Introduction: On the Nature of Situational Awareness

Situational awareness is a basic component in the prevention, identification, mitigation, and elimination of digital (cyber) threats. Situational awareness refers to gathering information, perceiving and understanding the state of the world, and predicting future states of the world. Without awareness, analysts lack a complete basis for recommending decisions, particularly in complex situations such as network security. This special issue presents a variety of research and practice related to knowing what digital systems should be doing, tracking what is happening, inferring when “should be” and “is” do not match, and acting on the difference.

Situational awareness derives from the foundational work by military strategist John Boyd and by Mica Endsley. Their work in physical-world awareness and decision making has been applied by analogy to digital networks. Their work led to describing processes of observing behaviors, orienting them in the operational situation of the observer, projecting possible future conditions based on available actions by the observer and the opposing actors, and then choosing the action leading to the best outcomes. In more modern work, this has been recast for information networks into a four-phase process described below.

Before we can understand the cybersecurity state of an organization, we need to understand what should be going on in that organization. In particular, we must know:

- Who are the legitimate users of internal and public-facing systems and devices?
- What devices are authorized and what are they used for?
- Which processes and applications are approved, where are they allowed, and how do they serve the organization?

The more precise the situational information available to cybersecurity personnel, the easier it can be for them to infer when there are security issues and do something about them. Precise information requires well-defined security policies, effective access controls, up-to-date inventories, and detailed network diagrams. Unfortunately, that information is often poorly documented, incomplete, or outdated. This forces analysts to infer the missing information, which at best provides only a semi-accurate picture of the intended organizational architecture and usage.

Knowing what *should be* and knowing what *is* are almost always different in practice. The former is about gathering information on organizational intentions (what organizations mean to do and are allowed to do to accomplish their goals). The latter is about examining the organization’s network to see what is really going on. Security teams cannot directly monitor all of cyberspace, and often struggle to obtain sufficient insight into their own networks; they must use a variety of tools to see into geographically scattered, and largely invisible, cyberspace areas. Accurate observations come from answering questions such as:

- Which observed devices, processes/applications, and users are active (see, for example, Ring et. al. in this issue)?

ACM Reference format:

Josiah Dykstra, Neil Rowe, Timothy Shimeall, Angela Horneman, and Marisa Midler. 2021. Introduction: On the Nature of Situational Awareness. *Digit. Threat.: Res. Pract.* 2, 4, Article 25 (October 2021), 3 pages.
<https://doi.org/10.1145/3462334>

© 2021 Copyright held by the owner/author(s).
 2576-5337/2021/10-ART25
<https://doi.org/10.1145/3462334>

- What vulnerabilities are known for the observed devices, processes, and applications (see for example, Samtani et al. in this issue)?
- How usage of systems and devices has changed over time?
- What the usage patterns and cycles are for systems, devices, and users?

Tools for achieving useful awareness typically aggregate information from sensing points and integrate that information in a way that makes it meaningful to analysts supporting security functions that infer when *should be* and *is* do not match. However, the sensing architecture required for maintaining awareness is costly and resource-intensive for humans and technology. Allowing processes and analysts to access and combine information effectively requires designing a robust federated or distributed system for situational awareness.

A security issue occurs when an event violates policy (perhaps implicit), e.g., a device is accessed by an unauthorized individual, a recording device taps a network, a crypto-miner is run on a Web server, etc. Some of these are easy to detect, others require linking several observations or indicators together. For example, if security logging is enabled on a device, you can see unauthorized accounts attempting to access the device in the security log. Another example is if all endpoint devices must use an internal domain-name resolver, any that do not can be found by checking outbound network traffic.

Unfortunately, many important modern cybersecurity issues require inference. For example, while security logging can track when a user successfully logs into a system, it cannot determine if a login is by the real individual assigned to the account or whether the individual's credentials have been compromised. That determination requires inference, which is a significant challenge. Inference can use many data sources, as with:

- Rule-based detection of direct policy violations
- Observations of volume or endpoint deviations from historical data (significant changes in *what is*)
- Unusual outliers in values such as ports or protocols
- Observations of new services or hosts generating network activity
- Matching traffic against described attack tactics, techniques, and procedures (see for example, Pruvine et al. in this issue)

Some traffic identified by these methods will be malicious (for example, phishing email). Other traffic may be inadvertent but harmful (for example, poorly formatted email that causes service software to crash). Other traffic may be harmful only by coincidence (for example, a flash crowd overwhelming a server). Still, other traffic may just be benign normal activity. Separating the harmful traffic from the benign-but-anomalous is a key part of differential awareness. Actionable differences include business and efficiency issues as well as security issues. Comparing information from “*track what is*” data collection against all relevant information from “*know what should be*” data collection can be technologically impossible or practically infeasible. Deciding how to choose which observations should be compared in what context (e.g., inbound responses to requests, scans, service probes, or brute force attempts) is a matter of priority and resources. The context of concern, traffic visible by the sensor, and capabilities to respond to any detected issues must accurately reflect mission and business priorities.

It does little good to know what should be, track what is, or infer differences if an organization does not act on the knowledge. Organizations will usually respond to explicit security breaches. They clean up malware infections, investigate potential data leakage (see for example, Happa et al. in this issue), report stolen resources, and report personally identifying information. Some responses are required by law, and some mitigate technical, financial, or reputational damage. Organizations are less likely to do something about differences between “what should be” and “what is” if they do not believe that the differences represent a security incident. Such oversights can make inferring security events more difficult in the future. The more items that do not match “what should be” (i.e., approved users, devices, and usage), the more noise that clutters and interferes with a clear picture.

Organizations must ensure that situational awareness findings get routed and resolved by the responsible parts of the organization for the assets involved, and that the owners identify ways to prevent such issues in the future. They can do this by maintaining productive communication channels (see for example, Aleroud et al. in this issue) throughout the organization, and quickly communicating findings, contextual information, and actionable intelligence to allow the parties responsible to resolve issues. However, success requires organizational accountability, managed relationships, and clearly defined areas of responsibility. Organizational politics, turf wars, and unclear product and process owners could interfere.

Valuable and high-quality research and practice are being done in situational awareness, as reflected in the papers of this special issue, but clearly much more work needs to be done in this area.

This material is based upon work funded and supported by the Department of Defense under Contract No. FA8702-15-D-0002 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center. (DM20-1030). The views expressed in this paper are those of the authors and do not necessarily represent the official policy or position of the Department of Defense or the U.S. Federal Government.

Josiah Dykstra
Neil Rowe
Timothy Shimeall
Angela Horneman
Marisa Midler