Cloud-Based, Intelligent Ecosystems

Redefining Intelligence & Driving to Autonomous Security Operations

Acknowledgments

Authors:

Paul Kurtz
Jim Reavis
John Yeoh

Key Contributors:

Ed Amaroso
Dave Cullinane
Dave Dervy
Ken Dunham
Ramses Gallego
Robert Gourley
Ernie Hayden
Sean Heide
Mat Henley
Jeff Huegel
Brian Kelly
Mark Kraynak
Nico Kseib
Anil Markose
Danielle Meah
Krishna Narayanaswamy
Rich Phillips

CSA Staff:

Stephen Lumpe (Cover)
AnnMarie Ulskey (Layout)

Table of Contents

Executive Overview

This paper proposes a call to action for security executives to break the endless cycle of iterative tool adoption and, instead, move to data-centric security operations, driving integration and automation leveraging cloud-based fusion. Section I unpacks "intelligence" and addresses the challenges of integrating data from internal security tools and external threat feeds. Section II leverages lessons learned from the autonomous vehicle industry's "sense, understand and act;" Section III, proposes secure, intelligent ecosystems to enrich data workflow and apply machine learning. Section IV, addresses security business analytics and the importance of measuring business outcomes for boards of directors, chief information security officers, and security operators. Section V proposes areas for further exploration and investigation.

> "Cyber is the most versatile instrument of power because it can be used overtly or covertly and with great effect against military and economic targets. It can also accomplish political objectives, recruit adherents, undermine confidence in institutions, affect decision making, and impact the personal lives of millions of people."[1]
>
> **- Robert Gates, Exercise of Power**

Former Secretary of Defense Robert Gates expresses the potential harm and power of cyber-attacks at all society levels. Despite the severity of the situation and significant contributions from the Solarium Commission's report released earlier this year, we have failed to rethink the fundamentals of cyber defense. We are in a cyber arms race that has precipitated a security tool-race with adversaries' evolving attacks forcing us to spend more to try to defend ourselves. Our default response is to adopt new tools to try to keep up. Over the past twenty years, tools have evolved from data loss prevention, unified threat management, managed detection and response, user behavior analytics to "next-generation" fill-in-the-blank. We are losing this race as adversaries continue to outpace defenders. In addition, we call for government action, increased criminal penalties, regulation, and new organizations. Action is warranted in each area, but success will be temporary until we break the cycle set twenty years ago and place a new cornerstone for cyber defense: cloud-based, data-centric defense. Leadership can break the cycle to defend companies through data-centric, integration, and automation of their tools and overall architecture.

[1] Robert Gates, *Exercise of Power*, Knopf, Borzoi Books, 2020 p.31

Section I. Redefining "Intelligence"

We can continue to look for the ideal mix of government action, new tools, and technologies to defend ourselves. Still, unfortunately, no combination will exist for the foreseeable future, given the rapid rate of change in technology and sophistication of adversaries. However, data integration and automation from existing security tools and intelligence feeds provide a far better chance to detect, respond, and recover from cyber events consistently. To succeed, we need to change our thinking in three areas:

> First, we must revise what we mean by "intelligence" in the context of cybersecurity. Intelligence can't be viewed only as external data about adversary tactics, techniques, and procedures. Instead, it must be defined as an organizations' capacity to normalize, transform, and automatically extract actionable insight and context from internal security tools and external sources to reduce the mean time to detect and respond. In this context, intelligence is defined as an organization's ability to know the threat and the operating environment to drive decisions.
>
> Second, we need to build "cyber memory." As Michael Kanaan states in T-Minus AI, "without memory, it's definitionally impossible for any individual, animal, organism, or other structure--computers included--to learn. It's just not possible"[2] Later in the book, Kanaan states, "our ability to learn requires our capacity to acquire data and to analyze it. Without data, intelligence just isn't possible, not at any level--and not any animal, individual, or computer."[3] Rather than responding to a stream of cyber events "playing whack a mole," we need to recall event data gathered from security systems seamlessly. Creating a virtual "memory" to absorb events will enable the use of Machine Learning (ML) to identify patterns to more effectively and efficiently address malicious activity. This memory would constitute data from both internal security tools and external threat data.
>
> Third, we must build and maintain secure, intelligent ecosystems. Secure, intelligent ecosystems are cloud-based memory banks that continuously fuse and enrich data from internal security tools and external sources. The enriched data can automatically update cyber defense tools or triage for further action by analysts. Data from an individual ecosystem can be shared with other companies or organizations to form a collaborative defense ecosystem.

This is not a call for a singular product but a new mindset to use "intelligence" to integrate and automate data workflows from security tools and sources used within and between enterprises to create intelligent ecosystems. Integration and automation of data workflows should increase the efficiency of existing tools and personnel. The results should lead to measurable outcomes for boards of directors, chief information security officers, and analysts.

[2] Michael Kanaan, *T-Minus AI*, BenBell Books, August, 2020, p 59.

[3] Michael Kanaan, *T-Minus AI*, p. 81.

How did we get here?

Over the past twenty-five years, we sought to secure ourselves from cyber-attacks through a series of security solutions and technologies. Each was built and deployed with the best intentions: to secure an increasingly computer-dependent society against more sophisticated threats and adversaries. Public and private sector companies, government organizations, and venture capitalists all contributed. Each new security problem was met with a novel or incremental improvement in technology. For example, firewalls became "new generation" firewalls, and intrusion detection became intrusion prevention. Novel approaches spanned a range of problems, from user behavior analytics to secure email gateways to technology to combat phishing. Threat intelligence platforms gave companies a better understanding of the capabilities and tactics of adversaries.

The number of security tools used today continues to increase though their efficacy is being questioned. A Poneman survey released last July states:

> - Companies deploy, on average, 47 different cybersecurity solutions and technologies.
> - Fifty-three percent of IT experts admit they don't know how well the cybersecurity tools they've deployed are working.
> - Only thirty-nine percent of respondents say they are getting full value from their security investments.[4]

A separate survey from December 2019, prepared by Market Cube commissioned for ReliaQuest's "Security Technology Sprawl Report," offers the following:

> - Security teams are deploying more tools than ever. Almost three-quarters (70%) of respondents say they've invested in more than five new technologies in the last year, including 19% who say they've invested in more than 20.
> - Teams are struggling to implement the tools. Seventy-one percent report they are adding security technologies faster than they are adding the capacity to productively use them.
> - The burden of tool maintenance compromises the threat response. Sixty-nine percent report their security team currently spends more time managing security tools than effectively defending against threats.
> - A majority of enterprises are less secure today as a result of tool sprawl. Over half (53%) say their security team has reached a tipping point where the excessive number of security tools in place adversely impacts security posture.[5]

When these points are taken together, we appear to be decreasing security and efficiency while at the same time increasing operations and personnel costs. Ironically, our complex and costly operations are increasing the probability of adversaries' success and exhausting defenders.

[4] https://www.businesswire.com/news/home/20190730005215/en/Ponemon-Study-53-Percent-Security-Leaders-Don%E2%80%99t

[5] https://www.prnewswire.com/news-releases/rapid-growth-of-cybersecurity-tools-in-enterprises-increases-risk-according-to-reliaquest-survey-300972019.html

After stepping back from the problem, the complexities associated with security tool adoption become more apparent. In IT infrastructure management, we have seen this problem before. The graphic below from Forrester's study on IT infrastructure transformation depicts the state of infrastructure before the adoption of greater integration and automation.

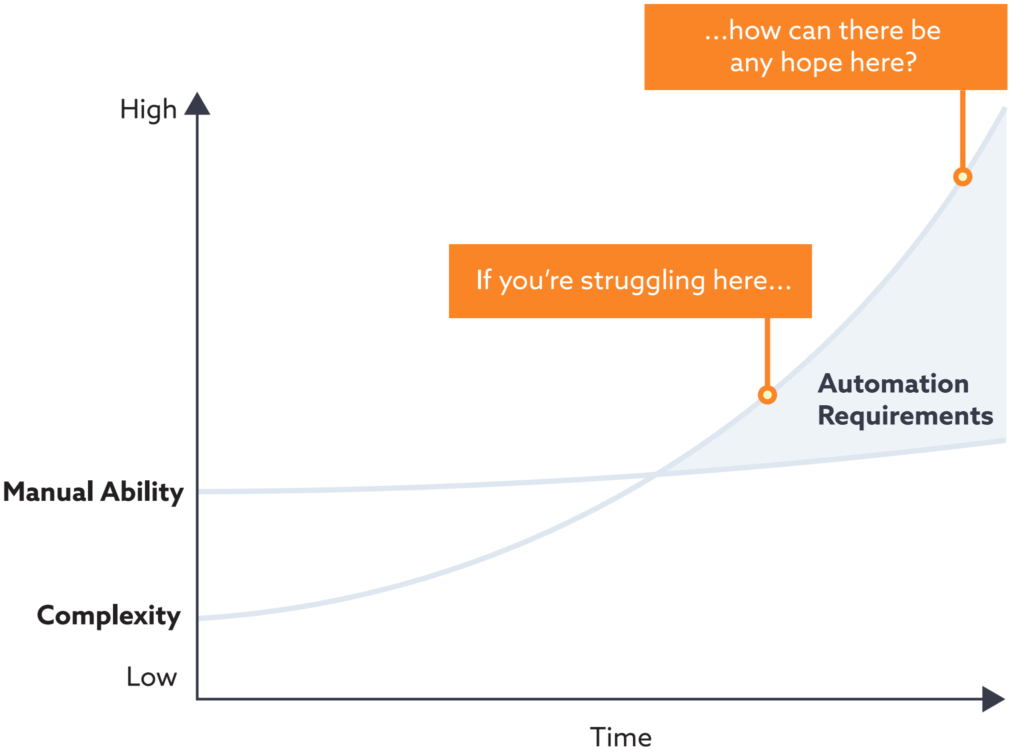

Figure 1

Forrester, Automation Drive the I&O Industrial Revolution, The Infrastructure Transformation Playbook, Glenn O'Donnell, and Chris Gardner, November 19, 2019.

The graphic above can be applied to the current state of cybersecurity. However, cybersecurity challenges differ from IT infrastructure transformation in three ways: the rapid pace of the evolving threat, the explosive growth of security tools, and the time-sensitivity to respond to attacks. Perhaps naively, we thought security with the right tools or intelligence sources would give defenders the upper hand. However, most tools and intelligence feeds—developed with the best intentions—have a short half-life, and adversaries stay at least two steps ahead. A high-risk, multifaceted war is underway, and that requires an expeditious change in strategy.

Rethinking Internal and External

A key consideration in the protection of enterprise has always been the difference between so-called internal and external data, assets, and resources. For many years, this was defined by physical boundaries in network perimeters enforced by firewall hardware that would try to maintain policy control on either side of its positioning. By default, internal assets were located on the inside of the firewall, and external assets were found on the outside. More recently, however, this has had to be fundamentally rethought. With cloud infrastructure, SaaS-based applications, social media usage, mobile apps, and other less traditional enterprise resources, the definition of internal and external data is less obvious. The difficulty comes from whether such an asset is defined as internal or external. For this paper, it might be best to view any security tool defending mission-critical operations as an internal asset; this could encompass SIEM, EDR, secure e-mail gateway, and vulnerability identification assets. External sources would be proprietary or open-source intelligence feeds or reports.

Traversing the Valley of Death - The Road to Integration and Automation

CSA took a step back to examine the challenge holistically and identified a critical gap: the absence of a capability to easily leverage and fuse output from security tools and threat intelligence deployed. While many tools exist and some are integrated, why is a data processing layer missing that interconnects security tools and intelligence sources? Solving hard data fusion problems is not new. Data challenges exist in counterterrorism, cancer research, and climate change, for example. At least five unique challenges surfaced with cybersecurity:

> First, as described above, both security technology and adversaries are changing fast. Keeping pace with new and emerging problems has made it hard to step back to examine the situation as a whole and the underlying issues that develop into more pronounced threats.
>
> Second, the vendor community has focused on a “single pane of glass” that visually represents event data. The idea is good, but there are at least two limitations. First, the wealth and diversity of event data are hard to represent, along with the pace of malicious activity. Second, buyers are reluctant to commit to a single pane, given the significant investment in training on major security products, such as SIEMs, case management, and orchestration platforms.

The result is analysts in each area engage in “tool toggle,” leaning on a platform of choice, and manually bringing in data from dozens of sites. Threat Intelligence (TI) teams focused on threat hunting seek their own tools, understandably. Security operations center (SOC) teams maintain security stack tools to monitor current operations. Fraud operations are often set apart from SOC and TI, creating at least three silos of activity. The knock-on effect is the creation of data silos. This situation creates inefficiency within and between teams using different sets of tools and little data integration. TI teams pass intel findings “over the fence” to the SOC, and, in many cases, it is unclear whether the SOC follows up or deems the data useful. The Poneman Institute’s recent survey on SOC operations highlights the problems: managing threat intelligence, malware protection, and defense, waiting on tools to respond to operations, and tool maintenance.[6] The situation is made worse, given the uneven value of threat intel providers and differing scoring metrics. Recent research showed that two of the largest threat intel providers could offer different views of the threat landscape.[7] This compounds the problem for TI analysts and SOC operators alike as they are left to assess value and decide whether or not to take further action based on conflicting reporting.

[6] 2020 Devo SOC Performance Report, A Tale of Two SOCs, Survey Independently Conducted by the Poneman, Institute, p.10

[7] Robert Lemos, Dark Reading, “Research Casts Doubt on Value of Threat Intel Feeds, August 14, 2020

> Third, the absence of a readily implementable exchange protocol and data-labeling ontology slowed progress. Common standards, such as Structured Threat Intelligence eXchange (STIX), sought to enable information exchange between companies in a common format. STIX 2.1 was recently deployed, but, given the pace of technology and attacks, it has proven difficult to implement except for large companies. The classification of adversary behavior also remained a challenge until recently, with the wider adoption of MITRE's ATT&CK framework by companies and vendors. MITRE's ATT&CK framework is easy to deploy, aided by quarterly updates to adversary behavior classifications.
>
> Fourth, normalization and transformation of disparate data sets from security tools and intel sources have represented the "valley of death" for integration and automation until recently. The integration and processing of data were difficult given different formats and protocols, managing duplicates and redactions, and the importance of understanding context. A classic example is differentiating a software version number for an IP address. The task has been manual and tedious, requiring significant personnel resources to support workflow processes in security orchestration, automation, and remediation (SOAR) platforms. The addition of more readily available computing capabilities on cloud-based assets enabled computation and fusion around much larger data sets.
>
> Fifth, a shift from a singular focus on software and products to secure systems to focusing on the data generated by security systems. The data science associated with extracting value from the data is critical to focus on the complex relationships within the data discovered with models.

The latter two areas--data normalization and transformation--and data science stand out as the most significant challenges. The next section will discuss these problems in greater detail as success here enables a transition to a holistic security ecosystem, including the selective application of machine learning.

Section II. Sense, Understand, Act

The cornerstone of a new paradigm is integrating and automating the output from internal security tools and intelligence sources we use today and those developed for future security challenges. In the business context, we are familiar with Observe, Orient, Decide, Act (OODA), where an external problem is identified and steps are taken to mitigate the challenge. The paradigm shift can also be illustrated by looking at a different industry: autonomous vehicles. The autonomous vehicle industry identified three functions: sense, understand, and act, which may more closely align with cybersecurity challenges given the dependence on computers and interaction with human operators.

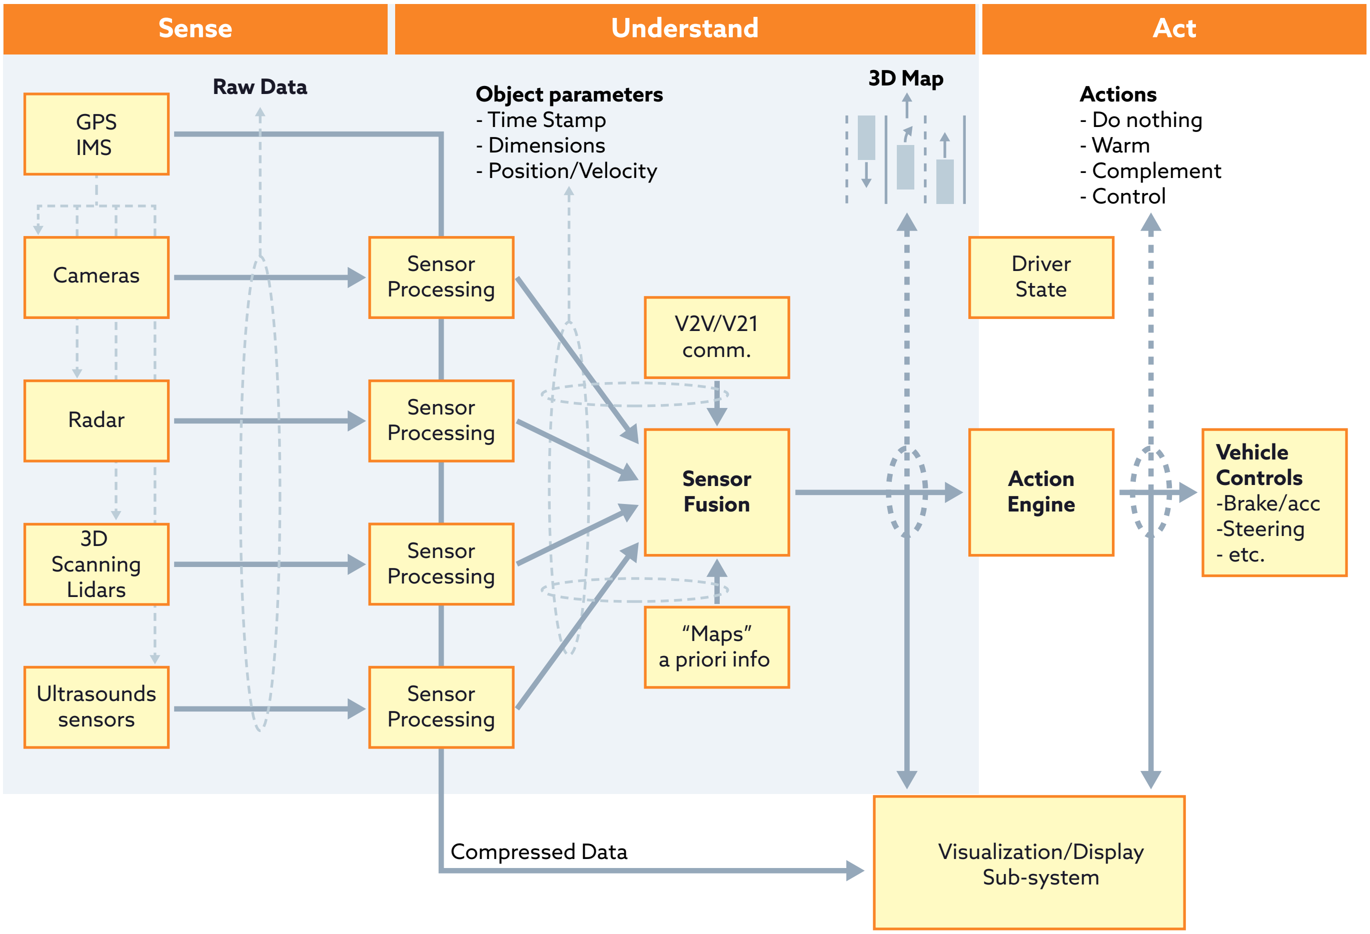

Figure 2

Papoutsidakis, Michail; Kalovrektis, Konstantinos;Drosos, Christos; Stamoulis, Georgios; 6/15/2017, Design of an Autonomous Robotic Vehicle for Area Mapping and Remote Monitoring, International Journal of Computer Applications

The challenges in cybersecurity are strikingly similar to those noted in the autonomous vehicle schematic above. The output of different sensors (cameras, radars, lidars, and ultrasound) deployed to process, classify, and fuse disparate data sets to automate actions for the auto's decision-making computer system and driver. The challenge is complex. For example, the perception and movement of external objects such as other cars, people, speed, and weather must be sensed, understood, and acted upon quickly.

In the cybersecurity ecosystem, endpoint detection and response (EDR) and security incident event management systems (SIEM) would be considered "sensors." Threat intelligence would include external data about real or perceived problems, and security orchestration and automated response systems would support decision management by computers and humans. While the comparison between automobiles and cybersecurity is not exact, the challenges associated with timeliness, false positives, false negatives are germane. Normalization and fusion of data to support computer-aided workflow decision-making processes are central to both. A safe, usable, fully autonomous vehicle is impossible without a holistic ecosystem to sense, understand, and act on data quickly. A cyber ecosystem is similar: data workflow is critical. The schematic below borrows significant elements from the autonomous vehicle industry and applies them to cybersecurity.

The practical application of "sense, understand, act" means vendors and organizations need to categorize their tools and operations. This represents a significant challenge for the security industry, as vendors often position tools as standalone products versus serving as part of a system, as an autonomous vehicle. This is separate but related to the problem outlined above on data normalization and transformation. However, buyers can accelerate the creation of a system by ensuring they buy tools that can seamlessly integrate with other tools and intelligence management systems. Here again, the autonomous vehicle industry provides an example for sensor systems modeling, which, in turn, has used biological examples such as boid's model of animal flocking or the Couzin model of animal motion.

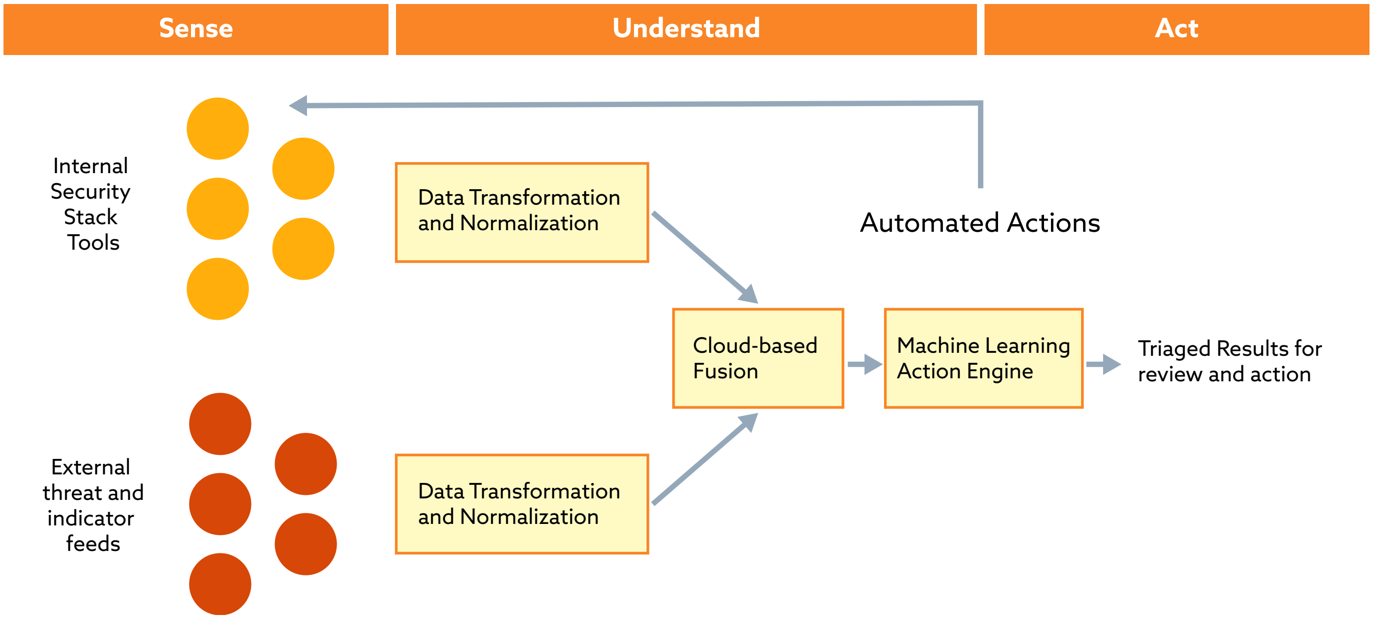
Sense
Understand
Act
Internal Security Stack Tools
External threat and indicator feeds
Data Transformation and Normalization
Data Transformation and Normalization
Cloud-based Fusion
Machine Learning Action Engine
Automated Actions
Triaged Results for review and action

Figure 3

What is the difference between process workflow and data workflow?

The concept of automated process workflows is not new, but data flow is new to cybersecurity. SOAR technologies and the adoption of orchestration have fundamentally changed the way we think about cybersecurity. No longer must the process workflow be manual with the application of playbooks. This helps SOCs make material progress in improving their organization's cybersecurity posture while balancing a shortage of qualified and experienced analysts. However, orchestration playbooks are not designed to manage your intel. Using them as a data workflow engine to solve complicated data problems such as extraction, normalization, whitelisting, and prioritization is not scalable.

For example, SOC analysts have to organize and catalog constantly changing cyber intelligence from different sources to operationalize for use in orchestration tools. Alert generation and triage, event prioritization, incident investigation, and remediation are distinct tasks for a SOC analyst, each requiring its own depth of intelligence context, enrichment, and structure. The need to account for new use cases can drive additional data enablement requirements from source to consumption.

Enterprise SOCs need capabilities that cover data needs for analysts to work through a spectrum of alerts and incidents, ranging from network intrusions to insider and privilege abuse, to point- of- sale intrusions and credential stuffing, all the way to identity breaches. Without accessible and integrated cyber intelligence, analysts are stuck manually collecting, curating, and cataloging intelligence and end up spending very little on triaging alerts, analyzing incidents, and threat hunting. Automated data workflows will become a strategic imperative as SOCs move away from self-service and ad-hoc efforts to accelerate their workflows.[8]

The challenge becomes data flow to support SOAR tools. As noted in Section 1, data is messier with the sprawl of security tools. Implementing workflow process automation and orchestration tools before your data is transformed and normalized can dilute, rather than enhance, SOAR tools, and SOC operations' efficacy. Simply stated, too often security leaders can't get their security tools to work because of data workflow challenges. As teams take on complex orchestration challenges, they must automate how to ingest, combine, curate, and deliver data needed by various tools. Automated data workflows become a crucial enabler for SOCs to realize the full potential of SOAR. In the absence of automated data workflows, analysts end up bridging data gaps with manual copy-paste operations, making them glorified data entry operators or hacking orchestration playbooks to perform complex data manipulations. Automated data workflows will become a strategic imperative as SOCs move away from self-service and ad-hoc efforts to accelerate their workflows.[9]

Section III. Building a Cloud-based Secure, Intelligent Ecosystem

Building on the previous sections, we can envision a cloud-based ecosystem that enables organizations to build a cyber "memory," and between organizations. The cloud allows organizations to bring the output from sensors (internal tools and external intelligence) together. Cloud-based computational software capabilities help organizations understand (transform, normalize, correlate, and enrich data). Machine Learning (ML) helps triage and automate actions by efficiently identifying event-patterns within organizations and between organizations. These capabilities can be applied on-premise, but far less efficiently and without the flexibility to work with other companies or sharing organizations.

Enabling Machine Learning

A focus on the data workflows enables the more efficient application of ML. Companies that have invested in multiple internal and external sources would want to "squeeze" as much information as possible to enhance their ability to detect new threats. Software tools that act as a data hub for historical intel should be investing in ML to improve the automation of detection and response tools.

[8] Shimon Modi, Why Automated Workflows are a Foundational Capability for Enterprise SOCS, February 21, 2020

[9] Shimon Modi, Why Automated Workflows are a Foundational Capability for Enterprise SOCS, February 21, 2020

ML can be helpful on multiple fronts:

- ML can help elevate the discussion around how we reason about errors. Concepts such as precision and recall will become more central to the efficient operation of the SOC.
  - ***Precision*** is a metric representing the total number of true positive predictions divided by the total number of predictions. It assesses the model's ability to identify a true positive out of its total predictions. Precision is related to the number of false positives; the larger the number of **false positives**, the less precise you are. In the context of cybersecurity, precision represents your model's ability to precisely detect a type of threat.
  - ***Recall*** is a metric representing the total number of true positive predictions made from the total number of true positives present in your population. Recall is related to the total number of false negatives; the larger the number of **false negatives**, the smaller the recall.
- By borrowing from these ML concepts, you can assign precision and recall metrics to your sources to assess their performance. In addition to that, in the context of Mean Time to Detect (MTTD) and Meantime to Respond (MTTR), higher precision and recall can help reduce both these metrics. Indeed, the larger the number of false positives and the smaller the precision of your detection systems would lead to inefficient use of resources. Also, the larger the number of false negatives and the smaller the recall means that your systems are missing out on detecting certain threats, which increases MTTD.
- Clustering and entity resolution techniques can help you leverage threat indicators to label your generated alerts. These techniques can go beyond a simple direct string match and allow you to use more advanced features from your already labeled indicators of compromise and use them to classify your internal data. Investing in these capabilities can help substantially improve your ability to automate the triage of alerts.
- Once your internal data is labeled (SIEM & EDR Alerts, suspicious emails reported by users, etc.), you can go back to your historical dataset and train Machine Learning models to score these alerts in real-time. These ML models can be trained using input features different from the indicators of compromise used for labeling your historical data. For example, you could use input features that identify language patterns used in phishing emails. We would like to pinpoint a significant benefit to this approach in reducing errors performed by your detection system. Using different features to detect in real-time makes it harder for attackers to avoid your "detection traps." The reason is that you are not relying on the same approach used for labeling your data, a process that relies on direct matches with specific indicators such as IPs, domains, etc. Instead, you are leveraging complicated features used as inputs into a mathematical model. This makes it harder for the attacker to reverse engineer your detection approach and, as a consequence, improves your detection performance. For a more in-depth discussion around this topic, please refer to the following book: "Machine Learning & Security" (Chapter X).
- Missing data problems. Many intelligence sources, usually open-source intelligence (OSINT), are scoreless and can lead to creating a large number of false positives, hence deteriorating recall and MTTR. Machine Learning trained on historically trained indicators can be used to fill-in missing scores and reduce false positives.
- Well trained ML models can help compress information in real-time to predict a score for events and new indicators without the need to search large databases of indicators. This helps reduce service-level agreements (SLAs)of software tools and, as a consequence, reduce the MTTD and MTTR of enterprises.

Key elements of a secure intelligence ecosystem:

1. Security tools (sensors) and threat feeds must be “API-first,” machine-to-machine ready to enable disparate data sets for processing, fusion, and enrichment.
2. Identify key datasets to be processed, fused, and enriched. This step can be subdivided into four parts:
    a. Identify critical internal security stack tools generating alerts for review, such as SIEM and EDR systems.
    b. A list of whitelist terms that should be bypassed during data processing and fusion. (This list can be built over time too.)
    c. Identify essential workflow process tools, such as case management and orchestration platforms, that require data to support workflows.
    d. Identify critical external threat feeds that you want to correlate against your internal alerts. This should include proprietary feeds, OSINT, sharing organization threat feeds, and data from government CERTs.
3. Identify a fusion platform with the following capabilities:
    a. The ability to ingest data from security stack tools via integrations or API.
    b. The ability to ingest email (this capability is particularly important to address suspect phishing emails)
    c. A private repository-- a cyber memory--to absorb internal security stack data and indicators
    d. The ability to normalize and transform diverse data sets, including material in different formats
    e. The ability to correlate data from internal security stack tools with external threat feeds
    f. The ability to extract and correlate observables as well as customized artifacts.
    g. The ability to enrich internal and external data sets and bi-directionally update security stack tools, including SIEMs, EDR, case management, and orchestration platforms
    h. The ability to apply machine learning to identify patterns and use appropriate tags
    i. Appropriate access and security controls for data – multifactor authentication, administrative control privileges
    j. The ability to quickly redact and share data with other parties

Ecosystem Maturity Model

The chart below shows a breakdown of a maturity model. Layer two in the chart is most critical as this is where the data is “understood,” or fused, labeled, scored, triaged, and sent back to secure applications such as SIEMs, EDR, case management, or orchestration platforms for further action.

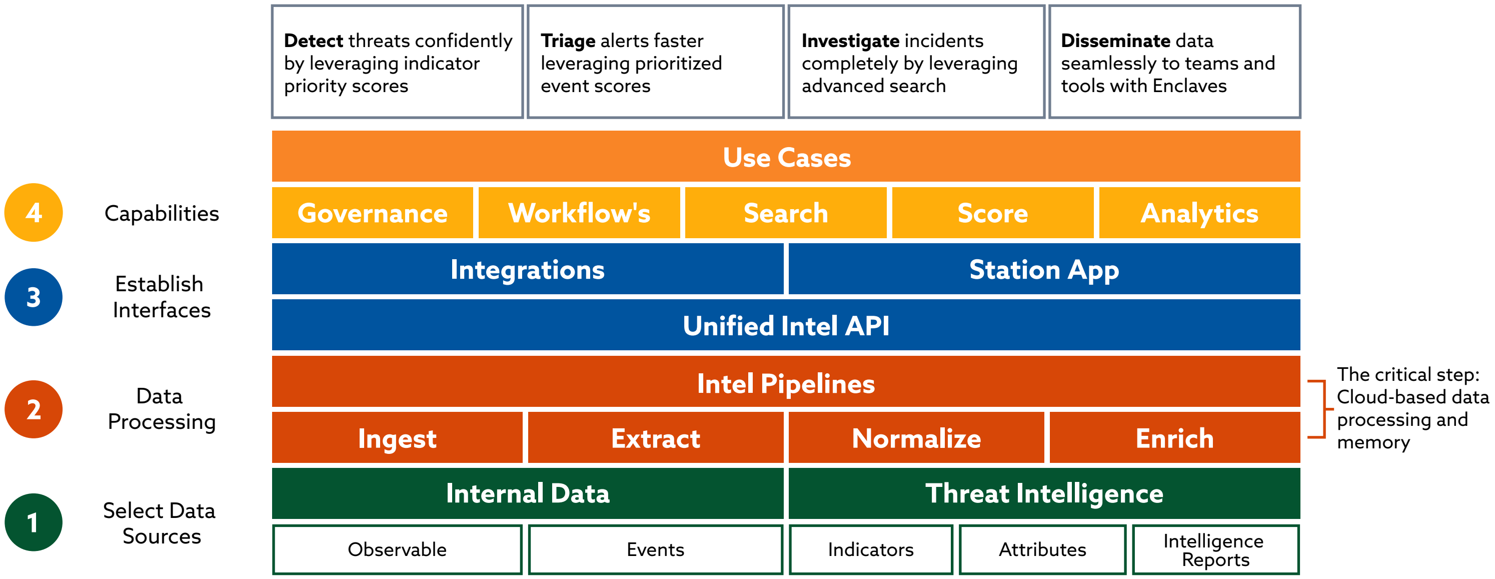
Detect threats confidently by leveraging indicator priority scores
Triage alerts faster leveraging prioritized event scores
Investigate incidents completely by leveraging advanced search
Disseminate data seamlessly to teams and tools with Enclaves
4
Capabilities
Use Cases
Governance
Workflow's
Search
Score
Analytics
3
Establish Interfaces
Integrations
Station App
Unified Intel API
2
Data Processing
Intel Pipelines
Ingest
Extract
Normalize
Enrich
The critical step: Cloud-based data processing and memory
1
Select Data Sources
Internal Data
Threat Intelligence
Observable
Events
Indicators
Attributes
Intelligence Reports

Figure 4: Ecosystem Maturity Model

Section IV. Security Business Analytics

The creation of secure, intelligent ecosystems provides senior leadership more insight into incident response metrics. The fractured cybersecurity market has hobbled the ability to capture such metrics given integration and automation challenges. Now key metrics can be captured and passed to leadership on their platform of choice. The example below shows some insights that can be captured and analyzed. Numerous API-ready business intelligence software tools exist today to query, visualize, and share results with others. Company executives and security operators alike can monitor MTTR and MTTD to understand trends.

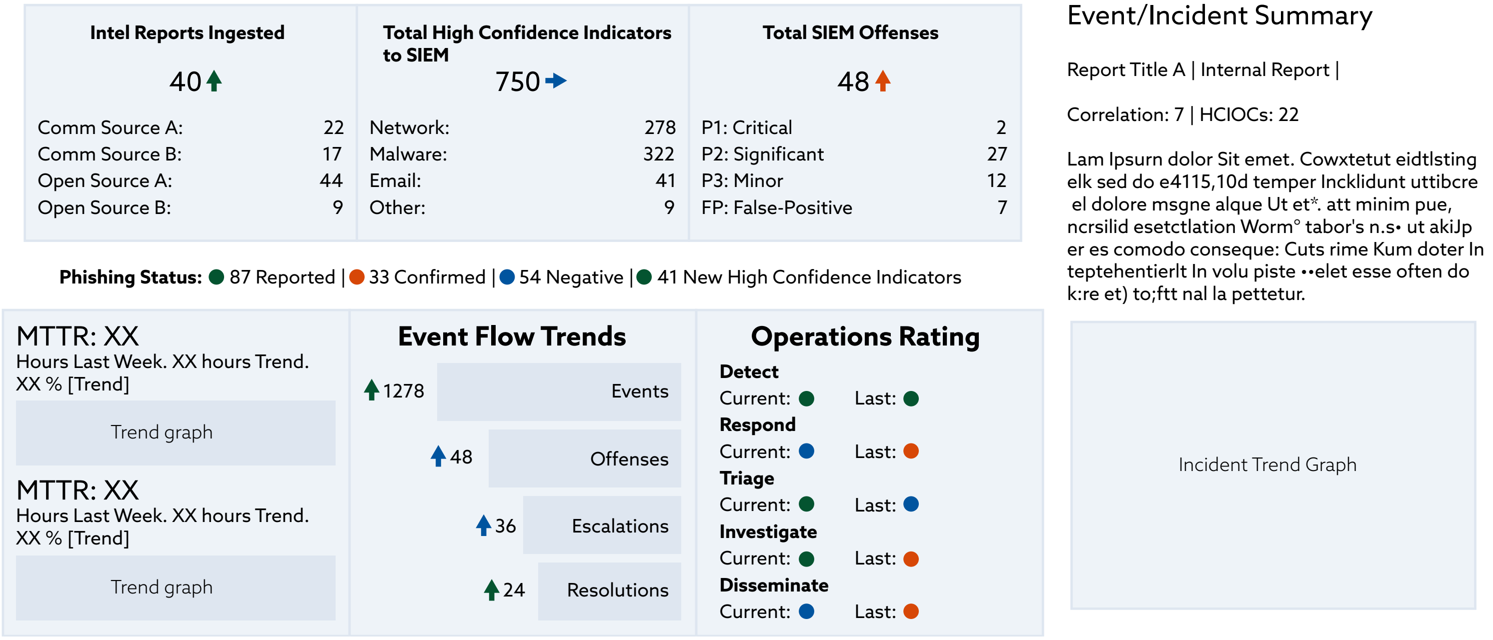

Intel Reports Ingested		**Total High Confidence Indicators to SIEM**		**Total SIEM Offenses**	
40 ↑		750 →		48 ↑	
Comm Source A:	22	Network:	278	P1: Critical	2
Comm Source B:	17	Malware:	322	P2: Significant	27
Open Source A:	44	Email:	41	P3: Minor	12
Open Source B:	9	Other:	9	FP: False-Positive	7

Phishing Status: ● 87 Reported | ● 33 Confirmed | ● 54 Negative | ● 41 New High Confidence Indicators

MTTR: XX
Hours Last Week. XX hours Trend.
XX % [Trend]

Trend graph

MTTR: XX
Hours Last Week. XX hours Trend.
XX % [Trend]

Trend graph

Event Flow Trends

↑ 1278	Events
↑ 48	Offenses
↑ 36	Escalations
↑ 24	Resolutions

Operations Rating

Detect		
	Current: ●	Last: ●
Respond		
	Current: ●	Last: ●
Triage		
	Current: ●	Last: ●
Investigate		
	Current: ●	Last: ●
Disseminate		
	Current: ●	Last: ●

Event/Incident Summary

Report Title A | Internal Report |

Correlation: 7 | HCIOCs: 22

Lam Ipsurn dolor Sit emet. Cowxtetut eidtlsting elk sed do e4115,10d temper Incklidunt uttibcre el dolore msgne alque Ut et*. att minim pue, ncrsilid esetctlation Worm° tabor's n.s• ut akiJp er es comodo conseque: Cuts rime Kum doter In teptehentierlt In volu piste ••elet esse often do k:re et) to;ftt nal la pettetur.

Incident Trend Graph

Figure 5: Weekly Dashboard

Section V. Areas for Further Research

Secure Intelligent Ecosystems seek to secure enterprises, which remain high value targets for attackers. As technologies evolve, the manner in which we are able to adapt to the many systems that come into existence will increase our ability to secure our cloud infrastructures as well as gather more actionable and tangible intelligence. Much of this research includes, but is not limited to;

- Additional analysis of the "sense, understand, act" classification model, including systems modeling applications. The use of the "sense, understand, act" model in areas, such as IoT, can be explored to better understand complexities in enterprise ecosystems. Data analysis between each stage of sense to understand, and understand to act will help the enterprise ingest, react, and respond with higher accuracy in more complex systems.
- The use of Blockchain applications in the gathering of threat intelligence as well as the use for securing cloud ecosystems. A strategy for mitigation against corruption and attacks is possible through the thoughtful use of blockchain. The data feeds that ultimately comprise the threat intelligence and security tooling should be protected to provide the greatest level of trust possible. By using the immutable logging inherent within distributed ledger technology, it should be possible to create an external source of truth to reduce the risk of degradation and errors within key systems.
- Leveraging Machine Learning (ML) and Artificial Intelligence (AI) for automation within intelligent ecosystems. ML and AI uses can also impact these systems negatively by corrupting the data set during the training process or creating evasion attacks on production systems generating both false positives and negatives. These and other use cases will need to be explored to properly integrate these techniques and better understand the sophistication of such attacks.
- Anonymizing trend analysis based on ML applications. The proliferation of security tools and systems will need to be anonymized into a unified field of view for the enterprise. The promises of ML can broaden the view of each enterprise security tool eliminating the single-pane view as services, tools, and solutions in the enterprise grow.

Conclusion

It is time for a paradigm shift. When describing the development of a chess-playing algorithm in SuperIntelligence, Nick Bostrom said it was thought that for a computer to play at the grandmaster level, it would "have to be endowed with a high degree of general intelligence." It turned out to be achievable through a "surprisingly simple" algorithm. Bostrom states:

> The fact that the best performance at one time is attained through a complicated mechanism does not mean that no simple mechanism could do the job as well or better. It might merely be that nobody found a simpler alternative. The Ptolemaic system (with the Earth in the center, orbited by the Sun, the Moon, planets, and stars) represented state- of- the- art astronomy for over a thousand years. Its predictive accuracy was improved over the centuries by progressively complicating the model adding epicycles upon epicycles to the postulated

> celestial motions.Eventually, the entire system was overthrown by the heliocentric theory of Copernicus, which was more straightforward and, though only after further elaboration by Kepler, more predictively accurate.[10]

Perhaps in cybersecurity, we are exiting the Ptolemaic era and are now someplace between Copernicus and Kepler. There is a less complicated model available--sense, understand, act--that breaks down cybersecurity challenges more efficiently and understandably. Redefining intelligence and building secure, intelligent ecosystems helps us rethink how we see and utilize data from our existing tools and leverage machine learning. No longer should security operations revolve around tools alone, but instead, data becomes the constant gravitational force at the center of our effort to achieve more predictable and autonomous security.

[10] Nick Bostrom, Superintelligence, p. 36-37