SWIMLANE

Swimlane ARMOR

The Swimlane **Automation Readiness and Maturity of Orchestrated Resources (ARMOR)** framework is the industry's first automation-centric maturity matrix.

→ Security automation has made the shift from "nice to have" to "need to have". Unlike threat detection, the security industry does not have established frameworks for security automation best practices. As a result, many organizations are overwhelmed by the people, process and technology considerations that go into automating security processes. This leaves security leaders asking questions like:

1. Is my organization ready for security automation?
2. What automation use cases achieve the fastest time to value for organizations like mine?
3. How do I compare with peers in my industry, location or company size?
4. How do I know if my investment in automation is optimized for maximum ROI?

Overview

Security automation has become a cornerstone technology that enables modern security operations (SecOps) teams to keep pace with the ever-growing demands they receive internally and externally. It is the only solution that can help security teams close the gap between their human capacity and the work that needs to be done. SecOps teams who adopt and mature their security automation programs can keep their organization secure while also meeting compliance and governance requirements.

Organizations have been leveraging automation for years. The ability to write scripts, macros, and develop homegrown applications has been part of security programs since the industry's inception. Now that automation is included in many enterprise security tools such as security information and event management (SIEM), extended detection and response (XDR), security data lakes, and purpose-built automation solutions such as robotic process automation (RPA), business process management (BPM), security orchestration automation and response (SOAR), and low-code security automation, security teams need to look for ways to measure and improve their automation outcomes. To accomplish this, security leaders need to understand if and how automation enables their teams to be more effective at achieving business and security goals.

Swimlane has been helping organizations of all industries and sizes mature their security automation programs for nearly a decade. We've harnessed our company-wide institutional knowledge and customer best practices to create the **Automation Readiness and Maturity of Orchestrated Resources (ARMOR) framework**. ARMOR is designed to deliver actionable insights for Security Operations Center (SOC) teams and any security function that currently or plans to leverage automation. Throughout the rest of this white paper, you'll learn:

- The **key metrics used to benchmark** and measure automation maturity according to the ARMOR framework.
- How the ARMOR framework has been successfully **applied to use cases** such as insider threat, data loss prevention (DLP), fraud, developer security and operations (DevSecOps), vulnerability management, application security, extended detection and response (XDR), and others.
- **Recommendations for how to leverage** the ARMOR framework in your organization.

Automation Readiness & Maturity

The ARMOR framework consists of two core components, a maturity matrix and a readiness assessment. The security industry has many threat detection and incident response (TDIR) centric maturity models like MITRE Adversarial Tactics, Techniques, and Common Knowledge (ATT&CK), Cybersecurity Maturity Model Certification (CMMC), National Institute of Standards and Technology (NIST), Capability Maturity Model Integration (CMMI) and Cybersecurity Capability Maturity Model (C2M2), but still there isn't a well-defined path for how automation should be applied in order to operationalize these approaches.

This perpetual complexity problem is what ARMOR aims to solve. The matrix aligns automation objectives, goals, tactics, and use cases to these industry standards. The assessment is a resource for beginners to evaluate their automation readiness and for more advanced organizations to benchmark their maturity. It is a straightforward multiple-choice self-assessment that allows you to understand how well-prepared your organization is to implement an automation program. The assessment asks a series of 20 questions in order to evaluate the maturity of the people, process, and technology resources needed for successful security automation. Most people can complete the assessment in less than 5 minutes.

Orchestrated Resources

The matrix uses several data points to measure the efficacy of technology capabilities and controls are applied to achieve their objectives and goals. These data points are used to map maturity levels to the five tiered matrix.

The ARMOR Matrix

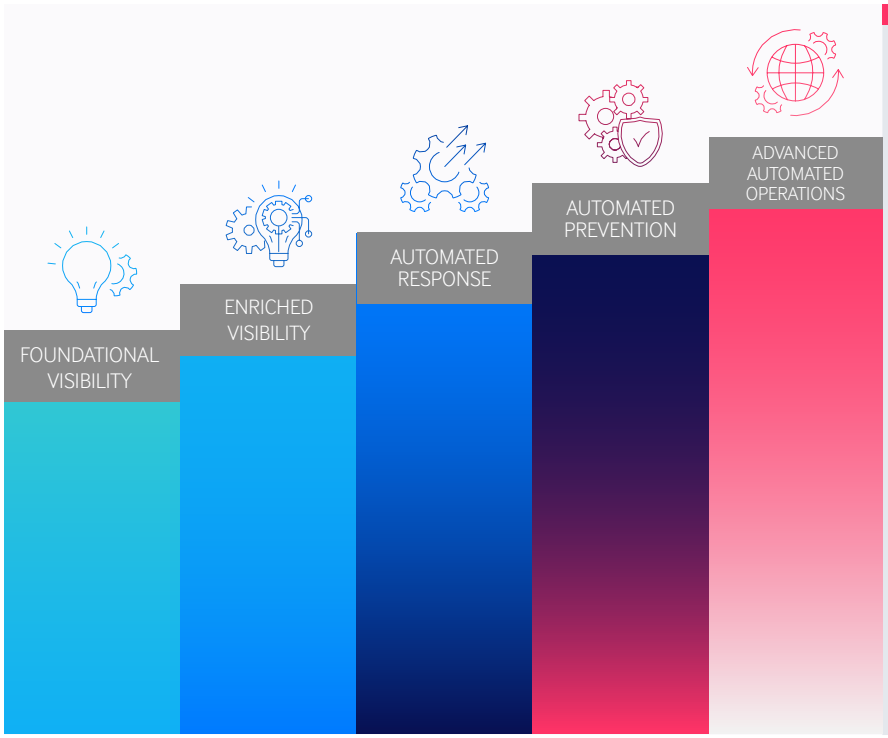

→ The ARMOR Matrix Data Points

- Automated versus manually created cases
- True vs. false positive alerts
- Percentage of enriched alerts
- Alerts mapped to MITRE ATT&CK
- Alerts mapped to MITRE D3FEND
- Alerts flagged for automation
- Alerts being automatically managed
- Manual cases that create an iteration

Automated Versus Manual Cases

At the highest level, security automation maturity can be thought of as how security teams leverage automation versus executing manual case management processes. A case is created when one or multiple alerts meet specific criteria that require the situation to be escalated. Once this happens, a team member assumes ownership of the case and executes processes for compliance and audit reasons.

The rationale for measuring the number of manual versus automatically created cases serves many purposes. The most foundational reason is that you can't improve what you can't measure. If cases are not being created at all, there is no method to learn, but if cases are being created automatically, they are much more likely to be created in the first place, they are created systematically versus ad-hoc, and the likelihood of them being missed due to being over capacity or human error is greatly diminished. The consistency in how they are created provides a foundation in which teams can report, visualize, and draw insights into many aspects of cases that are important for teams, including:

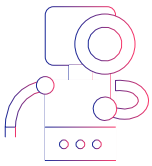

What are the sources of my cases?

- Are they sourced from threat detection alerts like EDR, NDR, CWP, SIEM, or other threat detection tools?
- Are they coming from active programs such as threat hunting or passive activities such as alert triage, SOC, and IR?
- Does the security team identify the source of the case, or is it being reported by system and application owners, IT, or users?

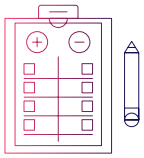

What types of cases are we managing?

- Are the cases we manage primarily true or false positives?
- Are we managing high severity or low/informational severity activities?
- Do cases require additional contextual information to process or not?
- Are my cases unique first-time cases, repeated, or correlated with other historical cases?
- Do my cases have related knowledge cases, or are they being worked on for the first time?

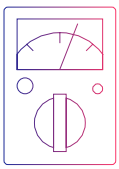

What measures are we capturing, and what are our baselines?

- How long does it take from detection to the time cases are reviewed (for manual cases)?
- How many cases are leveraging secondary or tertiary enrichment sources?
- How many cases and which types have available remediations or mitigations?
- How long does it take to put mitigations in place?
- What are my most commonly leveraged mitigations?

The ability to mature a security automation program directly correlates to your ability to measure what is happening repeatedly and systematically. The historical catalog and recording of what is typically hundreds of thousands of discrete data points over weeks, months, quarters, and years allow teams to have real insight and visibility into the ongoing state and changes of their operational capacity and capability, providing what, for many teams will be their first true insight into where and how they are spending their time and what they could do to improve the operational efficacy.

Swimlane ARMOR Best Practice

In your case management system, identify or add a tag, field, or other mechanisms that allows each case to be identified by its source as manually created or automatically created. After 30 days (or a month), calculate your manual versus automatic case creation percentage by taking the total number of automatically created cases divided by all the cases created during the same period.

True Versus False Positive Alerts

The ability to measure true and false positive rates as part of the alert management process is one of the most foundational ways to understand better if teams are working on the right activities and where they are having their time wasted. Teams are regularly understaffed for the amount of work required to meet best practices, so the idea of working on things that are not meaningfully improving the organization's security posture is a waste of time, money, and a huge morale burden for any security leader. The goal of measuring true and false positive rates is twofold, first, to ensure you are not wasting the time of your team, second is to understand where tools could improve from tuning or replacement.

Developing a true/false positive baseline, measuring it regularly, and putting programs in place to improve it over time will reduce the signal-to-noise ratio teams struggle with daily, and they will build a better relationship between security ops and security engineering teams because there are valuable feedbacks loops and actionable improvements that can be put in place to improve security and limited SecOps resources.

Swimlane ARMOR Best Practice

In your alert management system, which should be managed differently than your case management system, identify or add a tag or other mechanism to determine if the alert was a true positive or false positive. Once the measurement has been in place for 30 days (or a month), calculate your false positive rate, the total number of false positive alerts received in the period, divided by the total number of alerts. This percentage is now your baseline false positive rate.

Enriched Alerts

Measuring alert enrichment is a crucial element to understanding if you are using the breadth and depth of your security tools and telemetry to their fullest extent. The process of enriching alerts is extremely common in the processing and analysis of security alerts. Enrichment is the process of adding additional context based on related information to an alert not just to help make a decision but help make a better decision on if the alert being managed is malicious, its potential severity, and its impact on the organization, which teams if any should be involved, what procedures to follow and what limitations should be put in place.

Enrichment of alerts can take many sources, and your enrichment process can mature over time. Enrichment of alerts, in many cases, takes the form of looking up Indicators of Compromise (IOC) from alerts in secondary sources such as threat intelligence data sets, looking up targeted users or assets in systems such as a configuration management database (CMDB) or directory services, or searching vulnerability data associated with a targeted host to understand its vulnerability posture which could directly impact the severity and risk of a particular alert.

Measuring the enrichment process provides organizations with many interesting insights, including which alerts require additional context to mitigate, which enrichment sources are most commonly utilized, which sources lack the necessary data to enable analysis and many others.

Swimlane ARMOR Best Practice

Identify or create a tag or tracking method in your alert management system to understand if an alert was enriched and from which source. Once the measurement has been in place for 30 days (or a month), calculate your alert enrichment rate, which is the total number of alerts received that had some level of enrichment in the period, divided by the total number of alerts. This percentage is now your baseline alert enrichment rate.

Alerts Mapped to MITRE ATT&CK

The MITRE ATT&CK framework is a broadly adopted framework for understanding and mapping the Tactics, Techniques, and Procedures (TTP) used by threat actors to plan, initiate, and operate an attack on a target. The power of the MITRE ATT&CK framework is in the ability to have a standard method for understanding what visibility and monitoring capabilities are required to monitor and detect which types of attack TTPs. Tracking all the alerts that a team is managing will both help identify the most commonly used methods to target an organization but, over time, can help an organization understand where they may have security visibility issues and if their security program is effective at proactively identifying attackers and their techniques early in the attack lifecycle leveraging reconnaissance, resource development, and initial access techniques, before having a more established footprint and ability to inflict damage and loss with execution through impact techniques.

The number of threat detection tools that are providing MITRE ATT&CK tagging to their alerting is continuing to grow, but not all alert sources have embedded mappings, so implementing a method to ensure more and more of your alerts are mapped to MITRE ATT&CK is a crucial element of maturing your security visibility and understanding.

Swimlane ARMOR Best Practice

Identify or create a tag or tracking method in your alert management system to understand if an alert is mapped to one or more MITRE ATT&CK TTP IDs or "T Codes." (e.g.T1059.003 Windows Command Shell). Once the measurement has been in place for 30 days (or a month), calculate your MITRE ATT&CK mapping rate, which is the total number of alerts received that were mapped to one or MITRE ATT&CK TTP IDs in the period, divided by the total number of alerts. This percentage is now your percentage of alerts mapped to MITRE ATT&CK.

Alerts Mapped to MITRE D3FEND

While not as well known as MITRE ATT&CK, the complementary framework to MITRE ATT&CK is the defensive framework MITRE D3FEND (Detection, Denial, and Disruption Framework Empowering Network Defense) which is designed to enable teams to understand and implement security mechanisms to thwart, isolate, and remove an adversary from their environment. The framework provides teams with an extensive list of mechanisms to deter attackers that can be directly tied to the techniques found in MITRE ATT&CK. That ability to tie TTPs to defensive techniques found in MITRE D3FENSE allows security teams to understand better how they can leverage existing tools to deter, detect, and remove attackers, but also where they might have gaps in their ability to prevent a successful attack. An automated response system for security is only as strong as the mechanisms available to enact changes and take action.

The number of security tools categorizing their response and mitigation mechanisms to MITRE D3FEND is pretty limited, so leveraging a centralized system of record for tracking and managing alerts is critical to start applying tags and categorizing automated or manual responses to MITRE D3FEND.

Swimlane ARMOR Best Practice

Identify or create a tag or tracking method in your alert management system to understand if an alert is mapped to one or more MITRE D3FEND Defense Techniques (e.g. D3-CAA - Connection Attempt Analysis). Once the measurement has been in place for 30 days (or a month), calculate your MITRE D3FEND mapping rate, which is the total number of alerts response taken that were mapped to one or MITRE D3FEND Techniques in the period, divided by the total number of alert responses. This percentage is now your percentage of alerts mapped to MITRE D3FEND.

Alerts Flagged for Automation

At the heart of the ARMOR model is the ability to take automated responses. Building an enriched view of your security alert pipeline, ensuring it is at an acceptable level of false positive and false negative rate, and ensuring it is mapped to key frameworks puts a team in place to begin taking automated action to both thwart attacks and reduce the amount of work a team has to do manually. Many of the key metrics such as mean-time-to-respond (MTTR), mean-time-to-mitigate (MTTM), mean-time-to-investigate (MTTI), dwell time, investigation time, average analysts time per alert, case, or incident are directly tied the ability to automate parts of the analysis process and ultimately take an automated response.

The problem is that teams never just start automating these things out of the gate. There is a method and time when they need to understand how their automation system will react in a variety of scenarios and with varying data sets, ultimately they need to learn to trust the system. The measuring of the percentage of alerts that are flagged for automation gives teams a chance to implement a crawl, walk, run methodology without moving too hastily in the beginning and causing unwanted pain such as outages or disrupted business operations, but also gives them a mechanism to measure progress when full automation isn't yet an option.

Flagging alerts for automation is the way a team, early in their maturity, can see if they are getting enough fidelity in the alerts they are receiving, if the enrichment process is giving them enough context to make decisions, if the centralized data set is robust enough to determine the proper automation mechanism and all of that it is adjusted for both risk reduction but also potential business impact such as inadvertently disabling a VIP user, disabling business-critical functions, or taking revenue-generating services offline.

Swimlane ARMOR Best Practice

Identify or create a tag or tracking method in your alert management system to understand if an alert is mapped to one or more MITRE D3FEND Defense Techniques (e.g. D3-CAA - Connection Attempt Analysis). Once the measurement has been in place for 30 days (or a month), calculate your MITRE D3FEND mapping rate, which is the total number of alerts response taken that were mapped to one or MITRE D3FEND Techniques in the period, divided by the total number of alert responses. This percentage is now your percentage of alerts mapped to MITRE D3FEND.

Alerts Being Automatically Managed

The goal of a security team in pursuing SecOps automation and maturing their program is to be able to provide an end-to-end capability that can ingest, normalize, enrich, analyze, and take action with as little human intervention as possible, reducing the likelihood of human error, reducing the level of effort to manage telemetry and intelligence inflows, and ultimately reduce the likelihood of a successful attack. Beyond just identifying if an alert is automatable, the goal is to have the automation happen consistently and predictably, allowing teams to not "stare at the road" or, in the case of a security analyst, look at a dozen or more different security solutions in a myriad of browser tabs.

There will never be a future where automation is 100% responsible for every security alert. The techniques of adversaries change, the information provided by detection tools varies, and the ever-changing nature of technology will require human intervention and oversight of some percentage of security alerts. Still, the goal is to drive that number down to a manageable number that allows enough time for the team to think thoughtfully about what they should do with alerts that are not automated. Increasing your percentage of automated alerts will free up time and improve your team's morale.

Swimlane ARMOR Best Practice

Identify or create a tag or tracking method in your alert management system to understand if an alert was partially or fully automated. Once the measurement has been in place for 30 days (or a month), calculate the percentage of automated alerts, your automation rate, which is the total number of alerts fully automated in the period, divided by the total number of alerts. This percentage is now your automation rate. You can use the number of partially automated alerts and alerts flagged for automation to track improvements when your core automation rate might not be changing.

Manual Cases That Create an Iteration

When people discuss the advantages of automation, they generally gravitate towards the benefits of reduced time, reduced effort, reducing in critical metrics like MTTR, dwell time, and time spent per alert or case, but the hidden superpower of automation is that you are automating a tracking system that can be the most powerful tool for providing actionable and specific insights into the improvement of your security program. In an era of machine learning (ML) and artificial intelligence (AI), reinforcement learning is top of mind.

Before you get to AI or ML it's critical to first look at alerts that were false positives, were not enriched, didn't have the information required to take automated action, and were not mapped to the MITRE frameworks. Identifying these alerts will highlight the specific instances where engineering should tune their tools and orchestration. Once these visibility gaps are identified, security teams can validate the efficacy of their investigation or intelligence sources.

Each of these insights provides a security team an advantage they've never had, a structured and systematic method for making decisions on future efforts and changes can have the largest and most positive impact on their posture.

For every manual case in your system, you are probably seeing an alert that was reported by something or someone other than your security detection tools. That might be a user, a system administrator, or an auditor. Each manual or partially automated alert is a chance to initiate an iteration cycle. An iteration cycle is learning about your SecOps pipeline, taking that new knowledge, and feeding it back into the system to improve visibility, automation, and response and ultimately reducing risk. While the goal of a SecOps automation program is to do automation, a lot of the value in the program comes from the available time it provides to investigate the things that are not fully automated.

Swimlane ARMOR Best Practice

In your alert management system, identify or create a tag or tracking method to understand which alerts were partially automated or manually worked and if, as part of working that alert, there was feedback provided or a change to your security apparatus that might enable it to be fully automated in the future. It doesn't matter if the alert or case is or is not fully automated the next time, the tracking is focused on did it drive a change that is moving the organization toward a more automated future. Once the measurement has been in place for 30 days (or a month), calculate the percentage of alerts that did provide for a feedback iteration and divide the total number of partially automated for manually worked alerts. This percentage is now your improvement percentage.

Summary

Significant program improvements don't happen overnight. They are about aligning a team's effort with the proper measures and activities that will improve those measurements. You can't improve what you can measure, and having visibility and understanding of where you are and where you want to go is the first step in becoming a better SecOps team. The ARMOR framework provides structure, guidance, and best practices for tracking, measuring and improving automation outcomes over time. When organizations leverage the Swimlane ARMOR framework, they can realize the value of automation dramatically faster than if they build a security automation program on their own.

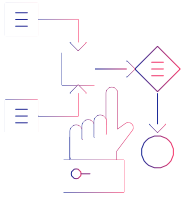

Level Up Today

Whether you're just getting started with security automation or maturing an established program, we are ready to help you level up. Visit swimlane.com/armor to take the assessment and get your own personalized maturity report.

SWIMLANE

Corporate Headquarters
363 Centennial Pkwy Suite 210
Louisville, CO 80027
1-844-SWIMLANE
swimlane.com

Swimlane is the leader in low-code security automation. The Swimlane Turbine platform unifies security operations in-and-beyond the SOC into a single system of record that helps reduce process and data fatigue, while helping security leaders overcome chronic staffing shortages and more easily quantify business value and the efficacy of security operations.