

## GDEVs - [SOTR075] SEARCH OFF THE RECORD - 75TH EPISODE - VIDEO

[00:00:10:27] - **John**

Hello, and welcome to another episode of Search Off the Record, a podcast coming to you from the Google Search team discussing all things Search and maybe having some fun along the way. My name is John, and today we're going to talk about trouble, and for that, we have Gary and Lizzi joining us.

[00:00:32:21] - **Gary**

Are you suggesting we are trouble?

[00:00:34:09] - **John**

No, no, of course not. I would never do that.

[00:00:37:12] - **Lizzi**

What kind of trouble? I feel that you must be more specific because this is a broad topic.

[00:00:42:02] - **John**

There's a lot of trouble.

[00:00:43:10] - **Lizzi**

Yeah, which trouble?

[00:00:44:27] - **John**

Do you mean we could get multiple episodes out of this?

[00:00:46:22] - **Lizzi**

Potentially.

[00:00:47:24] - **Gary**

Oh, God.

[00:00:50:14] - **John**

So, specifically, we're talking about Search, so maybe trouble in search, which could be a lot of things: maybe you get a bad search result; maybe you get no search results.

[00:01:01:10] - **Gary**

That would be trouble. That would be bad.

[00:01:02:29] - **John**

Or maybe you get really old search results.

[00:01:05:09] - **Lizzi**

Has the homepage ever been down? Like the entire google.com/search?

[00:01:10:29] - **Gary**

Okay, that's not the homepage.

[00:01:12:14] - **Lizzi**

Okay, sorry. Correct me.

[00:01:15:01] - **Gary**

The homepage is protocol google.com/. That's it.

[00:01:20:04] - **Lizzi**

Okay, so that page, has that ever been down to your knowledge?

[00:01:23:14] - **Gary**

I mean it can go down, but it's mostly like a routing issue. Not like our servers go down.

[00:01:29:17] - **John**

Oh, it's like someone else's problem, not ours.

[00:01:33:22] - Gary

Yeah. The service that hosts the homepage is the same thing that hosts the status dashboard, the Google Search Status Dashboard, and it has like an insane uptime number.

[00:01:48:16] - Lizzi

What's an uptime number?

[00:01:49:27] - Gary

How much is it reachable in a given period of time?

[00:01:53:24] - Lizzi

Like there's an 80% chance I can reach it, so therefore it's possible.

[00:01:59:23] - Gary

Yes, but the number is like 99.999 whatever.

[00:02:04:16] - John

Nine.

[00:02:05:18] - Gary

Nine.

[00:02:06:21] - Lizzi

Nine.

[00:02:07:24] - John

Nein, it's never down. Nein.

[00:02:11:21] - Gary

Okay. We have a German in the house.

[00:02:13:14] - Lizzi

So the possibility of the Status Dashboard being down is the same likelihood as the Google Search homepage?

[00:02:20:23] - Gary

I mean, ballpark. Yeah, ballpark.

[00:02:21:28] - John

Cool. Okay, so Status Dashboard. I think we're kind of getting ahead of ourselves. Like, when something happens that affects Search, like I don't know, like looking at the Status Dashboard, there were some issues with regards to, I think, indexing or crawling. Do you remember any of those recent issues?

[00:02:41:15] - Gary

Yeah. Why?

[00:02:44:10] - Lizzi

This is actually your performance review. We're going to go over how it went.

[00:02:48:10] - John

So what was it caused by? By you? Not on the record.

[00:02:53:21] - Lizzi

Off the record.

[00:02:55:03] - John

Off the record.

[00:02:56:16] - Lizzi

Were you just stumbling around in the--

[00:02:58:14] - Gary

Now I actually feel like I'm being interviewed for like keeping my job or something. I don't actually remember when it was, but the last one that I actively managed, I guess, that was on February 1st.

[00:03:11:06] - **John**

So what happens? Is it like someone in a data center stumbles over a cable and then something breaks?

[00:03:18:08] - **Gary**

Site Reliability Engineering org publishes their playbook on how they manage incidents. And a lot of the incidents are caught by incidents being issues with whatever systems. They catch them with automated processes, meaning that there are probes, for example, or there are certain rules that are set on monitoring software that looks at numbers. And then, if the number exceeds whatever value, then it triggers an alert that is then captured by a software like an incident management software. That's what happened on February 1st as well. Basically some number went haywire, and then that opened an incident automatically internally. Then we have to decide whether that's a false positive or it's something that we need to actually look into, as in like we, the SRE folk. And, in this case, they decided that, yeah, this is a valid thing. And then they raised the priority of the incident to one step higher from whatever it was. I think it was a minor incident initially and then they raised it to medium. And then, when it becomes medium, then it ends up in our inbox. So we have a threshold for medium or higher. Yeah. Looking back at the minor incidents, historically, there's never been anything that was worth talking about externally. So, when you look at those, it's like, "Oh, one of the racks ran out of RAM, and we need to supply more RAM to these tasks in those servers, for example. No one is going to notice externally that that particular rack is out of RAM or, I don't know, like one of the data centers is drained and we are still routing traffic to it. It's like it's still going to work.

[00:05:12:02] - **John**

What do you mean drained? Is it water cooling?

[00:05:14:20] - **Gary**

As in it's drained of traffic. Basically they divert traffic away from data centers so they can do maintenance tasks or whatever on the servers.

[00:05:25:01] - **Lizzi**

Oh, like scheduled maintenance.

[00:05:26:20] - **Gary**

Yeah. But if, I don't know, some configuration file is not updated, then it might be still sending traffic there and then that might cause problems for us internally. Because their data center is not receiving traffic, it will just divert it again to somewhere else. But, internally, it still looks like there's traffic coming in, so we can start the work on the whatever we are planning to do.

[00:05:52:27] - **Lizzi**

So, in this case, would we know that this is not a real problem or do we still have to investigate? We being like either the SRE team or us. We still need to go through the motions to see, is this actually a problem that is practically going to affect anyone? Or do we know at that point that this is fake news?

[00:06:10:00] - **Gary**

SRE would investigate everything. If they get a probe alert, for example, or an alert based on whatever numbers, they will look into it and will try to explain that to themselves. And, if it's something that is affecting users, then it almost automatically means that they need to raise the priority because users are actually affected.

[00:06:37:14] - **Lizzi**

How many users?

[00:06:38:24] - **Gary**

I don't have a number.

[00:06:40:06] - **Lizzi**

Do they have like some way to check this threshold? Like, how many is enough in order for it to be "Okay, this is considered a problem"?

[00:06:48:16] - **Gary**

I don't know. I never worked as an SRE, so I don't know.

[00:06:55:05] - **John**

What happens next? Like, the SRE team looks into it and says, "Oh, this is a real thing." Do we go off and just publish all of these on a Status Dashboard, or what happens as the next step?

[00:07:08:07] - **Gary**

No. If we publish everything, then it our dashboard would be very noisy. And that's because some of the things are not affecting users enough from our perspective, like from Search Relations perspective. Like, for example, recently there was an incident where some images were missing. If I remember correctly, then I stepped in and I said like, "This is stupid, and we should not externalize it because the user impact is actually not bad," right? Users will literally just not get the images. It's not like something is broken. They will just not see certain images on the search result pages. And, to me, that's just, well, back to 1990 or back to 2008 or something. It's like it's still usable and still everything is dandy except some images.

[00:08:01:08] - **Lizzi**

It's just a degraded experience.

[00:08:04:02] - **Gary**

Yeah, but I couldn't tell that the experience is degraded.

[00:08:08:14] - **John**

So, is the focus more on what users would see or what site owners would see? Because, as a Search Relations team, we would focus more on site owners. But it sounds like you're saying, for these issues, we would look at what users would experience.

[00:08:26:10] - **Gary**

So it's Search Relations, not Site Owners Relations from Search perspective.

[00:08:33:00] - **Lizzi**

But by extension, like the site owners, they would also care about their users. So, if we care about their users, it's the same group of people, right? Or is that too positive?

[00:08:41:19] - **John**

I see issues where people come to us and were like, "In my Search Console, this graph looks weird."

[00:08:48:04] - **Gary**

Right.

[00:08:49:08] - **John**

And then, first, I imagine we look into that and say, "Oh, is this Search Console," which is just weird, which sometimes is the case, "or is this actually something in Search?" which is weird. And then it could be something where a lot of people are seeing these weird things in Search Console, but an actual user wouldn't see, I don't know, a big site Search Console account. So, would that be something that we would flag, or is that something that Search Console should just like track?

[00:09:24:14] - **Lizzi**

What if we go back to the February thing?

[00:09:27:08] - **John**

Okay, let's go back.

[00:09:28:11] - **Lizzi**

So that one we did post about.

[00:09:29:16] - **Gary**

Yeah.

[00:09:30:19] - **Lizzi**

Why.

[00:09:31:22] - **Gary**

So that was a degraded experience for both users and for site owners. For users, if I remember correctly, it was results were getting old. We were struggling getting in new content, which meant that, for example, news portals or news sites were having trouble reaching their users, which also meant that, in Search, you would not be able to find

the latest news about your favorite topic, so it was affecting users and also site owners. And we, Search Relations, generally know how much effect one particular system's failure can have. For example, if the indexing system goes down completely, okay, that's going to be noticed. Or the crawler engine just goes away, okay, that's going to be noticed. And then you just make the call that I want to externalize this because it's impossible that people will not notice it and lots of people will notice it. There's also the flip side where we are like, "Well, we don't actually know if this is going to be noticed," and then two minutes later there's a blog that puts up something about "Google is not indexing new articles anymore. What up?" And I say, "Okay, let's externalize it."

[00:10:58:23] - **John**

Okay, so if there's more pressure on us externally, we would externalize it?

[00:11:03:12] - **Gary**

Yeah. For sure. Yeah.

[00:11:04:25] - **John**

So the louder people are externally, the more likely Google will say something?

[00:11:09:28] - **Gary**

I mean, in certain cases, yes, but it doesn't work all the time, because some of the things that people perceive externally as a failure on our end is actually working as intended.

[00:11:22:04] - **Lizzi**

Like what?

[00:11:23:06] - **John**

Oh, like we're not crawling your site as much as you want, but that's not a bug. That's because we just don't want to crawl.

[00:11:32:02] - **Gary**

Right. That's a feature.

[00:11:33:08] - **John**

Well, I don't know. I think as a site owner--

[00:11:35:25] - **Gary**

Well, from my perspective, from my perspective, it's a feature.

[00:11:37:25] - **John**

Well, as a site owner, you might be like, "This is very important content. I like when Google appreciates fresh content."

[00:11:46:21] - **Lizzi**

I think in this case it would have to be that you notice a trend, so there needs to be multiple, more than three, more than seven, sites that are having this problem.

[00:11:56:19] - **John**

Okay. I think that makes sense. Sometimes we get feedback after big ranking updates, like core updates, where people are like, "Oh, everything is broken."

[00:12:06:13] - **Gary**

Do we? We get feedback like that?

[00:12:09:15] - **John**

Well, yeah.

[00:12:13:06] - **Gary**

Oh.

[00:12:14:14] - **John**

I feel bad for them. I kind of understand that. I think those are the kind of situations where we would look at the examples and be like, "Oh, I see some sites are unhappy with this, but overall we're doing the right thing from our perspective."

[00:12:31:02] - Gary  
Right.

[00:12:32:09] - John  
And then we wouldn't see it as an issue, right?

[00:12:35:01] - Gary  
Yeah.

[00:12:36:08] - John  
Going back to that incident that you were looking at, basically, we noticed there was an issue. It was large enough.

[00:12:43:01] - Gary  
So that was actually a golden incident, from my perspective, because an alert was triggered internally that something is going bad, something is utterly wrong. And, at the same time, someone from the Google News org reached out saying exactly the same thing, that would happen if that particular system that was alerted for does something stupid. So basically we got proxy external alert and an internal alert, pretty much at the same time. The timing was like 10:13 in the morning the alert triggered. And then at 10:18, we got the email from the Google News org saying that publishers are writing in about this particular issue. Right after that, we need to make an assessment whether we want to externalize. In this case, my assessment was "This is probably not good," and that's a quote.

[00:13:37:27] - Lizzi  
Did you know right away that it was a problem with indexing.

[00:13:40:13] - Gary  
Yeah.

[00:13:41:21] - Lizzi  
How?

[00:13:43:00] - Gary  
Well, because the alert was triggered for indexing.

[00:13:44:24] - Lizzi  
For indexing, okay. So, in our first posting, we say like, "We're investigating an issue with indexing, but we don't know what is the root cause yet." So the first step is identifying the root cause, which is kind of interesting.

[00:13:56:11] - Gary  
So the system is named in the incident like the first email that we got. Again, I'm not reading out the email, but it's saying that indexing is pushing back on traffic. Basically it cannot process new documents.

[00:14:12:16] - Lizzi  
Okay.

[00:14:13:20] - John  
Oh, okay. Basically, I guess, intuitively, your first step would be where in the Search processes is this problem. And, in this case--

[00:14:16:13] - Gary  
It was named.

[00:14:23:21] - John  
We were lucky because the initial alert basically said "This system is experiencing a problem."

[00:14:29:24] - Gary  
Yeah, and then, two minutes later, I responded to that alert with, "This is probably not good," which was an understatement. And then Danny steps in--Danny, as in Danny Sullivan, the Search liaison--asking whether this alert can cause the issues that the news publishers are having. And then I just confirm that, "Yep." And then basically, from there on, I just started our internal process for externalization because, in my mind, it was noticed enough that we have to say something about it.

[00:15:04:06] - **John**

And then, after you started the process with externalization, you went off to fix the issue.

[00:15:09:27] - **Gary**

No.

[00:15:11:01] - **John**

No?

[00:15:12:18] - **Lizzi**

But you claim it. You say, "I am the captain of this issue."

[00:15:16:27] - **Gary**

Yes. Yes, so you're in charge, but not really? Well, in that case, I was in charge of the external communications. So basically I made the calls about what do we say on the dashboard and when. Our messages are typically templated, mostly to speed up the posting process so we don't have to run around like a headless chicken, seeking approvals from Legal and from PR and whatever to get out new messages. They just pre-approve messages that we can put up when we deem necessary. And that's basically what we were doing when I could not post something, because I think I was at a conference in Denmark or something, then John posted something or put something up, then I think you, **Lizzi**, you also put something.

[00:16:04:28] - **Lizzi**

This was, when it was going on for like a week, we were having to post. This is kind of a long running issue in our history of posting about these kinds of things. Usually it's like two or three days. But this was, I think, because it went over the weekend, the engineers who are working on the fix, they were also taking the weekend.

[00:16:21:23] - **Gary**

Yeah, I think it was also different from the perspective that we didn't close the incident as soon as it was mitigated, because the system that was affected in the indexing process was critical enough that if it regressed, then we would need to reopen it and that would just look stupid. I think I decided that we just keep it open and just say that it's mitigated. Like, "Right now, there is no visible issue related to this incident, but it's not actually fixed. We just put a Band-Aid on it for the moment, and we are still working on sewing up the thing to make sure that it actually gets fixed properly." Historically, we usually just close the incident because it was like, "Eh, it's not going to come back," and usually it didn't.

[00:17:11:05] - **Lizzi**

And I think it usually reaches a state where it's fixed enough. Like, even if we have to go and reindex everything, we still will close the issue and then say that, "Okay, eventually, you will notice that things are coming back."

[00:17:24:26] - **Gary**

Yeah.

[00:17:26:04] - **Lizzi**

But, for this case, we distinguish between fixed and mitigated.

[00:17:29:09] - **John**

I think that sometimes makes sense because some of these processes just take a long time. Where, if there is something, I don't know, stuck in crawling and then we fix that, it's going to take a bit of time for that to be visible in the search results and to be 100% sure that actually it's fixed for everything. Sometimes I think that makes sense. Basically, to recap, some alerting system rings a bell, the Site Reliability Engineering team looks at the details, figures out what actually is happening, if they can, or at least confirms that it's a real problem. Then it reaches someone on the Search Relations team or someone generally involved with communications who tries to estimate is it something that users and site owners will see. Based on that, we'll open an entry in the Status Dashboard. And then, basically, all of the fixes are happening in parallel. It's not something that the communications teams are involved with. It's like some engineers are running around trying to figure out what actually is happening, figure out what needs to be plugged in or changed and the configurations.

[00:18:45:10] - **Gary**

Yeah.

[00:18:46:14] - **Lizzi**

Can you share more about what the fix was or what the cause was for this issue that happened in January?

[00:18:50:21] - Gary

I don't remember. You don't remember.

[00:18:53:02] - Gary

I remember that it was a system that I previously worked on. Every now and then I ducked my head in the SREs chat room and I was like, "Perhaps we should look at this particular subsystem of the thing, because I used to work on that particular part and I know that it can like manifest an issue like this," but I don't think that that was actually the culprit. It was a red herring.

[00:19:24:06] - Lizzi

Oh.

[00:19:25:23] - Gary

Often what happens is you try to eliminate all the suspect systems.

[00:19:35:06] - Lizzi

Okay, and then if somebody launched something or pushed out some change and then that could potentially be the most recent thing that happened.

[00:19:41:22] - Gary

Yeah. Very often there are experiments that need to be rolled back because two experiments might interact very badly with each other. Like, some alert is triggering that something is going wrong. And then you root that issue back to faulty experiment or a bad interaction between experiments. I mean, sorry, that's not us. And then they would roll that back or ramp down the experiment. There's also like data pushes that need to happen every now and then. For example, if you think about, this is a wrong example, but the doodles that we have every now and then, they are part of data pushes. And then, if something goes wrong with the data push, like let's say we lose, on the wire, 100 bytes, and then the doodle is like in the middle, there's a line, a white line, that would be noticed by our monitoring systems that something looks really wrong on the search result page. And then someone would go there, would take a look, it's like, "Oh my God, what's up with that white line?" And then they are like, "Okay, so when did we push this doodle?" And then they would identify a time, what time is it, it was 2:33 PM push, and then they would roll back to the push before that.

[00:21:12:29] - John

But it's not always something that some engineer at Google pushes a button and launches something and then it breaks, right?

[00:21:18:25] - Gary

No.

[00:21:19:29] - John

Sometimes it's, I don't know, the internet changes slightly and then our systems aren't ready for that.

[00:21:26:06] - Gary

I mean, we also had issues with cosmic rays.

[00:21:28:04] - John

Cosmic rays.

[00:21:29:08] - Lizzi

What?

[00:21:30:13] - Gary

Yeah. Yeah.

[00:21:31:16] - Lizzi

What?

[00:21:32:19] - Gary

Yeah. True story.

[00:21:33:22] - Lizzi



What do you mean?

[00:21:34:24] - **Gary**

I don't know if I can talk about it because I don't know if we publicized it, but maybe we should look it up, and then we could.

[00:21:42:07] - **Lizzi**

Cut something in, In some bonus episode about the cosmic rays.

[00:21:46:14] - **Gary**

No. Maybe we could have an episode about insane incidents, but we would have to check what we publicized before. But, yeah, we did have an incident about cosmic rays, and we know that it can cause troubles. There was also election systems that were thrown off by cosmic rays. That's not a new thing to computers. It's just like, in this case, it was very unfortunate.

[00:22:16:00] - **John**

I read, I think, a non-google postmortem, which I think the these reports afterwards are called--

[00:22:23:04] - **Gary**

Yeah.

[00:22:24:06] - **John**

About an issue where playing some specific rock song very loudly in a data center caused hard disks to break.

[00:22:30:28] - **Lizzi**

What?

[00:22:32:14] - **John**

It was really weird. I imagine finding those connections afterwards is hard. It's like some engineer is looking at a dashboard and sees these hard disks are failing and doesn't realize that some technician at the data center is playing music really loud. Lots of weird, surprising connections.

[00:22:52:11] - **Gary**

Technically, anything can cause an incident.

[00:22:54:26] - **Lizzi**

Like a butterfly's wings.

[00:22:56:00] - **Gary**

Yeah, exactly. And we have fantastic people who can find these issues, or at least they can mitigate the effects of the issues. And then we also have people who can restore systems from, like crash code or crash dumps, because they can speak computer.

[00:23:16:05] - **John**

And the Status Dashboard itself is something that, from our perspective, also helps site owners to understand that, "Oh, something is weird at Google. It's not my fault."

[00:23:27:28] - **Gary**

Yeah.

[00:23:29:01] - **John**

I mean, it might still be their fault, but it's unlikely to be their fault.

[00:23:33:15] - **Gary**

Yeah. I mean, that's also one of the main reasons why we put up the Status Dashboard. It's not just about informing Search users about issues, but also informing site owners that we are rolling something out, for example, like with the ranking updates, whatever we call them nowadays. I would really want to extend that to more things. Like, for example, if we decide that now we want 40% more videos indexed, then we should tell people that.

[00:24:03:21] - **Lizzi**

So, for the future of the dashboard, would you be hoping for, I don't know, us posting more informational updates, not just something as broken.

[00:24:14:01] - **Gary**

Yeah, because then you can relate changes in how your site appears in Google Search to something that we've done.

[00:24:21:17] - **Lizzi**

It could be like a positive change. Not everything is like negative, broken, on fire. It could be that we've rolled something out that's good, or we're now indexing at a faster speed.

[00:24:30:24] - **Gary**

Yeah, I think still the majority would be negative.

[00:24:34:15] - **Lizzi**

Yeah.

[00:24:37:04] - **Gary**

But at least some of them would be positive. I think one of the main things that we could have covered last year with such an informational update, maybe not last year, but the year before, was how we started indexing short videos. But then that meant that we needed more space in the video index. So we decided that, "Okay, then we make space by kicking out low performing videos from our video index for the short videos." And then, for some sites, that was noticeable. But we could have said, an informational message, that, "We are decreasing our video index by 7%. It's not you; it's us."

[00:25:25:29] - **Lizzi**

Do you have any other wish list items of what you would want to see from the dashboard, John?

[00:25:30:11] - **John**

I would like to see a dashboard for Status Dashboard changes.

[00:25:35:03] - **Lizzi**

What? A copy of a copy.

[00:25:38:03] - **John**

Or do we call this our blog?

[00:25:40:05] - **Gary**

I'm leaving. I'm leaving.

[00:25:42:03] - **John**

All right. Cool. Yeah, I don't know, I'm happy to also hear from people externally if they want to send us some ideas or some thoughts on what we could be showing differently in the Status Dashboard.

[00:25:56:15] - **Lizzi**

I'm still waiting for someone to do something visual with the feed. I still haven't seen anyone pulling that in to then visualize it in their mobile-friendly way. Like however they want to improve that data. They could be reading it and then doing something.

[00:26:10:16] - **Gary**

Are you throwing shade on the mobile version thing?

[00:26:13:07] - **Lizzi**

I'm not saying anything. I'm implying something.

[00:26:16:05] - **John**

Oh well. Okay. I'm pretty sure that there are people out there, maybe at SEO round tables, that have a big dashboard that blinks and sets off alarms whenever something changes in our dashboard.

[00:26:36:23] - **Gary**

I'm pretty sure.

[00:26:38:15] - **Lizzi**

I will investigate.

[00:26:39:29] - **John**

Okay, cool. Well, thank you both for joining me here, for sharing all of this trouble.

[00:26:47:28] - **Lizzi**

Trouble.

[00:26:49:00] - **John**

Trouble. Yeah, it was good trouble. I liked it. Thank you. Thanks for listening, and goodbye.

[00:26:55:02] - **Gary**

Goodbye.

[00:27:00:13] - **Lizzi**

We've been having fun with this podcast and I hope you, the listener, have found it both entertaining and insightful too. Feel free to drop us a note on Twitter, @GoogleSearchC, or chat with us at one of the next events we go to if you have any thoughts. And, of course, don't forget to like and subscribe. Thank you, and goodbye.