



The Deloitte On Cloud Podcast

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

Title: A view from Red Hat on edge computing: What's now and what's next

Description: In this episode, David Linthicum talks with Darrell Jordan-Smith, senior VP of industries and edge at Red Hat, about how companies are learning to manage and grow their edge environments. They discuss new edge computing products and services and how companies are leveraging them to innovate and drive improved customer experiences. They also talk about current edge trends like containers, automation, open source, and data sovereignty, as well as the future of edge computing.

Duration: 00:23:60

David Linthicum:

Welcome back to the On Cloud podcast. Today on the show I'm joined by Darrell Jordan-Smith, senior VP of industries and edge at Red Hat. How are you doing, Darrell?

Darrell Jordan-Smith:

I'm great, David. How are you?

David Linthicum:

That's a cool title by the way. Industries, that means all the verticals and understanding how they work and the details with the industry-specific services and also the edge-based system, which is also very cool. So, tell us about you. What's your background and how you came to Red Hat? What's a typical day in the life of a senior VP of industries and edge at Red Hat?

Darrell Jordan-Smith:

Sure. My background is in technology, broadly. I graduated from university many years ago. I did an MBA like a lot of people in my position. I focused on electrical engineering, computer science, and mathematics. So, my background, naturally, originally started in engineering. I found myself doing a lot of work in the business, particularly for salespeople. Then, over time, it evolved in terms of wanting to understand more of the business dynamics, and that's what's really interesting about my role at the moment. I get to focus on all the different industries that we touch at Red Hat from an open source perspective, as well as going deep in a few areas where we see immediate and aggressive opportunity to add value.

David Linthicum:

What's an example of that? What's an example of an industry-specific engagement that Red Hat would be involved in. What role do you play? Also, what role does Red Hat play?

Darrell Jordan-Smith:

Sure. From an open source perspective, we try and drive open source communities to understand key use cases. Telecommunications is a great example. We're looking at use cases to create a common platform, a horizontal platform for telcos and operators around the world to really digitize their business, so they can deliver everything from their 5G core operating environments, through to their edge environments with their radio access if they're a mobile operator, through to their MEC, if you like, in terms of the multi-edge compute arena, into the enterprises that they're also trying to sell into, and therefore, also, extending out to the industry.

So, a day in the life of me is really focusing on helping our customers build out that platform, that infrastructure, doing it holistically, because in telecommunications, traditionally, they built vertically integrated platforms versus horizontally oriented platforms. In a cloud native world, if you're creating a telco cloud, it's a lot easier to do that in a common platform, a common set of technologies that you can scale out horizontally. So, you can actually have a common control plane, common automation capabilities, a common set of applications and services that you can deliver all the way out to the clients that you are trying to sell into, if you are a service provider, for example.

So, I'm very focused on those particular use cases. Then in the upstream communities, that touches things like sustainability. How do we control hardware, silicon? We call PNC state a silicon to look at how we reduce the footprint of power being used, through to how do we factor the application serving a RAN, for example, radio access network. We're looking at how do we tune the radios so they don't have to be on all the time? How do we rewrite those applications so we can drive sustainability topics there? Focused on the complete system availability with projects like the Kepler Project, which is an upstream project where we look at everything from the hardware all the way to the application and beyond in terms of how those things drive that sustainability. Other things in my role, which I'm really excited about, are really what we do around diversity.

By definition, open source is a very diverse community. We've really got to attract all those developers, millions of them out there, that are really focused from an engineering perspective on building the next best-in-class applications and services. We really act as sort of a broker between that developer community and traditionally what a regulated industry like telecommunications, as one example, is trying to drive to. So, really trying to help those developers understand those use cases, so they can actually create the best possible code. They can innovate. They can deliver that value into the telco. So, hopefully, David, that gives you a feeling for a day in the life of myself and what I do.

David Linthicum:

Wow. Busy dude. Awesome. Back in the '90s I wrote three books on integration, the EAI, B2B integration. That talks about kind of the next generation, where this stuff is going. Here we are in 2023 and we seem to have more silos around as we deploy in hybrid clouds, multi-cloud, different applications, edge-based computing. Or at least we're running the risk of creating these information silos and even process silos. Ultimately, the value is going to be our ability to create links with processes and data flows in between these various systems, and do so in such a way where it's going to be able to be configured and reconfigured as we need to adjust to the needs of the business. Tell me about that part.

Darrell Jordan-Smith:

Well, the work for us on that side of things, there's another layer that we're heavily involved in, and that's around data sovereignty. A lot of countries, because of geopolitical situations, and the importance of security associated with data, we're very focused on where the data sits. So, we're trying to create environments where data isn't in a silo. So, data itself can be interoperated between different applications in the cloud. So, you try and use that to break down the silos. Then you drive integration across the application sets in order to facilitate next-generation and innovation across the top of that. So, what that might look like might be a patient record. Look at health care. In the United States we have HIPAA, which is a regulated environment with all your hospital patient records. That data can't easily be crept out of, in this case, North America, or it can't be shared openly with lots of different people without the individual themselves giving permission to be able to do that.

As you can imagine, a hospital could be a silo in locking all that data in. So, we're trying to break that open by leveraging cloud-based technologies to have that data stored and protected in such a way in a data sovereign environment, but have access to it through different applications and services. We then double-click on that again and look at some of the applications that we're looking at there. Look at a simple X-ray. They might look at a particular cancer, for

example, and you layer AI across the top of that that looks at a particular cloud in an X-ray, and they give you some analytics around that. So, doctors can look at the fewer X-rays that they need to look at, to interpret that from a human perspective in terms of leveraging the AI components of that.

That data then needs maybe to go to a practitioner. It may need to go to be stored in your own records somewhere. It may need to go to different, other specialists. So, we're really trying to protect that data in the context of both the HIPAA compliance rules and laws associated with regulations around that, with multiple providers across multiple operators and hospitals and practitioners that would actually serve that particular patient in that particular example. We're using AI as a great example of how you might birth elements of that X-ray, for want of a better term, on that image, to a GPU somewhere that sits in a cloud. You don't want necessarily all the data to go to the cloud.

There's certain elements of the data you might want to go to the cloud. You need to control what that is because you never know within a cloud provider where that is being processed. Let's say if you're in the UK—where you can probably tell my accent is from—a patient record here currently in the UK could actually be transported to somewhere in the US, and it would leave the domain of the United Kingdom. People are worried about where that data goes and how it's being used. So, that data sovereignty, that focus and leveraging AI in the cloud is a really good use case and example of how we're looking at trying to break down those silos, and allowing maximum flexibility, expanding that using open source and software and innovation out there to solve some of those more gnarly problems associated with those types of things.

David Linthicum:

One of the observations that I've made over the last ten years and even the discussion that we're having here, ten years ago, twelve years ago there was a trend. Everybody wanted to centralize their information, centralize their data, so they could have a common access to it for lots of purposes, including running applications and now using it as training data for AI, things like that. We get up to 2023/2024, it seems that the push is now or the trend is now to leave the data where it is, where it resides. Chances are it can be secured and protected there. It's going to get us in the least amount of trouble because of the data sovereignty things you just mentioned, and the ability to leverage it in a ubiquitous way. In other words, no matter where it's physically stored, the ability to consider it and leverage it as a single source of truth. What do you think about that?

Darrell Jordan-Smith:

I think that's exactly where we're leaning towards. If you look at the next generation, if you like, of where telecommunications and computer services are heading, I'd recommend anyone listening to this to look at IOWN, which stands for Innovation Optical and Wireless Networking forum. It's a group of about 500 companies now, it started with about 50, where we're looking at how you disaggregate data, keep data in memory so you have fast access to it, but you can disaggregate storage, network, and compute using an all-optical network. So, really, the compute, or the telco itself, the whole network is actually being redesigned around looking like a massive compute platform, being able to leverage that data that sits in one place, and asks for elements of that data to be processed as appropriate. What that also means is that you're going to need to look at how you might rewrite applications.

So, think about databases at the edge, where the data isn't at the edge of that network, but it's in a core area where it's protected, but you can do inference across that data, and you can apply applications and services across the top of that. So, if you're a developer in the future, we're looking five years away from now, I think that we're going to look at ways where we're building new applications and services, leveraging that next generation of technologies. We're beginning to look at that now.

David Linthicum:

Isn't this a matter of decoupling the applications from the interfaces that deals with the abstract data? In other words, we're dealing with abstraction layers and looking at different physical databases, using one unified schema that we can change to adapt to the needs of the business. So, even though we're recoding some aspects of it, we should be communicating down to the common API that's able to resolve the data location and find the physical data that we need to carry out whatever operations that the application or AI engine is running. Is that the focus now?

Darrell Jordan-Smith:

I think that's a big focus of what we're doing. You'll have databases for databases. You'll have databases that will do certain computations really efficiently and others that will do it in other ways slightly different, from a different level of efficiencies, and different applications that will run across the top of all of that to help break down those silos. We have a project here at Red Hat called Service Interconnect, where we can actually, independent of Kubernetes platforms that you may be using or, container platforms where your applications may run, we can interconnect those containers so you can access elements and data. We're beginning to do that today. The limitations are obviously speed at which we do that, and the more we move out into the future we look at all optical networking and 6G technologies and Wi-Fi 6 and how all those things come together.

I think that's how the future of what we're going to be experiencing will be in the next five years. I also think, just touching on that, there are going to be data used in individual that you'll want to keep on your filing or around you very specifically. So, you might want to go from one country to another, and you might want to give people access to that in a controlled way, but giving you a little bit more control as an individual on that data, and letting you as an individual make the tradeoff between your privacy and how you want to utilize that data. So, applications and use cases like, I'm going to a theme park. I'm okay with having an image of my face, so I can get in and out of the park easily. I can do grab-and-go services when I'm buying goods in the park. I can get on the fast line without having to show a ticket or go through somebody and have a delay there.

I can use those environments in my phone or my device that connects with that environment, but I'm opting into that experience. I'm giving, in that instance, in that example, the opportunity to leverage certain data attributes about me, and my facial recognition is an example. So, I think people are going to see applications and services being developed at the edge along those lines, and we're going to be in a situation where we'll give certain credentials to give people access to their data for periods of time. And I think the regulators are going to get more involved in looking at that and what that actually means as we move towards it.

David Linthicum:

As I look at edge computing and some of the research that I've done and the articles I've written about edge computing, one thing that kind of strikes is the diversity of what those things are. Anything can be an edge device. It could be a mainframe in a system that's serving a particular role, but in many instances it's a purpose-built device that's doing a particular thing, and we're storing data on the device in certain ways and we're doing some processing on the device. So, obviously, we have to connect to those things in one way or the other. So, what are examples of edge connectivity and edge integration? What kinds of trends are you seeing going on in the industries right now?

Darrell Jordan-Smith:

Obviously, you hear a lot about 5G, 5G slicing, 5G private. So, you have the generic public 5G network. You might have a 5G private instance running in a factory, for example, connecting to telemetry and those devices. You might have, in addition to that, certain elements of containerization, which is really down to a very small piece of silicon. A single node on a silicon collecting or inspecting items coming off of factory floors is a good example. We're seeing those devices being connected over wi-fi and working in symbiotic ways between the 5G private and the 5G public environments in that way. Then we're seeing other carriers and other service providers saying, "Hey, we can do network slicing."

So, we can slice a network and give access to a service across multiple factories and environments which is consistent, so you as a consumer of that technology can build consistent applications and services, and we as an operator can operate that in the cloud. I see a lot of those things evolving at this moment in time. I think there's a lot of applications and focus and excitement, potentially hype in that space as well, but I think we're going to see some interesting applications and services enabled in those areas.

David Linthicum:

Yeah. I think we will as well. It's probably the most interesting part of what we do that's moving forward, because we're moving out of just dealing with generic servers for many instances. Obviously we're moving to quantum and some other areas and space and to these very elaborate, very cool, purpose-built devices, including robotics, including autonomous driving, including lots of things that we just want to do as human beings, and break computing power out of where it's found right now, even within the cloud, out to these edge devices there and carry out certain things and roles and things like that.

So, what types of products are out there now? What types of products do you guys deal with? What technology solutions are available to people who are looking to build edge-based systems and very distributed systems, and trying to knock down some of the data silos they have within their enterprises and within their problem domain? What works?

Darrell Jordan-Smith:

I think what we're seeing on our side, because Red Hat as a company is very focused and open source and all the different communities out there, there are millions of them literally, what's working at the moment, what we're seeing is Kubernetes. Obviously, Kubernetes started out as a big data center for these applications and services, mainly to put applications and services in containers. What we're actually seeing evolve, as you go out to the far edge as you mentioned earlier, is a lot of those containers are getting very, very thinner and they're working on different silicon. So, we're seeing companies this space creating silicon that really drives those elements and putting containers in those areas. So, you can actually have a consistent application environment, leveraging common data fabrics that enable you to execute new compute-based services in there.

We're seeing that beginning to work. The early applications for us in those areas, in telecommunications, are things like ORAN, Open Radio Access Networking, which is an open standards-based, disaggregated radio access networking technology. The interesting thing about that, once you start disaggregating these you actually create the opportunity for innovation, but if you're a big telco provider, you have then got to figure out how to integrate that.

So, the challenge then is how does that ecosystem of companies in the marketplace come together to help you do that integration, in a manner from a business perspective that doesn't create a situation where it becomes like a single flight to the moon in terms of cost. You really want to make sure that you can scale [and ~~cookie cutter~~ COSTS and drive that across your environment. So, that's one area. The other area is around automation. We hear a lot of automation. We hear about smart automation, really trying to actually use automation to secure environments to do the mundane tasks. There are lots of interesting automation technologies out there.

We're seeing agentless automation really gaining momentum, where you have an automation application sitting on a device, or in the cloud, that only comes on when you actually need to automate something. So, it's not using any computing power until such time that it starts doing that. So, we're seeing that gaining momentum in the marketplace. Then thirdly, from our perspective, you mentioned hyperscale earlier, the cloud providers. They're all looking at interesting ways of getting to these different markets and these different segments. It's going to be very interesting to see whether cloud companies, how their business evolves in this space. Are they going to become the next generation of digital service providers from a telecommunication standpoint? Are they going to look at how they interconnect what they're doing with satellite communications, so they can actually do new services in remote areas, and try to actually deliver those applications and services?

So, we're seeing that beginning to sort of show itself in the marketplace. That will, in turn, drive an increase in competition. You've then got the regulation elements of that in terms of how you maintain, protect, and potentially regulate those areas as well, which is a whole separate set of conversations. No one has really decided what's going on there, but people kind of feel it needs to have some sort of regulation. I'm sure the cloud providers—that it will be regulated in that way. So, there's obviously a lot of interesting discussion that's going to go on there in terms of where their services and how their services will evolve.

David Linthicum:

When I look at edge computing and where the hyperscalers are playing, they play huge roles in the edge. I remember a few years ago it was like cloud versus edge computing. Who is going to win? That wasn't the case. You've got to be the edge of something and typically you're going to be the edge of the cloud, and the cloud providers I think did some innovative development and developed some edge-based products. So, they participated in that market as well. So, where can we learn more about this stuff? Where do you typically go to learn everything we talked about?

Darrell Jordan-Smith:

You can go to RedHat.com and you can actually look at what we're doing around any of those topics that I mentioned. You can go to visit IBM and their website. They have a whole edge-based practice in their business. But also, a lot of our other partners are looking at edge-based products and services. Some really cool things are beginning to show there around robotics and retail and health care and autonomous vehicles, not necessarily for autonomous driving yet, but how you would actually collect telemetry and get more data out of the vehicle to drive more services and more of the software-defined vehicle type things.

So, those are the places I would go. There's a wealth of information out there. Obviously, I would default to the company I work for. I'm passionate about open source and what open source communities can do. We love to spend time and share that with anybody that has some deep interest there. But there are a lot of other companies that you can go and collect resources from as well.

David Linthicum:

Yeah, Key research and development you guys are working on is very, very forward-looking. Speaking of forward-looking, let's say we're doing this podcast in five years. What do you think we're going to be talking about?

Darrell Jordan-Smith:

I think we're going to be talking about two things. I think we're going to be talking about how do we accelerate the reengineering businesses to take advantage of some of the technologies that we're looking to build and deploy now. So, I think business process reengineering, changing cultures within companies to actually operate differently and think differently. That's what I mentioned earlier. I was heading towards diversity, not just gender diversity, but neurodiversity, those environments. I think those are going to be really interesting topics. I think the integration businesses out there, as we were talking about earlier, are going to be absolutely fundamental in helping businesses charter those waters and actually drive different innovation and look at different business models. Because I think we've got to facilitate that in the marketplace to support innovation and drive communities of interest that deliver value to both consumers and businesses, and therefore the overall economics of what we're doing with technology.

I love what I do, predominantly because I believe that particularly in telecommunications, financial services, manufacturing, all the industries I'm involved with, these are things that touch everybody every day. For me personally, I'd love when I retire to say I was part of that. I was part of doing this. I was part of that journey. I think everyone who's listening to this is going to be very much involved in all of that and working with an ecosystem with partners and the way that we go to market. So, in five years' time, I think we're going to be looking at how we partner differently in the marketplace and how that will mature.

David Linthicum:

Yeah. I join you with that aspiration. The ability to kind of drive where the next generation stuff is going to move, and the ability to kind of create building blocks, and IP and new ideas and new concepts become real solutions that people can implement change, companies change lives. I think that's why we're here. So, if you enjoyed this podcast—this was a great discussion by the way—if you enjoyed this podcast, make sure to like us, rate us, and subscribe. You can also check out our past episodes, including those hosted by my good friend, Mike Kavis. Find out more at Deloittecloudpodcast.com, all one word. If you'd like to contact me directly, you can e-mail me at dlinthicum@deloitte.com. So, until next time, best of luck with your cloud journey. Stay safe. Cheers.

Operator:

This podcast is produced by Deloitte. The views and opinions expressed by podcast speakers and guests are solely their own and do not reflect the opinions of Deloitte. This podcast provides general information only and is not intended to constitute advice or services of any kind. For additional information about Deloitte, go to Deloitte.com/about.

This publication contains general information only and Deloitte is not, by means of this publication, rendering accounting, business, financial, investment, legal, tax, or other professional advice or services. This publication is not a substitute for such professional advice or services, nor should it be used as a basis for any decision or action that may affect your business. Before making any decision or taking any action that may affect your business, you should consult a qualified professional advisor.

Deloitte shall not be responsible for any loss sustained by any person who relies on this publication.

Visit the On Cloud library
www.deloitte.com/us/cloud-podcast

About Deloitte

As used in this podcast, "Deloitte" means Deloitte Consulting LLP, a subsidiary of Deloitte LLP. Please see www.deloitte.com/us/about for a detailed description of our legal structure. Certain services may not be available to attest clients under the rules and regulations of public accounting.

Please see www.deloitte.com/about to learn more about our global network of member firms. Copyright © 2023 Deloitte Development LLC. All rights reserved.