EPISODE 821

[INTRODUCTION]

[0:00:00.3] JM: Intricately is a company that maps the breadth and depth of cloud infrastructure usage. Using a combination of clever algorithms, data engineering and web crawlers, Intricately derives information about how different companies spend money on infrastructure.

Fima Leshinsky is the CEO and Co-Founder at Intricately. In his previous job at Akamai, he began to study how cloud providers could figure out how much their competitors were charging certain customers. For example, if I'm Akamai, I'm a CDN. A CDN is a commodity with reasonably low-switching cost. Maybe I can convince a large customer of a rival CDN to switch over to Akamai if I know what that customer is paying. I can just offer to charge them less. Of course, the question remains, how do I figure out how much potential customers are spending on my competitors?

From his work at Akamai, Fima felt that there was a market opportunity to provide data services to the broader market of cloud providers. There are more cloud providers than ever before and the data that Intricately aggregates is highly useful to this competitive marketplace. Fima joins the show to talk about the modern landscape of cloud providers and how to build a system that maps the internet.

A few updates from Software Engineering Daily land, the FindCollabs first hackathon has ended. Congrats to Arhythm, kit space and rivaly for winning the first hackathon prizes. They got $4,000, $1,000 and a set of SE Daily hoodies respectively. The most valuable feedback award, as well as the most helpful community member award were both won by Vince Montgomery, who will receive the SE Daily towel and the SE Daily old school bucket hat. You can actually get your own if you want one of those. Those are linked too in the show notes along with everything else I'm mentioning.

FindCollabs which is the company I'm building is hiring a React developer. The details are in the show notes. A new version of Software Daily, our open source app and ad-free subscription service is available at softwaredaily.com. You can also contribute to the project on our

FindCollabs or on our github. Pod Sheets is our open source set of tools for managing podcasts and podcast businesses. We are looking for contributors that are interested in contributing to that podcast platform.

With that said, let's get on with today's episode.

[SPONSOR MESSAGE]

[0:02:46.4] JM: DigitalOcean is a reliable, easy-to-use cloud provider. I've used DigitalOcean for years, whenever I want to get an application off the ground quickly. I've always loved the focus on user experience, the great documentation and the simple user interface. More and more people are finding out about DigitalOcean and realizing that DigitalOcean is perfect for their application workloads.

This year, DigitalOcean is making that even easier with new node types. A $15 flexible droplet that can mix and match different configurations of CPU and RAM to get the perfect amount of resources for your application. There are also CPU-optimized droplets perfect for highly active frontend servers, or CICD workloads.

Running on the cloud can get expensive, which is why DigitalOcean makes it easy to choose the right size instance. The prices on standard instances have gone down too. You can check out all their new deals by going to do.co/sedaily. As a bonus to our listeners, you will get a $100 in credit to use over 60 days. That's a lot of money to experiment with.

You can make a $100 go pretty far on DigitalOcean. You can use the credit for hosting, or infrastructure and that includes load balancers, object storage, DigitalOcean spaces is a great new product that provides object storage, and of course computation. Get your free $100 credit at do.co/sedaily. Thanks to DigitalOcean for being a sponsor.

The co-founder of DigitalOcean Moisey Uretsky was one of the first people I interviewed and his interview was really inspirational for me, so I've always thought of DigitalOcean as a pretty inspirational company. Thank you, DigitalOcean.

[INTERVIEW]

[0:04:53.3] JM: Fima Leshinsky, you are the Founder and CEO of Intricately. Welcome to Software Engineering Daily.

[0:04:58.2] FL: Thanks, Jeff. Happy to be here.

[0:05:00.8] JM: You map the internet. Why would you want to do something like that?

[0:05:04.8] FL: Yeah, right. It's interesting, I first stumbled upon the opportunity, or the idea while I was at Akamai. I was actually working with Apple to help them scale their iTunes infrastructure. At the time, the sales team at Akamai rushed over to the engineering side of the house and we're preparing for a sales conversation with Apple trying to get reconnaissance on what infrastructure providers they were using in Southeast Asia.

I thought, “Wow, this is pretty cool. Okay, let's see if we can figure that out.” Next week, the same question came up, but this time they wanted to understand who was powering the same iTunes cloud in Australia? As the years went by, this just kept happening; sales and marketing appeared to be so heavily reliant on engineering to give them that visibility and prepare them for customer conversations. I thought, “Well, someone's building this.” This has got to be a better way to get that visibility in a more consistent and repeatable fashion. Sure enough, no one was, or at least no one that I could see, and so I thought what an interesting engineering challenge.

It's one thing to scrape and pull down HTML and JavaScript from websites to understand what the front door has, but an entirely different challenge to map out what's the backend infrastructure that's powering some of these huge companies? That was the very beginning of this idea.

[0:06:39.9] JM: This is a nuanced area. I think, we should actually go a little bit deeper into why you need to map the internet. Actually first, first I want to just address what you just said. Basically, the problem that you encountered was you were at Akamai. Akamai is arguably a cloud provider. I mean, it is a cloud provider. It's a CDN. It’s servers in the cloud, which they're

making these gigantic enterprise sales to companies like Apple, or companies like Netflix. Those contracts are for so much money.

If you're a salesperson at Akamai and you're trying to sell CDN support to Apple, first of all, they have no idea what is the ceiling on the price that Apple is willing to pay for that cloud hosting CDN technology. Second of all, you don't know what they're already – who they're already purchasing from, you don't know what they're sending, spending on those people. If you're a salesperson, you'd like to have some insight on these numbers before you engage in the sale, because the number at which a sale begins is extremely important. It is this anchor price that you're going to be entering into the conversation with. All the negotiation takes place around the anchor price and you got to figure out what that anchor is. That said, why do you need to map the internet in order to help out with this problem?

[0:08:08.6] FL: Yeah, yeah. You nailed it, right? There is big dollars at stake. We're talking about millions of dollars. Akamai is not the only company that struggles with this. It's really any company that's selling an enterprise class product that's a digital product; digital, or physical for that matter, right? Anyone powering digital delivery, whether that's storage, network, compute, or actually CDN has this this problem.

Is your question why do you need to map the internet in order to get to that answer? Because I think the business problem you've laid out clearly, right? You have to have that visibility to have a optimized sales conversation. How do you go about doing that and how do you do that at scale?

First things first, it's really what we refer to as entity mapping, which is understanding the relationship between things on the internet and actually businesses that are operating them. When we talk about mapping the internet, that's actually what it's referring to.

[0:09:15.9] JM: Describe what is involved in mapping the internet.

[0:09:20.2] FL: Yeah, right. Entity mapping is really about identifying digital or physical things at establishing a relationship to a business entity. Really, these digital things that we are monitoring and mapping are applications. It's any piece of software that's running somewhere that we can

tickle, touch, or see in some in some fashion. Step one is really just cataloging all of those things and trying to understand who's operating and who owns them.

Step two becomes really interesting, which is understanding their dependencies. Once you've identified that an application is there, and maybe you know who who's operating it, maybe you don't, but you want to try and understand what's powering that application? What are the dependencies that application has in order to deliver the experience that it was designed to deliver? Those are the two core pieces to mapping the internet as you put it.

[0:10:23.5] JM: A map can contain addresses in the physical sense. On the internet, we obviously have web addresses. We have URLs. Those URLs map to DNS entries. We can do lookups on those DNS entries, what exists at that lookup. We can we can ping those DNS addresses. We can do WhoIs lookups on those addresses. Given that the span of operations that you can make on any web address, how much information can you infer from those different operations that you can operate on address using?

[0:11:02.2] FL: Yeah. Now you're really unpacking this process. It turns out there's a tremendous amount of detail that you could gather from this type of interrogation, or scanning, or what have you. The thing that I really was turned on to at Akamai was that all of these different probes that you might said, whether that's a Whois lookup, or a ICMP ping, their response, the way that that probe might behave will depend on where you're sending the probe from, what time of day you're sending it from.

Understanding that oh, wow, if I really want to get a comprehensive view into the operation of an application or the behavior of an application, I actually have to inspect it from many different locations around the world. That really shed light on just the richness of data that you could collect. Really Whois and pings and DNS lookups, these are just scratching the surface of the analysis you can perform on an application, or an IP address.

[0:12:11.4] JM: Okay. Well, what else could you do?

[0:12:15.0] FL: Sure. For some applications, you can actually load them in a runtime environment and see how they behave, right? A web application, very simply you can load in a

headless browser, see what calls it's making. You can bypass DNS entirely and just start to examine IP space. Of course, you want to be a good citizen on the internet.

The good news is that there's a lot of very bad citizens on the internet. As long as you behave better than them, you can begin to map out from a TCP/IP standpoint, what port services and applications are running on IP space. Different vectors can get you to in sometimes the same place, but other times they really complement one another. We try to be as creative as possible in this mapping approach and try to identify as many endpoints and applications as we can.

[0:13:13.5] JM: How does this engineering problem compare to that of a web crawler?

[0:13:19.0] FL: Well, there's definitely some similarities, right? I think that a web crawler is probably the first place where people might look to start the process. A web crawler is really limited to HTTP-based applications. If you really want to understand what does, let's say Facebook's infrastructure look like, that's going to be scratching the surface. You really want to try and get deeper, and so lots of applications are going to be running – well, tension on different ports, but they might not understand HTTP at all. Yeah, I think there's a lot of similarities, but it's one data collection vector and there's many others.

[0:14:06.6] JM: Your task is similar to Qadium, or now it's called Expanse. We had them on the show in the past and they're mapping the internet to understand today, I guess more the physical and security infrastructure of what is commonly called – I think it's a – what is it called? Infrastructure sprawl. They're trying to help companies map what infrastructure they actually have. They seem to be concerned more with security than cost management. Do you have any insight into their approach and how your approach to mapping the internet differs from theirs?

[0:14:51.7] FL: Yeah. It's so interesting. There's actually I think a growing number of companies that are collecting very similar data, but using them for different applications, different business applications. I'm not familiar with Qadium specifically and how they're doing what they're doing, but I can tell you that I imagine, we're probably building very similar entity maps. We are also monitoring every company's infrastructure around the world, how it's changing, what different security solutions they have in place.

I could see how if we had the business acumen, or saw the opportunity, we could take the data platform that we're building. Instead of helping sales and marketing teams operate smarter, potentially we could build a product for security researchers, or IT ops infrastructure, management professionals. There's companies like BitSight, who are using very similar data to build credit scores for businesses, based on their exposure from an infrastructure standpoint. I'm sure, it's much more involved than that.

Yeah, it's really exciting to hear the different uses for what is underlying, like infrastructure or digital data collected from more or less public endpoints. We don't have any proprietary access to any of these companies' networks. We're just deployed in a highly distributed fashion and have tapped into all of the different wavelengths, if you will, that exists on the internet.

[0:16:43.8] JM: The end result of all this data crawling and data munging is well, I mean, one example is there's this article on your blog about how much money Netflix spends on AWS. This is basically inferred from this crawling of the internet that you're doing, this paying of the internet and prodding and basically deriving from external public sources how much money Amazon, or AWS is making. Well, not necessarily making, but how much money is being spent by Netflix on AWS. You came to the conclusion that Netflix spends around 25 million dollars per month on AWS. How did you infer that and how do you feel confident in that assessment?

[0:17:33.0] FL: Yeah, that's a great question and a really difficult problem. The way we got here was we had all of this low-level data around the breath of deployment that a particular product had, the amount of traffic that we saw; those different points of presence, or product deployments have. We were trying to figure out what's the right way to communicate this to a business person.

The obvious answer was well, let's try to estimate, let's try to estimate spend, right? Let's provide a dollar value, because that's the easiest thing to understand if you're in a sales or marketing team. It's next to I would say, I wouldn't say impossible, but it's very difficult to get something like that, especially when you're trying to do it for companies like Netflix, or Apple, or Microsoft that really would break every model you have.

Our spend estimates are intended to be directional. They are meant to give you a sense of investment, or size. We know for a fact that Netflix is actually spending way more money on AWS than that. The way that a marketing team or a sales team use our spend estimates is they'll bucket them and they'll say, "Hey, whether they're spending a 100 million dollars a month or five million dollars a month, that's a huge amount of spend." We want to bucket them in the huge enterprise category. Let's have our enterprise sales team reach out to them.

Hey, this other company that Intricately has provided a spend estimate on, well, they think it's $1,000 a month, or $10,000 a month. Okay, let's give that account to our lead gen, or outbound marketing team and let's treat them a little differently. It's really about arming sales and marketing to do account potential modeling, lead scoring. The level of precision there isn't quite the level that you might need for say some other type of use case.

[SPONSOR MESSAGE]

[0:20:00.1] JM: When a rider calls a car using a ride sharing service, there are hundreds of back-end services involved in fulfilling that request. Distributed tracing allows the developers at the ride-sharing company to see how requests travel through all the stages of the network. From the front-end layer to the application middleware to the backend core data services, distributed tracing can be used to understand how long a complex request is taking at each of these stages, so the developers can debug their complex application and improve performance issues.

LightStep is a company built around distributed tracing and modern observability. LightStep answers questions and diagnoses anomalies in mobile applications, monoliths and microservices. At lightstep.com/sedaily, you can get started with LightStep tracing and get a free t-shirt. This comfortable, well-fitting t-shirt says, "Distributed tracing is fun," which is a quote that you may find yourself saying once you are improving the latency of your multi-service requests.

LightStep allows you to analyze every transaction that your users engage in. You can measure performance where it matters and you can find the root cause of your problems. LightStep was founded by Ben Sigelman who was a previous guest on Software Engineering Daily. In that

show, he talked about his early development of distributed tracing at Google. I recommend going back and giving that episode a listen if you haven't heard it.

If you want to try distributed tracing for free, you can use LightStep and get a free t-shirt. Go to lightstep.com/sedaily. Companies such as Lyft, Twilio and GitHub all use LightStep to observe their systems and improve their product quality.

Thanks to LightStep for being a sponsor of Software Engineering Daily. You can support the show by going to lightstep.com/sedaily.

[INTERVIEW CONTINUED]

[0:22:12.0] JM: Let's go through a few examples. You have in this category, in this report on a Netflix spend, like one category that they spend money on is a content delivery network. They spend some money on AWS cloud front and they spend some money on Akamai. How would you figure out how much money, or how would you infer how much money Netflix spends on a content delivery network?

[0:22:40.8] FL: Yeah, absolutely. What we would do is and this is the general approach to how we estimate spends. We're looking at three, maybe four things the biggest influencer of spend, especially when it comes to a delivery product like a CDN is traffic. Ultimately, it's the amount of usage that that product is seeing. We have actually made a tremendous investment in trying to understand what does traffic for an application look like.

There's many third parties that collect this data and variety of ways. Alexa and other companies provide traffic data at the domain level. It turns out you can actually lift traffic data off of the public internet from analyzing public caching infrastructure. We augment the third party traffic data we collect with our own first party data. Traffic is a really key component to understanding spend.

Beyond traffic, you have to also understand the breadth of deployment. How much Akamai, or cloud front is being used in terms of the number of applications that have it enabled? Then how

is it being enabled? Because that can also influence spend, the type of integration that is occurring with that CDN in particular. Breadth of deployment is a important contributor to spend.

Then lastly, we maintain internally a rate card of sorts, where we know that certain providers charge more. Their products are more expensive and depending on the region that you're receiving that delivery from, you might be paying a premium as well. I'll give you a concrete example of that. If you want to do business in China and you want to improve the experience for your Chinese market, it's going to cost you significantly more money. In fact, companies like Akamai offer a separate dedicated delivery service to deliver inside of China. All of that factors into a spend estimate.

[0:24:55.4] JM: Let's say here there's a problem of breadth and there's a problem of the depth of spend. The breadth is how many instances are there. How many instances of AWS, elastic load balancing does Netflix use and then how much are they spending on each of those instances? It sounds like the breadth part of it is a little bit easier, just finding how many instances of ELBs are publicly pingable on the public internet. That seems an easier problem, than figuring out how much consumption of resources is going through those ELBs. Let's go on the on the depth side of things. How do you figure out how much data is going through the Amazon Elastic load balancers on Netflix's infrastructure?

[0:25:50.7] FL: Yeah. Now you're talking about probably the most challenging engineering problem that we have, which is understanding usage and understanding traffic. As I mentioned, so actually just taking a step back, traffic is one of those things that's really difficult to get visibility into, unless you are somehow in network.

Historically, the way that the rest of the world has gotten any ability to traffic is through products like Comscore, Nielsen or Alexa, where typically the way that traffic data is being collected is through a panel. That includes, or involves typically a browser plug-in, like the Alexa toolbar, or some piece of software that is sitting on an end-user's machine and tracking their browsing activity, reporting back. If you have enough of these user's reporting back data, you can start to get a sense for how much traffic at digital property, or a website is seeing.

Over time, as the rise of connected devices, gaming, mobile has grown, these approaches become less and less effective. There's a huge percentage of traffic that the traditional way of collecting traffic data just does not pick up. This forced us to build our own traffic measurement apparatus.

Conceptually, what we're able to do is monitor public hashing infrastructure and use the cache population of public caching infrastructure to understand the popularity of individual pieces of content that belonged to obviously an application, and infer the amount of usage or traffic that application sees in a particular location around the world.

We use that data to give us a sense of depth, to give us a sense of traffic or usage. Again, it's on a regional basis. You can start to imagine how if you know that say Netflix has a lot of ELB usage in say Russia, as for example, then okay, you can start to assign or allocate spend based on that distribution of usage.

[0:28:24.7] JM: What is public caching infrastructure?

[0:28:28.4] FL: Yeah. Most applications today rely on some form of caching, whether that's two categories here. CDNs are the obvious caching infrastructure that we're used to. DNS resolvers are another really important piece of caching infrastructure. When we load up a URL in our browser, as you mentioned at the front of the conversation, one of the first things that the browser does is it tries to resolve that host name to an IP address that your browser can then connect to.

Well, it does that using your local DNS resolver. There's a cache there. It operates exactly like a CDN cache does. These are examples of public caching infrastructure. We have instrumented that piece of the internet to give us a sense of – for that piece of content in that region how frequently is it being accessed and how does that map to overall usage?

[0:29:32.5] JM: Tell me about the instrumentation of the public caching infrastructure.

[0:29:38.4] FL: Well, effectively what you're doing is you have to check to see is that content actually actively being requested in region? Now we're tiptoeing along the lines of our own IP,

but really it's as simple as that. We're checking caches for specific pieces of content, to see if they've been requested within their TTL window.

[0:30:06.9] JM: When I think about the implementation of DNS infrastructure, it's probably something like an LRU. It's like, you can ask it, "Hey, what's the order of the things in your cache?" Based on how things move up and down in the LRU, you could derive how frequently is a given domain getting accessed.

[0:30:28.0] FL: That's exactly right. Yeah. TTLs inform you of how fresh a piece of content is in the cache. Using LRU as a for example, you can start to understand, "Oh, I requested this piece of content an hour ago and I got a cache miss." Well, that tells you is that for the past hour, there hasn't been any traffic, any demand for that piece of content.

[0:30:55.4] JM: Now, hopefully I haven't ruined your moat there. I don't think I have, because first of all, by the way, brilliant insight. That is a brilliant insight for how to figure out the depth of cloud spend. Tell me about some other tricks that are up your sleeve.

[0:31:13.8] FL: Yeah, and really quick here. For the first couple of years of starting the business, I thought I had to keep all of this so close to my best. The reality is that ideas are really just not worth that much. There's so much more involved in building a business here. Tricks of the trade, this isn't – it's much harder to actually execute and build something like this, especially at scale. Yeah, it is fascinating again to just learn about some of the depth that you can collect from the internet.

Let's see, what else is interesting, right? What other things are we able to do that maybe most people wouldn't be aware of? Another really big problem in the space is entity mapping. Being able to identify all of the different applications or domains that a company operates. If you think about a company like Netflix ,they're operating tens of hundreds of distinct domains. Most of them don't have anything to do with Netflix on the surface.

A problem we spend a lot of time trying to solve is how do we identify all of the different domains that Netflix operates? You mentioned who is. Well, that's an obvious first place to look. In many cases, Whois will not give you the information that you need. There's privacy services that

protect domain owners, identities. It's really easy to obfuscate your Whois records. You have to take a fundamentally different approach to understanding ownership. That was one of the first things that we built at Intricately as well.

You have to triangulate and look at different signatures tied to a domain to try and understand the relationships between those domains. That's another really interesting problem space that we spend a lot of time thinking about.

[0:33:33.7] JM: I'm realizing that this is going to be just like the Qadium show, because I had two and a half pages of questions prepared and we got to a fraction of the interesting discussions to be had, partly because this is such a – the problem has so many interesting engineering elements to it. Not only are you collecting all of this data, you have to do it in a cost-efficient manner, you have to do it in a manner that scales, you have to do all this data engineering in order to actually process this data in a manageable, highly-parallel fashion. Then you have to build this entire UI layer that lets people access this data effectively.

Before we get into some of that other stuff, you have a Chrome extension. There's even more opportunities to gather insights about how people are using various internet resources. Explain what your Chrome extension does.

[0:34:36.2] FL: The Chrome extension has been such a wonderful addition to our UI and has really empowered individual sales people to get spend reporting wherever they are on the web. That was one of the first things that we built, thinking about just ease of use and accessibility, and how can we deliver insights immediately to our sales users?

The Chrome extension is really straightforward. All that it's doing is depending on the page that you're on, depending on the site that you're on, it's pinging back home and when you click the Chrome extension, providing you with an Intricately report on everything that we see about that company.

We want to take a really conservative approach about just the data collection aspect. I think that's perhaps where you originally were headed with the question. We don't really do much

data collection as it pertains to the Chrome extension, because there's really not that much data we can collect. We can collect it in other ways at much higher scale.

[0:35:53.7] JM: Less loss of trust. I mean, if somebody was interested – this is the thing where if you were grabbing data non-discriminantly from it – I mean, first of all, it's just probably unpalatable for a lot of engineers that would otherwise want to work with you. Second of all, I think probably, if you would see that in the network requests perhaps, or maybe not.

[0:36:17.5] FL: Well, so you touched on the biggest reason why we took a conservative approach, which is our customers are the biggest enterprise, cloud and infrastructure companies in the world. Trust is such a critical piece of us building a long-term business here. Including passing, very stringent security reviews. We just chose not to play that game. Again, we can collect a lot of the same information in an asynchronous fashion using our sensor network. That was an easy decision to make.

[0:36:52.9] JM: All right, well I'd like to get into a little bit of the infrastructure that you used to perform this high-throughput, web crawling procedure. Now if I recall the Qadium discussion, they were using basically containerized – Kubernetes containers that were just reaching out into the internet, grabbing a bunch of different pieces of information, and then handing that data to a Kafka queue and then having various stream processing systems work off of that Kafka queue in order to process that data into a meaningful and munched set of data.

I remember when we were talking, that was like I think two years ago, they had not started using functions as a service yet, but this is what's one thing that's interesting about this problem is it's a perfect application for functions as a service, because you have these stateless web crawlers that can just grab some data and then spin down, you can spin them up really easily and cost-efficiently. I just wanted to lay out a bit of a map for the different pieces of infrastructure that I see is relevant to this problem. Tell me about your infrastructure.

[0:38:02.3] FL: Yeah. Fast is a really exciting proposition. It's like the new shiny kid on the block. We're excited to take that for a spin, but we're in the same position as the Qadium team, where we haven't quite taken the plunge there.

In order to do this type of data collection at scale, you not only have to be highly distributed, but you need to be running on many different providers. It can't be just across AWS, or Google Cloud. We're provisioned across 30 different – actually, over 30 different cloud providers around the world, a hundred different locations. We don't use containers nearly as much as we should. Really, we have a tool set that we deployed on these hosts and we hand-rolled basically a command and control center that's directing them in dropping instructions to them to perform a variety of different data collection.

Then we have another system that comes in and picks up the data from them. It's something that as I mentioned, we hand-built. It's been working really well and that's one of the core pieces of the sensor network. We have other pieces of the network that don't require massive distribution and can be more centralized. We'll let them do their sense of data collection, where they don't need to be close to those resources. That's at a high-level what the infrastructure looks like.

It took us a while to get there, because especially when you start looking at getting presence in regions where the major cloud providers don't exist, there's these regional providers that don't really have APIs, or their management portals are sometimes not working, or just the language barrier is very challenging. We spent a considerable amount of time trying to operationalize, being able to spin up resources in these lesser-known regions, so we can get this ability there.

[0:40:27.1] JM: Tell me about data engineering. What's the infrastructure for your data engineering look like?

[0:40:33.9] FL: Sure. We rely heavily on Elasticsearch. A lot of the data engineering that we do is I would say of band. What I mean by that is we'll use different ML frameworks or approaches to building models. We have a separate team that does that. Then to actually implement those models, we use Ruby.

A lot of the actual glue that takes a model and deploys that into production is implemented using Ruby. We've got a collection of actually PostgreSQL, Mongo and Elasticsearch all servicing different aspects of our data infrastructure. A lot of it is also home-rolled, so soup to nuts, it's

Ruby straight to PostgreSQL, or Elasticsearch and that tends to be where the data ends up living.

[0:41:42.3] JM: Is cost management an issue for you at all?

[0:41:45.7] FL: Absolutely. We've taken maybe a path that many startups have. We initially started on Heroku for really everything and very quickly realized that wasn't going to scale, so we moved AWS.

[0:42:02.6] JM: Sorry, real quick. Wouldn't scale in what sense?

[0:42:04.8] FL: Cost-wise. I mean, just from a data storage standpoint, we were using their PostgreSQL manner service. Very quickly, we saw just the amount of data we had to collect and what that was doing to our bill.

[0:42:19.9] JM: It was the PostgreSQL instance that were the PostgreSQL cluster that was the limiting reagent, or was the processing also too cost-intensive?
[0:42:28.2] FL: No. It was really all about data storage, it was all about PostgreSQL. We weren't storing any large objects in PostgreSQL. Just the pointers, the records, like –

[0:42:38.0] JM: The reason I'm asking – I'm really glad to hear that just because I have several apps on Heroku and I'm like, “Oh, no.” It’s a lot, “Please, because I don't have much data,” so it's okay. Anyway, go ahead. Please continue.

[0:42:50.7] FL: Yeah. Heroku is wonderful. I mean –

[0:42:52.4] JM: Exactly. Yeah.

[0:42:55.0] FL: It was such a pleasant experience and really taught us a lot about just infrastructure management and our deployment process. How do you make this really simple for developers? Yeah, just the cost of maintaining a large PostgreSQL cluster did not make sense. We moved it to RDS, on AWS and very quickly discovered the same thing, which was okay, cost are just going out of control.

We then move to DigitalOcean. We manage our own PostgreSQL cluster now. We're really thrilled with DigitalOcean. We're able to keep our cost manageable there. When you're collecting this type of data at scale, you have to take costs into consideration from the very beginning. We designed our – not just data collection, but data processing and data storage in such a way to be sensitive to that.

Really, there's a combination of just basic things like, hey, let's just make sure we compress things before we store them, to if some data can be sampled, let's sample it. Let's be smart about making sure that we're collecting the right amount of data at the right time to power the right data feature. So far, it's worked well for us.

[0:44:10.2] JM: Since you mentioned the eventual migrate – first of all, sorry to hear you had to migrate your core database twice. That sucks. That sounds really unpleasant. You mentioned DigitalOcean and you probably know about the cloud provider market, or then, or as much as anybody I've talked to. What do you see is DigitalOcean's, I guess strengths, or what do you feel is their core competence? Because I feel like – I look at the cloud – I mean, full disclosure, DigitalOcean is a sponsor of the show, but I feel they are this dark horse in the cloud space, but I honestly have no insight into how many customers they have, or what their usage patterns are like. Give me your framing of DigitalOcean in the competitive landscape.

[0:45:00.1] FL: Yeah, sure. Well number one, just go to Intricately and you can see how many customers they have. We give that to you for free. That's an easy one. DigitalOcean is one of those cloud companies that they really get out of your way. I mean, it's such a breath of fresh air when you compare them to the four, or the three big guys in the space, right? Like AWS, Azure, Google Cloud.

I don't know if you've ever attempted to provision anything on Microsoft Azure, but I would encourage you to do that and compare that to DigitalOcean. It's night and day. Number one, I think DigitalOcean's just approach to developer enablement and truly owning this notion that we're going to be embracing simplicity, that's really been the case throughout their product suite. They're wonderful from a economic standpoint, like pricing, especially when compared to an AWS.

Those are the two things that we looked at, which was hey, do our engineers enjoy using DigitalOcean? It's really easy for us to spin up instances and manage them. Then number two, what's the impact from a cost standpoint? They've won on both accounts.

[SPONSOR MESSAGE]

[0:46:30.7] JM: Testing a mobile app is not easy. I know this from experience working on the SE Daily mobile application. We have an iOS client and an Android client and we get bug reports all the time from users that are on operating systems that we did not test. People have old iPhones, there are a thousand different versions of Android. With such a fragmented ecosystem, it's easy for a bug to occur in a system that you didn't test.

Bitbar is a platform for mobile app testing. If you've struggled to get the continuous delivery in your mobile application, check out bitbar.com/sedaily and get a free month of mobile app testing. Bitbar tests your app on real devices, no emulators, no virtual environments. Bitbar has real Android and iOS devices and the Bitbar testing tools integrate with Jenkins, Travis CI and other continuous integration tools.

Check out bitbar.com/sedaily and get a free month of unlimited mobile app testing. Bitbar also has an automated test bot, which is great for exploratory testing without writing a single line of code. You have a mobile app that your customers depend on and you need to test the target devices before your application updates rollout.

Go to bitbar.com/sedaily and find out more about Bitbar. You get a free month of mobile application testing on real devices and that's pretty useful. You can get that deal by going to bitbar.com/sedaily, get real testing on real devices, get help from the automated test bots so that you have some exploratory testing without writing any code. Thanks to Bitbar. You can check out bitbar.com/sedaily to get that free month and to support Software Engineering Daily.

[INTERVIEW CONTINUED]

[0:48:41.1] JM: Not to go deeper on DigitalOcean. Well, despite the fact that they're a sponsor of the show, but I've spoken to Moisey, one of the founders a couple times. He really struck me by his vision. It was a vision that was pretty distinct from other cloud providers. I think that simplicity thing. It's just over the years, like I've seen the scale and the narrative of the other "hyper scale" cloud providers.

It became really hard for me to assess DigitalOcean in the competitive landscape when the other players are – I mean, first of all, they're able to land these tremendous enterprise contracts, because of their feature richness, but that feature richness is obviously the jiu-jitsu downfall of them relative to DigitalOcean. It's very hard for me to again, to assess the relative strengths of them.

I'd love to hear your thoughts on just the competitive differentiators of the different cloud providers. I see it as a very not zero-sum world, a very early days, not zero-sum world, but it is often framed in this competitive landscape mentality. I just love to get your philosophies on the cloud provider competitive landscape.

[0:50:04.3] FL: Sure. The famous quote by Marc Andreessen that software is eating the world.

[0:50:09.5] JM: Yes.

[0:50:10.7] FL: Well, actually it's AWS that's eating the world. That is the biggest insight that we have found from monitoring the cloud war. It's really astonishing. The pace and the size with which AWS is operating and moving from, just like a customer acquisition and customer growth standpoint. Azure and Google Cloud are such distant second and third place competitors. I think that we'll go one by one, right? Azure is interesting, because they have the benefit of a really large existing Microsoft enterprise customer base. I think that they're in a stronger position to capitalize on the market opportunity in front of them; just from that network effect.

I think Google Cloud is interesting and they've seen some shake-up from a leadership standpoint just this quarter. It'll be missing to see where they take their cloud business, but they definitely have a phenomenal brand and I think, especially on the AI side of things, they're doing some wonderful things. It's too early to say, but I think that the big insight is just how far ahead AWS is from everybody else.

Back to just DigitalOcean quickly, I think the big differentiator here when compared to the other providers is I think DigitalOcean is one of the few companies that actually eats their own dog food, right? They're actually using DigitalOcean and that's what makes such a big difference, and building products that developers actually love and can actually embrace, where you can tell that someone built this who's actually having to use this themselves. As opposed to doing customer panels and collecting usage information on how something is being used.

Yeah, I think that fundamentally different approach is what differentiates them. You're right, it's not a zero-sum game. The market is moving faster than any of these providers are able to move on their own. I think there's enough for everybody. Interesting to see in this frothy market, who's actually winning the most? It's pretty clear, AWS is just a complete monster in the space.

[0:52:45.5] JM: Do you have any insight into the second layer cloud provider market? You have the Netlifys and the Herokus and the Zeits of the world. Any sense of how much these things are growing and their unit economics?

[0:53:04.2] FL: It's interesting, where we are right now in the market really focused on the top tier, largest cloud providers. Outside of Heroku, we haven't really and I personally haven't looked at some of the other as you've called it, second layer providers. I'd love to take a look and follow up with you and give you a sense of what we see there, because we definitely have tremendous insight. It's just I can't actually recall when I looked at Netlify's growth, or for that matter, anyone in that space. If you don't mind, I'll take a rain check on that one and I'll –

[0:53:46.3] JM: Let's do it. Let's do a show on third-party assessment of second layer cloud providers.

[0:53:52.9] FL: A 100%.

[0:53:54.7] JM: Okay. Any sense of – this is really getting off topic, but any sense of how, just how profitable AWS is? Do you have any sense of – have you looked into the – thought about the unit economics of their different systems? That's pretty hard to assess, I imagine.

[0:54:11.2] FL: Wow. It's so interesting. We haven't crossed that bridge as yet, but one of the applications that we are really excited to power, right? We've started off really focused on helping sales and marketing teams operate smarter and really focus on operations teams. The financial services community, the investment community really wants to get a view into this specific question, which is not only how is AWS competing with the Azures and the Google Clouds and Oracles of the world, but unit economics, right? Margins. How much are they actually able to make off of all these users?

We just started thinking about that last year. That's a really difficult thing. Talk about a second order problem that is trying to build on top of the first one that you're really challenged by. I don't know if we have a great answer for that just yet. The way that we were thinking about attacking that problem is to actually look at the regions that AWS is servicing customers out of and the types of customers that they're servicing them out of. Yeah, you're really now doing lots of inferences. No great no great answer there just yet.

[0:55:34.0] JM: Of course, because we don't even know. They might be taking a loss on AWS lambda for all we know.

[0:55:40.1] FL: Well, that's exactly what we're seeing happening in the CDN space, where the way that Google – excuse me, AWS is using. CloudFront as a for example, is keeping their prices really low as a way to subsidize their compute and network and get users onto the platform, right? Really yeah, exactly. Very few companies are able to do anything like that, right?

AWS is in this crazy position where they can actually give away some of their products for free, as a way to get you onto the platform. We saw that firsthand when we were first fundraising, where we realized they were giving away a quarter million dollar worth of AWS credits to startups, just so that you could begin to build on top of them. They're doing that at a scale you wouldn't believe. The war that they're fighting, very different from most providers.

[0:56:33.6] JM: Getting even further away from anything that I thought this conversation would focus on, I would say that the thing that concerns me about AWS and Amazon is really the scope and scale and the complexity and the potential, perhaps susceptibility to black swans that this financial gamesmanship might lead to. I mean, I would never want to cast aspersions on

Amazon, or AWS, or anything, but when you start to do stuff like that, like take a loss on a highly complex product, does that ever come to mind to you? Does that concern ever come to mind to you? Like too big to fail thing?

[0:57:17.3] FL: It's interesting, one of the trends that we're seeing pick up is enterprises' use of multi-cloud strategies, multi-cloud provider strategies. That was always a thing in the upper echelon of enterprises. We're now seeing that move down market, where even SMBs, right? Companies in the mid-market are going with two or sometimes even three cloud providers.

Now, they might be doing it for different reasons. Not all of them are doing it, because they're concerned with respect to putting all their eggs in one basket, maybe out of necessity. I think it's actually – I mean, I'm taking a slightly different angle here.

[0:58:01.9] JM: Yes, you are. That's okay.

[0:58:03.6] FL: I think it's actually beneficial to the overall health of the market, where arising tide lifts all boats. I don't know enough about AWS as business specifically in this regard. Again, we're inferring based on you look at the government contracts, or the government relationships that they have, where some of the most conservative parts of the market have hitched themselves to AWS. You want to trust that they've done their diligence and they're doing the right thing there.
Of course, when you look at it from an outsider's perspective, it is a little alarming, right? I think there's been murmurs of regulation, or splitting up Amazon or any number of ways to try and manage that risk. We'll see what the future holds.

[0:59:01.9] JM: Yeah, yeah. It would be interesting to watch. Your customers from what I understand are cloud providers. I mean, mostly cloud hosting, data centers, content delivery networks. Just briefly, because I know we're almost of time, can you explain who you're actually selling to and how you make money?

[0:59:22.9] FL: Yeah, sure. Fundamentally, we help sales and marketing teams. Sales and marketing teams selling digital infrastructure goods. It's not just cloud or CDN. I don't know if I can announce this. Yeah, I probably can. Yeah. Snowflake is a customer of ours, right? For

those of you that aren't familiar with Snowflake, it's a data platform, right? They're not selling cloud hosting, or infrastructure. It's like AWS Redshift. It's a data warehouse in the cloud.

This is just an example of the types of products that we can help go to market faster, scale and help their marketing teams spend their marketing dollars in smarter ways. We're all about providing a map, or a compass to operations teams to help guide them with respect to where they should deploy sales resources, where they should deploy marketing resources. It's not limited to infrastructure. It's really anyone selling a large, expensive digital product.

It's interesting, for the last 20 years they've all been doing the same thing. They've all been using the same firmographic data to try and size their markets, to try and do territory planning, to build account potential scoring. It really doesn't make sense when you're selling a digital good and you're selling to companies that are cloud native, they're digital companies.

We're coming to the market with a fundamentally different approach to how you would do account planning, or territory segmentation. If you want to be effective in this digital age, you have to use digital consumption data, digital usage. We've boiled that down to a spend number and provide a lot of context on that spend, but we're really helping sales and marketing teams size and go to war in this new digital divide.

[1:01:28.1] JM: Fima, thanks for coming on the show. It's been a real pleasure talking to you.

[1:01:31.0] FL: Jeff, this has been a lot of fun. I'll definitely circle back with you on this second layer question that you have. Yeah, really enjoyed it.

[1:01:38.8] JM: We'll do it. We'll do round two. Thanks Fima.

[1:01:41.3] FL: Looking forward. Thanks, Jeff.

[END OF INTERVIEW]

[1:01:44.6] JM: This podcast is brought to you by wix.com. Build your website quickly with Wix. Wix code unites design features with advanced code capabilities, so you can build data-driven

websites and professional web apps very quickly. You can store and manage unlimited data. You can create hundreds of dynamic pages, you can add repeating layouts, make custom forms, call external APIs and take full control of your site's functionality using Wix code APIs and your own JavaScript. You don't need HTML or CSS.

With Wix code's built-in database and IDE, you've got one-click deployment that instantly updates all the content on your site. Everything is SEO-friendly. What about security and hosting and maintenance? Wix has you covered, so you can spend more time focusing on yourself and your clients.

If you're not a developer, it's not a problem. There is plenty that you can do without writing a line of code, although of course, if you are a developer then you can do much more. You can explore all the resources on the Wix code site to learn more about web development wherever you are in your developer career. You can discover video tutorials, articles, code snippets, API references and a lively forum where you can get advanced tips from Wix code experts.

Check it out for yourself at wix.com/sed. That's wix.com/sed. You can get 10% off your premium plan while developing a website quickly for the web. To get that 10% off the premium plan and support Software Engineering Daily, go to wix.com/sed and see what you could do with Wix code today.

[END]

