SANS

GUIDE TO SECURITY OPERATIONS

SEC450: **Blue Team Fundamentals: Security Operations and Analysis | GSOC**
sans.org/SEC450 | giac.org/gsoc

LDR551: **Building and Leading Security Operations Centers | GSOM**
sans.org/LDR551 | giac.org/gsom

Onto the Introduction ›

Introduction

While cyber defense is an enormously in-depth topic, there are certain mindsets, models, data sources, and techniques that can get any team started off on the right foot. This guide is a collection of some of the most useful information and models for those working in cybersecurity operations centers, as well as pointers to some incredibly powerful free tools, book references, and more to help build your team, skills, and defensive capabilities.

Onto the Table of Contents ›

Table of Contents

Table Of Contents ›

SOC Functions

When breaking down how a Security Operations Center functions, it can be helpful to decompose the complex workflow of events into atomic functions so that each can be studied. In any complex system (including SOCs), each function of a process has a specific set of inputs and outputs that can be generalized, understood, and measured to assess whether that component is performing as intended. Breaking down the components and tasks of a SOC in this way helps describe in clear detail exactly what a security team must do in order to succeed as well as how other groups and functions interact with the SOC core functions to provide all services necessary for success in cyber defense.

The following is a list of SOC core functions – tasks that commonly fall under the “SOC” organization, although this may, of course, vary from team to team depending on a number of factors. Each is listed with inputs, outputs, and that function’s goals, as well as interactions with other groups. “Auxiliary functions” are also listed. These are capabilities that less commonly fall under the SOC organization directly, but often operate very closely with the security operations team to ensure that the organization is secure.

Table Of Contents ›

SOC Functions Diagram

In the diagram below, core functions are drawn in solid boxes while auxiliary functions are in dashed outlines. The placement of the auxiliary functions in the diagram is not to be taken as an organizational chart recommendation, but merely to show the inputs from those groups/functions to the core tasks of the SOC. An important item to note is that since these functions form a serialized chain of inputs and outputs, failure to perform the earlier items in the process will have ramifications on capabilities further down the line. This means that functions earlier in the process should be optimized and focused on to ensure the best possible outcome.

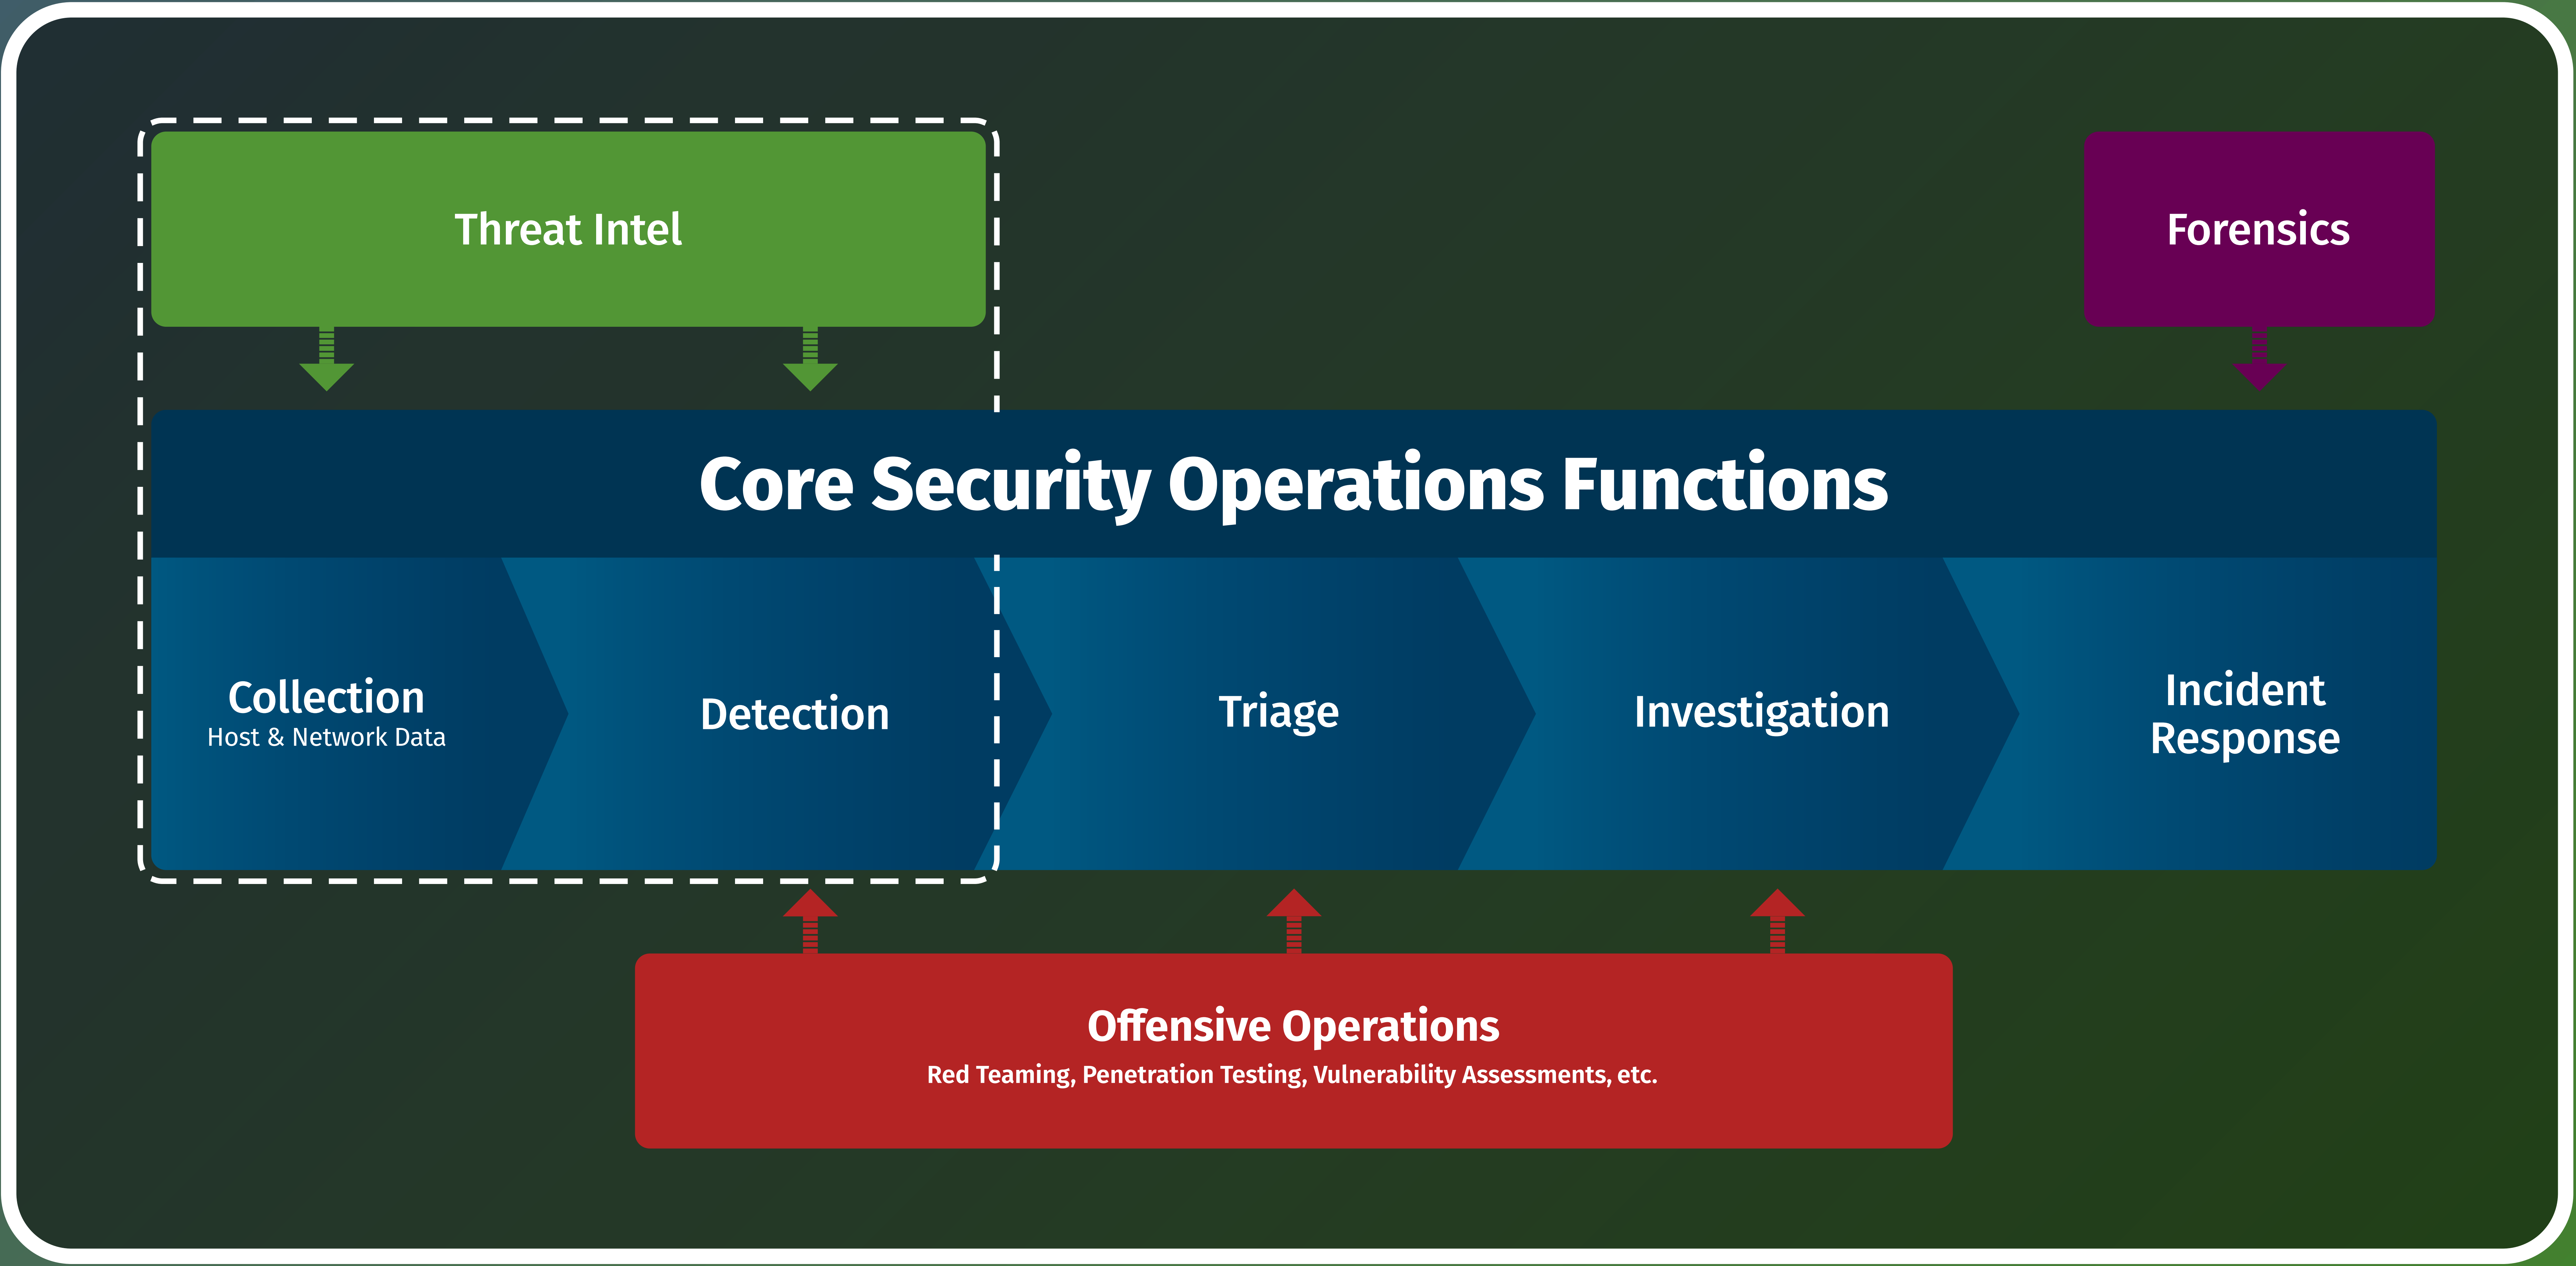

Table Of Contents ›

SOC Core Functions

The following are the deconstructed pieces of running a Security Operations Center and are listed as “core” activities that the average cyber defense team is responsible for. The section that follows will discuss each function, its goals, and how to predict that function's effectiveness.

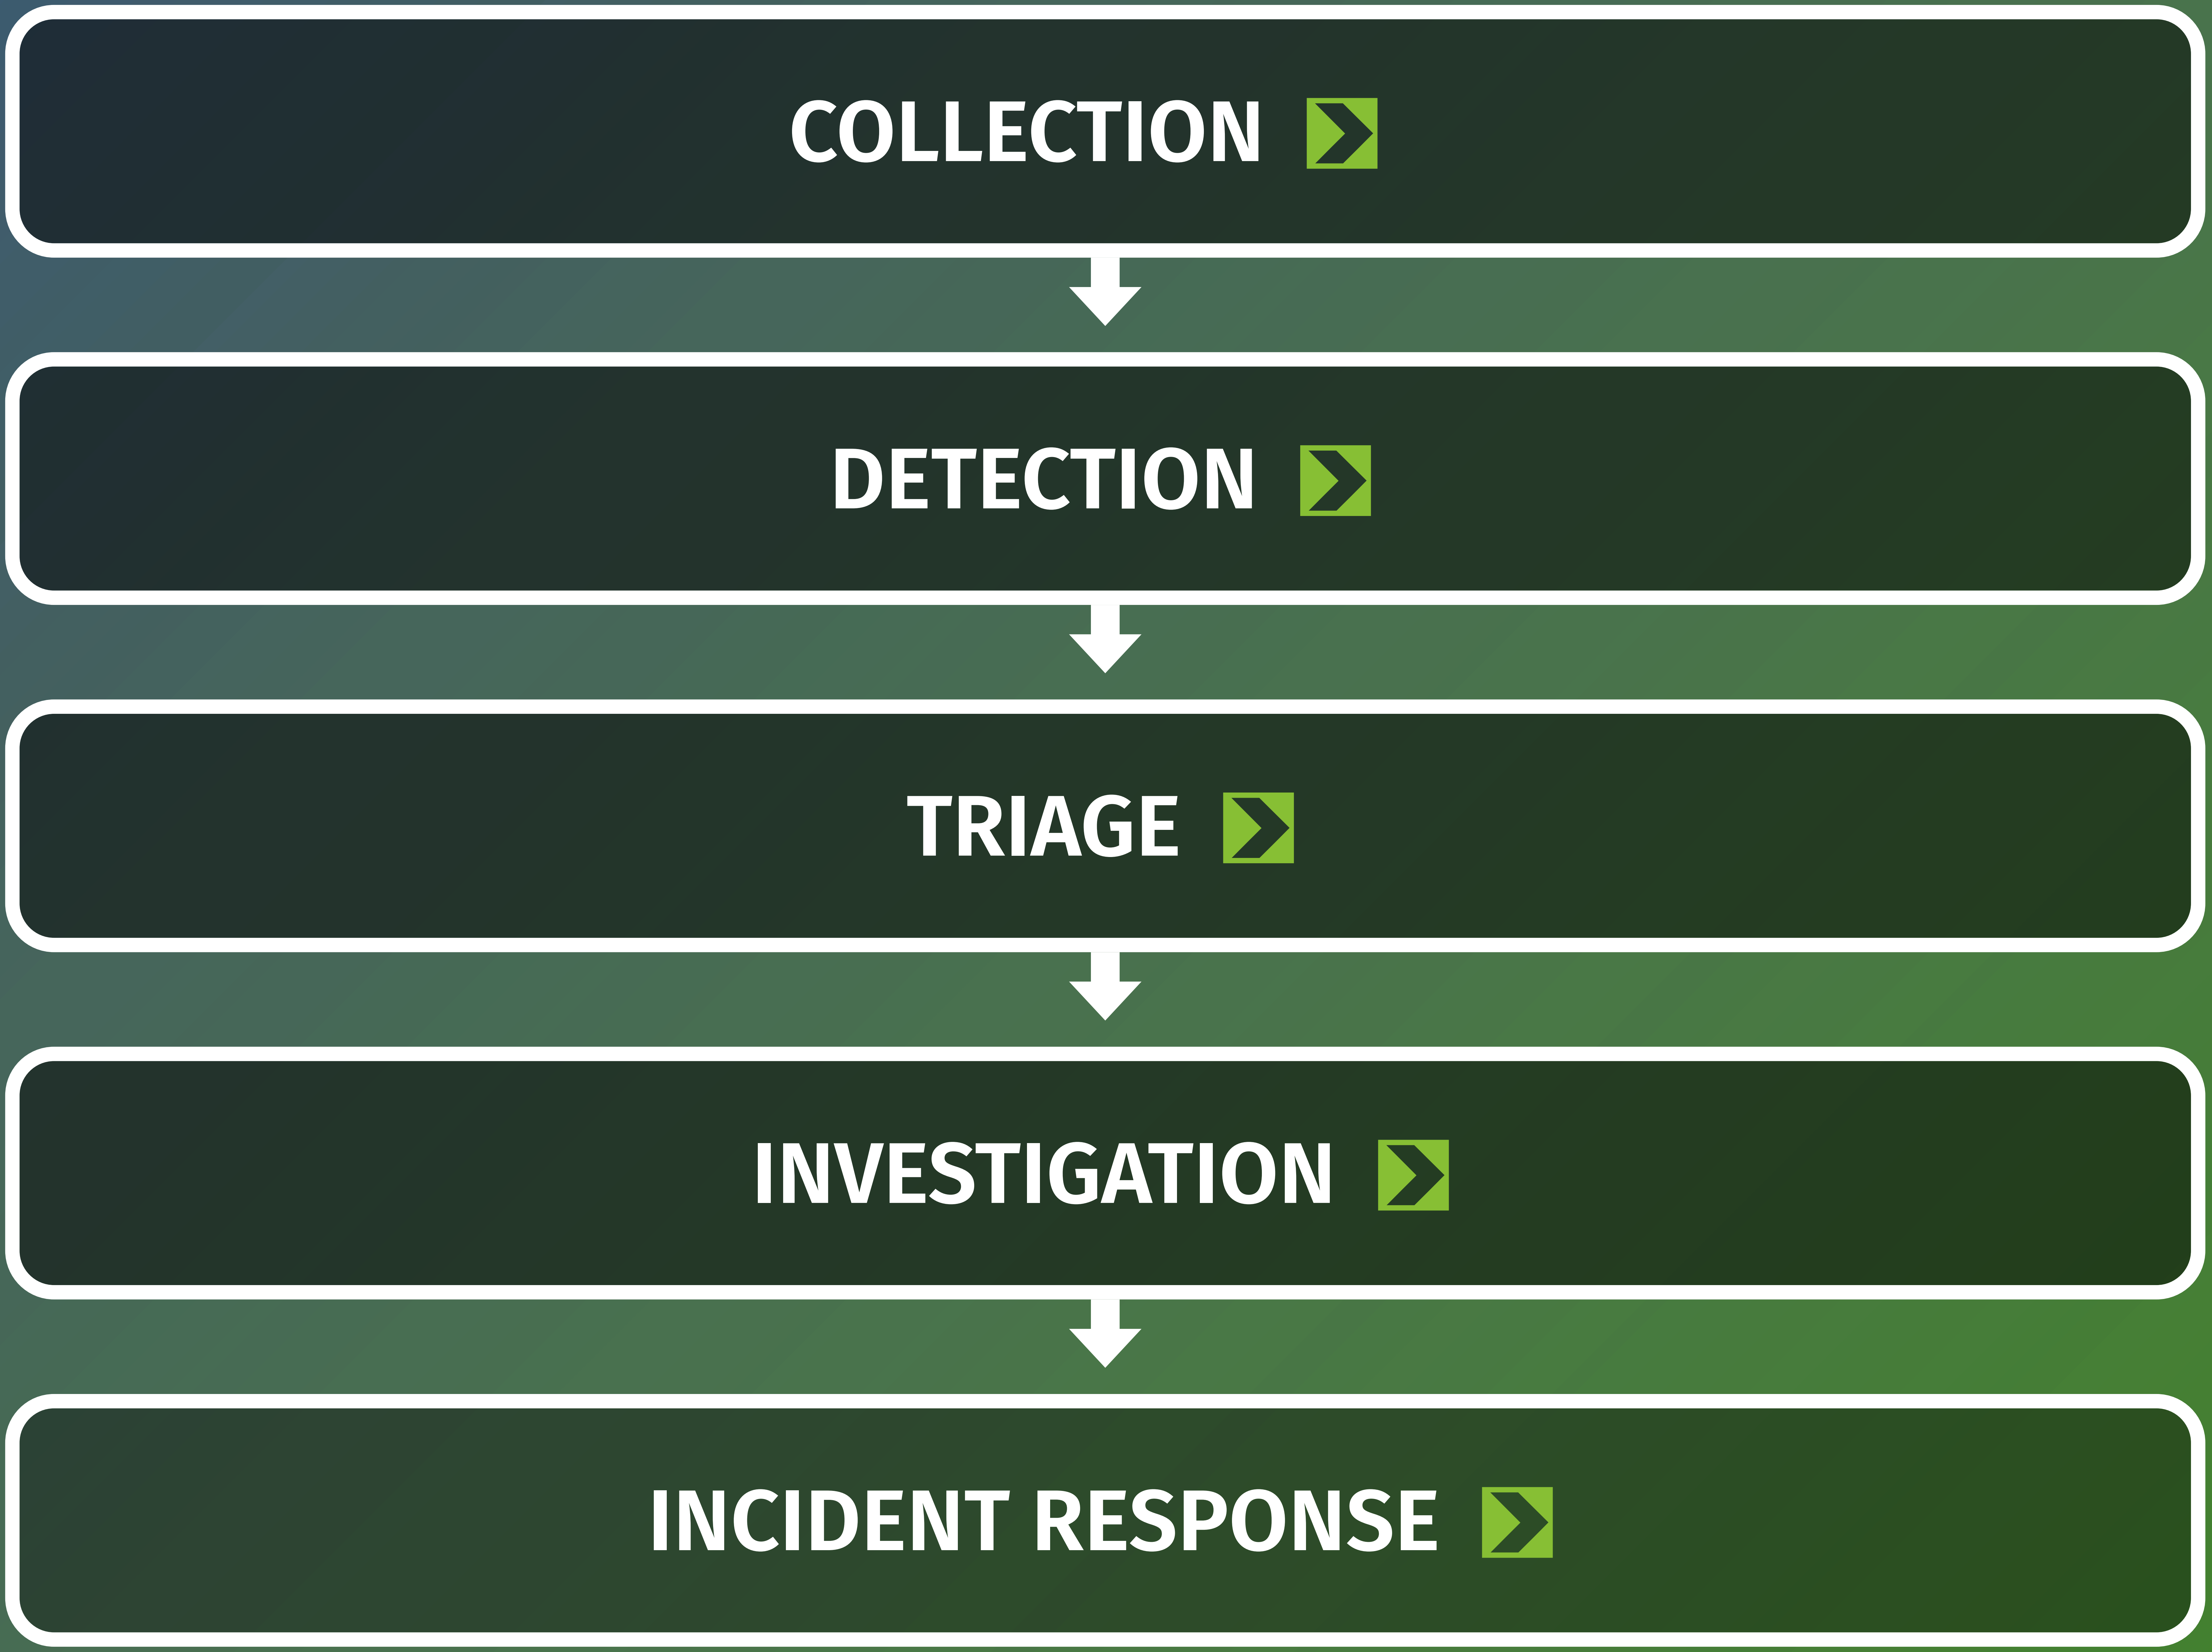

COLLECTION

DETECTION

TRIAGE

INVESTIGATION

INCIDENT RESPONSE

Table Of Contents ›

Collection

The first step in the process is collection. This step involves recording security-relevant events (any useful and observable but not necessarily malicious activity) in the environment. Recording all events such as web traffic, logins and more is required for spotting anomalous activity that may be used to identify attacks in progress. Specific data types you collect should be guided by your threat intelligence, which should inform you of what types of events and logs are required to detect attacks. In most SOCs, a thorough collection strategy will require cooperation across multiple teams and budget for the hardware and software needed to meet visibility requirements.

The output of this stage is events that may be logs, network traffic metadata, or other derived information about what occurred on a given device or network segment. This data is typically generated either on an endpoint device as a log or gathered/generated from network traffic that is observed directly by a network appliance or via network tap or switch mirror port. Collected data is ideally centralized and sent to a SIEM for correlation, application of analytic rules, and long-term storage. The goal of this stage is thorough collection of all security-relevant data that can then be used in the next stage for detection analytics. The effectiveness of a SOC at this stage will be directly related to its capability to identify security use cases and the key data sources that would indicate an attack is under way.

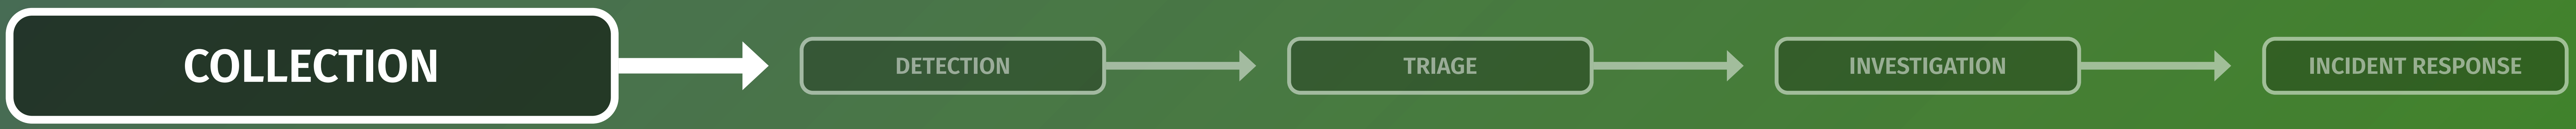

COLLECTION → DETECTION → TRIAGE → INVESTIGATION → INCIDENT RESPONSE

Table Of Contents ›

Detection

After the collection stage comes the detection stage. In this function the goal is to identify, as accurately as possible (without missing anything or generating false positives), all observed events that may indicate a potential attack. This happens two main ways, reactively and proactively. Reactive and automatic detections are applied by analytics engines in the SIEM, network, or endpoint sensors. Detections are also made by analysts proactively searching through the data via threat-hunting (see the threat-hunting section later in this guide for additional detail). The goal of this stage is to find all truly malicious activity and get an alert related to it into the triage queue for action by the security team.

The effectiveness of a SOC in this stage is correlated with the quality of the tools employed as well as the strength of the SOC's threat-hunting capabilities, threat intelligence information (both feeds and internally generated intelligence), and detection engineering functions. Successful detection naturally relies on the data being available from the collection stage. The better your collection function is operating, the better your detection function can work, and the same is true of the quality of signatures produced by threat intel and analytic engineers.

Triage

Once the detection stage has generated alerts on the events of interest, these alerts are all forwarded to one or more queues for triage. In this stage, SOC analysts must sort through all the potentially malicious activity that has been generated by the detection function and determine the order of importance in which to assess each alert.

Triage order is often based on factors such as how far the attack may have already progressed, the criticality of the system being attacked, the privilege of the account that may be compromised, and/or whether it appears to be a unique or targeted attack. Similar to a hospital emergency room, the analyst's goal in this stage is to correctly queue up items to be investigated in a logical priority order given the data presented. Effective analysts do this by combining their knowledge of concepts such as the Lockheed Martin Cyber Kill Chain and attacker TTPs like those in the MITRE ATT&CK framework tactics techniques – with their previous defense experience. Since alerts are often complex to interpret and hard to clearly understand in their raw form, effectiveness in this stage is largely driven by the level of detail and additional context provided by security tools to analysts as they are triaging alerts.

Table Of Contents ›

Investigation

Once an alert is selected from the triage queue, SOC analysts investigate the alert in more detail to verify if something bad is truly going on. As many SOCs suffer from overly sensitive analytics and untuned alert logic, this can be a step that often leads to a false positive determination and dismissal of the alert.

To get to the truth of the matter, analysts must take the provided alert data and use their analysis skills and any additional evidence they can find to assess whether the alert has correctly identified an attack. This may involve gathering data from additional network sensors or logs from multiple sources, or performing open-source intelligence research. The goal of this stage is the accurate verification of whether an alert is a true or false positive. When false positives do occur, it is wise for analysts to immediately identify the reason for the error and feed that information back to the detection engineering team to correct the alerting rules.

Analysts should be trained to perform the investigation stage in a structured and rigorous way, avoiding cognitive bias and other typical errors that can occur when intuitively performing an investigation. This can be one of the most difficult stages for new analysts, so eliminating known common investigation errors and pitfalls is one of the focuses of SANS SEC450: Blue Team Fundamentals – Security Operations and Analysis. Effectiveness in this stage is related to multiple factors – experience of the analyst, analysis technique, data availability, and technology and automation that helps generate and present evidence to analysts.

COLLECTION → DETECTION → TRIAGE → **INVESTIGATION** → INCIDENT RESPONSE

Table Of Contents ›

Incident Response

The incident response function (or team, where it is separate) receives notification of alerts that have been investigated and qualified as a situation that must be dealt with. The goal of this stage is to scope, contain, and eradiate the problem as quickly as possible, and ultimately recover from the incident with a minimum of damage. Depending on the severity of the issue, this activity may range anywhere from minor virus removal to months-long investigations with forensics, law enforcement involvement, and more. In the SOC, data collection and generation tools such as EDR/NDR and centralized SIEM logging help analysts and incident responders quickly query network and endpoint events and scope the problem and the extent of it, and then determine the timeline of the incident.

The outputs of an incident response function should be both the remediation of the incident and lessons learned on how to prevent that type of issue in the future. Post-incident attack details should be categorized for metrics and fed to threat intelligence for correlation with previous and future attacks in order to build profiles of threat groups. Tracking incident patterns and details over the long term helps give the SOC a tactical and strategic advantage in subsequent attacks.

COLLECTION → DETECTION → TRIAGE → INVESTIGATION → INCIDENT RESPONSE

Table Of Contents ›

SOC Auxiliary Functions

The following groups are often either part of the SOC organization, or work closely with it to help accomplish the cyber defense mission.

- **Threat Intelligence** – Threat intelligence teams must work closely with the SOC to inform them of what adversary groups exist, what they want, and which ones are most likely to be interested in your organization. Ultimately, they should help prioritize controls, defensive tools and detection strategies, and ensure that the team focuses on the right areas for defensive efforts.
- **Forensics** – Largely focused on assisting the incident response stage, forensics teams and specialists help find the ground truth during an incident using specialized knowledge of how activity on a machine may leave datable evidence. Since forensics is such a deep topic on its own, it’s often considered its own position and specialty knowledge area.
- **Penetration Testing / Red Teaming** – These groups help check the SOC’s people, process, and technology by simulating attacks. Tests may be announced or unannounced, and each type of testing can be focused on answering specific questions about the SOC's posture and ability to react to a realistic attack, how quickly incidents can be identified, and whether the team has received the training required to spot them.
- **Vulnerability Management** – Vulnerability management must work closely with the SOC, largely to make the results of scans available such that the SOC can detect when an exploit has been attempted against a system that may be vulnerable to it. Knowing an exploit attempt occurred that is likely to have worked is a key component of prioritizing triage efforts.

Table Of Contents ›

SOC Tools

In the SOC there are several systems that will receive near-constant use. The systems in this section are the tools that SOC analysts will be referring to, using, or searching for in nearly every alert they investigate or incident they cover.

SOC Tools Overview

The SOC tools work together to orchestrate the workflow of collected logs coming in, detection and investigation of suspicious events, and the tracking, working, and metrics collection for each incident. The diagram shows how each individual system described in this section works with other SOC tools and the type of data commonly exchanged between them.

Note that these are logical functions separated for separate discussion. In your environment they may be integrated into a single system – some SIEMs have incident management features built in, for example.

The analyst core toolset consists of:

- **SIEM** – The nexus of all the log data collected throughout the environment.
- **Threat Intelligence Platform** – Whether a dedicated solution or one built in to another of your other products, a threat intelligence platform should give analysts the context around any matched Indicators of Compromise (IOCs) found in the environment, and the adversaries behind their use.
- **Incident Management System** – The ticketing system analysts will use this system for doing alert triage and/or incidents, writing up reports on what occurred, and ultimately closing out and categorizing finished incident investigations.

Table Of Contents ›

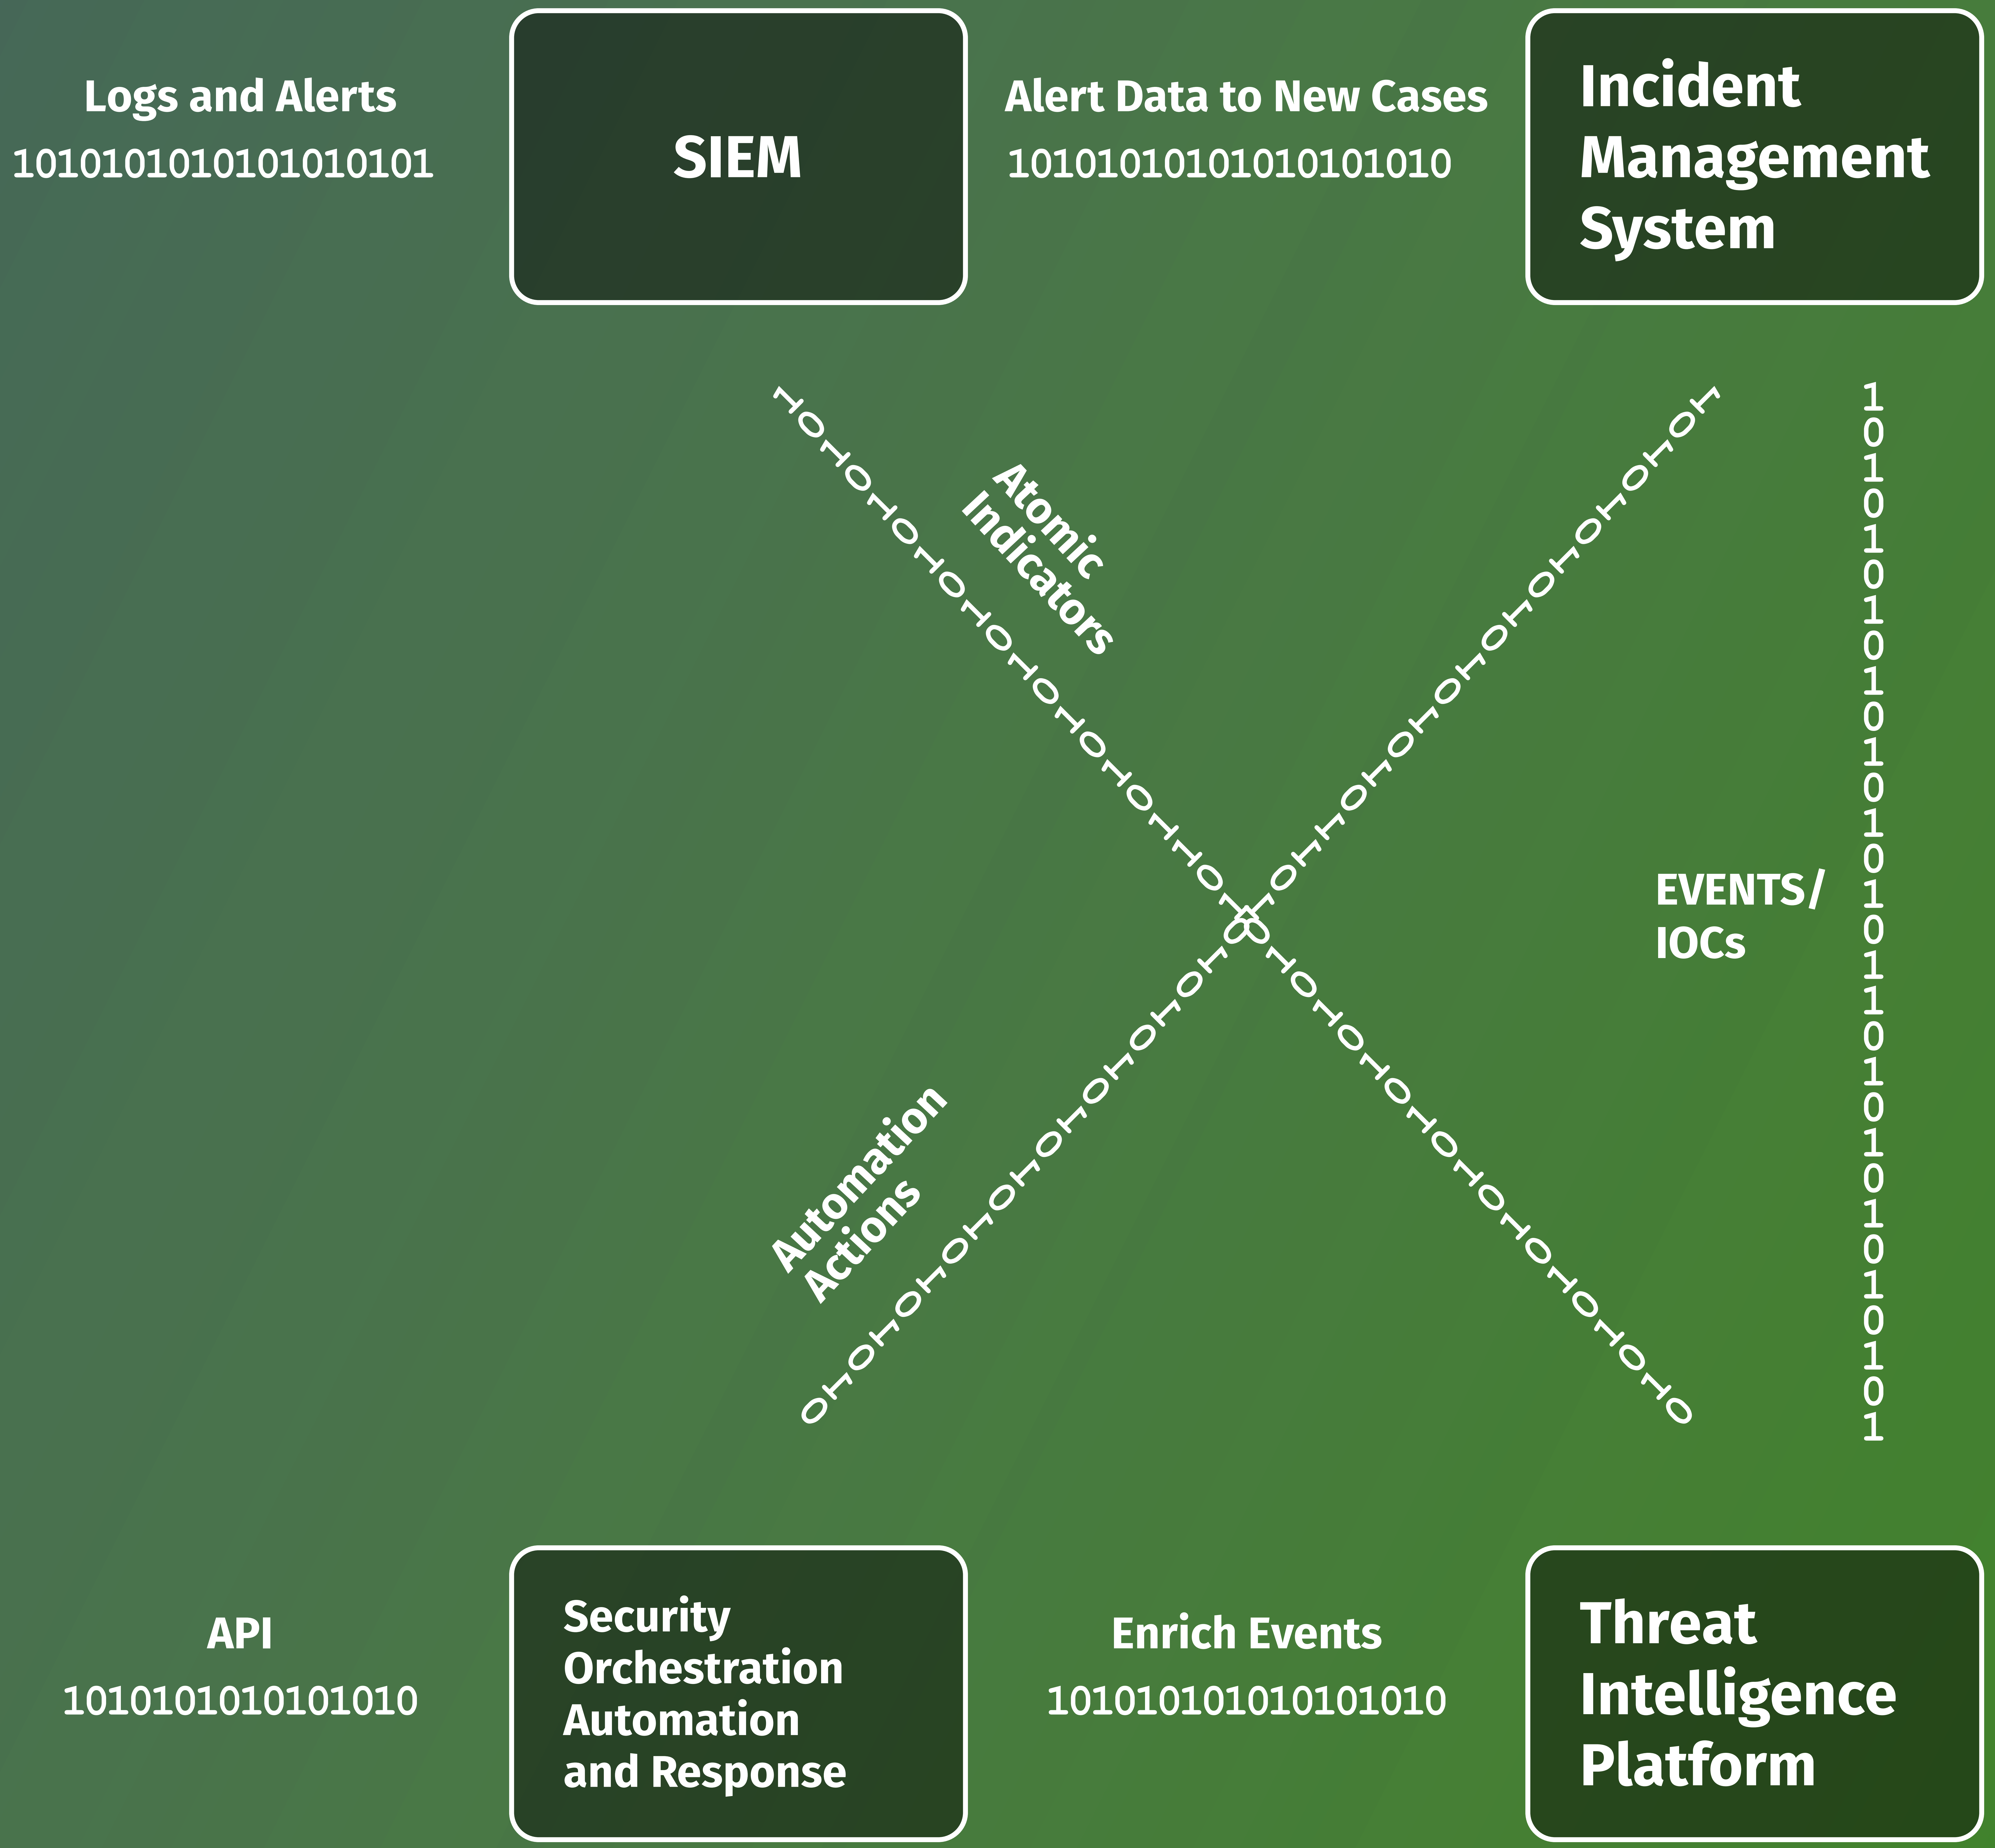

Logs and Alerts
1010101010101010101

SIEM

Alert Data to New Cases
10101010101010101010

Incident Management System

Atomic Indicators

EVENTS/ IOCs

Automation Actions

API
1010101010101010

Security Orchestration Automation and Response

Enrich Events
101010101010101010

Threat Intelligence Platform

Table Of Contents ›

SIEM

If one had to point to a single tool of particular importance for the security team to get right, it would be Security Information and Event Management (SIEM). The SIEM is in the best position to see and correlate nearly all the data from throughout the environment. It is the centralization point for all events and alerts recorded by security logging and often integrates with other tools to store context about asset information, vulnerability scan info, and more. No other single tool in the environment has this scope of data to work with, meaning the SIEM is in a uniquely powerful position in your security stack.

The SIEM's main job is to faithfully receive all logs and parse them correctly into the fields of interest, potentially enriching and correlating the information in the process. Afterward, the parsed fields are indexed into a database of some sort for quick retrieval. It is this data you can then quickly search through, alert on, or make visualizations and reports with.

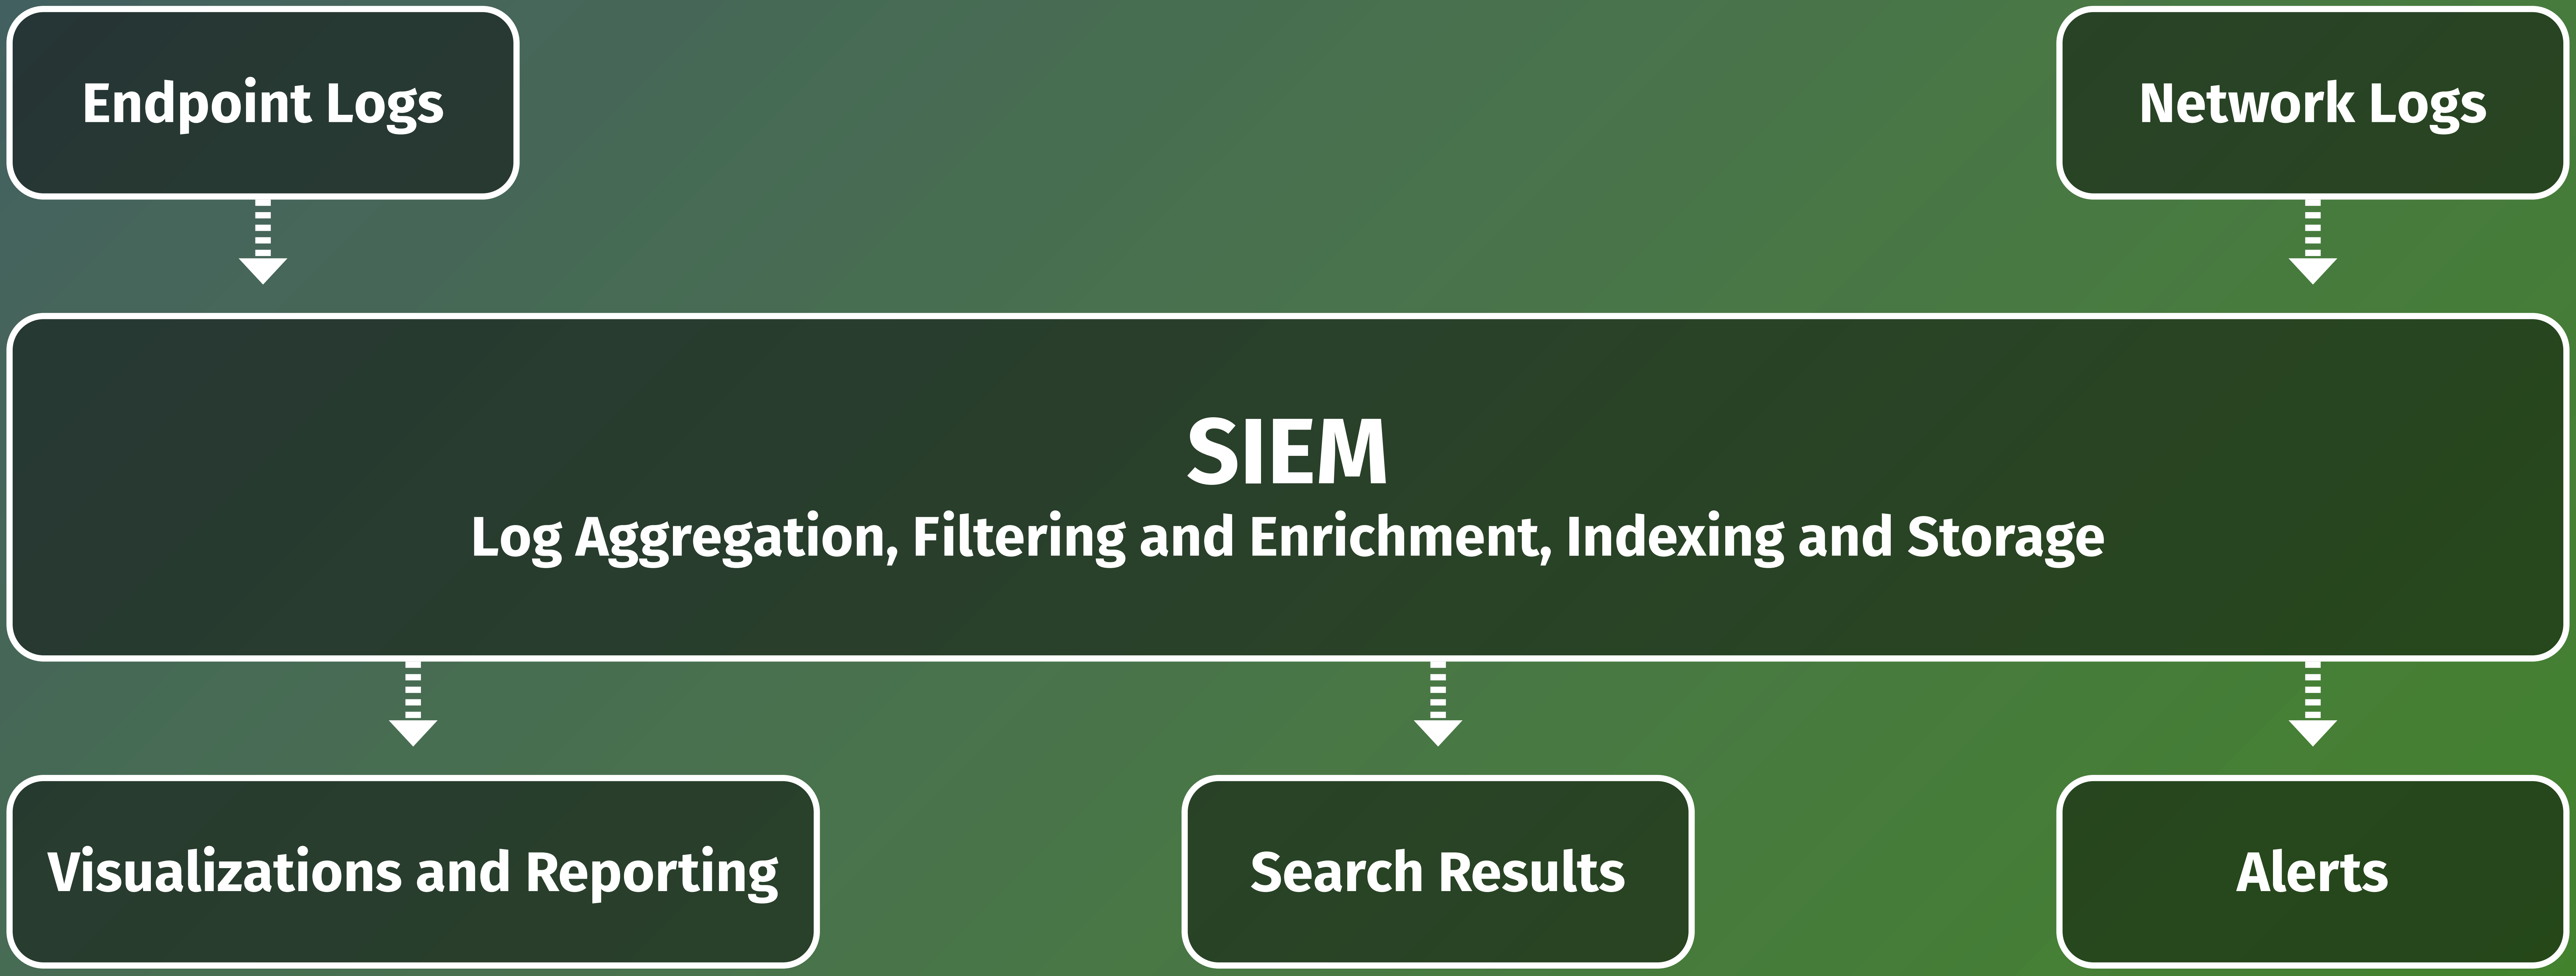

Endpoint Logs

Network Logs

SIEM

Log Aggregation, Filtering and Enrichment, Indexing and Storage

Visualizations and Reporting

Search Results

Alerts

Threat Intelligence Platform

Another key SOC tool is the Threat Intelligence Platform (TIP). In every SOC there must be a master list of all the known-bad domains, hashes, IP addresses, etc., that should be matched against all network and endpoint events. In many cases, the TIP is where this information is stored, and the TIP serves as the "source of truth" for IOCs to match against. To do so, all atomic indicators must be exported periodically to the SIEM, IDS, or any other tools being used to match the data, so that an up-to-date watch list can be applied to the events that are observed. At the same time, the threat intelligence platform will likely be taking in info from external threat intel vendors, as well as any information produced within the threat intel team or the SOC fed back from actual incidents.

A TIP must contain more information than just a list of atomic indicators. If an alert goes off that says "Threat Intel Match – Malicious IP Contacted", an analyst will need to find why that IP was marked bad. If the TIP is merely a list of atomic "bad" IPs with no context, they will have no direction to take the investigation. The TIP, therefore, ideally should contain information related to each atomic indicator that can inform an analyst how, when, and why each atomic item was marked as an IOC. If analysts can access the platform and see "This IP was related to APT1234 and was used for malware X in campaign Y on July 4th, 2020", they now have a clear next move to validate the alert and potentially correlate it with previous attacks.

Table Of Contents ›

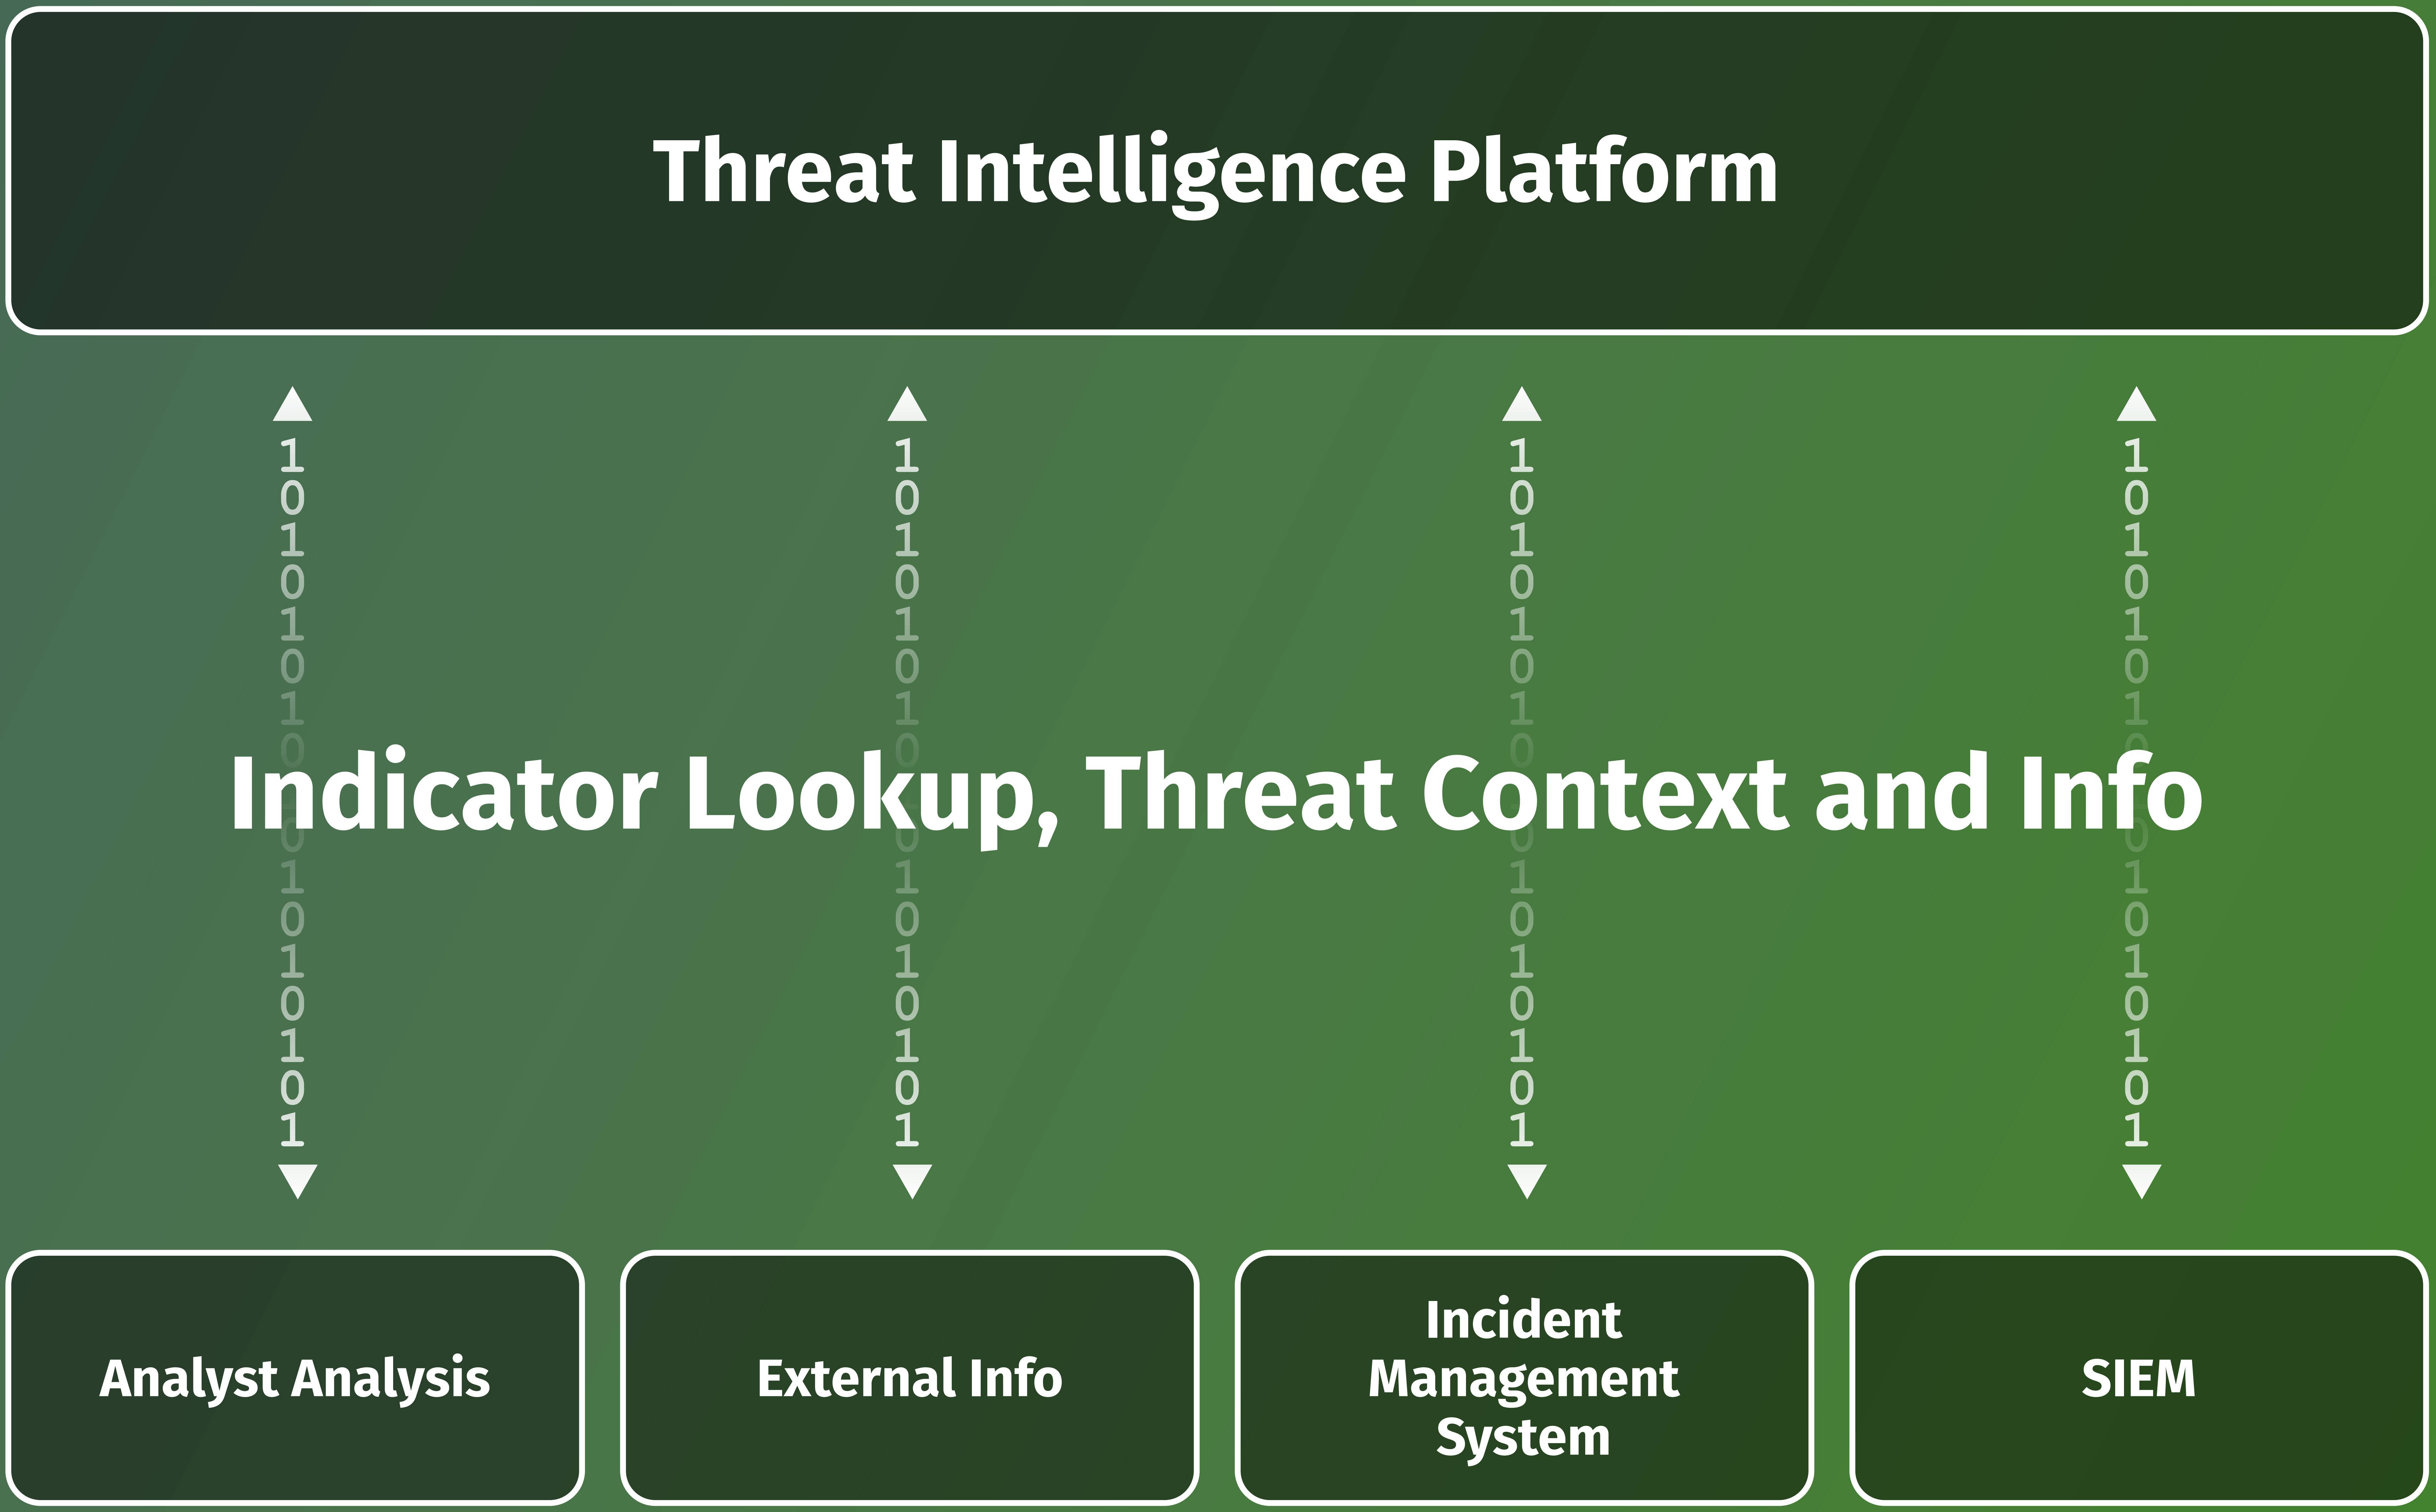

Threat Intelligence Platform

Indicator Lookup, Threat Context and Info

Analyst Analysis

External Info

Incident Management System

SIEM

Table Of Contents ›

Incident Management System

For the Information Management System (IMS), the main source of incoming data is alerts sent to it from the SIEM or other security tools (unless you triage alerts directly in your SIEM or SOAR tool, or separate security point products). Analysts use the IMS to triage alerts as well as work through active incidents. To do this as effectively as possible, integrations with your TIP and SOAR platform can help make additional information lookup and correlation quick and painless.

Once the incidents are investigated, remediated, and closed, the analysts close the associated ticket. In doing so, one key item that should not be ignored is the categorization of the incident for metrics purposes. These metrics can then be ideally automatically pulled at the end of each week and SOC managers can use the observed trends to provide feedback as to where additional budget can be best spent to protect the company. Since analysts often spend nearly all day in an IMS, it is another important piece of software to choose carefully.

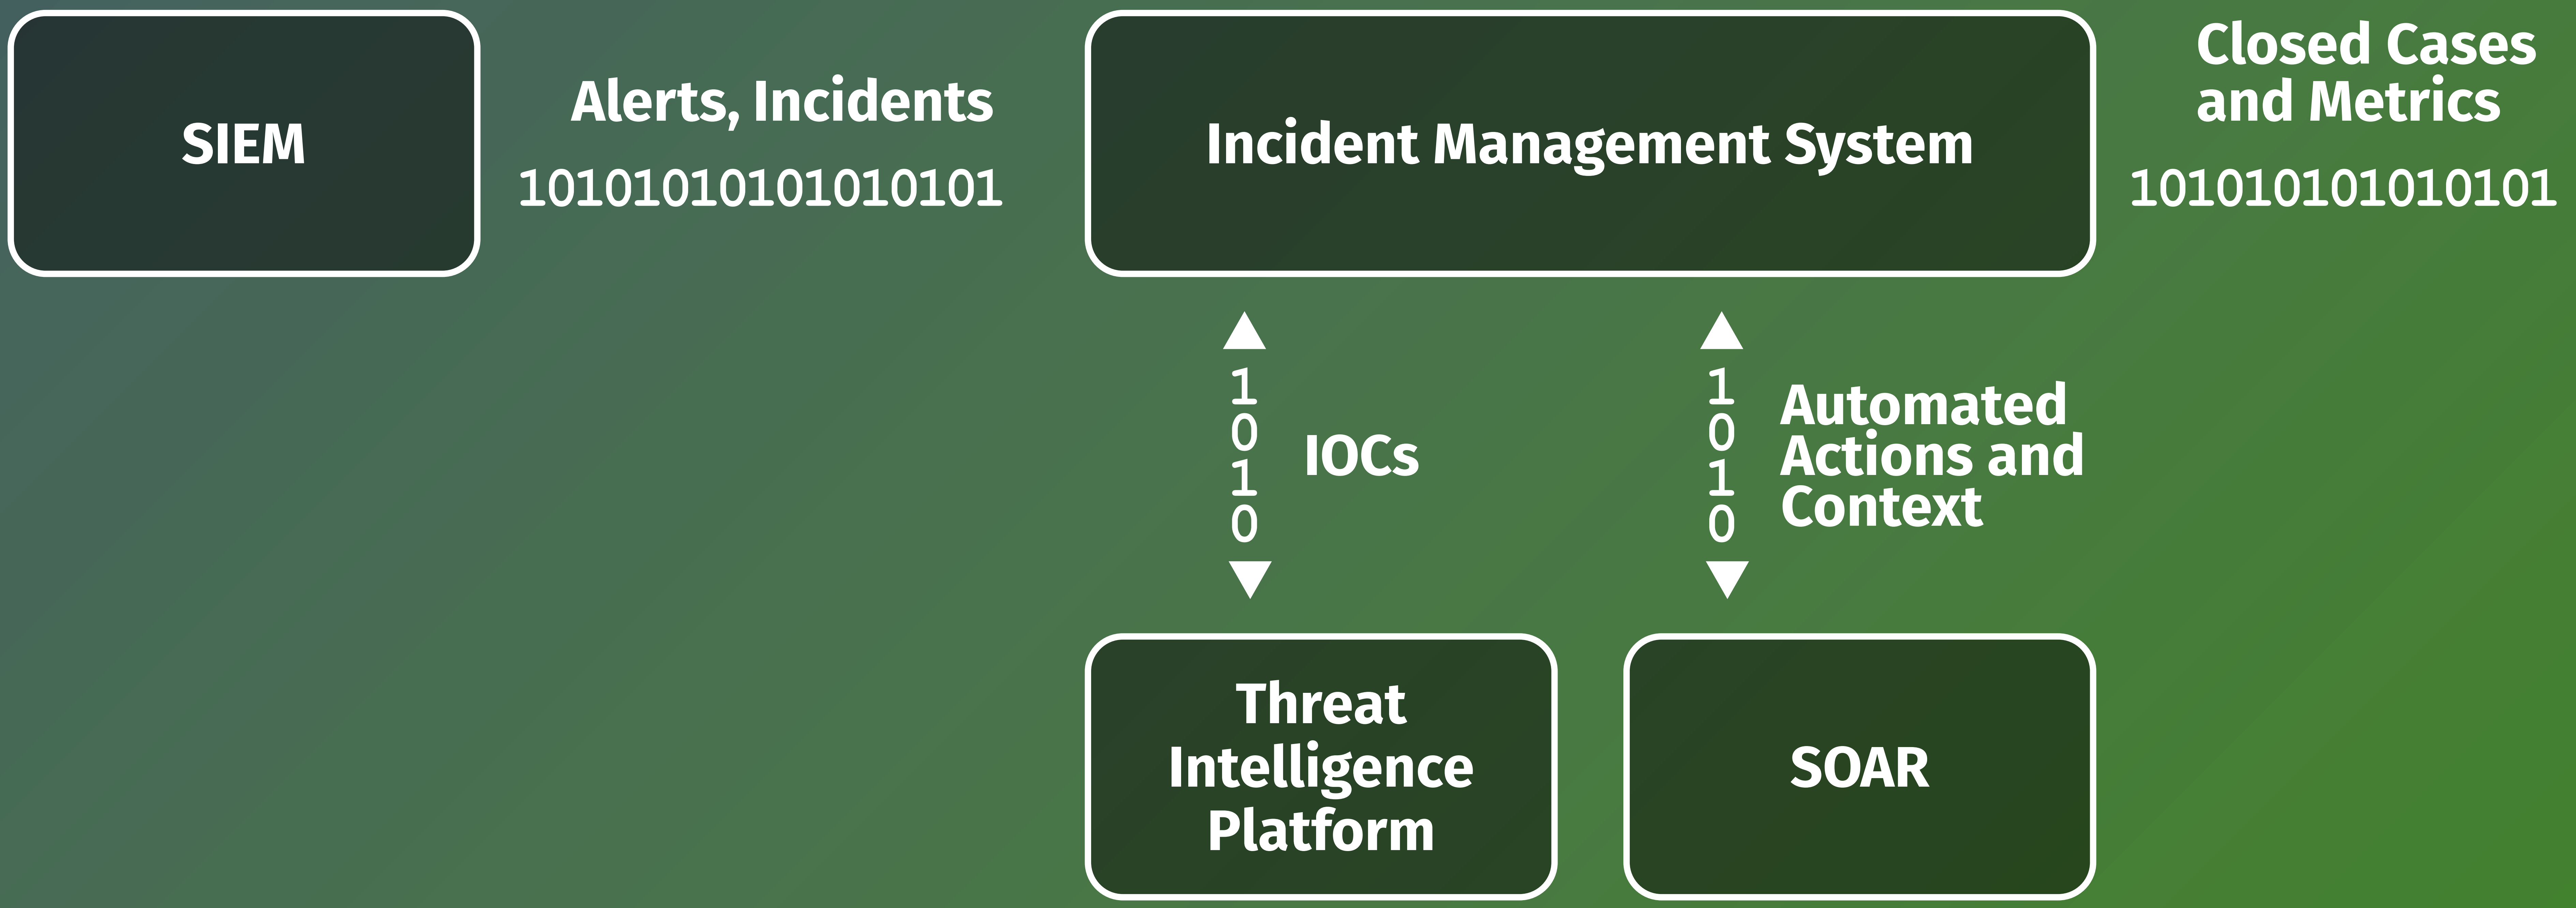

SIEM

Alerts, Incidents

1010101010101010l

Incident Management System

Closed Cases and Metrics

101010101010101

1010

IOCs

1010

Automated Actions and Context

Threat Intelligence Platform

SOAR

Table Of Contents ›

Collection of Key Data

High-performance cyber defense is by no means an easy task. Collecting and applying a constantly shifting set of detections to mountains of streaming data requires threat intelligence, a capable technology stack, and a skilled team behind it all. At the heart of any SOC is a massive and complex data collection operation. Since teams need data on both network and host-level events, getting the right data is one of the first challenges any team must face.

Data Types

Collected information can be broken down into two main camps – **network security monitoring** data and **endpoint monitoring** data (or continuous security monitoring as it is sometimes called). This data is collected from all points through the network and much of it is centralized to a SIEM for convenient searching, visualization, anomaly hunting and reporting. The following pages will further break down each type of data and describe the key sources for each that must be collected in order to give your security operations team the best possible chance at success for advanced attack detection.

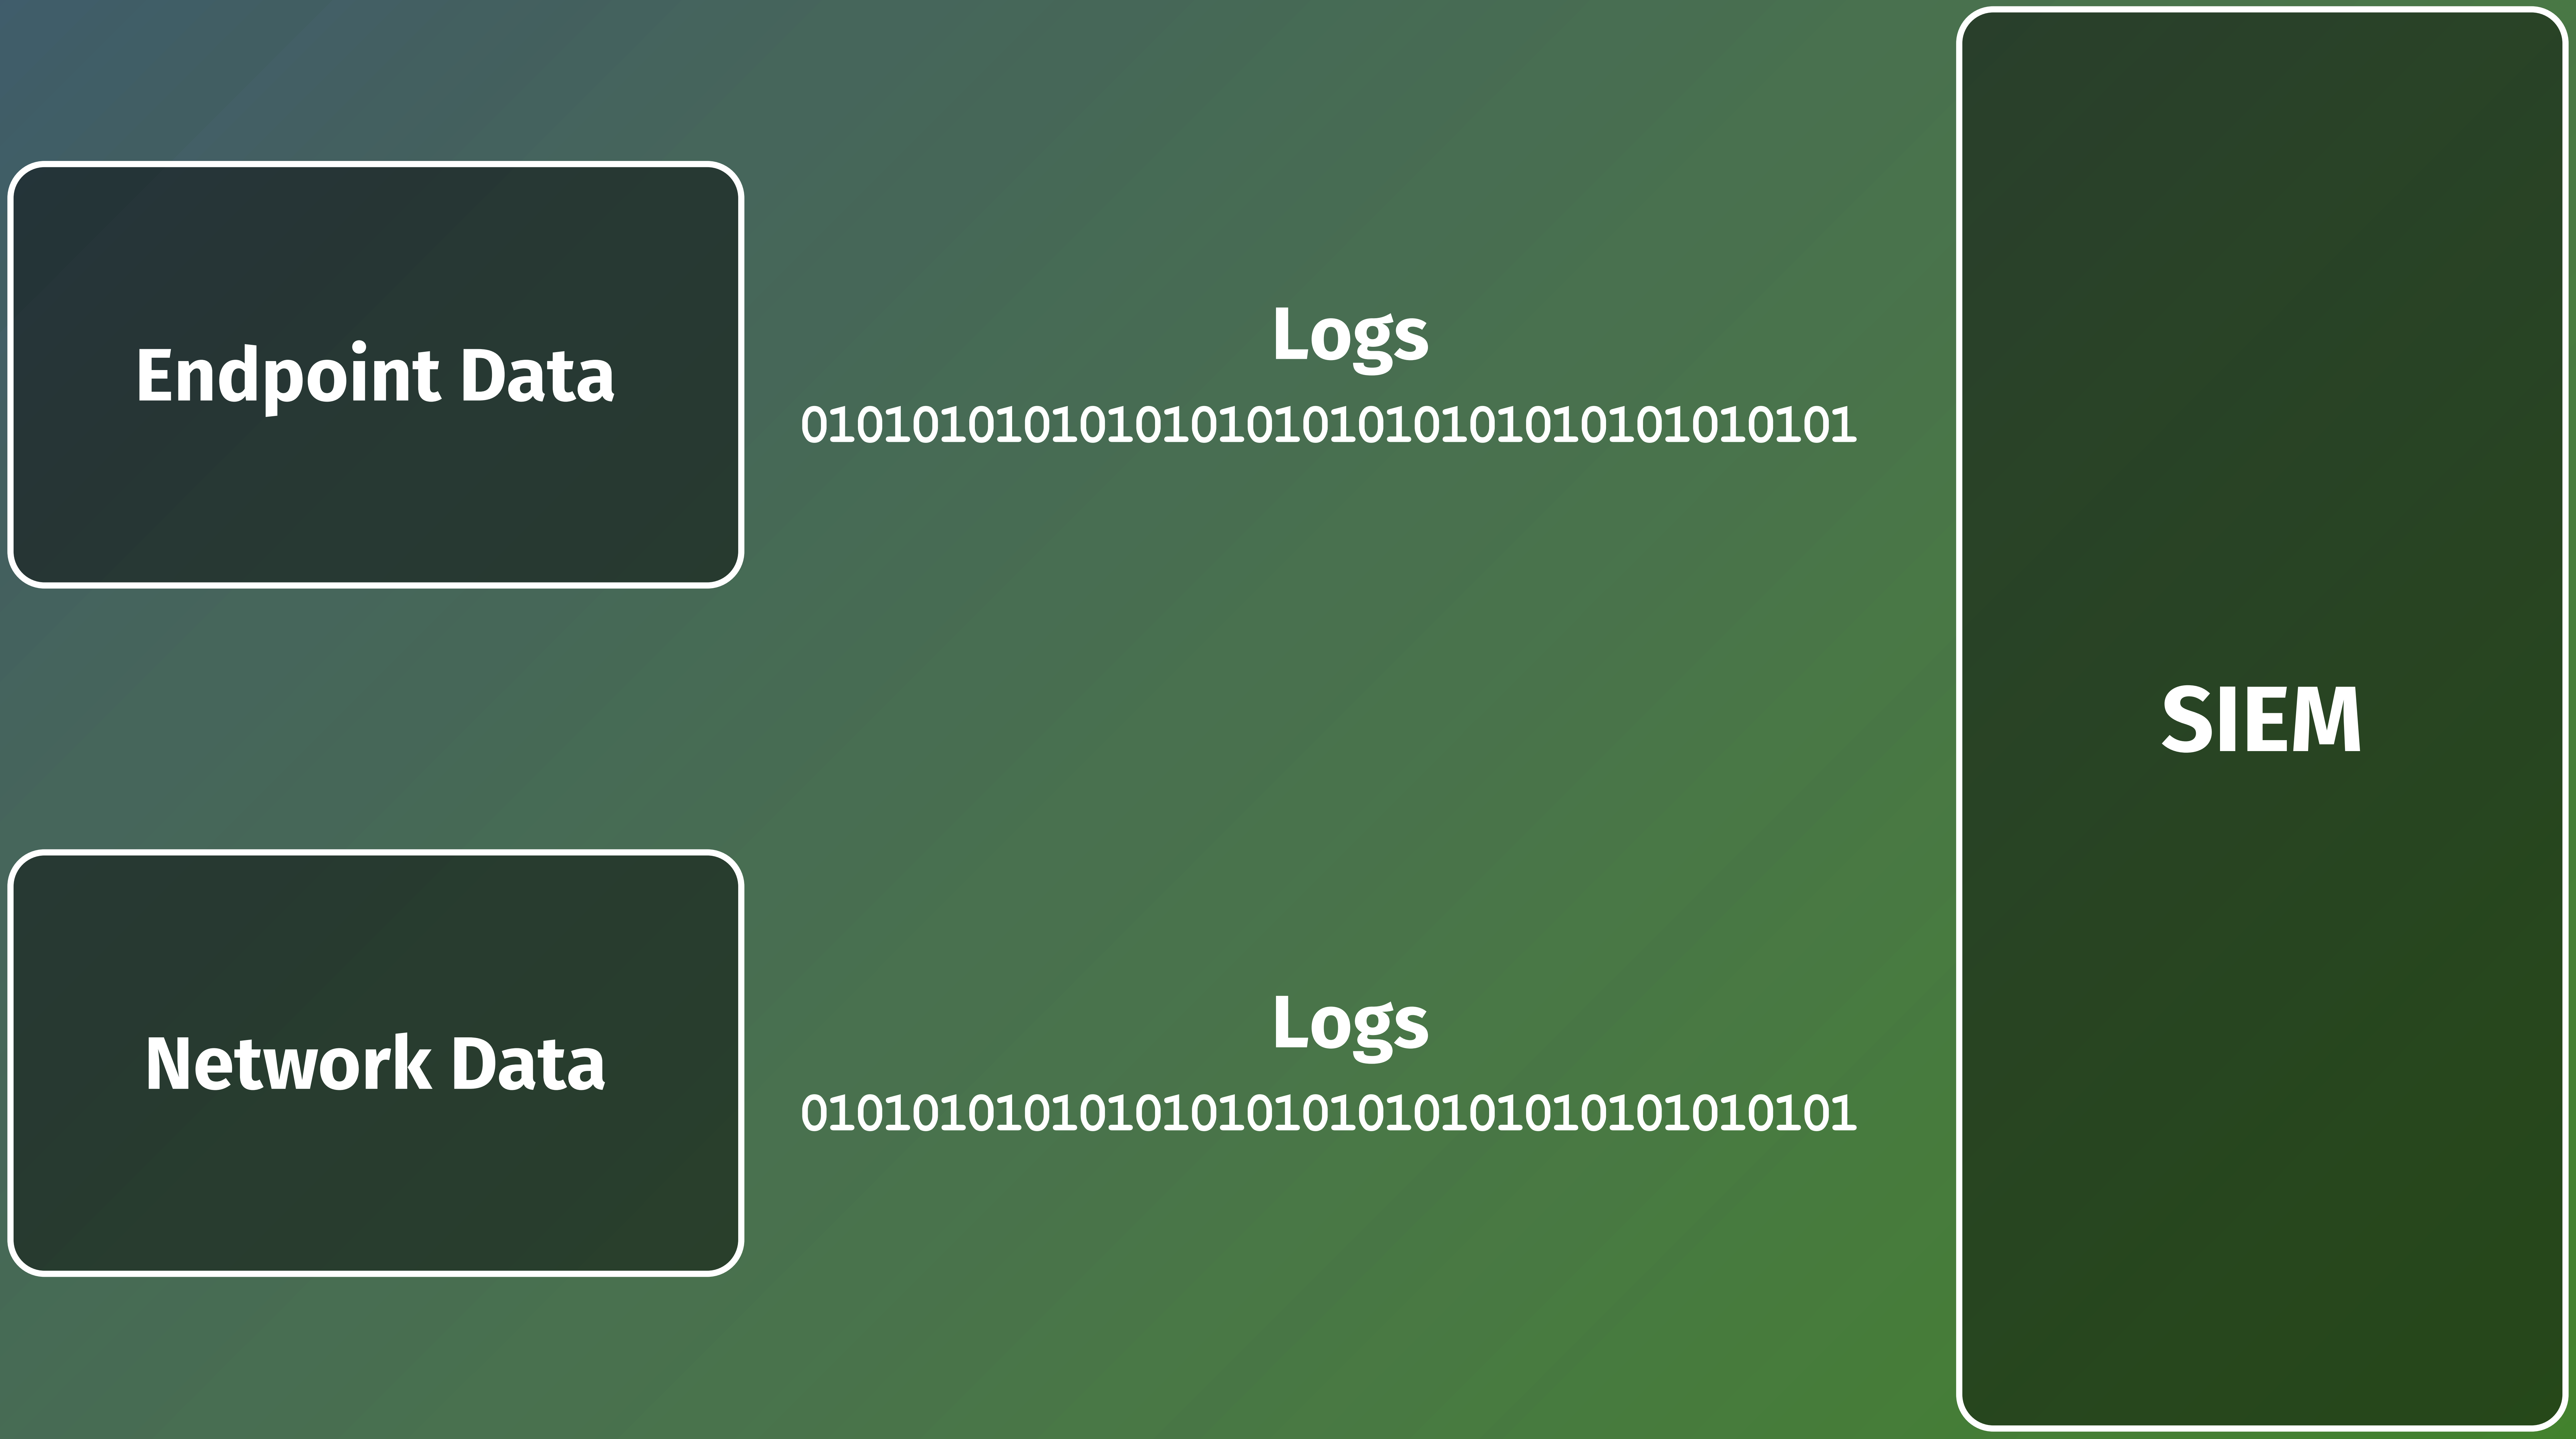

Endpoint Data

Logs

01010101010101010101010101010101010101

Network Data

Logs

01010101010101010101010101010101010101

SIEM

Table Of Contents ›

Flow Records

Flow records are the highest level of network security monitoring logs, describing mainly OSI layer 3 and 4 (TCP/IP) details as well as timing details. The pros of flow logs are that they are typically easily available to security teams because they are often used by network operations teams for performance monitoring and they take up relatively little space to store due to their limited data. The downside to flow logs is that they are often not detailed enough to truly determine whether or not there is a potential attack, unless they are being used to match IOCs such as IP addresses or to locate traffic anomalies.

Common tools and formats: Text-based logs produced by tools like Zeek or Suricata or proprietary formats such as NetFlow, JFlow, NetStream, Zeek conn logs, sFlow, etc.

Useful for profiling:

- Traffic volume / bandwidth
- Start, stop, and length of connections
- Conversation source and destination IPs
- Source and destination ports
- Protocols (under the assumption that ports are used with typical services – not a guarantee)

Transaction Data

Transaction data (also often referred to as network service logs) takes flow log-level data and extends it all the way to OSI layer 7, the application layer. This type of data is produced by tools that look at the full packet and do a true analysis of the protocols in use and the details of the information. This analysis provides information on TLS certificates, HTTP transactions, and more, and is much more useful for identifying attacks since many IOCs collected by a SOC for matching (such as hashes, domain names, TLS certificate details, user-agents and more) only become visible when analyzing traffic at this level of depth. SOCs should strive to collect transaction data as a minimum, since transaction data is still a text-based log that is only a bit larger than flow logs and thus can be generated and stored relatively inexpensively.

Common tools and formats: Text-based logs produced by tools such as Zeek, Suricata, numerous commercial network monitoring and NDR solutions.

Useful for profiling:

- Everything from the flow logs section to...
- Application-layer protocols in use
- Session and presentation layer details (example: TLS certificate details)
- Details of application layer conversations (example: HTTP methods, user agents, URLs, hostnames and more)
- Potential matches for known-malicious IOCs

Full Packet Capture

Sometimes even application layer metadata is not enough for making a true/false positive call during an incident investigation. In these cases, full packet capture provides every single byte sent over the wire and can be used by security teams to get to the ground truth of what happened in an investigation. There are two large common issues with full packet capture – encryption and the size of the data. Encryption can make packet capture much less useful unless TLS decryption is used to decode it before recording (if you do not have decryption capabilities, BPF-style filters can be used to not record data that will be otherwise useless). Due to the full-content nature of packet capture, the data is also significantly larger in size compared to flow and transaction logs. Therefore, retention periods for packet capture will typically be significantly shorter than for transaction and flow logs. Note that although packet capture was difficult to acquire for cloud services in the past, most platforms have now created solutions for packet capture for cloud assets.

Common tools and formats: PCAP and PCAP-NG are by far the predominant storage format for full packet capture data. Tools to record packet capture on a continuous basis are numerous both in open-source and commercially. Free options include tools such as Moloch, Google Stenographer, and netsniffing used by NSM distributions such as Security Onion. For cloud services **Amazon VPC Traffic Mirroring**, **Azure Network Watcher** and **Google GCP Packet Mirroring** can provide visibility to cloud network traffic.

Table Of Contents ›

Useful for profiling:

- Everything from flow and transaction logs to...
- Full content analysis of packets (example: HTTP GET/POST body content)
- Carving of transferred files and malware for analysis
- Detailed network transaction forensic analysis
- In-depth protocol analysis
- Advanced incident response and reverse engineering of attacks

Traffic Collection Opportunities

Since data types with higher levels of detail can be used to create the lower-detail-level data types, a single packet capture feed can be used to create full packet capture, transaction logs, and flow records. If your team does not have network taps or mirror ports available for full packet capture feeds there are still many network appliances and IDS sensors that can produce service-level and flow logs (such as next-gen firewalls or open-source tools like Suricata). Teams should strive to collect data in whatever way will be more effective and sustainable for them. Some of the options for network security monitoring data collection and how it can be converted from one form to another are shown below.

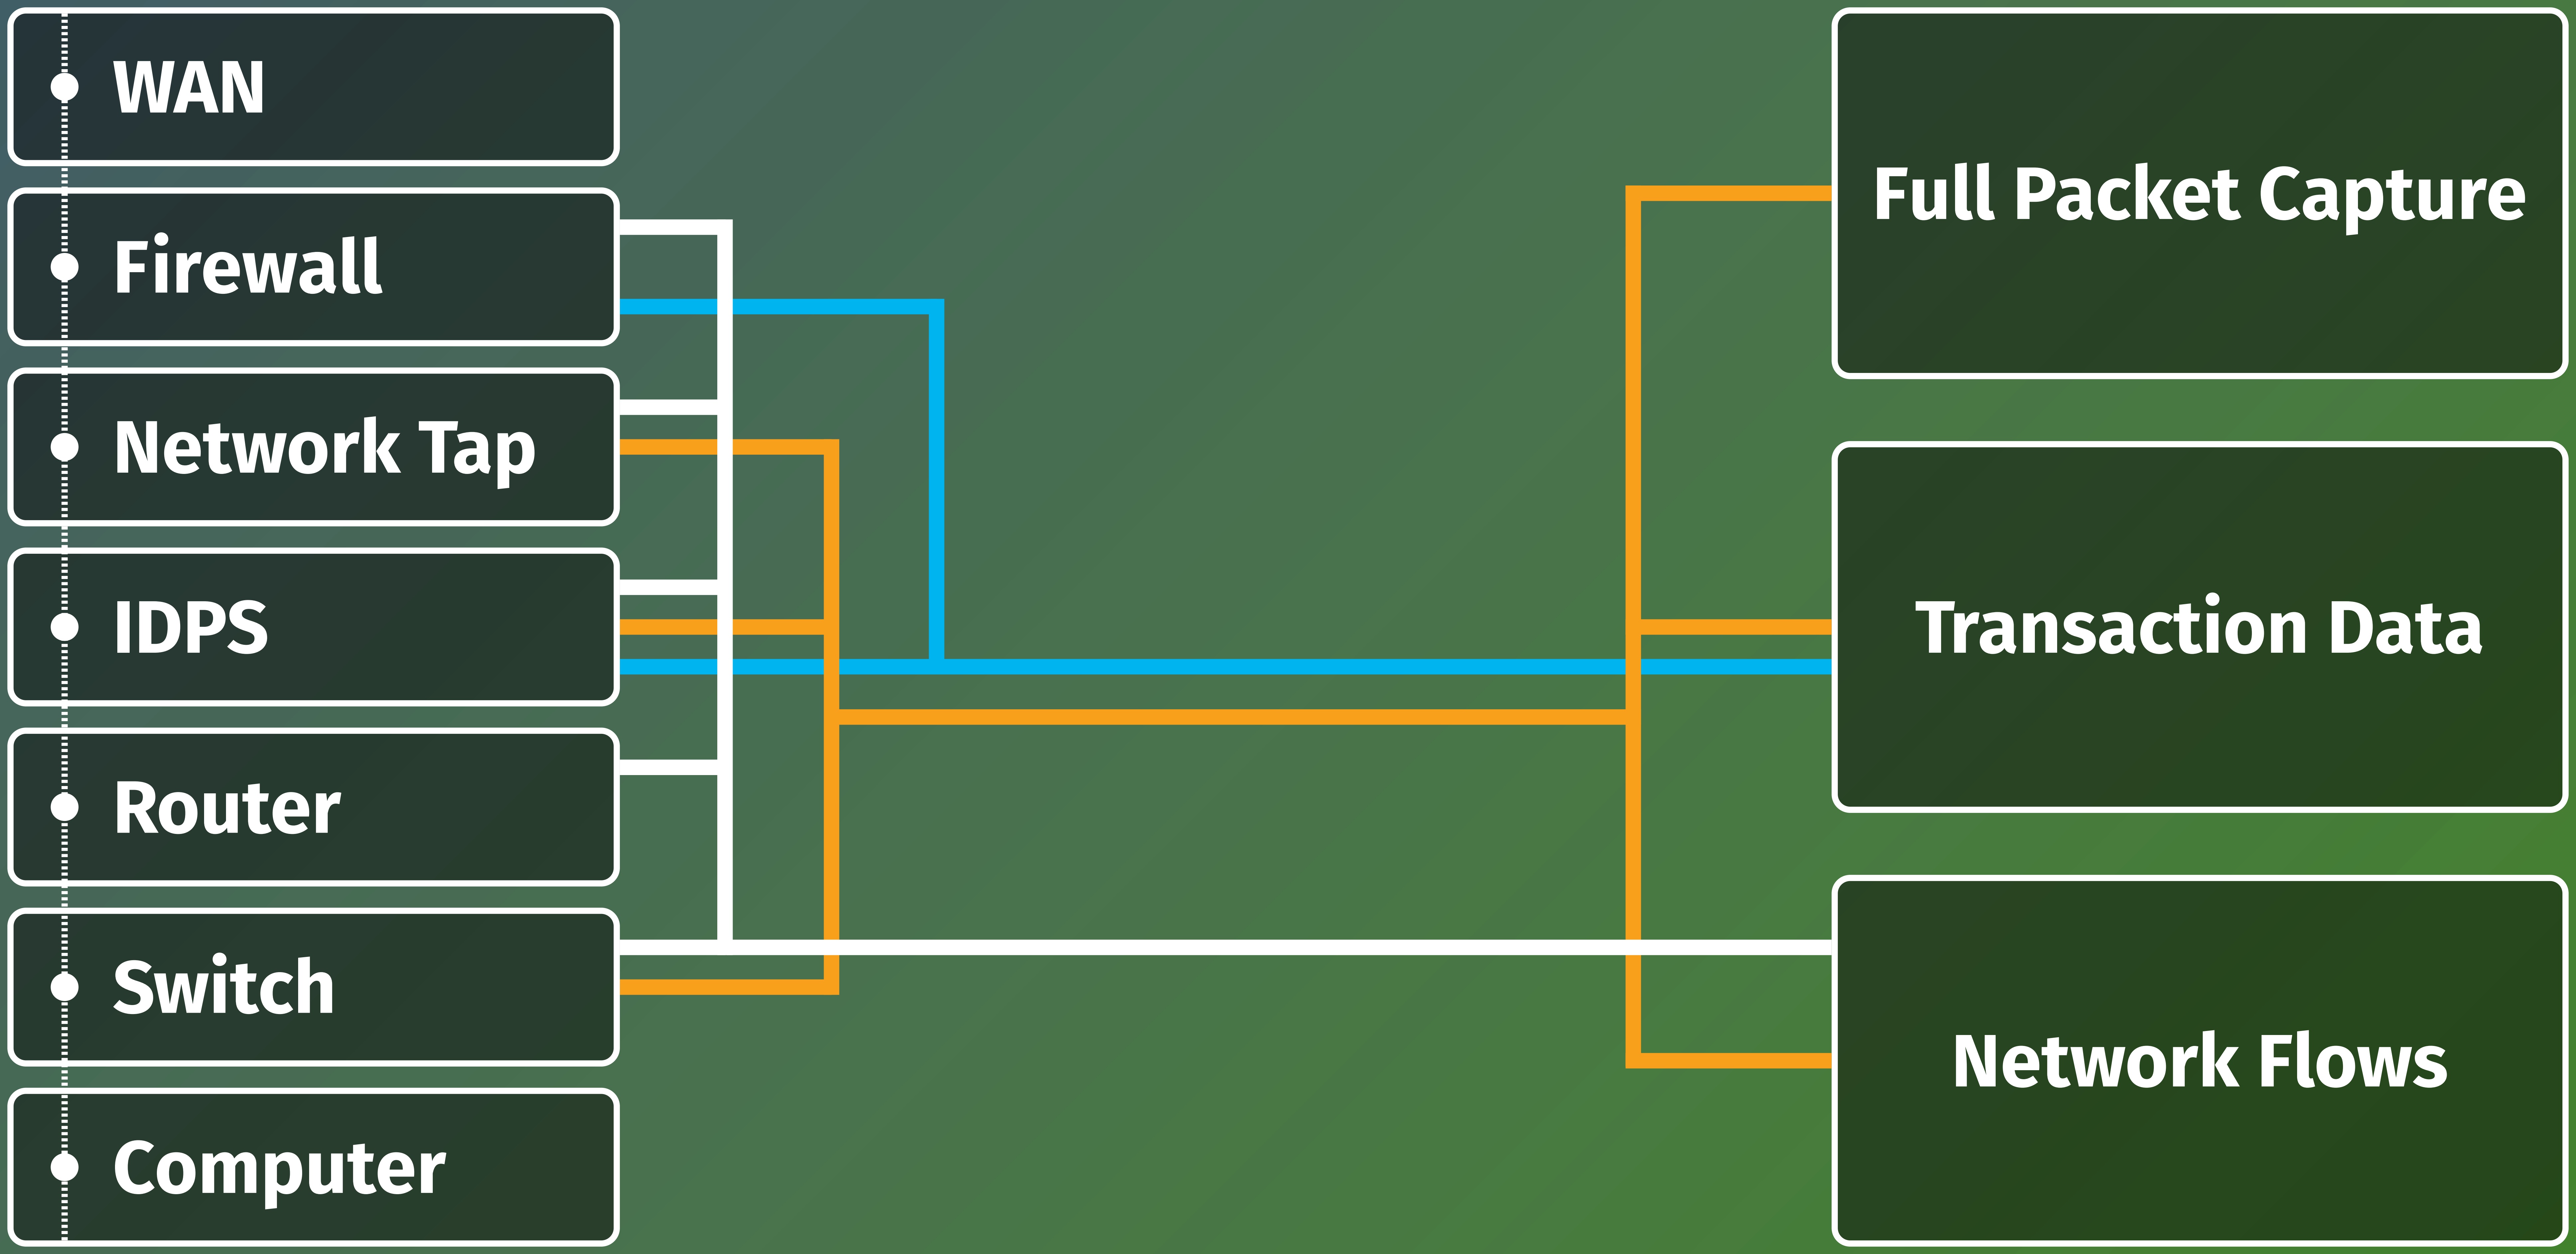

Table Of Contents ›

Endpoint Monitoring

Endpoint monitoring is the other type of monitoring data that must be collected by a SOC to ensure attacks can be identified and addressed in an accurate and timely fashion. Endpoint data collection (referred to as "continuous security monitoring" or CSM in many SANS courses) provides detail on the processes, logins, services, and other key info about what is happening on devices. This data, typically in the form of text logs, is generated by operating systems, applications, and security agents present on the device and is most often collected by log agents and centralized to a SIEM.

Logs to SIEM

Event Categories to Record and Collect

Program Execution

- Command line and arguments
- Program name and file path
- Parent process
- Hash
- Signature – signed or unsigned, signature details
- Application control list status (allowed or unallowed)

Script Execution

- Windows
  - cscript.exe, wscript.exe, cmd.exe
  - PowerShell – script file names, transcription logs, script block logs, module logs
- macOS - osacript (AppleScript)
- Linux – sh, zsh, bash commands etc.
- Cross-platform languages such as Python

Authentication Logs

- Which accounts are successfully logging in? Failing to log in?
- What type of login is being used? (local login, RDP, SSH, etc.)
- Authentication mechanisms – Kerberos vs NTLM, etc.
- If remote, where is the login originating from?

Table Of Contents ›

Configuration and Baseline Monitoring

- Running services
- Autorun items
- Scheduled tasks
- Sensitive file and registry additions, deletions, reads, and modifications

Vulnerability Status

- Operating system – version and installed patches
- Applications – What is installed and which version?

Vulnerability Status

Note: Host firewall logs are an incredible source of visibility, but often must be heavily filtered before collection – selecting only ports commonly used for lateral movement and sources/destinations where traffic would be unexpected

- Incoming/outgoing traffic to commonly exploited or unusual ports (SMB, RDP, SSH, PowerShell Remoting, VNC, WMI, FTP, etc.)
- Incoming/outgoing traffic to/from unexpected locations (looking for lateral movement such as user device to user device connections)
- Listening ports and the programs using them
- Outbound/inbound traffic by process (a key differentiator for network firewall logs)

Windows Log Sources

Windows log sources are broken up into channels in the Windows Event Viewer. While your Windows audit policy will determine which events create a record in one of the channels, it is still up to the security team to select which channels, and which events within each channel to collect. Note that while the most standard configuration is to start by collecting the Security, System, and Application channels, this is by no means enough to spot many advanced attacks. To better enable your team to catch attacks, the following Windows log channels and beyond (where applicable) should also be considered for collection. Note that many of these channels will require a detailed event filtering policy implemented by your log collection agent based on EventID and/or other fields in order to collect only security-relevant events.

Windows Log Channels of Interest

- AppLocker
- DeviceGuard
- EMET
- PowerShell
- Security-Mitigations/Kernel-Mode (Windows Exploit Guard)
- Sysmon
- Windows Defender
- Windows Firewall with Advanced Security
- WMI-Activity

Table Of Contents ›

References for Specific Windows Event IDs

For specifics on event IDs of interest see the following resources:

- NSA Spotting the Adversary with Windows Event Logging
- Sean Metcalf's AD Security Blog: adsecurity.org/?p=3299, adsecurity.org/?p=3377
- Malware Archaeology Cheat Sheets
- Ultimate Windows Security Log Events Encyclopedia

Table Of Contents ›

Linux/Unix Log Sources

Unix-based system logs are typically recorded either in text files as shown below, or in the system journal for distributions using the systemd journal. The following items are starting point suggestions for building a Linux logging strategy.

- `/var/log/auth.log` or `/var/log/secure`
- `/var/log/syslog` or `/var/log/messages`
- `/var/log/audit/kern.log` – Kernel logs (highly verbose, potentially limited value)
- `/var/log/audit/audit.log` – Auditd logs
- `/var/log/audit/ufw.log` – Firewall logs
- `/var/log/apache2` (or httpd) `/access.log` – Apache logs
- `/var/log/httpd/mysqld.log` – MySQL logs

Cloud Logging

Logging and monitoring for cloud-based systems can be done in a number of different ways depending upon how the cloud systems are managed. The main two areas are watching the administration of the cloud platform itself, and monitoring the servers, platforms, or applications being run in the cloud.

IaaS Logging

For systems that can be considered IaaS, logging and monitoring comes down to using the same mechanisms you would use as if you were running that system on-premise. Since you are in almost complete control of the system, logging can be implemented like any other virtual machine, often using log agents that collect information on system, OS, and application activity, and then sent to a centralized SIEM.

Cloud Management Plane Logs

One major area of concern for all cloud services are the administration interfaces themselves. Each organization has individuals who have permissions, for example, to log into the Amazon AWS console and read, modify, create, or destroy services and data. An attacker that gains access to these administrative functions can easily cause damage in many different ways. Therefore, cloud system monitoring must always include auditing of who is logging in to the administration platform, and which actions they are taking. Cloud management plane logging sources include items such as Amazon CloudTrail and Azure Activity Logs.

Table Of Contents ›

PaaS and SaaS Logging

Even though your organization does not control the underlying system in a platform or software as a service setup, concerns about their security should be just as high-priority as any other system. PaaS and SaaS vendors typically enable log collection via API calls or service logs written to a special area within that vendor's cloud platform (such as to an S3 bucket, Amazon CloudWatch, Azure Event Hub, or similar). These APIs and log collection points must then be set up in your SIEM as a log source that is automatically polled. The logs gathered should include activity for service as well as the users interacting with the software, and should audit the actions that were attempted/taken (changes, login/logoff, actions, etc.).

The difference between Saas and IaaS/PaaS is that the security of these systems is out of your control, leaving these logs as the only option for monitoring, and only including what your vendor is willing to share with you. Organizations utilizing SaaS solutions should carefully consider use cases for attacks on these systems and potential attacker goals, and map out how these events would look given the visibility available to you from the vendor API. Analytics and alerting rules around these use cases must then be created, tested, and continuously verified to ensure complete coverage.

Table Of Contents ›

Models & Metrics for Security Operations

This section contains some of the most important references for understanding and modeling cyber attacks, building threat intelligence, and executing and measuring your security operations.

Table Of Contents ›

Attack Mental Models and Reference Frameworks

The following is a list of some of the most useful mental models commonly referenced in InfoSec as well as the incredibly useful frameworks that are being developed to standardize and organize key cyber attack and defense techniques, threat intelligence and data.

- The Lockheed Martin Cyber Kill Chain
- Incident Response Cycle (NIST SP800-61r2)
- David Bianco's Pyramid of Pain
- The Diamond Model of Intrusion Analysis
- MITRE ATT&CK – list of attacker tactics, techniques, procedures, tools, threat groups, mitigation and detection options, and much more!
  - ATT&CK Navigator – Supplementary visualization tool
- MITRE Shield – Tactics, techniques, and a knowledgebase for active defense
- ATC RE&CT – A framework, collection and data source for incident response techniques
  - RE&CT Navigator – Supplementary visualization tool

Table Of Contents ›

Knowing Yourself and Your Enemy

Defining your enemy in various capacities is key for aligning defensive budget against the most likely attacks. Your enemy's goals, capabilities, tactics, techniques, procedures, and specific IOCs all should be extracted from incidents and external intel, and tracked as closely as possible. This will help produce a well-defined picture of who and what your team will be up against. The better this knowledge can be collected and organized, the more effective the team will be, and the more efficiently the defense budget can be allocated. This section presents some of the important models and concepts a blue team should be familiar with in order to make the most of threat intelligence.

Table Of Contents ›

Three Levels of Threat Intelligence

Threat intelligence comes in multiple “flavors” and all will be required for an effective defense. SOCs and SOC analysts often operate at the lower levels of “operational” and “tactical” threat intelligence, but that alone isn’t enough. A successful SOC needs to produce, acquire, or purchase all three levels of threat intelligence in order to best defend its organization.

Strategic Level

- *Consumers*: Executives and policymakers
- Looks wide at threat landscape, drives investments, policy, risk

Operational Level

- *Consumers*: Senior responders, managers
- Goals and trends, campaign tracking, adversary capabilities, attribution data

Tactical Level

- *Consumers*: SOC analysts, threat intelligence analysts, incident responders
- IOC level: IPs/domains, host artifacts + analysis
- The most common type for analyst-level usage

Table Of Contents ›

Threat Modeling

Every organization and individual should consider their threat model details as a first step in building out their defense. General guidance on this process can be found in the wonderful Ars Technica blog post referenced below.

Part of building a threat model involves answering questions like the following:

1. What are you protecting?
2. Who are you protecting it from?
3. How likely is it you will need to protect it?
4. How bad are the consequences if you fail?
5. How much trouble are you willing to go through to prevent these consequences?

Reference: sec450.com/threatmodel

For those looking to develop a personal threat model, see the EFF's guidance at the following two links:

1. eff.org/keeping-your-site-alive/evaluating-your-threat-model
2. ssd.eff.org/en/module/your-security-plan

Attack Trees

Attack trees are a method of visualizing the details of your threat model and how attackers might approach achieving their goals in your environment. Creating one for various assumed attacker goals is a useful exercise in brainstorming paths attackers may follow and reveals key evidence collection points to help improve your defensive posture.

Creating an attack tree is easy, simply start by listing the assumed goal of an attack in a box at the top of the page and then “take one step back”, listing all known methods that would allow that to happen. Then repeat down the page into as much detail as you’d like. You can place probabilities, rankings, or likelihoods on each item if desired, but this is not required.

When completed, you will have produced, in reverse, a map of how attackers may feasibly start their attack, work their way through your environment, and ultimately achieve their objective at the end. Doing this thought exercise may illuminate blind spots and other weak points in your organization's defense before the attacker finds them, giving the blue team time to fix the problem before it is discovered during a real attack. The diagram is an example of an attack tree for a bank robbery, provided by Bruce Schneier in the blog post referenced below, where you can find much more detail on how to work through this process.

Attack Trees Reference:
schneier.com/academic/archives/1999/12/attack_trees.html

Table Of Contents ›

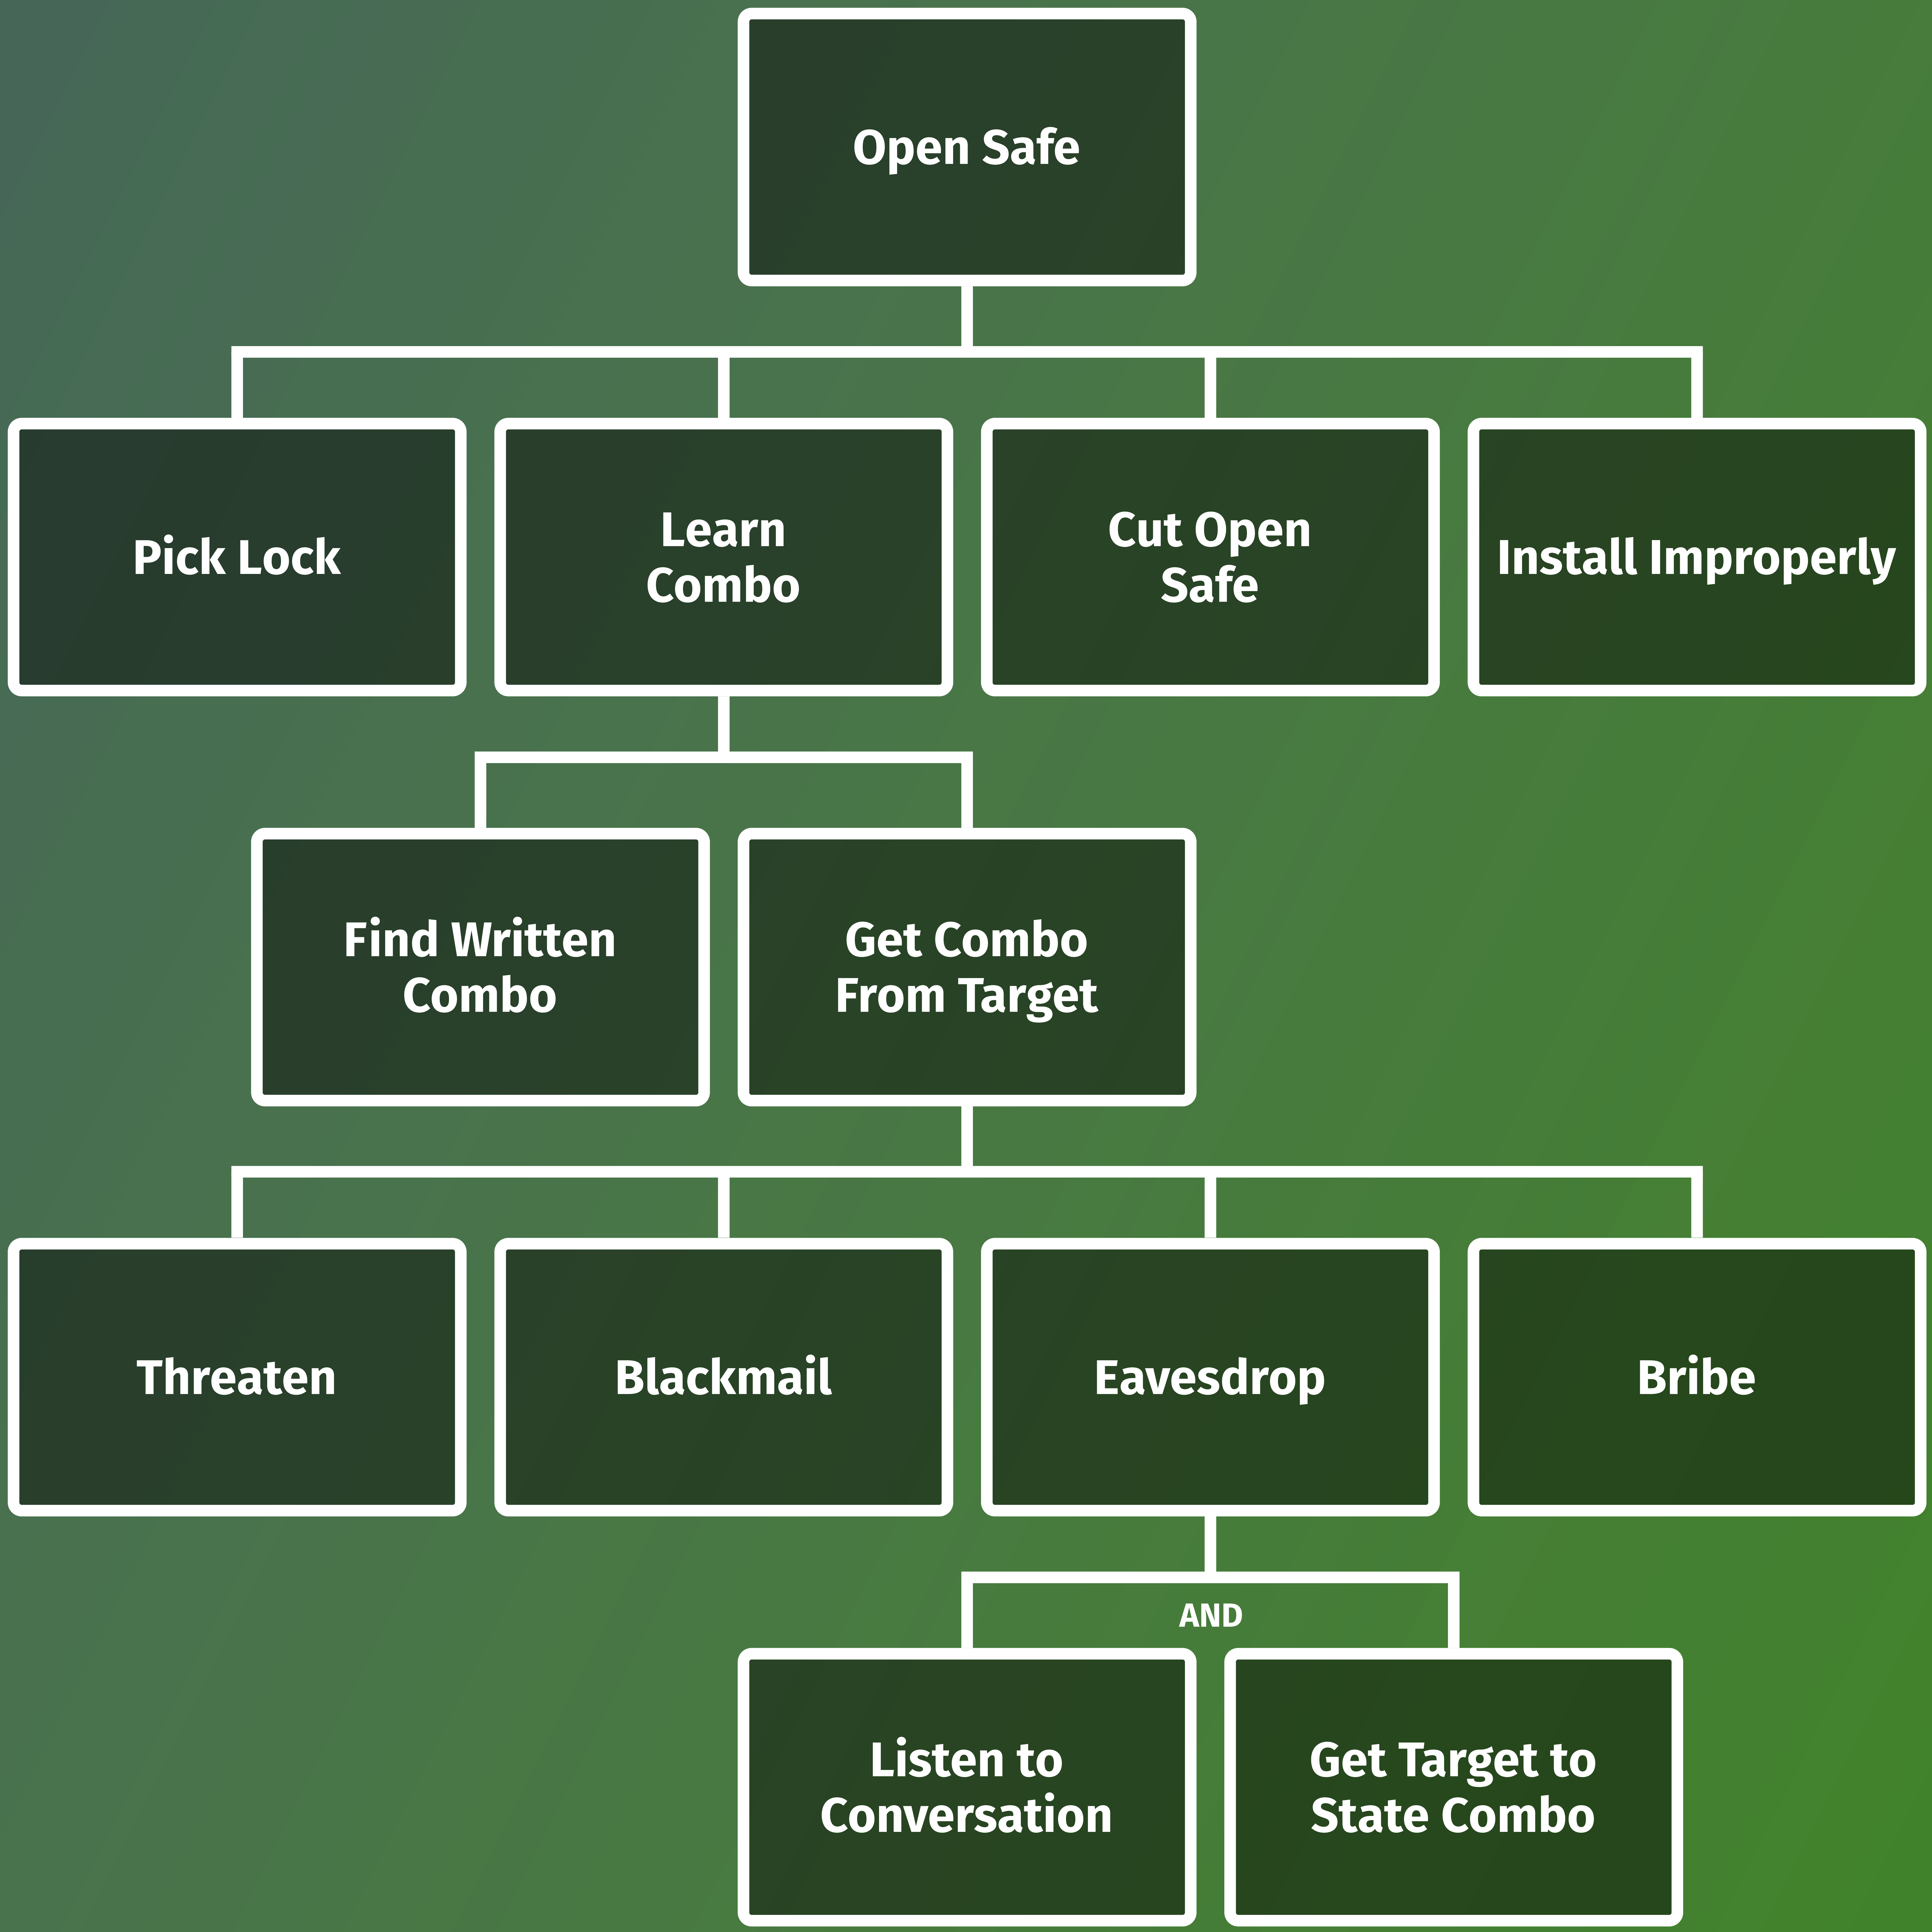

Open Safe

Pick Lock

Learn Combo

Cut Open Safe

Install Improperly

Find Written Combo

Get Combo From Target

Threaten

Blackmail

Eavesdrop

Bribe

AND

Listen to Conversation

Get Target to State Combo

F3EAD Cycle

The F3EAD cycle, explained in detail in the blog post referenced below, is a representation of how threat intelligence and security operations must work together to produce continuous improvement in cyber defense.

Reference: sec450.com/f3ead

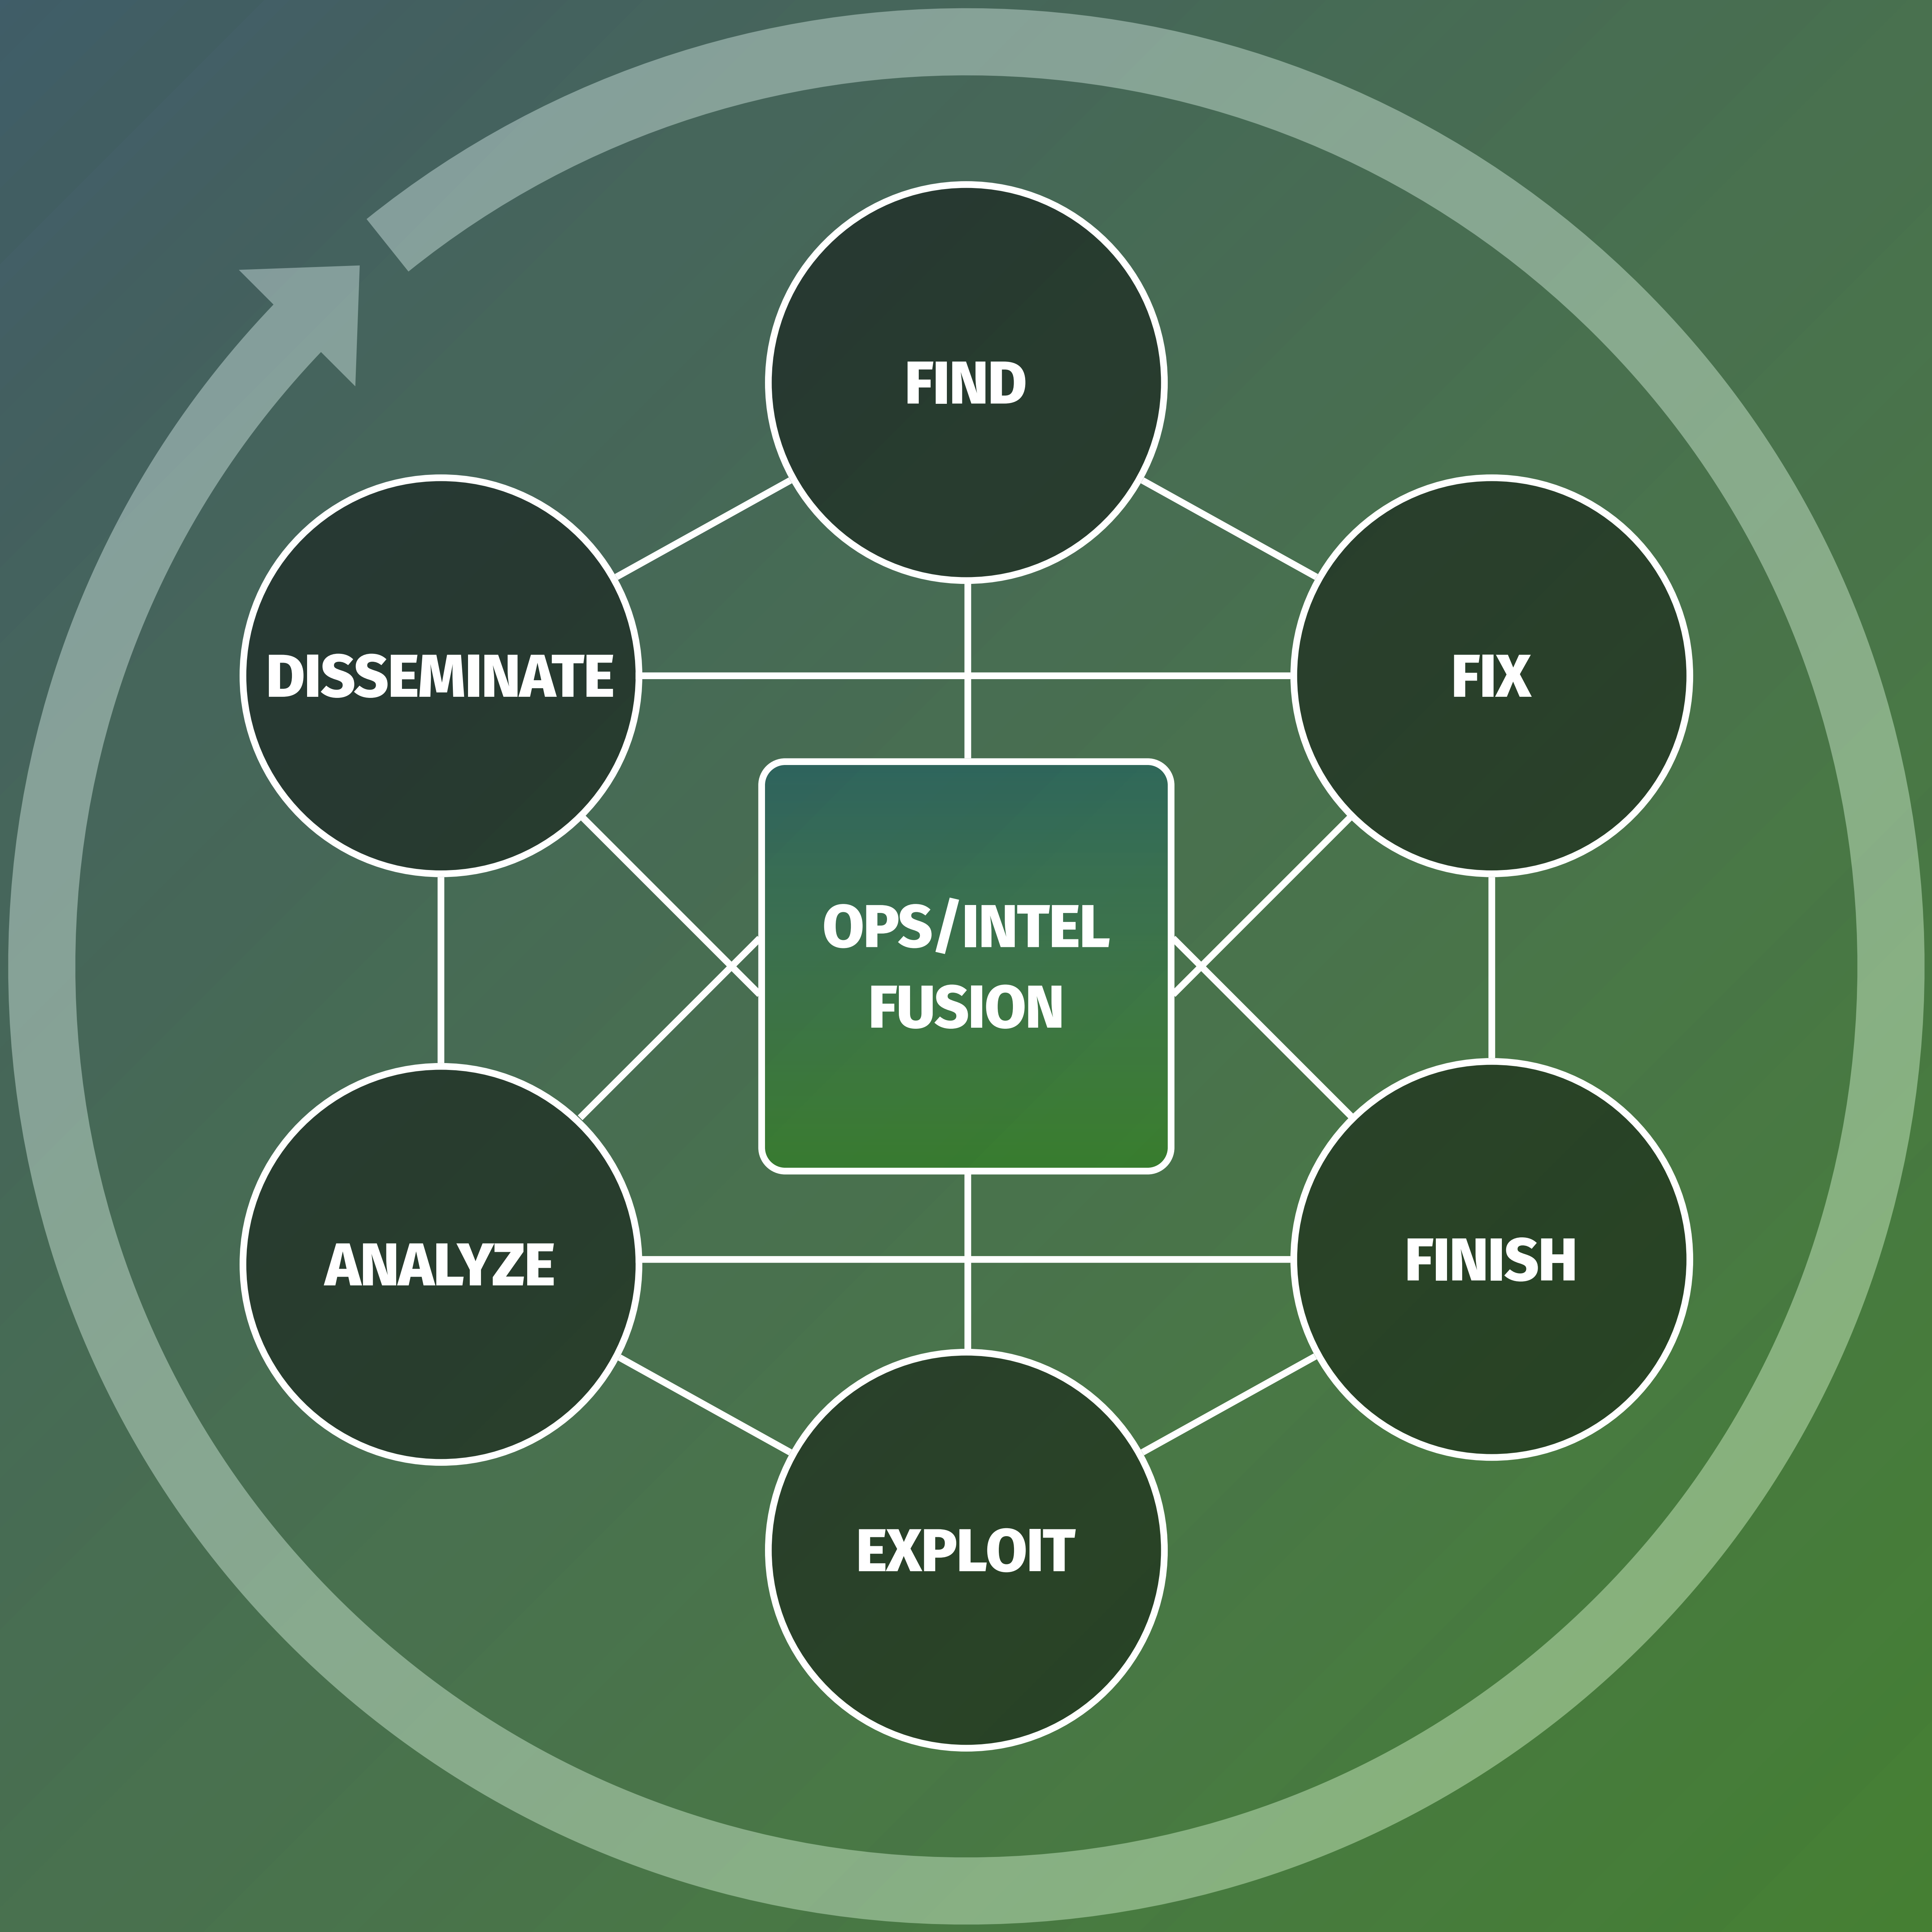

Table Of Contents ›

The OODA Loop

The OODA loop, which was designed by fighter pilot John Boyd in the 1940s, is meant to be a mental model showing the four stages any two parties engaged in any head-to-head competition are constantly going through. It also predicts who may emerge the winner based on the speed at which each side can iterate through the loop. For security operations teams, this model is a useful way to think about the importance of operations tempo, and how speed through each SOC function will contribute to your success fighting off an attacker.

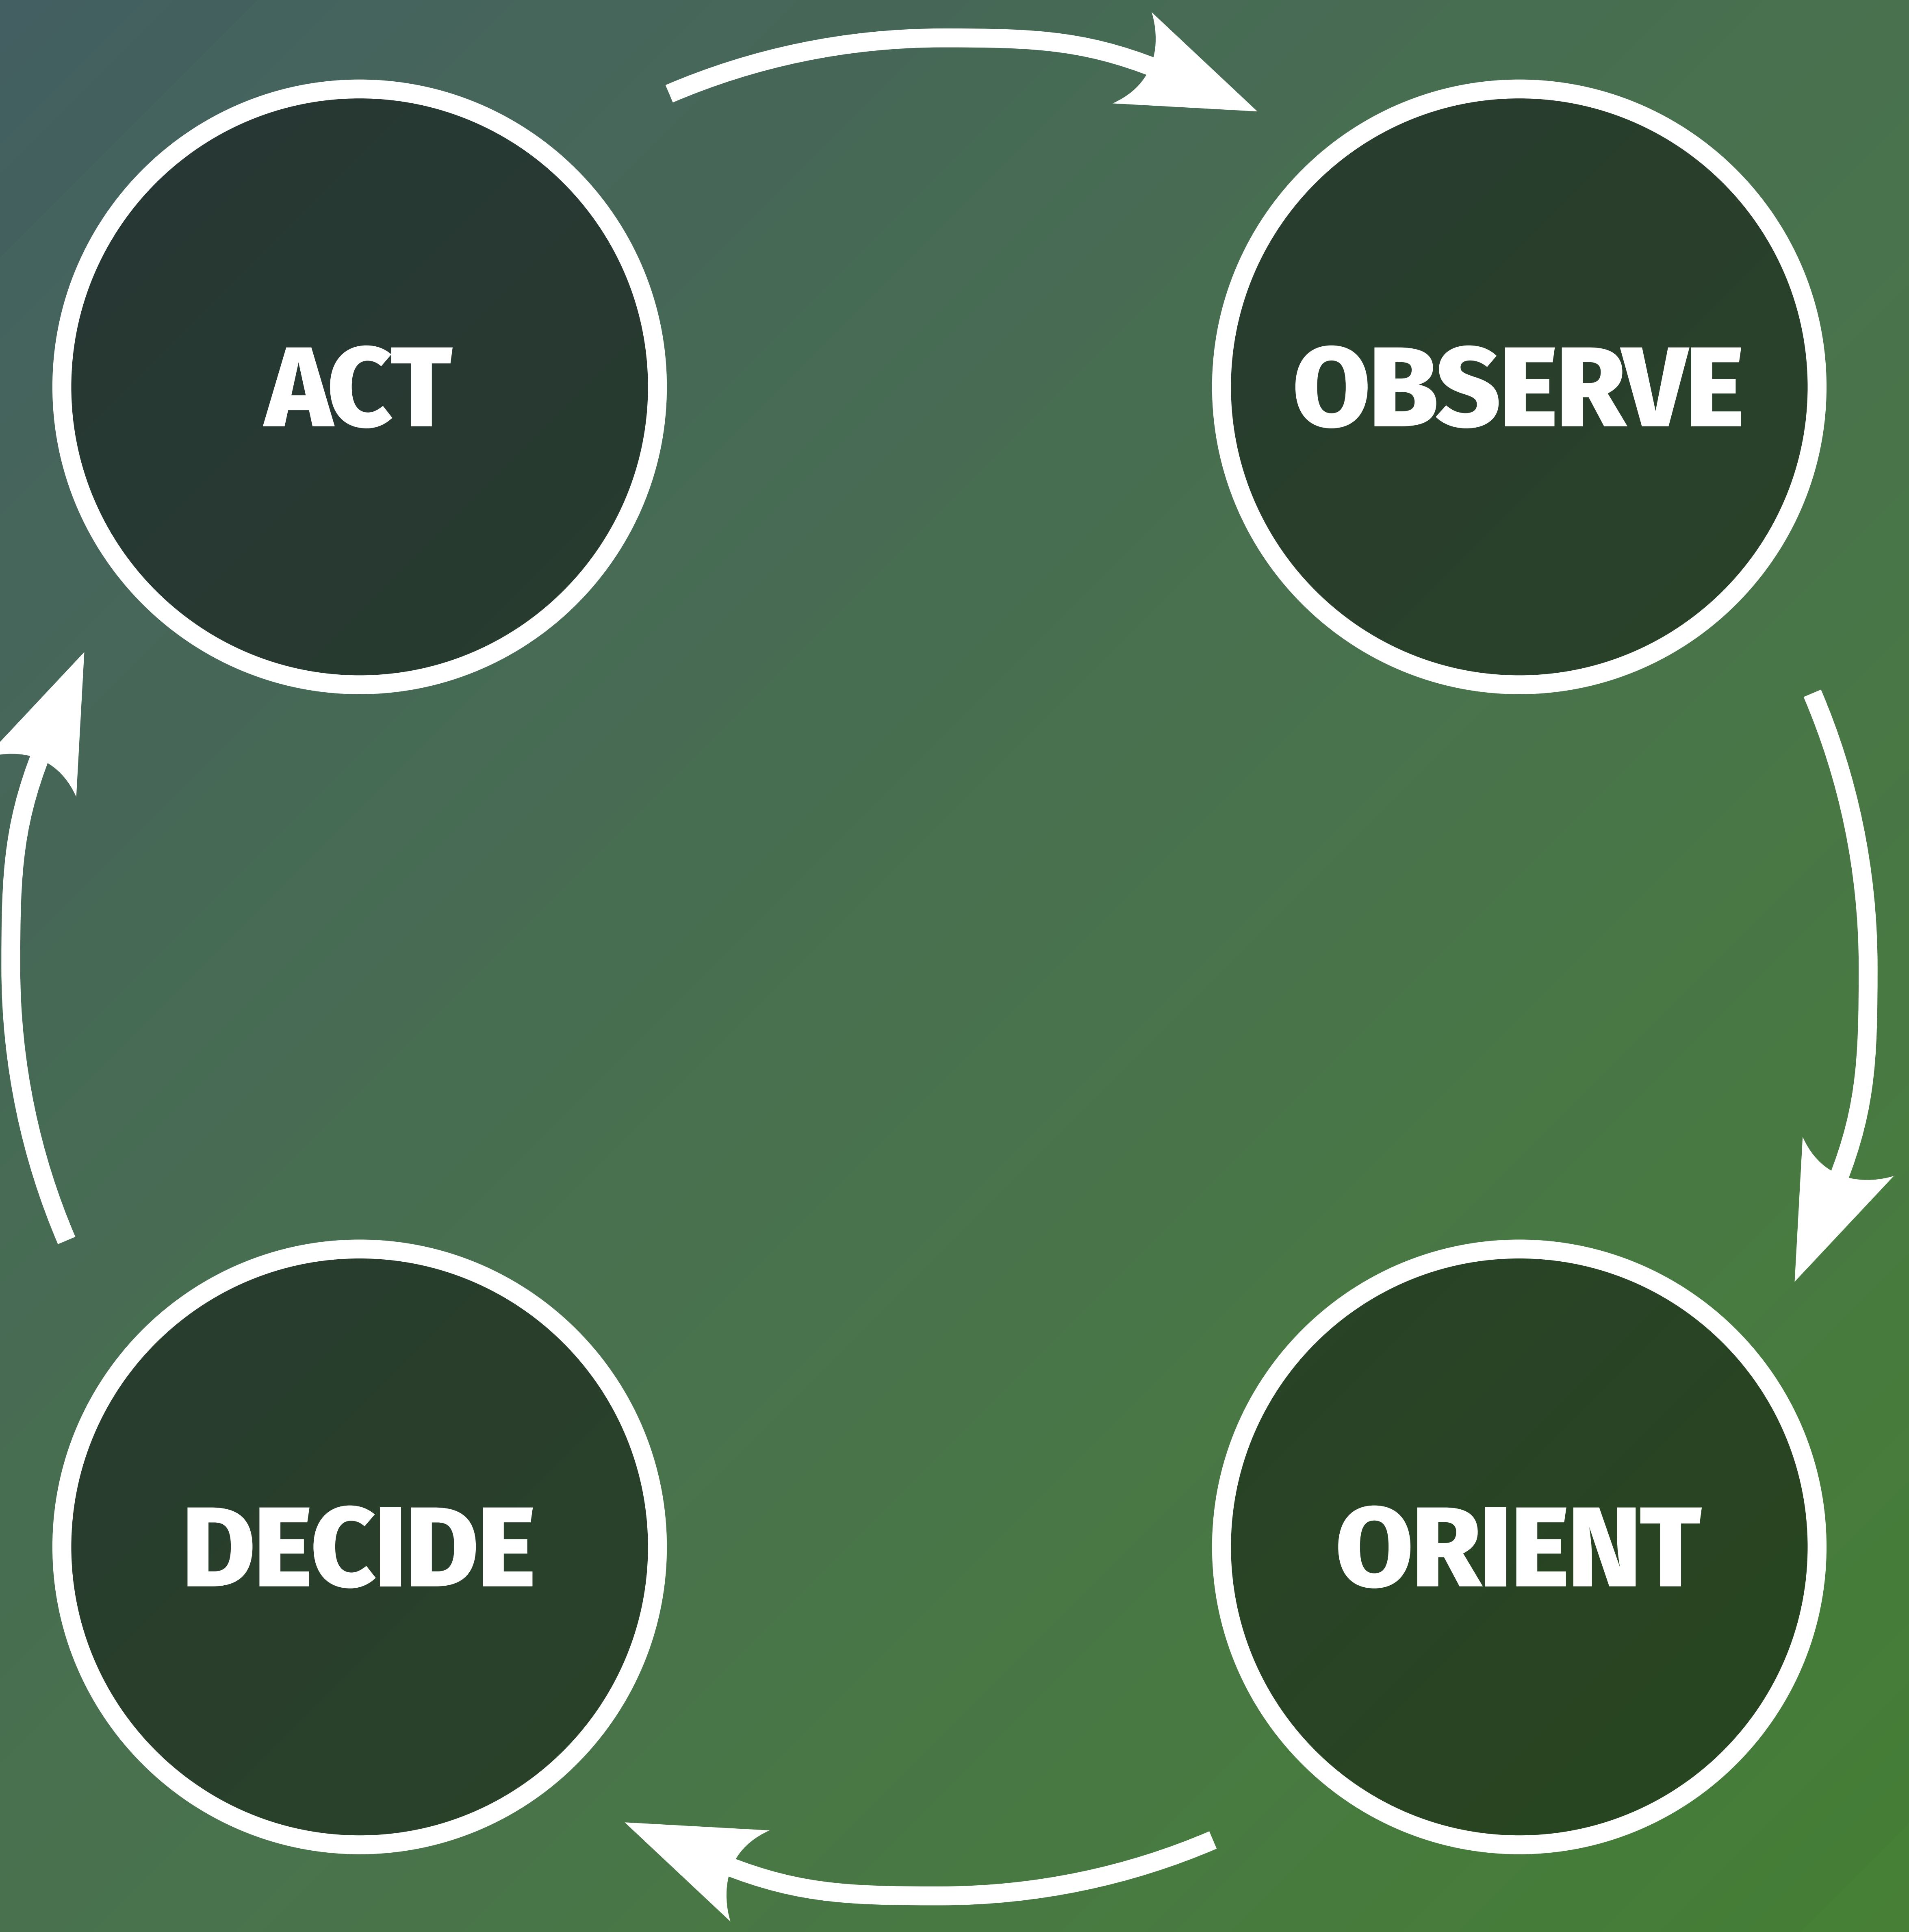

The OODA Loop and Operations Tempo

Observe
This stage is all about the collection of information from the environment. For a SOC, this translates to the collection function, the ability to see the network and endpoint data required to spot an attack.

Orient
Of all the phases, the orient phase has one of the largest impacts as it is where we try to interpret the situation given the data presented to us. For security operations, this is taking the content from the available network and endpoint data and using it to form a hypothesis of what might be occurring. Without a team trained to understand attacks or the available information, this stage may lead to an incorrect interpretation of the data, or a slower than necessary response.

Decide
In this stage it is time to decide, given your interpretation of events from the orient stage, the next best move to take against your opponent. For security operations, this is deciding on how to best disrupt attackers such that we will be able to have an advantage over them in the coming loop iterations. For a SOC, effectiveness in this stage often comes down to experience and training on how to best respond to any given situation. Playbooks and incident response procedures may be of large assistance at this stage.

Table Of Contents ›

Act

The act stage is all about following through on the course of action decided upon at the decide stage. In the SOC, this can be translated to your ability to quickly respond to an attack. Do you have the permission and tools and procedures in place to act fast and minimize damage? The act stage leads to feedback of the impact of your action to the next iteration of the observer step, where the SOC will need to see if the action taken had the desired consequences.

Threat-Hunting

Threat is a key piece of daily operations and is the other key piece of the “detection” function as previously described. While any SOC will have a large library of known attacks and signatures to identify the attacks, it is guaranteed that there are plenty of unknown pieces of malware and attack methods out there as well for which no signature exists. Therefore, threat-hunting is a necessary activity for all SOC teams, regardless of size or experience.

Threat-hunting should be viewed as the “other half” of the detection operation. Instead of being reactive based on matching IOCs, it is a proactive, “assume compromise”, “we know we’re owned so let’s go look for it” type search. Hunting without known malware signatures therefore often involves gathering and carefully analyzing large amounts of data or activity for anomalies that may lead us in the right direction. While all teams can threat hunt, the nature of threat-hunting tactics often means that those with extensive and well-planned data collection and parsing, as well as solid threat intelligence to use as reference, will be the most effective at finding evil. The process can be broken down into the following stages:

Table Of Contents ›

Formulate Hypothesis	What Might Our Adversaries Do?
Define Evidence	What Evidence Would This Type of Attack Leave?
Data-Source Identification	Which Data Sources Are Available in Our Environment That Would Identify This Type of Person?
Gather & Examine Data	Does Our Data Show Compromise?
Respond to Finding	Invoke Incident Response if Needed
Analytic & Protection Improvement	Close the Loop!

Table Of Contents ›

Metrics

Metrics are used in a SOC as an incredibly important feedback mechanism. This feedback is for both the SOC to measure itself, and as a communication mechanism between the SOC and upper management. Internal metrics tracked and watched by those inside the SOC need to show both measurements telling those in the SOC if things are operating in the range of "business as usual", as well as how improvement initiatives and projects are progressing. External metrics must focus on giving upper management the information it needs to make risk and budget decisions, as well as clearly demonstrate the return on investment being produced by the security team. This section contains concepts and assessment considerations for your metrics. Remember, success in the SOC relies on effectiveness in both day-to-day operational tasks as well as continuous improvement.

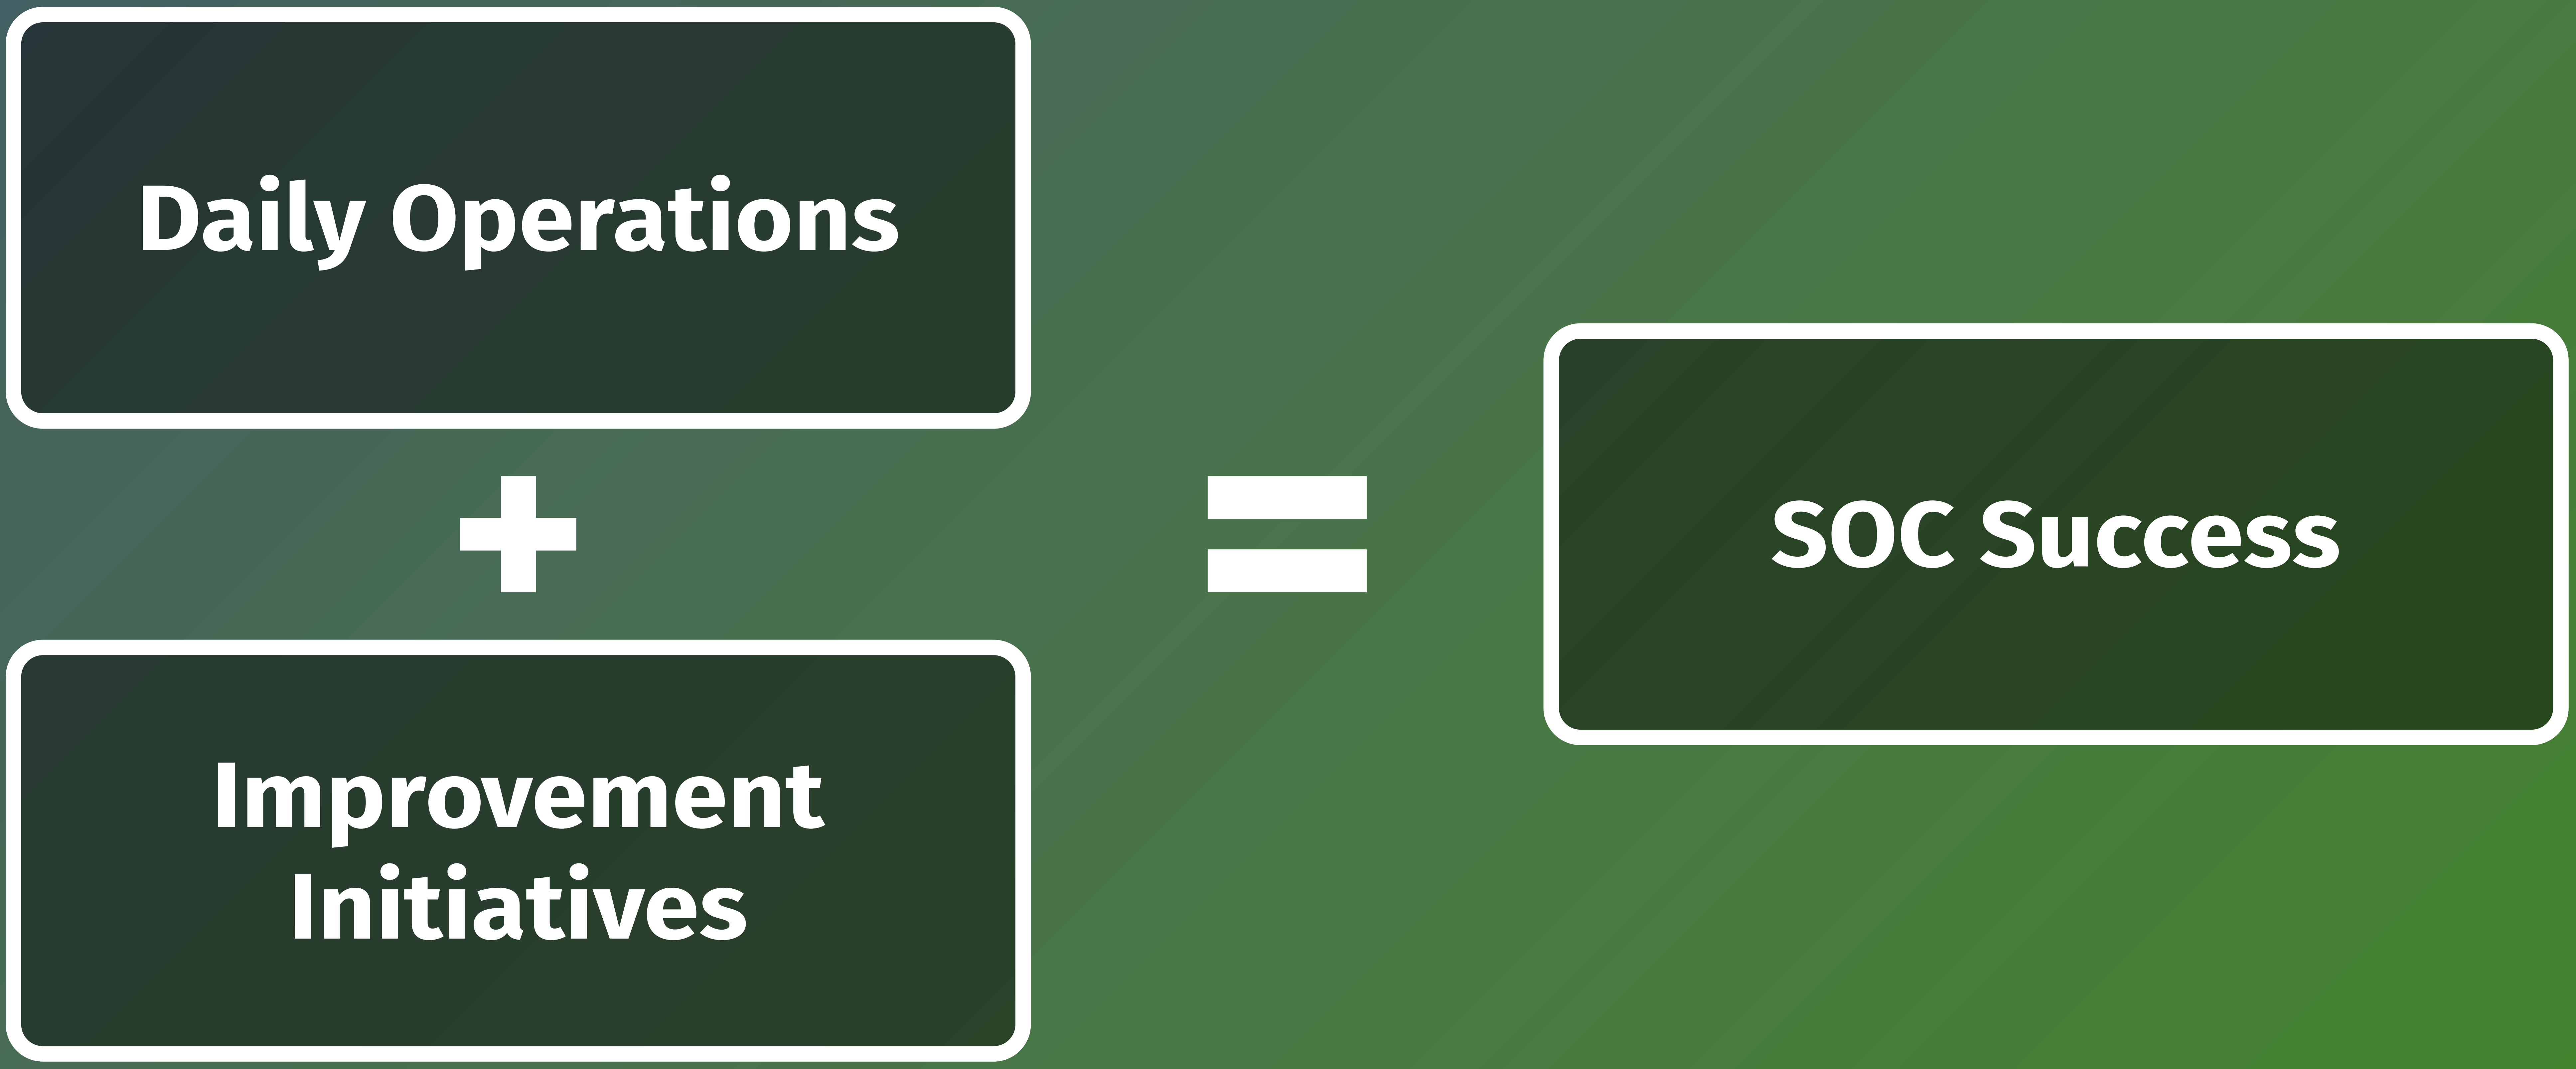

Daily Operations

+

Improvement Initiatives

=

SOC Success

Metrics Types

Metrics
A measurement of a business process, does not include a reference point or context. Metrics are used as key component of the next two items (example: alerts in the triage queue).

Objectives and Key Results (OKRs)
A strategic goal management framework for measuring advancements and initiatives. The OKR system can be a way of generating metrics for improvement initiatives in the SOC.

- **Objective** – The goal to be achieved
- **Key Results** – Specific and measurable ways you know you're making progress toward the objective
- **Initiatives** – Specific actions that will be taken to "move the needle" from the start to ending goal
- **Measurable Components**
  - **Start value** – Metric of interest at the start of the project
  - **Current value** – Metric of the status as of the most recent measurement period
  - **Target value** – Metric value when project will be considered complete

Key Performance Indicators (KPIs)
A metric + the desired limits or boundaries of that metric. Used for measuring and maintaining status quo or "business as usual". An out of range KPI usually means some action must be taken to correct the issue.

KPIs should be developed and used for SOC daily operational items to know that security tools and processes are operating as expected (examples: telemetry health, incident rates, false positives, and any other area where a goal, or an “normal”/“average” can be established).

Examples

- Objective: “Reduce the risk and impact of successful phishing attacks”
- Key Result 1: Less than 1% of spam email delivered to inboxes
- Key Result 2: 30-minute maximum response time from phishing wave alert to containment and protection in place
- Key Result 3: Less than 3 instances per week of malicious emails affecting users
- Initiatives: Malicious email sandboxing implementation, user awareness training, SOC email monitoring and threat-hunting

Reference: “Measure What Matters” by John Doerr

Practical Metrics Considerations

While specific metrics, KPIs, and OKRs of interest may change from team to team, what is largely agreed upon is that all metrics should be goal aligned. In other words, does the metric you are generating ultimately help answer a question about whether you hitting an operational or improvement goal? If not, that metric may not be a good selection to collect. Even goal-aligned metrics can have problems though. Some additional considerations for each number you generate and collect are below:

- How long does it take to generate that metric? Is it a manual or automated process?
- Is the method of gathering the metric well defined? Will two people produce the same result if asked to collect it?
- Is it clear what to do when that metric is outside of the desired bounds? In other words, is it actionable?
- Is the metric collected frequently enough to act fast enough if it falls out of desired bounds? Is it collected too often?
- When reporting a metric, are you needlessly removing fidelity and substituting a color or qualitative description that may make it harder to interpret?

Reference: “GQM” goal setting system

Table Of Contents ›

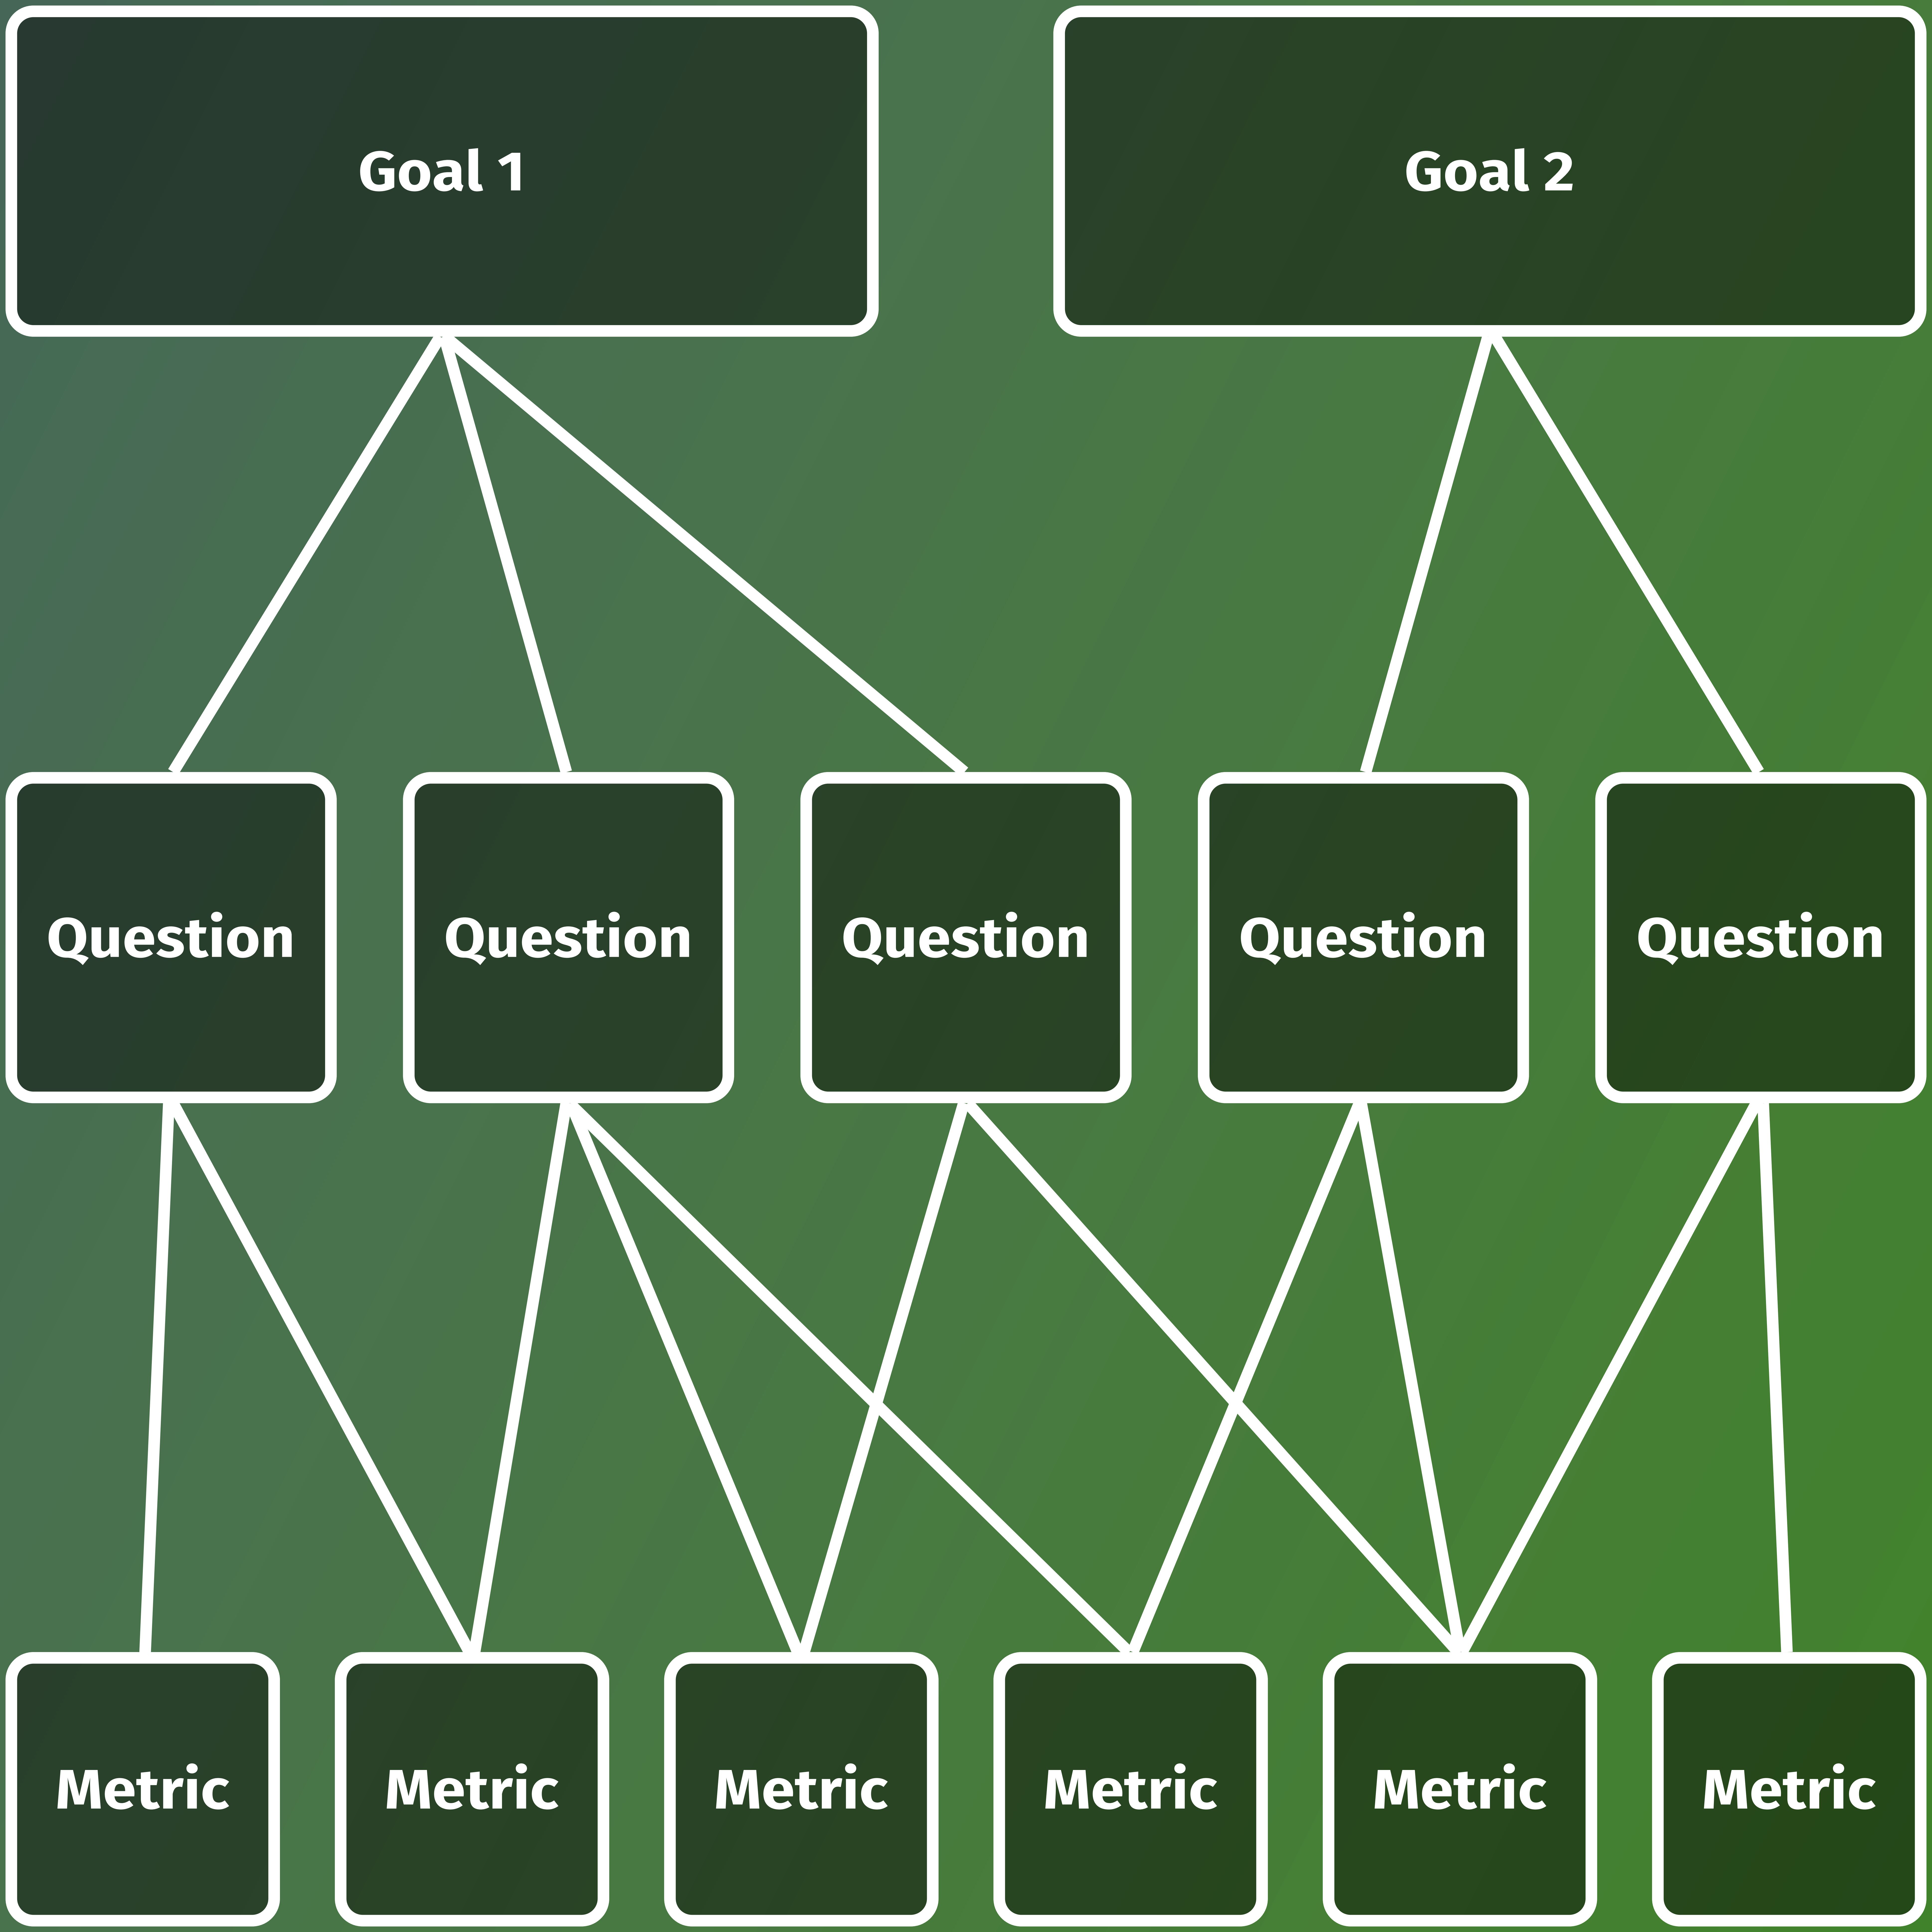

Table Of Contents ›

What to Measure

Conceptually, the SOC should aim to measure the inputs and outputs of processes to ensure they are operating as expected. Think of how you can quantify success in the SOC functions – collection, detect, triage, etc. – and how metrics can be easily, automatically, and continually produced that will give a simple view of whether any key items have changed or broken, or whether anomalies are occurring in incident type or frequency. This forms a small and quickly reactive feedback loop that can promote the fast discovery and correction of issues. Consider both effectiveness metrics (**are we doing the right things?**) as well as efficiency metrics (**how well are we executing the things we have decided to do?**).

Table Of Contents ›

SOC
References

Open-Source Tools

While there are many commercial cyber defense tools for all purposes, teams with constrained budgets need not worry. In the hands of a driven and motivated team, a world-class cyber defense can still easily be crafted using the plethora of outstanding open-source and free solutions that are now available.

Here are some SOC team favorites across various categories:

Incident Management Systems

- TheHive
- FIR

Network Security Monitoring

- IDS, network metadata
  - Suricata: IDS, network metadata, and PCAP capable
    - EVEbox: Alert triage
  - Snort
  - Zeek
- Full Packet Capture
  - Moloch
  - Google Stenographer
  - Netsniff-ng
- Distributions
  - Security Onion
  - RockNSM

Table Of Contents ›

Endpoint monitoring / HIDS

- NXLog Community Edition: logging agent
- OSQuery
- OSSEC: HIDS
- Sysmon
- Wazuh: HIDS

Incident Response

- Kansa
- Velociraptor

Malware Sandbox and Malware Analysis

- Cuckoo Sandbox
- REMnux: malware analysis tools Linux distro

Threat Intelligence Platforms

- MISP
- OpenCTI

Purple Team Testing and Reporting

- Vectr

SIEM / Log Management

- Elastic Stack
  - Elastalert: Alerting Engine

Security Orchestration Automation and Response

- NSA Walkoff
- Shuffle
- IBM Node-Red: generalized automation / orchestration framework

Podcasts

There are many high-quality podcasts available revolving around the different aspects of cybersecurity, and this is by no means a complete list. But it is provided as a great place to get started.

SANS/GIAC-Produced Podcasts

- BLUEPRINT – John Hubbard
- GIAC Trust Me I'm Certified – Jason Nickola
- SANS Internet Storm Center – Johannes Ullrich
- Wait Just An InfoSec
- Cloud Ace

General

- Beers with Talos
- Brakeing Down Security
- Cyber Security Interviews
- The CyberWire Daily
- Darknet Diaries
- Defensive Security Podcast
- Hacker Valley Studio
- FireEye State of the Hack
- Paul's Security Weekly, Enterprise Security Weekly
- Security Now

Books

Security Operations/Reference

- Crafting the InfoSec Playbook: Security Monitoring and Incident Response Master Plan – Jeff Bollinger, Brandon Enrich, and Matthew Valite
- MITRE Top 10 Strategies of a World Class CSOC – MITRE/Carson Zimmerman - FREE
- Blue Team Handbook: Incident Response Edition: A Condensed Field Guide for the Cyber Security Incident Responder – Don Murdoch
- Blue Team Handbook: SOC, SIEM, and Threat-Hunting: A Condensed Guide for the Security Operations Team and Threat Hunter – Don Murdoch
- Blue Team Field Manual – Alan White and Ben Clark

Malware

- Practical Malware Analysis - Michael Sikorski and Andrew Honig
- Malware Data Science - Joshua Saxe and Hillary Sanders

Incident Response

- Applied Incident Response – Steve Anson

Honeypots / Active Defense

- Intrusion Detection Honeypots: Detection Through Deception – Chris Sanders
- Offensive Countermeasures: The Art of Active Defense – John Strand, Paul Asadoorian, and Ethan Robish

Table Of Contents ›

Threat Intelligence

- Intelligence-Driven Incident Response: Outwitting the Adversary – Scott J. Roberts and Rebekah Brown

For Getting Started in InfoSec

- The Linux Command Line – William Shotts - FREE
- Linux Basics for Hackers – OccupyTheWeb
- Practical Packet Analysis – Chris Sanders

Table Of Contents ›

SANS SOC Training Courses

SEC450: Blue Team Fundamentals: Security Operations and Analysis

Author: John Hubbard

This course provides students with technical knowledge and key concepts essential for security operation center (SOC) analysts and new cyber defense team members. By providing a detailed explanation of the mission and mindset of a modern cyber defense operation, this course will jumpstart and empower those on their way to becoming the next generation of technical blue team members. **sans.org/sec450**

GIAC Security Operations Certified (GSOC)

The GIAC Security Operations Certified (GSOC) certification validates a practitioner's ability to defend an enterprise using essential blue team incident response tools and techniques. GSOC-certified professionals are well-versed in the technical knowledge and key concepts needed to run a security operations center (SOC). **giac.org/gsoc**

LDR551: Building and Leading Security Operations Centers

Authors: John Hubbard and Mark Orlando

This course is for SOC managers and leader looking to unlock the power of proactive, intelligence-informed cyber defense. Learn how to combine SOC staff, processes, and technology in a way that promotes measurable results and covers all manner of infrastructure and organizational requirements. LDR551 will give SOC managers and leaders the tools and mindset required to build the team, process, workflow, and metrics to defend against modern attackers by building the processes for continuously growing, evolving, and improving the SOC team over time. **sans.org/ldr551**

GIAC Security Operations Manager (GSOM)

The GIAC Security Operations Manager (GSOM) certification validates a practitioner's ability to effectively manage a technical team and strategically operate a Security Operations Center (SOC) to align with an organization's business goals and security requirements. **giac.org/gsom**

SANS Free Resources

Free Educational Resources
sans.org/free

SANS Blogs
sans.org/blog

SANS Newsletters
sans.org/newsletters

SANS Reading Room
sans.org/reading-room

SANS Webcasts
sans.org/webcasts

SANS Posters
sans.org/posters

SANS Internet Storm Center
isc.sans.edu

About the Author

John Hubbard | @SecHubb

John is a Security Operations Center (SOC) consultant and speaker, a Senior SANS Instructor, and the course author of two SANS courses:

- SEC450: Blue Team Fundamentals: Security Operations and Analysis
- LDR551: Building and Leading Security Operations Centers

John also teaches additional SANS Blue Team courses such as SEC511: Continuous Monitoring and Security Operations; and SEC555: SIEM with Tactical Analytics. Through his years of experience as a Lead Cyber Security Analyst and SOC Manager for a major pharmaceutical company with over 100,000 employees and global operations, John has developed real-world, first-hand knowledge of what it takes to defend an organization against advanced cyber-attacks.

Learn More About John ❯

SANS The most trusted source for cybersecurity training, certifications, degrees, and research