

Survey

SANS 2023 SOC Survey

Written by [Chris Crowley](#), [Barbara Filkins](#), and [John Pescatore](#)

June 2023

Executive Summary


Welcome to the 2023 SANS Institute SOC Survey. In this, our seventh annual survey, we added many questions but didn't really take any away. Our new areas of focus include operational threat hunting, threat intelligence, data ingestion into the SIEM, and SOAR, as well as more detailed questions relevant to staff hiring and retention. Thank you again to the respondents who generously spent their time, a mean of 59 minutes (Q5, $n = 641$) based on the Qualtrics reported duration, to answer our barrage of questions.

The lead author (Crowley) has heard many times that just taking the survey is good for SOC staff and managers, because it is challenging and thought-provoking. If you're reading this and it provides value, please be sure to take the survey in 2024! We're already planning enhancements and updates. We're also hoping to hear from you what you'd like to read about in the future. If you have analysis that you'd like to perform, the *deidentified* raw data set and a Jupyter notebook (Python) is available for download and analysis at <https://soc-survey.com>. Among this year's top findings:

- More than 75% of respondents detected incidents before external notification, 9% via proactive threat hunting (Q3.31, $n = 327$).
- 84% of SOCs collect and expose metrics, including these top three:
 - Quantity of incidents
 - Time from detect to eradicate
 - Ratio of incidents from known/unknown vulnerabilities
- Continual tuning of SOAR by skilled analysts is needed to obtain value—SOAR as a work style, not throwing a switch.
- SOAR work style increases effectiveness more than it reduces staffing needs.
- Monthly review of SIEM data ingestion proves valuable to vet sources.
- SOC funding follows a traditional IT model: SOC budget requests go up, allocations come back down.
- SOC outsourcing tends to be pen-testing (and variants) and forensics, whereas in-house tends to be security system architecture, engineering, planning, and administration.

Figure 1 on the next page provides a snapshot of the demographics for the respondents to the 2023 survey.

Top 4 Industries Represented




Each gear represents 10 respondents.

Organizational Size




Each building represents 25 respondents.

Operations and Headquarters



Top 4 Roles Represented



Each person represents 10 respondents.

Figure 1. SOC Survey Respondent Demographics


What We Learned in 2023

We included several questions this year that weren't present in previous surveys.

We asked about visibility into data and ingestion choices made for data into SIEM. Everything (Q3.5, 171/600, 28.5%) and some selectiveness based on risk (Q3.5, 169/600, 28.2%) were top explanations, with a monthly review (Q3.7, 105/239, 43.9%) being the most common frequency for those who said they reviewed ingestion (Q3.6, 256/597, 42.9%) on a periodic basis.

Aligned to VERIS structure of detection sources, we asked respondents to identify the ranking of incident discovery. A little over two-thirds of respondents (Q3.31, 246/327, 68.3%) indicated that monitoring/alerting was most frequently responsible for detection. See Figure 2 on the next page.

Rank the following techniques for incident discovery from most to least frequent.



We asked about SOAR use, because it is now a technology fixture in SOCs. (Next year, we'll introduce a similar line of questioning for use of AI/ML.) It seems that people are taking SOAR as an ongoing change and adjustment project (SOAR as a work style), and they most commonly (Q3.33, 16/46, 34.8%) allow the users/analysts of the SOAR to tune as they go. Important to note, the people who got to answer the question on how they tune the SOAR were only those respondents who indicated (Q3.32, 48/398, 12.1%) that SOAR was their primary method for event data correlation and analysis. See Figure 3.

We included a number of questions on what SOC managers focus on when hiring; look at the SOC staff section for more details.

Figure 2. Ranked Incident Discovery (Q3.31, n = 327)




Figure 3. SOAR Updating and Tuning (Q3.33, n = 46)

Key Findings

“Do more with less” is a hallmark clarion call, trite and honest. There are only limited resources in the organization, and SOC managers who can show connections from increased investment in the SOC to improvements in business-relevant metrics are in the best position to benefit from that increased spending on cybersecurity. Nonetheless, in the past 10 years, cybersecurity budgets have increased substantially.

There is still no universal equation for budgeting for an adequate SOC. Deriving estimates from overall IT spend is a common practice in trade literature but carry the caveat that it's no way to set a budget.¹ Figure 4 shows how respondents allocate budgets.

The most common (Q3.69, 126/300, 42%) answer was that SOC management prepares budget input, and then higher-level decision making allocates funding. That seems rosy compared to those (39/300, 13%) who stated that budget decision makers pay little attention to the SOC management's recommendations.

The budget shows minimal correlation to organization sector and size. We'll explore this more in a later section. Most people want to see the money, so see Figure 5 for the responses of what reported annual budgets are. Important to note, the most common answer was: Unknown! (Q3.68, 68/307, 22.1%).

Metrics are regularly used in SOCs—only a small portion (Q3.47, 39/349, 11.2%) said “No” they don't provide metrics. Of those 11%, we wonder how this is possible, and come to the depressing conclusion that the audience who would be receiving the metrics likely doesn't care. Interesting to note, when industry vertical is cross referenced, government (Q2.2, $n = 69/641$) is the top industry that said “No” metrics (Q2.2 x Q3.47, $n = 9/41$).

The respondents who do use SOC metrics were generally satisfied with their effectiveness. Only 23% (Q3.48, 28/124, 22.6%) expressed being “not satisfied” with current metrics. A later section will discuss metrics in more detail.

To understand how SOCs can use metrics to move to better performance, we asked if they have a method for calculating the value the SOC provides. This is a tricky calculation, because it expresses the value of *something not occurring*. It's no surprise that people are trying, but there isn't clarity or consistency in doing this, partially because it isn't easy. As a result, more than half (Q3.55, 184/327, 56.3%) of the responses indicated people aren't trying to calculate this.




Figure 4. How Funding Is Allocated (Q3.69, $n = 300$)

What is your estimated annual budget for new hardware, software licensing and support, human capital, and any additional costs?




Figure 5. Estimated Annual Budget (Q3.68, $n = 307$)

¹ <https://www.gartner.com/newsroom/id/3539117>

For those (Q3.55, 83/327, 25.4%) who said they are, we asked what the result was. Most (Q3.56, 64/76 & 55/77) of the responses indicate there was a 50% or less reduction in handling and incident impact cost. Reducing time to detect/resolve/restore directly is the second-most-common metric in use by SOCs, and that correlates directly to reducing the overall cost of an incident and demonstrating SOC value and business-relevant progress. See Figure 6.

To facilitate the estimate of reduction we're asking about in Q3.56, there's typically a value assigned to assets: when those assets are challenged by an actual attacker, the SOC gets to claim a reduction due to its preparation (reduced handling costs) and ability to intervene (reduced incident costs) to minimize damage.

So, what's the basis of the value claimed? We asked respondents if they have a cost per record. Most (Q3.53, 160/323, 49.5%) said no, but there's a high (Q3.53, 63/323, 19.5%) percentage who don't know if they have a cost per record or not. See Figure 7. We'll delve deeper into these record types and costs later in this paper.

It takes qualified people to run a SOC. This has been a consistently reported aspect for the past six years of the survey. Again this year, we asked many survey questions related to staff and appropriate qualifications. But the most common question encountered in the authors' experience related to SOC staff is "How many are required?" This is typically in the form of something like, "If [company] in [industry] has [number of employees], then how many people are needed to staff the SOC?" The cynical author has started simply answering, "around 25," because in the survey data, the most common (Q3.58, 83/335, 24.8%) SOC size is between 11 and 25 staff. See Figure 8.

For the estimated incident with SOC vs. incident without SOC, select the best match for estimated relative handling and loss costs on a per-incident basis.




Figure 6. Estimated Handling and Incident Cost Reduction (Q3.56, n = 76,77)

Have you calculated a "cost per record" from an actual incident?




Figure 7. Cost per Record Based on Incident Data (Q3.53, n = 323)

What is the total internal staffing level (i.e., all related positions) for your SOC, expressed in terms of full-time equivalents (FTEs)? What is the number of FTEs specifically assigned to the management of your SOC systems, not just to analysis of the data from your SOC systems?




Figure 8. Total Internal Staff Levels (Q3.58, n = 335)

If we only look at large (50,000 FTE or greater) organizations, the most commonly reported size is 26–100 SOC staff (Q3.58 v Org Size, $n = 8/26$) followed closely by 11–15 (Q3.58 v Org Size, $n = 7/26$) SOC staff. See Figure 9.




Figure 9. Number of SOC Staff by Organization Size

Key Challenges

Now that we've identified some key elements within the results, let's address the key challenge identified by respondents in 2023. "Lack of context related to what we are seeing" was the most popular (Q3.79, 50/313, 16%) response. See Figure 10. It's worth noting that "Lack of context" was almost at the bottom of the list in last year's survey. (See Figure 4 in the 2022 Survey if you want to compare.²)

There are other challenges at roughly the same numbers. Lack of automation and orchestration (Q3.79, 40/313), lack of enterprise-wide visibility (Q3.79, 43/313), and lack of skilled staff (Q3.79, 44/313) are also top challenges.

Skilled security staff is needed to solve that "lack of context" problem, as well as to make the "SOAR workstyle" effective, but experienced security analysts are still hard to find.

Lack of management support has decreased as an obstacle, indicating funding is available to make progress on increasing context/visibility and automation.

Key challenges are illustrated in a word cloud in Figure 11.

What is the greatest challenge (barrier) with regard to full utilization of your SOC capabilities by the entire organization? Select the best option.




Figure 10. Key Challenges (Q3.79, $n = 313$)




Figure 11. Key Challenges Word Cloud (Q3.80, $n = 138$)

² www.sans.org/white-papers/sans-2022-soc-survey

Expanded Content

Things get curioser and curioser as we delve deeper. The full set of responses and an accompanying Jupyter notebook in Python to assist with performing analysis is available from <https://soc-survey.com> if you choose to do some of your own analysis.

Design/Development/ Implementation

“Do you actually run a (cyber)security operations center (SOC)?” is a reasonable question to start from as we look into design, development, and implementation topics. Or put another way, “How does the survey define a SOC?” The way this survey characterizes a SOC is broad. It is a cybersecurity operations center (SOC) if an ongoing mission of an operational team to centralize cybersecurity activity is authorized and funded.

A more detailed approach to assessing “SOC or NOT?” is reviewing capabilities the team and outsourced partners perform. This reveals a continuum of basic SOC to learning SOC with almost all capabilities. As has been consistent across years in the SOC Survey, the respondents largely agree that the capabilities we inquire about are done. Figure 12 shows a list of capabilities sorted on if they’re done, regardless of whether they’re done internally, outsourced, or both.




Figure 12. Capabilities Performed Sorted by Total (Q3.13, n = 545)

The survey doesn't explore why people outsource, although we'll speculate on this in later sections on staffing and funding. But what we do know is that SOCs outsource consistently. Taking the same data from Figure 12, we see that the more commonly outsourced activities are variations on forensics, threat intel, and penetration testing, which are less commonly done activities for SOCs (see Figure 13), as the lower ranks of activity performed are at the highest ranks of outsourcing.

Activities that are more likely to be outsourced are also slightly less likely to be done *overall*. This could be due to lack of budget, meaning those specialized items are simply not done at all. Or it could reflect the sentiment that forensics and pen testing are not a requirement and, hence, not done.

The authors see the categories more likely to be outsourced—forensics, pen-testing, and threat intelligence—as specializations that require substantial training and experience but aren't used consistently within most SOCs. Of course, looking at alerts, triaging them, and performing handling requires knowledge, skills, and abilities (KSAs). However, these categories are broader in scope of KSAs, so it is more difficult to identify an outsource partner where a transactional basis assures a value proposition over leveraging more flexible internal staff.

Hunting and threat intel are important capabilities of the SOC, in the opinion of the authors. The SANS Institute has analyst papers on these areas, but we included similar questions here to see what the SOC respondents had to say about their operational performance of these activities.




Figure 13. Capabilities Performed, Ranked and Sorted by Outsource Rank (Q3.13, n = 545)

First, we consider threat hunting the investigation of available data, presuming that other alerting-based mechanisms have failed. We do not consider looking in logs for something specifically known to be malicious (such as a known malicious domain name) as hunting. Others might include a malicious domain search within the grouping of threat hunting.³ Although valuable, we define this as historical/retroactive analysis. Threat hunting entails a distinctive attribute that we don't have a specific value to match. We're looking for outliers and new indicators by identifying objectionable behavior or patterns. These refined semantic distinctions ultimately have limited value because we need to perform both of these approaches to identify problems. The distinction is drawn to implore SOCs to do the more difficult form of hunting (looking for the unknown) in addition to historical analysis.

Of course, we wondered how our respondents' SOCs performed threat hunting, so we asked them. Figure 14 depicts partial automation as the most common answer across all of the categories we inquired about: historical, analyst-driven, and technology-driven hunting. Manual isn't very far behind in all categories.

Where responses don't make sense in the survey, we think it is appropriate to address them with some potential explanation. In this case, the 48 responses (Q3.15, 48/491) indicating that analyst-driven threat hunting is fully automated is baffling to the authors. Bots; human respondents without a clue, comprehension, care, or consideration; considering "fully automated" as a query or script created by a person as analyst-driven threat hunting; and a linguistic misunderstanding of the question options are the hypotheses we've articulated but not assessed for this oddity. See Figure 14.

Threat intelligence is the study of adversaries with the intention of optimizing the use of scarce resources. We use it to improve our defensive posture, identification capability, and post-detection response readiness and capability.

There's an abundance of threat intel data out there. So, we asked how people use it in an actionable way in a SOC. We separated feed consumption, production of threat intelligence, and attribution threat intel categories for the question (Q3.18, n = 309). Figure 15 shows that internal only (in house: 127, 132, 113) activities outnumber either purely outsourced or mixed internal/outsourced when taken distinctly. But, adding outsourced and both results in a higher aggregate number in all categories (out+both: 155, 141, 153). So, we assess that more people outsource threat intel entirely or partially (out+both) than do it entirely on their own (in). From the same numbers, it's clear that more people do it internally in some respect (in+both: 208, 201, 187) than do it externally in some respect (out+both). See Figure 15.




Figure 14. Threat Hunting Categories (Q3.15, n = 491)




Figure 15. Threat Intel Activities (Q3.18, n = 309)

³ www.sans.org/security-resources/glossary-of-terms

SOC Architecture

We consider something a SOC based on mission and capabilities, not architecture. But the architecture of SOCs is still worth exploration. In our survey questions, we consider the physical locations, staffing arrangements, and what is protected to be part of the architecture.

Centralized, all in one physical location SOCs might still allow work from home, for example. But the physical work location where the staff “sit” is still one geographic region. The other aspect of this centralized and distributed notion is where the data used by analysts to view alerts resides. The centralization of all data into a SIEM from cloud resources doesn’t always make sense from a value proposition.

So, where the people are and where the data is are not necessarily the same. Related, some jurisdictions and industry verticals prefer (or are legally obligated) to keep data within the country or within organizationally owned systems. This makes architecting the SOC systems complicated. What’s more, SOC staff may have strong opinions on working together as a team—meaning being together in one place. If scarcely available staff insist on a specific arrangement, it is likely to manifest in the SOC architecture.

In Figure 16 we see the continuation (in this survey’s history) of the single, central SOC dominating (Q3.8, 271/557, 48.7%) the responses.

In Figure 17 we see the same signaling we’ve seen for the past three years. “Cloud-based services” is projected to be the architecture next year. (Q3.9, 130/527). But, based on the percentage of “current” in 2021 (12.9%), 2022 (15.2%), and now in 2023 (19.8%), we’re seeing only a modest change represented in the responses in the survey. We don’t track individual responses from year to year so we don’t know if people are saying they will change but not doing it, or if the respondent composition year after year has the same forward-looking thought but doesn’t achieve the change.




Figure 16. Structure of SOC (Q3.8, n = 557)




Figure 17. SOC Infrastructure (Q3.9, n = 546, 527)

Another architectural attribute is whether the SOC operates 24 hours a day, every day of the year. Overwhelmingly, the answer is yes, with only 18% (Q3.23, 80/434, 18.4%) indicating they do not run 24 hours a day, as shown in Figure 18. The architectural decision of running non-stop drives quite a bit of outsourcing, with 49% of the overall answers (Q3.23, 213/434, 49.0%) and 61% of the yes answers (Q3.23, 213/349, 61%) indicating outsourcing was used in whole or in part to accomplish non-stop operations.

We'll describe composition of staff and staff roles in a moment. First, keeping with the architectural focus, we consider remote work for SOC staff as an architectural attribute. When we dive into the staff section, we will describe what factors enable individual employees to work remotely. Almost three-quarters of respondents (73%) say staff are allowed to work remotely (Q3.24, 318/435, 73.1%). See Figure 19.

Necessarily, some of the respondents who said they work in a centralized SOC also responded that the SOC allows remote work. So, we delved into this set. Of the respondents (Q3.8, n = 271) who say they have a single central SOC, 58% (Q3.24, 157/271, 57.9%) indicate that remote work is allowed. See Figure 20.

The structure (Q3.8) of the SOC doesn't appear to have a substantial influence on whether SOC staff can work from home. See Figure 20, which depicts SOC structure and whether staff are allowed to work from home.




Figure 18. SOC 24/7 Operations (Q3.23, n = 434)




Figure 19. SOC Staff Remote Work (Q3.24, n = 435)

SOC Staff

How many staff are there currently in the SOC, and is that the right number? This is an important question with multiple attributes to explore.

Each SOC could probably do more or operate at higher quality with more *qualified* people. The SOC is a space where adding people risks detracting from performance despite added expense. Most SOCs struggle to effectively incorporate new or junior staff. They can only tolerate onboarding a small volume of staff who need substantial on-the-job training to develop the required knowledge, skills, and abilities. Why? It is the opinion of the authors that:

1. SOCs aren't designed, built, or operated to address the human capital cycles that actually occur; and
2. SOCs are chronically understaffed to the degree that tasking those busy people to help address the shortcoming is essentially loading on one more new skillset (training others) for the SOC staff to try to master.




Figure 20. Remote Work by SOC Environment Structure (Q3.8 vs Q3.24, n = 435)

In the key findings section, we described the typical size of a SOC. Let's look in more detail at the job roles of those staff. We asked about specific roles, such as monitoring analysts, systems administrators, and incident handlers. Then, we asked how many of these were on staff. Figure 21 doesn't account for the overall size of the team, just their reported numbers for each role. Some of these values are surprising—seeing someone report there are more than 1,000 threat intel analysts, for example.




Figure 21. Staff Role Count (Q3.59, n = 337)

Of course, SOCs want to understand what the *right* number of analysts is. What appears to be happening is that people are calculating workload based on actual time worked on tickets. Responses indicate that existing workload calculation is used for staff count justification (see Figure 22). More maturity in SOAR and overcoming the cited obstacles in visibility and context are needed to see reduction in average analyst time/incident and reduction of headcount need.

We hear a lot about staff hiring issues and staff turnover. So, we asked questions around what the average duration of employment is. As shown in Figure 23, three years or fewer is most common. Further, fewer than or equal to five years (Q3.60, 227/311, 73.0%) is the reported average tenure for about three-quarters of respondents. This is in line with overall IT turnover and should be expected.⁴

Select the best description of how you calculate per-analyst workload.




Figure 22. Per-Analyst Workload (Q3.52, n = 311)

What is the average employment duration for an employee in your SOC environment (how quickly does staff turn over)?




Figure 23. Average Employment Duration (Q3.60, n = 311)

⁴ www.inc.com/business-insider/tech-companies-employee-turnover-average-tenure-silicon-valley.html

We delved further into hiring this year than in previous surveys. We asked what hiring managers are looking for. We asked about this in several attributes including technical skills (see Figure 24) and non-technical skills (see Figure 25 on the next page). The top answers were “Information Systems and Network Security” and “Risk Management,” respectively. These responses can be utilized by hiring managers and prospective candidates alike.




Figure 24. Technical Skill Focus for Hiring (Q3.64, n = 308)

The takeaway is that SOC managers are looking for “T-shaped” analysts with deep skills in one or more technical areas augmented with broad communications, risk management, and business knowledge. Because such analysts are in high demand, job satisfaction is key to reduce turnover.

When the workload is too high (and probably also when it is too low) analysts decide they don’t want to be at the organization. There are plenty of reasons a person leaves a job. So, we asked the question about how to retain staff. This might reduce the hiring to only those new staff positions you’re able to secure. The most commonly cited retention method? “Career progression” (Q3.66, 93/314, 29.6%) topped the list. It seems people employed within the SOC are seeing this as a journey and are looking for personal development and meaning.

Meaningful work (Q3.66, 64/314, 20.4%) was the second most cited response, barely eclipsing money. Money (Q3.66, 63/314, 20.1%) works, too. See Figure 26.

After hiring and retaining staff, there’s a lot to explore about how they work. A current hot topic is the concept of working from home. It has long been these authors’ opinion that working from home is viable for all SOC roles given appropriate data protections.

What is the single most important non-technical skill deficit you are trying to address when you hire staff for organizational roles?




Figure 25. Non-Technical Skill Focus for Hiring (Q3.62, n = 306)

What is the most effective method you have found to retain SOC employees?




Figure 26. Employee Retention (Q3.66, n = 314)

To illustrate the idea of remote work, monitoring a DOD Classified network could be done using remote workers, if the transport networks for the protection were themselves protected at the appropriate level. Some people balk at this, but would simultaneously accept that there are networks that provide video conferencing, with people remotely discussing the content stored on the computers under monitoring by the SOC. Of course, the cost of securing the necessary transport and storage network of a SOC like this becomes onerous at some point. So, organizations decide on boundaries where SOC staff must be within a physical location that can be secured at the appropriate protection level. Hence, most SOCs monitoring classified networks require on-premises staff.

Most readers of this document don't need to meet the strict physical, technology, and administrative requirements of a USDOD Secret (or higher) network. So, they're left to decide if the data can go to a SOC analyst's home computer or not—usually without any rigorous standard of quantitative assessment. We suggest using the aforementioned "value of a record calculation" as a start to this effort, but regrettably can't provide a simple equation to do this risk vs. expense calculation.

See Figure 27 describing the factors involved in allowing a SOC analyst to work remotely. Tied for first (Q3.25, 125/289, 43.3%) were the role of the SOC staff, and if secure access to data is feasible. We presume that some SOCs deal with data that is considered "on premises only," and analysts supporting those SOC-monitored systems aren't allowed to work remotely.

There might also be a rationale for citing that the data doesn't need to be on premises per se, but remote access technology is not adequate for the security sensitivity. Most SOCs would err on the side of caution within these parameters.

Caution for appropriate work-life balance and avoiding expectations of constant availability is the counterargument to working from home. A certain way to drive employees to a breaking point is to enable work from home, then foster an environment that drives expectations that the person is always available.

Making sure the work-life balance is appropriate for the long term and adding tooling to relieve analysts of needless and frustrating tedium are likely to give them a sense of career progression, wherever they work.




Figure 27. Remote Work Factors (Q3.25, n = 289)

Technology

Our technology question section is very long and optional. Those who made it this far in the survey (roughly half) opted to skip this part, with roughly 50% saying yes (Q3.37, 188/380, 49.5%).

Those who persisted were subject to an extensive set of questions on how much they like technology and how completely the technology is deployed.

We first show the technology on a GPA-ranked basis. Figure 28 has the GPA-ranked tech based on our respondents. Top of the list is host-based EXDR (Q3.39, GPA: 2.89). Bottom is a tie between network-based packet analysis and artificial intelligence/machine learning (Q3.39, GPA: 2.18). Note that no GPA was above a C ... unless we grade on a curve!

Next, we cross walk the phases of deployment, and respondents' satisfaction with them. In the past two surveys, we've presented a similar picture, and the correlation seems to hold that technology that reaches production has higher satisfaction. Although we do not claim this as a causal relationship, it speaks to the reality that the tech in use has a higher score when it is fully deployed. See Figure 29 on the next page.




Figure 28. GPA Ranked Technology (Q3.39, n = 194)

Technology	A	B	C	D	F	GPA	Total
Host: Vulnerability remediation	38	61	41	11	6	2.7	157
Net: Network intrusion detection system (IDS)/intrusion prevention system (IPS)	40	54	42	15	4	2.7	155
Net: Next-generation firewall (NGF)	50	54	34	9	7	2.9	154
Host: Malware protection system (MPS)	44	56	40	10	5	2.8	155
Net: VPN (access protection and control)	53	51	34	14	5	2.8	157
Net: Email security (SWG and SEG)	54	49	36	10	6	2.9	155
Analysis: SIEM (security information and event manager)	47	49	41	15	3	2.8	155
Host: Endpoint or extended detection and response (EDR/XDR)	57	50	31	16	4	2.9	158
Log: Endpoint OS monitoring and logging	45	62	37	12	2	2.9	158
Log: Log management	43	58	36	13	6	2.8	156
Net: Ingress filtering	36	54	34	11	11	2.6	146
Host: Ransomware prevention	45	47	47	14	5	2.7	158
Net: Network segmentation	39	44	41	20	16	2.4	160
Net: Web application firewall (WAF)	38	57	34	16	11	2.6	156
Net: Web proxy	40	54	33	19	5	2.7	151
Net: DNS security/DNS firewall	34	57	39	12	9	2.6	151
Net: Network traffic monitoring	33	45	38	23	9	2.5	148
Log: DNS log monitoring	48	50	36	14	8	2.7	156
Net: DoS and DDoS protection	45	51	32	21	7	2.7	156
Analysis: Customized or tailored SIEM use-case monitoring	39	40	46	16	10	2.5	151
Log: Endpoint application log monitoring	44	51	33	19	11	2.6	158
Host: Continuous monitoring and assessment	41	52	33	18	6	2.7	150
Net: Asset discovery and inventory	35	43	46	13	14	2.5	151
Net: Egress filtering	35	59	37	13	6	2.7	150
Net: SSL/TLS traffic inspection	37	52	36	13	14	2.6	152
Analysis: Threat intelligence (open source, vendor provided)	29	43	46	19	10	2.4	147
Net: NetFlow analysis	29	43	38	22	14	2.3	146
Host: User behavior and entity monitoring	34	52	43	9	16	2.5	154
Net: Network Access Control (NAC)	31	41	50	13	14	2.4	149
Analysis: Attack surface management	27	48	42	16	14	2.4	147
Host: Behavioral analysis and detection	44	47	40	10	14	2.6	155
Host: Data loss prevention	27	45	53	11	17	2.4	153
Analysis: Threat hunting	32	48	41	16	14	2.5	151
Analysis: Threat intelligence platform (TIP)	34	47	38	17	10	2.5	146
Analysis: E-discovery (support legal requests for specific information collection)	32	41	39	20	15	2.4	147
Net: Malware detonation device (inline malware destruction)	32	45	40	11	20	2.4	148
Analysis: External threat intelligence (for online precursors)	29	47	39	23	6	2.5	144
Host: Application whitelisting	23	49	50	14	18	2.3	154
Net: Full packet capture	34	39	36	21	22	2.3	152
Net: Packet analysis (other than full PCAP)	23	47	36	16	25	2.2	147
Analysis: digital asset risk analysis and assessment	26	44	40	18	16	2.3	144
Analysis: SOAR (Security Orchestration, Automation, Response)	34	35	47	13	17	2.4	146
Analysis: AI or machine learning	25	38	43	23	19	2.2	148
Analysis: Frequency analysis for network connections	28	37	38	21	21	2.2	145
Net: Deception technologies such as honey potting	25	44	45	13	21	2.3	148
Other (Please specify)	15	13	14	5	10	2.3	57

Figure 29. GPA-Ranked Technology (Q3.39, n = 194)

Related, staff is required to use the technology because, at least for now, computers don't install or operate themselves or analyze the data they contain without human oversight. The top two technologies desired by hiring managers (by a substantial margin) were SIEM Analysis (Q3.65, 81/306, 26.5%) and EDR/XDR (Q3.65, 83/306, 27.1%) products.

See Figure 30 for the big jump (more than double the next lower value) and the ranking of the rest of the items.

In the long-form qualitative responses, SOC managers' most common need was analysts with broad technical knowledge vs. individual product or technology experience. The general feeling that an analyst who understood both how business process flows worked and how threats were likely to attack them could quickly learn how to use and extend SIEM and EDR/XDR management consoles and tools.




Figure 30. Technology/Tools for New Hires (Q3.65, n = 306)

Evaluation

As mentioned in the key findings, only a small portion (Q3.47, 39/349, 11.2%) say “No” they don’t provide metrics. Figure 31 shows this, plus that reporting SOC-related metrics regularly to board of directors and organization executives, both within and outside of cybersecurity management hierarchy, is common, but not done in the majority (Q3.47, 128/349, 36.7%) of cases.

We asked what metrics are in use, and Figure 32 shows the answers sorted on the value of “Used” for outsourced capabilities.

In some cases, the idea of these metrics being “enforced” seems untenable, but people answered that way, nonetheless. For example, enforcing a metric of “monetary cost per incident” would mean that incident handling is terminated once a certain amount of resource is expended. Perhaps this is what people are reporting they do. The authors sincerely hope this is not the case.

Does your SOC provide metrics that can be used in your reports and dashboards to gauge the ongoing status of and effectiveness of your SOC’s capabilities?




Figure 31. Metrics Reported Audience (Q3.47, n = 346)

For outsourced functions (or capabilities), what KPIs (key performance indicators) and/or metrics do you request or receive from your MSSP for tracking performance?

Indicate whether these metrics are used to enforce service level agreements (SLAs) and whether your SOC consistently meets the service level represented by that metric. Indicate N/A those that are not used.




Figure 32. MSSP Metrics/KPIs/SLAs Used, Enforced, and Consistently Met (Q3.50, n = 256)

Another aspect of metrics we asked about was cost per record, discussed briefly in key findings. Details are provided here for comparison to your in-house calculations (see Figure 33).

Based on popularity, the top values for each type are:

- Internal user account \$1–\$5 (Q3.54, 24/103)
- Customer account information: \$1–\$5 (Q3.54, 23/103)
- Credit card: \$5–\$10 (Q3.54, 22/103)

The definition of “cost per record” varies and is hard to estimate—a high percentage of respondents indicated they were not using cost per record. Large incidents can be the most damaging, but actually show the lowest cost per record. Conversely, ransomware attacks can disrupt an entire business by encrypting one key file with a small number of records, if any.

The SOC metric of time to detect/response/restore represents the only part of cost/record that the SOC actually owns. Having accurate estimation of that metric enables the SOC to support business needs for a cost/record estimate.




Figure 33. Cost per Record
(Q3.54, n = 103)

Budget and Funding

How much does all this cost? Figure 34 shows the responses at varying budget sizes. That the most popular answer is “Unknown” by a wide margin (Q3.68, 68/307, 22.1%) indicates that most SOC staff don’t have accounting duties for what this all costs. It’s as though the topic of the SOC’s expenses isn’t shared or they don’t ask about it.

Accountability, reasonable expense, frugality, and alignment come from an understanding of the real cost of resources for the SOC. The SOC offers potential loss prevention. SOC staff should understand this cost of doing business comes at the expense of other potential business expenditures. It is necessary, and should be based in financially sound practices. Understanding this financial reality is an important visibility into the business context for the SOC. If you flip back up to Figure 10, you’ll see that the top response (Q3.79, 50/313, 16.0%) is “Lack of context related to what we are seeing.” There’s a disconnect here between the business owners, the SOC cost and expenses, and the information systems used by the business. There’s no simple solution to this; it requires diligence and ongoing effort to gain context for awareness and understanding.




Figure 34. Estimated Annual Budget
(Q3.68, n = 307)

Returning to the notion of metrics, we asked what metrics were used to justify budget for the SOC. Figure 35 identifies the responses. The number of incidents handled is the most common response (Q3.75, 131/282). Hopefully, this doesn't reinforce bad behavior to show management that the SOC is needed by running up the count of incidents. The more likely use of this metric is showing how many issues arise that need appropriate detection and response, and possibly using this for a "cost per record" type justification.




Figure 35. Metrics Used to Justify Funding (Q3.75, n = 282)

The next two most highly cited metrics (time to detect/eradicate and percentage of incidents exploiting unknown vulnerabilities) are much more valuable for both corporate management and SOC operations. The board is not interested in how many raindrops are hitting the roof; they want to know if we are getting better at finding the leaks and fixing them before the business damage occurs.

Summary

We asked a lot of questions, but we also wanted to know what respondents would ask other SOCs. Here at the closing, the authors have selected their favorite question: "How have you managed to be effective despite heavy staff and resource constraints?"

It's worthwhile to note that there were many questions around people, AI, and risk. How about a word cloud to summarize it? It's in Figure 36.




Figure 36. Word Cloud of Questions for Other SOCs (Q3.81, n = 101)

This year's SOC Survey covered many points. Threat hunting and threat intel are important parts of the processes of the SOC. The most popular response on the question of the key challenge to SOCs is that there is a lack of context of the systems that are being protected. Hiring and retaining staff continues to be a challenge. Most SOCs are using metrics, and most are reporting to entities outside the SOC itself.

In 2023 there will be several additional discussions of the survey and the data. It also should be noted that a deidentified data set and Jupyter notebook is provided by the lead author (Crowley) for follow-up analysis. This is intended to help readers and respondents answer their own questions. If you have specific questions that you would like answered, the authors are interested in understanding how to improve the report for the future, and what additional information would be valuable to the community.

Sponsor

SANS would like to thank this survey's sponsor:

