



Daniel Newman: Spiros, welcome to the 2022 Six Five Summit. I am very excited to have you joining me today.

Spiros Xanthos: I'm excited to join you, Daniel. Thank you.

Daniel Newman: Observability is a really big topic right now. We are seeing more and more companies entering the space and we are seeing more and more enterprises adopting the technology as they're trying to understand better how their data, their systems, processes are up working, running, managed, et cetera, and so Splunk is at the center of this. And while over the past few years, the company has pivoted, it's changed its business model a number of different ways. It's definitely gained a strong foothold and momentum as being one of the biggest players in the observability space. Now, having said that Spiros, I sometimes think the word observability, it goes without full understanding. People say it and they pass it off as, oh, it's an observability thing, but I'm not certain that everybody fully gets it. And then those that do, I still think a lot of us could use a refresher. So let's start right there. Someone hits you up, walks up and says Spiros, what is observability? How do you answer that and make it simple for people?

Spiros Xanthos: Yes, definitely. Like you suggested, I guess observability is a term that few vendors have been using for a while, but I think now users have started adopting it and it has become definitely mainstream in the last two years, and we see this with our customers. Now, one, but what it does it mean? Essentially what has happened is with the evolution of IT systems, infrastructure and applications, and the move to the cloud, and especially the acceleration that we have seen of all that with the pandemic, we're dealing with a lot more complexity than we used to deal in the past. We used to have servers, some network devices, some monolithic applications running on them and simple monitoring systems which were exclusive to each one of these. Let's say I had a network monitoring system, I had an infrastructure monitoring system and I had my logs, that was sufficient to keep the systems up and running.

But on the cloud where you have oftentimes hundreds, thousands of components each changing multiple times a day with complex interdependencies, the tools of the past don't work anymore. So that's where observability comes into play, which is this idea that where combining, let's say, all these types of telemetry together so we can monitor infrastructure and application in one place and actually even monitor end users in their interaction with these systems as a whole so that we can much more effectively understand and troubleshoot these systems. Oftentimes, observability requires combining logs, metrics, traces, real user monitoring into one solution so that it's truly effective and can allow us to monitor and observe these complex infrastructure and applications.

Daniel Newman: Everyone out there can see why I asked you that question, because even when you try to take observability and make it as simple as possible, it still has a lot of complexity. And I think that complexity is largely rooted in the exponential data volumes that organizations have and of course the significant shifts in the way infrastructure and architectures are being modernized. You've got workloads being distributed everywhere. You've got new architectures, everything from your on-prem and legacy workloads to running cloud workloads on bare metal and everything in between. Of course, you've got containers and then of course all the data that resides.



And all of this has to be organized in some way that a company can hopefully see it and streamline it, manage it and make sure that systems are running and organizations are uptime and secure. In a nutshell, you just hit on it. You've alluded to the value of observability, but with what I just mentioned, you got the unprecedented data volumes and you've got this dynamic shift in the infrastructure, how do companies approach improving that visibility and really taking advantage of the opportunity that observability creates?

Spiros Xanthos:

Very good question. As you described, we have now a lot more complexity, but oftentimes, let's say, as an end user or as a business owner, my real goal here is to accelerate my velocity, focus essentially on my business and achieve my business outcomes faster and with confidence. So more change and more complexity at the same time. That's where, I guess, observability vendors come into play. Our job, first of all, observability is a data problem and we have both complexity and a lot of data, and that's where vendors come into play. We need to help, let's say, the end user achieve their goal of velocity and confidence while dealing with all these huge volumes of exponentially growing data. At the end of the day, what a system should be able to do is ingest all the data and then not just simply dump it back to the user.

It's not a matter of ingesting all of that and then maybe giving to the user just a simple user interface where the query is data and they validate or invalidate their own hypothesis, because that still would rely on them having the expertise. We're not doing much. I think what's happening with observability is that in addition to being able to handle this large volume of data, there is a structuring and, let's say, the analytics on top of it, what we oftentimes call AIOps. We can ingest all this data, structure it and start suggesting back to the user maybe where the problem might be or what the next action should be. It's not just a matter of ingesting, and, as I said, dumping the data back, but it's also a question of can we be more intelligent now and start guiding the user to where the problem might be. So they still have to have their own expertise and understanding of their infrastructure application, but observability solutions, modern observability solutions should be able to start suggesting to the user what the problem might be, as I said, and what the actions they should take is.

Daniel Newman:

It's interesting you mentioned that. And by the way, curious, but given your expertise, do you feel sometimes observability SecOps, ITOps, AIOps get conflated? As you're talking about it, how do you simplify that for folks? Because is observability the way that all that rolls up into observability or is there a way that you... Because I sometimes hear the people using ITOps, SecOps, AIOps and then observability almost interchangeably, and, of course, I don't think that would be an accurate depiction.

Spiros Xanthos:

Correct. I don't think that's an accurate depiction. And I think to some extent because observability has become, let's say, a trend, let's say, that a lot of now users are trying to follow, everyone is studying on it regardless of where they come from. In my opinion, to say that somebody's following observability practices, there are a few principles that have to be in place for that to be true. In my opinion, observability actually starts with actually open standards for data collection, if, say, vendors who do static monitoring or proprietary data collection that only works in their platform, I don't think they're true observability vendors because they operate in a closed ecosystem where the data cannot be connected. Secondly, I



think that as we rely on these open standards, we have to be able to correlate infrastructure and application and even end users.

So we have to break up the silos and see our entire infrastructure and application as one thing, because that's way more effective way of troubleshooting the applications. And finally, we should be able to move the needle with what we do with this huge volume of data. The tools should be able to assist us in monitoring and troubleshooting these systems proactively, not just provide, as I said, the data for us to do all the troubleshooting in our head. I believe that these are the principles of what modern observability solutions look like. And of course then there is some overlap.

AIOps, in my opinion, was a concept that was around for a while because it's the idea that systems help us identify and even resolve issues automatically. And I think modern observability tools make that actually simpler and actually possible because we have a lot more data that is structured where essentially analytics and machine learning can now be a lot more effective in identifying problems. Now, observability in the future, in my opinion, will start also overlapping with security as well because all this data we're collecting is very rich and it's not just useful for telling us that we have, let's say, a performance issue or some downtime, but it's also useful in telling us if an incident we're facing is potentially a security incident. It's a lot richer data to use for security analytics as well. We're going to see this, let's say, overlap more and more in the future, especially with modern security practices like SecOps.

Daniel Newman:

I think that's a good way to tie it together. I reemphasize that because one of my thesis has been that observability, we need to back up a little bit and break down these pieces, and so you did a really nice job there. You've mentioned a few times, I've mentioned a few times this exponential data issue and the growth of data and the challenges it's presenting. So as a vendor that's working with many of the largest companies in the world and beyond on their observability strategy, how are you helping your customers understand that first of all they have the right information to act upon and then second of all that they're investing in the right technology stack to build out, to fully embrace the opportunity of observability?

Spiros Xanthos:

Splunk has in some way always been in the observability business, even before we called it observability, especially when it came to IT operations monitoring, because Splunk as a platform has this ability to ingest all the data, let's say, no matter where they're coming from. And that, in my opinion, one of the most important factors in having effective observability. You don't know in advance what data you're going to need and when are you going to need them when you have a security incident or when you have some availability problem. In my opinion, one of the most important things that we should be doing as vendors, and that's a principle in which we operate as a company, is collect all the data. So unlike, let's say, traditional APM, that solutions that sample heavily, they collect maybe something like one in 1,000 transactions, which is good enough to understand the trends of your environment, but obviously if you're looking for something specific, chances are you're not going to have it because you have sampled it out.

We believe in what we call full fidelity. So we collect all the logs, all the traces, we collect high fidelity metrics as well. So you have all the signals in one place so that when something goes



wrong, you can quickly connect, let's say, the symptom to the problem and to the root cause with all this data. And that's something we can do really at any scale as a vendor. And I think that's a requirement and something that customers should be looking for when it comes to observability. And one more thing I want to point out here is another thing that we're proud about at Splunk is that we're the co-creators and the biggest vendor in terms of support in OpenTelemetry, which is a CNCF project that tries to standardize data instrumentation and collection for, let's say, observability use cases, try to standardize traces metrics and log collection into a single set of standards and implementation.

And OpenTelemetry is today the second most popular project in CNCF, second only to Kubernetes in terms of contributions, and it's becoming the standard. We are big believers in standardizing all of that for the benefit of the users so that they have control of their data and they can choose whatever vendor brings the most value in terms of processing and handling that data.

Daniel Newman: It's really important these days. I think lock in is one of the things a lot of companies get concerned about, having the opportunity to maximize data, be flexible. Lock in and agility are like opposing forces for a lot of companies, and so when you're a DevOps environment, you need to stay flexible. That doesn't mean you don't want to be sticky as a vendor, but it just means you want to make sure you're constantly adding value and the customer never feels like they're trapped with a inferior technology or solution as the improvement keeps happening.

I got a couple more questions for you, Spiros, and thanks for spending the time. I think you've added a lot of clarity here and I'm hoping everyone out there is getting a lot of value out of this. The talent in this particular space has to be a challenge right now. We're seeing the economy slow down a little bit, but labor supply is very tight and this is a very specialized capability. How do you advise your enterprise customers to build a pipeline of talent that can take advantage of observability technologies like what you're offering at Splunk?

Spiros Xanthos: By the way, we just did a study called the state of observability, where we surveyed more than 1200 professionals in the space to understand the correlation, let's say, between up time and observability practices, or talent retention observability practices. And what we found actually in this particular topic is that companies who have advanced observability practices, who, let's say, have adopted observability and have a mature implementation of it generally are better at retaining and attracting talent. Observability itself is a newer thing and I think it's attractive to smart engineers and SREs who want to stay current and improve, let's say, their careers. What I see is that companies who, let's say, do even simple things like create an observability team and advertise that they have one and try to recruit people to work in observability and of course are serious about their observability practice certainly tend to attract and retain talent better because the best engineers want to work on things that will be relevant in the future and find ways to improve their knowledge.

Definitely, adopting observability and having an advanced practice, it helps, and more broadly, I would say, obviously we see the great resignation that has happened in the last year, I think that being obviously in technology these days and having a critical skill like observability definitely



helps you land a great job. So for us as maybe managers or companies, I think what we should be doing, in addition to offering great benefits and compensation and all of that, is focusing a lot on the culture and focusing on actually offering great opportunities for advancement to our teams. In addition to the things that everyone else can offer as well.

Daniel Newman:

I think some of it's technical, Spiros, and some of it's just common sense leadership. People in almost every field, there's some discrepancies and differences in how people want to be managed, and there's some things that are consistent. They want to develop, grow, have chances to earn, have chances to learn, to get more certified, so I think there's some things here that are somewhat leadership common sense, but I do like the stuff on a technical basis that definitely can create this. And I think for people out there, this is a field to think about going into that's going to continue to grow and be in really high demand. And I know that when you're out there thinking about which college route you want to go to or what you want to study after, it's good to think about where is there going to be a lot of demand.

And that the demand around data, data science, DevOps pipeline, observability, is going to continue to grow. I like to take a moment at the end of these conversations, Spiros, to just look at the future. And so while there's more vendors entering this space, more competition, there's also going to be a lot of growth in the overall TAM, the adoption that the utilization of this technology. For companies that are either looking to get into observability, what are some of the things you suggest they prioritize to get ready, and what advice would you give them in terms of getting the most out of their observability initiatives?

Spiros Xanthos:

For companies that haven't yet started going down this journey, my position is to start, I guess, as soon as possible, because it is a journey and it's going to take some time for the practice to mature within an organization. I think an opportunity to start is usually when maybe they start migrating an application they have, some workload they have to the cloud. So they start dealing with, let's say, the complexity of the cloud and they require better tools. They can start there. They can start adopting, let's say, OpenTelemetry for these new workloads and maybe looking into combining, let's say, infrastructure, application monitoring into one solution. And obviously complexity is only going to increase, so I think that it's a requirement if they want to maintain, let's say, velocity and confidence. In evolving their applications, it's a requirement that they improve their monitoring practices as well, and start moving into observability.

And as I said, based on the observability report that we conducted, we see that it makes actually a huge difference in how much money it costs for companies who have adopted advanced observability practices, actually save a huge amounts of money in terms of downtime, because they're much more proactive in troubleshooting and maintaining their uptime, they're much better at attracting talent, and eventually they're actually much more successful in their digital transformation initiatives, because I guess they're better covered with the tools to be able to make the changes they need to make. I think observability is here to stay and I think it's beneficial really to everybody to start understanding what it is, and if they have already started down this path, to keep maturing their practices so that they can keep up with the complexity of the cloud and modern apps.



Daniel Newman: That's a great way to end it. We agree. We see this as a significant growth area, significant expansion, companies are going to need to be able to maximize and understand the data at scale much faster, and this is going to be one of the key tools. Spiros, thank you so much for joining us here at The Six Five Summit 2022. Look forward to having you back on the show sometime soon and congrats with the role and a lot of success in your future. We'll see you soon.

Spiros Xanthos: Thank you. Thanks for hosting me.