GUTSY

Process Mining: The Security Angle

Upgrading IT Security With Process Mining

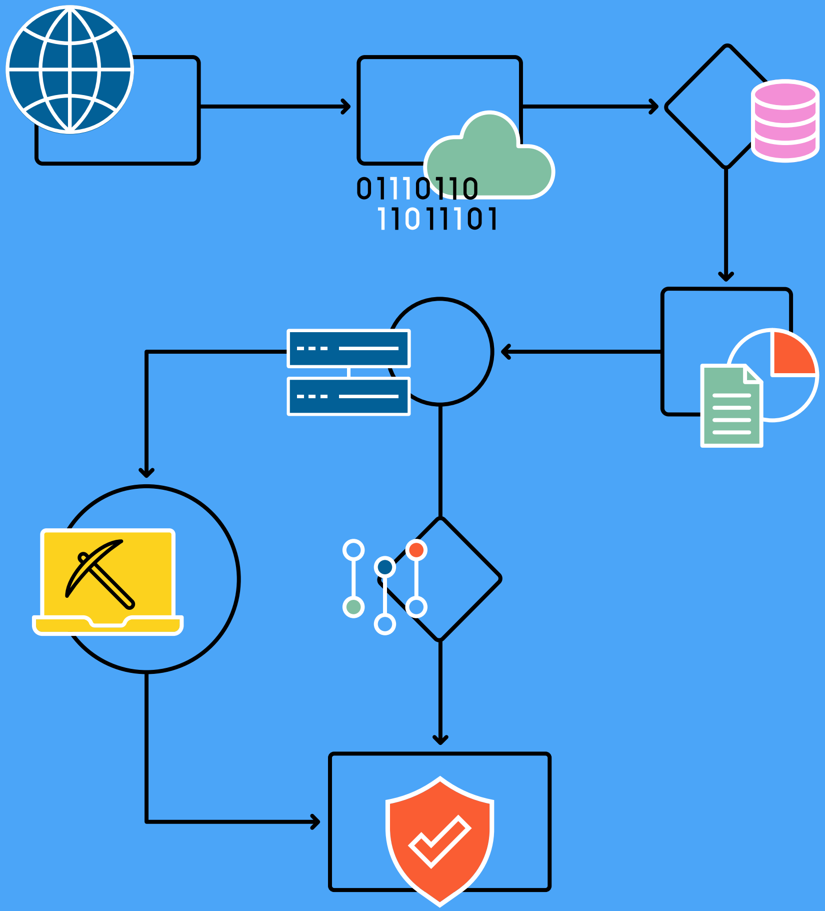

May 2024

AQSA TAYLOR

About this Book

“Process Mining: The Security Angle” is your introduction to the application of process mining techniques to your cyber security practice. Our premise is that security operations can be broken down into sequences of steps that form distinct processes, which can be designed, managed, and optimized.

Processes form the backbone of everything we do.
Processes establish routines and are repeated every day. For example, there’s a process for remediating a software vulnerability, a process for updating a container image, and a process for offboarding staff. As security professionals, everything we do should adhere to a prescribed process. Having a better understanding of your processes will help you keep an eye on your goals and ensure you can consistently reach them.

Process mining in itself is not a new concept. It simply refers to the practice of extracting data from multiple sources and arranging the data in a way that is conducive to the holistic study of processes. This book explores the importance of process mining from a security perspective. It assumes the ongoing need to continually improve security operations. The book provides real-world examples to demonstrate why process mining is instrumental to security efforts and why process mining holds the promise of outperforming other security improvement approaches. Read on to learn how you can use data-driven visibility to optimize your processes.

About Gutsy

Gutsy is a data driven security governance platform that helps you understand how your security teams, tools, and services work together, so you can lower risk, accelerate auditing and compliance, and drive accountability. Gutsy connects to all your existing systems, automatically correlates how they work together, and identifies inconsistencies that lead to risk and noncompliance.

Visit www.gutsy.com to learn more.

GUTSY

About the Author

AQSA TAYLOR
DIRECTOR OF PRODUCT MANAGEMENT

Aqsa Taylor is Director of Product Management at Gutsy, a cybersecurity startup specializing in process mining for security operations. A specialist in cloud security, Aqsa was the first Solutions Engineer and Escalation Engineer at Twistlock, the pioneering container security vendor acquired by Palo Alto Networks for $410M in 2019.

At Palo Alto Networks, Aqsa served as the Product Line Manager responsible for introducing agentless workload security and generally integrating workload security into Prisma Cloud, Palo Alto Network's Cloud Native Application Protection Platform. Throughout her career Aqsa helped many enterprise organizations from diverse industry sectors, including 45% of Fortune 100 companies, improve their cloud security outlook.

Aqsa holds an MSc in Electrical and Computer Science Engineering from the University of Texas at San Antonio where she authored "Securing cloud containers using Quantum networking channels" (IEEE Conference Publication). Outside of work, Aqsa enjoys reading and writing poetry, exploring new places, and spending time with her mom, sister and husband, who have been her unwavering supporters throughout.

Target Audience

What's your interest in process mining? This book is intended for leaders in a cyber security practice in a decent-size enterprise in any industry who are determined to improve their company's security.

The book shares insights about the power of process mining and explains its potential for leveling up governance and cyber security execution management.

APPLY THE THEORY

By the end of this book , you can expect to understand:

- What process mining is and what it can achieve
- The need for process mining in cyber security
- Application of insights from process mining tools
- The importance of conformance and risk analysis
- How process mining aids audits and overall governance strategy
- What to look for in a security focused process mining tool
- How to start implementing process mining in your organization

What's new in this version

This version introduces a new chapter (Chapter 3) on security governance and how governance can be transformed when process mining is applied. The chapter dives deeper into the definition of security governance, core principles and practical examples that demonstrate advantages from application of process mining.

CHAPTER 3: SECURITY GOVERNANCE

What is Governance

Govern is the core of cybersecurity framework

- Core Requirement
- Challenges
- Transforming with Process Mining

Government mandates on Governance

- Core Requirement
- Challenges
- Transforming with Process Mining

Audits and Certifications

- How it works
- Transformed Auditing

Transformed Governance

- Data driven
- Measured
- Continuous
- Actionable

Book Outline

Book Outline

01

Applying Process Mining to Security

CHAPTER 1

From the moment we wake up, to the choices we make, to the routines we follow, almost everything we do is part of a ‘process’: “the way something is done.” In security operations, a process is defined as a sequence of events that produce a certain outcome.

Consider the simple process of getting to work – for me, I grab a coffee and head into my home office. While for another, it may be to get ready and drive to their office across town. Although the objective remains the same, getting to work, “how” that objective is achieved, is quite different.

Organizations, like individuals, follow processes to conduct their operations, and these may differ widely across organizations. Every company tends to follow its own set of procedures, with limited overlap or common practice shared across organizations. Moreover, organizational processes tend to be complex, involving people and equipment working together to complete distinct tasks that together deliver larger organizational outcomes.

For IT security teams, the overall goal is to ensure the organization’s data and systems are safe and secure. To achieve this goal, security teams must execute and properly coordinate many complex processes, such as, the regular removal of accounts and access of staff members who have left the organization, the detection and remediation of vulnerabilities and various incident response activities. While the particular steps each organization takes to carry out these processes may vary, the cadence of actions across different people, technologies, and services to achieve outcomes is common to all organizations.

While countless books and conference presentations on security technologies exist, there’s been relatively little focus on the processes they’re part of. After all, technologies alone are just tools. Unless properly integrated with people and other technologies, stand-alone tools cannot deliver the outcomes security teams need. This book introduces a new approach to security governance by leveraging process mining techniques.

In this chapter you will learn the definition of process mining and, more importantly, how it can be expected to optimize security outcomes.

Good Security is Built on Good Processes

Much of the current thought leadership on cyber security focuses on the technical capabilities of security tools. The industry buzz around "EDR", "CSPM", or "CNAPP" all too often emphasizes the tools and technologies over the overall processes they are part of. As a result, the burden of designing and optimizing processes that coordinate and synchronize people, tools, and technologies to deliver security outcomes falls entirely on the customer.

CISOs and security leaders are well aware of this conundrum and the sad reality that acquiring all the right tools doesn't always guarantee reliable security. A survey from IDG and ReliaQuest reveals that on average, enterprises maintain 19 different security tools, of which only 22% are vital to their organization's primary security objectives. The same survey, highlighted in Fig 1.1, shows that 85% of security decision makers believe they are adding technologies faster than they can productively use them, with 71% admitting most of their existing tools are underutilized.

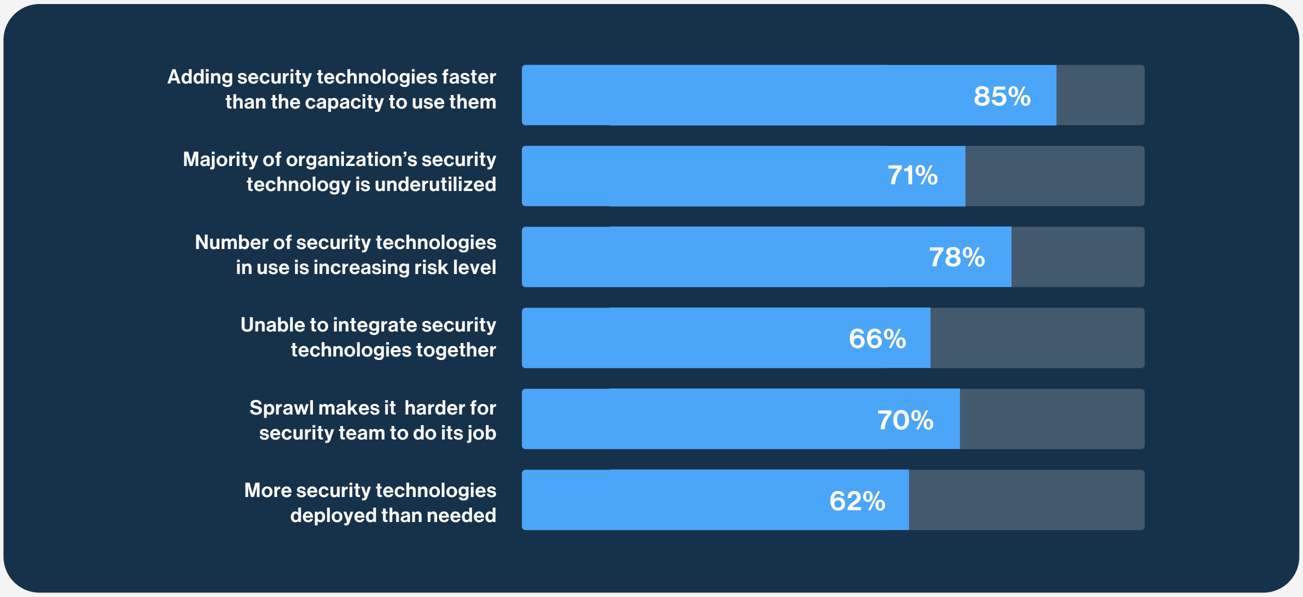

FIG 1.1 SECURITY TOOLS: VOLUME AND UTILIZATION IN TODAY'S ENTERPRISE [Source: IDG/ReliaQuest]

Despite spending countless dollars on tools, organizations still struggle to achieve desired security objectives. In reality, a lack of tools has rarely been the problem. In fact, it is more likely to have tool sprawl and alert fatigue caused by having too many tools "deployed" yet not fully operationalized. Further complicating matters are all of the non-security technologies that are key parts of security workflows today.

For example, your HR software is probably a key part of your identity management processes even though it wouldn't be considered a security software. Is it still required to maintain records of offboarded employees, their offboarding date and other information which is relevant when removing access of those employees. For those processes to work well, the security team has to use data from this software in tandem with your directory services and federation services, to maintain a secure offboarding process.

Historically, it's been challenging for security leaders to understand how these systems and teams can effectively work with each other, to identify problems, and to continuously improve on them.

Common Misconceptions

The excessive focus on tools and the disregard for the importance of the underlying processes leads to damage that often goes unnoticed. Let's examine some common misconceptions about security and their consequences.

"My security tools provide full visibility"

Security processes usually involve multiple security and non-security tools. Buying a security tool is much easier than using it effectively. Normally, security tools produce alerts on incidents, but they do not show how these incidents were managed or what caused them.

Results: Siloed views of what's really happening

Each tool provides just a partial picture of what's going on within its own narrow context. The more tools, the more difficult it is to grasp a holistic picture and understand how the tools work together and with your teams.

"We have a documented process for that"

A prevailing false assumption is that an organization's "official" processes are followed consistently. In reality, every team, even every team member, may contribute to variation in the process, which may lead to inefficiency and risk. For example, if a ticket is assigned to the wrong team it will need to be manually reassigned. If the user accounts of terminated staff members are not fully deleted, this introduces risk. Despite everyone's best efforts, the consequences of even simple mistakes, human errors, or shortcuts may be magnified, even when the tools are working as they are designed to.

Results: Inconsistent outcomes

There may not be a right or wrong way but it's far from ideal to have many different ways to accomplish a security task. Variations lead to wasted time, wasted investment, increased security risks, and escalating frustration.

"I already get a lot of metrics from my security tools"

Security teams depend on technical metrics from security tools to measure performance. However, each security tool in the stack only provides information about its own results, such as how many threats it has detected or how many vulnerabilities were detected. These are good data points, but they don't connect to what happened after the notification was sent.

Results: Not enough context for real systemic improvements

Technical metrics alone cannot tell us what improvements need to be made in an overall process. Identifying how one step influences other steps is crucial when evaluating the effectiveness of security processes.

"Our periodic auditing is good enough"

For many organizations, the closest thing to understanding process effectiveness today is periodic security audits, but even this doesn't tie all operations together. Instead, it focuses more closely on compliance than conformance. Further, such audits usually involve expensive external consultants and rely on small samples of data and personal anecdotes. Most importantly, though, they're simply narrow point-in-time snapshots. Their role is to show you what happened during a certain period of time based on the samples of data provided. They fail to show how your organization works over time or find all hidden inefficiencies and risks.

Results: Lack of continuous data-driven visibility

The larger the organization, the more teams and tools involved in operations, and the higher the likelihood of variations in day-to-day processes. Periodic audits only provide reports for a specific point in time and are reliant on opinions rather than continuous operational visibility based on real data.

Introducing Process Mining

Let's turn our focus back to security objectives. As someone responsible for their organization's security is well aware, it's crucial to know if everything you've invested in — tools, people and services — are all working together as expected, consistently.

This question is best answered with a "process map", by which I mean an end to end mapping which illustrates what steps were taken to achieve a particular result. For security teams, some example processes are offboarding privileged users, mitigating vulnerabilities, and responding to incidents. Mapping out these processes helps create standards to consistently deliver the outcomes you want. Process mapping helps to illustrate a holistic perspective and provides a blueprint of how everything "should" work together.

In contrast, process mining ingests real time data from systems and correlates all activity, in order to visualize processes. It offers a holistic perspective and provides a blueprint of how everything "actually" works together. The key concept here is the idea of visualizing processes as they are, reflecting reality as opposed to theory.

DEFINING PROCESS MINING

Gartner defines process mining as "a technique designed to discover, monitor and improve real processes by extracting readily available knowledge from the event logs of information systems." Process mining is distinct from process mapping, despite the fact that both produce process maps. Process mapping is the method of manually-constructing a representation of how a process works or should work in theory. In contrast, process mining, as depicted in Fig 1.2, ingests event data from IT systems and provides fact-based insights into how a process is actually working based on the raw data. Process mining not only shows a map of how things work most of the time, but also surfaces all of the process 'variants', where the process diverges from the prescribed path. These variants provide key insights into hidden risks and inefficiencies in your security operations.

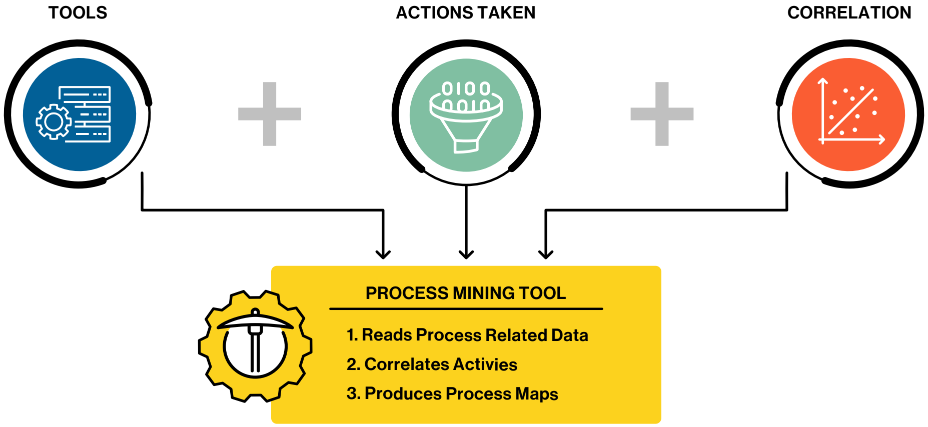

FIG 1.2 BASIC PROCESS MINING STEPS

How Process Mining Helps Optimize Business Processes

Process mining itself is not a new concept. Business process optimization teams were early adopters of process mining, using it to analyze a variety of complex business activities. Complex processes such as quote to cash, supply chain management, and loan origination are typically difficult to visualize and require extraction of events from many siloed tools. By applying process mining techniques, organizations have been able to identify opportunities to reduce waste and optimize processes to achieve faster results.

Here are some examples of how process mining has helped large organizations improve their processes:

Advanced auditing

In a case study, researchers found that process mining revealed a host of errors. These findings were considered audit-relevant, such as payments made without approval, violations of segregation of duty controls, and violations of company-specific internal procedures. These key findings were missed during a regular internal audit which also took place at the same time.

Higher return-on-investment

Following the pandemic and its effects on consumer behavior, a fashion company had to accelerate its digital transformation. In the process of transforming, it became evident that the inefficiencies of their back-end operations would have a significant negative impact on customer satisfaction. In preparation for serving a growing customer base, customer service processes needed to be dramatically tightened up. Utilizing process mining, the company evaluated which process fixes would yield the highest ROI and prioritized them accordingly. Aided by simulations, they determined which repetitive steps could be automated and calculated that the automation of these key steps in the process could reduce customer service resolution times by 90% and cost per resolution by 46%.

Discovering new data

A subsidiary of a well-known bank expected their loan application processes to be fully automated and consistent. However, after applying process mining, it was discovered that only 45% of loan applications were processed automatically and that manual steps added hours of hidden delays to each application.

There are obvious similarities between the business problems discussed above and the kinds of problems security leaders are faced with. Specifically, these business processes involve many discrete tools from different vendors with siloed data and the involvement of multiple teams performing a variety of manual steps. Think about the security processes in your organization along these dimensions and you'll probably recognize many underlying similarities.

How Process Mining Can Improve Security Outcomes

If everything worked perfectly and every tool were well integrated, the world would witness fewer security problems. But that's a far cry from the reality of things today. Even with the best tools and people in place, failure to follow good processes consistently is a major cause for gaps that increase risk and waste time and money. Understanding how processes are really working in your organization based on real data is fundamental to getting more value out of your existing security investments.

Applying Process Mining to Security

With increasing digitization, evolving infrastructure, and demanding compliance requirements, security cannot be stagnant. Consistent improvement is needed to keep up with business needs and an ever-changing threat landscape. However, even with the most advanced security tools and staff, many companies still fail to achieve the outcomes they expect from their security program.

All too often, organizations jump to the conclusion that their people or tools are the cause of these failures, following the latent belief that if only they had the right tools and the right people, their security outcomes would be better. What they often don't account for, is that process variations significantly affect their end results, even with the best people and the best tools in place.

Example: A Process for Application Security

Let's examine a simple app security pipeline, as shown in Fig 1.3. The main goal of this process is to consistently ensure that app vulnerabilities are quickly identified and remediated. The application is regularly scanned for vulnerabilities, the vulnerabilities are reported on, and developers work on remediating the vulnerabilities. It's a continuously repeating process.

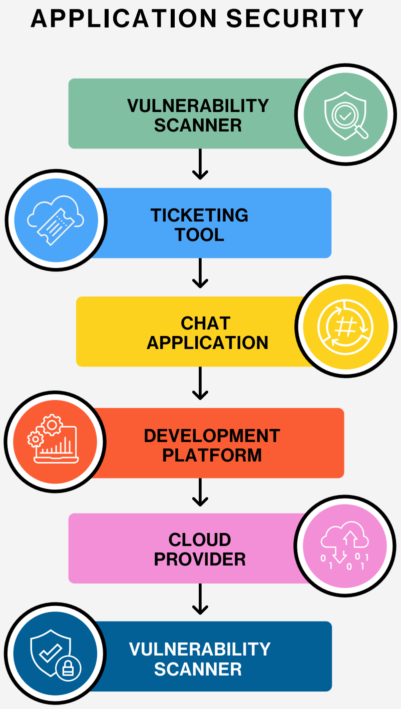

FIG 1.3 SIMPLE APP SECURITY EXAMPLE

LET'S ARRANGE THE PROCESS MAP FROM FIG 1.3 IN ORDERLY STEPS:

1. Security team uses a vulnerability scanning tool to scan the app at build time
2. Scanner reports vulnerabilities by creating tickets
3. Tickets are assigned to the respective owners of the app components
4. App owners use a chat tool to discuss the tickets
5. Developers build a new version of the app and deploy it to the cloud provider
6. Security team uses a different security tool in runtime environment to verify that the reported vulnerabilities were remediated

Observations

Let’s draw some observations from the process map in Fig 1.3:

No tool or team works in isolation

We see from our example that even a “simple” security process involves many steps, tools, and teams. Almost any security process of substance requires coordinating work across multiple tools and teams.

Security processes involve non-security tools

We see that the teams involved in the process use a combination of security and non-security tools. For example, while a ticket management system and a chat app aren’t security tools *per se*, they’re integral parts of this security process.

Security is about designing and executing processes consistently

Any misstep in the process, such as assigning tickets to the wrong team or not upgrading the application at runtime to the right version, has the potential to leave an insecure application in production or waste time rectifying the mistake. It’s difficult to identify these errors without correlating actual event data across all steps in the process.

Processes are “hard to see”

Unless you talk to every team involved in the process, and collect logs from all relevant tools and carefully correlate them, you cannot create an accurate end-to-end mapping of the process. As a result, many security processes are effectively opaque to organizations. Leaders only see negative outcomes after a problem arises, and often don’t clearly understand its causes.

Need For An Upgrade: Shifting the Security Focus to Processes

Today, if you went to your vulnerability management team and asked each of the team members to describe your process for remediating vulnerabilities, you would probably hear slightly different answers from each one, and few if any would reflect the actual reality of what’s really occurring. Even with an official document outlining the process, it’s unlikely that the steps are followed with absolute consistency, not to mention human errors and system errors.

Any attempt to improve security outcomes without a clear view of the processes behind them is akin to flying a plane blindfolded. Even if you know where all the controls are and how to use them, if you can’t see what’s really going on, you can only feel the controls and hope you’re flying safe. A pilot needs to see how changing pitch affects altitude to land safely. Similarly, a security leader needs full visibility to see how the tools and people in a process work together to improve governance strategy and deliver better outcomes.

Maintaining consistent visibility into your security processes based on actual data is vital to achieving good outcomes, and this is where process mining comes in. Process mining shows exactly how people and tools are working together, identifies risks introduced by variants, and helps you see correctable inefficiencies to get more out of your existing tech and teams.

Some key benefits of applying process mining to security are:

Data-driven understanding

You already know the importance of having visibility into where your workloads are deployed, what services you are using, where your data is stored, and how it is secured. Similarly, security process mining ingests data from your systems and automatically correlates activities across your operations to give you unprecedented understanding of how your organization actually works.

Identification of hidden risks

Security is not a one time event. It is a set of continuously occurring processes, happening every day. Even with great detection and response tools and a detailed incident response plan, you still face unnecessary risk if responders forget to isolate systems during investigations. Tools can be great for helping to identify specific technical vulnerabilities and risks, but how you respond to and mitigate them is a larger process that's made much easier with process mining.

Continuous improvement

Gaining an end-to-end understanding of processes helps identify bottlenecks and remediate them faster. Why did it take a month to deploy a critical security update? Is it because your deployment software is unreliable or because it takes weeks just to identify what systems are impacted? Process mining can help you identify the root causes of problems and delays and can help reduce them continuously.

Stronger compliance

Compliance monitoring and enforcement is a well-known security objective. While many posture management tools can alert if you don't have the settings recommended by NIST, CIS, and others, process mining can help you understand how you got to be non-compliant.

Standardization

Process mining helps your teams work towards their objectives with consistent goal setting and performance measurement based on real data, removing bias and other influences.

To summarize, comprehensive cloud security goes beyond purchasing the right tools or hiring the right personnel. To achieve your stated security goals, you need to know if the processes delivering them are effective, consistent, and secure. In the following chapter, we will learn how process mining works.

02

How It All Works

CHAPTER 2

In the previous chapter, we defined process mining and discussed its application to improving security outcomes. To recap, process mining helps organizations visualize and understand how their complex internal processes are actually working, based on real data. Process mining can also be applied to identify areas where processes can be enhanced, such as by reducing manual steps or automating tasks. When applied to information security, it can also identify and analyze unseen risks.

This chapter explains how process mining works, what types of data are involved, how the data is ingested, and what factors are important for visualizations. There are many sources to collect the relevant data from, including cloud platforms, cloud security platforms, and XDR platforms. As noted earlier, even non-security tools involved in security processes need to be mined for security process data, for example ticketing systems and HRIS platforms.

As discussed in chapter 1, there are usually multiple teams and tools involved in any given security process. To better understand the implications of this, let's examine a simple app vulnerability remediation process, as shown in Fig 2.1.

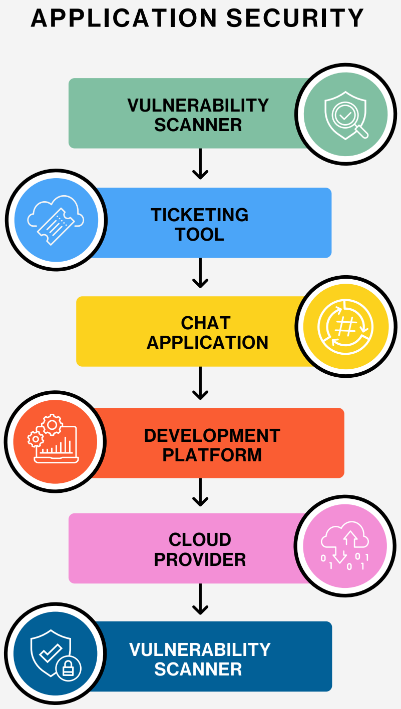

FIG 2.1 SIMPLE APP VULNERABILITY REMEDIATION PROCESS

Ingesting event activity from Data Sources

Processes are composed of steps, each of which leaves a distinctive digital footprint on the systems involved. Every tool retains records of the activities and events that took place in the tool, such as a timestamp of when a vulnerability was discovered, or when it was assigned to an owner. The magic of process mining is its ability to mine the individual steps and automatically correlate the records in order to construct a holistic visualization of the process, thereby making it possible to see and understand processes comprehensively and with a high degree of accuracy.

Process mining is like reading a book; each page reveals clues that, when taken together, tell the story of what happened. Similarly, by examining the digital footprints in information systems, we can gain a comprehensive understanding of how the whole process unfolded. An event in process mining is analogous to a page in a book. A book's pages tell its story in a unique sequence. In order to get the full picture of the story one must read all the pages in the book sequentially. A single page provides only a partial understanding to the larger plot. The book's story is analogous to a security process. Pages in the book correspond to activities recorded by tools. A detailed picture of the overall process can be compiled by combining all activities from all tools, in the order in which they were executed.

Just like a person's unique footprints can be used to trace their path. In the same way, process mining can trace unique events for a process' path, known as a "variant" of a process. Furthermore, just as many people can tread the same path, processes can be performed in a similar manner and follow the same pathway many times. Each such workflow, or pattern, is identified as a process variant. Process mining tracks every relevant case, and automatically categorizes it by the process variant it matches. Typically, a process will exhibit multiple variants, with most cases falling into just a few common variants and following a predictable pattern, with some cases diverging from the norm, exhibiting unusual variants and paths that are rarely taken.

WE MUST BEGIN MAPPING PROCESSES BY INGESTING DATA (MINING) FROM A VARIETY OF SOURCES. GENERALLY, DATA SOURCES CAN BE CATEGORIZED AS FOLLOWS:

1. Security tools
2. Non-security tools
3. Manual tasks

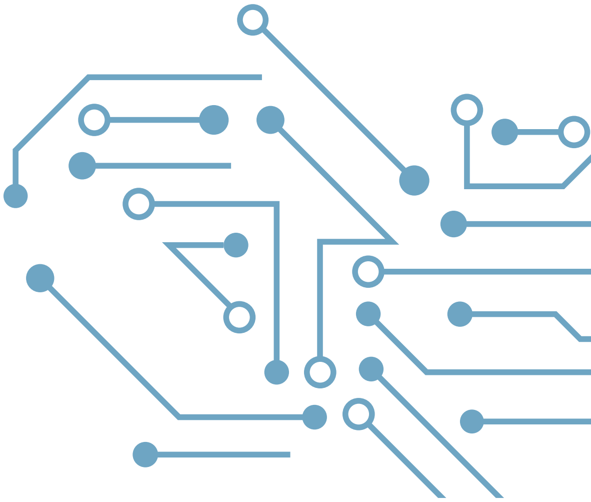

Security Tools

Collecting data from your security tools is key if you want to have visibility into security processes. Security platforms are anything you use to protect your systems, from vulnerability scanners to firewalls to code scanners. Information provided by these tools is essential to understanding various processes. Some examples of data retrieved from these systems are: an incident detected by an EDR sensor, a newly detected vulnerability in an operating system component, or a phishing attempt directed at a high value target. In our example in Fig 2.1, the security team uses vulnerability scanners during build time and at runtime.

Non-Security Tools

Many security processes also involve non-security tools, often in critical roles within the process. There are many examples of this in familiar, common practice. For example, ticketing systems are used to track vulnerability remediation efforts, HRIS platforms trigger the process to create or remove users, and messaging platforms are used to receive alerts about security incidents. Mining data from such non-security systems is critical to understanding security processes comprehensively.

In the example shown above in Fig 2.1, the ticketing tool and the cloud service provider hosting the application are non-security tools. Such tools still support the security process, despite the fact that they don’t inherently provide security capabilities. They are essential for tracking whether risks are being mitigated, who is responsible for mitigating them, and how tasks are progressing. The security aspects of the application, however, fall outside the scope of their responsibility.

File Imports

Lastly, processes can involve offline steps which could enrich the overall picture if added to the automatically mined data sources. Take for example a user offboarding case that involves an HR staff member manually preparing a spreadsheet of dismissed staff for a given week and sending it to an account management team for processing. While the spreadsheet itself may not be an online resource that can be automatically ingested, it can be imported into the process data set to make the overall data set even more complete.

Data exported from systems must be correlative and include at least three pieces of information:

Identifier

The data must contain an identifier associated with a particular case. Each case is a single, unique occurrence of the process. An identifier is needed in order to track multiple activities related to the same case across different systems. An identifier helps tie together the data from all of the different systems and correlate it to the case.

Event message

The data must describe the activity that took place.

Timestamp

The data must include the time associated with the event. Activity timestamps are needed in order to construct a process map that arranges the data in the proper sequential order.

Process mining does not require you to connect to everything at once. If you want to start with less complexity, analyze a category of processes such as identity management, which only requires connection to systems associated with it, like your directory service and federation service. Connecting tools such as application scanners and firewalls that aren't part of identity management processes is not necessary. By starting with a specific process as a first step, you can build your way up to analyzing more processes and systems and scale as needed.

Ingestion

Process mining involves correlating events from a wide variety of data types from multiple systems.

When ingesting data, three factors should be taken into account:

Data format

As discussed above, system records should contain at least three important data points — the identifier, event message, and timestamp.

Frequency

To ensure accuracy, data should be ingested at regular intervals so that the different times and ways processes occur can be taken into account. Continuous synchronization of process data can be achieved by ingestion of real-time or near real time data from systems. When automatic ingestion is not supported, data can be supplemented by uploading batch files periodically.

Data transformation

Transforming the data collected from different systems, tools, and uploads into normalized and correlated information is a necessary prerequisite to process mining analysis. The backend of any process mining system should be capable of consolidating the data collected, removing noise, aggregating and normalizing the data as needed, and automatically constructing a process map from the results.

From Ingestion to Visualization

It's possible to ingest a wide range of data from different systems and tools for process mining, including log files, audit trails, alerts from security tools, and ticket metadata from workflow systems. Using a simple example, let's examine how events from different data sources are correlated.

Figure 2.2 illustrates a vulnerability management process based on data from three sources — a vulnerability scanner, a ticketing platform and a development pipeline.

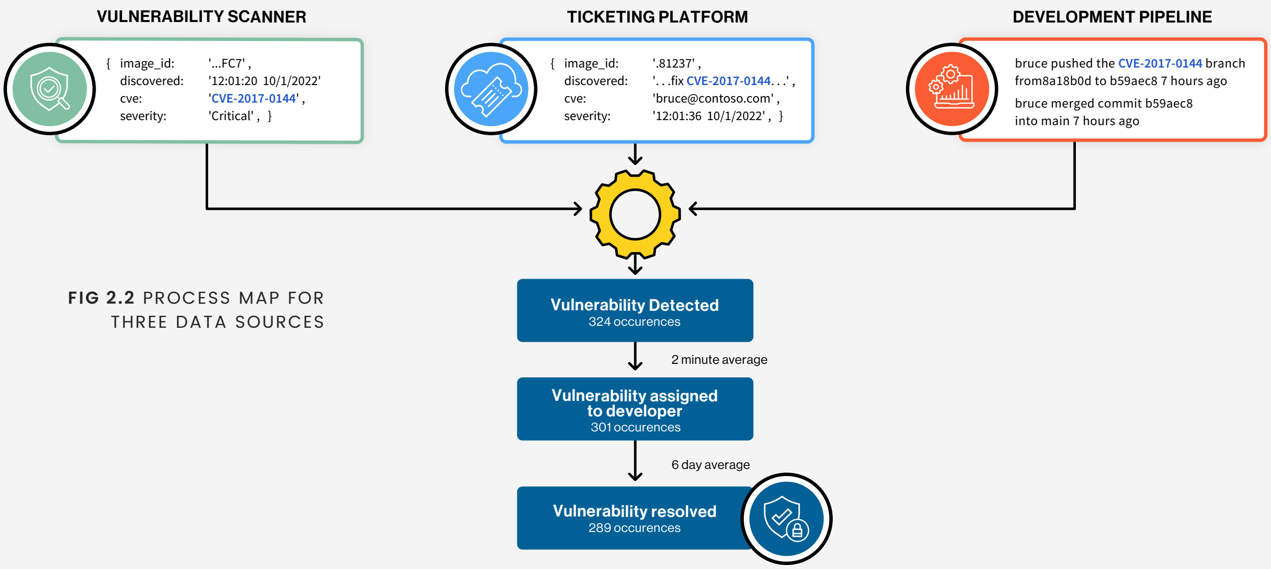

FIG 2.2 PROCESS MAP FOR THREE DATA SOURCES

```
{  image_id:       '...FC7' ,
   discovered:     '12:01:20  10/1/2022'
   cve:            'CVE-2017-0144' ,
   severity:       'Critical' ,  }
```

```
{  image_id:       '.81237' ,
   discovered:     '. . .fix CVE-2017-0144. . .' ,
   cve:            'bruce@contoso.com' ,
   severity:       '12:01:36  10/1/2022' ,  }
```

```
bruce pushed the CVE-2017-0144 branch
from8a18b0d to b59aec8 7 hours ago

bruce merged commit b59aec8
into main 7 hours ago
```

1. **The vulnerability scanner reports a vulnerability on an image**

   Identifier: CVE-2017-0144 and resource_type:image_id

   Event message: "..discovered CVE-2017-0144"

   Timestamp: 10/1/2022, 12:01:20

2. **The ticketing system indicates that the vulnerability alert has an associated developer ticket**

   The ticketing platform holds the following information:

   Identifier: CVE-2017-0144 and resource_type:image_id

   Event message: "ticket created"

   Timestamp: 10/1/2022, 12:01:36

   Other info: Assignee bruce@contoso.com

3. **The development platform shows that the developer pushed a commit to resolve the CVE-ID**

   Identifier: CVE-2017-0144 and resource_type:image

   Event message: "pushed CVE-2017-0144 branch and merged commit"

   Timestamp: 10/1/2022, 20:00:00

   Other info: dev user bruce@contoso.com

By correlating just three records from three different tools, a process mining tool can automatically construct a process map that shows what happens when a security bug is discovered. We see here an example of how process mining involves automatically correlating events from different tools.

Process Variants

A security process mining tool first maps each occurrence of the process, (i.e. case) by correlating the timestamps and identifiers for all of the events ingested from the data sources. Every case, which is to say end-to-end process execution, is mapped as a sequence of events. Cases follow repeatable patterns, and each such pattern is identifiable as a unique process variant. Every process variant is recognizable by the type of events it includes and their sequence.

In the above example shown in Fig 2.2, we see that to successfully manage vulnerabilities, all three steps must be followed precisely in the same order. In other words, the sequence of the events is crucial to the outcome. This diagram shows the ideal process flow, which is sometimes referred to as the target or desired variant. Ideally, all cases of a process should follow the pattern of the target variant. But in reality, we find that cases often deviate from the target, with methods often changing between executions, and leading to many process variants. Not surprisingly, some variants lead to undesirable outcomes, which in our example would mean that the vulnerability would remain unresolved.

Next, in Fig 2.3, we see how a process mining tool can help detect the presence of multiple process variants.

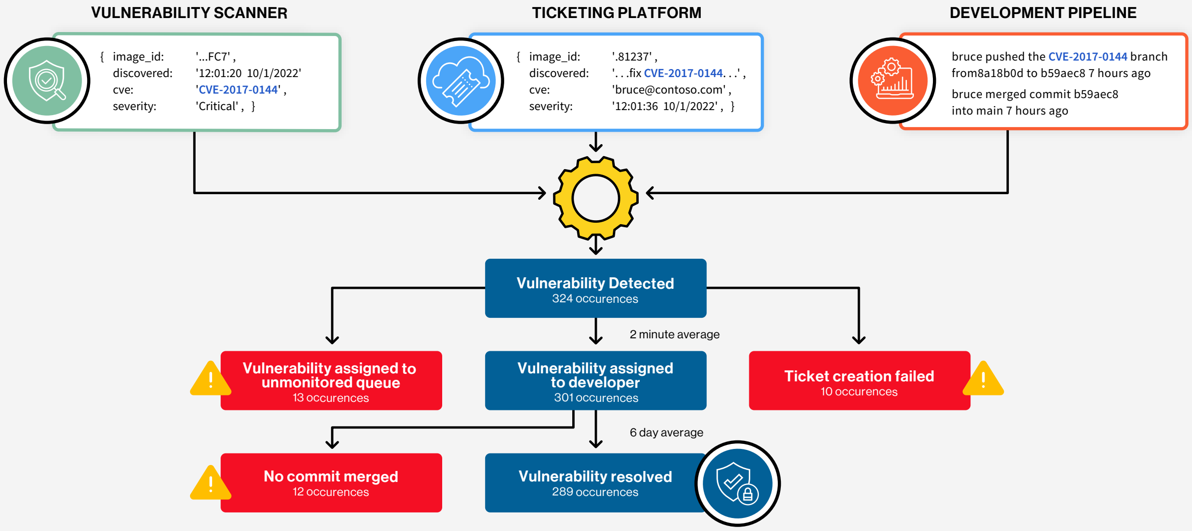

```
{  image_id:      '...FC7' ,
   discovered:    '12:01:20  10/1/2022'
   cve:           'CVE-2017-0144' ,
   severity:      'Critical' ,  }
```

```
{  image_id:      '.81237' ,
   discovered:    '. . .fix CVE-2017-0144. . .' ,
   cve:           'bruce@contoso.com' ,
   severity:      '12:01:36  10/1/2022' ,  }
```

```
bruce pushed the CVE-2017-0144 branch
from8a18b0d to b59aec8 7 hours ago
bruce merged commit b59aec8
into main 7 hours ago
```

FIG 2.3 PROCESS VARIANTS

Some Observations

In our example, a total of 324 instances of vulnerabilities were discovered, yet only 289 were actually resolved. After being successfully discovered, those vulnerabilities remained unresolved because other steps in the process failed or were not pursued as expected. Overall, process mining reveals that 10% of the cases failed in achieving the desired outcome. More importantly, process mining uncovers the reasons why 10% of the cases failed:

1. **In 13 cases** the ticket was assigned to an unmonitored queue and went unnoticed.
2. **In 10 cases** the vulnerability scanner tool discovered a vulnerability but no ticket was created in the ticketing system. Failure to create tickets resulted in discovered vulnerabilities not being tracked for remediation.
3. **In 12 cases** a developer was assigned to the ticket, but the work was not completed. This was indicated by the fact that the development pipeline did not register a code commit or merge action to indicate that a fix for the vulnerability was pushed to production.

Conclusions

As we can see from the variants in Fig 2.3, the visibility afforded by process mining can be quite significant.

Tools can work without errors and still fail to meet the desired objective

Most variants of the process involved the use of all the tools without errors —

- CVEs are successfully discovered by the scanning tool
- Tickets are created and assigned using the ticketing tool
- The development platform has commits for all code pushes

Despite the fact that all tools work correctly, the intended result of remediating vulnerabilities was not consistently achieved.

Processes can be defined yet followed inconsistently

In this example, the process for vulnerability management is already established. Once a vulnerability is discovered, a ticket is created and assigned to a developer, who will work on it and push a fix to the code repository and then close out the ticket as done. This process may even be clearly documented in runbooks or a wiki. However, even if the desired flow is written down, it's not always followed consistently, because of human error or even personal preference.

Process variations are usually expensive

Inconsistencies in a defined process can often lead to unexpected delays and, in some cases, complete failure to meet the objectives. Therefore the universal adoption of a well-designed process should be considered a primary goal for any security leader. In the absence of visibility into what is actually taking place, it is impossible to measure conformance, nor understand what's leading to inconsistent outcomes.

What Process Mining is Not

We have so far discussed what process mining is, but in order to gain a better understanding of the concept, let's take a look at what it isn't.

"It's just a data lake"

Process mining does not try to create a common data lake for detailed data from all your tools and systems. A good process mining tool should remove noise (uncorrelated events) and show and ingest only information relevant to the processes being tracked. It's expressly not designed to just aggregate data, but rather to present a contextualized visual representation of actions.

This distinction is important because it helps to focus on just the most important data elements that compose the process. It also helps to reduce the amount of data that has to be collected and analyzed, making the process more efficient and less resource-intensive.

"Process mining has to collect all data from all tools to be useful"

It's not necessary to have all events from every system for process mining tools to make correlations. As shown in example 2.2, process mining requires only enough metadata to correlate events. Metadata can include information such as the date and time of the event, the type of event, and other aspects related to it. For instance, a process mining tool could be used to detect process bottlenecks by correlating events with the same type, such as "User Created" or "User Added to Group", without needing to access the details about each user profile or mine other data.

"It's the same as posture management"

The purpose of posture management systems is to tell you if something is configured correctly. For instance, "Make sure your cloud users have 2 factor authentication enabled." A posture scan results in a Boolean value — PASS or FAIL. The individual setting (posture) is configured either one way or another at any given point in time and the posture management system doesn't offer any context or insights into how or why that's the case.

In contrast, process mining shows how an overall process works, rather than just how settings are configured. Instead of just looking at the 2FA settings, process mining helps you understand the entire flow of how users are created, how their accounts are configured, and what led a user to not have the 2FA setting enabled.

This is not to say posture management isn't important; In fact, a posture management tool may be a data source that's mined as part of an identity management workflow. However, while the posture management tool tells you what a setting is, process mining helps you understand how it came to be. For example, process mining may help you identify that in the onboarding variant where users are created by an outsourcing partner, 2FA settings aren't being configured. Posture management helps monitor settings while process mining helps provide understanding of broader workflows, of which individual settings may be a part.

Now that we have a better understanding of what process mining is and isn't, as well as how it works, let's explore the actionable insights that process mining can provide.

03

Security Governance

CHAPTER 3

In the previous chapters I defined how process mining works and why there is a need for process mining in security. In this chapter I elaborate on this concept further by diving into security governance, what it means, its current state, and what impact the application of process mining can have in transforming governance.

What is Governance

Gartner defines “Security Governance” as a process for overseeing the cybersecurity teams who are responsible for mitigating business risks. An essential function of security governance teams is to develop and maintain the risk management strategy of an organization. This essentially determines an organization’s readiness to cyber threats.

In this chapter we will dive deeper into the significance of governance functions and elaborate on how the application of process mining makes governance more measurable and efficient. To begin with, let’s understand what security governance is in terms of people, process, and technology. I’ve placed it in this equation:

(People + Technology) * Process = Security Outcomes

Governance is how you balance the equation to improve and deliver desired security outcomes. An important point I’d like to highlight in this equation is that ‘Process’ is a multiplier to the investments in people and technology. What it implies is that the sum of strong investments in people and tech can be negated when process conformance is poor; conversely, even when an organization invests leanly into people and tech, they can get good outcomes if they have strong process discipline.

Now that we had a refresher on what governance is, let’s take a look at some of the core aspects of cybersecurity governance and see how process mining can help with practical examples.

Govern is the core of Cybersecurity Framework

You can define a governance strategy from scratch but it is highly recommended to adopt an industry wide established cybersecurity framework such as NIST's Cybersecurity Framework (CSF). Frameworks assist organizations by providing guidance on how to establish the processes and procedures that your organization must take to assess and mitigate cybersecurity risk. Let's take a look at one such well known cybersecurity framework in the industry and its guidance for governance.

The NIST Cybersecurity Framework, which is a leading cybersecurity standard guidance for organizations world wide, recently released a new 2.0 version of the framework. The major change in CSF 2.0 is the introduction of "govern" as an additional sixth pillar of the program as shown in Fig 3.1. Not only did they add govern as a core function but CSF also positioned it as one that touches all other functions - Identity, Protect, Detect, Respond and Recover.

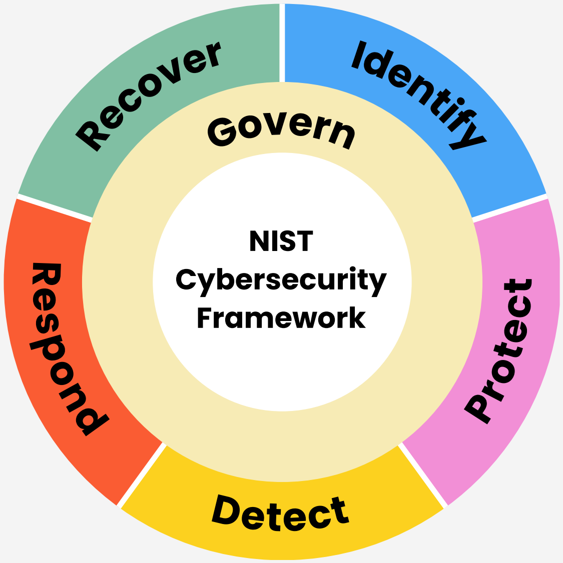

FIG 3.1 THE "GOVERN" FUNCTION, WHICH EMPHASIZES THAT CYBERSECURITY IS A MAJOR SOURCE OF ENTERPRISE RISK AND A CONSIDERATION FOR SENIOR LEADERSHIP.

REFERENCE: CREDIT: N. HANACEK/NIST

NIST highlights how the governance function in an organization essentially informs how an organization will implement all the other core functions. It places cybersecurity as a major source of enterprise risk among other risks such as finance and legal. It means that if you're improving governance, you're improving the foundation of all other functions. If governance is weak, then all the other areas are also negatively impacted. To put it simply, governance should align how all the people and technology that organization has invested in work together in processes to deliver consistent security outcomes.

Ideally, all your people and technology would work together in consistent processes, delivering optimal security outcomes across all your security programs. But we know in reality, that is rarely the case. This is why measurement is such a key function within governance. Being able to measure your organization's process conformance against industry frameworks and your own internal goals can help provide a clear picture of where you currently are, what is missing, and what hinders you from reaching your optimal conformance rate.

Core Requirement:

Established frameworks aim to provide guidance to organizations to reduce cybersecurity risk. CSF 2.0 introduces 'Govern' as an additional core function of cybersecurity posture that impacts and is impacted by all other core functions - identify, protect, detect, respond and recover.

By following the framework, organizations can assess their current maturity profile, accurately identify gaps between their current profile and target profile, and be able to implement procedures and processes to improve their cybersecurity maturity.

Challenges:

Although the framework offers comprehensive guidance in understanding an organization's current maturity and improving cybersecurity risk management, the practical implementation of processes to adhere to CSF has several challenges:

> **Assessing current cybersecurity maturity profile** is not easy when the data required to assess is distributed across many teams, technology and processes.
>
> **Even when gaps are identified**, implementation of new governance strategies is often resisted because of factors like technical inertia, budget, and business justification.
>
> This results in governance that is often **manual, slow and reactive**.

Transforming with Process Mining:

Here's how the current challenges could be addressed when process mining is implemented:

> **Instead of relying on manual analysis of your company's cybersecurity profile, tracking your maturity through automated process conformance measurement dashboards that are based on continuous data ingestion.**
>
> **Example:** Process variants that show all process executions (cases) where you failed complying with key security controls like patching critical vulnerabilities and the exact event in the process chain that caused the non-compliance. Remember though, that this is not a compliance scanner to show you a TRUE or FALSE state of a setting. It shows you the actions or inactions that lead to your TRUE or FALSE compliance state. For example, if the fix for the vulnerability has been built and tested but is awaiting approval to be deployed. In this example, your people and tech may have found and begun to address the problem, but the process delay results in an undesirable security outcome.
>
> **Continuous visibility into the health of the overall process and its outcomes, not just tracking initial detections or alerts.**
>
> **Example:** Instead of waiting for audits to show you where your processes fall short, process mining with alerts integration can proactively send alerts based on process condition failures in real time. For example, an alert sent to the security team whenever a critical vulnerability ticket remains in unassigned queue for over a week, or a pending fix in production is delayed by an unusually long time.

Government mandates on Governance

Software is eating the world. Many crucial and day-to-day functions today are taking place digitally. Your identity, wealth, health, education, and many things you own have digital traces stored somewhere in the cloud, one breach away from exposure.

With increased reliance on software, the threat and liability of the organizations utilizing and storing this data is also heightened. Given the impact cybersecurity incidents can have on a country's citizens, it is predictable for governments to mandate rules and regulations on cybersecurity practices. And that's exactly what we are witnessing with elevated government imposed cybersecurity regulations across the globe.

Security governance has become not just an important pillar for the organization itself, but also for external stakeholders such as governments and end users. As a governance leader, you are held accountable to reputational risk, financial risk, and legal mandates. You're also responsible for communication of breaches to impacted parties based on accurate information.

The European Union recently introduced the Cyber Resilience Act and the Cyber Solidarity Act, with a goal to mandate increased cybersecurity preparedness in organizations. The existing EU umbrella law, the GDPR, has been used to levy heavy fines on companies that failed to inform customers about the compromise of personal data during cybersecurity incidents. In 2023 alone, the EU levied around $2 billion in GDPR fines against major global companies.

Regulatory drivers include data privacy concerns, nationalism, and the economic value of data

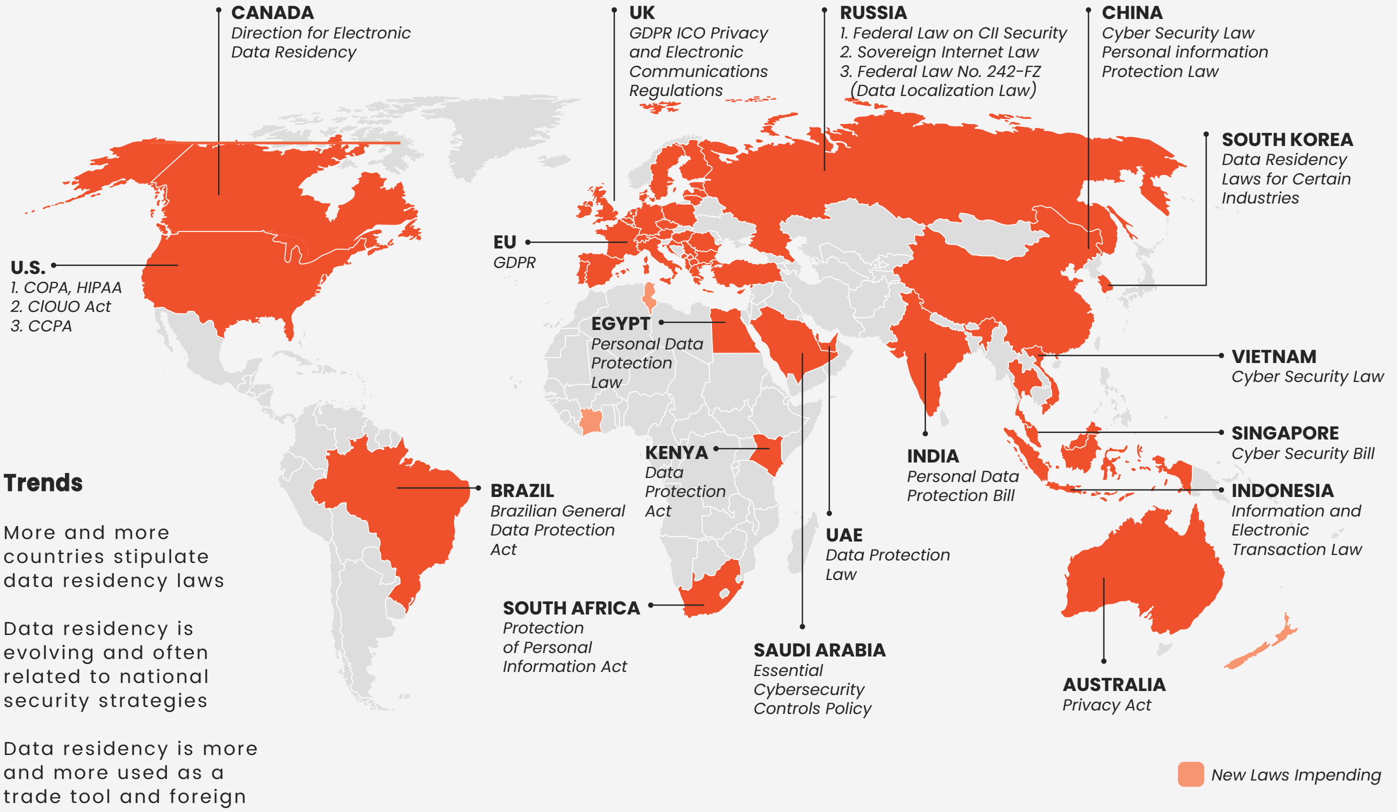
CANADA
Direction for Electronic Data Residency
UK
GDPR ICO Privacy and Electronic Communications Regulations
RUSSIA
1. Federal Law on CII Security
2. Sovereign Internet Law
3. Federal Law No. 242-FZ (Data Localization Law)
CHINA
Cyber Security Law
Personal information Protection Law
SOUTH KOREA
Data Residency Laws for Certain Industries
U.S.
1. COPA, HIPAA
2. CIOUO Act
3. CCPA
EU
GDPR
EGYPT
Personal Data Protection Law
VIETNAM
Cyber Security Law
KENYA
Data Protection Act
INDIA
Personal Data Protection Bill
SINGAPORE
Cyber Security Bill
BRAZIL
Brazilian General Data Protection Act
INDONESIA
Information and Electronic Transaction Law
UAE
Data Protection Law
SOUTH AFRICA
Protection of Personal Information Act
SAUDI ARABIA
Essential Cybersecurity Controls Policy
AUSTRALIA
Privacy Act
New Laws Impending

Trends

More and more countries stipulate data residency laws

Data residency is evolving and often related to national security strategies

Data residency is more and more used as a trade tool and foreign investment

FIG 3.2 CYBER SECURITY REGULATIONS ACROSS THE WORLD.
IMAGE CREDITS: UNSECURE.IO

In the United States, the FDA (Food and Drug Administration) levied a heavy fine against a healthcare organization for failure to comply with HIPAA guidelines due to a phishing incident. This marked the first phishing-induced HIPAA violation penalty in the industry. Apart from nationwide rules, more and more states in the US are also adding their own cybersecurity policies and regulations, increasing liability and security governance complexity on organizations.

The United States has also seen an uptick in lawsuits against CISOs, some even leading to criminal charges for their response to cybersecurity breaches. This shows an accelerating trend towards more complexity and potential liability in security governance. Security leaders are held to stricter expectations of maintaining risk management strategy along with responsible disclosures.

In 2023, the SEC (Security and Exchange Commission) in the United States announced new cyber security rules with much more strict cybersecurity incident disclosure requirements, putting companies on a tighter four day timeline. SEC's new policies also emphasize periodic disclosure of cybersecurity risk management strategy, security processes, and governance posture.

Core Requirement:

Governments across the world aiming to legislate and mandate cybersecurity regulations all have similar interests in mind.

> **To help organizations bolster their governance strategies** and thus improve preparedness against cyber attacks.
>
> **To enforce responsible disclosure in the event of a cybersecurity incident**, both to the government and to the victims whose data might be compromised due to the incident.
>
> **To describe and maintain disciplined security specific processes** such as those involved in assessing, identifying, and managing material risks from cybersecurity incidents.
>
> **To hold organizations and their leadership accountable** with increased liability respective to security governance practices in an organization.

Challenges:

While the policies impose heavy mandates, companies face serious challenges in abiding:

> **Governance processes and procedures may be documented** but are often outdated and based on human assumptions. Static documents often don't reflect operational reality and aren't sufficient to make decisions when heavy liability is involved.
>
> **The ongoing evolution of government mandates** requires faster adaptability and transformation. However, governance that is not fully automated and data-driven is slow and resistant to change.
>
> **Lack of security culture and accountability** across the organization. The flow of information and communication can often get mingled and confused with many teams, layers, and responsibilities spread across an organization.
>
> **Accuracy of reports.** Security leaders are increasingly held personally accountable to disclosure and regulatory requirements. This means that the information they provide to stakeholders needs to be accurate and data driven. However, gathering data at scale across the many systems and teams that support an organization's security programs is challenging.

Transforming with Process Mining:

Here's how the current challenges could be addressed when process mining is implemented:

Automatic process visibility to provide data driven, process centric understanding of how your governance functions really work in an org. Being able to compare your actual process performance against the target goals enables clear understanding of program effectiveness.

Example: Being able to provide a process "blueprint" or ideal process workflow into your process mining tool and then compare how your actual processes do against this ideal documented process. Assess maturity based on objective ingested data instead of anecdotes and assumptions.

IT is being significantly impacted by trends like cloud, AI, and decentralized IT decision making. Process mining can be used to improve governance by providing clear alignment between organizational priorities and security programs, based on objective KPIs that process mining can measure.

Example: When deciding to replace a security tool or add automation in the stack, process mining can measure the effectiveness of new investments. For example, the detailed correlated data from process mining can clearly track what impact the change has had on the overall security objective and how the change can improve / degrade security outcomes.

We all know that expecting to drive security accountability through HR mandated security training can only go so far. To drive security culture, you need to provide visibility into staff members' actions in security related workflows and the outcomes they result in.

Example: With proper ingestion of metadata, process mining can show you not only how all events (steps in the process) are executed but also who owns each particular step's execution. So the next time an SLA is breached, instead of blind blaming, you can show the process owner exactly which step caused the failure, with responsible disclosure of action that must be taken to rectify it, thus enforcing security culture with event visibility.

An important and often stressful aspect of SEC cybersecurity rules is its requirement to disclose any cybersecurity incident determined to be material within four days of discovery. Recent lawsuits against CISOs for failure to do so within acceptable timelines show the personal liability an organization's security leaders carry in providing such information. It is hence of utmost importance that the reports you present as a security leader are accurate based on real events instead of anecdotes from humans. **Process mining can automatically compare** the case identifier (specific process execution in which the incident was determined)across all tools and activity and show you progression of the attack and response so you can have accurate information for reporting.

Audits and Certifications

The average organization spends 301 hours, or 37 work days, gathering data for an audit. Still, recent cybersecurity news shows that among the enterprises that suffered a breach, almost all of them were certified with a known compliance standard and validated by an audit. The problem isn't with the compliance certificate checklist but in the way that these certifications are often obtained.

Audits today rely heavily on a laborious manual effort. Even then, the data that is collected for audit is only a small sample of the actual volume. Hence the certification becomes a mere representation of the sampled, point of time data rather than active, comprehensive coverage of security governance.

The following diagram shows some of the typical steps that take place during a generic audit in Fig 3.3

How it works

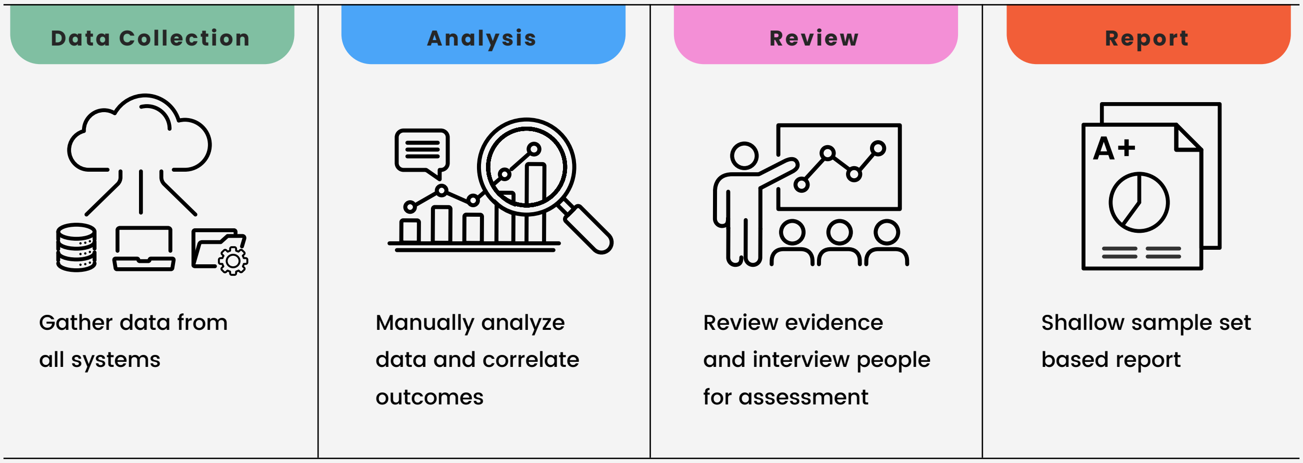

Data Collection

Gather data from all systems

Analysis

Manually analyze data and correlate outcomes

Review

Review evidence and interview people for assessment

Report

Shallow sample set based report

FIG 3.3 STEPS IN AUDIT PROCESS

1. **Gather data** from the systems and teams involved in operations. Samples of data within a specific time period are taken for review. Questionnaires, exceptions, and manual explanations are documented for auditors.
2. **Manually analyze** large amounts of data, often using spreadsheets or other generic tools. Significant effort is usually required just to normalize the data from different systems and identify the flows between them.
3. **In the end**, you have a findings report that is based on a shallow, narrow, sampled data set based view of what your organization is doing.

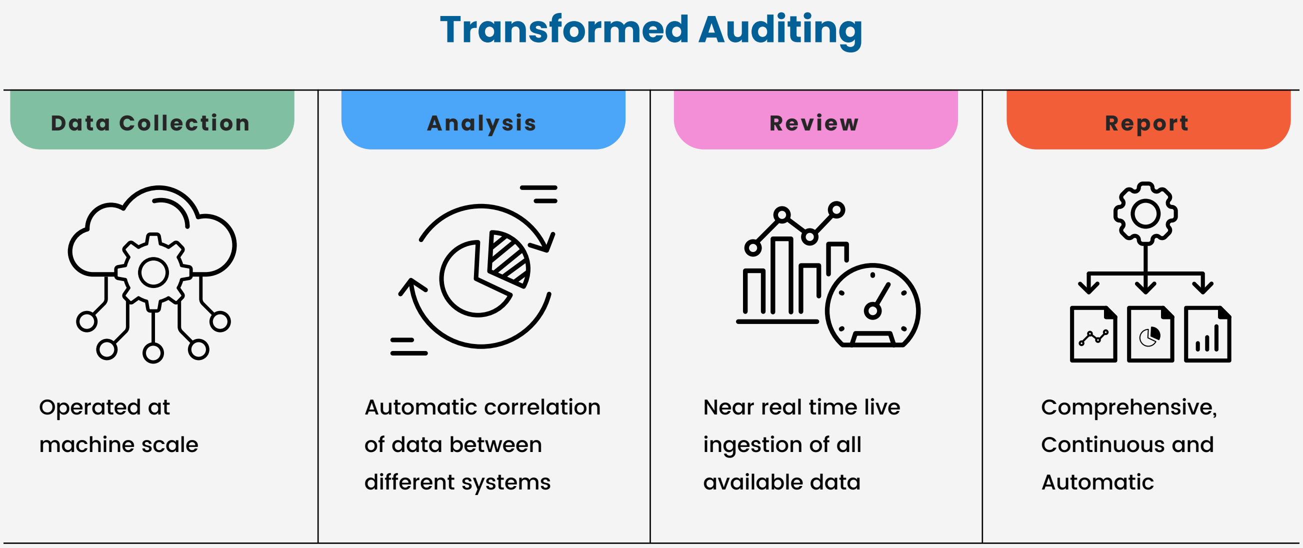

Transformed Auditing

Data Collection	Analysis	Review	Report
Operated at machine scale	Automatic correlation of data between different systems	Near real time live ingestion of all available data	Comprehensive, Continuous and Automatic

FIG 3.4 STEPS IN AUDIT PROCESS WHEN PROCESS MINING IS APPLIED

1. **Instead of taking data dumps from each system and team,** directly integrate all systems for automatic ingestion so it's always available. Collected data from all systems is available for auditors within minutes, lessening human dependency.
2. **Automatic correlation of data is done by process mining,** matching case and event identifiers across all systems and measuring them against security outcomes.
3. **Near real time live ingestion of data going forward.** Even after the audit is complete, process mining gives you proactive visibility into how your organization is really doing and if you are truly audit ready.

Transformed Governance

The goal of applying process mining to governance is to make its functions faster, more flexible, and data driven. Today, security governance is often opinion or human assumption driven. Often, monetary investment is focused on getting "a better" security tool to get better results. However, the fundamental questions that need to be answered to improve security governance are not derived from one single tool. Developing a strategy and executing on it requires visibility and comprehensive understanding of ALL tools and the security team's performance using these tools. That is what process mining provides to security leaders.

To summarize:

Process mining directly impacts security governance in the following ways:

Data driven strategy changes

Governance underlies everything in security and is at the foundation of delivering successful security outcomes. It's important to have visibility into all security processes, to set goals, measure performance, fix problems, and communicate results. With process mining, you don't have to assume where to invest resources or make changes for improvement. You have objective data correlated from across all the tools and systems you already own that can be used to measure performance, identify problems, and communicate results.

Measured governance goals

How do you measure your security governance today? To answer this question you would probably need to look at your last audit report, spreadsheets with data derived from different teams, presentations made in board meetings based on data from different tools ,and yet it will only be a partial picture. Setting goals for security outcomes is easy but measuring goals when the visibility is siloed across different tools and teams is difficult without process mining.

Continuous performance measurement

Current ways of understanding security maturity in an organization through audits and certifications is based on samples of point in time data. It's not easy to make or measure improvements when performance analysis is not continuous. With process mining ingesting data continuously you can see immediate results and see the effect of changes, improvements over time, and remaining bottlenecks to overcome in real time.

Actionable alerts for audit preparedness

Instead of waiting to know what's wrong by auditing, a process mining tool can alert you on process deviations, so you can make changes immediately and improve your conformance.

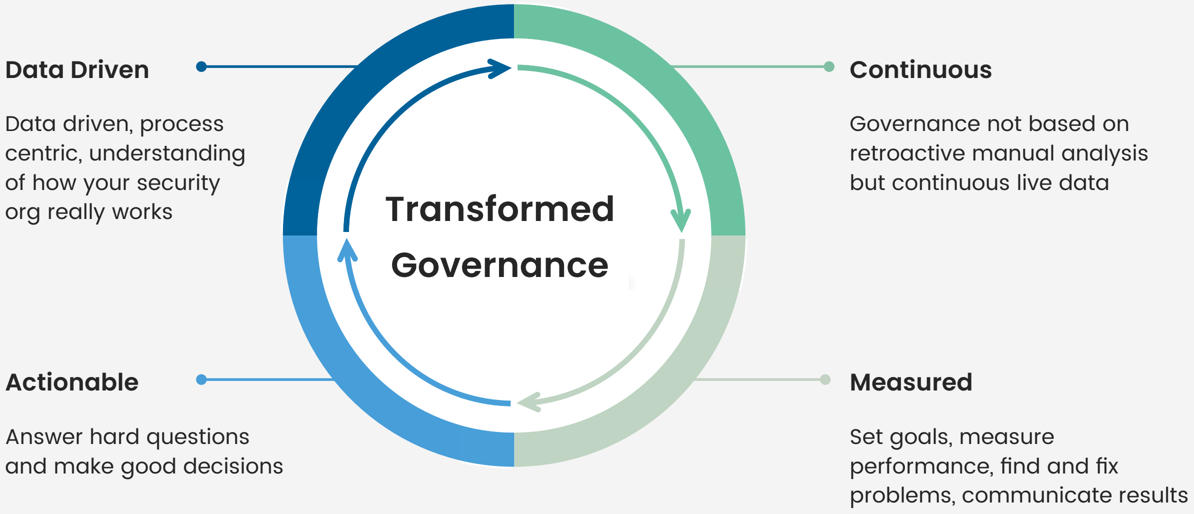

Data Driven

Data driven, process centric, understanding of how your security org really works

Continuous

Governance not based on retroactive manual analysis but continuous live data

Transformed Governance

Actionable

Answer hard questions and make good decisions

Measured

Set goals, measure performance, find and fix problems, communicate results

FIG 3.5 IMPACT OF PROCESS MINING ON SECURITY GOVERNANCE

In the next chapter I will elaborate further on actionable insights that only a process mining tool can provide for cybersecurity.

04

Actionable Insights from Process Mining

CHAPTER 4

In our previous chapters, we discussed what process mining is and how it works. The purpose of this chapter is to explore how process mining insights create real business impact. Here we'll survey how process mining can assist in managing risks, inspecting variances, identifying inefficiencies, and improving outcomes.

Further, we will discuss how to use process mining tools to track internal goals and compare performance to best practices and industry averages.

Understanding Variants and Variation Tolerance

In the last chapter, we defined what a variant was — a unique path taken in the execution of a process. Although not ideal, it is common to have several variants of a single process. Tolerance is a term used to describe how much deviation from the ideal execution path (desired variant or target variant) of a given process is acceptable. The ability to keep security operations within acceptable levels can be difficult without knowing what tolerable variances are. To explore this more, let's go back to our vulnerability management example in Figure 3.1.

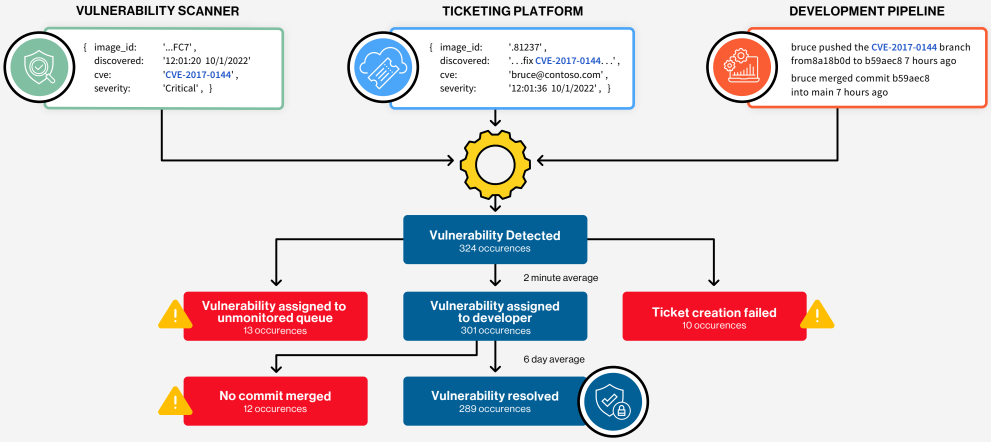

```
{  image_id:     '...FC7' ,
   discovered:   '12:01:20  10/1/2022'
   cve:          'CVE-2017-0144' ,
   severity:     'Critical' ,  }
```

```
{  image_id:     '.81237' ,
   discovered:   '. . .fix CVE-2017-0144. . .' ,
   cve:          'bruce@contoso.com' ,
   severity:     '12:01:36  10/1/2022' ,  }
```

```
bruce pushed the CVE-2017-0144 branch
from8a18b0d to b59aec8 7 hours ago
bruce merged commit b59aec8
into main 7 hours ago
```

FIG 3.1 VULNERABILITY REMEDIATION PROCESS VARIATIONS OUTSIDE TOLERANCE

In this image, we see some variants of the vulnerability management process are highlighted in red. These variants all resulted in unsatisfactory outcomes, where at the end of the process, the remediation of the vulnerability was not achieved. In cases where a ticket was assigned to an unmonitored queue, there was a failure to create a ticket. In cases where the developer did not commit the fix to the repo, the vulnerability remained. Therefore, these variants lead to vulnerabilities that were not addressed, leaving the organization exposed. In other words, these variants in themselves are "risks" in the process.

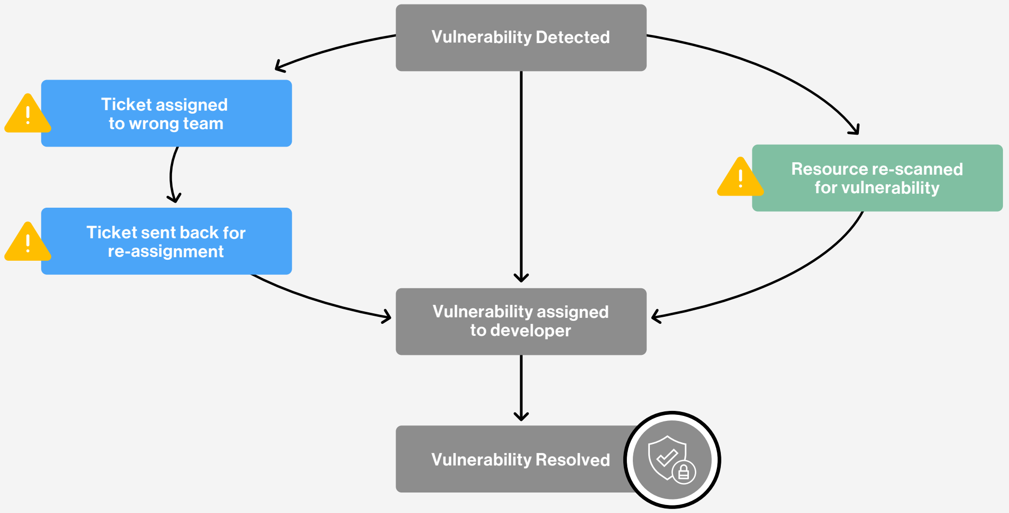
Vulnerability Detected
Ticket assigned to wrong team
Ticket sent back for re-assignment
Resource re-scanned for vulnerability
Vulnerability assigned to developer
Vulnerability Resolved

FIG 3.2 VULNERABILITY REMEDIATION PROCESS VARIATIONS WITHIN TOLERANCE

Now consider some other variants we might encounter in a vulnerability management process.

We see two variants in this example Fig 3.2 are colored blue and green:

> **Variant 1:** A vulnerability is detected. A ticket is successfully created for the vulnerability, but assigned to the wrong team. The ticket is routed back to the queue for reassignment to the correct team, then reassigned to the correct developer. The vulnerability is then successfully resolved.
>
> **Variant 2:** A vulnerability is detected. However, the scanner fails to create a ticket after the scan. When the next periodic scan occurs, the resource is rescanned and the vulnerability ticket is created and assigned to the developer after re-scanning.

In both cases, the vulnerability is eventually resolved and the overall goal of the process is still achieved. However, both of these variants were a cause of redundant steps, wasted time, and extra effort. In security, wasted time often means increased exposure, as is the case here. The longer it takes for the ticket to be assigned to a developer, the longer the environment is susceptible to the vulnerability.

Variation tolerance is the permitted range of deviation from the standard process flow, i.e. the desired variant. When a process crosses the threshold and has too many undesired variations, it is said to be "out of bounds." Variation tolerance is important because it helps to measure whether the process is being followed correctly and that the end goal is achieved. In our example, if such variants occurred frequently, the process may be considered to be out of bounds. While there's no universal tolerance rate that's applicable to all processes — each organization needs to determine what tolerance levels are acceptable for a given process.

Desired Variants

While not all process variations result in negative outcomes, they can lead to other issues. In a security context, they may lead to inaccurate reporting of vulnerability metrics or underestimating the frequency of SOC incidents.

A lack of consistency can make assessing overall performance difficult due to variations in the process. The variations can also lead to a decrease in quality and an increased number of errors if they are not properly monitored. For this reason, it is essential to understand variations in a process so that they can be managed effectively.

The "desired variant" of the process is the ideal and most efficient flow. Basically, the desired variant is what you **want** the process to be every time. By identifying and following a desired variant, organizations can ensure that quality and accuracy are maintained, and that the process runs efficiently. Without defining a desired variant, it is impossible to measure the effectiveness of a process or the impact of variance on it.

The desired variant for any given process is likely to differ from company to company and even within the same organization across different teams. My definition of an ideal process may differ from yours for the same task. If I were writing a book, I'd outline it first, then write the chapters based on my outline. Other authors might start with an idea and develop the outline as they write the chapters. Our ideal paths to completing the book need not match even though we share the same end goal. Similarly, we expect the desired variants for the same processes to differ between organizations, teams, and individuals.

There is no one benchmark for defining a desired variant for your processes, but here are some criteria that can help define and identify it:

Does it minimize risk?

A desired variant in security processes is often measured by a simple question — does it minimize risk? For example, do you require multi-factor authentication when onboarding a new member of the team? If you skip this step, does it introduce a security risk?

Can it be reliably repeated?

A process path that is both secure and efficient, but difficult to repeat consistently due to other factors, is unlikely to be followed consistently. When pressed against timelines, it's human nature to take the easiest path. Setting an unrealistic path for a process is a recipe for failure.

Does it finish the task in minimal steps?

Let's assume you found a path that minimizes risks. It is also important to consider how many steps are required to complete the task. A path that requires too many redundant or unnecessary steps will cost in other ways — we'll learn more about this in the next section.

Does it introduce other bottlenecks?

If your desired variant introduces bottlenecks in your overall strategy, then it probably isn't an ideal flow, regardless of how secure and time efficient it is. For example, you can directly deploy a vulnerability fix to a running application by updating packages, rather than pushing the code to a repo and waiting for approvals before deployment. Despite the fact that this does not entail any direct risk and fixes the vulnerability using minimal steps, it may cause problems in your development process, which requires quality assurance before apps are deployed.

The desired variant of the process should be consistent and efficient, with minimal steps. It should be easily repeatable and should not introduce any bottlenecks. Following a desired variant can help organizations maintain quality and accuracy, and ensure efficient and reliable processes.

Risks

The term “risk” in cloud security probably conjures up images of security alerts. Over-privileged permissions, for instance, are a risk in identity management. Vulnerabilities pose a risk to application security. Historically, risk is associated with an event or alert rather than with the underlying process which led to the issue. Identifying potential threats in your environment is important, but the security team’s job covers more than just identification.

In most cases, an alert is just the beginning of a process. An alert is meant to raise awareness and trigger a series of tasks that should culminate in its remediation or dismissal.

Focusing on threat or vulnerability detection techniques without understanding the underlying processes is like watching and following workout videos on YouTube and wondering why you aren’t getting same results. Unless someone watches your movements and provides guidance, you can’t be sure if you’re doing the exercises correctly. Skilled guidance makes a big difference.

Security operations are similar — if you can’t see how you’re doing the process overall you can’t see how to improve it, or how to do it more consistently.

Inefficiencies

According to a report from Formstack and Mantis Research, on average companies lose $1.3 million a year due to inefficient tasks weighing on employees. Inefficiencies are often not readily apparent, which makes them difficult to detect. In many security metrics, data is only collected on whether a task has been completed. For example, whether a vulnerability has been remediated and whether an employee has been fully off-boarded. As a result, inefficiencies in the process flow often go undetected.

Organizations can reduce their costs and improve the operational efficiency of their security teams by discovering and correcting non-conformant variations, leading to time and resource savings.

Let’s go back to our vulnerability remediation example. Consider any variant that does not align with the desired variant. In the process shown in Fig 3.3, the desired variant shows the following metrics:

- **Scan time:** 2 minutes to scan the application
- **Ticket assignment:** 1 day to assign the ticket
- **Remediation:** 7 days to remediate the vulnerability

Total time to complete the process when following the desired variant = 8 days

When the ticket is assigned to the wrong team instead of the application owner, the process is significantly longer:

- **Scan time:** 2 minutes to scan the application
- **Ticket assignment (to wrong team):** 1 day
- **Ticket assessment:** 3 days before the ticket is sent back for re-assignment
- **Ticket reassignment:** 7 days to assign the ticket to the right owner
- **Remediation:** 7 days to remediate the vulnerability

Total time to complete the process in the event of a wrong assignment = 18 days

Delays accumulate when a ticket is assigned to the wrong team, averaging 3 days for the ticket to be

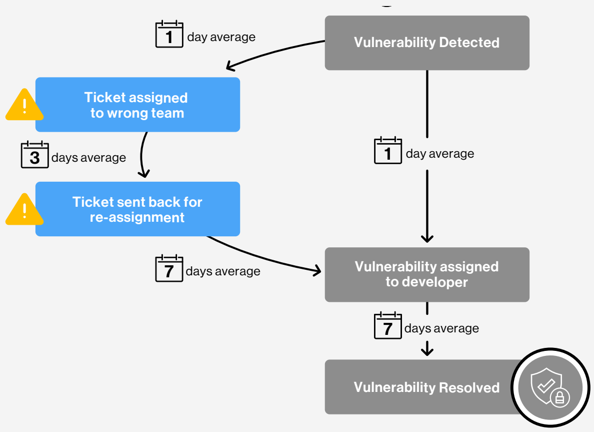

FIG 3.3 INEFFICIENCY ADDED BY PROCESS VARIANT

validated and sent back for re-assignment to the correct team, which can add 7 days to the overall process. All in all, the additional steps add up to a delay of 10 days on average. It's like trying to get from Point A to Point B while having to drive around the block a few times. The delays add up, eventually resulting in far more time being spent than necessary.

Process inefficiencies in security can lead to increased attack surfaces and exposure windows. Delays in identifying and remediating vulnerabilities can lengthen the time malicious actors can exploit these weaknesses. To boot, such delays also increase the cost and complexity of the security process.

Enabling Process Improvement

The first step to securing your resources is having visibility into all of your resources. It's impossible to protect what you can't see. Processes are similar. You cannot identify risks and inefficiencies in processes if you cannot observe their origins, the paths they follow, or the time elapsed between steps. Process mining provides a data-driven approach that can help identify process inefficiencies, delays, and bottlenecks.

Process mining can reveal gaps and offer insights into where process improvements can increase efficiency. If assigning tickets to the wrong team was found to result in delays, an automated step to validate the application's owner should be added prior to assigning the ticket. In our example, adding such an automated step would add a mere 3 minutes to the overall process time, which already averages 8 days. However, the automation step helps to quickly identify the appropriate team to assign the ticket to, reducing the amount of time needed to manually assign a ticket to the right developer. This helps to reduce the overall amount of time needed for the process significantly, from 18 days to 8 days and 5 minutes. Carefully adding well-planned steps to the process can drastically increase its efficiency and reduce the total time to execute.

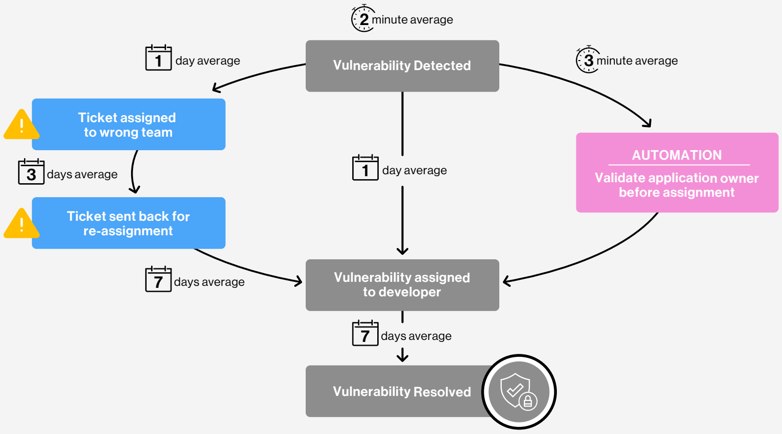

FIG 3.4 PROCESS IMPROVEMENT DISCOVERY

Process Understanding Helps Drive Security Maturity

There are several ways process mining can help with assessing your organization's security maturity:

Assessing the Value of Your Technology Investments

A tool is an instrument used to deliver a result. Understanding the realized value of tools is crucial to evaluating your return on the investment from them. However, lacking processes surrounding tools can lead to their under-utilization and even tools becoming bottlenecks. Process mining can help you measure the value of your security investments and understand how effectively you've integrated your tools with people and with other tools. Often, organizations don't realize the potential value of what they've bought because the process around a tool isn't well defined or consistently followed.

Looking Beyond the Dashboards

Many security tools have some kind of dashboarding capability but these dashboards are usually narrowly focused only on specific data points covered by the tool. As discussed, nearly everything in security is a process involving multiple tools and teams. Thus, to understand how your organization is really performing, you need to understand how the processes are performing, across the different teams, people, technologies, and tools involved in the processes. Process mining can give you a deeper, data driven understanding of the overall outcomes your organization is delivering, helping you orient and track performance more comprehensively.

Improving Compliance

Organizations can implement best practices for security by following compliance standards like NIST, CIS, and other industry benchmarks. Failure to meet compliance standards introduces insecure configurations and increases the attack surface. CSPM tools can perform a compliance scan and show you which individual settings you failed to adhere to, while analyzing your process flow can show how you got there.

Tracking Internal Goals

To achieve a goal, you must be able to measure its progress. Process mining is particularly helpful in looking at goals from an end-to-end perspective, helping you focus not just on the outcome but how it's being delivered.

The following are metrics often used to measure the effectiveness of security processes:

Conformance

How consistently do you follow the desired path of a process? For example, if you execute a vulnerability remediation process 10 times, but only follow the desired path or target variant 5 times, you have a 50% conformance rate.

Risk

What is the degree of risk introduced by process variance? For example, when offboarding a user, if user access keys are deleted only 8 out of 10 times, 20% of cases introduce risk.

Efficiency

Calculating process efficiency allows you to identify areas for improvement, reducing costs and saving time. A common measurement for efficiency is a comparison of the execution time of a variant relative to that of the target variant.

Compliance standards

Instead of simply looking at system configuration settings when comparing processes against industry benchmarks, you can now understand their causes. For example, NIST SP 800-53 recommends monitoring and responding to logon attacks. Many tools can check across your identity systems to see whether you have particular settings in place to mitigate risk against these attacks. However, only a process mining approach can show how these settings work in concert with your attack detection tools, SOC team, and other components of the overall attack response workflow.

WHEN USED PROPERLY, A PROCESS MINING TOOL CAN PROVIDE MEASURABLE ANSWERS TO QUESTIONS SUCH AS:

- What was the number of times the process failed to achieve its goals?
- Which variants of a process resulted in added delays?
- How many variants of a process resulted in indirect risks for the organization?
- During the execution of a process, what step was most often missed?
- In order to improve efficiency, what steps could be automated?

Process understanding is necessary for improving and evolving your organization's overall security. By using a process mining tool, you can quantify the efficiency of how your tools and teams work together using data-driven metrics. Furthermore, it assists with identifying risks that have been introduced into your environment. Our next chapter will cover case studies on risk identification using process mining.

05

Case Studies for Risk Analysis with Process Mining

CHAPTER 5

In previous chapters, we discussed how process mining can provide security teams with actionable insights. To explore the concept in more detail, we will use real-world examples of security processes in this chapter.

Let's recap some highlights from the previous chapter as we begin.
In order to deliver on security teams' process outcomes, it's vital to monitor:

Conformance rate

The conformance rate is the ratio between the number of times the process adhered to the ideal, desired process and the number of times it failed to do so. Ideally, every security operation would always follow the right process to reduce risks and operate efficiently. With a process mining tool, it's possible to analyze how frequently the right process is followed and when it deviates from the expected flow. By tracking this metric over time, you can measure the overall conformance of your processes.

Efficiency

While not all delays cause inefficiency, some delays are significant to the overall process and caused by missteps that could be avoided. To identify bottlenecks, it is necessary to have visibility into the entire process along with its average duration and compare it to the duration of cases that are delayed due to mistakes.

Risk introduced by a variant to the process

When a variant results in a security risk to the organization.

Additionally, process mining can provide additional data points, such as:

The most common undesired variants of a process

By clearly mapping variants, process mining highlights which steps are missed most often, which manual steps are responsible for the longest delays, and which steps could be automated for the most ROI.

Case Study: Offboarding a Privileged User

When an employee leaves an organization, offboarding is the process of removing the employee's access to systems and data. This process is particularly important if the user has elevated privileges, such as being an administrator or root user. The security team must ensure that the former employee does not continue to have access to sensitive data and systems. Offboarding is a complex process that involves multiple systems and teams, and process mining therefore offers unique insights and advantages that can be used to identify risks in such complex workflows.

Common Systems and Steps

Organizations typically grant access to tools and systems via enterprise directories, SAML providers, and cloud platforms. Offboarding aims to prevent terminated employees from retaining any access to the organization's assets after termination.

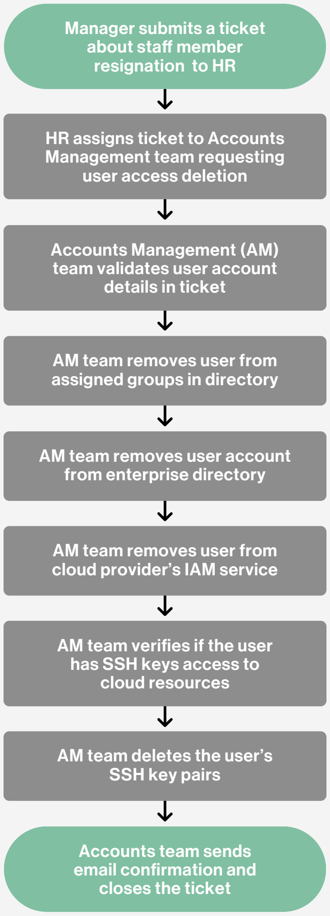

Manager submits a ticket about staff member resignation to HR
HR assigns ticket to Accounts Management team requesting user access deletion
Accounts Management (AM) team validates user account details in ticket
AM team removes user from assigned groups in directory
AM team removes user account from enterprise directory
AM team removes user from cloud provider's IAM service
AM team verifies if the user has SSH keys access to cloud resources
AM team deletes the user's SSH key pairs
Accounts team sends email confirmation and closes the ticket

FIG 4.1 OFFBOARDING PROCESS IDEAL FLOW

HERE IS AN EXAMPLE OF AN OFFBOARDING PROCESS IN FIG 4.1:

1. Manager submits a ticket to HR about the resignation of one of his staff members.
2. HR assigns the ticket to an account management team requesting deletion of the user associated accounts across all systems.
3. Accounts management team removes the user from security groups in the directory.
4. Account management team removes the user account from the directory.
5. Account management team removes the user from the cloud providers' identity and access management systems.
6. Account management team checks to see if the user has SSH keys to access systems in cloud providers.
7. If so, they also delete the user's SSH key pairs from the cloud providers.
8. The accounts management team closes the ticket confirming the user has been removed.

While the organization may believe that the desired process is followed accurately and consistently, it may not realize the extent of cases that follow undesired variants and introduce unexpected risks. In most cases this does not happen because of the failure of a tool but because of the failure to follow the ideal process. Without visibility into the process itself, these failures could go undetected.

To successfully accomplish the goal of a process, both the tools and people assigned to use these tools must remain coordinated and work cohesively to follow the intended process each time before the offboarding date is reached. A good process mining tool can provide a data driven view to show where the bottlenecks are and, more importantly, what causes them.

Residual Access Variant

In Fig 4.2, we see an example where the accounts management team skipped an offboarding step and failed to check whether the user had SSH keys to access the cloud service provider. Due to this oversight, the terminated employee was still able to access cloud resources post- termination as the SSH keys remained in the cloud service provider. Since this step is not caused by any error in a tool, no alerts were generated. It is therefore possible for this risk to go undetected unless a full audit is conducted.

There are many metrics that can be obtained from a good security process mining tool when applied to monitoring an offboarding process:

Conformance rate

Although the process for offboarding an employee was executed 100 times, verification of SSH access was only executed 90 times. The conformance rate for this process is 90%.

Risk

Of the 100 requests for employee offboarding assigned to the accounts management team, only 90 cases followed the right procedures for verifying SSH key access. Variants introduce 10% risk to the process.

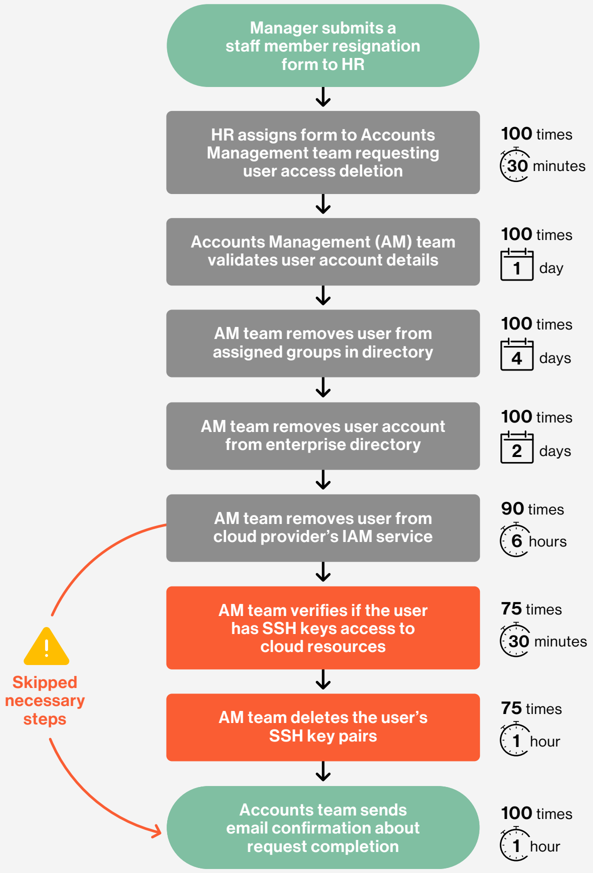
Manager submits a staff member resignation form to HR
HR assigns form to Accounts Management team requesting user access deletion
100 times
30 minutes
Accounts Management (AM) team validates user account details
100 times
1 day
AM team removes user from assigned groups in directory
100 times
4 days
AM team removes user account from enterprise directory
100 times
2 days
AM team removes user from cloud provider's IAM service
90 times
6 hours
AM team verifies if the user has SSH keys access to cloud resources
75 times
30 minutes
Skipped necessary steps
AM team deletes the user's SSH key pairs
75 times
1 hour
Accounts team sends email confirmation about request completion
100 times
1 hour

FIG 4.2 OFFBOARDING PROCESS SKIPS SSH KEYPAIRS' REMOVAL

KEY LEARNINGS

As a result of a variant in the process, 10 staff offboarding tickets did not undergo a verification step to delete SSH keys. This posed a risk to the organization since the terminated employees could still access the cloud provider with SSH keys.

Missing Validation Variant

In some cases, despite following all the necessary steps in the process and removing all access, additional unintended steps were added due to missing important information required to offboard an employee which caused additional delays. As a result, the process is inefficient, and the risk of a terminated employee still having access to the system during this delay increases. Let's take a look at the variant in Fig 4.3.

A process mining tool can help measure metrics such as:

Number of times a pattern occurred

The accounts management team sent back 20 tickets, quoting insufficient employee identification information was provided by the HR team.

Additional delays

The process was delayed by an average of two days.

Risk introduced

On average, terminated employees had 2 additional days of access to secure cloud accounts after leaving the company.

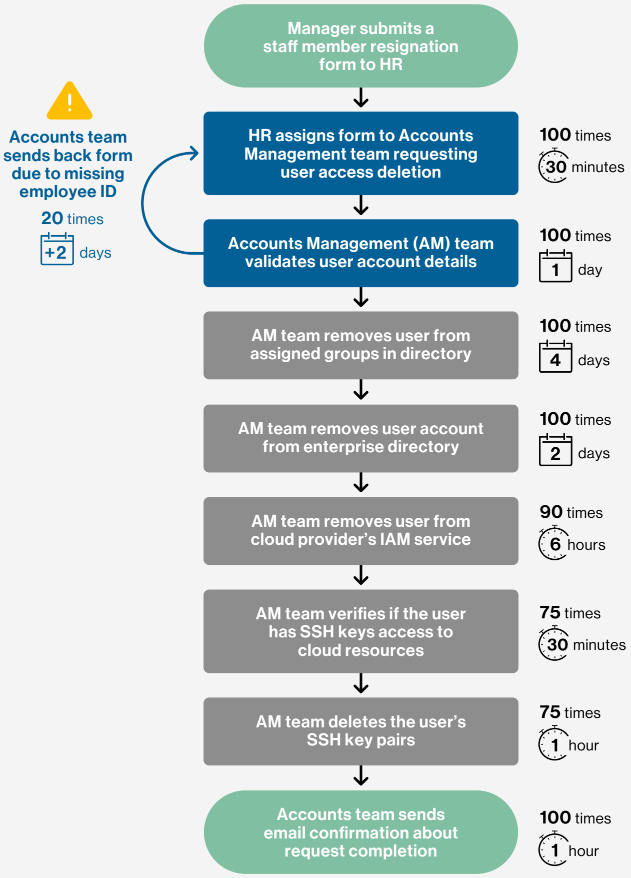
Manager submits a staff member resignation form to HR
HR assigns form to Accounts Management team requesting user access deletion
100 times
30 minutes
Accounts team sends back form due to missing employee ID
20 times
+2 days
Accounts Management (AM) team validates user account details
100 times
1 day
AM team removes user from assigned groups in directory
100 times
4 days
AM team removes user account from enterprise directory
100 times
2 days
AM team removes user from cloud provider's IAM service
90 times
6 hours
AM team verifies if the user has SSH keys access to cloud resources
75 times
30 minutes
AM team deletes the user's SSH key pairs
75 times
1 hour
Accounts team sends email confirmation about request completion
100 times
1 hour

FIG 4.3 VARIANT INTRODUCING ADDITIONAL DELAY IN THE OFFBOARDING PROCESS

KEY LEARNINGS

In this example, a process inconsistency caused a delay that gave terminated employees 2 additional days of access to secure cloud assets after finishing their last day on the job. This introduced risk that a terminated employee might exploit their access to exfiltrate data. A process mining tool can help identify such delays and show how they might be avoided by adding automated validations to the HR form, to check that mandatory details such as the user's employee ID and email address are not left out.

Case Study: Business Email Compromise

Let's take a look at how an organization addresses business email compromises in their cloud email system. If suspicious activity is detected on an email account, the SOC team is notified. The SOC team then takes a series of steps to investigate the incident and respond if necessary.

Common Systems and Steps

A process for resolving email security incidents typically involves email platforms, ticketing systems, mail gateways, and directory services.

Let's take a look at one such example process in Fig 4.4.

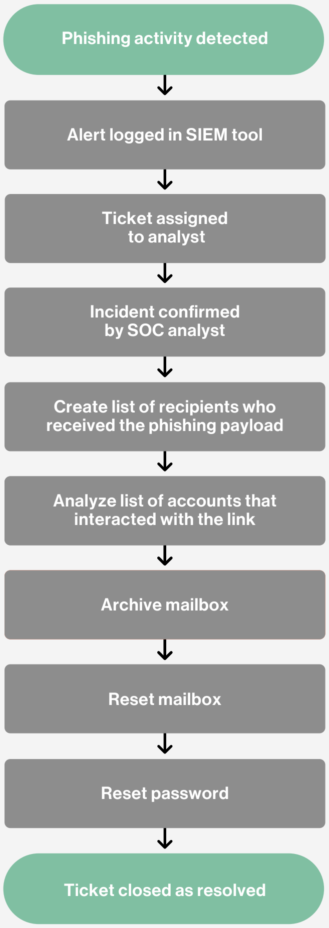
Phishing activity detected
Alert logged in SIEM tool
Ticket assigned to analyst
Incident confirmed by SOC analyst
Create list of recipients who received the phishing payload
Analyze list of accounts that interacted with the link
Archive mailbox
Reset mailbox
Reset password
Ticket closed as resolved

FIG 4.4 BUSINESS EMAIL COMPROMISE IDEAL PROCESS

THE STEPS IN THIS PROCESS ARE:

1. Phishing activity is detected by the cloud provider's mail security system.
2. An alert is logged by the SIEM tool.
3. A ticket is generated by the SIEM tool and assigned to the SOC team.
4. A SOC analyst triages and confirms the incident.
5. As part of the remediation process, the SOC analyst creates a list of people who received the phishing email as recipients.
6. The SOC analyst further narrows down the list to those users that clicked on the link and interacted with the phishing payload (assuming user interaction is required).
7. The SOC analyst archives mailboxes of users affected by the incident.
8. The affected mailboxes are reset.
9. Passwords for all affected accounts are reset.
10. Phishing-specific security training is provided to affected users.
11. Ticket is closed as resolved.

Incomplete SOC Response Variant

Sometimes, the process for remediating incidents is not followed as it should. The variant shown in Fig 4.5 leaves the incident only partially resolved.

A good security process mining tool can provide metrics such as:

Conformance rate

When email incidents occur, the organization executes its desired response plan only 79% of the time.

Number of times the process failed

There were 40 cases where SOC analysts failed to reset passwords when they responded to an incident.

Impact of process failure

620 tickets were opened for email compromises, but only 490 lead to password resets. Where the passwords remained unchanged, the organization was more vulnerable to attack.

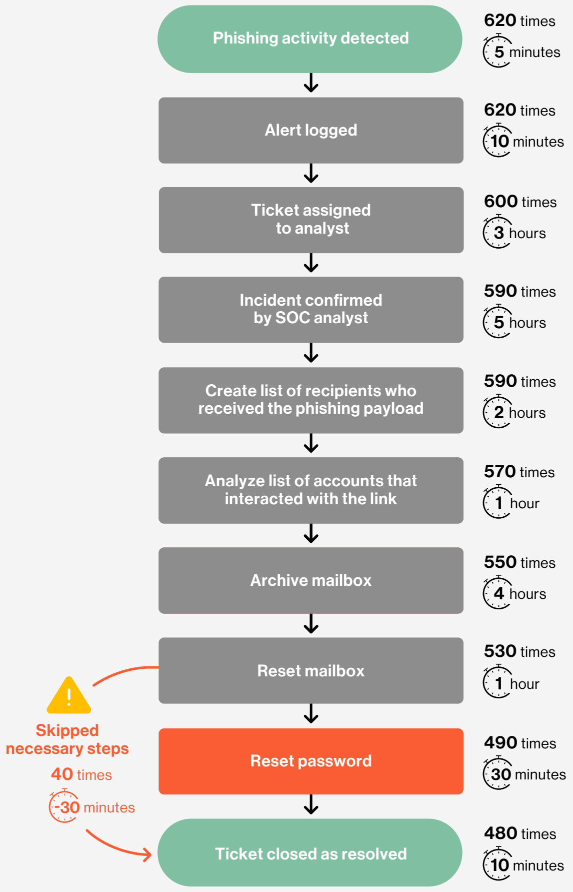

FIG 4.5 BUSINESS EMAIL COMPROMISE PROCESS VARIANT THAT LEAVES A RISK

KEY LEARNINGS

This variant shows that sometimes the SOC analyst will reset the affected mailboxes without resetting the passwords on affected accounts. If the malicious actor managed to gain access to the passwords through the phishing attempt, this leaves a vulnerability that can be exploited later. This failure is not the fault of the tools being used; the scanner detected the malicious activity, and the ticketing system created an alert successfully. However, the incident remains only partially resolved because a critical step in the process was missed.

A near-real-time view of end-to-end processes can help to drive process improvement. Teams can use the data to identify the most frequently missed steps in the process and add guardrails to prevent them from occurring again.

Inefficient Variant

Let's take a look at another interesting variant of the same process, shown in Fig 4.6, which introduces unnecessary additional effort to the process. The ideal process for resolving an email incident is for the SOC analyst to filter down on impacted accounts and only work on securing those accounts that were actually affected. The purpose of this step is to help the SOC analysts focus their efforts only on areas that require their attention.

When this validation is not performed, the SOC analyst resets all of the email accounts that received the phishing email, even if they were not compromised. Aside from the money, time, and effort invested in archiving accounts unnecessarily, this delay puts impacted users at risk by prolonging their exposure. Additionally, it inconveniences other users who did not interact with the phishing payload or have had their accounts compromised but still need to wait for their accounts to be reset.

A process mining tool can:

Display variance

Show the step that deviated from the desired flow and resulted in an additional delay of 11.5 hours.

The number of times this occurred

With the process mining tool, it's easy to view how many times this step was missed — 50 times.

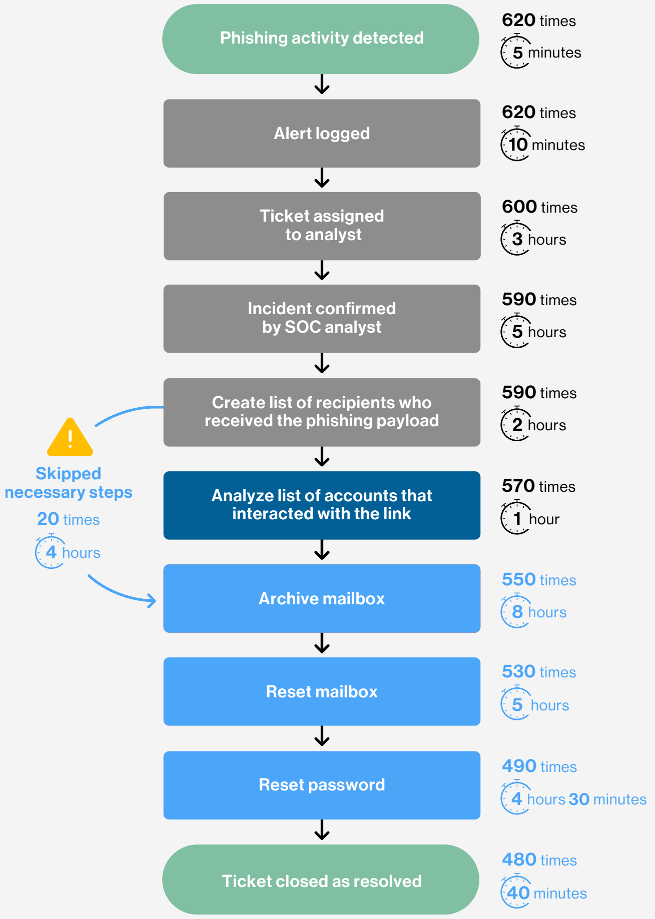

FIG 4.6 BUSINESS EMAIL COMPROMISE PROCESS VARIANT THAT ADDS UNNECESSARY DELAY

KEY LEARNINGS

Without process mining, it would be difficult to expose the cause for the delay, namely a failure of the SOC analyst to narrow down the impacted users. This misstep can go undetected and leave users frustrated. With process mining, you can see how many times this step was missed and measure the delay it caused to the process and ensure that this process deviation does not occur in the future.

Case Study:
Externally Reported Vulnerabilities

Today, almost all organizations create software, to serve their constituents, sell to their customers, or manage their supply chains. It is increasingly common for organizations to have vulnerability reporting programs that encourage external researchers to responsibly disclose their findings. Here, we examine how a company triages and remediates externally reported vulnerabilities.

Common Systems and Steps

Vulnerability reporting programs generally involve vulnerability disclosure platforms, ticketing systems, code repositories where software is stored, and cloud service providers where the app runs.

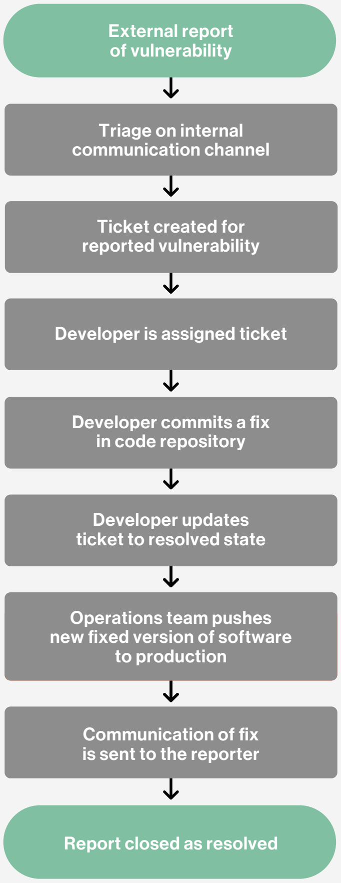
External report of vulnerability
Triage on internal communication channel
Ticket created for reported vulnerability
Developer is assigned ticket
Developer commits a fix in code repository
Developer updates ticket to resolved state
Operations team pushes new fixed version of software to production
Communication of fix is sent to the reporter
Report closed as resolved

FIG 4.7 DESIRED PROCESS FOR EXTERNALLY REPORTED VULNERABILITY

EXAMPLE OF A VULNERABILITY EXPOSURE PROCESS IN FIG 4.7.

1. An external researcher submits a report using a vulnerability reporting platform.
2. The organization's response team triages the report and communicates about it over an internal chat channel.
3. After the report has been validated, the team opens an internal ticket for the reported vulnerability.
4. A developer is assigned to the ticket.
5. The assigned developer works to remediate the vulnerability and submits a patch to the code.
6. The operations team pushes the fixed new version of software to production.
7. The developer closes the internal development ticket as resolved.
8. The response team updates the external researcher about the fix using the vulnerability reporting platform.
9. The external report is closed as resolved.

Unfixed Vulnerability Variant

Consider the following process variant, shown in Fig 4.8. Here the developer assigned to the ticket closes the ticket after committing a fix to the code repository and does not wait for the operations team to deploy the fixed version to production. As it so happens, the fix is never deployed, hence the report is closed as resolved while the vulnerability remains in production. A process mining tool would depict this variant as a sequence deviation that is also missing a step.

A security process mining tool can provide insights such as:

Process failure

Before examining the data, we might expect this process to be followed consistently. However, the data reveals that of the 350 times that an external vulnerability was reported, a fix was pushed to production only 220 times.

Conformance rate

Only 77% of cases adhere to the target process.

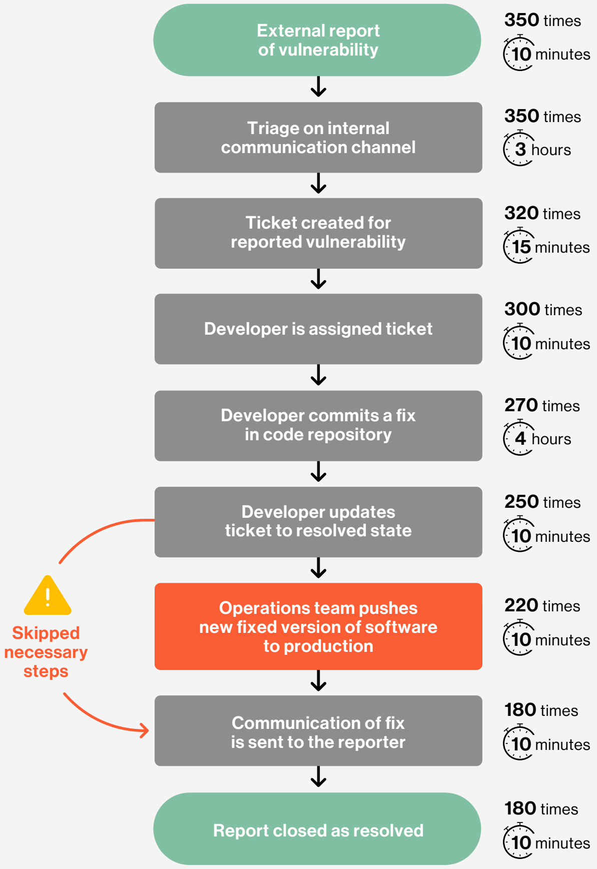
External report of vulnerability
350 times
10 minutes
Triage on internal communication channel
350 times
3 hours
Ticket created for reported vulnerability
320 times
15 minutes
Developer is assigned ticket
300 times
10 minutes
Developer commits a fix in code repository
270 times
4 hours
Developer updates ticket to resolved state
250 times
10 minutes
Operations team pushes new fixed version of software to production
220 times
10 minutes
Skipped necessary steps
Communication of fix is sent to the reporter
180 times
10 minutes
Report closed as resolved
180 times
10 minutes

FIG 4.8 PROCESS VARIANT FOR EXTERNALLY REPORTED VULNERABILITY

KEY LEARNINGS

Process failure can lead to a false sense of security. In this example, the tools were all working correctly yet the vulnerability remained unresolved. To make matters worse, the organization was misled by the fact that the development ticket was closed as resolved. This can result in the same vulnerability report being reopened after failing validation by the reporter, or worse, it can result in a threat actor exploiting the vulnerability.

Lengthened Exposure Window Variant

Consider another variant of the same process. In Fig 4.9, we see cases that suffer delays which leave the system vulnerable for longer periods of time.

A process mining tool will uncover the following insights for this variant:

Delay caused by human error

We see that in 20 cases, after the ticket was opened, instead of assigning the ticket to the team responsible for fixing the vulnerability, the ticket is assigned to an unmonitored queue. Thus, the ticket goes unnoticed until the next sprint planning, adding a delay of 27 days until the ticket is assigned to the right team.

Impact of process failure

The entire process is delayed by 27 days. This not only introduces inefficiency, but also leaves the system vulnerable during this period.

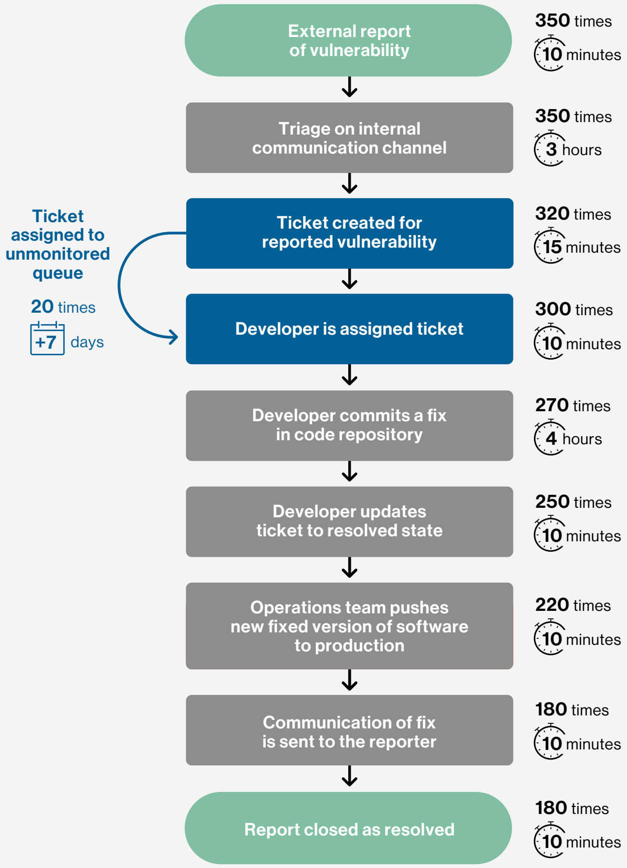
External report of vulnerability
350 times
10 minutes
Triage on internal communication channel
350 times
3 hours
Ticket created for reported vulnerability
320 times
15 minutes
Ticket assigned to unmonitored queue
20 times
+7 days
Developer is assigned ticket
300 times
10 minutes
Developer commits a fix in code repository
270 times
4 hours
Developer updates ticket to resolved state
250 times
10 minutes
Operations team pushes new fixed version of software to production
220 times
10 minutes
Communication of fix is sent to the reporter
180 times
10 minutes
Report closed as resolved
180 times
10 minutes

FIG 4.9 PROCESS VARIANT FOR EXTERNALLY REPORTED VULNERABILITY RESULTING IN LENGTHENED EXPOSURE WINDOW

KEY LEARNINGS

Introducing automation to alert on tickets that haven't been worked on for over a day can reduce delays and surface such misaligned tickets faster.

Conclusions

Although each of the case studies had a desired process defined, this desired process was not followed consistently. When dealing with complex security processes traversing many different systems and teams, such errors are common and can easily go unnoticed.

WHAT WE'VE LEARNED FROM THE REAL WORLD SCENARIOS EXAMINED SO FAR:

1. In reality, the desired workflows for security processes are not followed consistently.
2. Even with the best security tools in use, risks can be introduced if the process surrounding these tools is not well-defined or not enforced.
3. More tools do not necessarily provide better security. Tool sprawl and inconsistent use can become a hurdle, rather than an advantage for the personnel involved in the process. A process mining tool can correlate data from tools and systems to automatically highlight process inconsistency.
4. Many process variants involve missed steps or steps that could be automated. Process mining can help identify where investing in automation will provide high returns.

Now that we've seen examples of the impact a process mining tool can bring to security operations, in the next chapter we will highlight what characteristics a security focused process mining tool should have.

06

Key Characteristics to Look for in a Security Process Mining Tool

CHAPTER 6

In previous chapters, we discussed the importance of applying process mining to your security operations. Now that you understand the benefits of process mining for security, you may be considering using it in your own organization.

By using the right process mining tool, you can gain detailed insights into your organization's processes, identify risks, find areas to automate to improve efficiency, and measure overall conformance. Process mining for security processes differs from process mining for non-security processes because it needs to have an awareness of information security risk.

Thus a security process mining tool will need to have the following characteristics:

Broad Library of Integrations for Security and Non-Security Tools

As noted in Chapter 2, security processes often involve both security tools, such as vulnerability scanners, and non-security tools, such as ticketing systems. For process mining to be effective, it must have a wide range of integrations not only for security tools, but also non-security tools. When direct integration with the system is not possible, there should be an automatable API and file upload capability to get event data ingested.

Automatic Risk and Efficiency Analysis

Security processes have different goals than non-security processes and their impact is expected to be different. Whereas unresolved software vulnerabilities can result in compromised data and breaches, failure to follow the right steps in customer support cases may lead to dissatisfied customers. Both processes are certainly important for the organization, but they require different kinds of understanding and analysis.

When looking for a security process mining tool, you should look for one built specifically to analyze security processes for risks and inefficiencies. Security best practices must be embedded in the tool for it to be able to find these risks. For example, the tool should be able to identify cases where missed or delayed steps create security risk for the organization.

Built-In Analytical Capabilities to Drill into Process Data Directly

Visualization is a fundamental capability, but a security process mining tool must also enable users to interact directly with the data for further analysis.

Library of Security Benchmarks and Standards for Comparison

Industry guidelines, such as NIST special publications, CIS benchmarks, SOC2, PCI, and HIPAA compliance standards provide organizations with standards and best practices. To compare internal processes against these industry standards, process mining tools must have an internal library of security benchmarks and provide the ability to compare your actual processes to them.

Low Friction Implementation

To show the full end-to-end process, a process mining tool needs to integrate with multiple tools and systems. It is essential that the tool offer a low-friction implementation and maintenance process. It should be a resource for security organizations, not another bottleneck. For example, whenever a new tool is introduced into your security process, connecting it to the process mining tool should be as simple as adding credentials.

IDEALLY, A SECURITY PROCESS MINING TOOL SHOULD ENABLE ANALYSES LIKE THESE:

1. Conformance measurement against a desired variant.
2. Comparison of multiple variants overlaid or side-by-side.
3. Risk and inefficiency identification in variants.
4. Sorting and filtering capabilities. Such as by frequency, rate of occurrence of a step, and date of past process execution.
5. Exporting of graph data for analysis and usage in other tools.

Least Privilege Access

It is important for a security process mining tool to ingest only data specific to the process it is mining. For example, when analyzing privileged user offboarding, process mining tools do not need access to everyone's personal details like home addresses or phone numbers. A security-focused process mining tool should be able to identify the minimum metadata needed to correlate Active Directory activity with other platforms and ingest only the metadata necessary for such correlation.

Foundational Understanding of Security Processes

To be able to ingest metadata that is specific to a security process, the process mining tool must have a foundational understanding of what each process would generally look like. It must be able to distinguish between the different security processes like access management, vulnerability scanning, incident response, and be able to understand the identifiers associated with each, so it can ingest only the least privileged permissions and metadata. This foundational understanding of security processes is what separates a security focused process mining tool from a generic process mining tool.

Strong Security with Accredited Certifications

Because a security process mining tool will be ingesting sensitive data, such as metadata about vulnerabilities, incidents, and employee details, it is crucial that the tool maintains high standards of privacy and security. Accreditation with well known standards like SOC 2 and ISO 27001 is a good validation of a process mining tool's security controls.

Continuous Data Ingestion

In order to stay ahead of rapid changes in environments and user activity, a process mining tool should support continuous data ingestion. Continuously ingesting data enables near-real-time insights into your processes, so you can improve efficiency and take action quickly. For example, continuous data ingestion allows you to recognize when an automated process has failed or a manual process has deviated from its standard operating procedure, alerting you to take action or investigate the issue.

In the next chapter, we summarize what we've learned throughout the book and how a process mining tool can influence your organization's security goals.

07

Conclusion

Hopefully, you've gained a deeper understanding of process mining, how it can be applied to security, and how it can improve the outcomes that security organizations deliver. This chapter gives a quick synopsis of the book, highlighting key learnings from each chapter.

Applying Process Mining to Security

In Chapter 1, we discussed the importance of looking at security from a process perspective.

Today's security thinking is often focused on the technical capabilities of security tools. Customers are left to design and optimize processes that orchestrate the tools along with the people that use them. Organizations can maximize the value of their investments in tools and people by focusing on how they work together to form processes.

Process mining is the method of extracting readily available data from the event logs of information systems to discover, map, monitor, and improve processes. Process mining provides data-driven insights by analyzing event data from IT systems. It shows how things ideally work much of the time, and exposes all of the process variations when things work differently. These process variants can help you identify hidden risks and inefficiencies in your business so you can improve your overall security.

Process mining drives security improvements by helping to determine if people are following proper procedure and using the provided security tools as expected. In doing so, process mining can help to ensure that security tool investments are optimized and justified and that security operations are being conducted appropriately and efficiently.

As a methodology, process mining can be a valuable addition to any organization that wants to assess its overall security operations, measure its outcomes, and identify areas for improvement. When organizations invest in the right security process mining tool, they can save time and money by quickly identifying and addressing non-conformant security processes, and make progress towards improving their security maturity and compliance.

How Does Process Mining Work?

Chapter 2 details how process mining works, how data is ingested, and what factors are critical to visualizing these processes.

As shown, every activity that occurs in an IT system leaves a record that we call an "event".

A process mining tool ingests these events and categorizes them by unique identifiers that it uses to correlate event activity between different systems. Using timestamps, it organizes this data into a process flow mapping to visually depict the different steps involved in achieving the desired end goal.

Since this visualization is based on data and not human observations, process mining can provide real insights into multiple variants of a process, while simultaneously identifying bottlenecks and opportunities for improvement. Due to the nature of security processes, these gaps can result in inefficiencies and palpable risks to the organization. For example, any delay in remediating vulnerabilities can increase the risk of the software being exploited.

A wide variety of data types must be incorporated into security process mining. Many security processes involve both security tools and non-security tools. Consider one such process: vulnerability management. This involves a security tool, such as a vulnerability scanner, which sends out alerts of vulnerabilities. It's important to understand that to take action on the alert and remediate the vulnerability, you would use non-security tools such as ticketing systems and code repositories.

Hence, it is important for a process mining tool to be able to integrate with a wide variety of security and non-security tools. In addition, it should support the ingestion of data from open file formats when direct integration isn't possible.

While a process mining tool may ingest data from multiple tools, it is important to note that it is designed not to be a security data lake nor a posture management tool. A security-focused process mining tool is purely focused on metadata relevant to tracking activities related to processes.

Actionable Insights from Process Mining

In chapters 3 and 4 we see the application of process mining on security governance and how process mining tools can provide actionable insights for security leaders.

All these insights are driven by data, including risk awareness, inefficiency measurements, return on investment KPIs for tools, and security maturity. Utilizing process mining tools can help to both track internal goals and compare performance against industry standards.

The following are some actionable insights that a security process mining tool can provide:

Risk

Process mining can measure the risk impacting a process due to process variance. For example, if deleting user access keys is a vital step in offboarding a user, non-conformant cases introduce needless risk.

Conformance

How consistently do you follow the desired path of a process? For example, if you execute a vulnerability remediation process 10 times, but follow the desired variant only 5 times, you have a 50% conformance rate.

Inefficiency

Calculating process efficiency allows you to identify areas for improvement or automation where you can reduce costs and time. A common efficiency measurement compares the median case duration (or execution time) of a variant to the median duration of the desired variant.

Compliance

Beyond simply looking at the configuration of settings when comparing processes against industry benchmarks, we can surface the causes of non-compliance. For example, NIST SP 800-53 recommends monitoring and responding to logon attacks. Many tools can check to see whether the correct settings are in place across your identity systems related to these attacks. However, only a process mining approach can show how these settings work in concert with your attack detection tools, SOC team, and other components of the overall attack response workflow.

Variant Analysis with Process Mining

In Chapter 5, we saw the application of process mining to real world security processes.

We explored three security processes:

- Offboarding a privileged user
- Responding to business email compromise
- Remediating externally reported vulnerabilities

In each of these we see how lack of visibility into the process and raw case data left important questions unanswered. These include but are not limited to the cause for an open risk, the unnecessary delay caused by a misstep, and the additional effort wasted on correcting process mistakes. Without performance indicators, the overall outcome of security processes will be negatively impacted. With process mining, we learned you can measure outcomes and improve your overall strategy to optimize the available tools and staff already within your organization to gain more efficiency and decrease friction.

Characteristics of a Security-Focused Process Mining Tool

In Chapter 6, we highlighted some of the necessary characteristics to look for when choosing a security process mining tool.

- A broad library of integrations for security and non-security tools
- Automatic risk and inefficiency analysis
- Built in analytical capabilities to drill into process data directly
- A library of security benchmarks and standards for comparison
- Strong security of the tool itself, with accredited certifications

Security cannot remain stagnant in an era of increasing digitization, evolving infrastructure, and demanding compliance requirements. In order to keep up with changing threats and business needs your processes must constantly evolve. Many organizations struggle to achieve their security goals despite having the best security tools and staff, simply because they never really understand the process. Process mining is the lens which helps organizations understand the reality of how they actually work and how to improve using data-driven insights and quality analytics.

TAKING THE FIRST STEP

To get started with process mining, you do not need to connect to every system from the get go. You can start small to reduce complexity and scale as you go. For example, if you want to start with less complexity, analyze a category of processes such as identity management. That way, you will only need to connect to systems associated with identity management, like your directory service and federation service, without needing to connect to tools that aren't part of your identity management processes, such as application scanners or firewalls. You'll get a chance to experience the advantages of process mining with minimal effort, and once you feel ready to scale your process mining operations, you can choose to expand and build your way up to analyzing more processes and systems.

Conclusion

Many organizations believe that failures in security outcomes are caused by their people or their tools. They believe that they can only achieve better security outcomes with the right tools and the right people. The simple truth, though, is that processes have the greatest impact on security outcomes. Under-resourced security organizations with great process discipline often outperform well-funded organizations that lack process discipline.

A company may have the best EDR tools to detect incidents or the best vulnerability management tools to find CVEs, but if they don't have a data-driven understanding of their processes, then alert fatigue, false positives, and a chaotic security organization will ensue.

Security operations are driven by a combination of people, processes, and technology. Process mining enables you to understand how they accurately interact and work together so that you can optimize them to consistently deliver better outcomes.

Glossary of Terms

Key terms defined in the context of process mining.

Process mining

A methodology involving the ingestion and analysis of event data from IT systems to provide data-driven insights into how a process is actually performed from a broader organizational viewpoint.

Process

Simply put, a process is "the way something is done." Within the scope of security operations, a process is a workflow involving people and tools with a designated security outcome. A process mining tool will visualize a process as a sequence of events depicted as a map.

Case

A single process execution. Each case is an instance of the process, for example, an incident report or a detected vulnerability. A process mining tool automatically constructs cases from event data and catalogs cases by process variants.

Event

A record of activity in a tool or a system. Examples include a ticket status change in a ticketing system, an alert sent over a chat message, or a DEV task opened by a SOC team member.

Variant

A unique process flow. Each variant reflects a particular process path, or sequence of events. A process mining tool automatically analyses case data to identify the different process variants and group the cases accordingly.

Desired variant

The expected workflow. The desired variant is expected to be followed consistently.

Conformance

Rate of adherence to the desired variant of the process Cases that follow the target variant are considered to be conformant.

Risk

Risk may be introduced to the organization due to process variation.

Efficiency

A measure that compares the execution time and/or path of a variant to that of the desired variant.

Variation tolerance

The permitted range of deviations from the desired process, i.e. standard process flow.

Ingestion

The automated consumption of data from systems. Ingestion can be achieved by integrations, custom APIs, and manual data uploads.

Data sources

The tools and systems that a process mining tool connects to and harvests data from. Examples include vulnerability management tools, EDRs, cloud platforms, and ticketing systems.

ISBN 979-8-218-28515-9

90000>

9 798218 285159