Deloitte.

For Cloud Professionals
November 2019

For Cloud Professionals, part of the On Cloud Podcast

David Linthicum, Managing Director, Chief Cloud Strategy Officer, Deloitte Consulting LLP

Title: **Modernizing your architecture? Think hybrid and multi-cloud**

Description: As the panoply of cloud vendors and products grows ever larger, so does the configuration and management complexity of the cloud. For many companies, this complexity can reduce the value of their cloud-computing investments. In this episode of the podcast, David Linthicum and guests, Intel's Lisa Davis and Deloitte's Doug Bourgeois, discuss ways companies can tackle cloud complexity and leverage multi-cloud architectures to help optimize cloud deployments in terms of performance, cost, and agility. Lisa and Doug also argue that companies need to realize that hardware isn't a commodity and that cloud architectures—and talent acquisition decisions—must be driven by mission-critical business needs and workloads.

Duration: **00:27:17**

Operator:
The views, thoughts, and opinions expressed by speakers or guests on this podcast belong solely to them, and do not necessarily reflect those of the hosts, the moderators, or Deloitte. Welcome to On Cloud, the podcast for cloud professionals, where we break down the state of cloud computing today and how you can unleash the power of cloud for your enterprise. Now here is your host David Linthicum.

David Linthicum:
Welcome back to the On Cloud Podcast, your one place to find out how to make cloud computing work for your enterprise. This is an objective discussion with industry thought leaders who provide their own unique perspective around the pragmatic use of cloud-based technology. Today on the show we have Lisa Davis and Doug Bourgeois. Lisa is Vice President and General Manager for IT Modernization at Intel Corporation's Enterprise and Government Data Center Group, where she's responsible for growing Intel's data center business and working alongside enterprise and government CIOs to create IT modernization strategies for their organization. Doug Bourgeois is a managing director in the Federal Technology and Strategy [practice at Deloitte] and has

a long history in both areas. Lisa and Doug, thank you for joining me on the show, but let's get behind the scenes here. Doug, you and I met a long time ago at a VMware show way back when, and I've been following your career ever since, but give the audience perspective for how you got to Deloitte, what you typically focus on, and then, Lisa, we're going to dig big time into what you're doing right now.

Doug Bourgeois:
Cool. Thanks, David. I started in cloud a little over a decade ago. I was actually still in the federal government running a $2.5-billion-a-year shared services organization and stumbled across this thing called cloud and realized that it was an imperative to drive all of our back office applications and systems over to a cloud architecture to save money and improve agility for our clients. And from there, I went over to VMware, as you mentioned, to be their CTO for the public sector in the US, helping our clients adopt cloud based on the experience that I'd had at the time. For the last four years now, I've been with Deloitte as a managing director in our government practice, leading our cloud strategy and cloud migration offerings to our clients, helping them adopt, plan for, and transition over and adopt cloud models across their enterprises.

David Linthicum:
So, Lisa, love to get a perspective on what you're working on today, so what are you doing over at Intel, and how are you working with some of the CIOs that are out there? I'd love to get maybe even an example story of what you've run into in the last few years, specifically as government CIOs are taking on cloud.

Lisa Davis:
Yeah, absolutely. I would start with for the CIOs listening today is I've been a CIO really my entire career, starting as a government CIO for both the US Marshals and for Counter-Intelligence, with a full career in government. Went over to Georgetown University as their CIO, and then came to the Intel Corporation about four years ago, starting in IT before transitioning to the business side of the house in data center group. And within running this business, I'm really fortunate, because I get to do the best of both worlds. I engage in digital transformation strategies or IT modernization strategies, and that really is helping our CIOs understand what is the architecture that's going to be required from a current state to a future state perspective and really driving their modernization efforts and really ultimately helping, from a government standpoint, the mission succeed. How do we enable the mission? How does technology enable that mission? And on the corporate side, how do we ensure that your business or your company is going to remain relevant today in this digital era and this data era that we live in? So, we engage with our CIOs in helping them figure out what that strategy looks like, what that architecture looks like, and why – as we're going to talk about today – hardware does very much matter when we make those decisions.

But I also get to engage with our partners, and those partners, whether they're SAP, Microsoft, Oracle, VMware, all of our OEM partners – Intel has an incredible ecosystem of partners that we engage with, not only from a software optimization standpoint on our hardware, but also how we can build solutions that allow our CIOs to adopt, integrate technology faster.

David Linthicum:
That sounds great. I think that, by the way, being CIO of a government organization, whether it's on the military side or the civilian side, is probably one of the most difficult jobs in town. In other words, your ability to deal with the underlying need for resources, the ability to deal with very hard and fast budgets, and the ability to, in essence, keep up with the modernization of the existing technology and do so in a pragmatic way, that's – I always feel for everybody who's out there fighting the good fight on the government side. So, as we're moving forward, I'm focusing a lot on multi-cloud computing, specifically cloud complexity management stuff, which is out there in Deloitte.com if you want to check that out, but I'm watching the market explode in front of me. The ability to leverage hybrid cloud computing, hybrid IT, the ability to leverage this in a multi-cloud scenario, private and public clouds working together, multiple public clouds working together, and the ability to get this in a pragmatic march, and so we are helping the government CIOs and the civilian CIOs out there in moving their path forward. So, how extensive is the market interest in hybrid cloud, Lisa, and what are you seeing in the space right now on the Intel side?

Lisa Davis:
I completely agree. I think the latest data point – 70 percent of organizations or – IT organizations – have a hybrid multi-cloud strategy, and of course we look at these organizations, really your medium to large enterprises, of course many applicable within the federal sector, in which I have workloads that I need to make decisions on. I need to protect my data. I need to understand where I'm going to leverage my data to enable the mission, to increase my business relevance. And it's a hybrid multi-cloud world that we live in, and unfortunately it has only gotten more complex, so the complexity has continued to rise, frankly, as more choices have come into the market with our CSPs, our cloud service providers, and many other business models. So, it's good to have choice. It's good to have diversity of solutions in the market, but it also drives the complexity, in terms of what decisions do I need to make that best suits my enterprise and my agency, and not only best decisions – and those decisions are based on cost, security, protection of data, access to data. All those decisions need to be incorporated. But it also is important in terms of the complexity increasing where, ultimately, I decide to do business with and where my data will lie. So, multi-clouds become front and center, because we know most enterprises today, both commercial and federal, will distribute their workloads across multiple cloud providers.

David Linthicum:
So, Doug, what are you seeing in this space? Do we see an increased interest in the hybrid cloud and multi-cloud computing in the government space? And what specifically are they interested in?

Doug Bourgeois:
Yes. I completely concur with Lisa's assessment there. I think on the one hand, the way I've internalized it the IT industry has come full circle. There was a time, not too long ago, when everybody believed that infrastructure and cloud were commodity resources, and the situation is rapidly changing. And the reason why it's changing is that, one, people's knowledge and understanding of the complexities and security requirements and performance characteristics of their own workloads as increased, and it's been driven by their desire to actually adopt cloud. And, so, they have taken the time to look into how – what are the implications of driving their workloads and their data over to the cloud. So, what are they interested in? I think, based on that is the backdrop and the understanding that enterprises and government have come to terms with, they're not realizing that it's not as simple as they thought it was going to be, that the workloads and the performance characteristics and the security requirements are driving really hardcore, substantial strain into the platform, and as a result, the platform needs to be architected in a way to address those characteristics.

And at the same time, because of those requirements that I was referring to, they realize that not everything is suited to going to the cloud. I think even industry-wide, whether it's government or non-government, you're going to find overall, particularly from an enterprise standpoint, that no more than 20 percent or so in the aggregate of the workloads have gone to the cloud, and there's valid reasons for that.

But because people are looking for an elasticity to scale up, they're looking for agility for DevOps type workloads, they know that they need to connect their on-premises systems and solutions to cloud solutions and they need to have them run in a seamless manner. And, so that's something that's really driving the architecture. And then the last piece that I think is really, really important is where the industry is going. Big data, analytics, and IoT from an improved customer experience and capability at the edge perspective are also accelerating this push to the hybrid cloud, and those specific workload types and use cases are putting even more strain on the underlying platforms and driving the hybrid cloud design and architecture as well.

Lisa Davis:
Yeah, if I could just add to that, I couldn't agree more to what Doug just said. If you just look back – just three years ago, where there was a mindset that all workloads would move to the public cloud, and this is fascinating to me. And, of course, those of us that were CIOs were like, "How can that be possible? That doesn't make sense." So, now we have – and sometimes I call it, "Have it your way." It's like Burger King. I have all of the major CSPs – I'm not going to go into names – that now offer hybrid cloud solutions, and I think what's most important for our listeners to understand is that all of those solutions are not created equal. So, within those hybrid solutions that I have with each major cloud service provider, I have a very different approach, and I might want to call it a walled garden, in terms of once I place my workloads with that particular cloud service provider, I have the ability to orchestrate between their public cloud environment and a private cloud environment. It doesn't necessarily give me the opportunity to orchestrate across multi-cloud environments, so it's certainly – in the job that I have at Intel, part of what we engage with is those partners, whether it's a Google Anthos, whether it's a Red Hat IBM OpenShift, whether it's a Dell VMware Cloud Pivotal model, in terms of multi-cloud orchestration, because I truly believe, as an enterprise customer, as an enterprise CIO, I should demand common standards and APIs to multi-cloud orchestration.

David Linthicum:
So, I'm a CIO, and let's say I'm a government CIO looking across the table at both of you, and you're, in essence, having a discussion with me as to ways in which I can modernize my infrastructure, and we're considering – and by the way, I love what you said, Doug, about really starting out with a self-assessment. I think that's what you have to do to actually get this stuff right, because we're going to map a solution set into a problem set, and our ability to get to an optimized solution means understanding where we're coming from. I think that the fact of the matter is that everything out there with enough time and money will work. It's a matter of optimization and the solution that you're getting into, and that's why we're getting the hybrid and multi-cloud, because we're able to leverage different options and different best-of-breed technologies to make that happen. What benefits are we going to expect? If I'm a CIO, I'm asking that question. How do you respond to that?

Doug Bourgeois:
Well, several different ways. First, I'll put a big exclamation point on what Lisa said about all of the cloud service providers have hybrid cloud solutions. None of them are the same, and the projection from that is the adoption pattern – to your point, David, optimization – how am I optimizing these capabilities in my overall architecture to solve my business and my mission problems, which is taking all of our enterprise and government clients to the multi-cloud model? That is a foregone conclusion that's playing out before our very eyes, and that's why multi-cloud is becoming much more important. So, when you boil that down to the benefits and why you would want to do this, you mentioned optimization. It's optimization from several different perspectives. There's a critical element of performance optimization here. So, the architecture delivers the capabilities and demands of the workloads, and so it drives from the workload capabilities, and, as I alluded to earlier, there's a large diversity of the workloads here, and the characteristics and the performance requirements that they place on the platform are very different, depending on those workloads. Some are going to drive to memory optimization. Some are HPC-driven. Some are really around data protection and security, and those are not going to lend themselves to exactly the same architectural solution, so that's one.

Financial optimization – many of our clients learned that the sins of the past in their architectural decisions that they made in their traditional legacy infrastructures – in other words, throwing more resources to make sure I could scale and deliver the performance – bringing that over to the public cloud was not an economically-efficient decision. And, so, when you look at it in the context of hybrid cloud and the multi-cloud environment that's evolving before us, it's really, really important to optimize the architecture to deliver financial performance as well as technical performance. And then finally I would say – so those are two benefits right there. The third is really the agility. So, the architecture itself, and it's not just at the physical level. It's up to the management layer and the abstraction layer. It's providing the ability to extend your use cases from the on-prem to the cloud and from the cloud to the on-prem. So, let's look at DevOps from that vantage point. Some clients are adopting DevOps from a DevTest standpoint in their data center, moving to production for the cloud. Others are going the opposite way – DevTest in the cloud, QA and production in their data centers – and the hybrid cloud architecture is enabling that to happen more seamlessly and much more economically, so that's agility as a benefit.

Lisa Davis:
Yeah, I think those are three great points. For CIO in making these decisions, I think we could simplify that by saying one size doesn't fit all. If we look at those workloads – and Doug mentioned 20 percent of the workloads – and I've referred to those as commodity workloads – email, many times CRM, dev test is a workload that I – disaster recovery, possibly – that I have decided to move to a public cloud environment. But those workloads that are mission critical – core IP workloads, workloads that I'm going to make a decision that I need to keep on-prem because of security reasons, probably front and center from a federal government standpoint – then a one-size-fits-all, or commodity hardware, is probably not your best answer. And this goes back to the point that Doug made around AI – artificial intelligence – and analytic workloads, and these workloads that are now moving to the edge, and certainly at the edge from a war fighter standpoint, where I had zero latency requirements.

I need data that is occurring or being computed, analyzed at the edge, maybe feeding back into, for example, a common operational picture from a war fighter or government perspective. How am I going to move that data and those workloads across my distributed ecosystem? And this is where we talk about where hardware matters. Do I need to leverage Intel's new Optane data-center persistent memory that gives me the access to put more data next to my CPU that is persistent, lower cost, that can accelerate and drive performance with my data set across my environment? These are decisions where hardware features, capabilities, optimized for the particular workload, really drive – and I'm going to use the word performance, in terms of a broader context of performance, versus CPU performance, but relevance of the data set accessibility, agility, security of that data set that is really relevant for our government CIOs.

David Linthicum:
So, Lisa, you said something incredibility profound. The hardware does matter, and the thing is – and I'm a software engineer by training, and the reality is that in the last five years, since I've been in the cloud and moving into IoT and edge-based computing, I've learned more about hardware optimization and the ability to write software correctly for certain hardware platforms. So, I agree with you wholeheartedly. The hardware really does matter. But that's really a concept that I don't think is well enough known out there. Certainly, I'd love to hear from the listeners around this particular angle of it. So, Doug, if the hardware matters, how can we educate the public into understanding how to the leverage hardware more effectively and how to pick hardware effectively?

Doug Bourgeois:
Yeah, that's a great question. Like I said, it's come back around full-circle, but in the context of modernization in today's current architectures and technical capabilities. So, the first is – and this is very fundamental –stop thinking about the hardware and the underlying platforms as commodity, because it's just not the case, particularly with the diversity and the evolution of the workloads, and Lisa even referred to an IoT scenario that is front and center in driving even more stress onto the underlying platform. So, recognizing let's get out of that fantasy land when we all wanted to think everything was going to be simple and it's an easy button. There ain't no easy button. The platform's not commodity. Let's just recognize that. So, the second is then focus on the optimization as I referred to earlier. Look for better performance. Look for higher security. Look for financial optimization to deliver the requirements of the complex workloads that we really are driving into our enterprise to solve our more complex mission problems, including AI and edge computing, and allow for a longer-term performance and cost optimization around that.

And the third part is something that Intel and Deloitte are working on together in the context of our hybrid and multi-cloud architectural approach, which is there's a talent situation here, too. Solution architects need to be familiar with underlying key capabilities. Lisa referred to Optane. OpenVINO as a convolutional neural networks capability. We have engineers that are using the cloud to build neural network models, leveraging the OpenVINO capability, then deploying those models onto computing infrastructure at the edge that has Optane plugged in, in order to help optimize those – the performance of the infrastructure to support those needs.

So, the point is these architects need to understand scalable processors. They need to understand field-programmable gate arrays. They need to understand what a nonvolatile memory express with solid state drives is, and how it fits into the architecture. Software guard extensions from a security standpoint. They need to know how to use OpenVINO and leverage the capability of convolutional neural networks to drive these workloads into solving business problems. Otherwise, what's the point, if we're not solving the problems that our mission and our business users really are looking to leverage this technology for? So, it's not about the architecture itself. It's about solving those business problems in a way that the users are going to adopt and get value out of it.

David Linthicum:
I love it, Doug. Your propeller is spinning at full speed.

Lisa Davis:
It is. He's fired up now. And I love how he's talking about all those Intel capabilities. It goes back – we used to think of a data center. If you think about your enterprise data center as on-prem workloads, we are so beyond thinking about just what lives within your data center, because as a CIO, first and foremost – and Doug was calling this out – you have to understand the mission requirements. You've got to understand the business requirements and how technology will enable that to provide better outcomes from a mission standpoint or a business standpoint. If you don't understand that, that's job number one, because that is why the IT organization should exist today. Second of all, we need to understand that, as CIO I am managing a distributed ecosystem in which I will make decisions of workload placement, whether that is public cloud, on-prem private cloud, or sitting at the edge. And ultimately, those decisions need to be informed by my data strategy. Where is my data? Why is that data relevant? What data is needed for the mission or needed for the business to remain relevant and competitive in today's market?

So, understanding that and taking a step back and realizing, "Oh, this is much bigger and much more complex than it's ever been before," and how partners and certainly us working together, can really sit down with our CIOs, sit down with our enterprise partners, and really help them understand how best to architect. Why does infrastructure matter still? What does hardware matter? And how do I leverage, if look at Intel's broad portfolio of capabilities –How do I leverage that broad portfolio of technology, whether it's security feature sets, OpenVINO, AI capabilities, accelerant capabilities, new memory capabilities? How do I leverage that broad portfolio of capabilities to differentiate my mission and business by those most critical workloads? That's the job that needs to be done.

David Linthicum:
Indeed. So, let me have you guys do an exercise for me. Let's give the listeners a homework assignment. So, Lisa, where can they find out about Intel stuff on the web, and what do you typically read to keep up with the industry? And Doug, what do you typically read to keep up with the industry? Lisa first.

Lisa Davis:
Oh, what do I typically read to keep up with the industry? I read a little bit of everything. I read Wall Street Journal¸ Cloud Wars, federal government in the news. I'm a huge podcast fan, so I'm excited to do this podcast as well. In keeping current, I'm fortunate and have the ability to engage in a regular basis with key enterprise partners that all of our CIOs and enterprise engage with today, whether that's SAP, Microsoft, our OEM partners, so I have the ability to engage with them and really work on solutions and capabilities that ultimately help our enterprise customer in the long run. And you can always go to, certainly, Intel.com. We have a broad lineup of what we call Intel Select Solutions that allow you to take advantage of all of the cloud services for your on-premise data center and its simple and accelerated deployment plan. These are verified solutions that opportune and optimize infrastructure for your private cloud strategies, whether it's VMware vSAN, Microsoft Windows Server, Microsoft Azure Stack, HBI, Red Hat OpenShift container platform. Think about how a solutions approach and capabilities that Intel now provide with many of our other ecosystem partners, of, "How I can buy a performance-tuned, optimized hardware platform that I can plug and play into my infrastructures to start modernizing the capabilities within my ecosystem."

Doug Bourgeois:

Yeah, I'll take a stab at it. It's hard to cut through it all, because I get a ton of feeds, and I just monitor my feeds all day long. I couldn't tell you where half of them come from. But I'm a big Wall Street Journal fan as well, as well as Forbes online. That's where I try and get grounded in a business perspective as well as Business Insider. I read those on a daily basis. All of my gov feeds, in particular GovExec, but I pour through – you can tell I do connect with my old inner-self in the propeller-head world and read all the white papers from some of our key partners, I pour through all of those. But I also suggest that there's a blog post that is associated with this podcast, "Why Hardware Matters for the Cloud." I would suggest that the listeners go ahead and take a look at that. And then also I coauthored a white paper recently called Designing Hybrid Cloud Architecture for the Future. I would suggest that—text string will pop it up right off of Deloitte.com, and I think that's probably a pretty good, quick read as well.

Lisa Davis:
Yeah, it's an excellent – I highly recommend Doug's white paper on that. I thought it was great. We're in that white paper with them. We jointly talked about as we build our hybrid cloud solutions, multi-cloud solutions for our enterprise customers, so totally agree there.

David Linthicum:
Wow. Check it out. Great references. So, if you enjoyed this podcast, make sure to like and subscribe on iTunes or wherever you get your podcast. Also, check out our past episodes, including On Cloud Podcast hosted by Mike Kavis and his show, Architecting the Cloud. He has a book by the same name. I recommend that as well. And if you'd like to learn more about Deloitte's cloud capabilities, check out DeloitteCloudPodcast.com, and if you'd like to contact me directly, you can reach me at dlinthicum- it's a tough one, D-L-I-N-T-H-I-C-U-M - @deloitte.com. So, until next time, best of luck in building your cloud-computing solutions. We'll talk to you guys next week. Cheers.

Operator:
Thank you for listening to On Cloud for Cloud Professionals with David Linthicum. Connect with David on Twitter and LinkedIn and visit the Deloitte On Cloud blog at www.deloitte.com/us/deloitte-on-cloud-blog. Be sure to rate and review the show on your favorite podcast app.

Visit the On Cloud library

www.deloitte.com/us/cloud-podcast