RAPID7

2015 Incident Detection & Response Survey

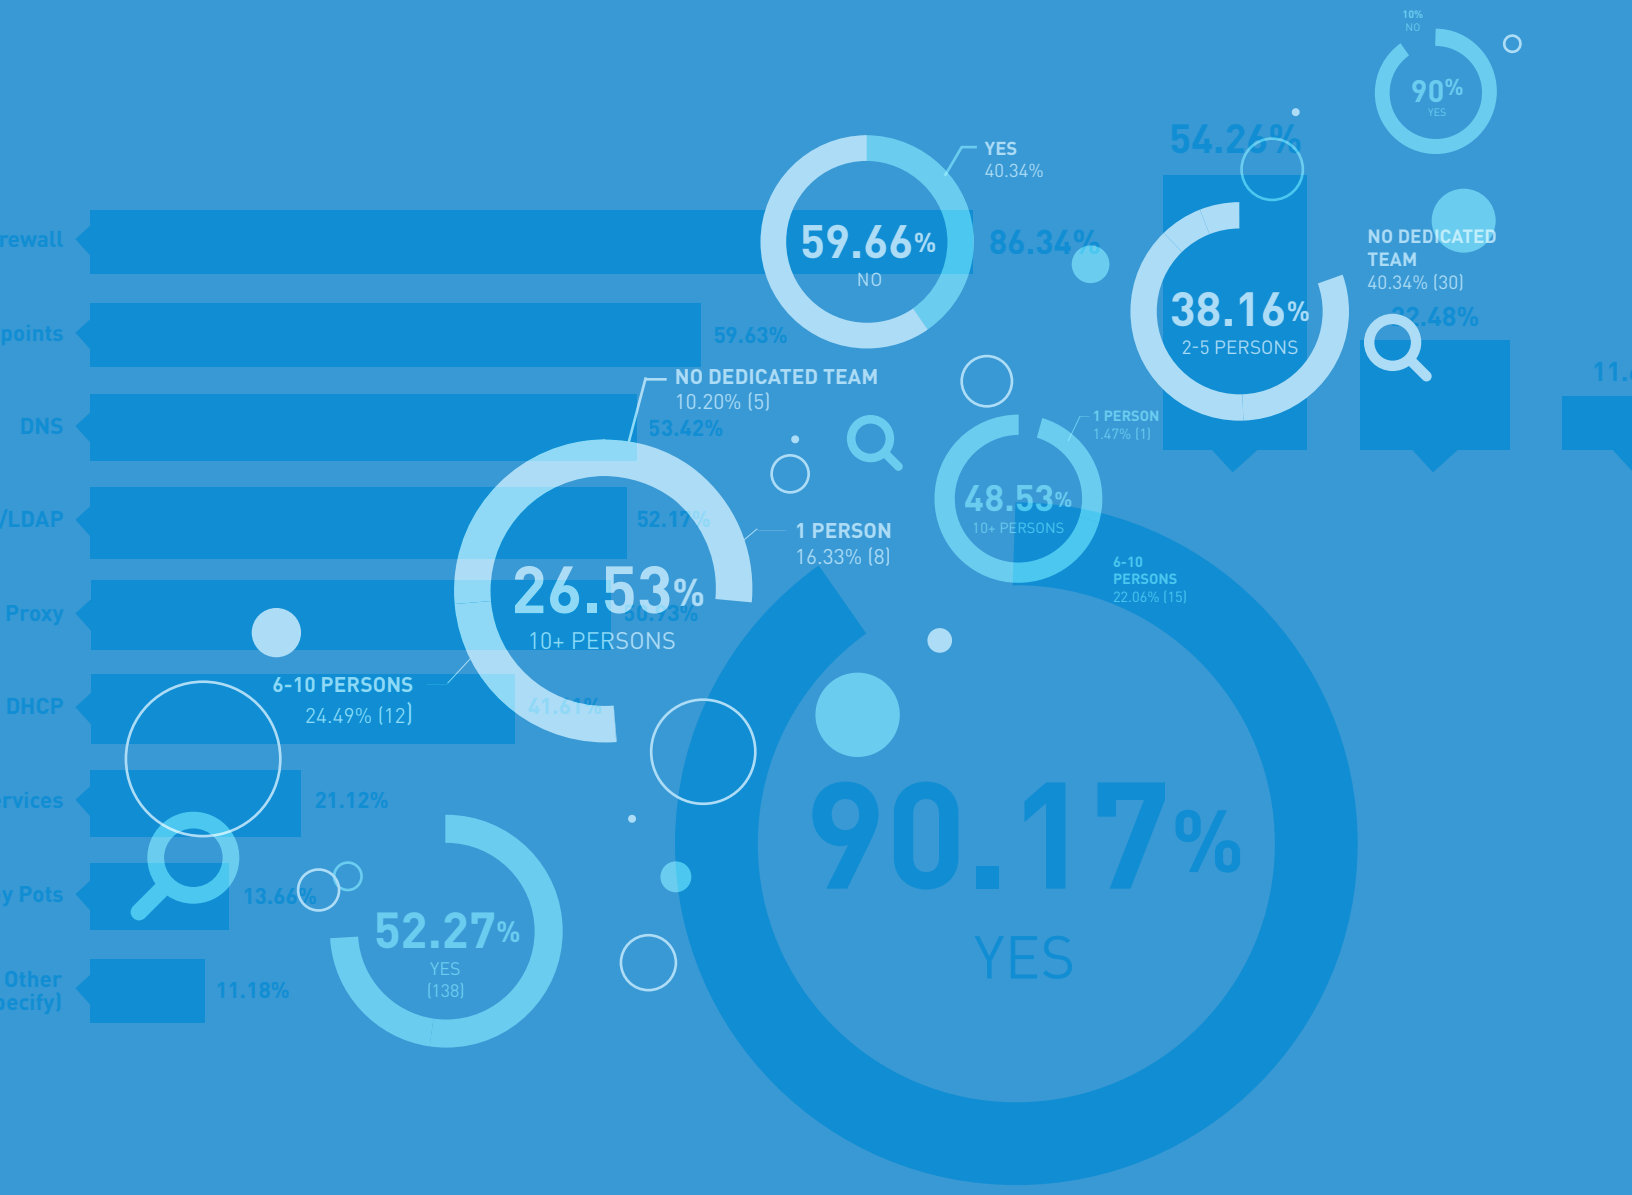

RAPID7

2015 Incident Detection & Response Survey

Contents

01 Executive Summary 3

02 Approach and Scope 4

03 Size of Security Teams, Security Stack, Challenges 5

04 2015 Top Security Initiatives 8

05 SIEMs: Who has one, and how are they using it? 9

06 Are We All Moving to the Cloud? 12

07 Top Takeaways 14

08 How Rapid7 Can Help 15

09 About Rapid7 16

RAPID7

01

EXECUTIVE SUMMARY

Today's information security teams are expected to mitigate risk in environments where employees are accessing critical and confidential data from anywhere, at any time. The network perimeter has expanded to include cloud services, mobile devices, and global forces that encompass partners and contractors, making it impossible to completely lock down the ecosystem and prevent all attacks. At the same time, preventative solutions are failing to cover the entire spectrum of attack vectors. As a result, security teams are investing in incident detection and response to detect and contain compromise as soon as it occurs.

Rapid7 conducted a survey regarding incident detection and response, in order to gain insight into today's security teams, including strategic initiatives, current tools used, and challenges. This report details and highlights the responses from 271 security professionals across the globe from organizations ranging in size and industry, including healthcare, finance, retail, and government.

The high-level results were as follows:

- Security practitioners are reporting a gap between the number of alerts generated and the amount that can be feasibly viewed, investigated, and remediated.
- 76% of respondents are not comfortable managing more than 25 alerts, yet 29% are receiving more than 75 alerts every day.

While this underscores the need for context, correlation, and prioritization, organizations are still not satisfied with their ability to detect the number one attack vector behind breaches[1], compromised credentials. 90% of respondents are worried about attacks using compromised credentials, though 60% say they can't catch these attacks today.

Security teams of all sizes face the same challenges: too many alerts; investigations take too long; and lack of visibility into user context and risk across their network. As a result, respondents said the top three initiatives for 2015 have been to (1) deploy and maintain a SIEM; (2) expand on their vulnerability management program; and (3) improve or replace their network firewalls. 52% of organizations already use a SIEM, with a further 21% looking to purchase one in the future. The flexible ability to aggregate and correlate logs enables organizations to simultaneously monitor firewall, endpoint, and DNS data. However, there are still gaps in cloud services, DHCP-to-user mapping, and honey pots.

The reality is that attack surfaces will continue to expand. 79% of respondents use at least one cloud service, with Office 365, Google Apps, and Salesforce reported as the most commonly used. The challenge is that with cloud services, attackers merely need to steal credentials to access confidential records. Currently, 59% of organizations report a lack of security visibility into their cloud services. Moving into the new year, security teams must prioritize detecting compromised credentials and the resulting lateral movement, not only on the network, but locally on endpoints and across cloud services.

As organizations continue to build out their Incident Detection and Response programs, they must settle on an approach that will improve the accuracy of detected alerts, accelerate incident investigations for more efficient investigator bandwidth, and highlight user risk from the endpoint to the cloud.

Figure 1: Are you worried about attacks using compromised credentials?

Answered: 234 (86%) Skipped: 37 (14%)

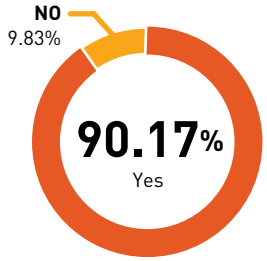

Figure 2: Can you detect attacks using compromised credentials today?

Answered: 233 (86%) Skipped: 38 (14%)

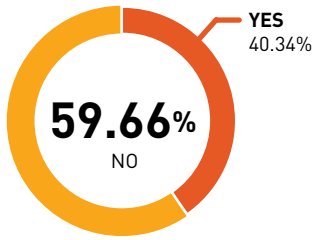

[1] Verizon, 2015 Data Breach Investigations Report, April 27, 2015

02
APPROACH AND SCOPE

The goal of the study was to answer the following questions:

1. What are the sizes of today's security teams? What about tomorrow's?
2. Is there any disconnect between incident alerts and the ability to respond?
3. What are the next major priorities for security teams for 2016 and beyond?

After launching the Incident Detection & Response survey via e-mail, Twitter, LinkedIn, and other channels, we collected 271 responses, providing a significant sample size to learn how today's security teams are operating today.

Figure 3: What size is your organization?

Answered: 271 (100%) Skipped: 0

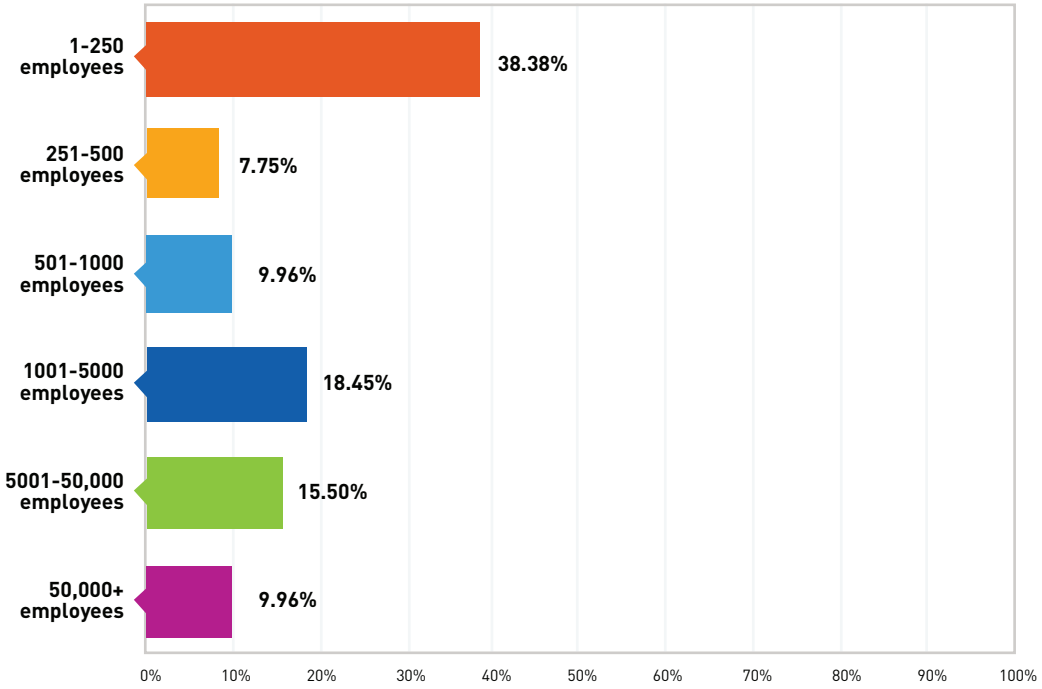

03

SIZE OF SECURITY TEAMS, SECURITY STACK, CHALLENGES

We received responses from security professionals around the world, with more than 20 countries represented. Out of the 210 that provided contact information, 45% were from the Americas, 33% Asia-Pacific (APAC), and 22% Europe, the Middle East and Africa (EMEA). Our first questions asked about the size of their organization and their security team. See below for a break-down by company size.

No big surprises here. The most interesting takeaways are that in the 1001-5000 employee bucket, 26% of organizations are working with one security person or less. At the same time, 50% have six or more dedicated security team members. This large variety in team sizes shows that (1) the security industry is still in flux, and (2) vendors need to be aware of security team resource constraints. Even as organizations look to expand teams, it's a challenge to find top security talent to bring into the people-process-technology equation.

In the 1001-5000 employee bucket, 26% of organizations have one security person or no dedicated team.

Figure 4: What size is your security team?

1-1000 Employees

Answered: 152 (56%) Skipped: 0

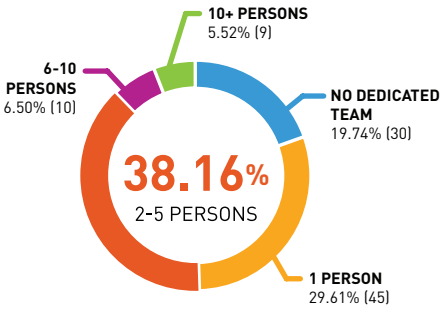

1001-5000 Employees

Answered: 49 (18%) Skipped: 1 (~1%)

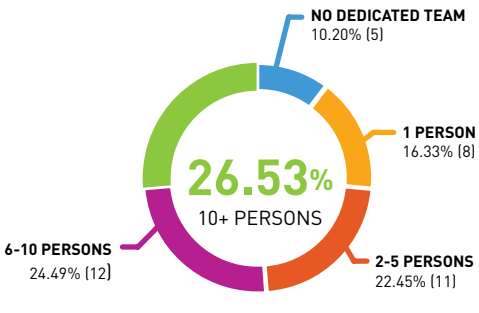

5001+ Employees

Answered: 68 (25%) Skipped: 1 (~1%)

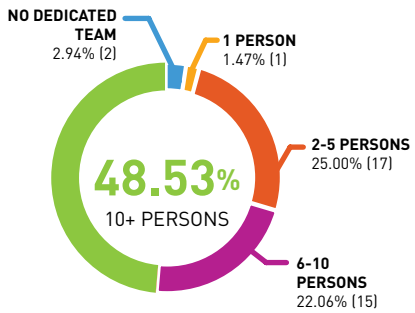

Figure 5: Tell us about your incident response spending.

Answered: 270 (99%) Skipped: 1 (1%)

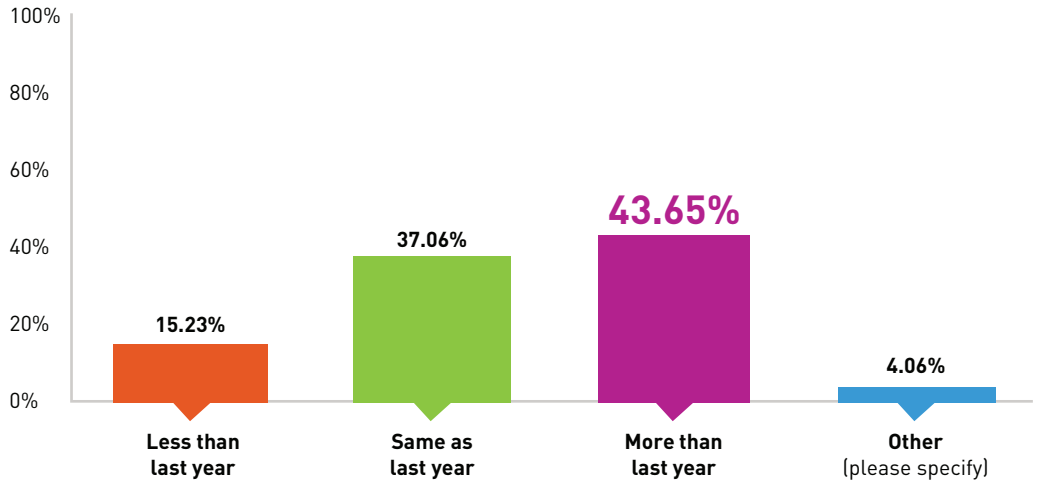

Incident Detection & Response spend is a closely related topic. Here, there was consistency across the board (figure 5).

Only 15% of respondents are cutting on budget; the free-form fields ranged from "We are putting together a plan now," to "Millions every year." On a high level, this looks to be driven by a combination of both growing security budgets and increased focus on incident detection and response. Across the industry, the decreasing effectiveness of prevention-based systems has led to the use of active, analytics-driven approaches to reduce and manage organizational risk.

Employees jump between IPs, assets, and services. We may check our email on a smartphone coming into work, then plug into ethernet upon arrival. Meetings take place on Wi-Fi across the company, coffee shops, airports, and hotel rooms. As cloud services provide anywhere, anytime access to critical information, the network perimeter now revolves around the user.

Figure 6: What security products do you use for Incident Detection & Response?

Answered: 249 (92%) Skipped: 22 (8%)

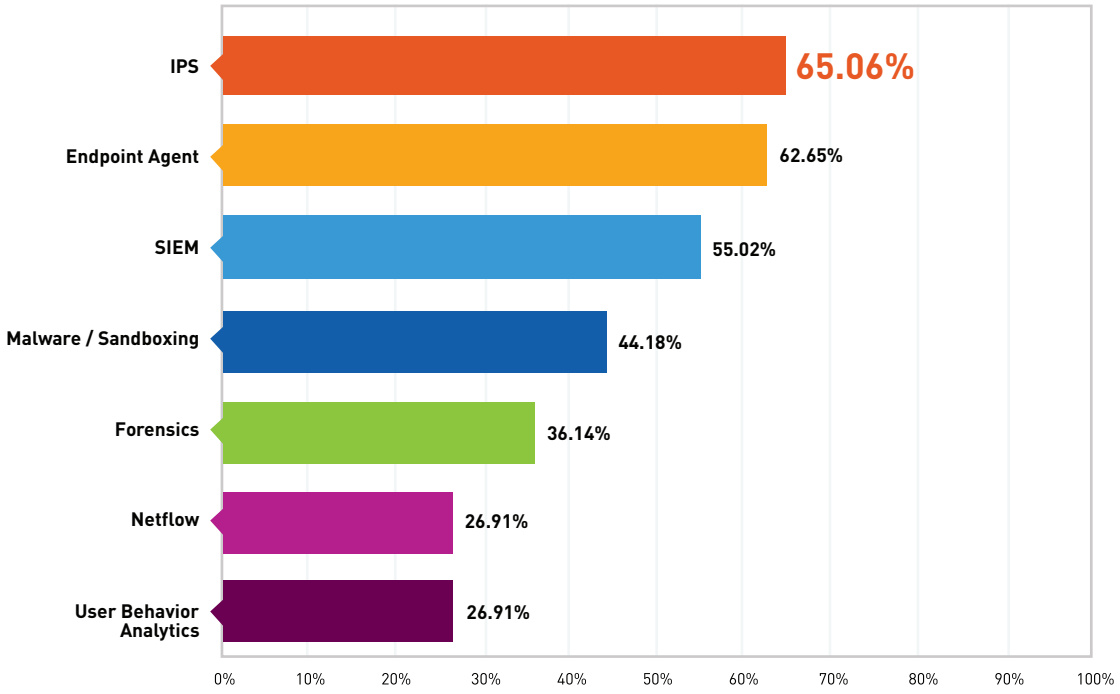

Figure 7: What are major sources of pain with your Incident Response Program?

Answered: 181 (67%) Skipped: 90 (33%)

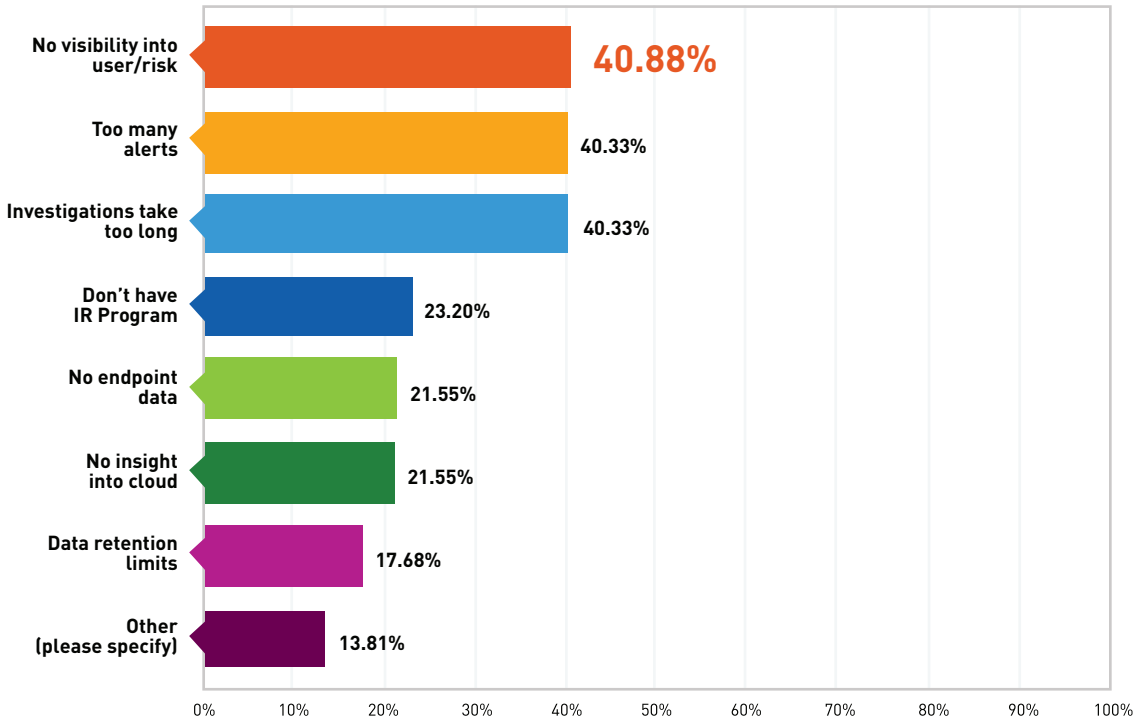

To learn more, we asked security practitioners about the security products they use, and the pain points in their Incident Response program.

IPS, Endpoint Agents, and SIEMs top of the list. Endpoint Agents range from commodity anti-virus to solutions that detect advanced malware. These three solutions are capable of generating massive amounts of alerts, but still leave gaps when it comes to detecting compromised credentials, phishing, and attacks on cloud services. Endpoint Agents are effective at detecting malware on assets, but may not be able to detect local lateral movement, log deletion, and privilege escalation exploits. All of these attacks are not only frequent, but center around malicious or negligent user activity. As a result, we expect that the use of User Behavior Analytics solutions, currently used by 27% of respondents, will continue to rise.

THE TOP THREE PAIN POINTS

- **No user context/risk visibility.** Alerts don't show "who" is affected, which can take time to figure out. No holistic view of risky behavior in the organization.
- **Too many alerts.** Security analysts must weed out false positives from multiple systems, each generating alerts. There is a distinct possibility of user-based attacks still going undetected.
- **Investigations take too long.** Teams must dig through disparate sources of raw log data and retrace users' steps across assets, IP addresses, and services. Worse, this must be done for many alerts that end up being false positives.

04

2015 TOP SECURITY INITIATIVES

We then asked respondents about their top three security initiatives for 2015. Answers spanned a wide range, as shown in the word cloud to the right.

The top three security initiatives for 2015 were:

1. Security Information and Event Management (SIEM): Deploying and maintaining their SIEM
2. Threat Exposure Management: Penetration testing, expanding programs, and web application scanning
3. Firewall: Tuning, replacing and deploying next-generation solutions

As SIEMs are expected to help detect and investigate incidents, this further highlights the investment organizations are willing to make regarding Incident Detection and Response. These large-scale initiatives require significant resources, including professional services, creation and maintenance of detection rules, and other deployment efforts. It is therefore not surprising to see SIEM management on the list of respondent feedback.

As today's primary tool for Incident Detection and Response, we also included questions for those using a SIEM to better understand their experience.

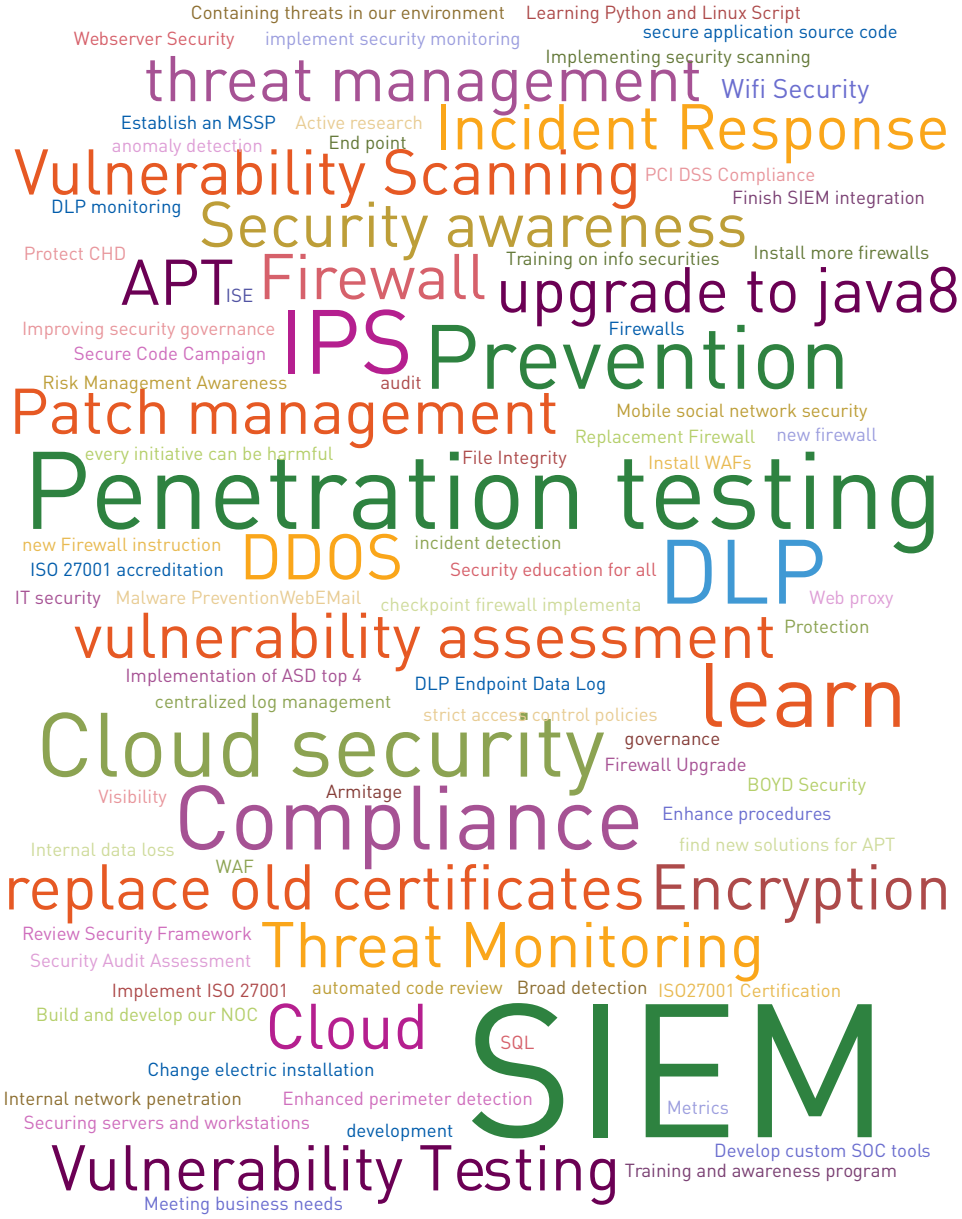

05

SIEMS: WHO HAS ONE, AND HOW ARE THEY USING IT?

52% of our respondents use a SIEM; another 21% are looking to purchase one.

Out of those that responded "Yes", the primary drivers for purchasing a SIEM were the desire to improve (1) Incident Detection, (2) Compliance, and (3) Log Search.

SIEMs certainly can provide security value by collecting and analyzing data across the entire organization. They do this by writing rules to alert on situations. If, for example:

- An administrator logs in from a new device
- There are authentications from suspicious, non-standard locations
- Domain credentials are used to authenticate to multiple critical assets over a short period

In these instances, anomalies can be detected and converted into actionable alerts. If needed, the raw log data can be quickly brought to the forefront for further security investigation.

We wanted to investigate the scope of how SIEMs are being used for detection, and if there is any disconnect between the number of alerts generated and the feasible number teams could respond to.

Figure 8: Does your organization use a SIEM?

Answered: 264 (97%) Skipped: 7 (3%)

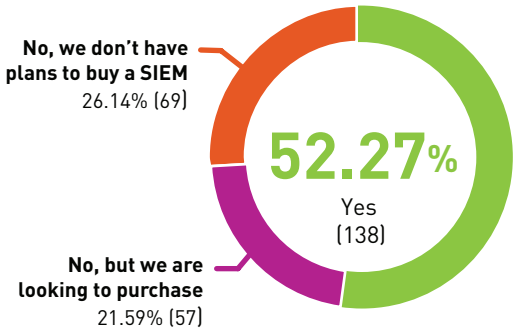

What are we monitoring?

In Figure 9, the top and bottom three data sources collected tell an interesting story. **Firewall & DNS** traffic is great at identifying anomalies throughout the network, including cloud services. Combined with threat intelligence, the two sources can identify phishing attacks and traffic from malicious URLs or IP addresses. **Endpoint** data is fantastic to send to the SIEM, such as anti-virus scan results, vulnerability scans, and Windows Management Instrumentation (WMI) logs. It's vital to detect not only malware on the endpoint, but also attacks stemming from compromised credentials, such as local lateral movement, local log deletion, and privilege escalation exploits. If performed without malware, attackers may stay under the radar and remain undetected.

The bottom three sources highlight today's coverage gaps. **DHCP Logs** identify *what asset* had *which IP* address at a particular time. If an alert only provides an IP address, it can be time consuming to retrace this back to the asset, and then to the user behind it. For this reason, many User Behavior Analytics solutions consider it a foundational source of information in order to provide the user context for any given incident alert.

Only 21% of organizations are monitoring **Cloud Services**, despite the fact

that they often provide access to client records, key financials, and other monetizable data. There are at least two reasons for this gap: (1) Cloud Services must expose a security API that provides authentication logs, and (2) SIEMs must be able to ingest and parse this data. As organizations continue to benefit from today's cloud architectures, we expect more security teams to combine cloud authentications with logins from the rest of the ecosystem.

Honey pots are a great "intruder trap" that can detect attacker reconnaissance, such as network scans, typically performed after an attacker lands on the network. A triggered honey pot is a high confidence indicator of an attack. However, these have been historically difficult to setup, and would also require an integration with SIEM. With few organizations purchasing and deploying 3rd party trap technology, it's not surprising that despite the efficacy of honey pots, they make the bottom of the list.

Figure 9: What do you currently monitor with your SIEM?

Answered: 161 (59%) Skipped: 110 (41%)

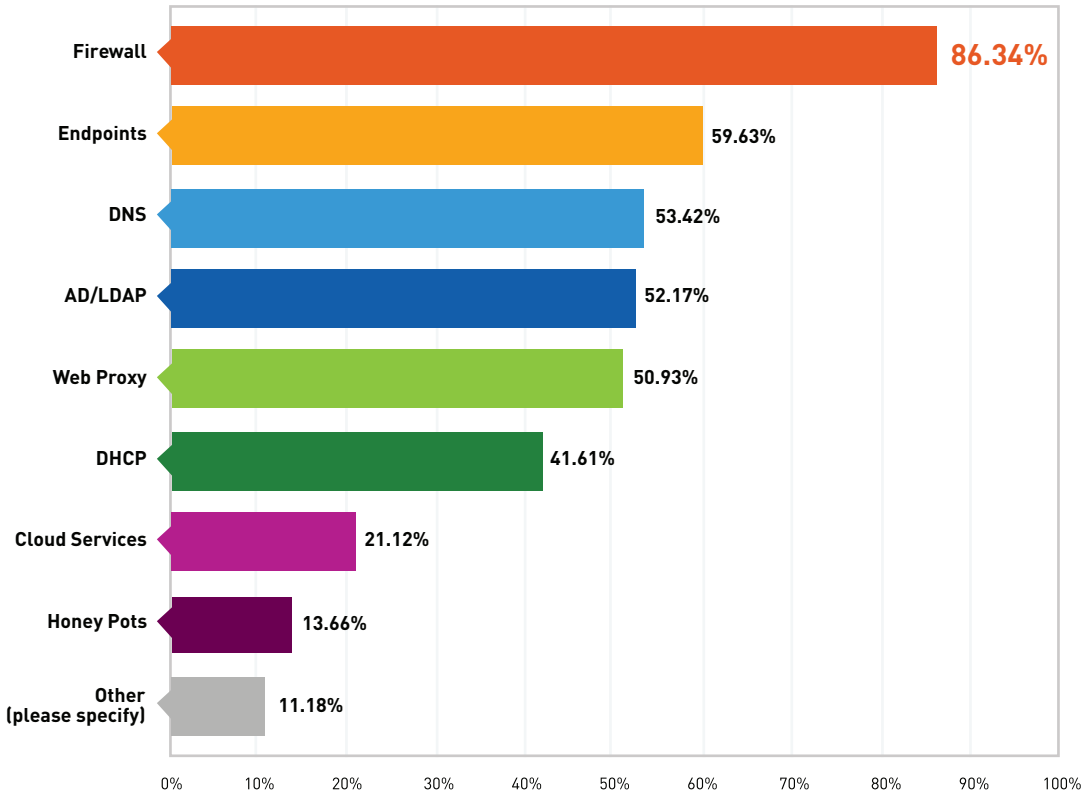

SIEM Alert Quality: Are they manageable?

Some SIEM vendors also provide investigation features to help security analysts dig deeper into incidents. However, these often require the investigator to "know what they are looking for" and it still takes time to understand the severity of the alert, the number of users affected, and remediation steps. More than three quarters (76%) of security teams are uncomfortable investigating more than 25 alerts daily, given the high time commitment associated with each one.

After comparing the answers from these two questions, we found that 62% of organizations are receiving more alerts than they can feasibly investigate. Further, these questions are only for SIEM alerts. Sophisticated security stacks mean there are many other tools that need to be maintained and referenced. Second, the constant need to tune detection makes deployments a continued work in progress. Too many alerts is just as worrisome as too few alerts, which could mean that certain attacks remain undetected.

The "Other" category included:

- "Still needs tuning. Currently, 1000+ alerts"
- "Still tuning"
- "Lots, but they are categorized so one alert isn't equal to another."

The takeaway for security vendors is that: solutions must strive to (1) prioritize, correlate, and generally reduce the number of incoming alerts, (2) simplify and automate tuning and detection, and (3) enable security teams to investigate incidents faster to bridge this gap.

Figure 10: On an average day, how many alerts can your team investigate?

Answered: 179 (66%) Skipped: 92 (34%)

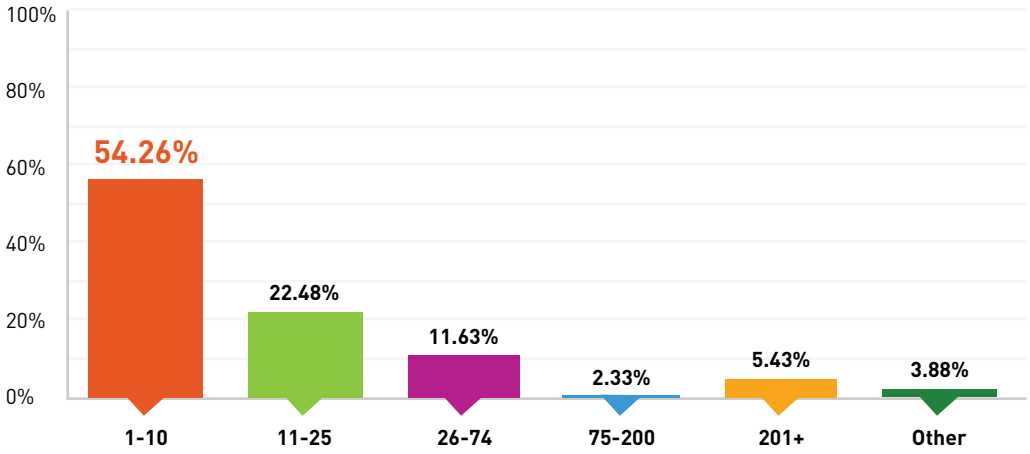

Figure 11: On an average day, how many alerts do you receive from your SIEM?

Answered: 173 (64%) Skipped: 98 (36%)

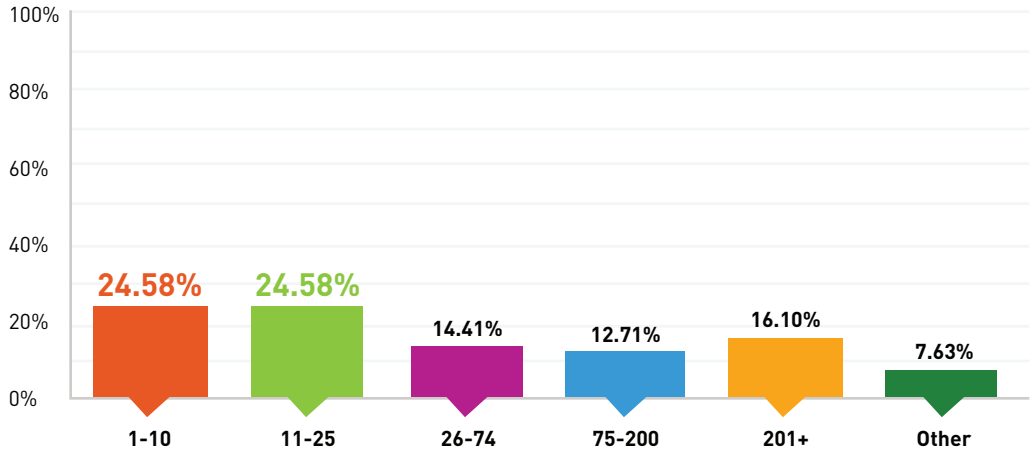

06

ARE WE ALL MOVING TO THE CLOUD?

79% of responders are using at least one cloud service, and it's no surprise that the top used cloud services are Office 365, Google Apps, and Salesforce.com. The top three answers in the "Other" category are:

1. Dropbox
2. NetSuite
3. Microsoft Azure

While cloud services have exploded in both demand and usage, teams have struggled to improve transparency and security. A huge step in the right direction is that all of the named cloud services offer a Security API with authentication data. Security companies like Rapid7 can now integrate this data with their existing offerings and extend coverage from the traditional perimeter to the user.

Figure 12: What is your company's stance on cloud services in the organization?

Answered: 235 (87%) Skipped: 36 (13%)

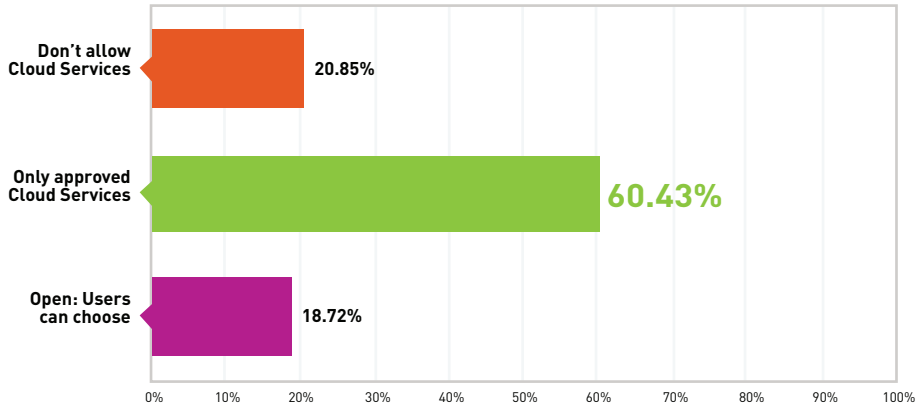

Figure 13: What cloud services do you use? Please check any of the following:

Answered: 196 (72%) Skipped: 75 (28%)

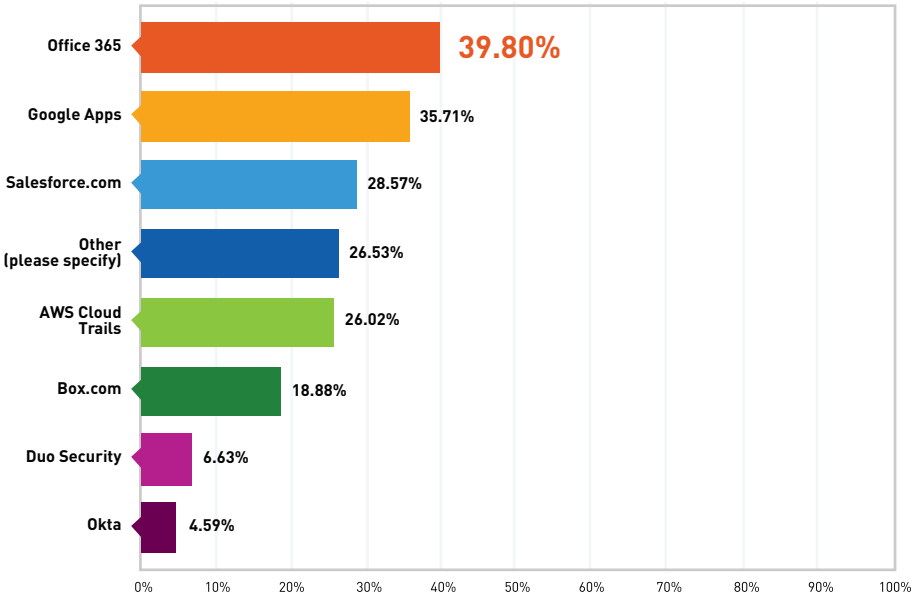

Figure 14: Do you have security visibility into cloud services running in your organization?

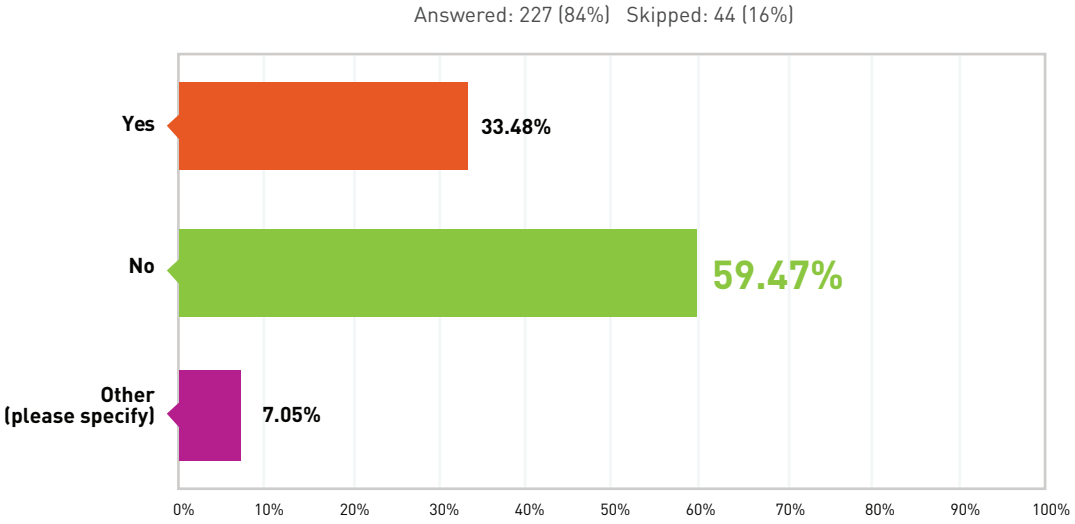

The "Other" category included responses of, "Some, partially, maybe", and "Only on a few. Software engineering teams routinely push back on security visibility measures and best practices."

Today, only 33% of organizations have security visibility into cloud services.

The challenge is not only having visibility into cloud authentication, but also incorporating this new source of data into existing workflows. We've seen that security teams already receive a myriad of alerts from multiple solutions. An ideal solution not only ingests information from Cloud APIs, but combines those authentications with the rest of the network to identify compromise. Let's say one of your employees locally logs into your US headquarters, and 20 mins later, signs into Office 365 from Japan. Would you want to know about it? Is that something you can detect today?

07
TOP TAKEAWAYS

From the insight we gain into attacker methodologies from Rapid7 Metasploit and our experienced Global Services teams, we can confirm that an emphasis on prevention is no longer a sufficient approach to security. Organizations continue to increase their reach and productivity through partners, cloud services, and mobile devices, all of which increase risk. For example, 79% of respondents use at least one cloud service, but only 33% have security visibility into those services today.

This survey confirms that organizations recognize this challenge and are incorporating incident response in their people, processes, and technology. Over the next year, only 15% of organizations will reduce their incident response spend, and 73% of security teams either already have deployed SIEM, or are planning to do so.

Gartner agrees, predicting that "by 2020, 60% of enterprise information security budgets will be allocated to rapid detection and response approaches – up from less than 10% in 2014."[2] However, further investment does not mean smooth sailing; organizations with larger security teams still report challenges in identifying user context and risk visibility, too many alerts, and lengthy investigations. Indeed, 90% of organizations are worried about compromised credentials, but only 40% can detect these types of attacks today.

The top three attack vectors behind breaches continue to be compromised credentials, malware, and phishing.[3] Both security vendors and practitioners must ensure that attacks leveraging these methods can be detected immediately across the entire network ecosystem. Further, this must be done while taking into account the realities of the security world: limited time and resources, very low tolerance for false-positive alerts, and the desire to receive alerts in a centrally managed system that covers all IT assets, from the endpoint to the cloud.

[2] Gartner, Designing an Adaptive Security Architecture for Protection From Advanced Attacks, February 12, 2014

[3] Verizon, 2015 Data Breach Investigations Report, April 27, 2015

08

HOW RAPID7 CAN HELP

Rapid7 can help fill your security needs whether they are in people, processes, or technology.

UserInsight: Our User Behavior Analytics solution shows you the intruder clearly without the guesswork. It analyzes user behavior, detects account takeovers, and provides visibility of user context and risk, giving you the confidence to know where intruders are and what you need to do to repel or recover from an attack. UserInsight detects common attack vectors such as compromised credentials and lateral movement that malware detection can't see. Unlike SIEMs, UserInsight is easy to use, provides few false positives, can be deployed in days, and is cost-effective to operate. By analyzing user behavior, UserInsight also shows you which users are putting your company at risk through bad security practice. Applying Rapid7's unique combination of Red Team and Blue Team expertise has made UserInsight the most effective solution for detecting intruders.

Analytic Response: This fully managed service is available through our 24/7/365 SOC. Analytic Response combines the UserInsight technology with real-time threat intelligence, and world-class security analytics with deep experience hunting for and validating threats. Once a possible threat has been identified, Analytic Response validates and prioritizes it based on your unique environment. Customers are only notified if the threat is indeed real.

The result: quickly identify nearly 100% of attacker activity without wasting time investigating a mountain of false alerts. Remediation is also available with immediate expert guidance, from the same team that already understands the architecture and intricacies of your network.

Global Services: In addition to the security analysts available for Incident Response, the Global Services team is available to provide Program Assessment and Development Services, Penetration Testing Services, Security Awareness Training, and more. Rapid7's experience in building and managing security programs enables us to work with organizational leaders to drive executive alignment.

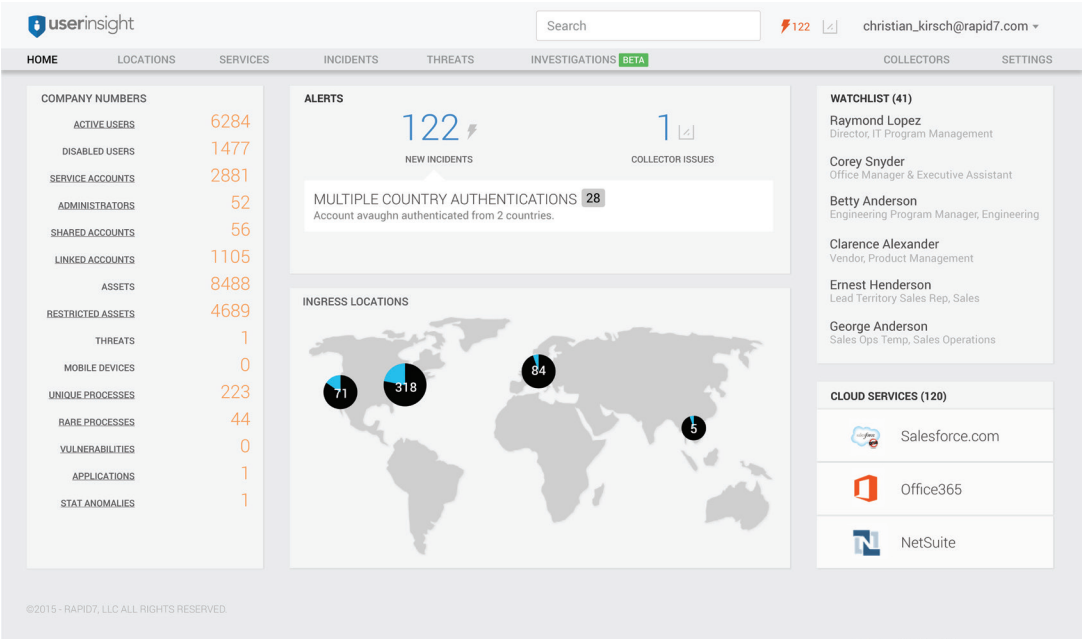

09
ABOUT RAPID7

Rapid7 is a leading provider of security data and analytics solutions that enable organizations to implement an active, analytics-driven approach to cyber security. We combine our extensive experience in security data and analytics and deep insight into attacker behaviors and techniques to make sense of the wealth of data available to organizations about their IT environments and users. Our solutions empower organizations to prevent attacks by providing visibility into vulnerabilities and to rapidly detect compromises, respond to breaches, and correct the underlying causes of attacks. Rapid7 is trusted by more than 4,400 organizations across 90 countries, including 35% of the Fortune 1000. To learn more about Rapid7 or get involved in our threat research, visit www.rapid7.com.