



DevOps Open Source Monitoring Helps Reduce Costs and Increase Uptime

Q and A with Sean Mack, CEO, xOps

Question: **How does xOps work with a customer?**

Sean: We first send in team which includes a management consultant and a DevOps engineer, who have implemented DevOps in large enterprise environments. Our team works with the business to figure out what DevOps can mean for the company.

We discuss organizational structure, and talk about all the different moving parts related to application roll-outs. We offer a change communicator to work with the organization, and then we've got engineers who can start to work with the teams, and maybe even off-shore some 24/7 work. We also build a roadmap to look at what can be more automated and where we can supply machine learning. We're building the whole CI/CD pipeline for our clients, everything soup to nuts.

This is the standard approach for large engagements but we do unbundle DevOps, and offer point solutions to enterprises trying to do DevOps. Often large enterprises don't actually have people who are full-stack engineers, and we help provide some of the best DevOps engineers on the planet.

We try to cut down on effort and expense for clients. We automate a lot of operations and we've got a great offshore team who can offer the same services at a lower cost.

Question: **How do the open source tools help customers become more secure, self-heal, find problems faster, and keep the business running at a lower cost. And what kind of people do we supply that are associated with those tools?**

Sean: Our open source presentation and correlation tool, xView supports the core DevOps cultural shift, which is basically, the idea of a cultural bias to working collaboratively. In the modern enterprise we often have very siloed data and information. xView helps to bring that data together. Operators can see it in a cohesive way that's shared with not just the ops team, but the dev team as well. And maybe not just the dev team, but the business team. So you don't have a report that's generated off sale data that goes to the business. You have a real-time report that goes to the business, the dev teams, the ops teams. They are all looking at the same data at the same time.

And you implant this notion of collaboration, of DevOps, in the tools you build. What that means is you've got quicker time to resolve issues. I don't know how many times I've been on a call where the network guy's saying, "Oh, it's not my problem!" And the database guy's saying, "Oh, it's not my problem!" If you can ensure they're all looking at the same screen, that shows the same red dots. That cuts down on that conversation a lot.

And that's just the basics. The other piece of that is integrating business data in that same platform. You're looking to all your log data, all your metric data, all your customer data including Tweets. For example, if all of a sudden, there's a spike in Twitter users going, "Hey, I can't access...", "Anybody else having trouble?" the tech team isn't looking at that today. But they should be. It's a good indicator of a problem. You can feed that data into the tool. This results in reduced time to identify infrastructure issues. We have all of our alerts, all our data in one place.

Question: Don't some big players already offer this?

This isn't new. It isn't something that hasn't been done before. What we offer is similar to what IBM and HP are doing, and have been for years. But they want you to drop a million dollars of hardware in each data center. They want you to have a team of five people just to maintain the system. And they're going to tie you into their other products. I discussed this with HP and they said, "Well, you don't really need our other products, but you do need to copy all your data into our CMDB." And I said, "But I'm using ServiceNow for our CMDB." And they said, "That's fine, just also put a copy in our CMDB." I said, "Guys, this is crazy! I'm not going to have two copies, two CMDBs, just to use your tool."

We are offering an open source version. We put it right on top of the Elastic stack. It's built on top of something that's open-sourced so it's rapidly implement and scalable.

What really ramps up efficiency, is when you start to deploy machine-learning on top of it. You take all of your incident data, which also usually isn't considered when you're considering all the data you've got in your operations, but ServiceNow logs when the incident happened. Now we can we start to get the machine learning to look for the last time this was trending this way, this was trending this way, and this was doing this, we had this major key won. Voila – we can predict outages.

The machine can now say, "I'm seeing these trends. You may have an incident if you don't do something now." So now you to predictive analytics, where you're actually fixing the incidents before they even happen.

Question: How does xView help reduce downtime?

Sean: xView provides tremendous value to customers. For example, we had a situation at a client where every Tuesday, at about 10 A.M., their machines would start flapping, and we had to adjust everything and redistribute the load manually, and we couldn't figure it out. We had to do a deep dive into all the data, from all the machines, in order

to find out the exact situation root cause. You can only do that if you had all the log data and all the metric data, and all the delivery stats in one place, to really analyze that and prevent it from happening again.

#