EPISODE 1031

[INTRODUCTION]

[00:00:00] JM: A large cloud provider has high-volumes of network traffic moving through data centers throughout the world. These providers manage the infrastructure for thousands of companies across racks and racks of multitenant servers and cables that stretch under seas connecting network packets with their destinations.

Google Cloud Platform has grown steadily into a wide-range of products including database services, machine learning and containerization. Scaling a cloud provider requires both technical expertise and skillful management.

Lakshmi Sharma is the director of product management for networking at Google Cloud Platform and she joins the show to discuss the engineering challenges of building a large scale cloud provider, including reliability, programmability and how to direct a large hierarchical team.

As always, we're looking for show ideas, and topics, and guests for upcoming shows. If you have ideas, send me an email, jeff@softwareengineeringdaily.com. We're always looking for important topics to cover in the world of software.

[SPONSOR MESSAGE]

[00:01:14] JM: When I'm building a new product, G2i is the company that I call on to help me find a developer who can build the first version of my product. G2i is a hiring platform run by engineers that matches you with React, React Native, GraphQL and mobile engineers who you can trust. Whether you are a new company building your first product, like me, or an established company that wants additional engineering help, G2i has the talent that you need to accomplish your goals.

Go to softwareengineeringdaily.com/g2i to learn more about what G2i has to offer. We've also done several shows with the people who run G2i, Gabe Greenberg, and the rest of his team. These are engineers who know about the React ecosystem, about the mobile ecosystem, about GraphQL, React Native. They know their stuff and they run a great organization.

In my personal experience, G2i has linked me up with experienced engineers that can fit my budget, and the G2i staff are friendly and easy to work with. They know how product development works. They can help you find the perfect engineer for your stack, and you can go to softwareengineeringdaily.com/g2i to learn more about G2i.

Thank you to G2i for being a great supporter of Software Engineering Daily both as listeners and also as people who have contributed code that have helped me out in my projects. So if you want to get some additional help for your engineering projects, go to softwareengineeringdaily.com/g2i.

[INTERVIEW]

[00:03:04] JM: Lakshmi Sharma, welcome to Software Engineering Daily.

[00:03:06] LS: Thank you.

[00:03:07] JM: You work at Google Cloud on the networking team. Describe the modern networking challenges of high-volume data center infrastructure.

[00:03:16] LS: The challenges that we see in data centers, so I'll put them in like two categories. One is like cloud provider themselves managing their own application and services, right? Because you do want to continue to put to your own services into the same cloud infrastructure or data center infrastructure. Then, also, like you're bringing enterprise customers over the cloud.

From data management perspective, the biggest challenge that I would see is that of how do I manage my data? As a customer, as an enterprise customer, how do I manage my data and what data to retain? What data to throw off? Then how do I intelligently make decisions managing that data? Also, there are compliance related questions. Then as a cloud provider, I want to make sure that the data that I have to manage my infrastructure on which enterprise customers are coming, how do I continue to use that usage of large customers themselves

where I can only have access to how much of the interest that they're using so that I can do write capacity planning, I can build infrastructure into places where I may not have been.

So getting that insight from the usage of the enterprise customer, so that is what as a cloud provider I would like to continue to build up on. But as an enterprise customer coming to my data center, then coming to my cloud, how do I give them the tools and how do I give them the capabilities so that they can make the right decisions about what kind of data do they retain including the compliance and other aspects of risk and governance aspects of how they manage the data as well as give them intelligence about the data so that they can make right decisions about what services to choose, what locations to choose and all of that.

[00:05:00] JM: It's always been hard to build networking infrastructure for a company that is running a high-volume of servers and that Google has been doing after a long time, but since Google has become a large cloud provider, you have to think about networking in a whole new way. You have to develop products that not only are offering the right level of networking infrastructure quality, but you also have to give this level of control and programmability into the network. You have to give the remote developer programmable interfaces into the networking infrastructure that might have previously been just handled by internal Google engineers. How do you determine what aspects of networking infrastructure to expose to a remote operator?

[00:05:54] LS: That's and excellent point and I will answer that in two ways. One category from enterprise perspective is one is their own operators. Let's say administrators and operators, the folks that are networking administrators, the folks that are security ops people. That's one category of folks that we are building products for and those people are – Their job is to make sure that network is always up [inaudible 00:06:16] just the network infrastructure on-prem. Now you're thinking about the connectivity to the cloud. Now you're thinking about how you design a new virtual network and sell services around it including security and firewall. Giving the capabilities and APIs and the self-serve managed infrastructure do manage these networking services, security services, is one aspect of it. That's kind of one category of services that we offer.

The second we're talking about that the developers themselves, right? How are the developers who are using these network services, their network infrastructure, and in this hybrid

environment and who are doing migration? They're migrating from on-prem to cloud, and how do we give them that abstracted view of this infrastructure, which is again, underneath it, if we're talking about VPC, if we're talking about firewall, on-prem and in cloud? How do we make it easy now for those developers and for those administrators at the same time so that they give an abstracted view of the services? What does abstraction mean? Taking an example, if an enterprise is doing a VM-to-VM migration or VM to, let's say, container migration. We have tools like, say, Anthos for migrate. We have tools like migrate for GCE, which is coming from like the acquisition Google did some time back, which was [inaudible 00:07:40]. There are other acquisition that we did, for example, CloudSimple.

If you put these kinds of products in place, then what we hear from customers is that I have the specific way of designing. My DMZ, which includes DBI, which includes packet metering, which includes firewall, which includes site to site VPN connectivity. I have certain way of designing my on-prem infrastructure, but that is becoming a bottleneck to my developers. Why? Because in cloud, the same kind of DMZ does not exist. If it were to, and let's say if they were to bring the virtual instances of this as a DMZ from on-prem to cloud, they would be using methods of approval of API access and the cloud resources access, which would not be comfortable or which would not be easy for their developers to use.

For example, they may be using a jump host to come to the cloud. That means somebody in the security ops on-prem need to open up a pool to go to the cloud for every single API that they use. The team that is doing the VM management, the team that is doing container management, they all may be different teams, and what we know is that people like who are doing transformation, sometimes they'd build an organization called cloud organization, or cloud platform organization, while there are still other organizations like networking organization, CICD organizations.

What we want to do as part of our hybrid enablement of the infrastructure is that how do we give you consolidated view of the configurations that work on both the site? When I say configuration, I'm still talking about firewall services. I'm still talking about node balancing services. Our goal is that how do we build a common view of these load balancer across? How do we give you APIs so that they can work with your on-prem environment as well as like in the cloud? But also make it easy for you to migrate that. How do we make it easy to migrate? If you

have a firewall policy, let's say you're using VMware, for example, on-prem, and then you have certain way of introducing security policies using NSX and any third-party sitting on top. Now when you're migrating those same set of VMs over to cloud, then the firewall policies in cloud are not similar.

If we could translate those policies for you in an automated manner and present to you while you're migrating, detect your application to application relationship and use that also to add this recommendation that we are providing, it makes customers' lives very easy. We are kind of going multiple ways. As I describe that, we are looking at what do we offer to network operators? What do we offer to application owners and then how do we make it easy also to the cloud teams who are given the charter of migrating?

[00:10:39] JM: Can you take me through the design and development of a networking product that you worked on recently?

[00:10:46] LS: Yeah, I would love to. Recently, at London Next, we launched a product called Network Intelligence Center, and that ties into a lot of things that I just talked about. Network Intelligence Center helps you leverage your logs, your metrics and the services that you're using, it looks at all of them. It looks at your network capabilities and their configurations and it takes you from, let's say, the reactive way of operating and managing your network and services to a proactive and intelligent way of managing and operating your network infrastructure and services.

As part of that product, we launched three products, tools, let's say, as part of the product category. We launched three products, and those products are network topology graph. One product is called connectivity test and one product is called firewall intelligence. The way it helps, the way the product is used today, when you're migrating, when you're in an hybrid environment and you're migrating applications, you have different ways of configuring, as I described. Like your cloud base connectivity is very different from how you have connected like things with an on-prem.

When customers come to cloud, they think cloud is a black box, "Oh! Cloud providers don't give access. They just describe firewall in a certain way. They talk about network function

virtualization. Talk about SDN. Oh! They have made it so complicated." Our customers used to think like, "They made it is so complicated. They made it black box, and I don't understand how my services are even connected." They had no comfort level in adding more unknown network capabilities, because they could not trace the services and they could not understand the network at all.

Especially in case of network related issues and typically even [inaudible 00:12:30] application not working right, people always blame network. That's very common issue that we hear. If network is to be blamed all the time, then how do we make it easy for customers to troubleshoot that network and have very clear guidance on where in the network or where in the cloud network or in the hybrid connectivity, where exactly the problem lies?

There are two things we introduced. One is that network topology graph. If we give you a visualization of how your services are connected, right? How your, say, VPN coming from on-prem to cloud, and that then going through a public IP, private IP, then it is going through a load balancer, then it goes through an ingress firewall. All of these things happening in software. There are no kind of boxes doing this work.

But if it could represent that in a manner, which looks like boxes connecting to each other or circles connecting to each other similar to how these enterprise customers have been used to looking at them on-prem, so then it gives them so much comfort to understand where the bottlenecks are, and because they can understand it in the same way as they could understand it on-prem. This product, network topology graph, gives you very good visualization of your services how they're connected in cloud all the way from visualization into where your services are deployed, which regions are they deployed? How different regions are connected? How your VMs are connected? What is the bandwidth utilization? What is the latency between your different instance types? It gives you all.

Once you have that visibility, then it can – Then, still, you believe that your connectivity is right and everything is great, but then you're still seeing a problem in connectivity. Let's say you cannot offer GCS bucket or your one VM in one region cannot access VM in another region or maybe your on-prem instances are not able to access parts in the cloud. Then we have a tool called connectivity test, which is very simple to use, and what you do is that [inaudible 00:14:39]

to this tool is the two endpoints, endpoints could be IP addresses and port. Endpoints could be instance names. Basically, you provide that information and then you run the test and the tool will exactly point you to where the problem is, where your traffic is blocked, because the tool basically does reachability test. These are some of the things that we have brought in.

The third aspect of it is intelligence. Let's say everything is working fine and now you want to continue to add your firewall rules. You continue to add more load balancer policies, backend policies. Overtime, your configuration becomes so unmanageable that you're even afraid to actually touch which firewall to takeaway and then you know that managing firewall rules and also keeping them under compliance is a huge kind of security requirement for customers.

Then if we could give you a tool where we can exactly tell you, "These are the firewall rules which are being matched with your traffic overtime," and there are these others who are always shadowed or give you a recommended configuration of how you can optimize, let's say, you thousands rules to hundred rules and we guarantee that there would not be any pattern reachability, and you could in fact use that report to go back to your compliance team and say, "Hey, it actually still complies."

This has made our security operators life very easy and compliance teams from enterprise customers, they really love these kinds of tools where we can give them a config log or a trail, and with this optimized configuration so that they do not have to worry about going to line by line management and as well as getting approval from the compliance. There are lots of these things we have done. These three products have already come and they're more coming under the same suite.

[SPONSOR MESSAGE]

[00:16:34] JM: Gauge and Taiko are open source testing tools by ThoughtWorks to reliably test modern web applications. Gauge is a test automation tool that makes it simple and easy to express tests in the language of your users. Gauge supports specifications in markdown, and these reusable specifications simplify code, which makes refactoring easier and less code means less time spent maintaining that code.

Taiko is a node library to automate the browser. It creates highly readable and maintainable JavaScript tests. Taiko has a simple API. It has smart selectors and implicit weights that all work together to make browser automation reliable. Together, Gauge and Taiko reduce the pain and increase the reliability of test automation.

Gauge and Taiko are free to use. You can head to gauge.org to know more. That's G-A-U-G-E.ORG to learn about Gauge and Taiko, the open source test automation tools from ThoughtWorks.

[INTERVIEW CONTINUED]

[00:17:49] JM: I'd like to understand a little bit more about how networking has changed in recent years, I guess, with – Well, for example, a move towards containers. The increased popularity of containers. Does that create any networking challenges that are significantly different from VM-based infrastructure? Just as an example, if we compare networking infrastructure where it's mostly just people spinning up applications on VMs versus applications hosted on containers, how does that create different networking challenges?

[00:18:22] LS: If you were to ask me this question like two years ago, I would say there were a lot of challenges. I will address this in two ways. One is that how generally people perceive it and what Google has done. Two years ago, when we started on this journey of GKE, which is our managed Kubernetes infrastructure, we decided that Google will be a container, say, container native company for enterprise customers, while our infrastructure is already container native.

What we have done is that, since then, every product that we built, whether it's load balancer, firewall net, public IP, private IP, we build it for both containers as well as compute or VMs at the same time. An example, if we are building a load balancer, the same load balancer that we use to load balance to VMs would be used as an ingress for a GKE port, right?

If the same policies that we are using for firewall in compute are utilized in the GKE or a GKE environment. From that aspect, we have evolved and we are very mature in terms of our

understanding and the implementation for managed container or managed Kubernetes or, say, GKE.

From that aspect, just from a Google perspective, we are pretty mature in our integration because we have learned very early on as we decided that any product that we build – Another example is like, say, CloudArmor. Another example is CDN. Another example is Cloud DNS. Cloud DNS is another one we are at this time natively integrating. Yes, there were challenges and we learned about them very early on integration with GKE, and that's what I would say that we are a part in terms of every product that we have built in offering to containers as well as to VMs.

Yes, there are a lot of known challenges in an outside world if you were to bring in these open source environment for Kubernetes outside managed infrastructure that we offer, you would run into a lot of plugin challenges, CNI challenges. How do you distribute load instances within a part versus like cross to parts? They would be all kinds of challenges that you will hit into.

[00:20:44] JM: Let's say I'm a company that is entirely in the clod. Let's say I started my company entirely in the cloud. I'm doubling every year. More and more people are using my product. I'm scaling. How is my network going to look different overtime and what kinds of networking challenges am I going to encounter as I scale?

[00:21:04] LS: As a company who is growing in cloud, if you have not, a couple of things which will happen. How did you plan your, say, demand for IP addresses? That's typically we see quite often. People like who come from – If you're cloud native versus, let's say, a legacy enterprise. This question I'm addressing is if you are asking for the legacy enterprise, or not legacy enterprise, or legacy applications coming to cloud versus cloud native, because cloud native companies, we see that they have evolved with us because Google has been the first choice for cloud native companies for years. We really have done very well learning from them, building them and scaling them.

I believe that for cloud native, these similar challenges may not apply, although some of them. IP addressing is the scheme that we hear quite a bit. Customers might have used – Large enterprise might have bought some public IP address pools which they have been using. They

might have used some email services where they had to expose an IP address to a DNS mapping and a URL mapping outside their environment. Then how cloud providers do that mapping in the DNS or in their IP addressing scheme. Which regions do they map? What IP address to? We see that IP address-based mapping and questions quite a bit.

For those kinds of offerings, we [inaudible 00:22:31] product called bring your own IP, where we try to help customers by saying, "Hey, if you already managed a block which you had used in certain regions or certain part of the world, then you can bring that IP address for this and we will help you design." That's one challenge we see.

Then second one we have seen that since cloud enables customers to do the things they could not do. If you had to go like, say, from one region to another, it's just so easy to [inaudible 00:22:57] your traffic to test it out and then you introduce your new service. We have certain gaming customers, for example, they were only in Southern America, and then they are trying out in, let's say, Asia. Completely different zone for them and territory for them, and they can do things that they have not even imagined.

Then for those kinds of customers, helping them out with networking design, that what is the best networking design for them to access, for example, interregional resources? How do they manage their egress cost when they're going outside the region? How do they manage their load placement, which also requires understanding of what kind of instance types do you need across the regions what are available? Those are the challenges I see. In fact, some of them are good ones, right? Because you realize then that there is more potential to your business than what you could potentially manage before. With that comes these hidden things that you need to be aware of, how would you scale, replicate your data, for example? Then when you're replicating your data, are you aware of your data localization requirements? What kind of APIs would you use? How would you use identity? How would you use, as I said, like the IP addressing scheme or how would you use the identities or geo-location understanding of what data should be mapped where? These are things which customers would not have thought before. They end up thinking about when it comes to scaling their infrastructure or applications in cloud.

[00:24:22] JM: How does Google decide where to open up new data center regions?

[00:24:29] LS: There are a lot of kind of factors involved. To be honest, it's really a standardized process. The way you look at it is that you're looking to your services that you're offering in different parts of the region. We look at the demands that are coming. They are coming from feel for our services and usage. Once we understand that – Remember, like one thing about Google is that we try to be – When we launch a service, our goal is to launch that service in 100% of our regions almost at the same time. That helps us get signals to where else do we want to add new services and where else our customers are coming from.

The second thing that we also look at our customers that we already have, for example. We continue to work with them and understand that where are other opportunities they have globally? Some customers, as I mentioned, they maybe only US-based customers, but they want to launch – I'll give you an example of gaming. They may only be offering gaming services in one part of the globe. Then we learn signals from our customers. We learn the usage of our services and once we understand that based on the usage and the requirements for the customers and the proximity of those customers, like these new customer demands, new services, is how we choose our regions and decide based on that.

[00:25:51] JM: As your role, you're director of product management, I think, for Google Cloud Networking.

[00:25:57] LS: That's correct. Yeah.

[00:25:58] JM: Is it important for you to go to the new physical data centers or is it sufficient for you to communicate with them remotely?

[00:26:06] LS: I would have loved to. Coming from like my networking background, I always loved to go to these data centers and parts, but no. For my work, like I'm on the cloud networking site. I don't have to go there and visit there. Also, even for ourselves, like internal employees, the guidelines and the security requirements to go these places are pretty high. Unless until there is something crucial that somebody has to, we don't go for security reasons. At the same time, my work does not require me, because are like a completely SDN, completely software-driven. We don't have to go.

[00:26:46] JM: All right. Well, on the product development side, it's 2020. What are the hardest parts of building network infrastructure? What are you focused on right now?

[00:26:57] LS: I think the biggest challenge at this point is the data related regulations and asks that are coming from different parts of the world. When we built Google's infrastructure and we built Google Cloud, we built services that could be accessible to every customer in any part of the world with almost same latency and absolutely same experience. Any service that Google has built, whether it's Gmail, GSuite or any other part of that is sitting in cloud.

Now, in the enterprise world, you have enterprises. Like, say, some of your [inaudible 00:27:31] enterprise who wants to do a backup. Now, they're asking for backups within the same country, for example. Same thing in Asia. They would like to keep their data. Some countries would like to keep their data only inside. Whereas, Google as a company that has been designed wants to take care of your data in past, replicating data in as many regions and zones with as many copies as possible so that we give you very reliable, very resilient infrastructure. But these requirements which are required now to keep your data within the same country or within certain boundaries and requirements, from a product side, we have think about requirement and how all the way from identity, to the country and regions and zones, how you virtualize, how you copy, and all of that is networking and security in addition to how we build the infrastructure. There are a lot of requirements and, I would say, pressure in how do we think about security and data guidelines in these geopolitical scenarios.

[00:28:40] JM: Do you think those concerns, the regulatory concerns, are they well-founded or does it feel like in a sense what you're doing there with that work is just sadly not effective and just kind of catering to the unfortunate regulations that have risen?

[00:29:00] LS: It's both. In some cases, in some areas, they're clearly defined, to be honest. Then rules and regulations change, right? But we trust that, in some cases, when our account teams and our partners, they're bringing those requirements, we trust them and we do build it based on that. In some areas, we learned directly that they're still being built, they are still being revised. They're both the categories and we try to work towards both.

[00:29:29] JM: Have machine learning workloads introduced new networking challenges for you?

[00:29:33] LS: I would say actually it has really, really helped with the products that I described like Network Intelligence Center and other intelligent products we are building. Machine learning has really helped a lot. For folks like me who have been in network, enterprise networking and networking security for all my life. One of the things that is always challenging, is always challenging for network folks, is that the cost of managing infrastructure, the cost of managing appliances like firewall, IDS, IPS and how do you replicate your data and how do you manage. There are lots of cost of involved and a lot of products that are used for managing a network and building your network.

The cost and also the capacity, and there is a capacity involved with this. With machine learning, we can optimize the usage, we can make recommendations about what the right configuration would be for your set target OPX, right? It can be so much valuable to network operators and networking teams and they love to see where my costs are going in, let's say, 7 years. But at the same time, do I have write capacity to manage my applications? Do I have visibility also at the same time in what you're suggesting and what you're doing with my configurations? Giving the visibility, but at the same time, making sure that you're giving them right recommendations to optimize a network and infrastructure while they can still operate at the same availability and reliability is very key. Machine learning has helped us a lot with that in that space.

[00:31:10] JM: On the data collection side, much of your work is looking at logs and telemetry to understand what could be improved across the infrastructure, what's the process of aggregating and collecting data across the networking infrastructure and how do you figure out what to do with all that data?

[00:31:30] LS: There are multiple levels at which we do and how we expose. I will take examples of our products. For any customers that come in, we give them options to enable their logs, right? We ask the customers when they to enable like what kind of logging do you want. There are logs like VPC logs and there are VPC logs. There are logs from firewall. There are

logs from net. We expose all different kinds of logs after customers can go and click those options in [inaudible 00:32:01] UI or through APIs.

Customers have access to the logs so they can aggregate those logs with options to buy integrating third-parties. There are a lot of CM companies that are companies like network performance monitoring diagnostic space, like Sumo Logic, and there are companies like [inaudible 00:32:19]. You can go and put those products and use them to aggregate that information. Splunk in another one of the products, which can give you insights through your logs. That's one way of doing it.

Another way of doing and aggregating your logs is that you can use stack driver, where we send all the logs. There are lot metadata with a lot of telemetry information. Then you can create your custom views of those logs and then you can understand how your infrastructure is doing.

Then other thing that we are doing is the product that I described, is that through Network Intelligence Center, we can use passive monitoring through your infrastructure or through your cloud infrastructure, we do active monitoring of your infrastructure. We can do a data plane level, data plane level telemetry, and that we expose using this product that I described. Network Intelligence Center, and then you use that telemetry to get your network topology information, latency information, packet loss and whatnot.

Then there's other aspect of it, which is config-based modeling, what people call like intent-based networking. Our customers had been asking that since we moved to cloud, I don't know, is this the right configuration that I had? Because on-prem, I had full control. I have full visibility. I had full team operating my configuration. In cloud, you have reduced it to so many few lines. I don't know whether this is the right configuration as I intended to be. Also, our time as I'm scaling, I'm adding more policies. I'm adding more routes and I'm adding more regions and more services. How do I know whether the configuration that I have is similar to what I intended?

Then we have what we call config-based modeling. We model the entire configuration, and based on the current configuration of the customer, which includes your hybrid connectivity as well. Then we expose that config model to the customers. There are multiple ways in which

customers have a visibility into their logs, into their data plane level telemetry information and through third-parties and our own native tools.

[SPONSOR MESSAGE]

[00:34:37] JM: If you are selling enterprise software, you want to be able to deliver that software to every kind of customer. Some enterprises are hosted on-prem. Some enterprises are on AWS. There might be a different cloud provider they use entirely, and you want to be able to deliver to all of these kinds of enterprises.

Gravity is a product for delivering software to any of these kinds of potential environments or data centers that your customers might want to run applications in. You can think of Gravity as something that you use to copy+paste entire production environments across clouds and data centers. It puts a bubble of consistency around your applications so that you can write it once and deploy it anywhere. Gravity is open source so you can look into the code and understand how it works.

Gravity is trusted by leading companies, including MuleSoft, Splunk and Anaconda. You can go to gravitational.com/sedaily to try Gravity Enterprise free for 60 days. That's gravitational.com/sedaily to find out how applications can run the way that your customers expect in their preferred data center. That's gravitational.com/sedaily.

Thanks to the team behind Gravity, the company Gravitational, for being a sponsor of Software Engineering Daily.

[INTERVIEW CONTINUED]

[00:36:02] JM: I'd like to talk through a case study. When I think about all of these big enterprises that are adapting cloud and they're becoming hybrid cloud or maybe they're moving entirely to the cloud, I'd like to know what that experience is like. What is the migration process like? What kinds of issues do they encounter? What is life like before and after a cloud migration for a big enterprise?

[00:36:30] LS: Sure. Without naming a customer, I will talk about the journey. This customer is a retail customer. They migrated their application in multiple. We use [inaudible 00:36:43] for how an enterprise will categorize the application that they want to migrate and something with a priority going up first in a sense they first minded, which are non-critical, not so much business impacting to something which is really keep the lights on for the business, right?

The first ways would be typically I want to migrate my, say, catalog services, my catalog and item services. Then they migrate, then how people will search through that catalog and items. Let's say that is one set of service. I'm taking some very simple example. Let's say those were VMs sitting on-prem and they want to migrate to VM sitting in the cloud. We will talk about VM to VM migration scenario.

In this scenario, what are the things that they will have to think about when it comes to networking? They will have to think about if this VM, what is the instance size? What is the connectivity that I need for cloud? Is it IP site VPN? Do I need a very dedicated Impala circuit? Do I need to go to my managed service provider or network service provider? Need to buy-in to connect, or whatnot?

But among all of them, the most important thing for them is what is the reliability for that link, right? If that link goes down for some reason, how do I reach out to this cloud provider? Who do I call and how do I make sure that it comes up? It comes up in the time and SLAs that I have given to my ops team. There are SLAs involved for connectivity. There are SLOs involved for connectivity. In addition to connectivity, then the biggest, biggest thing that you have is DNS. Without DNS, you really cannot reach any site. Then there are products that you think about, like from cloud mapping perspective, customers will have to think about how they are configuring DNS on-prem. Where are they putting the resolver? Where are they putting forwarder?

Once you do that, all of that, then you'd also think about like where is my database of all these services, which means like some kind of registry, some kind of directory. Then they will have a product like let's say service directory. It's a product that we are just going beta. It is already in Alpha, which we'll have an information about all kind of endpoints, VMs or any kind of service

load balancer endpoints into that registry or into that database so that you keep adding your catalogue of services.

What kind of challenges then come? Challenges would be – So you're migrating. You're deciding to go with three VPN links or one interconnect, and while migrating you saw a surge in your access of the services and certainly your pipe was completely filled. Then, initially, you thought this was going to take, let's say, a day or two days. Now you don't even know why that surge happened, for example. Then second is that your link is down. Your full connectivity is down. Instead of one or two days, it took you maybe 7 days, or 8 days.

What are the challenges involved here and how would you fix it? What our customers had been asking, so while I'm migrating – So during migration, can you give me a view of utilization of my current link and can you suggest to me as this link utilization goes beyond a certain threshold, do send me a notification, an alert, so that I can create a new VPN internal or I can request a new connectivity.

Customers typically who migrate, they ask this question all the time, so for which like we do have product animation, the network intelligence center with the connectivity would be able to help you. What it cannot do, for example, today, what we are building, is that how do you give an alert so that alert is also in an automated way setting up a tunnel for you. It's like – Because you're migrating and things don't happen, right? You don't know the utilization, right? Then you can't really do things.

Now the second aspect, you are used to accessing your services. Let's say you're already using some GCS, which is a storage service. You are using BigQuery, which is our data analytic service and you are using them through public APIs. But now you have a private set of services in cloud, which are on VPC. You're using RFC 1918 space, but now your challenge is that how do I go from my private space to this public IP address? For that, then you would like to create a similar private boundary for these services so that there is no exfiltration of data. There is no public access or user coming from a public IP address now getting into a private service space and exfiltrating information. Then we created a product like VPC service control where you can create a parameter of a bunch of these services together, so your BigQuery, and GCS, and

CloudSQL, and then VN, your VNs, all of them together so that you are 100% sure of not losing your data inadvertently to any user who's coming from.

Security challenges, design challenges related to that, and also do I know whether I have enough – My utilization at any point is going up and down. How do I map my DNS entry so that I'm not losing access to any of the commerce or any of those services. Those are a common set of challenges, I would say, that we hear and run into, especially for VM to VM. To be honest, they apply to VM to microservices environment too. In fact, in case of microservices, if you're going from a monolithic application to microservices, the scale of the challenge becomes a little higher, because now you have many containers. You are optimizing them to be running on one host or across hosts because of how you want to design your applications. Then you need to be more careful about your inter-zone application design, your inter-region design, all of that.

[00:42:28] JM: How much of the product development for the networking team is driven by these interactions with large enterprise? Is it mostly like you have the correct solutions and you just need to get them correctly configured and deployed for a large enterprise, or do particular customer use cases lead to the development of brand-new products?

[00:42:51] LS: I love this question. Now I would say that – Now means like in 2020. We are enterprise ready. What that means, like almost all verticals, let's say, we have prioritized or we're working, whether it is retail, healthcare, education, manufacturing, we can serve all of them for most of the demand. Let's say any kind of demand today, any known demand today.

Yes, we are ready. We can help them and we can leaders especially in like transformation with microservices and also with the introduction of [inaudible 00:43:27] tools, we have brought in both our power of GCE, which is VM, and containers together in seamless manner. So, absolutely.

But then there are still new things that we learn all the time, right? Let me take an example of, let's say, manufacturing. You talk about anything people think about manufacturing, or let's say construction, construction companies. These two industries – And logistics. In my opinion, these three have disrupted over past 3, 4 years a lot because how cloud is evolving, because of how like retail is evolving. While these verticals – And automotive as well, right? While these verticals

have transformed from their current model of serving their users to a completely new model, like how logistics happen and how automotives are and how drones are and how logistics, like people use all kinds of means of delivering, say, delivering the packages to home and other services. In the case of manufacturing, people used to think only IoT as the way of getting into manufacturing, but all these devices are being managed somewhere. A lot of them are coming to cloud.

We learned a lot from these new kinds of customers and new kinds of industries who are disrupting. Some of those things, I would say, at Google we are really best suited to play in because of our global reach, because of how distributed we are in terms of our application design. Our cloud networking design, like we have Andromeda, we have Espresso. I would say we are truly innovators and also implements for enterprise-based applications for these globally-distributed massive scale applications, which are either running parallel at the same time, or in many, they are cohesive set of services talking to each other. Yes, we learned, and we continued to develop and build more for serving them better.

[00:45:28] JM: All right. Well, just one more question on management, because you direct product management for the networking team, and I'd be very curious to understand your perspective on leadership. What's the most important lesson in leadership that you've learned in the last year?

[00:45:44] LS: I'll speak because it's last year I will only talk about my experience in Google Cloud. In past 1-1/2 year, I would say we have scaled our response to customers. We were building products. We were building really awesome products prior to that and we had a lot of customers, but we were in the cloud native leadership space. We were seen and perceived as a company who was cloud native first. But we wanted to be both first enterprise and cloud native.

When you do that transformation, that transformation of what kind of customers are you addressing also changes your team dynamics. Also changes the people you bring in the team, their experiences. Then having people who can serve cloud native, and I'm talking about product management and product management leaders. In Google, a product manager is considered as the CEO of the product. How a CEO of a cloud native company will think? Let's say, a very large gaming company will think versus how a CEO of a very large, very established

in hundreds of years, maybe that's the amount of time they have taken to be where they are today. When you're talking about these two sets of markets, then you need CEOs of the products, the product management leaders who can think like that, right?

I think from a leadership perspective, what I've realized is that what kind of people do you hire and how do you bring that knowledge together so that you can serve your customers the best way and you do not look like two different organizations at all. You are going together. You're learning from each other and you're serving your customers collectively, learning from experiences of each other. How do you drive that? As a leader, how do you break those siloes? How do you bring that knowledge to rest of the organization? Because we are the people like who can bring and converge those words for Google and its customers. How we do that, how we evangelize that and bring those synergies is where I spent a lot of time.

[00:47:53] LS: Lakshmi, thank you for coming on the show. It's been great talking to you.

[00:47:55] JM: Thank you so much.

[END OF INTERVIEW]

[00:48:05] JM: Today's show is sponsored by Datadog, a modern, full-stack monitoring platform for cloud infrastructure, applications, logs and metrics all in one place. Use Datadog's rich, customizable dashboards to monitor, correlate, visualize and alert on data from disparate devices and cloud backends to have full visibility into performance.

Datadog breaks down the silos within an organization's teams and removes blind spots that could cause potential downtime. With more than 250 integrations, Datadog makes it easy to collaborate together and monitor every layer of their stack within a single platform.

Try monitoring with Datadog today using a free 14-day trial at softwareengineeringdaily.com/datadog, and they'll send you a free t-shirt. That's softwareengineeringdaily.com/datadog. You sign up and you get a free t-shirt. Check it out at softwareengineeringdaily.com/datadog.

[END]

