EPISODE 1344

[INTRODUCTION]

[00:00:00] KP: Interest in autonomous vehicles dates back to the 1920s. It wasn't until the 80s when the first truly autonomous vehicle prototypes began to appear. The first DARPA Grand Challenge took place in 2004, offering competitors $1 million to complete a 150-mile course through the Mojave Desert. The prize was not claimed. Since then, rapid progress has been made in autonomous driving, fueled by advances in sensor technology, software, and the hardware, which runs it. Infrastructure has become a serious consideration for autonomous vehicle companies. In this episode, I speak with Vinoj Kumar about infrastructure at Cruise, the company helping Walmart do all electric self-driving grocery delivery.

[INTERVIEW]

[00:00:52] KP: Vinoj, Welcome to Software Engineering Daily.

[00:00:56] VK: Thank you, Kyle. It’s my pleasure.

[00:00:58] KP: Well, before we get into infrastructure at Cruise, can you tell me a little bit about your background?

[00:01:04] VK: As a brief introduction, I spent over a decade at Google helping them scale their technical infrastructure that covered compute, storage, networking, machine learning products, data center stuff, all the stuff that goes inside the data center. So that's primarily the last 12 to 14 years I've been dealing with.

[00:01:26] KP: Google is a company that's pretty famous for scaling things and also scaling them in multiple different complex directions. Can you talk about what part of the scalability story you contributed to?

[00:01:37] VK: Indeed, when I came in there, the technical infrastructure group was very tiny, it was in the hundreds. Back then, Google was building their own servers, networking. When I first

came on board, I was working on building what's called Watchtower, the first 160 port 10 gig servers. And that kind of grew. And if you look at – There are some papers published around this. If you look at the epochs of networking, I kind of have played a key role in every single one of them, like the Jupyter networking stuff scaled when and data center networking related products.

[00:02:15] KP: What are some of the major innovations in infrastructure you've observed over that time?

[00:02:20] VK: Yeah, it's about feed and speed. For example, on the networking side of things, it's about going from 10 gig, 100 gigs, and the poor densities as well. And also building compute farms from, say, 10,000 servers to 30,000 servers, connecting campuses and network in the wide area networking side. I saw a lot of innovation there. Like for example, on the artistic side of things. And also on the compute side, going from single socket servers to multi-socket servers that powers cloud, and also into GPUs, and TPUs, because Google's kind of innovating in that area with Tensor for processing units.

[00:02:59] KP: Well, for any listeners not yet familiar with Cruise, tell us a little bit about your products and mission.

[00:03:05] VK: Yeah. Let me start with Cruise's mission. Cruise's mission is around building the world's most advanced self-driving vehicle to transport people and things. It's built around three pillars; safety first, powered by clean energy, right? Protecting the environment, and freedom of movement that is very affordable. To cruise's story, Cruise is building an advanced self-driving vehicle from the ground up.

[00:03:31] KP: I'm reminded of the expression, if you want to create an apple pie from scratch, you must first invent the universe. Where does building it from the ground up begin?

[00:03:41] VK: All the way not only just technology. I mean, technology is a big part of it, right? Being able to conquer it in the most complex environment, for example, Cruise very consciously chose San Francisco as the training or testing ground, right? If you look at San Francisco, it's one of the most dense urban area and a very good testing ground for autonomous vehicles with

pedestrians, cyclists, you name it, right? Even for the most experienced driver, San Francisco is very daunting experience.

So Cruise made a very conscious decision to actually go test in that dense environment so you can learn quicker, have better training models and stuff. So technology is one aspect of that. But again, if you think about it, it is not just technology, right? You have to build trust with users. You have to build this whole ecosystem around autonomous vehicles to transport people. That's what Cruise has been working for the last five to seven years.

[00:04:37] KP: There's interesting work to be done in both hardware and software for this application. Where are some of your main areas of focus for innovation?

[00:04:46] VK: From the hardware point of view, yes, it includes both hardware and – The three components, if you look at it, right? The hardware that powers the autonomous vehicles. And sort of big part of it, the software that actually makes the vehicle move. And the third component or dimension is the infrastructure that supports motion of this vehicle, right?

If you look at from the hardware side of things, it has lots of sensors that actually collect data, which needs to be analyzed and trained. And when the vehicle is in motion, you're going to have to interpret it, right? Perceive, predict and act upon it. So the software stack has to be very resilient. I mean, it has to be cutting edge. Supporting all these is on the infrastructure side of things, where my specialization is, right?

If you think about infrastructure, that infrastructure needs to be built on multiple pillars. One, something that's very efficient, something that scales when you go from five vehicles to 500 vehicles. You go from one city to tens of cities in a very cost effective manner. So that is where a lot of the innovation is combined. If you think of it, it's a multi-dimensional. When you put together cutting edge hardware, cutting edge sensors, with cutting edge software, combined with an infrastructure that supports it.

[00:06:02] KP: All those cutting edge sensors must produce a lot of data, which means there's a lot of things to be processed. Where do you need to be strategic about data management?

[00:06:13] VK: So when it comes to autonomous vehicles, it's all about data movement. On a daily basis, for example, as you have a fleet of vehicles, they have 40 plus nodes or sensors collecting data. And then you ingest around four plus petabytes of data, right? That data has to be backhauled, stored somewhere. And then you analyze – And AV engineers analyze, right? You go through it. You have segments. You look at scenarios and figure out what went wrong? Or how can we improve? And this loop of look, analyze, make changes, run simulations, execute in a simulated environment, or perhaps on the road, and that loop has to be titled low, right?

So what it comes down to all about data movement. So you have petabytes of data. You're running around millions of simulation tasks a day. And the data is stored in various, for example, various backend data storage components within the cloud. The question is, “How do you move that efficiently so that your loop is very tight?” This inner loop of look at the data, analyze, test, execute, is very tight.

[00:07:20] KP: Moving the data is important. Do you also face any challenges giving the right people access to query it?

[00:07:26] VK: This could be a huge bottleneck, right? For example, to give you an example of scale, we leverage BitQuery. And for example, we process in excess of four exabytes per month, right? This is around 20 plus million queries across 500 plus users that we have to process. What this means is this impacts a lot more in the AV engineering iteration cycle, right? The time it takes for an engineer to find out if your changes work, if your code improves the vehicle as intended. And the cycle can take up an entire day or so. The question is how do you mount these simulations? Wait for your turn to use the constraint GPU or cloud pool, analyze the output, or maybe perhaps use a hybrid between automation and hand analysis, and make improvements. So these are some of the challenges that you have to work with.

[00:08:21] KP: GPUs are expensive when you buy them in the cloud. Although I guess they're also expensive to purchase the hardware. But once you've done that, you can run it till the device fails. It's yours after that, although you have to power it. How do you look at decisions around allocation in the cloud versus perhaps doing something on-prem?

[00:08:39] VK: So AV's, in general – Let me take a step back. AV's, in general, involved off the shelf as well as custom hardware components, which means you can have both on-prem and cloud leverage infrastructure, right? When you do need a solid on-prem infrastructure for custom target hardware, and perhaps very low latency, high touch hardware in the loop type of testing, right?

So going back to your question, the way you manage each is very different. With cloud, we don't have to be concerned with physical hardware. It's all about API's and abstractions for everything, right? But it also adds complexity. For cloud providers also, what we've noticed, there are not so quite published contention points that you don't have visibility to. And you have to rely on the cloud provider for assistance. Again, I talked about cloud data movement. For example, when you have data spread around different regions or different subsystems, when you move data to do computation on GPUs or CPUs, these contention points come into play either through egress bandwidth, bottlenecks, or some other throttles that are there, right?

On the flip side, for on-prem, you're very clear about how the data flows, but you have the challenge of managing soup to nuts including supply chain, right? For example, provisioning and managing task could take a very different technology stack. You need to plan space, power well in advance. And given the long lead times for all these, that could be quite challenging.

On the flip side, though, with cloud, you'd need to worry about, for example, multi-region support. Design your system so that it's very resilient if one thing goes down, so you can failover to another subsystem. And for example, figure out where your data and compute needs so that you can – Your simulations or even trainings are very efficient. And most importantly, considering all this is the cost impact. For example, if you don't design your architecture, so to speak, very well in the cloud, you could see our cloud bills go up very significantly very quickly.

[00:10:48] KP: I see a lot of reasons why simulation would be an important part of your process. It's very safe for one thing, which is nice. How does the simulation contribute to the things you're working to improve in the product?

[00:11:01] VK: If you think about in the real world, right? When the fleet dries around, you could have lots of scenarios, nice edge cases or nuggets, like for example. Our bicycle is not yielding,

cutting across. And these type of kind of corner case or edge case is very hard to recreate in real world, kind of to test your scenario, right? That's where, for example, simulations come in handy. For example, you can actually improve upon a situation like that golden nugget, where you have a corner case, and have a whole bunch of tests that's around this type of scenario running simulation. You can actually create kind of these type of scenarios and test them, right? They're not really cost effective. You can actually test cases where it's kind of difficult to do it in real world. So that's one area.

So that is where we spend – Also, if you think about like all the – I mean, Cruise mentioned, has announced Origin, which is a specially designed AV fleet that doesn't have steering wheels. So which means a lot of your testing improvement has to be done in simulation, right? That's where we expect to see in the next couple of years the infrastructure around simulation to go significantly. And that's where in lies the challenge for infrastructure.

For example, your simulations are going to go up because you're going to test a lot more scenarios in the virtual world, which means from the infrastructure point of view, you got to make sure that even your cost structure and your infrastructure is not linear, right? It has to be sub-linear so your bottom line is not affected.

[00:12:32] KP: How much cost forecasting and prediction work do you have to do as you think about budgeting for the future?

[00:12:38] VK: Well, we have to absolutely work with them. It's like a team work. For example, we have dedicated capacity teams. Actually, there are some models you have to build, for example, organic growth in terms of simulations, or how they can grow, right?

On top of that, like I said, architecturally, we can have a lot of influence on what type of tests to run, or what kind of metadata that we choose to pull. Segmentation, for example, like I mentioned, we have petabytes, petabytes of data that's collected on a daily basis. This is where you can actually be intelligent in terms of how to segment it, how to load appropriately, right? And kind of utilize your infrastructure in a very efficient way. So in that sense, my team works very closely with the simulation team, for example, AV engineers, ML engineers, kind of figure out how do we plan capacity? How do we look ahead? What models do we use? And how do

we use a very prudent sensitivity analysis to kind of see where our median points are in terms of utilization, right?

[00:13:37] KP: What are some of the scalability challenges you think a lot about?

[00:13:41] VK: So the biggest challenge is data movement can be a bottleneck. I mentioned this, right? And I also mentioned that we run millions of simulation tests a day, and it gave you a flavor of millions of– It's tens and millions of VMs running these computations. So where it impacts is on the AV engineering iteration cycle. And our goal here is how do we make it frictionless? How do we shorten that cycle?

You can think of data movement is one of the biggest bottleneck, right? As I mentioned earlier. The other bottleneck today is, for example, the global supply chain. For example, as you scale in terms of your own target hardware, number of cities and the cars. Today, that's becoming a huge problem as well. As you scale from R&D to commercialization, managing this global supply chain of both hardware and systems associated with this is a bottleneck.

[00:14:32] KP: What sorts of technologies are in the typical autonomous driving stack.

[00:14:37] VK: Autonomous vehicle involves lots and lots of sensors collecting data, and acting on it. Hence, there is lots of data movement, computation and also backend components, custom components involved.

For example, AV's need to securely connect to backend services, the broker communication to a set of microservices that handle everything from dispatch, remote assistance, mapping, role planning, etc., right? And once the drive is done, there is more data movement. For example, ingesting data into our data lakes for post processing, then you do analysis, continuous improvement and stuff like that.

For these type of stuff, it's the micro service architecture, distribution of these. And also, we rely a lot of custom tools for analysis and visualization of data. For example, Webiz is one that we open sourced as an example, right? So these are some of the components that are in the stack that supports autonomous driving. Again, not to mention the business side of things. When you

have these vehicles roaming around, you need to have a business, a good ecosystem of business components, like CRM tools, and stuff like that. For example, if you're doing ride hail, rideshare, being able to look at doing market analysis, payments, etc., that go alongside this ecosystem of autonomous driving.

And infrastructure needs to be very reliable, right? You're putting cars on the road without drivers behind it. It needs to scale. As you grow the number of vehicles, as you grow into more cities, it needs to scale. And equally importantly, as I mentioned, it needs to be cost effective, otherwise, the business is not very viable, right? So the first and foremost focus is on automation. Historically, infrastructures has been fairly manual ordeal. Now it's all about automation and using software to solve a lot of process and scale problems.

Part of that is focus on API-driven independent services, and leverage open source with Kubernetes, Istio and other type of mesh architectures. And one of the things that, for example, we at Cruise, focus on this latest technology in terms of high scale processing, right? Leverage GPUs, perhaps custom inference engines, and … etc., right? I mean, it's an evolving technology to leverage those.

And the key for all this is knowing to know you know – Kind of knowing when to build versus buy. And you collaborate with large industry trends with open software. Lastly, you're going to need to leverage multi-cloud at some point. And at Cruise, we're doing a lot work with partners such as GCP and Azure to optimize our usage and cost.

[00:17:18] KP: I'm aware of people being interested in multi-cloud for redundancy reasons, and also for cost reasons. Which is interesting to you?

[00:17:25] VK: Both, right? Budgetary is one aspect. The second aspect is also kind of leveraging the strengths of each of the vendors, for example, right? Each cloud vendor has their own unique strengths. Kind of leveraging both is to our advantage.

[00:17:41] KK: It's no surprise that you'd be making good use of GPUs. I'm also aware that there's some interest, just broadly speaking, in specialized hardware. Maybe the automated vehicle industry is probably an exemplar of pushing the limits of certain things. Maybe there

should be custom hardware, or do you think things will standardize on just the standard compute systems?

[00:18:05] VK: Usually what happens in these type of industry, you start with a generalized solution. As you can attack problems, as you stall the general problems, you start getting to these corner cases and differentiated cases where you could potentially end up with a customized solution. Again, this comes down to cost versus what's unique to you, and what's unique to your business plan, so to speak. And that's where you kind of have to make your choice.

[00:18:32] KP: Autonomous vehicles is a up and coming industry. Certainly a good deal of what you're working on would be considered R&D. Can you speak to the role of R&D at Cruise?

[00:18:44] VK: As we move from R&D to commercialization, then it becomes about scale, right? Then you start – There are two things need to happen. When you start, when you are at scale, these development cycle or a kind of fixed cycle, the tight loop that we talked about needs to be sharper, which means the infrastructure needs to support the ML and AI engineers so that the infrastructure is frictionless, very efficient, right? For example, they can quickly analyze the data, make changes, put in the code.

For that to happen, the CI/CD pipeline has to be very efficient. Again, we talked about on-prem versus cloud. For example, your CI/CD pipeline needs to be very abstracted and efficient so that you can support both on-prem and on cloud infrastructures, where you're handling millions of builds. I'll give you an example, when we run simulations, or when vehicles are on the road where all the backend micro-service computation sometimes we run, right? They run into millions of VMs and GPUs, for example, right? Both training as well as perhaps pure computation. For these kinds of stuff, your developer platform needs to be very robust, right?

For example, we want the AI, ML engineers to have best in class tools in terms of compilers, debug tools, visual tools, etc. Infrastructure is not purely focused on the R&D side or supporting the development infrastructure. It plays a key role in supporting the product as well as.

[00:20:12] KP: How do your customers interact with the product? Under what opportunity might I be able to get into a Cruise vehicle?

[00:20:19] VK: The initial goal, like at least from a Cruise point of view, is rideshare, right? Making transportation affordable to everyone. So the focus has been rideshare, moving people, and with our arrangement with Walmart, delivering goods at a very cost affordable cost point. Who knows? It could end up in a driveway near you.

[00:20:41] KP: What are some of the ways you have to give a special consideration for security?

[00:20:45] VK: Great question. We take security and safety very seriously, right. Again, when it comes to security is defense in depth. It's not like one solution that fits all. Like you said, it's about best practices. And not only about best practices, it's about what frameworks are in place to be secure. For example, you're collecting data. Data has to be secure. You need to selectively – As cars move around, they connect to the backend. They talk to a whole bunch of micro services in the back for the automobiles to move. So in that sense, every segment has to be secure. And overall, security is sort of ingrained as a DNA within us at every step of the way. Like I said, it's defense in depth. At every layer, it needs to be secure.

So talking about commercialization, it comes down to capacity at scale, right? So the capacity to acquire, commission and operate sites at scale becomes critical. In that sense, for example, network monitoring and automation is an area. For example, when you monitor different cities, how do you backhaul data through the network? Do you go through public network? Do you have your own backbone? What does it mean to be secure and have a cost effective connectivity? Those are some areas. And tied to that, for example, analytics in terms of data movement. For example, as do data moves around, not only does it have to be secure. Depend on where you ingress and egress, there's a cost associated with that as an example.

As you get into the realm of commercialization from R&D, it needs to be reliable. It needs to stay up. So hence, enhanced visibility into network capacity, connectivity failures, failure detection responses, that is an area we are focusing right now. And a lot of work has been happening from that end.

[00:22:33] KP: At the scale you're describing, it seems you might be testing the limits of some of the cloud providers you're working with. Maybe see edge cases that I as an independent developer would never see at a smaller scale. Do you ever find that you are testing those limits?

[00:22:48] VK: Yes, we do at times. GPU capacity is an example, right? We push to the limits, because the amount of data that you're processing, in terms of analyzing and acting, that's an example where you push the computation capacity of cloud providers sometimes. For example, we use cape to train our models. And as a data point, we process around roughly 30 million frames a day for continuous improvement of our car models.

When you process that amount of data sometimes, and depending on the availability of the GPU, sometimes you push it to the limits. Same thing with computation with compute power as well. If you're running a million simulation tasks a day with tens and millions of VMs, you push it to the limits.

[00:23:37] KP: Well, as you do that, you're investing heavily. So you need to see some return on that investment. What are some of the metrics you look at to determine if you're really well-optimized in allocations?

[00:23:49] VK: Yeah, there are several metrics that we track. There are lots of metrics. For example, from an AV point of view, you kind of measure miles per safety critical event. Or from a stability point of view, how many miles per safe staff, and you measure comfort, as an example from AV point of view? Again, from an infrastructure point of view, something that's very important to us is your total cost of ownership per bill. For example, how long does it take for a developer to make a change, right? Because it translates into efficiency and also cost, right? So we look at those KPIs.

Again, from a developer point of view, we look at how many engineering are saved, right? You could apply those to actually make changes to the AV stack with the savings. So at every point in time, at every stage of the pipeline or the process, there are metrics that we track to make sure we're always striving to improve.

[00:24:42] KP: Do you face any observability challenges?

[00:24:46] VK: Yeah, other use cases in terms of – Again, it's another challenge or another big investment we are making is in terms of observability, since you brought that up. So this is where we want to make sure that the systems are architected such that there's a resilience built and reliability built in. Especially as you move from R&D to commercialization, there is a huge focus on reliability. Systems have to be reliable. Since we have tens, to 20, 30, 40 microservices each dependent on one another, we want to make sure each of these micro services are architected well. So that brings in experts in observability. SREs, we embed them within the engineering teams so they can work with the engineering teams in developing highly reliable and scalable microservice architectures.

[00:25:35] KP: Can you comment on the size of the team and some of the efforts to make them effectively collaborate at scale?

[00:25:41] VK: Yeah, we have hundreds of development, and we have ML and AI engineers working collaboratively. Usually, the cycle is very – It's a very collaborative environment, so to speak, because especially for something that's cutting edge, we want to move quicker. The loop of develop, make changes, build, integrate, roll it out, how do you test? It's a combination of process systems or frameworks in place and, also, the collaboration within the team that we build within. Our term for that called Cruisiness. In other words, developers help one another. If something's blocked, they can actually walk over. Of course, now we're doing virtually, they can Slack over and say, “Can you help me with that?” It's a combination of a common mission or goal towards accomplishing what we want. Supported by tools, infrastructure, processes, and framework is what differentiates in my mind from anywhere else I've been,

[00:26:40] KP: It’s surprise that the players in your industry are tackling will require massive amounts of data. Is it that clean of a competition? Is it whoever has the most data is likely to produce the best products?

[00:26:53] VK: I can't speak to the other players in this area. But I do know what we value. For example, at Cruise, our focus is on quality of data versus quantity of data, not the volume of data. Going back, for example, I mentioned Cruise consciously decided on San Francisco, as

an example, or high dense urban environment where the quality of the data is a lot high. Because you're going to run into in these situations that you wouldn't see in a sparsely populated suburb, right? So our focus has been using that quality of data and then working on the toughest problems so that we can get to the market faster and scale faster as well.

[00:27:34] KP: In terms of a launch, when might someone be in a Cruise vehicle or getting something delivered via Cruise vehicle?

[00:27:41] VK: Cruise has already announced partnership with Walmart. We're actually doing pilots with their delivery. And pretty soon, in months, not in years, that we'll be launching a commercial service in San Francisco. And Cruise already also announced a partnership with Dubai coming out in 2023 as the first international market. So pretty soon you'll be able to hop on to one of those rides within a Cruise vehicle.

[00:28:07] KP: Well, it's very exciting. Are there any particular technical challenges you're working on in preparation for that launch?

[00:28:14] VK: Yeah. This falls into kind of three different – For an operation like this to be successful, it's not just a technical aspect, right? For example, you need to have an ecosystem of commercial subsystems in place, your capacity to scale operated sites. I mentioned networking systems and stuff and the backend systems to be reliable. And also, you need to have a whole bunch of partners in place to support your commercialization effort. For example, we are backed by leading cloud provider, Microsoft. And also from a hardware point of view, we have strong relationship with Honda and GM, who have hundreds of years of combined experience in building hardware, which is an important aspect when you're building cars for commercial use, right? And not to mention our agreement with Walmart, a retail giant.

So as we move from self-driving, from R&D to commercial, it's an inflection point. It's pretty exciting time. I mean, at Cruise especially, AV in general, and Cruise in particular, we're at an inflection point. This is the most exciting phase where you're moving from R&D to commercialization. As we hop on, you can see autonomous vehicle without drivers behind the wheel driving around the city. A lot of exciting work both on the AV stack as well as in the

infrastructure side of things. So I'm super excited to be at this time, especially in an area in our company that's at the cutting edge.

[00:29:48] KP: Well, Vinoj, thank you so much for taking the time to come on Software Engineering Daily.

[00:29:53] VK: Thank you very much. It's been a pleasure talking to you, Kyle.

[END]