Blind Spots of Security Monitoring in Enterprise Infrastructures: A Survey

Florian Skopik, Max Landauer, and Markus Wurzenberger | Austrian Institute of Technology

Cybersecurity monitoring today is laborious but straightforward: dump network traces at chokepoints in a network, collect log files from services, feed both to a proper security information and event management solution, and get alerted if something suspicious happens, right? Wrong!

Prevention ultimately fails, no matter how meticulously security measures are designed, deployed, and operated. There is always a way around this firewall, a nasty trick to elude application whitelisting, and a new technique to circumvent user authentication or authorization solutions. Intrusion detection is therefore an indispensable part of every professional security architecture. However, initial intrusion detection will often fail too, that is why it is important to gain the capabilities to spot adversarial activities after initial intrusion in any state of the cyber kill chain. Better late than never.

Continuous monitoring at all layers, in all protocols, and on all systems related to critical company assets is a vital prerequisite to effective detection and response—we can't fight what we can't see. Unfortunately, these comprehensive monitoring capabilities are widely underdeveloped in most organizations, specifically those not native to the security domain. "We are safe, because … we have antivirus, … we have a firewall, … we have this awesome security product," or even "… we have insurance" are statements that are still common these days.

Security information and event management (SIEM) and security orchestration, automation, and response solutions as well as security operations center as a service are applied by matured security teams on the other side of the same scale. These advanced solutions promise unprecedented security levels to businesses by taking monitoring data such as NetFlows, network packet captures, and log data of various core services to spot adversarial activities. However, these technologies are just tools that depend on the data we feed in and require highly skilled and well-trained analysts to operate effectively.

Threat hunting in complex organizations has become an art. Threat detection is no longer only concerned with spotting malware samples on critical systems or the exploitation of a single machine, it focuses on complex attack chains that harness multiple attack vectors spread across an infrastructure to reach the ultimate goal (usually the exfiltration of sensitive data or destruction of assets). To spot and interpret these malicious activities across services, monitoring needs to lay out a dense net where nothing slips through unnoticed. Tracking the transfer of potentially malicious files across several nodes in a complex network is one example for doing exactly that; tracking user behavior across different services—being both in house and in the cloud—is another one.

But how do we achieve this? There is no "silver bullet" to this problem, no single security solution that solves all of our troubles. In fact, we need to apply a wide variety of monitoring techniques to create visibility in the network, on the hosts, in the cloud, and at different technological layers, starting from basic networking up to high-level, web-based applications. A multitude of security monitoring solutions exist to aid us, and although the adoption of these technologies has the potential to increase security tremendously, even then,

Digital Object Identifier 10.1109/MSEC.2021.3133764
Date of current version: 4 January 2022

numerous blind spots remain. We need to be aware of the existing limitations of monitoring techniques so as to not create a false sense of security.

Monitoring Needs a Proper Architecture

In the early age of the Internet, all networks were flat. This design soon failed catastrophically with the emergence of the first Internet worms. As a consequence, networks were segmented using proxies, gateways, and firewalls, starting approximately 30 years ago.[1] From thereon, the concept of a perimeter, a protective border that surrounds a network was born. In 1990, Bill Cheswick described this concept as *a sort of crunchy shell around a soft, chewy center*. A lot has happened since then, but still, many organizations apply this "candy bar" design, characterized as a hard-on-the-outside but soft-on-the-inside, flat network with limited segmentation, and a hardened perimeter with weaker, internal systems.

Hundreds of publicly discussed cyberattacks have demonstrated that organizations applying such a network design fail miserably once attackers find a way in. This drives the need for advanced monitoring solutions. If we are not able to keep attackers out, we at least need to spot them once they are inside. Unfortunately, many miss the important fact that cybersecurity monitoring can barely be established on top of an existing architecture. It works quite the other way around: we need to carefully plan and design a proper security architecture that allows us to perform detailed monitoring. Applying the right design patterns enables the monitoring of data in motion (i.e., packets on the wire) and at rest (namely, log files and databases on hosts and servers).

A newer concept discussed in the last decade is the zero-trust model.[2] Although the traditional approach to cybersecurity relies upon barriers such as firewalls that control traffic coming in and out of a network zone, zero trust assumes that there is no perimeter that protects a network. Instead, it assumes that the system will ultimately be breached and designs security so that no entity trusts anything else by default, even in the same network, which could be compromised already. It is therefore of paramount importance to detect adversarial activities in any stage before they reach their goals and cause major harm. The foundational principles of zero trust support this goal by

- ensuring that all resources are accessed securely, regardless of their location
- adopting a least-privilege strategy and strictly enforcing access control
- inspecting and logging all traffic, hosts, and service logs.

Following these principles enables us to come up with an infrastructure design spanning the network layer and hosts up to third-party services, which allows effective monitoring and logging using state-of-the-art technologies.

The aim of this article is to make a larger audience aware of the complex issues of security monitoring. This is more than picking a random SIEM solution and dropping it into an existing infrastructure, which is a common misconception. Security monitoring needs to be "designed into" a network and cannot be established on top; otherwise, it creates blind spots. The article justifies exactly that view on the design, rather than the implementation aspects.

When realizing a monitoring strategy, we roughly run through the following phases:

1. *Phase 1*: Decide what to monitor, depending on what the high-risk areas are and which types of attack vectors need to be discovered.
2. *Phase 2*: Investigate which data sources to utilize and how to make them accessible.
3. *Phase 3*: Pick technical solutions for data collection, forwarding, extraction, aggregation, and normalization.
4. *Phase 4*: Select feasible data analysis approaches.
5. *Phase 5*: Create reporting and alerting mechanisms.

We argue that phase 1 is the one most prone to blind spots. Of course, we can always make a bad technical decision or have struggles with unsupported software; however, neglecting an important monitoring area in the first place is something we cannot mitigate in later phases. In that sense, in this article, we mainly focus on monitoring blind spots on a methodological and conceptual layer, not on a technical layer.

Approach to Identifying Blind Spots

We started our investigation of blind spots with a literature survey with the keywords (and variations thereof) *cybersecurity monitoring*, *defensible network architectures*, *network visibility* and *intrusion detection* in IEEE *Xplore*, ACM Digital Library, Scopus, and Google Scholar. The aim was to identify the commonalities of modern enterprise networks, and suitable cybersecurity monitoring and logging solutions as well as their typical use cases. As we discuss the application aspects of these solutions in operational environments, we did not conduct our survey with common academic sources alone, but instead further focused on industry reports, white papers, and the collected wisdom of industry experts. Numerous industry best practices are available in the domain, including those from large equipment vendors such as Cisco;[3] recommendations from governmental bodies, including the Australian Cyber Security Center (ACSC);[4] and nonprofit organizations, such as the Center for Internet Security,[5] just to name a few.

Furthermore, large standardization bodies, such as the National Institute of Standards and Technology, have prepared and documented their framework for defensible networks.[6] We studied those in detail, and extracted common topics they address and typical questions they ask, to identify the important aspects of modern cybersecurity monitoring and logging.

In the course of our research, we were particularly interested in the following two questions:

1. How does the academic sector and industry define modern network architectures?
2. What are common limitations of monitoring in the context of these architectures?

Specifically, we were interested in case studies that apply modern network security principles and that further reveal issues and pitfalls as well as limitations of certain monitoring solutions.

To structure our further discussions, we use the reference architecture in Figure 1, which is inspired by the mentioned sources and carefully designed applying the following design rules:

- *Segmentation*: Although the traditional approach to network security clusters all machines/services of the same importance (trust levels) into groups (e.g., there is a segment for servers, employees, and guests), zero trust goes a step further and usually applies microsegmentation, where no direct interaction between any two machines in a segment is allowed, and every data exchange takes place via dedicated proxies only. Private virtual local area networks (VLANs) and station isolation modes are two means for wired and wireless devices, respectively, that implement this concept.
- *Controlled data flows*: On top of a segmented infrastructure, only specific data flows justified by business needs are allowed (while everything else is denied). In reality, this is often defined for whole segments, e.g., the marketing staff is allowed to access the web server, but the office is not. However, in a zero-trust network, this is enforced even on a single-machine level. Software-defined networks (SDNs) that enable flexible, centralized management of the control layer is a cornerstone to achieving that.
- *Hardened and uniform configurations*: In a mature environment, machines use carefully designed software baselines for operating systems and service configurations that account for foundational principles of cybersecurity, such as the need to know and least privileges. A uniform environment tremendously simplifies the analysis of monitoring data.
- *Extensive monitoring*: To close the control loop, we need to regularly monitor whether our preventive measures (firewalls, access control lists, application whitelists, and so on) are still effective and haven't been circumvented the one or the other way. For that reason, we monitor all data flows on the network layer to identify bandwidth occupancy, communication patterns, and utilized protocols as well as user-authentication events, service usage, endpoint behavior, web traffic, and so forth.

Following the current state of the art resulting from our survey, and specifically the stated design rules, we come up with an example infrastructure (see Figure 1), which closely resembles networks of well-matured organizations. Note that a part of almost every modern architecture are also different web-based services that offer infrastructures, platforms, or software as a service (SaaS) via either public, private, or hybrid clouds. In our design, we foresee several network zones that are well separated through firewalls. Within some network zones, we further restrict data flows through private VLANs and build chokepoints with proxies. We also deployed proper monitoring equipment, such as network taps and switches with port mirrors, to dump raw packets and network security monitoring (NSM) solutions that extract information from these packets. Endpoint detection and response (EDR) solutions enable close monitoring of hosts. Next-generation firewalls with Transport Layer Security (TLS) inspection, a sandbox to investigate suspicious executable files (malware sandbox), a central log store to collect log data from hosts (e.g., either pushed through agents or fetched from machines via network shares), and an SIEM solution complement this setup.

Common Blind Spots

Even after carefully applying the wealth of state-of-the-art principles, best practices and recommendations for defensible network architectures, and deploying the newest cybersecurity monitoring solutions, a number of blind spots remain. Figure 1 visualizes seven issues, depicted with star symbols. In this section, we take a closer look at them.

Encryption Is a Double-Edged Sword

The encryption of data in transit is an important means to ensure confidentiality, especially in insecure or public networks; but it is also a key principle of zero-trust architectures. TLS is the way to go. Today, virtually all web-based services use TLS, including common web servers and mail or file servers. With the emergence of free certificate services, such as the popular Let's Encrypt offering, the web has tremendously changed its appearance in the last couple of years. Large websites dealing with sensitive data are not the only ones using TLS, but

also private blogs and pages, which is both good and bad for our internal security. Of course, proper certificates reduce, to some extent, the risk of successful impersonation or man-in-the-middle attacks (however, investigating this fact in detail would justify an article on its own). On the other side, end-to-end encryption significantly stalls our monitoring efforts and makes us blind in one eye. By recording the Domain Name System (DNS) logs, we can still see which domains the clients in our network request. Additionally, with a web proxy, we have a perfect chokepoint to still see the initial communication attempts and the TLS handshake, including certificates presented by the web servers, but not any further into the actually exchanged data.

Anyway, we can still argue that often it is not required to investigate connections in depth, and that we can still see which clients connect to which external servers, in which frequency, and for how long as well as the amount of exchanged data. Although this could allow us to spot, e.g., command-and-control traffic, the next issue is just around the corner. Thus far, the DNS provided us a wealth of information, specifically the domains of requested sites—but what if external DNS servers are used? Then we can still detect traffic through our NSM solution and inspect that to learn about requested sites. And what if DNS traffic is encrypted too? DNS-over-HTTPS (DoH) and DNS-over-TLS aim to do exactly that. DoH especially is horrifying from a security monitoring perspective. With DoH, DNS is not only encrypted but also flows through the standard HTTPS port 443, usually to an external provider, such as Cloudflare or NextDNS, in the case of Mozilla Firefox. Nothing to see here.

Mitigating actions. Many organizations have therefore started to break open TLS connections "at the perimeter" (which, as we learned, is diminishing), i.e., use a web application firewall (WAF) or proxy, which terminates TLS connections, inspects them, and then re-encrypts the content in the course of new TLS connections. To run smoothly, certificates compatible with the WAF's or proxy's private key are deployed at the clients, otherwise they would report a broken trust chain. Note that data privacy

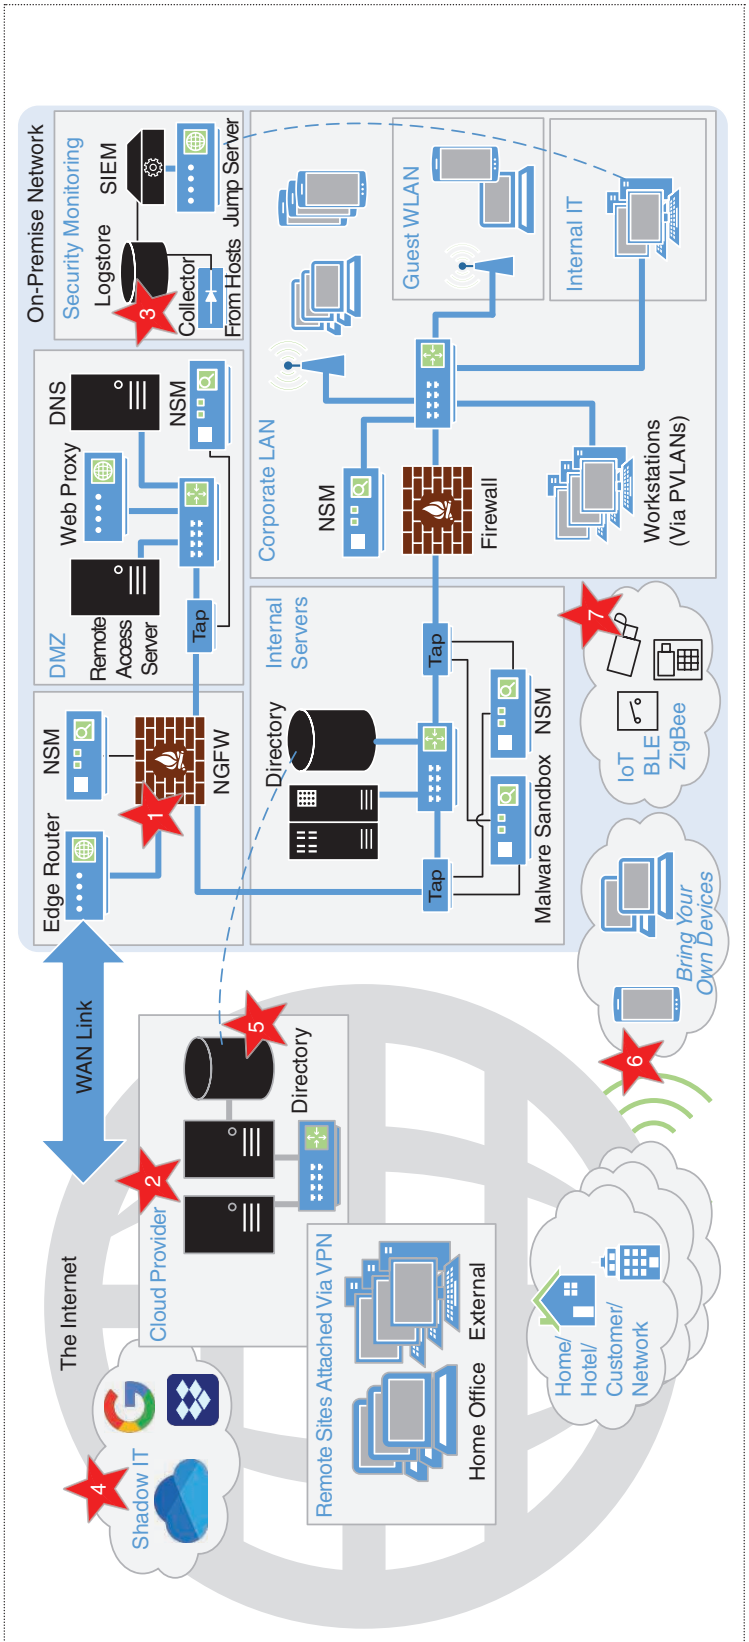

Figure 1. A typical network infrastructure, which also accounts for equipment and services off premise, and the seven identified issues of cybersecurity monitoring: 1) encryption, 2) the cloud, 3) data volume, 4) shadow IT, 5) scattered identities, 6) multihomed devices, and 7) nonenterprise IT. VPN: virtual private network; BYODs: bring your own devices; NSM: network security monitoring; NGFW: next-generation firewall; IoT: Internet of Things; BLE: Bluetooth Low Energy; LAN: local area network; WLAN: wide local area network; PVLANs: private virtual local area networks.

usually does not allow for breaking all the encrypted connections; commonly, certain site categories, such as finance- or health-related services, are excluded.

The Cloud Is Somewhat "No Man's Land"

Outsourcing is a major driver of business growth: focus on your core competencies and do what you do best, while others take care of essential IT services. Regardless of whether infrastructure as a service (IaaS), platform as a service (PaaS), or SaaS models are used, they all pose different challenges to monitoring if others run the infrastructure (or platform and software, respectively) for you. In these cases, we do not have access to monitoring data from the layer underneath the rented services. For instance, if we rent a number of virtual machines, we do not have access to the underlying networking infrastructure; if we obtain whole software services, such as Microsoft Office 365, we do not get monitoring data from the machines that run the instance we use.

Major IaaS/PaaS providers, such as Amazon, Microsoft, or Google, have realized that this creates a monitoring blind spot and have therefore started to provide connectors to retrieve curated logs from their environment and allow the analysis of cloud audit trails.[7] Additionally, user entity behavior analysis (UEBA)[8] is often part of an additional service offering.

However, utilizing monitoring and audit data from cloud services implies that we know about them being used by our company's employees. That is certainly true for large cloud services strategically purchased by the company but not so much for so-called shadow IT.[9] These are usually free IT services, unauthorized by the company but used by employees out of convenience, such as noncommercial (personal) offerings of cloud storage or collaborative web editors.

Even private cloud setups, making heavy use of containerized or virtualized setups, e.g., using Kubernetes or OpenShift, may impose a blind spot. Monitoring an environment where containers are designed by solution developers rather than classic system administrators and dynamically instantiated in a multilayered virtualized setup is a whole different story. Even if this environment is entirely operated on premise, we might overlook that the monitoring strategy needs to be carefully coordinated between the developers and operators of the containers, otherwise we are not able to monitor what's going on inside a container.

Mitigating actions. Processing data in containerized or virtualized environments as well as entirely externally operated infrastructures that we can't properly monitor obviously poses a certain risk. Tight coordination between all involved entities, mainly the infrastructure operators, container designers, and application developers, is key. Moreover, the deployment of honey tokens[10] in the course of data-loss-prevention activities are an interesting means to reduce the risk of large-scale data loss (or at least increase the timely visibility of these losses). Honey tokens are fake words, projects, users, and so on, which are unique, but artificial strings distributed in certain documents. If one of these strings is spotted outside the intended environment, a data leak—intentional or unintentional—might be the probable cause, such as an employee who copied data out of a secured environment, sent it to the wrong recipients, or accidentally put them online in shared public folders.

Data Volume and Quality Is Key

Collecting just the right amount of data is a tough job. In a perfect world with unlimited resources, we would like to collect all data from all available sources and analyze them with all available means to achieve maximum visibility. However, we are not living in a perfect world, and everyone is constrained by economic pressure. We must therefore carefully decide where we invest our scarce resources.

Kernel-level logs (e.g., from a verbosely configured audit daemon) and full-network packet captures always contain the traces of attacks on technical systems, regardless of how advanced and stealthy the attack is. The pity, however, is that there are no economically feasible analysis techniques that allow us to sift through these massive amounts of data to discover them reliably, and that preferably do this in near-real time. The amount of data is simply too much, and additionally, data structures are often quite complex or not fully documented nor understood (especially for application-specific logging). As a consequence, most attack-detection approaches rely on just a few quite-specific data sources only (e.g., authentication logs; compare Table 1) and try to discover malicious activities from a rather limited set of data features as this is fast and easily automatable. However, the risk is that we either collect tremendous amounts of data but do not properly exploit them, neglect important data sources completely due to misconfiguration or misunderstanding of their importance, or apply ineffective analysis techniques and create blind spots as a result.

Mitigating actions. Numerous guidelines and industry best practices make well-applicable and justified proposals for data collection, correlation, and analysis techniques.[4–6] Besides the general recommendations at the conceptual layer, detailed technical guidelines help to avoid troubles with inappropriately configured log data agents, overloaded network taps, mismatching time stamps, varying data formats, mixed encoding styles, and truncated log lines.[11]

Shadow IT and Other Unknown Assets

With increasing remote work via virtual private networks (VPNs), often from private devices, our monitoring opportunities are shrinking. From these unmanaged devices that are directly connected to an enterprise's network, no host logs can be collected, and barely any policies can be enforced. Of course, we see connection attempts of devices on the remote access server or VPN concentrator and authentication events against a central directory (e.g., active directory), but this amount of information is significantly less compared to what a company-managed and monitored device would reveal. A similar challenge arises from guest Wi-Fis and bring your own devices (BYODs), which are therefore usually confined to a dedicated network segment with strict access control lists.

Personal email accounts being used to conduct business, unsanctioned BYODs, and third-party SaaS applications outside the purview of the IT department are well-known examples of so-called shadow IT. For these assets, no monitoring data exist and no log data are collected.

Mitigating actions. For overcoming this challenge, organizations need to take a hybrid approach to continuous monitoring. Combining passive, real-time monitoring that collects logs and network traffic with an always-on active scanner can provide both clarity of vulnerable endpoints and detection of previously unknown assets. By actively scanning the network for new devices as well as investigating network traffic (including that of these devices), we gain the ability to detect the network activities of unknown devices that aren't present in the official asset list. This is how an enterprise can effectively create a more modern security infrastructure for endpoint activity.

Even social media is a form of shadow IT if users put potentially confidential data about their work online. Again, although the social media platform cannot directly be monitored, active scanning for such contents helps to reveal data breaches early.

Scattered Instead of Federated Identities

SaaS, hosted by cloud providers but seamlessly integrated with software and services deployed on premise, has become standard in recent years. Microsoft's Office 365 is a good example for that. This seamless integration relies on a federated identity concept, where identity information is shared with, and synchronized between, a directory on premise and one in the cloud. In general, a trust relationship among separate organizations, including application vendors or partners, enables them to share identity information and authenticate users across domains. This is the basis of single-sign-on (SSO) for services across organizational boundaries.

Although SSO seems to work mostly satisfactory in terms of user experience, federated identities are challenging from a security perspective. On the one side, different copies of user identities need to be kept in sync from a technical point of view, which is complex and sometimes flawed. On the other side, user-behavior analytics cannot exploit its full potential because the logs concerning a user's activities are not available for all services at a single place.

A somewhat special case to consider is when the authentication mechanisms of well-known and trusted web platforms are used to authenticate against third-party sites. For instance, a company could decide to use Github for its development activities, and employees receive dedicated Github accounts linked to the company. However, said employees could also use their accounts to authenticate via Github toward other platforms using OAuth,[12] and that would go completely unnoticed by the organization as none of their own monitored services would be involved in that process.

Eventually, even today, truly working federated identity management solutions are rare and can never be fully enforced across all relevant services. In fact, a LastPass survey from 2017[13] discovered that the average business user stores 191 passwords (presumably for an equal number of accounts), which reveals a long way to go until we reach a status of seamless SSO.

Mitigating actions. In light of these facts, effective user-behavior analytics can only work for isolated services or within well-defined boundaries. Organizations need to define and enforce these boundaries and

Table 1. The monitoring areas and their data sources.

Monitoring area	Examples
Low-level network monitoring	Network equipment (switch, router, and firewall), network intrusion detection systems, and extractions from pcaps
Application layer (server and client landscape)	DNS, HTTP, Simple Mail-Transfer Protocol, authentication logs (Windows Active Directory and Lightweight Directory Access Protocol), TLS certificate inspection, malware-detonation chamber, next-generation firewall, WAFs, and Windows event logs
Virtualization and the cloud	Virtual infrastructure logs, container logs, hypervisor logs, and cloud audit trails
User and entity behavior monitoring	Authentication logs, activity logs, and application-layer logs
Open source intelligence analysis	Social media activities and information shared on public platforms

must not become reluctant out of convenience. Several different system reports will need to be analyzed to ensure that data are clean, complete, and of high quality. Accounting for detection of misuse before roll out of a service, ideally, even in its design phase, is tremendously helpful and eases the lives of security personnel later.

Multihomed Devices and Roaming Users

Dual-homed devices are those with two network interfaces, used simultaneously, which effectively bridge two disparate network segments. A common example is a laptop connected to the company network that at the same time uses an LTE modem to establish a direct connection to the Internet. An interface can even be virtual; for instance, a computer residing within an employee's home network that at the same time is connected via a VPN to a company network. This might cause security issues as such a device poses a connection point between two often carefully separated networks or network zones.

A different, but related challenge, are devices that frequently change among networks. For instance, employees connect their laptops to the company network but take their machines with them anytime and also plug them into their home networks, airport, or hotel Wi-Fis as well as into customer infrastructures. Apart from rather static, classic network designs, this issue is specifically important in dynamic environments, where devices may be connected through SDNs and flexibly roam among dynamically defined network segments.

The issue with roaming devices is that their monitoring is often severely limited. Either the device itself consists of a monitoring agent, which might, however, not be able to communicate with the backend monitoring system via public/foreign networks, or monitoring takes place passively through the inspection of network traffic. Then, however, it is important to identify and match devices on different interfaces to gain a complete picture about a device's activities.

Similarly, this problem exists on a user level too. Users may use different devices to fulfill certain tasks. For an effective UEBA and user profiling, these activities need to be combined, e.g., regardless of whether a user performs activities on a laptop, mobile phone, tablet, or home computer, profiles need to be matched and completed—a tricky task if identity management is not properly configured. Eventually, either on the device or user level, an improper combination of monitoring data from different devices (or the same devices but collected on different interfaces) hinders full visibility of all relevant activities of an entity and thus poses the risk of a blind spot.

Mitigating actions. Numerous solutions exist for dealing with that issue, e.g., "Always-On VPN." As the name implies, this is a mandatory VPN connection that always places the machine logically into the company network (given that split tunneling is deactivated). Network access control is a solution that allows access to a network only if certain constraints are fulfilled, e.g., a monitoring policy is enforced on an endpoint or a certain EDR solution is active. These technologies, together with clean user identity management that correlates user activities from different devices, is key to centrally collecting distributed user events and matching them into one profile for inspection with UEBA.

Ignored Nonenterprise IT

Enterprise IT focuses on servers and client PCs that run the essential applications and services to conduct business and the networking components that connect the machines. Some of them might be virtual components, others could be mobile components, but they are usually all managed by a company's IT department. Even smartphones are commonly included in enterprise IT management through mobile device management solutions. Nothing stays under the radar, in theory.

The Internet of Things (IoT) has massively transformed the IT landscape.[14] Numerous protocols and transmission technologies, including Bluetooth Low Energy, Z-Wave, and Zigbee, among others, are a part of most modern environments. They are used to control locking systems; heating, ventilation, and air-conditioning (HVAC); Internet Protocol-based video-cameras; and much more. Usually these systems are not controlled by the IT department but by separate organizational units in the course of building management. However, these systems, enterprise IT and nonenterprise IT, almost always share infrastructure components. If they do not use the same wires (e.g., using VLANs), wireless IoT components share at least the same transmission medium as an IT-managed 802.11 Wi-Fi device. This could be a threat to the availability of a network link but is barely taken into account, because enterprise IT management is often insufficiently informed about other IT installations.

Mitigating actions. Monitoring often stops where traditional enterprise IT ends. Although we all agree that monitoring the server landscape is essential, what is that good for if the HVAC system of a data center is being compromised and threatens the availability of IT services? Therefore, we need to break open the unfortunately still-common silo mentality and connect often desperate organizational units of building management and enterprise IT. If technologies under their supervision interact, responsible people need to cooperate too.

BONUS: Regulatory Constraints and Privacy Issues

Regulatory constraints, such as those imposed by the General Data Protection Regulation in the European

Union but also sector-specific regulations around the world, e.g., protecting employees' rights, can lead to unintentional side effects. Limiting the amount of collected data, the time span of storage, and the type of processing (specifically when it comes to correlating data across systems) may have unpleasant side effects on security. If data anonymization hinders our correlation effort, or if data was deleted or not even recorded in the first place due to legal and compliance constraints, modern UEBA approaches may be stalled and create an "artificial" blind spot. Thus, monitoring and analysis solutions that respect these regulatory requirements need to be employed to follow the requirements and maintain a high level of security at the same time.[15]

Modern solutions make large-scale cybersecurity monitoring a reality. Full packet captures on 10-Gb links and detailed kernel-level logs on virtual machines are within reach. WAFs, TLS inspection, and malware detonation devices enable deep insights into network operations. However, even with these tools, several blind spots remain, and malicious activities may go unnoticed in the context of cloud services, shadow IT, encryption, and multihomed devices.

Specifically insidious is the combination of these blind spots that nevertheless exist in almost every enterprise today. Imagine a company frequently buying goods from a distributer via an external web shop. The purchase department uses a company account that is shared by three employees to order goods. One of them quits. Obviously, the shared password should be changed immediately for security reasons; however, it rarely is. The external service is not a part of any internal asset list, and the account is neither managed with an internal directory (like Windows Active Directory) nor directly linked to an individual person. If the now-former employee enters the platform after leaving and makes orders on behalf of the company, reroutes goods, or changes payment details, no one will recognize it immediately. The platform is not a part of the company's assets, the account activities are usually not monitored, and authentication is completely externalized.

This is an example of a simple form of supply-chain attack, one that has become so prominent in recent years. This case alone shows that there is still a long way to go to effectively closing remaining loop holes in today's security monitoring architectures.

Help us to fix our blind spots! Please tell us which blind spots we did not mention in this article; email aecid@ait.ac.at. ■

Acknowledgments

This work was supported, in part, by the Austrian FFG project CADSP under grant 873425, and in part by the European Union H2020 project GUARD under grant 833456 and the EDIDP project PANDORA under grant SI2.835928.

References

1. L. L. Peterson and B. S. Davie, *Computer Networks: A Systems Approach*. Amsterdam, The Netherlands: Elsevier, 2007.
2. J. Kindervag and S. Balaouras, *No More Chewy Centers: Introducing the Zero Trust Model of Information Security*. Cambridge, MA, USA: Forrester Research, 2010.
3. "Network security baseline," Cisco, San Jose, CA, USA, 2008. [Online]. Available: https://www.cisco.com/c/en/us/td/docs/solutions/Enterprise/Security/Baseline_Security/securebasebook.pdf
4. "Essential eight," Australian Cyber Security Centre, Canberra, Australia, 2021. Accessed: Jul. 12, 2021. [Online]. Available: https://www.cyber.gov.au/acsc/view-all-content/essential-eight
5. " The 18 CIS critical security controls," Center for Internet Security, Washington, DC, USA, 2021. Accessed: Jul. 12, 2021. https://www.cisecurity.org/controls/cis-controls-list/
6. "Cybersecurity framework," NIST, Gaithersburg, MD, USA, 2021. Accessed: May 28, 2021. [Online]. Available: https://csrc.nist.gov/Projects/cybersecurity-framework/nist-cybersecurity-framework-a-quick-start-guide
7. R. A. K. Duncan and M. Whittington, "Enhancing cloud security and privacy: The power and the weakness of the audit trail," in *Proc. 7th Int. Conf. Cloud Computing, GRIDs, and Virtualization,* Rome, Italy, 2016. [Online]. Available: https://www.iaria.org/conferences2016/CLOUDCOMPUTING16.html
8. M. Shashanka, M.-Y. Shen, and J. Wang, "User and entity behavior analytics for enterprise security," in *Proc. IEEE Int. Conf. Big Data (Big Data)*, Dec. 2016, pp. 1867–1874, doi: 10.1109/BigData.2016.7840805.
9. M. Silic and A. Back, "Shadow IT–A view from behind the curtain," *Comput. Security*, vol. 45, pp. 274–283, Sep. 2014, doi: 10.1016/j.cose.2014.06.007.
10. M. Bercovitch, M. Renford, L. Hasson, A. Shabtai, L. Rokach, and Y. Elovici, "HoneyGen: An automated honeytokens generator," in *Proc. IEEE Int. Conf. Intell. Security Inf.*, Jul. 2011, pp. 131–136, doi: 10.1109/ISI.2011.5984063.
11. A. Chuvakin, K. Schmidt, and C. Phillips, *Logging and Log Management: The Authoritative Guide to Understanding the Concepts Surrounding Logging and Log Management*. Rockland, MA: Syngress, 2012.
12. D. Hardt, Ed. "The OAuth 2.0 authorization framework," Internet Engineering Task Force, Fremont, CA, USA, RFC 6749, 2012.

13. "The password exposé," LastPass, 2007. https://lp-cdn.lastpass.com/lporcamedia/document-library/lastpass/pdf/en/LastPass-Enterprise-The-Password-Expose-Ebook-v2.pdf
14. Z. K. Zhang, M. C. Y. Cho, C. W. Wang, C. W. Hsu, C. K. Chen, and S. Shieh, "IoT security: Ongoing challenges and research opportunities," in *Proc. IEEE 7th Int. Conf. Service-Oriented Comput. Appl.*, Nov. 2014, pp. 230–234, doi: 10.1109/SOCA.2014.58.
15. F. Menges *et al.*, "Towards GDPR-compliant data processing in modern SIEM systems," *Comput. Security*, vol. 103, p. 102,165, Apr. 2021, doi: 10.1016/j.cose.2020.102165.

Florian Skopik is head of the Cybersecurity Research Program at the Austrian Institute of Technology, Vienna, 1210, Austria. His main research interests are centered on critical infrastructure protection, intrusion detection, and national cybersecurity. Skopik received a Ph.D. in computer science from the Vienna University of Technology. He is a member of various conference program committees, including the Symposium on Applied Computing, International Conference on Availability, Reliability and Security, International Conference on Critical Information Infrastructure Security, editorial boards, and standardization groups, such as the ETSI Technical Committee on Cybersecurity, International Federation for Information Processing Technical Committee 11 Working Group 1, and OASIS Cyber Threat Intelligence. He is a Senior Member of IEEE. Contact him at florian.skopik@ait.ac.at.

Max Landauer is a scientist at the Austrian Institute of Technology, Vienna, 1210, Austria. His main research interests are log data analysis, anomaly detection, and cyberthreat intelligence. Landauer received his master's degree in computer science from the Vienna University of Technology, Austria. Contact him at max.landauer@ait.ac.at.

Markus Wurzenberger is a scientist and project manager at the Austrian Institute of Technology, Vienna, 1210, Austria. His main research interests are log data analysis with a focus on anomaly detection and cyberthreat intelligence. Wurzenberger received a Ph.D. in computer science from the Vienna University of Technology. Contact him at markus.wurzenberger@ait.ac.at.

Digital Object Identifier 10.1109/MSEC.2022.3211752