EPISODE 1445

[INTRODUCTION]

[00:00:00] JM: A company's response to an incident defines how that company responds to an adverse unexpected scenario. Kintaba automates the incident management process so that teams can quickly report, respond, resolve and reflect on major incidents collaboratively. Kintaba's Incident Response platforms gives you instant access to incident management processes based on the practices of companies like Facebook, Google, Stripe and other successful organizations.

John Egan is co-founder and CEO at Kintaba and joins the show to talk about incident response.

[INTERVIEW]

[00:00:31] JM: John, welcome to the show.

[00:00:32] JE: Thanks, Jeff.

[00:00:35] JM: I want to talk to you today about incident response. And to start us off, can you just explain what a typical incident response process looks like at a typical company?

[00:00:47] JE: Well, so an appropriate incident response process, at a company that practices this and practices, it well, looks like a series of steps that are repeatable each time. So, around declaring the incident, the group response to the incident itself, including the sub roles, that those responders have, a step around mitigation, and making certain that everyone knows that status, and then the actual review on the incident itself that identifies root cause and make sure that it doesn't happen again. That's the really basic flow that any organization that's practicing incident management and incident response should be following.

[00:01:24] JM: How does incident response vary in practice, from company to company?

[00:01:28] JE: So, despite that being sort of the best practice, the reality is, most organizations sort of deal with incidents more as an outlier unexpected moment, without too much structure around it. So generally, when you talk to a person or most organizations about the last incident they remember, they tend to remember more of the unfortunate parts of it, right? The panic that comes into it, the multiple communications channels of where are we talking about this? Where are we fixing this? What is the status of this? What are we doing? By their nature incidents, as we remember them, and as they're defined, tend to be these kinds of unstructured moments of worry inside an organization where things are going badly fast. So, that's kind of more often what we see, in industry. It's only really recent history, that sort of putting structure to it in a meaningful way has started to happen, and that's really come from organizations like Google, like Facebook, like Netflix, that have started to become much more public about their internal practices, maybe since about 2016.

[00:02:32] JM: There probably hasn't been as much coverage of the area of incident response as is due. I know you're putting on a conference around incident response. Can you describe what you're looking to cover at that conference?

[00:02:48] JE: Yeah, so what we found is incident response often gets buried underneath other practices within the organization. And its sort of one of the challenges for the space historically, is when you're dealing with an incident because it tends to be a company critical moment, it's actually something that cuts across lots of different organizations, right? It's not just the SRE team, or just the engineering team, or even just the customer success team. It's all of these teams having to come together in concert. And something we noticed when we started Kintaba was that organizations, if you were to walk up to an entire organization, say who here works in incident management or practices, incident management, you wouldn't actually see that many hands go up. But if you stand in front of that same company, and you say, who here has been involved in an incident? Who here has had to be a responder to kind of accompany critical events over the last six months? You'll actually find a lot of hands start to go up.

That sort of spurred this idea internally, that we probably ought to go and create a place for these folks to come together where we can start to build community around it, because we couldn't find an existing conference out there. So, we started this new thing. It's called IR Conf. It's really the first incident response conference and it's meant for anyone who's ever been part

of a major incident at their organization to kind of come together and talk to other people, and start to share best practices and stories and approaches that people use to deal with these emergency events at their organizations. And we've gotten some really great people. We've got Emily Freeman from AWS is speaking, Dave Rensin, who's ex Google and SRE and is now at Pendo. Folks, co-founders from Honeycomb, Commodorem got speakers from Wix. It's going to be a really interesting conversation that finally just gets to focus on this very specific practice of dealing with major critical situations at organizations and really just helping us all come together as a community and accept that we're all dealing with this. This isn't a unique thing, the Facebook's and Googles of the world.

[00:04:46] JM: Can you talk to me about the tooling of the modern incident response? Where does tooling fit in and what are the primary tools that are used?

[00:04:58] JE: So, the tooling looks an awful lot like orchestration tooling, mixed with collaboration, tooling. Because these are the two things that really come together for a successful incident response. You need a tool that will help you to push that process that I described earlier, that declaration, response, remediation and review. You want an orchestration tool to manage that. And then you want that orchestration tool to be plugged into a collaboration tool, that's actually letting your responders come together, pull all of that relevant metadata together and work towards determining root cause, or at least mitigating steps that can be taken to get the incident under control as quickly as possible.

Kintaba, our company, really tries to bring those two things together, and give you a really simple and easy to use platform as a tool that can be deployed across the organization, as opposed to a tool that maybe two or three people from your SRE team would have access to. You really want this tool to be accessible by everyone. Because when things go badly, when metrics fall South really quickly and start to affect customers, you're bringing together as SREs, and engineers, and product managers, sometimes you're even bringing legal into the fold, and certainly PR and customer success depending on the situation.

So, the tooling really looks like that combination. And when you look at it, as it's wired into the tooling that already exists, some of those pieces are probably tools you already have. For example, we wire really closely in with Slack to make it so that if that's already your

collaboration tool, all of the orchestration that's happening inside of Kintaba in terms of making sure you move through this process is also connecting in to your Slack channels, so that the place where your customer or your employees already are, becomes the place where you can practice that response effectively. You really sort of see those two things coming together when you think about tooling, and then you want that like wide deployment across the organization.

[00:06:54] JM: You have been working on Kintaba, which is a tool that is dedicated to incident response. Given what you've covered about incident response thus far, describe why a dedicated tool is needed for incident response?

[00:07:09] JE: Yeah, so this is sort of a big shift that's happening here over the last couple of years and is continuing, is that in maybe the old days, we would say, okay, incident response, maybe can live inside of our task management system or maybe our project management system above that. Maybe we can just handle incident response directly inside of Slack, and we can record the incidents somewhere like Google Sheets, and we can start to write our automation processes for things that are predictable that we want to do during incidents happen inside of like Zapier, or write them as shell scripts and kick them off.

What you start to find really quickly is that all of the existing tools that are out there have to sort of be used individually and independently for pieces of incident response, and it gets really messy really quickly. So, a simple example of that is, if you were to just try to use Slack for incident response, you end up with different parts of your organization running to different channels at different moments, trying to figure out where the right places to have the conversation, because of the fact that the incident cuts across. If you have a site outage, your entire site goes down, and your SaaS business, it's everyone who's going to be running to the forefront, saying from the CEO, all the way down to a frontline salesperson, trying to figure out what's happening, where to go, where to have that conversation. We hear this over and over again, when we go talk to organizations, especially unicorn startups, companies that are growing quickly, where that panic of trying to use a tool that isn't really purpose built for the system, or this situation, just causing more confusion and a lot of administrative overhead.

Where you actually end up in a world where, trying to put together the management of the incident ends up taking up more time than just the actual incident resolution itself. You really

don't want teams stepping on each other. You really don't want to be defining the roles for the incident response at the moment of the incident. You want those predefined. You want them predefined well, and you want them visible. And you really don't want to be running to different places for video communications as well. Where are we going for the Zoom room that's attached to this thing?

You also have a lot of these moments where you want to say, “Well, who from organization I've never worked with before? Who from PR needs to be involved in this incident?” It's big enough, that we're going to have to go and actually talk to press about it. Who is that person? I'm an SRE. I've never had to deal with PR before. Who's the on call? Do they even have an on call? All of these questions come up. And you realize that the existing tooling, while effective, PagerDuty is fantastic probably for your technical on call rotations that you already have. Slack is doing a great job of day-to-day collaboration work. Google Sheets is doing a good job of capturing sheets and table data that you're putting together. But when an incident happens, you need all of this stuff really happening in concert, in an orchestrated fashion, and that's where a single tool really solves that problem.

There's a saying in industry, which is, “If you want to process, build a tool.” When we were starting to build Kintaba, we went to all of our friends who had left, mostly who had left Fang and gone to other organizations and said, “You know what, what are you missing from your tool stack?” And consistently, that was the feedback. It was, “Well, we don't have that place to go, that orchestration tool for when a major emergency happens, and we really wish we could just pull one off the shelf.” And that's really kind of the genesis of Kintaba was there ought to be a very simple off the shelf solution for this, that provides that real time orchestration, as well as the long-term learning and remediation tools.

[00:10:42] JM: Describe if I'm a user of Kintaba, and I have access to it. What does that change about my incident response process?

[00:10:53] JE: So, first and foremost, it's going to give you visibility into the incident response world inside of your company. So, at a very simple level, what is currently on fire? And what has been on fire recently? It's a very simple sounding question, but it's pretty tough to answer in most companies across organizations. If you're a random engineer, company like Facebook,

you can go and see that you can go to a dashboard at any moment and see what are all of the fires happening, everything from SEV1, company critical, everything's really going down. All the way down to SEV3, like lower customer impact, but still, something that's an emergency that we need to deal with in a real time factor. First and foremost, that's what you get access to, by having something like a Kintaba installed in your organization.

Beyond that, you'll also get access to setting up the roles for the responders in your company, so you can see at any given moment, who's my on call for PR? Who's my on call for engineering? Maybe who's on call for operations? Are broken down by different sub orgs, depending on how my organization is structured. Really critical data that you need for when dealing with a response to an incident. And then forward, what are all the learnings we've had for incidents, right? What are the pieces of knowledge that have been gained? These are published, we call them post mortems. The document that the engineer has written up after the incident that helps us all understand what was the root cause, and what is the process piece about our organization, the systemic change that we're going to make, to make sure that that incident never happens. All of that information goes into document libraries that everyone in the company has access to it.

And then going deeper, things like defining automated remediation steps, right? Every incident we have, say, if it's tagged security needs to follow certain steps, we have an automation tool set that lets you go in to find things that happen based on metadata for that incident. So, I have a SEV1 security, we make sure that we post in the security process documents directly into that chat channel, we assign the correct on call rotations and ping them and make sure that they show up. Maybe we even go and call some internal tools through webhooks. All of that is stuff that you get out of the box with a Kintaba, and then ultimately at the end reporting, to go and look at this stuff from a trending standpoint. And it's reporting not just about what are the incidents that I've had, what are my meantime resolutions, but also the people reporting. Who are the people in my company who are responding to these incidents? How often are they doing it? What is the heatmap look like for when my incidents happen? Are they happening Monday at 2 AM? Or are they happening midday, on days, when we do a push. All of that information kind of gets pulled together into a single space, so that you can react from a systemic standpoint to make sure your organization gets better and better at having fewer and fewer of these sub ones.

[00:13:39] JM: As an incident is occurring, describe the more technical and more technical level, the usage of Kintaba and, like at a lower level in more specific terms, is it assisting me? What information is it presenting to me? What is it allowing me to collaborate on?

[00:13:58] JE: So again, from an orchestration standpoint, Kintaba is going to, if you're a technical responder, is going to alert you to the fact that the incident is happening. Or if you're the person filing the incident, it's going to give you that interface to go and file the incident. It's going to direct you to the place where the work is being done to go and solve that incident. So, if that's a Slack channel, or if that's in the Kintaba chat experience, either one, it'll direct you there. And then it will give you all of the relevant metadata about the incident that's known. So, who are the other responders who are here? What roles are they playing? Are these people from a specific org? Are they representing a different part of engineering, a different part of operations? It will show you all of that information. It'll give you an interface to ping other responders.

So, if you've been brought in and you recognize that we need to bring someone new in who has an expertise in a specific area based on the incident happening, the interface is there to go and ping that responder, and make sure they're notified and also brought into the incident. And then other information around status of the incident, is it open? Is it mitigated? Where are we? What's the current status that's been reported? It's going to give you all of that information around the incident collaboration space. Kintaba won't solve the incident for you. That's not really incident management's role. Incidents, by their definition, are black swan events. They're unexpected. There isn't really a preset set of steps that will solve the problem. But it will give you the space and put you into the right context, to rapidly be able to work that problem and log that process towards remediation, so that as a technical responder, you have access to all of that data around what did you do? What were the important moments? What were the milestones that were hit? And then let you pull all of that after the fact and do a post mortem that can almost be pre written for you, where you can then record that change that you're going to make systemically so that the incident doesn't happen again.

[00:16:00] JM: Can you talk about what the metrics that I should be tracking around my incidents, in my incident response, in order to get an understanding of how my organization is responding to incidents on average?

[00:16:18] JE: Yeah, so there are kind of two sets of metrics. There's sort of the traditional metrics, which we sort of call the, the MTT stars, right? So, all the meantime twos. When an organization tends to get involved in incident management, they start to immediately think about things like, how quickly are we responding? We call this mean time to resolution. We're not huge fans of those numbers, actually, at Kintaba. In a world of black swan events, trying to work out an average looks pretty messy, and it actually encourages some pretty poor behavior, right? If you're measuring an average on response time, it really disincentivizes things like filing the incident at all, until you get farther through to understanding what the remediation is because you want to keep that number low.

The metrics that we really encourage, that I think are much more positive and much more effective for organizations, number one is you actually want to see your total number of incidents being reported going up. Because the likelihood of you missing incidents in an organization that hasn't really implemented incident response practices yet, is pretty high. You're probably having a lot of little mini maybe SEV3, SEV2, kind of low high priority incidents inside of the organization, and you're not capturing them. The problem with not capturing them means you're not learning from them.

So, what you really want to see is an increase. We have a talk we give called more and more and more. Why you want more incidents in your organization. And the counter to that is while you want to see the total number of incidents being submitted and increase, you want to also see a decrease in the most critical incidents, your SEV1s. So, if you're catching incidents earlier, if you're catching these SEV2s and SEV3s earlier, and you're seeing an increase in overall incident volume, you should also start to see a decrease in major catastrophic situations, your SEV1s. And that's kind of the chart you want to see, increase in overall, decrease in SEV1s. That's kind of a top-level view, you want to watch over time to make sure that your practices are getting better.

Outside of that, you also really want to be watching metrics around incident occurrence. I mentioned earlier that we have a heatmap for when incidents happen, right? Very similar to if you've used Google Analytics, and you're looking at sort of when people visit your website. Conversely, to that, you really want to be knowledgeable about when incidents tend to happen

to be able to identify any trends there. Your incidents tend to happen within 24 hours when you do your weekly push, do they tend to happen at midnight, because you're allowing code pushes to happen all over the place that are unchecked at hours of the day that force people to wake up and things go badly? All of that knowledge is super important in order for you to be able to take kind of top-level systemic steps to say, I want to make some changes here when incidents are happening, to prevent them from happening at particularly inconvenient times. Because incidents are human events, right? They require by their nature, humans to wake up to do things about them.

Similarly, you'll want to track metrics around who the people are, who are responding. Traditionally, organizations don't really track incident response as part of performance review for employees, which is a really dangerous practice, right? Because incident response is something that employees are having to practice. And if you're waking up the same employee every night as a responder at 2 AM, and then not tracking that as part of their job, you're really not doing a great job as an organization of understanding that person's work and effort that they're putting in. So, the other metric you want to have outside of response times, counted incidents, watching incident, SEV1s going down, and incident occurrence is you want to understand the trends around the people. Are you putting the burden of incident management unreasonably on specific individuals in the company? Are there things you can do to go and spread that burden to make sure you don't cause burnout? Or frankly, just cause a decrease in productivity from those people and happiness overall.

[00:20:10] JM: You worked at Facebook, can you give me some context for how your perception of incident response was shaped by Facebook?

[00:20:19] JE: So, Facebook, I think, is one of the more impressive organizations when it comes to incident management, because they've done such a good job of opening awareness to the process across everyone inside of the company. There’s an internal tool there, it's called SEV Manager that everyone in the company has access to, that provides this great thousand-foot view into what's currently on fire. And it's both engineering fires, things like site outages in certain parts of the world data, centers going offline, problems with databases becoming corrupt. But it's also major incidents that are non-technical. So, policy situations, PR situations, and other types of legal and non-technical things that are currently putting the business at risk.

And what was so great about that culture of openness around things that are currently critical and sort of emergent in terms of risk to the organization was it really helped other people from the company come in and help where they could. If you don't know about something happening and legal where you could potentially be helpful, even if you're not in that organization, you'll never go and offer that help. But the openness there really helped the organization, practice kind of this responders who we know are necessary to go and deal with the incident as well as attracting in other people who could be helpful and giving them access to know where to go, and who are the people currently working that problem where they can be helpful. I think that openness around the process really helped Facebook become what I would consider probably one of the most reliable organizations out there, right? The site simply doesn't go down, by and large, and when it does, it's a worldwide event.

[00:22:01] JM: So, what about the tooling side? And what kinds of tools does Facebook have to handle incident response when you're in the line of fire?

[00:22:10] JE: So, the tools are very similar to what we've actually recreated with Kintaba. The tooling looks like a set of dashboards, collaboration, tooling, on call rotations, automations, for dealing with repeatable situations, post mortems, and how they're written and where they're stored and distributed. It's that entire flow, has all been written into tooling to make sure that it's done in a clear and repeatable fashion. The front of that tooling really being that a SEV Manager interface where you can declare the situation and provide visibility across the organization for everyone to come together and work that problem.

[00:22:48] JM: How important is it for the incident response tool to be connected to other platforms like Datadog or any other monitoring platform, logging platform? Is important to be connected to metrics and logging and analytics? Or is it more like a just a kind of a standalone place for information around incident response?

[00:23:16] JE: I think in most cases, it's actually the ladder. It's really helpful to have data from external sources. Datadog is a good example, around metrics, right? Especially top-level metrics, where say, for any incident, having to do with the engineering org, you really want to pull in, say, an egress chart. Being able to pull that in quickly, and having it available is pretty nice, and is pretty helpful. That said, when incidents are happening, the types of data that you

need access to tend to be the ones that are a bit unpredictable. And the reason for that is, if you were able to catch the fact that the incident was happening through an existing metric you were tracking, odds are you probably already have an automation or process in place to deal with the fact that that metric is moving, right?

So, if you have an overload on specific servers, and that's being tracked in your core metrics, you probably already have an auto scaling script that ought to go out there and deal with it. Where the incident would come into play would be that auto scaling script has fallen over and all of your servers are now overloaded, and the incident has been loaded up and launched, and we've thrown a couple of these metrics into it. But most of what needs to be happening in the incident space is more ad hoc work. It's the unusual actions that the responders have to take to dig into what's going on. And they'll share I think a lot of that data in real time as it's discovered. You might laugh at this, but we find a lot of that data comes in almost better as just screenshots, as point in time data of what they've captured and what that responder is trying to show and share.

So, for example, we have a really nice interface with Datadog, where you can go and pull this metric data. And it's always been a great selling point, it looks really great in demos. And when we see it get used from time to time for responses. But to be quite honest about it, the act of responding to major incidents is more human, then it necessarily is, kind of metrics driven, or sort of like expected metrics. Where I actually see more interesting relationships is some of the kind of more investigative tooling, Honeycomb has a really great product, I think, for this, and sort of observability that actually think does this a little bit better, where you'll go and dig, dig, dig, dig, dig, and then want to share that piece of metric data. And as you found it, for maybe where the root cause is and what's occurring. That's a little bit of a strange answer to that question, because I think the natural response to folks going into incident response, when they're building out their automation tooling is they think I need to go and throw as many charts as I can into this thing. And the reality is, you've already got your chart dashboards, and they're probably already really well built, and they're probably in different tool. And in those worlds, you're almost better off linking out to that information than you are trying to pull it all into live collaboration space, where it's going to go and get pushed off the fold.

[00:26:14] JM: How does an organization properly categorize and I suppose a annotate and create a historical repository for incidents for future learning and for post mortems and stuff?

[00:26:30] JE: This is one of the challenges that actually prevents organizations from adopting proper incident management practices in the first place is, when you start to think about how will I store everything that's ever happened, you start to get a little bit crazy about all of the structured data you want to demand if you're incident responders, right? The natural thing to do is to say, let's get almost a Microsoft Access type list of 10,000 different dimensions of information about this incident that we want you to write down, whatever what happens. And based on that, we'll get these great trend charts and understanding of everything that's happening at a really tight resolution.

What we've really found is, that's not a very healthy way to do things, because ultimately, again, people will avoid filing incidents because of it. They'll say, “Oh, this situation is not important enough, I don't want to fill out these 10,000 pieces of structured data.” It turns out, the better way to do it, is to try and capture high level data in real time that can be used for kind of high-level bucketing of information, right? So, for example, in Kintaba, we use tags really effectively. As an incident is being created, and as it evolves, what are the different dimensions around the organization you want to attach? And these tags can look like things like security, in terms of the organization being impacted, or they can look like things like data types. You can say, like PII and have a tag for that. And those tags can then trigger automations, which can do things like add other tags, or add other pieces of metadata to the incident.

But it turns out, if you do a pretty good job tagging, if you do a pretty good job of filing into the right severity level, and then you audit and store the full conversation that's happening, while marking important moments and milestones, you actually get out the other end of it. A pretty great database of searchable information, where your search actually becomes the more important factor, right? Can you full text search against all of the incidents that have ever happened? So, if an incident is currently going on, and you want to check a phrase, or a word or an area affected, that's going to be the best way to go search for it, right? It's sort of the Google approach to finding information in a haystack of incidents, as opposed to, kind of the old web approach to things which was way back when you use like an Alta Vista and you had to go and burn things down by category.

So, we've really found that you want to keep that barrier low for how hard is it to put basic information in, and then do a really aggressive job of indexing it. And if you do that, you'll get a nice library. We also provide a library view into post mortems, which is sort of a live document view of all the post mortems that have been written, which also has like a real time full text search against it. So, you can bury that down to say, okay, let's find everything where we ever really talked about, Dev cluster 105, and type that in and pull everything that comes up that has that in its category.

Because a lot of the time, there's sort of two things you want to do with your incident repository. And thing, number one is operational use of it. An incident is happening, and you want to be able to reference back to your incident repository aggressively to go and find similar incidents. That's one. And then number two is trending, and trending needs to really happen kind of at that higher level, right? Where are my tags? And where are my severity levels? And can I track those at a really high level? You don't have to go super deep when it comes to actual trending. So, if you do those two things, you can get a really effective repository of information that can help you make decisions, both real time during an incident that's happening, as well as kind of high-level company decisions to see if you're moving the boat in the right direction.

[00:30:03] JM: Can you walk me through the engineering stack of Kintaba? I'd like to get a sense of some of the technical problems that you've had to tackle and what's been hard about building it.

[00:30:17] JE: Yeah. So, at its core, Kintaba is a web interface. It runs on node and next. And it's a real time chat and collaboration interface that has a bunch of services attached to it, to go and help it do things like reach out to responders in real time. I think we use Twilio to go in do like our SMS and phone call outreach, as well as Twilio's SendGrid integration to do our emails and make sure everyone's notified. But underneath all of that, we're actually built on Azure, and we use a lot of technology in Azure to be resilient ourselves. We're built on Kubernetes. We cross multiple regions. We have a very controlled release process through you know, development, staging, test branches, as well as automated CI and deployment processes. We employ just about every kind of modern CI and resilience engineering tactic in our own builds, to kind of make that product be able to do all of the things it needs to do.

It's interesting, in incident management, people will tend to interact with it. I talked about the collaboration platform before often being Slack. It's also often that people want to interact with their incidents through email. So actually, the email interface and integration becomes one of the really critical portions of the product, because you'll have emails flying back and forth to Kintaba that are then being synced back into the chat system and being then forwarded into Kintaba's chat and then synced out to Slack. So, you end up with kind of some interesting pieces to the product there where it's just making sure that all those integrations are tying into a real time experience successfully. So, that no matter how anyone is absorbing the information, whether it's through email, through SMS, through a Slack integration through the Kintaba UI, all of those things are staying in sync really successfully.

We also have a pretty well built out, we call it an Async Tier, it's just processing these requests and keeping up to date with them and has fallbacks and various bits of resiliency on itself. So, that even if something like the web interface were to go down, those jobs continue to run, and people still maintain updates around the incidents that are happening. I can't give you a perfect overview. I'm not our CTO. We'd have to bring Cole on here. But that's kind of a high-level overview of how the things put together. The Async Tier, also drives the automation system, which is what does sort of, if this, then that triggering inside of incidents to – we call it turning people into cyborgs, right? To make people more effective and make incidents be able to update themselves as appropriate, where there are predictable steps to be taken.

[00:32:47] JM: How do you handle design? What's your process for product design?

[00:32:53] JE: We try to use really human and accessible design approaches. I mentioned before, that a good incident response process isn't just used by the engineering team. We try really hard to build human oriented design. I think one of the almost backhanded compliments someone had given us on Hacker News, when we first launched, it was that looked a little bit like Facebook, which I thought was funny. I think it was meant as an insult. But we took that as a bit as a compliment, right? Because you enter your incident management systems should be accessible and easy to use, as a social network or something that's expected to be used by anyone in the world.

We've worked with a series of designers, primarily out here in New York and in Brooklyn, who have all just been really focused on making this as easy as possible, making things one or two clicks away, making sure that interfaces for doing edits don't look like table data, and it doesn't look like a technical requirement, and you're never writing code, right? You're always like using drop downs and selecting people and roles through type of heads. So, we've tried really hard to keep the UI as friendly and interactive as possible, which isn't always the case for development tooling. I think incident management tends to start its life in the SRE and orgs, even if it spreads to the rest of the organization. And so, we have to walk that line a little bit. But I think we're always making changes there to kind of just find the sweet spot there for our customers.

[00:34:19] JM: Can you give me a sense for, I guess, what it's like to design the product, and kind of your ongoing way of looking at how customers interface with it. So, I'm sure you have kind of access to how people are using the platform, and how people are annotating their incidents and moving through their incidents. I just love to get a sense for how the feedback loop between seeing incidents being responded to in action, translates to changes in the product.

[00:34:56] JE: So, we try to live our own PR. So, when I talk a lot about more incidents is better, and we try to sell that as something that's really valuable for organizations. It's kind of the core metric that we watch for customers that come in. We actually expect to see a pretty rapid ramp up of incident creation, inside of the organization, when they come on board, and we expect to see that kind of happen in tandem with a growth of the user base. And if we're not seeing that, then we start to look into metrics a little bit deeper into why. If we are seeing that, we look pretty deeply into metrics of where those incidents are originating from.

I think a lot of people assume incident management tends to originate from automated systems, right? Like a metric happens and an incident is kicked off for it. But that's not really how it happens. Most incidents are human declared, and they tend to be a reaction to a dashboard somewhere either being bright red, series of metrics coming back as bad, or worse than that, a dashboard being green and a customer's reporting that things aren't going well and things are going actually not what you think they are. So, a lot of our work is a balance of how do you make the barrier to entry for incident creation low enough, that people are willing to start to use a tool like this, if they haven't really practiced in the past? How do you make sure that initial

page that they see that asks them to file an incident as simple as possible? And how do you really kind of march people towards that point? Does the product work when there are no roles set up, because maybe you just came in when you had an incident, you needed to do it right away, and so now you need roles in real time, right?

A lot of that kind of data is really what helps us understand how to get you up and running. It's more common for a company to be coming in and starting to use Kintaba to have really had no tooling before that other than running to a Slack channel. It's much more rare that a company comes in that is migrating from an existing set of tooling. So, we have to also be the educators as much as we can be without throwing wizards at you, or like, let's learn about incident management. We have to be the educators visually, in terms of what's important here in this incident management tool now? What's important, and is it easy for you to do that?

For most people, what's important is get your employees in there, get the roles defined, and get an incident filed, but just not always in that order. Sometimes it's get that incident filed, and while you're responding to it, get your employees added and then start giving them roles. That's really the stuff we watch. Everything beyond that, once the incident is up and running, once we know that you're filing them, and you're kind of moving through that process, we certainly don't look into data around what customers are filing as incidents. We stay out of that. We're compliant with PII the way we ought to be, and we're certified for that.

So, we're mostly looking again, at high level data beyond that. Are you filling out your post mortems? Is the barrier to hide there? Can we lower that barrier as much as possible? Even a one sentence, post mortem is better than nothing. So, we're really watching that kind of zero to one onboarding, and then sort of the follow on actions of ensuring that you are able to feel comfortable post incident of sharing information with the rest of the org. Beyond that it's high level metrics, right? Are people opening the reports page? Are they using things? Are they finding them? Do the wordings make sense? Is the documentation up to date? But we really care about that zero to one. Incident management is an emerging industry, and whenever you're working in an emerging industry, you can't assume that people even know the steps to take. You have to provide those to them in the UI and through the design.

[00:38:29] JM: Well, as we begin to close off, I'd like to get a sense for the destination of Kintaba, and where do you see the future of incident management going?

[00:38:41] JE: So, I think we're a bit unique, and that a lot of companies in this space, see the world moving deeper and deeper into the technical side. Becoming a deeper and deeper SRE, sort of a practice of incident management. I think we really kind of see it as what we call almost a shift left of incident management, which is more and more portions of the organization participating. A success for us has always been an entire organization adopts Kintaba and begins to practice it, not one or two people, but the whole company. And we're seeing a lot of success there. Our North Star has always been all organizations having incident management and all people at those organizations having access to it. That's the direction we're running. I think what it really does is it brings incident management not just closer to the non-technical portion of the company, which is pretty critical, but also it brings incident management closer to the customer. Not our customer, but our customer's customer, right? The closer we can bring incident management to the customer of our customer, the better we'll be at being a tool that just makes organizations better at a top level, like more successful, happier, operating more resiliently. If we can do those things, I think that's very much what the future of this industry looks like. We've certainly talked to folks who think that incident management has to be a subpart of another tool, right? Or needs to be like within a department and I very much disagree with that. I think incident management as a company practice, and that's really the shift that's happening in the world today that Kintaba is riding the wave of.

[00:40:26] JM: Cool. Well, John, it's been a real pleasure talking to you and I look forward to seeing Kintaba develop.

[00:40:32] JE: Great, thanks, Jeff. Just to say it, we mentioned IR comp at the beginning, would love it if other folks can sign up for it. It's a free conference. It's at irconf.io.

[END]