

EPISODE 1085

[INTRODUCTION]

[00:00:00] JM: Kubernetes continues to mature as a platform for infrastructure management. At this point, many companies have well-developed workflows and deployment patterns for working with applications built on Kubernetes. The complexity of some of these deployments may be daunting, and when a new employee joins a company, that employee needs to get quickly on-boarded with the custom development environment.

Environment management is not the only issue with Kubernetes development. When a service gets updated, that update needs to live and usable as fast as possible. When Kubernetes-related errors occur, these problems need to be easily accessible in a UI for triage.

Dan Bentley is the CEO of Windmill Engineering, a company that makes a set of Kubernetes tools called Tilt. Dan joins the show today to talk about the workflow for deploying Kubernetes infrastructure and the role of Tilt, the product that he has been working on.

If you want to reach 30,000 unique engineers every day, consider sponsoring Software Engineering Daily. Whether you are hiring engineers or selling a product to engineers, Software Engineering Daily is a great place to reach talented engineers. You can send me an email, jeff@softwareengineeringdaily.com if you're curious about sponsoring the podcast. And we're also looking for writers and a videographer. If you're interested in working with us, you can also send me an email, jeff@softwareengineeringdaily.com.

[SPONSOR MESSAGE]

[00:01:32] JM: Our thanks to Datadog for sponsoring this episode of Software Engineering Daily. In their latest report on container trends, Datadog analyzed 1.5 billion containers and found that in orchestrated environments, container lifespans averaged about 1 day. To easily monitor this dynamic infrastructure, use Datadog's container-specific monitoring features. The container map provides a bird's eye view of your container fleet and the live container view searches groups and filters your containers with any criteria, like tags, and pods, and

workspaces, and you can try Datadog in your own environment with a free 14-day trial. You'll receive a complementary t-shirt. Just go to softwareengineeringdaily.com/datadog to install that agent and get a complementary, warm, comfortable, cool T-shirt.

[INTERVIEW]

[00:02:34] JM: Dan Bentley, welcome to the show. Hey, great to be here.

[00:02:35] DB: Hey, great to be here.

[00:02:38] JM: You were at Google for 10 years, from 2004 to 2015. How did Google deploy services when you were there?

[00:02:48] DB: Oh! Yeah, that's a great question, and it changed a lot and I joined as an intern on the release engineering team as Google was getting ready to go public and you started caring about are there nefarious software engineers who are able to insert the bugs to round pennies and siphon it to their accounts. That was a cool entry point.

I honestly spent most of my time at Google not working on services that ran in production as much as the developer tools that ran on laptops. I'm trying to swap it in. I worked on Google Sheets at the end and then it was – We had nightly builds that got cut and pushed out and I was working on particularly nasty server crashing bugs that only showed up under heavy loads or kind of had to get my change in that day so that the next day at about 3PM Eastern I could see everything crash in the way that I wanted it not to.

Deploying services was really up to teams. There were different patterns and you could see that people copied what had worked from teams that they had been on or teams that they were friendly with. Yeah, each team's culture had some effect on how they deployed to production.

[00:04:21] JM: That's obviously the Google way of doing things, which was back then pretty Google specific. I think Google was pretty far ahead of everyone else and Facebook had their own thing. After Google, you were at Twitter for a year and a half, and Twitter was more

representative of the broader software ecosystem. How did Twitter manage its infrastructure when you were there and how did that contrast with what you saw at Google?

[00:04:48] DB: It was really interesting leaving the Google cocoon and learning about different kinds of patterns. Twitter famously moved to Scala, which is a functional language on the JVM, and I was impressed with the way that they thought of their systems in a more functional programming kind of way. When you look at – I think the paper by Marius Eriksen, which is fantastic, of your server as a function that gets to how you compose things. I realized that I had just been taking for granted the kind of storage systems and the way you built things on top of them that was so different than more popular, like event stream modeled systems.

At Twitter, it was I think similar-ish, because a lot of people had been there, but with a different attitude of we really want to enable teams to communicate with other systems in a better way. I think when you look at today, a lot of cool service meshes and some that weren't popular, but were brining stuff from Twitter. When I think of Linkerd from Buoyant, that's some ex-Twitter people putting that way of thinking about how do you communicate with other services? Really getting that out to the wider world.

[00:06:22] JM: When you started working on your own company in 2017, you had these past experiences. You had been at Google. You had been at Twitter. What was your original thesis for windmill engineering when the company got started?

[00:06:37] DB: Yeah. I think our original thesis was that cloud could make the developer experience, the inner loop, as you're typing better. That could mean a lot of things, like web IDEs. But I think from the beginning, we were really focused on, "Oh! There's a great tweet that's like debugging is like being the detective in a murder mystery where you're both the victim and the killer.

We're excited to at least save you from having to also be the lab tech at the CSI that's having to perform the experiments. How can we show you faster the thing, the piece of data of this build broke or this test just switched to passing that's going to give you the aha moment of, "I know what's happening. Now I have to figure out how to change what's happening," but to get you

back to that, for me at least, really, fun, imaginative part of software engineering. That was our original thesis.

[00:07:51] JM: How has that thesis of making debugging, making the software development experience easier? How has that evolved overtime and become refined?

[00:08:00] DB: Yeah. We built a product that I'm incredibly part of technically that each time you hit save would upload just the file contents you changed and your run build and unit test in our cloud, in parallel, in seconds. When we showed it to people, a lot of the response was, "Hey, yeah! I've had times in my career when that was what was blocking my productivity," but now it has a lot more to do with building container images and pushing them and pulling them and scheduling them in a cluster.

How I think of it is that our vision that cloud could improve the developer experience ran smack into the reality that cloud is already hurting the developer experience. What I mean by that is with cloud native architectures or microservices or distributed systems, you're not building your app in the same way you used to. You used to build your app as one binary. For 40 years, apps got bigger by binaries getting better and you needed new tooling to shove off those bits into one binary faster and more consistently.

Then with this switch to cloud native architectures, suddenly you don't just have one binary. Your app gets bigger by having more pieces of it. The tools that were meant to build single binaries, well, that's all they do. When you're thinking of your app as composing these little pieces, suddenly you're having to maintain all of these workflows in the state they're in in your head so that when you save a file, you have to remember which microservices you do or don't need to restart and what shortcuts you can use to get them into the proper state.

[00:10:09] JM: All right. Well, it's 2020, and the Kubernetes ecosystem has had some significant time to evolve at this point. Can you tell me how you see the Kubernetes ecosystem today and how it intertwines with the products they you're building?

[00:10:28] DB: Yeah. I mean, I think the Kubernetes ecosystem, really, the key piece of it is the community. The way that it went from being a project that was written by Googlors and

published to now you have all sorts of companies as well as individuals and just some really great people, Stephen Augustus, Ian Coldwater come to mind. That's fantastic. I think some of the genius of Kubernetes is in knowing who their customer is.

You think of some ways to run data centers or clusters that might target the cluster ops people of how do you just get servers into the racks and being able to accept traffic at all? You think of the ops of is my app actually serving 200s or their problems? You think of the developer experience. I think Kubernetes kind of realized, there's so much to do here that we can't solve all of those. If we try to solve all of them, we're going to have too much to talk about before we actually deliver value.

I think one of the great things about Kubernetes is that it goes, "Here're these orthogonal concepts," and so you're able to, really, as the Cooper Mideast developer, be able to be the preferred customer. You're able to get in there and add things and build on top of it, and that's really powerful. It gives them a defined scope so that you know what is and isn't Kubernetes.

There are lots of experiences that you kind of want out of the box, except then it's opinionated and it might not match your opinions. You see all of the different tools to set up a Kubernetes cluster. That's kind of outside the scope of the Kubernetes project. I think the core developer experience of as I'm writing code that runs on Kubernetes in production, what do I do as I'm in my IDE before I have a pull request. As I'm hitting save every couple seconds or couple minutes, what's that experience? That's really the question that we're trying to answer with Tilt.

[00:12:56] JM: By the way, why are there so many templating and build tools for Kubernetes?

[00:13:06] DB: Oh, yeah. There are so many different layers to answer there from the technical, to the opinions, to the organizational structure questions. I think it's that people see a hole, and maybe there's not actually a hole. There was something else out there, but there's something that's missing. The templating is a great one, a great example. There was Jsonnet and Ksonnet. There was a – Helm is very popular now. Customize got pulled into kubctl with a little drama around it. [inaudible 00:13:46] is a new one that's out there.

I think it's that, really, Kubernetes is building blocks. How do you put those building blocks together, right? Why are there so many different Lego models? Well, because they're building blocks and different people see them fitting together in different ways in their head. I'm not sure if that's a particularly satisfying answer, but I think it's kind of a lowest common denominator – Let's get stuff going and let's see what sticks. I think it's exciting. Time, also time, that can be nerve-racking because you have to think, “When I pick up this tool, is it still going to be around in two years?” That's not particularly fun as an end-user.

[00:14:37] JM: Indeed. Let's talk more about Kubernetes usage. Tell me about some of the problems that you are seeing Kubernetes teams encountering.

[00:14:49] DB: Yeah. What we see is that the incentive, right? When the incentives aren't great, every run responds to the broken incentives in a slightly different way. So the problem that we see most, that we're most focused on and try and dig into and hear about is how do you start your app and change it? Some people who run in Kubernetes for production, they just run all of their apps outside containers on their laptops, and so you just in one tab start up the backend in your terminal. In another tab, you start up the frontend.

Some people use Kubernetes as they are developing on their laptop using Docker for desktop or Minikube, but then you have problems. You have to wait, because building an image, maybe it's fast and it takes 10 seconds, but that 10 seconds is different than the two seconds that it was working outside containers.

Some teams say, “You know, we're not going to run our whole app at all. Instead, while we're developing, we're just running one micro service at a time,” and as result, if you're trying to add a new field to like the user in the app. Step one is you change the database schema. You make a pull request, you land it, you wait for it to propagate. Then the next day, you add the code to the backend. The day after that, you add to the frontend, and if you realized you needed something different in the schema, well, you're back to square one.

How do you see your app as you're developing it and with what tradeoffs do you have to make between speed and consistency? Because everyone kind of wants to be testing the same way that they run in production, but you don't want each developer to have to wait as long as it takes

to get things into production. You don't want to have to be paying for as many replicas for each developer. Those are the problems that I'd say I am most attuned to. Is that answering the right question?

[00:17:16] JM: Yeah. I mean, there's no right answer to the problems that people were encountering. One way I think we can explore the problems of developing with Kubernetes is through the lens of a developer who is new to working with Kubernetes. If I'm a developer that's new to working with Kubernetes, what is my workflow? Is it hard to get started? How am I getting started?

[00:17:44] DB: Yeah. Just to kind of set the context of what life was like on your previous team that wasn't on Kubernetes, at least as I understand it to kind of a platonic ideal. You come in and you type. Maybe it's makes server or a Bazel build or something, you get a server, you start it. When you change a file, maybe you hit control C and then you rerun it. When there's a problem, it crashes, and the stack trace is the last thing in your terminal window. You know where the problem is.

When you come into Kubernetes, you probably have a `start.sh` script that handles doing a Docker build to get an image from the source that's on your file system. Then it will take the image that was built, template it into YAML in some way, and then `kubctl` apply it, and that will start one service. You probably have a few or many different services and they are all running. Because Kubernetes is really built for production, those logs are just being logged in Kubernetes. But when you run into a problem, you have to go play 20 questions with `kubctl`. App is running. You're interacting with it. You see something weird happen. How do you find that out? Well, it's up to you to go ask each individual service, "What state are you in?" I think every Kubernetes developer has an experience with crash loop back off, right? There's just a problem as something starting up and you wish you'd known about it before you started to interact with it, to try to interact with it.

It's really that difference of going from, "When I hit save, I have one command that brings my app up to date in the pre- Kubernetes monolith world. To now, you're having to keep these different plates spinning and remember what parts of the workflow to do and it's very custom to your team. It was maybe written by a teammate whose maybe not there anymore or working on

a different part of the project. It can be really hard to just get a sense overall of what's happening in your app. Where's the problem?

[SPONSOR MESSAGE]

[00:20:24] JM: JFrog SwampUp is an online user conference with more than 30 sessions from cross-industry expert, including Google, Microsoft, capital one, Adobe and more. You can get a unique look into the broad DevOps market, not just point solutions. There will be tracks for cloud native DevOps, enterprise DevOps, DevSecOps and digital transformation.

Participate in expert DevOps training classes across DevOps tools, security, CICD and more, and you could join from two time zones, June 23rd and 24th for the Americas, and June 30th and July 1st for EMEA and APac to suit daylight hours across the globe. Go to softwareengineeringdaily.com/swampup to learn more and sign up

JFrog will be donating all conference registration proceeds to COVIT-19 research. Go to softwareengineeringdaily.com/swampup and check out SwampUp.

[INTERVIEW CONTINUED]

[00:21:27] JM: Are there mistakes that commonly occur where a developer is working with Kubernetes and they accidentally deployed to production?

[00:21:40] DB: Oh sure. Absolutely. Yeah. The new model of what cluster you're talking to, it can be very useful, but it's also easy to do something in the one terminal window then switch back to where you were doing your development. You checked on what's happening in production and then you switch back and you go, "Oh yeah. This is where I'm developing." you do some command and only then do you realize that asking about production over on the left means that now the window on the right is also talking to a production.

We're giving you these very low level commands, and that's great so that you can compose them. But it also means that everyone has to compose them, that every developer has to be not a Kubernetes expert, but familiar enough to know what the abstractions are even when you're

saying, "I'm just trying to change from some JavaScript. I'm just trying to update the constant in this Go code. Why do I have to go to boot camp for a cluster system just to make a change?"

[00:22:48] JM: When I've already deployed a service to production, I'm very frequently going to want to update that service, whatever service I've deployed to production with new code. I might write my changing, and that change is ready for release. What commonly goes wrong in new releases to production?

[00:23:10] DB: Oh, yeah. I think there are in many orgs – An di probably safer good reason, a lot of difference between – There's almost a handoff between eng and ops, and there are great things about DevOps and having it be that everyone can be on call, but there is also I think a lot of orgs where that's not the case. I'm not sure if it should be the case.

One thing that goes wrong is just it takes too long for you to wait. So you merge a pull request. Then there's some CI system that's starting to send it out, but you're not waiting for it. When there's a problem, you've been not just checking email or Reddit, but actually starting something else, because if it's taking more than a minute, you're probably not going to wait. You've already swapped into another context, and it's not like you get a little notification at most places that there was a problem.

One is just you don't even know that there's a problem. Then there is all sorts of issues of, "I was just changing code and something in the way that the cluster was set up, suddenly we started using too much RAM or too much CPU," and that wasn't the level I was thinking at when I was making the change. So now I have to investigate these new kinds of problems.

There's often even the social divide of there's more people coding than interacting with production, and so you're kind of getting through this semi-porous membrane where different people have different levels of comfort with dealing with production and touching it. I know lots of people who like to delete their tokens that touch production except when they really need to, because that's a great way to get in trouble.

[00:25:23] JM: Let's start to talk about what you have built with Tilt. Can you explain what the Tilt product is?

[00:25:31] **DB:** Yeah. I hope. You tell me in a couple minutes.

[00:25:35] **JM:** Okay. Sure.

[00:25:37] **DB:** What we are trying to do with Tilt is make it that multi-service development is easy, that you can get back to where you were where when you come in in the morning or when you join a team, you just type Tilt up and it takes all of the source code for all of your different services and it spins them up. And so you have a developer instance that could be running on your laptop, could be running in a cloud cluster, and you're able to, when you hit save, then see that code running for real, not in a production environment, but you can load it in your web browser and you can interact with it and you can change any line of code, whether it's Go code in the GraphQL server or Typescript code that's running in the frontend or a JVM service, right? Kind of getting live reload for your entire stack.

So what's important there? Well, one thing is the speed of from when you hit save, how long does it take until your code is running? Developers are this fantastic and evil kind of lazy where if they see a shortcut, they'll start doing it and it will end up being unsafe at some point, and then you've saved yourself some time because it's faster as you do it. But when it gets out of sync or askew, you then have to spend hours figuring out what's happening.

I'm really proud that Tilt supports live update, that you can actually update containers that are running in the cluster even in the cloud in less than a second. You would never want to do that for your production workloads, but it's the way that you're able to get the same kind of feedback loop that's low latency as you had before you move to Kubernetes for development.

Then I think the thing that we've realized is you don't just need to get the code out there. You also need to get the feedback. When there is a problem, when a pods in crash loop back off, when it dies, when it can't even accept the YAML, you're able to see it in the UI, because Tilt runs as a web app. It's running on your laptop. You're seeing it in your browser at local host, but you're able to get a holistic view. In addition to the logs in the center of your screen, over on the right, there's a sidebar that shows you each microservice in the constellation that makes up your app and a little red icon and a badge showing you what the problems are when something goes

wrong. So you're freed from having to hunt it down yourself. You just sit there, and when something goes wrong, you get told it via this peripheral vision.

[00:28:51] JM: That's a bunch of nice benefits. The example of the instant build, kind of hot reload experience for when you edit something in Kubernetes and you get kind of a hot reloading experience, what's involved in engineering and building that?

[00:29:11] DB: Yeah. Are you talking from the Tilt side where we built it, or like if you're trying to use it, how you how you describe it and set it up?

[00:29:22] JM: From your side. How did you build it? What did you have to do to build that?

[00:29:25] DB: Yeah. It turns out, the Kubernetes API server supports this kind of operation. You can run `kubctl exec` to run a command in a running container. There is `kubctl cp` to copy files into a container that uses that infrastructure. What Tilt is doing is when it `kubctl` applies then waits till the pod is ready, and then it is watching your hard drive.

When you save a file, it takes that file and it sends it to the running container and updates it in place there. Also, it can run any build steps you have so that it can work, not just for interpreted languages, but also for compiled languages.

Where I think kind of productionizing these things that are part of Kubernetes that may be they shouldn't be for production systems. You might want to lock it down a bit more, but you're able to use Tilt just as this almost conveyor belt or pneumatic tube. Well, those are dated references. A transporter from Star Trek to just have your new code running in the container and Tilt's handling, figuring out which container to connect to and how to do the port forwarding and how to set it up.

[00:31:00] JM: In order to use Tilt, there's actual Tilt instance, right? There's an actual Tilt application?

[00:31:08] DB: Yeah. It's a Go binary with a JavaScript frontend that runs on your laptop. You can just download it from the website, and then in your terminal, you type Tilt up, and then Tilt is

running on your laptop. It's using your Kubernetes credentials. It's not having to like be something that your cluster admin has installed in the cluster.

[00:31:33] JM: As a developer that's working with Tilt, I am configuring my Tilt usage with a file called a tilt file? Is that right?

[00:31:45] DB: Yeah. Exactly. What goes into Tilt file config?

[00:31:52] DB: In this simple case, you pointed at your Docker file and you pointed at your Kubernetes YAML and then Tilt figures out, "Oh! This image is referenced in this YAML, and so I know everything from each source file to where it's running in production." To use the live update, we're talking about also that kind of information isn't just in the Docker. So then you add some lines to the Tilt file where you say, "My Docker build, I don't just want to build the image each time I save. I want to actually sink the files from this directory on my laptop file system over to that directory in the container file system. I want to run these commands." There's your little more work to get the speed.

We also have found that that kind of description of, "Oh! You just have Docker files and you just have Kubernetes YAML. Isn't what most teams have?" I'm really excited that we let you also say, "Hey, my YAML isn't just a file on my hard disk. Actually, we run this Python script that generates YAML, because maybe we have some way of specifying a service within our company."

That's really what we think of as Mr. Rogers onboarding, where when you're a project like Bazel or Kubernetes, can say, "Yeah, I'm going to get you somewhere great, but it's going to take a lot of work to set it up, but it's worth it, because like we've spent three millennia building this." We're a small company. We don't have that ability to kind of convince you to do all of this work upfront. Mr. Rogers onboarding is like Fred Rogers, a Pittsburgh resident, Presbyterian minister, children's television show host. He always said, "I like you just the way you are," and Tilt likes your project just the way it is. Even if you have grungy scripts to generate YAML, or even if you're using some other tool to build the container image. You can point Tilt at whatever you have set up already and use that without having to re-architect it to the way we think you should

do it because we understand every project has different constraints and we embrace that diversity.

[00:34:32] JM: The developer experience for using Tilt, is it like I am working with an IDE or is it like a console tool? Tell me more about that the developer experience.

[00:34:45] DB: Yeah. You keep editing in your IDE. You start Tilt in your terminal. It opens up a tab in your browser where you get the UI and it's watching the file system. When you hit save in your IDE, it goes, "Ah! I know what to do." You just changed main.go in the shopping cart service. I know all of the things I need to do to get it that your personal dev instance matches it.

So you can see the logs of the build and see how long it took. You can see the logs from the shopping cart service as it's running. You can point-and-click. We have had a UI that's really all in the console and we found that people – There aren't idioms for how do you navigate between different tabs. It's not really discoverable, whereas in the web, people know how to pick and choose. So you can say, "Hey, maybe I want to see all the logs from all the services. That's cool to kind of see what's happening, but there's probably a lot. I can also click on the sidebar on the right on just the shopping cart service and I can see, "Boom!" It's red. I click it. I see the stack trace. that's what's causing the problem that I'm seeing in my app right now.

[SPONSOR MESSAGE]

[00:36:28] JM: Scaling a SQL cluster has historically been a difficult task. CockroachDB makes scaling your relational database much easier. CockroachDB is a distributed SQL database that makes it simple to build resilient, scalable applications quickly. CockroachDB is Postgres compatible, giving the same familiar SQL interface that database developers have used for years.

But unlike older databases, scaling with CockroachDB is handled within the database itself so you don't need to manage shards from your client application. Because the data is distributed, you won't lose data if a machine or data center goes down. CockroachDB is resilient and adaptable to any environment. You can host it on-prem, you can run it in a hybrid cloud and you can even deploy it across multiple clouds.

Some of the world's largest banks and massive online retailers and popular gaming platforms and developers from companies of all sizes trust CockroachDB with their most critical data. Sign up for a free 30-day trial and get a free T-shirt at cockroachlabs.com/sedaily.

Thanks to Cockroach Labs for being a sponsor, and nice work with CockroachDB.

[INTERVIEW CONTINUED]

[00:37:50] JM: Would you call the work you've done around Tilt, was it inspired by Docker Compose heavily?

[00:38:00] DB: Oh! That's a really great question. Docker Compose is fantastic. It is a popular way to run different services. It has a really low overhead way to describe what you're trying to do in a way that no one can just kind of sit down and write a Kubernetes YAML from memory, because there's all of these your finicky details you just copy and paste.

Tilt, actually, we mainly talk about it sitting on top of Kubernetes. It actually can also sit on top of Docker Compose. If you're using Docker Compose, you can start everything up. You can see all of the logs together, but without you adding hacks, it's not updating. You have to, after you hit save, remember, "Oh, I need to go rebuild and restart that service." You're not able to dig into the logs as easily.

Kind of both Kubernetes and Docker Compose and Docker Swarm are about orchestrating containers. Tilt is almost at a level slightly above that of orchestrating the development, the workflows from when you hit save until the code is running and then back to when you know what happened when the code ran.

[00:39:30] JM: Just so we can kind of level set here, could you rephrase what are the problems that Tilt is solving for the Kubernetes developer?

[00:39:43] DB: Yeah. One problem is having to deal with a bunch of different tools between the Docker build in the kubctl apply, and having to know those tools for each developer is a lot of

cognitive overhead to add. Tilt understands the workflow and automates it in the same way that your build tool automated combining the different libraries.

With a monolith, you didn't have to know every library from day one, whereas with Kubernetes development, you need to know each microservice because you deploy them all separately. The other problem is speed that dealing with building container images slows you down. Then another way you're slowed down is having to go hunt for problems instead of having them popup so that you can see them when they occur instead of when you think to ask, "Hey, is the service X broken?"

[00:40:58] JM: You also have CI tool, right? Like Tilt CI?

[00:41:02] DB: Oh! No. We say that we're like your IDE or for when you're in your IDE. We've thought about the CI space and we think the same way that you have a different way to release the final of binary versus what you're using as you're editing. We think that we're in the space of as you're editing. Not of how you release to production.

[00:41:32] JM: Got it. If I am deploying using Tilt, there is this new tool that I – Or this tool that I hadn't heard about when I was reading through the Tilt documentation called KIND, Kubernetes in Docker. This is the ability to deploy Kubernetes within a Docker container. Does this how deployments work with Tilt?

[00:41:57] DB: Tilt works with any Kubernetes cluster. So you can point it at your AKS, or EKS, or GKE, or you can point it at Minikube, or Docker for desktop. Now, also, as Kubernetes is becoming more popular, there are more ways to run a Kubernetes cluster on your laptop. KIND is one, Minikube is one, Docker for Desktop is one. MicroK8s is another. Ben Elder is doing great work on KIND. So it's just a really easy way to get started. It has really few dependencies. So we recommend it. Because we hear from a lot of people, I just have to reset my local Kubernetes cluster on my laptop all of the time. I'm having to delete VM's, and KIND seems of just kind of gotten it right.

[00:42:56] JM: Cool. The business around tilt is Tilt Cloud. As I am using Tilt, if I am managing my deployments this way, managing my infrastructure, I might want to share the data that's

being generated off of my infrastructure. What's the data that comes off of Tilt and why would I want to share that with other developers?

[00:43:25] DB: Yeah. We see the world as having app devs who are writing the code that runs in the cluster, and that's the vast majority of devs and of users of Tilt. Then there's also the developer experience team, the dev X team whose setting up that tilt file. So it's not each user has to set it up. It's checked into your repository and it's maintained. The kinds of questions of what's my developer experience like? What's the data? You can ask questions like that of what's running in the cluster, because you have observability. For the workflows that run on coworkers' laptops, it can be really hard to actually know what's build time actually like. What's the ground truth? It can change a lot, because different people have different generations of laptops or have crowded hard disks.

Tilt cloud, for now, is mainly for that developer experience team. We think it's the infrastructure for when production is a coworkers' laptop. Laying you answer the questions that observability gives you in the data center and CICD give you about releasing with confidence to the data center, Tilt cloud is for the developer experience team to actually be able to show the impact they're making.

[00:45:12] JM: How have you seen it be used by developers?

[00:45:18] DB: One of the features we're really proud of is snapshots, which is when you're in the Slack channel in your company's Slack and someone comes in and says, "Hey, is anyone else seeing a problem with the shopping cart service?" You as the developer experience engineer enter into this, I think frustrating for everyone, game of telephone where you ask what are you seeing? No, not that line. Like go to this other log file and pull this.

With snapshots, because Tilt has a web UI, we can just upload a snapshot of what's in that UI not as a screenshot, but as all of the data so that if you're helping me in Slack, I can just send you a link. You can open it and you can click around and get the answer to the question of what's blocking me and you can get someone going again in 20 seconds instead of trying to set up a screen share.

[00:46:30] JM: Was the snapshot tool, was that hard to implement?

[00:46:34] DB: Because of the way that the tilt UI is implemented, really, snapshots are almost – I feel like I'm giving away all the secrets here. It's almost paste bin, if you remember that, where we take the JSON data that everything that the Tilt UI knows about. We upload that to a server. Then when someone else comes and looks at it, we give them the same JavaScript and the Tilt web back and the JSON that's the data that was uploaded and you're able to see it pretty easily.

[00:47:17] JM: Very cool. It's not like an actual like distributed systems level snapshot of everything going on at any given instance in your infrastructure.

[00:47:29] DB: Exactly, and that's also something that we hear and is on our roadmap. I think one reason I'm excited to be in this space in addition to the community is that I think for a lot of the reasons we've mentioned, being in the best Kubernetes dev tool is like being the meanest moppet. It's just not a very high bar right now

We're going to have to get better, and that's exciting. But right now, just giving someone that access of what are the logs? What's crashing? Solves the same way that you null pointer errors are 80% of bugs in languages with no pointers. You can fix 80% of bugs really quickly in a way that you haven't been able to before.

[00:48:19] JM: I think the way to describe your – The vision that you're going for, is basically in your ideal world, if you're managing a Kubernetes cluster, you also have this Tilt node in your Kubernetes cluster and your Tilt node does things that are useful to you. And there are plenty of generically useful things that that node could do.

[00:48:45] DB: Yeah. Yeah, you have your own instance of your app and a UI that makes sense of it.

[00:48:52] JM: Right. Yeah. UI is something we didn't really delve into that much, but the idea that like you have all these services running across your cluster, and the tilt UI gives you just a perspective into what all these different services are doing.

[00:49:09] DB: Yeah.

[00:49:10] JM: Well, tell me little bit more about how you expect the space to unfold? How do you expect the Kubernetes ecosystem to look in the next 3 to 5 years?

[00:49:20] DB: Yeah. I think the best description I've heard of it is Kubernetes has an air of inevitability, and I think it's still not what everyone is using, but I think there's really exciting things happening. I hope that we see more and more pieces of infrastructure that are accessible in Kubernetes. I think you see – Traditionally, cloud providers have had some console that you can go to to spin up new resources. Now you're starting to see, especially Azure and GCP, I know best have Kubernetes customer resources so that you can say, “Hey, I want a cloud SQL or CosmosDB instance as Kubernetes YAML,” and then it gets spun up. So you're able to control more of your infrastructure from one common interface. I think that's pretty exciting.

I think we're going to see more and more – It's almost like Kubernetes is zipping up the stack, right? Everyone kind of had VM's for, but what you did on those VM's, how you deployed to them was different. I think you're seeing that now people are building, especially at large orgs, their own layer above Kubernetes.

Knative is an open source one. I think you could go to lots of orgs, they have a way. What's the next layer above, right? Even if Kubernetes is there, how much do people have to know about it? I'm really not sure what the answer is. I know Git, for instance, claimed to do this. They have this whole notion of Porcelain so that people have a better interface. But they were all such leaky abstractions that people still have to know about Git. I'm really unsure but excited to see what happens with what's that next layer above Kubernetes. Is there one? Is there three. Are there 500? I just don't know.

[00:51:31] JM: Very reasonable question. Dan, thanks for coming on the show. It's been great talking to you.

[00:51:35] DB: Yeah, you too.

[END]