CSO Insights
FROM IDG

Issue: **November 2018**

THE COMPLETE GUIDE TO IT MONITORING

Credit: iStock

Produced by IDG

Credit: iStock

Monitoring station

It's a broad term, but monitoring in all of its forms has come on leaps and bounds over the past decade, with IT teams now able to collect and make sense of vast amounts of data into how their infrastructure, systems and applications are performing, with the boundaries increasingly starting to blur as enterprising business users look to get their hands on this technical data.

Monitoring may not be sexy, but it is integral to keeping key applications up and running and the best tools on the market are increasingly being enriched with machine learning capabilities to get ahead of issues before they hit users.

Log data and monitoring capabilities have also become key components outside of the traditional IT function, with many security operations centres (SOCs) reliant on monitoring tools to spot hacks and breaches in record time, saving on costly damage to the organization had they gone unmonitored. Read on for an example of how the Bank of England completely rebuilt its SOC around a set of cutting-edge monitoring capabilities.

Here we assess the monitoring software landscape, focusing on machine data analytics and application performance management (APM) vendors and use cases, honing in on a couple of key commercial providers such as Splunk, Dynatrace and AppDynamics (which is now part of Cisco), some real world case studies and why good monitoring is so important to UK organizations as varied as Just Eat and Vodafone. **Scott Carey**

Contents

Credit: iStock

Best APM software

We run down the best application management software tools

Application performance management (APM) software is increasingly important as developers shift towards continuous delivery and customers expect consistent service from their apps.

These APM tools allow developers to monitor and track the performance of apps to spot and resolve performance issues. Modern day tools also veer into the automation space, using machine learning to identify critical issues and even resolving them before a human has had a chance to be alerted. These tools can also be leveraged to help improve the user experience, spotting pinch points or areas of the app where customers generally drop off, allowing developers to fix bugs or tweak the application accordingly.

Typical APM metrics include the transaction volumes, response times and errors, often visualized in dashboards, scorecards and reports.

Here are the best APM tools on the market, from established cloud vendors to pure play solutions.

Amazon CloudWatch

Amazon is well-positioned because so many developers build their apps on AWS – so picking a monitoring platform that's part of that stack means everything, in theory, works in step. CloudWatch can be used for monitoring your AWS services, but also the applications running on those services. The AWS tool is the leader in G2 Crowd's Grid for APM software, with positive reviews pointing to its ability to autonomously react to app events with triggers or Lambda events. It can get expensive when used at scale, though.

Pricing: There is a generous free tier for CloudWatch for three dashboards of up to 50 metrics, 10 alarms and 1 million API requests each month. After that it is priced by what you use starting at $3 (about £2.38) per dashboard per month, and $0.10 (about 7.6p) per alarm per month.

AppDynamics

AppDynamics can be deployed either on-premise or as-a-service (SaaS). The company was acquired by Cisco in January 2017 for $3.7 billion. The software can be used for baseline monitoring and alerting as well as end-to-end monitoring to track customer journeys. AppDynamics has proved popular with UK enterprises, with Barclays, Paddy Power, the British Medical Journal and EasyJet being public exponents of the software.

Pricing: There is a free Lite tier and AppDynamics Pro costs $300 (about £228) per unit per month, according to G2 Crowd.

Google Stackdriver

Stackdriver is a collection of APM tools, including the newly added Stackdriver Profiler, which lets developers profile and explore how code executes in production, as well as standard Monitoring, Trace, Logging and Debugger tools for analysing, logging, alerting and debugging apps.

Profiler is in public beta at the time of writing. Google has built sound integrations between Stackdriver Debugger and GitHub Enterprise, and offers code mirroring for GitHub, Bitbucket, Google Cloud Repositories and locally-stored source code. Stackdriver can be used on GCP or AWS cloud infrastructure (but not Azure) or even on-premise.

Pricing: Stackdriver is free at the basic tier for up to 50GB of log data per month with basic email alerting and read-only API access. The premium tier is $8 (about £6.10) per resource per month.

LightStep

One of the newer entrants to the market is LightStep, which came out of stealth mode in November 2017.

The start-up was co-founded by Ben Sigelman and Daniel Spoonhower, who both used to be software engineers at Google, and Benjamin Cronin, who was the CTO at design agency Pusher.

The company's first product is called LightStep [x]PM and is built to monitor and diagnose anomalies across web, mobile, monoliths and microservices, regardless of infrastructure. LightStep can also be deployed in staging environments to keep issues out of releases.

Early customers include Silicon Valley technology companies Twilio, Lyft and Yext. LightStep has also been funded by Silicon Valley venture capital giants Sequoia and Redpoint.

Splunk

Machine data and IT logs specialist Splunk also provides an application management tool as part of its eponymous platform, so if you are already using Splunk for logs you have a good head start. It will give you a view of baseline app performance and alert developers to reduce mean-time-to-resolution (MTTR). The firm is investing heavily in machine learning across its entire platform, so expect smart alerting for issues and proactive resolution guidance.

There is a Lite version of Splunk for smaller clients or for those wanting to take the software for a test run.

Pricing: Application management comes included with Splunk Enterprise, which is priced according to data consumption, and starts at $2,070 (about £1,578) per one-year term for 1GB of daily index volume (including support) and down to $690 (about £560) at 100GB of volume.

New Relic

New Relic is considered a leading SaaS APM vendor by both Gartner and G2 Crowd. It is easy to get up and running and can be used to monitor Ruby, PHP, .Net, Java and Python apps, plus it provides analytics to improve the user experience of applications with metrics for crashes and user drop off. New Relic is also investing in machine learning to deliver smarter alerts and quicker routes to issue resolution under the brand New Relic Applied Intelligence.

Pricing: New Relic is priced according to instance size and is calculated in relation to CPU Cores + GB RAM x hours used. The Pro tier typically ranges from as little as $18 (about £13.70) per month, up to $200 (about £152) if you are running apps on AWS, for example.

Dynatrace

Dynatrace software can be used either on-premise or as-a-service to proactively identify and fix problems with app performance through the whole technology stack, and spot end-user experience issues. Dynatrace has been investing heavily in AI through its platform to bring smarter analytics, full automation and self-healing capabilities to customers. It is a flexible solution that works across all major environments and technologies.

The vendor merged with fellow APM firm Keynote in 2015. UK customers include Thomas Cook, Travis Perkins and Thomson Reuters.

Pricing: Dynatrace starts at $0.035 (about 0.26p) per host per hour for the self-serve solution and scales up depending on data volumes. There is also an enterprise version, or monitoring solutions for web-scale organizations, where pricing is available upon request.

Stackify Retrace

Retrace from monitoring pure-play vendor Stackify specializes in performance and error monitoring, and active alerting to reduce the time of resolution for bugs and errors. You can also use Retrace to profile an app's performance during development.

The SaaS solution is lightweight and easy to get up and running across a variety of cloud environments. Plus, it can be used to monitor SP.NET and Java web apps.

Pricing: Retrace is priced $25 (about £19.60) for a single CPU or $50 (about 39.20) per multi-core CPU.

SolarWinds

SolarWinds provides software for both server and application monitoring. For APM it's a flexible solution that works out of the box for major cloud environments, and across a wide range of enterprise applications across data centres, remote offices, and the cloud. SolarWinds also lets users proactively monitor and alert on performance issues and simplify troubleshooting from a single dashboard. The dashboards and UI aren't the best in class and deployment can prove tricky.

Pricing: Starts at £2,275 for up to 150 monitors for one year. **Scott Carey**

Credit: Splunk

Splunk expanding beyond IT and security users

The machine data specialist is looking to grow beyond its core user to 'bring the power of Splunk to more people'

Machine data specialist Splunk is looking to expand its user base beyond IT and security professionals with a new approach to product development under what it calls Splunk Next.

Speaking to us before the vendor's .conf event in Orlando last month, head of product marketing Jon Rooney said: "One of the guiding principles is to bring the power of Splunk to more people."

While Splunk CEO Doug Merritt said in a press release: "We are in the midst of the data revolution, and these product updates ensure the Splunk platform evolves as our world does to deliver business outcomes no matter the organization, team or data set."

Splunk started life as a big data analytics platform for IT professionals looking to index and make sense of their machine data. It has since pivoted to security, helping teams do the same to combat cyberthreats.

"A lot of our success has been bottom up through IT and security practitioners," Rooney revealed, "which is great and still the core of our DNA, but there are a lot of types of users and functions across a business that maybe don't know about Splunk."

Building an experience layer

In terms of actually delivering this, Splunk has made announcements for things like a revamped iOS mobile app and natural language search functionality. The new app includes alerting and simplified authentication, and the ability to take actions directly within the app.

"If you're at a sporting event, you get an alert from Splunk and you don't have to open a laptop, VPN in and do stuff," Rooney said. "You can look and interact with dashboards and take an action like shutting down a port, for example, on the phone."

Natural language search "abstracts away the complexity of working with our search language, or things that are a bit more hands on keyboard for our power users," he revealed. "We want to create an experience layer for the non-super users to get value out of Splunk." This includes allowing users to interact with Splunk data via a voice assistant or a chatbot on Slack.

This shifts Splunk into an analytics market that is far more competitive, with the likes of business intelligence (BI) pure plays Tableau and Qlik to reckon with.

"We always want to go from our position of strength because the value of Splunk is never going to be a beautiful dashboarding experience on top of tightly structured data," Rooney argued. "Our special sauce is our ability to correlate and make sense of data you couldn't jam into a relational database and put another BI tool on top of. It is still the idea of bringing insight to chaos." However, he added: "We will always have the notion that people who want access to the bare metal get access to the bare metal. We are always going to have the super user experience and a lot of the enhancements to Splunk Enterprise are focused on that."

One example of Splunk moving into different business areas is Splunk Business Flow. Currently in beta, it is aimed more at the traditional BI space. Rooney used the example of an online product manager who wants to know how a customer is progressing through a product, and being able to surface that "in a way that is more drag and drop and more visual".

Splunk's recent foray into the Internet of Things (IoT) is another example of the vendor looking to reach new audiences, focusing its first product in the industrial IoT space with predictive maintenance for manufacturing customers such as BMW.

"We did a lot of market research and talking to customers on what is the right beachhead for us and where we could get started," Rooney explained. "Obviously there is tonnes of opportunity, but we didn't want to be an all things to all people player in IoT."

Splunk Next

Rooney was speaking more broadly about what the company is calling Splunk Next, a more forward-looking approach to product releases. This will bring “the power of Splunk to more data sources and more people no matter where, when or how they access that data to deliver limitless insights,” he enthused.

“We have updates to our entire portfolio, including [new additions to the portfolio] Phantom and VictorOps, but we are also doing a series of initiatives that are really about expanding the footprint of what we can do from a product perspective,” Rooney added.

As a result, the first two beta products announced at .conf were Splunk Data Stream Processor and Splunk Data Fabric Search. The first allows customers to do more with their data while it is in motion, instead of waiting for it to hit the index. The latter is a highly scalable search functionality that works across indexes.

Leveraging acquisitions

Splunk announced a lot during .conf, and Rooney put that down to its recent raft of acquisitions, which has aided growth of the product team at Splunk by 55 percent in the past year.

“It’s twice as much product as we announced last year because through acquisition we have picked up a lot of product teams,” he said. “So not just Phantom and VictorOps, but companies like SignalSense and Rocana. We wanted people that were born in the cloud and big data, so that’s given us a ton of product velocity.”

VictorOps is a collaboration tool for devops teams to speed up issue resolution, while Phantom brings automation and orchestration capabilities to security teams.

The next thing for Splunk is taking those acquisitions and helping them broaden out their target audience. “I think the plan for all of those acquisitions is to expand it more," he said. More specifically: "While Phantom as we acquired it was a security company, we are going to take those underlying capabilities and apply it to multiple use cases starting with IT. Likewise with VictorOps, you think of it as a collaboration and incident response platform for IT, we want to make sure we leverage those same capabilities for our security customers."

Machine learning

Splunk is also continuing to invest in machine learning across its portfolio to help users get to their insights faster. "We will continue to make investments into machine learning with new versions of user behaviour analytics, which does insider threat detection, anomaly detection for security folks," Rooney explained.

"Which is really important as the talent shortage continues and people can't hire enough, so how do you start to leverage technology to take care of the first line of defence? That model of interaction and the idea that you only boil up the action for the security professionals where you only really need a human brain."

Splunk is also looking to give users of its IT Service Intelligence product more predictive alerts and actions with its 4.0 release. "People want to be able to see the health of their systems, looking ahead and trending forward to alert you ahead of time so that you can remediate that problem before it happens," Rooney said.

Lastly, the vendor is expanding its Machine Learning Toolkit by allowing customers to share, shape and build algorithms from GitHub community contributions, extend contributions and functionality from TensorFlow, and use a new connector for Apache Spark to tap into the MLib library. **Scott Carey**

Credit: Dynatrace

Dynatrace focusing on software intelligence

Vendor shifts away from APM to software intelligence

In a bid to meet customer demand, application performance management (APM) software vendor Dynatrace says it is shifting from app management to focus on software intelligence at scale.

Speaking on stage at its first pan-European Perform conference, John Van Siclen (pictured above), CEO at Dynatrace said: “We all know digital transformation is driving every company to become a software company.”

He gave the example of GE Digital, the General Electric technology arm that focuses on the Internet of Things (IoT). However, it is critical to understand that not only are companies shifting to software but there is an urgency to make sure the software works perfectly, Van Siclen explained.

"We're reinventing the platform and the application model at the same time with microservices, and we're thinking: 'well that's not hard enough, so let's start writing code and releasing it every hour, maybe every minute and to keep up with Amazon, maybe every few seconds'," he added.

Now Dynatrace wants to create a platform for speed, agility and scale that will develop with companies as they go through that transformation.

"Not only is digital transformation cloud-first, but as it relates to us, the cloud has changed everything because all the monitoring that we used in the past… it actually doesn't work, it doesn't deal with scale of the cloud, it doesn't deal with the dynamism of the cloud or the new application models," he revealed.

What next?

In order to deliver in this 'cloud-first' world, Dynatrace has developed its entire software stack to adapt to this.

From what may have previously been a $3 billion APM business, the company is now setting its sights on becoming a $15 billion monitoring, business intelligence, devops space.

"Software taking over the world is a massive change, and at the core of all this, we believe intelligence is at the centre. Software intelligence is where it's heading and we believe that Dynatrace can be that software intelligence for the autonomous platform," he enthused.

Already, the firm boasts of over 500 percent growth of its new all-in-one performance management platform since the launch in 2017, with 70 percent of recent bookings being for the platform.

At the conference, the company made two new product announcements – management zones and log analytics – putting the company into direct competition with machine intelligence software vendor Splunk.

Management zones is a new capability built to provide easy to consume software insights for enterprise teams. The product is designed with automation at its heart, with the ability to automatically provide environment information from the orchestration layer.

The second is the launch of log analytics, an all-in-one solution that is automated and built-in context. It includes log files, which are picked up automatically. Dynatrace is also keen to tout its AI capabilities, surfacing the most important insights automatically.

During a live demonstration, Florian Ortner, chief product officer at Dynatrace said: “All of that, more or less can be solved with Dynatrace out of the box. Dynatrace AI can pick up the root cause of the problem.”

The two products are now available for user access online. Customers can sign up for a free 15-day trial at **fave.co/2OtrzE1**. Hannah Williams

Credit: iStock

Vodafone standardizes monitoring around Splunk

Telco looking to leverage more machine learning to boost uptime

Vodafone is rapidly centralizing its IT monitoring and event management globally around Splunk tools, enabling IT operations teams to get better uptime from mission critical applications and to leverage more machine learning capabilities to avoid incidents.

Speaking at Splunk Live in London earlier this year, Luke Bradley, senior manager of engineering and operations for Technology Shared Services at Vodafone, said: “IT is increasingly becoming more and more the mechanism by which we differentiate ourselves within the market.

"Being a telco is being a telco, the services you can offer over the top of that is what makes the difference. We are very much on our digital transformation journey, this introduces an additional set of requirements to really put analytics and data at the centre of what we do."

The technology shared services group is an internal division of Vodafone group, which provides IT services globally, such as service desk, infrastructure management and application operations, across 26 geographies with 8,500 employees alone.

Over the past four years Vodafone has been standardizing its infrastructure and application performance management (APM) monitoring around Splunk tools, including all event management onto a single IT service management (ITSM) platform.

"So we have been trying to create a single operational analytics platform that sits on top of all of this stuff. We are really trying to get to the point where we have a single store of operational data," Bradley added, regardless of user group, geography or use case.

That re-architected monitoring programme now covers 40,000 servers and 3,500 applications, with 430TB of data capacity, but also saw Vodafone flip the way it approaches monitoring.

Bradley explained: "We have traditionally taken a bottom up approach to monitoring, as most organizations have – the operations team has traditionally struggled to map issues to something of real significance – so as part of that monitoring transformation we have turned things on their head. We are taking a top-down view of business services for all markets. We are standardizing the concept of a business service across those countries, so the act of visiting a mobile phone shop in Düsseldorf is logically the same as in Dublin."

ITSI and machine learning

That single ITSM platform is now monitored using Splunk's IT Service Intelligence (ITSI), so that the operation team has full visibility over the project. Talking in more detail about the ITSM platform, Stefan Ciobanu, product owner for analytics and big data solutions at Vodafone, said: "We are now running one of the largest IT service management platforms."

Vodafone's ITSM platform has around 13,000 daily users generating 2,000 tickets per day. "So this platform can't go down," he said. "Having this level of tickets on a global scale, you have to have a solution to monitor it perfectly and ensure uptime is 99.99 percent minimum."

"This empowered operations to predict downtime and make sure the platform is running at higher capacity," he said. "By doing so we gained our objective of a higher availability platform, so giving operations teams the correct vision inside all of our services. By using the predictive solutions in ITSI we managed to do preventative maintenance and helped avoid incidents."

Vodafone is now looking to roll ITSI out onto other mission-critical services. Now that Vodafone has this single monitoring capability, it is looking to start leveraging more machine learning for smarter and more predictive alerting and issue resolution.

"With the machine learning toolkit on Splunk we are building a community of data analysts to deploy more and more predictive maintenance on our tools and applications," Ciobanu said. **Scott Carey**

Credit: Just Eat

How Just Eat avoids costly weekend outages

The food delivery service uses AppDynamics monitoring to avoid outages at weekends and upsetting hungry customers

Popular food delivery service Just Eat has turned to AppDynamics to get a high-level view of its web and mobile apps as it looks to cut down on costly outages at peak times.

Just Eat has more than 10 million customers in the UK, and with the growing popularity of food delivery apps, as well as its recent acquisition of rival Hungry House, that figure is only set to rise.

Just Eat is, amazingly, more than 12 years old. In a busy market where upstarts such as Uber Eats and Deliveroo are aggressively expanding, outages at peak times can cause a huge amount of damage. These apps pride themselves on the convenience factor, meaning customers often respond to issues with an outsized level of outrage.

Take Just Eat's reported outages on New Year's Day, where a tweet from the company reporting technical issues across its website and app received more than 200 responses, some verging on the dramatic side. "You're ruining my life!" one user tweeted.

The ideal situation for Just Eat, then, is to get ahead of issues before they hit customers, and a solid monitoring function is a key part of that.

Adopting AppDynamics at Just Eat

This is where AppDynamics comes in. Just Eat adopted the application performance management software – which was acquired by Cisco in January 2017 for $3.7 billion – in the spring of 2017 and got it in front of the relevant users within three months.

Just Eat's infrastructure consists of approximately 400 microservices running across thousands of AWS instances in the cloud, with an open source monitoring and logging service, which has been tweaked in-house.

This rolls up to an alerting system, also open source, which allows engineers to reduce downtime by spotting issues before they affect customers.

AppDynamics then "fits on top of that stack," Just Eat's director of technology Richard Haigh told us. "We wanted a high-level view of all of those systems and understand which areas things go wrong," he revealed.

The aim was to have a “joined-up view across the whole estate to improve triage”.

Where AppDynamics really stood out from rival solutions for the Just Eat team was its ability to pinpoint ‘architectural dependencies’ or in other words, pinch points. This is important within a tech function that spans, according to Haigh, a complex estate with “hundreds of monthly releases” and “engineering teams all working at their own cadences”.

“So visibility is important and that dependency mapping allows us to spot things that might have made sense a year ago,” he explained. “We now have data for that.”

Previously staff relied on a range of monitoring systems and dashboards, without a ‘single pane of glass’ view. Now they can see the “most important systems in as good as real time, where they communicate and how they perform in one screen,” Haigh said.

So through AppDynamics, the operations team gets a performance baseline, whether that’s for a Tuesday night in August or a Sunday night in December, so it can spot any latency or potential performance issues ahead of time and drill down directly from the tool.

For example, Just Eat doesn’t use autoscaling for its AWS infrastructure, instead opting for scheduled scaling “as demand can be cyclical,” Haigh said.

This means that any unprecedented spike in demand or a software release that’s using up a bit more capacity than expected can be spotted early, allowing engineers to scale the infrastructure before users get cut off.

AppDynamics is now part of the toolset of the service operations team, who use it as one of their key dashboards. Developers get a view of their systems to assure what they have deployed is working as it should be as well as being able to test in pre-production – and lastly the tech managers use it to look at high level or business-focused dashboards.

What next?

Haigh did say he would like to see more automation within AppDynamics in the future. The dream scenario, he explained, would be: “To come in and see we have had an issue that could have affected customers and an automated system has mitigated it, and supplied a clear list of actions we can take the next morning.

“So if I don’t have to wake up a member of staff in the middle of the night, and no customers have been affected and we are left with a clear list of actions.”

This is where John Rakowski, director of technology strategy at AppDynamics, jumped on the line to talk about the company’s recent efforts to make the system more automated and intelligent.

The Cisco acquisition of Perspica in November will play a role here, too. The San Jose-based company specializes in machine learning and data processing technology and has joined the AppDynamics team.

“Our path is looking at systems of intelligence and automated response to get enterprise customers to that automated response point quicker,” Rakowski told us. Scott Carey

Credit: iStock

How the Bank of England built its 'SOC 2.0'

The bank has shifted how it runs its security operations centre in a bid to protect nearly £1 trillion in assets from 'unknown' cyberattacks

The Bank of England has drastically shifted the way it runs its security operations centre (SOC) from being reactive to more proactive, employing more data science techniques to answer one fundamental question: "How do you spot an attack when you don't know what it looks like?"

Attempting to answer that question last month during security analytics vendor Splunk's annual .conf in Orlando was Jonathan Pagett, head of the security operations centre at the Bank of England.

The UK's central bank, among other functions, is responsible for the UK's payment infrastructure, both by acting as the settling agent to allow financial institutions to exchange funds and by operating the CHAPS payment network, meaning the bank's SOC is protecting the trillion dollars (about £700 billion) or so that moves through those systems every day.

Two years ago Pagett and his team within the SOC decided that simply being reactive to security alerts would not be sufficient for an organization that is being targeted by highly sophisticated attacks that may not be known by their researchers or threat intelligence tooling.

"We knew attacks can bypass those security controls, so we asked how to detect those," he told the audience during a breakout session. "We also took a moment to ask what our strategy is, not just what technologies you have, but how it all fits together."

SOC 2.0

The result is what Pagett calls SOC 2.0, which was established in earnest late last year and comprises three elements: technology, people and process.

First there is the tech platform, which is underpinned by Splunk solutions. "We needed a tool to proactively search and spot behaviours of attacks," he said.

Next is people, with the current SOC consisting of 10 analysts, with a skill set that now varies from traditional IT to more data science practices to help identify the threats from within that massive data set.

This remains a major hurdle for the bank, though. "Recruitment is the biggest challenge," Pagett explained. “We have the tech and tools to be able to do this, but someone has to put that logic into the tools.” This includes recruiting from non-conventional backgrounds and participating in programmes such as the Cyber Security Challenge.

Pagett believes that the bank’s new proactive model does make it a more attractive destination from a recruiting perspective. “Our SOC people have a problem not a task,” he told us.

The result is that security analysts at the bank now spend around 80 percent of their time building up what it calls ‘attacker profiles’. This consists of modelling attack behaviours, so “using Splunk to write analytics that look for those behaviours,” Pagett explained. The other 20 percent of time is spent on incident response.

This is all underpinned by a new operating model that focuses on discovering unknown attacks and creating a repeatable method to defend against them. “We wanted a SOC doing daily, continual improvement to create new analytics to detect those attacks,” Pagett revealed.

This process starts with acquiring data (NetFlow, DNS, endpoint logs, access controls). Next is threat intelligence and research to create hypotheses for how an attacker might behave, instead of what the specific piece of malware might look like.

Then there is data mining, using bespoke Splunk searches and machine learning algorithms designed using the Splunk ML Toolkit. The last phase comes down to alert triage and incident response processes, followed by a wrap up that is focused on making that analytic repeatable.

The result is a culture of continual improvement within the SOC. “We are never going to finish, which is why we liked Splunk, to build a devops team within the SOC, so we can develop our own Splunk apps to extend that functionality," he said, citing a new voice assistant interface recently developed by a member of the team.

Results

Pagett admitted that it is difficult to quantify the success of its new SOC model as it is a process of continual improvement, but so far the team has developed 273 different Splunk searches, each associated with different actors.

These are reviewed daily in red team exercises with internal pen testers. They also triage their threat intelligence database, using a bespoke Splunk app to identify threats that haven't fired in over a year or been hit by their pen testers, flagging that vector for review.

In terms of impact on the business Pagett talked about the 2016 Bangladesh Bank cyber heist, which targeted $1 billion (about £770 million) worth of assets from the bank using vulnerabilities in the SWIFT payments network. When an attack like that hits, Pagett tends to get a knock on his door asking: "Could we be hit by this?" Having the ability to turn around and say they had it covered is certainly a positive result for the head of the security operations centre and his team.

Road map

The next thing for the bank is building more automation and orchestration into the SOC's practices, specifically what Pagett calls "the contextualization of security incidents." Expanding on that, he explained: "When you get an alert there will be questions you want to answer: have I seen anything like this before, or this exact thing? That means we can build out that threat and see what other incidents could be part of this puzzle, giving us an instant triage platform to bring that straight to the analyst."

Fortunately, the Splunk road map, especially after the recent acquisition of security automation and orchestration specialist Phantom, aligns nicely with the bank's priorities.

"We will continue to make investments into machine learning with new versions of user behaviour analytics, which does insider threat detection, [and] anomaly detection for security folks," head of product marketing at Splunk John Rooney told us.

That being said, Pagett did have some words of advice during his breakout session, telling the audience: "Don't be driven by your vendors, I know that's a strange thing to say at a vendor conference, but it's your business and you know it better than Splunk and RSA does, so invest in good people to make those decisions." **Scott Carey**