

Evidence-centric threat hunting

How data fuels informed decisions
and AI/ML models to empower a
new generation of threat hunters



Introduction

As threats become more sophisticated, so do the cybersecurity technology innovations that help us prevent, detect, and respond to them. While these tools and techniques provide valiant security analysts and threat hunters the chance for a winning advantage, they also create complexity and a “data dilemma” problem.

Enterprises generate, ingest, and store ever-increasing amounts of data from users, devices, sensors, and systems daily. Estimates range from gigabytes to terabytes per day, and by 2025, this amount is expected to reach 180 zettabytes.¹ This is inclusive of data generated by cybersecurity tools and solutions. Beyond just sheer volume, business data — and metadata — gathered from multiple disparate sources, in different forms and formats, provide varying depths of information, yet the efforts

to turn raw data into actionable, connected business intelligence is like moving mountains.

But the rewards are self-evident. Superior data utilization goes well beyond capturing transactional data and scanning for unwanted activity. By unlocking the power of data, an enterprise can examine its past, react more effectively to the present, and chart a course for the future that’s more secure and better protected.

¹ “Big Data Statistics in 2023: How Much Data Is in the World?” First Site Guide, 2023

The foundation for advanced threat hunting

Data is the currency of the 21st century. And unlike gold, no one is short of data. For it to fuel AI/ML models used in predictive user analytics or advanced threat hunting, you've got to start at the beginning of the data upstream. When datasets, cleansed and normalized, are shared and used by multiple teams, security and data engineers can craft AI/ML models that learn from your data, adjusting baseline behaviors using your business context.

While generic out-of-the-box detections can help less-mature security organizations get up and running, nothing beats using data directly sourced from your organization to provide significant contextual awareness.

Once data is normalized, the real cybersecurity work can commence, but looking for threat signals in all that data is like looking for a needle in a haystack. How can enterprises sidestep these challenges while reducing risk and empowering highly effective threat hunting?

DataBee® from Comcast Technology Solutions enables data consumers and security teams to build improved and optimized data to aid threat hunting. This guide will explore the details of using an evidence-centric approach for improved threat hunting.

Building AI/ML models with optimized datasets to fight cybercrime is within reach, but challenges remain.



Buying security products to mature your security program can quickly become expensive to do, and to maintain.



Cybersecurity experts with the right know-how are in short supply.



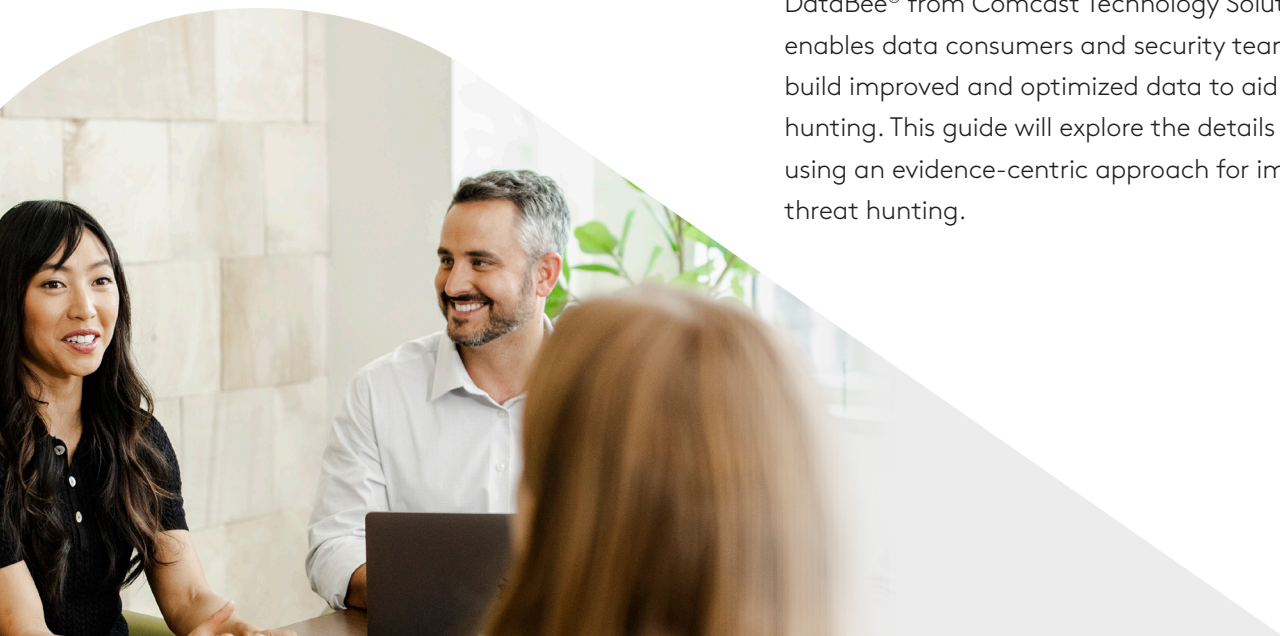
Data analysis is often complex, as each source may use unique syntax to represent and structure the data.



Data parsing is an intricate process, making any analysis or threat hunting challenging.



Normalizing the data is time-consuming but necessary to get its structure consistent for analysis.





Getting data right for AI/ML in threat hunting


Let's start with level-setting on some terms:



Artificial intelligence (AI) refers to the processes and algorithms that simulate human intelligence, including mimicking cognitive functions such as perception, learning and problem-solving, visual perception, speech recognition, and decision-making.



Machine learning (ML) is a subset of artificial intelligence and is a way for a machine to imitate intelligent human behavior without explicitly being programmed. Machine learning is one way to use AI.



Optimized data is data that has been parsed, flattened, and normalized. This “clean” data can be enriched, integrated, and shared with other applications, including AI and ML models.

Why is AI/ML being talked about in virtually every industry — especially in the cyber community? It's the ability to thoughtfully cull through terabytes of data in a fraction of the time it would take a human or team of humans, coupled with increasing levels of comprehension and context with continued use, to make for a superior way to make smarter decisions faster — including how to stay ahead of threat actors' innovations.

GETTING DATA RIGHT FOR AI/ML IN THREAT HUNTING

Threat hunting

Threat hunting proactively seeks to identify adversary activity that existing detections or incident response programs miss.

Large enterprises often have internal threat hunting teams, while smaller organizations may outsource the function to a third party due to cost.

Today, in most organizations, threat hunting is still a largely manual process and restricted by technology vendors' specific capabilities. Threat hunters typically begin their investigations by developing a hypothesis following an actual or suspected security event or breach. Hunters then need to find evidence or artifacts of that incident using data pulled from multiple sources. The aggregation and triaging of related alerts from the varied technology tools can result in investigations spanning days, weeks, or months. It's not uncommon to miss multiple low-level alerts.

A fresh approach to threat hunting

The more logs and relevant data a threat hunter can review and correlate, the more successful they will be. As data streams to its destination, DataBee detects threats and accelerates investigations by applying vendor-agnostic Sigma rules and automated detection chaining.

DataBee pulls Sigma rules from a community- or company-driven GitHub repository and natively translates the detections to apply them on data optimized to the Open Cybersecurity Schema Framework (OCSF). When alerts are triggered, the logs are sent downstream to SIEM/SOAR and other tools.



Speed and scale: Active detection streams from DataBee reduce the noise and junk data from your logs while retaining the logs that triggered a Sigma alert. In addition, raw and processed data can be stored in a data lake — an ideal place to test and run AI/ML models. Modern cloud-based data lakes running on elastic compute engines can provide powerful performance and enable collaboration.



Context: Security events without business context require arduous event triaging and prioritization. Sigma rules and AI/ML models for security detections facilitate the evaluation of multiple baseline scenarios with deeper context, leading to better decision-making and fewer false positives, empowering threat hunters to maintain greater focus on higher-priority threat signals.



Data optimization: Ensure the most effective and accurate AI/ML models by focusing optimization early in the data pipeline. Data consumers, like threat hunters, must trust its integrity when testing queries and conducting investigations. Data can be parsed and deduplicated at scale using AI/ML so it can be reviewed quickly to provide increased visibility and early threat detection.



Reduced cost: Reduce SIEM operating costs, IT support, and threat hunting by ingesting high-value logs and non-traditional logs. Using optimized data and AI/ML to integrate and create cleaner datasets helps reduce the computation strain on existing security analytics and provide faster outputs. Ultimately, this alleviates pressure on CISO and CIO budgets.

Workflow for AI/ML modeling

Threat hunting requires a collaborative, yet programmatic, workflow approach. Data must be specially prepared and manipulated by data scientists and analysts in order for threat hunters to receive optimal value. The workflow stages include steps to:

- **Develop use cases.**
- **Capture and store data.**
- **Cleanse data.**
- **Transform data: parse, normalize, and enrich.**
- **Create baselines.**
- **Perform threat hunting.**

Develop use cases

Threats and threat actors are strategic, and your threat hunting program should be as well. Your security team should run tabletop exercises to identify relevant use cases. Threat hunters use AI/ML outputs from many sources: operating system logs, application logs, endpoint security logs, telemetry data, alerts, cloud logs, cloud platform and accounts information, user access logs, etc. They can use this data to pinpoint unusual or suspicious behavior and hypothesize use cases. These use cases may include specific or suspected threats not previously detected by in-house incident response teams or off-the-shelf security tools.

Capture and store data

To create effective AI/ML models for investigations, you need data. Data science teams strive to capture as much data as possible and store them for as long as possible to find trends and create relevant intelligence about the business. Multiple data sources are usually involved, including the organization's pertinent data sources, cloud-based storage services like AWS S3 or Azure Blob Storage, and data lakes or warehouses.

Cleanse data

Once the data is captured, it needs to be cleaned. Data cleansing is the process of deleting or fixing improper, corrupt, or incorrectly formatted data. This is also the process whereby data is deduplicated and data fragments are addressed. For data teams, clean data is a top priority because the better the cleansing results are, the better and faster it is to extract insights from the data.

Transform data: Parse, normalize, and enrich

At this point, collected data from disparate sources is ready to be transformed into a standard format so it is more searchable and usable for data consumers. While there are many schemas to choose from, DataBee maps optimized data and security telemetry to OCSF, a collaborative, open-source project that aims to standardize security taxonomy and data types to improve consistency, analysis, and collaboration. The value of OCSF is that it can be used by multiple users interested in security analytics and for multiple personas. DataBee is an active member in the community and has built extensions so that the schema can work across security, IT, and non-traditional data sources, reducing time and effort when analyzing security alerts and investigating events.



Data ingestion:

The collection of raw data from multiple sources into storage mediums



Data parsing:

The process of identifying patterns and extracting data from large datasets into another format



Data flattening:

The process of converting data, whether structured like firewall logs or unstructured Word docs, into fewer, more manageable datasets



Data normalization:

The process of turning data into a standardized, consistent schema. Examples of popular schemas are Elastic Common Schema and OCSF.



Data enrichment:

The process of enhancing data with relevant information and attributes to improve the analytical and operational value of the data

Create baselines

Once the data is transformed, data scientists collaborate with threat hunters to develop and define baselines. Typically, this means describing the environment or situation “right now” and searching for deviations to indicate malicious activity. Creating proper baselines requires expertise to know what attributes and data points to use and how to use them. For example, a baseline may have many attributes, but not all attributes may benefit the use case — or one data point might overlap. Therefore, security experience and expertise are vital to asking the right questions, interpreting the results, and creating the final baselines.

Perform threat hunting

With the baseline defined, threat hunting starts in earnest. The hunters launch queries against the data, looking for anomalies against the baseline. These queries might return hundreds or thousands of results that must be further parsed and analyzed. The next step is to validate the efficacy of the data and hunt. This involves the threat hunters trying to mimic the threat actors' behavior. And often, the testing implies that the baseline needs to be adjusted, so the whole cycle repeats.



Why engaging with data early in the pipeline is important

Engaging early with data sets the tone for the entire AI/ML interaction, improving data quality and model scalability.

Specific workflow steps like data cleansing and enrichment are long-pole activities, so allowing extra time can only be beneficial — and even necessary. Early involvement permits more time for data extraction and modeling. Thus, early engagement ensures ample focus by data scientists to optimize the data and for the threat hunters to engage with it.

Early engagement can also help set the parameters for budget and cost allocations. The cost of storing and analyzing massive amounts of data can be prohibitive and can escalate along with increased volume. Establishing a budget provides the guardrails and guidance to teams that help to control costs, and to prioritize projects and use cases.



Threat hunting use case example:

Lateral movement

For threat hunters, lateral movement refers to attackers exploring a compromised network to find vulnerabilities or assets of value. It is called lateral movement because once attackers gain access, they explore and pivot between systems undetected to achieve their objective, which may be espionage or a more sinister goal like installing remote access tools. An attack might begin with malware placed on an employee's device from a successful phishing email.

Detecting lateral movement without the help of AI/ML has traditionally been a time-consuming, repetitive process. First, the hunters would take inventory of their organization's attack surface, giving special attention to how a threat actor might target points of vulnerability. Hunters would then consider existing techniques observed in frameworks like MITRE ATT&CK that would help them write, test, and run various anomaly detection techniques to detect lateral movement.

With AI/ML, threat hunters can quickly automate and accelerate hunting techniques for lateral movement, even in massive, hectic, noisy environments.

Access to an integrated data layer consisting of large connected datasets enables threat hunters to cover attack vectors previously unconsidered. Without access to that data, there are visibility gaps where an attacker could hide. And this might prevent a threat hunter from a successful hunt altogether.

While the benefits of using AI/ML for threat hunting far outweigh manual detection techniques, it is not a panacea. It can assist with hunting and speed up the process, but it is only as good as the threat hunters' experience, models, and baselines. Hence, having a seasoned threat hunting team coupled with experienced data science and analytical teams is critical to success.

For example, hunters might look for suspicious lateral movement using Microsoft's Windows Management Instrumentation (WMI). Hunters will first identify the processes on all hosts with a parent process consistent with lateral movement. From there, they will use grouping and stacking to identify rare command line strings within the environment for further review and deep manual analysis.



Data, threat hunting, and Comcast

Comcast is a Fortune 30 company that produces terabytes of data from diverse sources daily.

The Comcast Cybersecurity team actively threat hunts against an expansive and diverse data landscape by leveraging a security, risk, and compliance data fabric that merges disparate data sources and feeds with organizational insights and business policies. With the help of a modern data lake, the result is proactive threat detection with near-limitless hunts, thus reducing the company's overall risk profile.

For those looking for similar capabilities and outcomes, DataBee enables complex hunts to be conducted on large-scale, historic, and integrated datasets collected, normalized, and enriched into an authoritative source. DataBee integrates with your security data lake to remove compute constraints so you can schedule simultaneous and intensive queries whenever necessary.





How DataBee® can help organizations of any size with modern threat hunting

Up until now, AI/ML modeling for threat hunting has been cost-prohibitive. The initial cost and resource investment often exceeded the budget of many security organizations. On top of that, recruiting and training experienced data scientists and threat hunters adds another challenging complexity.

DataBee is designed to overcome many of these financial and logistical challenges. DataBee is a cloud-native security, risk, and compliance data fabric to accelerate AI/ML initiatives. DataBee modernizes security data management and provides integrated and enriched insights by connecting disparate data sources and feeds with organizational

data and business intelligence. With simplified and enhanced traceability for data ingestion, flexible enrichment options, and automated normalization, DataBee produces an enhanced dataset ready for high-performance analysis and reporting. Customers can stop choosing between capacity and security and ensure data quality by storing optimized data in a data lake.

DataBee benefits

-  Use data to drive cross-department collaboration.
-  Maximize data efficiency without sacrificing quality.
-  Gain back control and ownership of your data.
-  Stay ahead of threats and changing data privacy regulations.

DataBee was initially inspired by the internal Comcast security, compliance, and data team. Through the implementation of a security data fabric at Comcast that spans across 150,000+ employees and over 1M endpoints, Comcast was able to maintain >50TB a day of data throughput and retain 10PB of security data for over a year while reducing \$10M in annual operational costs. Eliminating point products and optimizing SIEM ingestion with cleaner data enabled 65% higher fidelity alerts and 3X faster threat detection.

DataBee helps you get more out of your data

Comcast Technology Solutions (CTS), a division of one of the world's leading media and technology companies, is proud to offer DataBee, our cloud-native security, risk, and compliance data fabric platform.

DataBee helps detect and stop threats with vendor-agnostic detections as the data streams toward its destination. The security, risk, and compliance data fabric platform weaves together security and non-traditional data and business context to help you solve the most complex challenges. Take down technical barriers for more immediate and connected insights so you can collaborate with data you can trust using business intelligence tools your people already know and love. Your data remains your data with better performance and low-cost data processing.

[MORE INFORMATION ON DATABEE →](#)



DataBee® is a registered trademark of Comcast Corporation.