public sector business and it's doing amazing things, obviously, with big-government contracts.

>> Yeah.

>> But you raise a really interesting point, Larry. You almost describe what I would say is really a hybrid application in this hotel example that you use, because it's you know, breaking the application and leveraging microservices to do things around the core that allow you to take advantage of some of this agility and hyper-fast development, yet still maintain that core stuff that either doesn't need to move, works fine, be too expensive to refactor.

>> Yeah.

>> That's a real different way to even think about workloads and applications and breaking those things into bits.

>> And we see that pattern all over the place. I mean, I gave you the hotel example, where, but you know, look at financial services, retail banking, so open banking, allow all those retail applications are on the mainframe. I mean, insurance claims, and you look at it, the business value of replicating you know a lot of the regulatory stuff, the locality stuff, that doesn't make sense to write it, there's no real inherent business value, so if I can wrap it, expose it in, you know, a microservices architecture now do a cloud-native front end that's going to give me a 360 view of a customer, change the customer experience, you know, I've got a much, you know, I can still get the agility, the innovation cycles by public, by wrapping that legacy environment.

>> And Prasad, you're ready to jump in, I'll give you something to react to, which is the single plane of glass, right? Now how do I manage all of this stuff? Not only do I have distributed infrastructure, now I've got distributed applications and the thing that you just described, and everyone wants to be that single pane of glass, everyone wants to be that app that's on everybody's screen. How are you seeing people deal with the management complexity of these kind of distributed infrastructures, if you will?

>> Yeah, I think that's an area that's actually very topical these days, because you know, you're starting to see more and more workloads go to private cloud and so you bought a hybrid infrastructure. You're starting to see movement from just using VMs to you know, containers and Kubernetes and you know, we talked about serverless and so on. So all of our clients are looking for a way and you have different types of users as well, you have developers, you have data scientists, you have operators and so on, so they're all looking for that control plane that allows them access and a view to everything that is out there that is being used in the enterprise, and that's where I think, you know a company like Accenture, we're able to use the best of breed to provide that visibility to our clients.

>> Right, yeah, I mean you hit the nail on the head. It's becoming, you know, with all the promise to cloud, and all the power in these new architectures, it's becoming much more dynamic and ephemeral with containers and Kubernetes, with serverless computing, that one application for the hotel, they're actually starting, they've actually got some now running in AWS of their containers and looking at serverless, so even a single application can span that, and one of the things we've seen is first you know, a lot of our clients used to look at you know, application management, you know different from their infrastructure, and the lines are now getting very blurry. You need to have very tight alignment. So you take that single application, you know, if any, my public side goes down, or my mid-tier with my, you know, open shift on VM ware goes down, or my back-end mainframe goes down, or the networks that connect it go down, the devices that talk to, it's a very, well, despite the power, it's a very complex environment, so what we've been doing is, first we've been looking at, how do we get better synergy across what, you know, our application services teams that do the application management optimization. Our cloud and infrastructure, you know, how do we get better

alignment with our embedded security, you know, with our managed security services, bringing those together? And then what we did was we looked at, you know, we got very aggressive with our cloud first strategy, and you know, how do we manage the world of public, but when looking at the public providers, the hyper-scalers and how they're hitting incredible degrees of automation, we really looked and said, "Hey look, you've got to "operate differently in this new world, what can we "learn from how the public guys are doing that?" And we came up with this concept we call "Run different." You know, how do you operate differently in this new multi-speed, you know, very hybrid world across public, private, and even legacy environments? And start to look and say, "Okay, what is it that they do?" you know, first they standardize. And that's one of the big challenges, you know, I go in to almost all of our clients, and there's this sprawl, and whether it's application sprawl it's infrastructure sprawl,

>> "My business is so unique though, Larry."

>> Yeah, yeah (laughs)

>> "No other business out there "has the same processes that I have."

>> So we started to make, you know, how do we standardize our Accenture hybrid cloud solution you know, partnered with HP and IBM, where we, you know, how do we, that was an example, so we can get to, you, 'cause you can't automate 'em unless you standardize. So that was the first thing, you know, standardizing our service catalog, standardizing that. You know, the next thing is the operating model. They obviously operate differently. So we'd be putting a lot of time and energy in what I call a cloud and agile operating model, and also a big part of that is truly, you know, you hear a lot about dev ops right now, but truly putting the security and operations into devsec ops, the bringing, you know, the development and the operations much tighter together so spending a lot of time looking at that and transforming operations. Rescaling the people, you know, the operators of the future aren't eyes on glass, they're developers, they're writing the data ingestion and the analytic algorithms, you know, to do predictive operations. They're writing the automation script to take work, you know tasks work out, and you know, over time they'll be tuning the AI engines to really optimize the environment. And then finally as Prasad alluded to, is that the platforms, the control planes are doing that, so you know, we, what we've been doing, is we've had significant investments in our Accenture cloud platform, our infrastructure, automation platforms, and then the application teams with my wizard framework, and we're starting to bring that together, you know, an integrated control plane that can, that we can plug into our clients' environments to really manage seamlessly, you know, and provide, you know, automation, analytics, AI, across apps, cloud, infrastructure, and even security,

>> Right.

>> And that, you know, that really is AI ops, I mean that's delivering on, you know, as the industry starts to define and really coalesce around AI ops, that's what we view AI ops

>> So, just so I'm clear, it's really your layer, your software layer, kind of management layer, that integrates all these different systems and provides kind of a unified view, control, AI, reporting, et cetera?

>> Right. Exactly, and then can plug in and integrate, you know third-party tools, to do specific functions.

>> I'm just curious, 'cause one of the themes we hear out in the press right now, is this kind of pullback of public cloud apps. Some of 'em are coming back, where maybe it was, kind of a rush, a little bit too aggressively. What are some of the reasons why people are pulling stuff

back out of public cloud? Is it just the wrong application, the costs were not what we anticipated to be? We find, you know, what are some of the reasons that you see apps coming back in house?

>> Yeah, you know I think it's a variety of factors. I mean, certainly cost, I think is one. So you know, as there are multiple private options, and you know, we didn't talk about this, but the hyper-scalers themselves are coming out with their own, different private options like Anthos and Arkos and Azure Stack and, Alibaba has Up Stack and so on, and so you see a proliferation of that, then you see many more options around private cloud, so I think the cost is certainly a factor. The second is, data gravity is a very important point, because as you're starting to see how different applications have to work together, then that becomes a very important point. The third is just about compliance and you know the regulatory environment. As we look across the globe, you know, even outside the U.S., you know, we look at Europe and other parts of Asia, as clients are moving more to the cloud, you know, that becomes an important factor. So as you start to balance these things, I think you have to take a very application-centric view, you see some of those, some apps moving back, and I think that's the part of the hybrid world, is that you know, you can have an app running on the private cloud, and then tomorrow you can, you know, move this since it's been containerized to run on public. And you know, it's all managed at that level.

>> Yeah, I mean cost is a big factor. I mean, if you actually look at it, most of our clients, you know, they typically, your big cap ex businesses and all of a sudden they're using this consumption, you know, consumption model, and they weren't really, they didn't have a function to go and, you know, to look at the thousands or millions of lines of,

>> Right

>> You know, of Azure builds

>> Of the statement

>> Yeah, exactly. So I think they misjudged, you know, some of the scale and you know, I mean that's one of the reasons we started it's got to be an application-led modernization that really, that will really dictate that, and I think in many cases, people didn't, may not have thought through which application, what data's there, the data gravity, I mean data gravity's a conversation I'm having just about with every client now, you know, if I've got a 64-terabyte HANA and that's the core of my crown jewels of data, you know, how do I get that to TensorFlow, how do I get that,

>> Right. But if Andy was here, Andy would say, "We'll send down the snow, the snow," well I can't remember which version snowplow

>> Snowballs

>> Well there's snowballs, but I don't think you've seen [crosstalk] tractor trailer that comes down, and he'd say, "Take that and stick it in the cloud," because if you've got that data in a single source right now you can apply a multitude of applications across that thing. So, you know, they're pushing, "Get that data in this single source." Of course then, to move it, change it, you run into all these micro-lines of billing statements.

>> Well take the hotel, I mean their data's still in the mainframe, so if they may need to expose it, yeah they have a database cached and they move it out, so you know, particularly as the data sets get larger, it becomes, that data gravity becomes a big issue,

>> Right.

>> 'Cause no matter how much, you know, well Moore's law might be, might have elongated from 18 to 24 months, the network will always be the bottleneck, so ultimately we're seeing, you know, as we proliferate more and more data out, data sets get bigger and better, the network becomes more of a bottleneck on that, so a lot of times you got to look at your applications and say, "Hey, I've got some legacy database "I need to get to," you know, "I need this to be approximately, you know, somewhere "where I don't have, you know, "high bandwidth or high latency type connectivity"

>> Right.

>> Or so, egress costs, a pretty big deal. So my data's up in the cloud and then I'm going to get charged for pulling it off, you know, that's been a big issue.

>> That's funny, and I think a lot of the issue obviously, complexity billing is a totally different billing model, but I think too, a lot of people will put stuff in the public cloud and then operate it as if they bought it and they're running it in the data center in this kind of turn it on, turn it off when you need it. Everyone turns, everyone loves to talk about the example of turning it on when you need it, but nobody ever talks about turning it off when you don't, but to kind of close out our conversation, I want to talk about AI and Applied AI, 'cause you know, it's a lot of talk in the marketplace about AI and machine learning, but as you guys know probably better than anybody, it's the application of AI in specific applications which really unlocks the value, and as we're sitting here talking about this complexity, I can't help but think that, you know Applied AI in a management layer like your Run Differently setup to actually know when to turn things on, when to turn things off, when to move data, when not to move data. It's going to have to be machines runnin' that, right? 'Cause the data sets

>> Absolutely

>> and the complexity of these systems is just going to be overwhelming.

>> Yeah, yeah, absolutely, completely agree with you. In fact, at Accenture, we actually refer to this whole area as applied intelligence, and that's our AI, right? And it is absolutely to, you know, add more and more automation, more everything, to where it's being run by the machine rather than, you know, having people really working on these things.

>> Yeah, I mean, I think you hit the nail on the head. We're going to, AI's, I mean, given how things are getting complex, more ephemeral, you think about Kubernetes et cetera, we're going to have to leverage AI. Humans are not going to be able to get, you know, manage the environments going forward, and what's interesting about AI, we've used it quite effectively for quite some time, but it's good at some stuff, not good at others, so we've found it's very good at like, ticket triage. Like ticket triage, ticket routing, et cetera, you know any time we take over account, we tune our AI in, 'cause we have ticket advisors, etc, and that's where we got probably the most, you know, the most bang for the buck. We tried it in the network space less success to start, even with commercial products that were out there. I think where AI ultimately bails us out of this, is if you look at the problem, you know, a lot of times we talk about optimizing our own cost, but then performance, I mean, and it's somewhat, you got to weigh 'em off each other, so you've got a very multidimensional problem on, "How do I optimize my workloads?" Particularly when I got a Kubernetes cluster and some's in Amazon, you know, some's running on my private cloud, et cetera, so we're going to get some very complex environments, and the only way you're going to be able to optimize across multiple dimensions of cost, performance, and service levels, you know, and then multiple options to do it public, private, you know, what's my network costs? Etc, is an AI engine, you know tuning that AI engine. So ultimately, I mean and you heard me earlier on the operators, I think, you know, they



write the analytic algorithms, they do the automation scripts, but they're the ultimate ones who then tune the AI engines that will manage our environment.

>> Right.

>> And I think Kubernetes will be interesting, 'cause it becomes the link to the control plane to optimize workload placement, you know, between

>> Well and the best thing, too you then have dynamic optimization, 'cause you might be optimizing costs right now, but you might be optimizing for output the next day

>> Exactly

>> So it's really a, you know, a neverending matrix

>> When you got to see them together, which, and multi-dimension optimization is very difficult, but I mean, humans can't get their head around it.

>> Right.

>> Machines can, but they need to be trained, so.

>> Well, Prasad, Larry, lots of great opportunities for Accenture to bring that expertise to the table. So thanks for taking a few minutes to walk through some of these things.

>> Our pleasure.

>> Our pleasure thank you.

>> He's Prasad, he's Larry, I'm Jeff. You're watching theCUBE. We are high above San Francisco in the Salesforce tower at the Accenture Innovation Hub in San Francisco. Thanks for watching, we'll see you next time. (outro music)