ISSN: 3067-1140

AMERICAN ONLINE JOURNAL OF SCIENCE AND ENGINEERING

VOLUME: 02 ISSUE: 01

RECEIVED: JANUARY 08 REVISED: JANUARY 26 ACCEPTED: FEBRUARY 22 PUBLISHED: MARCH 06

Cloud-Native AI Frameworks for Proactive Cybersecurity Threat Detection

Independent Researcher
Dhanaraj Sathiri
dhanrajsathiri@gmail.com

Abstract

Overreliance on signature-based sensor data, automation, and information sharing often enables threat actors to perpetrate cyberattacks such as ransomware, cryptojacking, and other undetected activities for months. Even without prior knowledge of an attack, security teams have the opportunity to locate advanced persistent threats that have already evaded existing security controls. Organizations must thus embrace threat hunting—proactively seeking active adversaries in their environments using tailored detection capabilities. Information technology teams operating in public cloud environments can now empower hunting by observing attack paths, developing detection capabilities in a serverless, on-demand, and cost-effective manner, and storing massive amounts of data in a cloud-native architecture.

Threat hunting is defined as the hypothesis-driven exploration of an enterprise's environ- ment to detect and isolate undiscovered malicious activity. The process relies on advanced knowledge of attackers and their techniques to create detection-focused intelligence that is then fed back into observability solutions. Conducting efficacious threat hunting involves defining clear objectives; selecting appropriate data sources, collection pipelines, and storage solutions; and developing detection capabilities that specifically target areas of interest. Despite the widespread adoption of threat hunting in traditional enterprise networks, cloud operating models introduce significant new characteristics and considerations that warrant development of a corresponding enterprise cloud threat-hunting framework.

Keywords : Cybersecurity; Cloud Computing: Machine Learning; Threat Detection; Deep Learning; Serverless Computing; Knowledge Graphs; Federated Models; Explainable Models; Observability; Telemetry; DDoS Attack Analysis; ICS System Verifications; Digital-Machine Considerations; Model-Reaction Engines; Anomaly Detection; ChatGPT Models; Supervised Learning; Anomaly-Detection Segmentation; Computer Network Datasets; Ground-Truth Dataset; Cyber-Attack Data; Cybersecurity.

1. Introduction

Digital transformation processes tend to undermine the use of conventional detection mechanisms based on rule matching for blocking known signatures. Next-generation security solutions typically integrate multiple detection strategies, such as file package aggregation to capture compound and multi-stage attacks. Yet, these measures cannot completely eliminate blind spots, as sophisticated malware may not be detected in server-side execution environments. In addition, detection delays can open up windows for adversarial exploration, including evasive actions that go undetected for days or months after a breach. As a result, threat detection as a mission is evolving from preventing breaches to hunting for indicators of compromise (IoCs) and adversarial

presence. Threat hunting is an active process to seek out undetected threats using heuristics and knowledge of attack patterns. Cyber-security teams no longer rely solely on detection engines but assume an adversarial mindset to actively seek exploits lurking in their cloud stacks.

The fundamental workflow of a threat-hunting mission typically includes knowledge gathering, data collection, hypothesis formulation and validation, and reporting. Commercial threat intelligence services, such as CrowdStrike's Falcon X and Mandiant MAZE, can be leveraged to enrich hunting missions with up-to-date IoCs. Hypotheses are usually guided by threat intelligence and are designed to reveal specific exploits or behaviours that might otherwise escape detection. Public cloud service providers offer out-of-the-box observability and telemetry for their environments, making the data-gathering phase easier than in private or hybrid stacks. However, cloud operations often leverage auto-scaling groups and containers with ephemeral lifetimes, collecting architecture, data, and application layer IoCs becomes complex and hard to automate.

1.1. Research design

Cybersecurity incidents often arise from known vulnerabilities or newly discovered zero-day exploits. Security patches for commonly used software packages, web servers, email applications, and user platforms (such as browsers) are released either by their vendors or as part of security-relevant Linux distributions. Why, then, do we still read about data breaches in the news? Cyber intrusion detection and prevention systems (IDS/IPS) should shield us from known attacks by applying rules based on previous incidents, such as patterns of traffic associated with some malware. Yet data breaches continue to occur, often resulting in years of costly, drawn-out lawsuits, regulatory scrutiny, and loss of public trust. These observations suggest that classical security measures are no longer sufficient to prevent online data breaches or similar incidents. Many institutions now undertake threat detection and proactive preemptive threat hunting even in (or perhaps especially in) environments that implement these classical measures.

This form of intelligence-driven detection deviates from the conventional prevention-and-detection operating model, and can be thought of as the inverse mirrored model. Intelligence-driven detection seeks to unblock traffic and address the attack aggressor in the vulnerability space. It not only looks for known indicators of compromise (IoCs), but also for the absence of indicators of benign behavior. The need for such a shift is perhaps more critical in large and complex environments where there may be multiple agents with the ability to detect these events. Intelligence-driven detection relies on external threat intelligence sources, internal data-sourcing mechanisms, and talent skilled in exploiting both. It may use traditional machine-learning (ML) detection techniques, or a combination of supervised and unsupervised anomaly detection approaches. It may use directed-testing models that dynamically test the strength of detections linked to exploited-entry kits in multiple environments. But key in all techniques is an understanding of delay and exposure risk associated with expected change points in the environment. Its growing importance, meanwhile, adds complexity and operational overhead for a security operations center (SOC) team that is often already stretched thin.

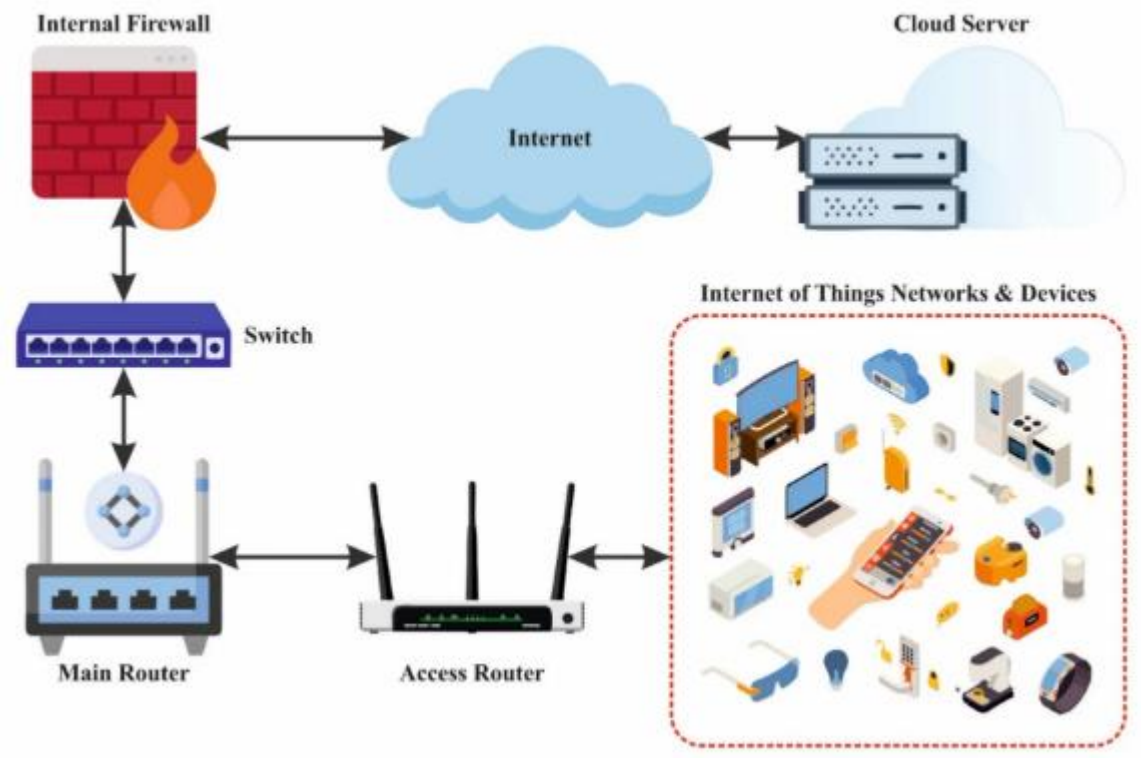

Fig 1: Intelligent Cybersecurity Threat Hunting

1.2. Background and Significance

Proactive threat hunting enables security analysts to discover potential cyberattack indicators before exploits occur, proactively investigating these hypotheses using internal and external data sources. Advanced AI techniques facilitate threat hunting tasks using diverse approaches, including deep learning, reinforcement learning, and knowledge representations, within public-cloud, private-cloud, or hybrid-cloud settings employing serverless architectures and microservices.

Malicious entities discover and pounce on system vulnerabilities within minutes or hours after exposure, yet many organizations have been tardy in discovering incidents and determining root causes. Deploying security solutions is ineffectual without reviewing the alerts they generate. Proactive threat hunting enables security analysts to discover potential cyberattack indicators—whether malware, ransom, or a denial-of-service—well before planned exploits occur, thereby improving incident-detection capabilities. With a hypothesis-based approach, security personnel take a more proactive stance, using intelligence to identify and catalogue warning signs of threats. By integrating and investigating several possible indicators using internal and external data sources, security staff can determine whether their hypotheses are accurate.

Advancements in artificial intelligence have rendered it possible to apply advanced techniques in algorithmic search and AI techniques from classical-AI and machine-learning domains to alleviate some of the tedious, time-consuming, and embarrassing aspects of threat hunting. Facilities exist for applying deep learning in tasks such as data-centric cyber-situational-awareness tasks. Yet there remains a need to deploy advanced threat-hunting techniques based on other AI paradigms such as reinforcement learning, knowledge representation, knowledge graphs, and deep reinforcement learning within a controlled environment that matches the active-hunting capabilities of the best human analysts in a specific area of interest.

Equation 1: Threat score from weighted telemetry features

Let the telemetry feature vector for an entity be

$$\mathbf{x} = [x_1, x_2, \ldots, x_n]^T$$

where each $x_i$is a feature such as:

- unusual IAM calls,
- abnormal API rate,
- privilege escalation signals,
- suspicious login location,
- process anomalies.

Define a weight vector

$$\mathbf{w} = [w_1, w_2, \dots, w_n]^T$$

Then the raw threat score is the linear combination

$$S = \sum_{i=1}^{n} w_i \, x_i$$

Step-by-step derivation

Start with the assumption that total suspicion is the sum of feature contributions:

$$S = w_1 x_1 + w_2 x_2 + \cdots + w_n x_n$$

Using summation notation:

$$S = \sum_{i=1}^{n} w_i \, x_i$$

Using vector notation:

$$S = \mathbf{w}^T \mathbf{x}$$

2. Foundations of AI-Driven Threat Hunting

Threat hunting is a term that encompasses a broad range of different human-centered activities and methodologies in cybersecurity. It refers to the proactive searching for threats that are present within or have made it past an organization's standard security defenses. While most security operations centers (SOCs) focus on detection and response to threats, threat hunting seeks to identify advanced threats that may be eluding detection. These threats use stealthy and cunning means of attack designed to evade common detection measures, and seek a foothold from which to launch more devastating attacks against the enterprise. Strategic threat hunting increases the probability of early discovery of these hidden threats before damage to the enterprise occurs. Factors such as timing, the general state of the enterprise, and large-scale background noise from insecure network configurations can make hunting for such threats complex—or even futile—at times. Despite their appearances of being wholly on-the-offensive, cybercriminals still try to adopt some level of covert defensive behavior so as to avoid detection.

Threat detection and prevention are constantly evolving in an ever-fighting cat-and-mouse game. To maintain a competitive edge, skilled and knowledgeable cybersecurity personnel are a key element of any organization's defenses. There are tools available to assist in this, yet false positives hinder their effectiveness and solutions designed to prevent detection can remain invisible to these tools. As malware continues to adapt, so must active defenses. Threat hunting is an active exploratory process in which a hunter systematically scours the network for undetected and confirmed threats. Anomaly detection is a normal part of any

intelligence-gathering procedure; if true anomalous behavior in an enterprise—malicious or otherwise—is discovered, steps can be taken to further analyze and trace the root of the event. Threat hunting is more than simply finding a host of compromise or whether a command and control (C2) connection is visible in traffic; it is taking the time to look for the processes and behaviors, localized or system-wide, that correspond with those found in successful breaches elsewhere.

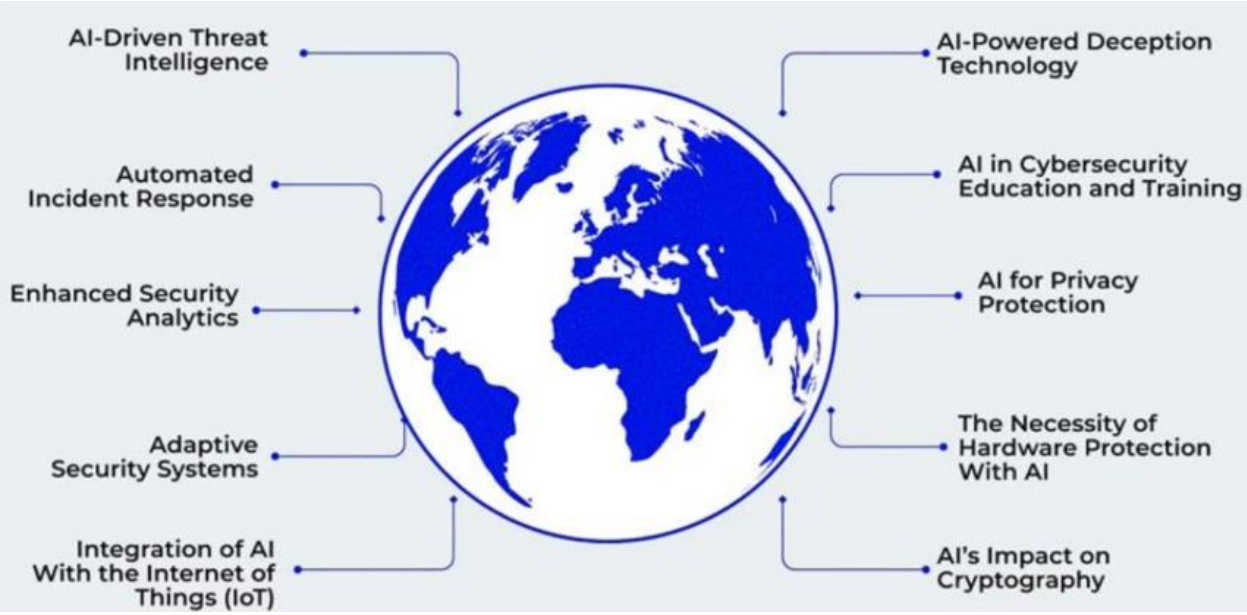

Fig 2: AI-Driven Threat Hunting

2.1. Definitions and scope

Intelligent threat hunting can be defined as the proactive exploration of a cloud environment's data with the goal of detecting advanced threats that evade existing security solutions. The cloud environment can have threat hunting capabilities enabled by cloud security providers that maintain multi-tenancy solutions for security tools in their data. Threat hunters can then hunt for threats in their cloud environments using these tools but may not have extensive security knowledge nor the easily interpretable data availabilities. Model security can be ensured by using differential privacy when training machine learning models. Data from threat hunting can be shared with other enterprises to help future threat detection and mitigation, but privacy of sensitive data should be ensured. Evaluation and effectiveness metrics should be established to validate the threat hunting process. Cloud environments that deploy adopted threat hunting data, pipelines, and AWS and Azure environment reconstructions.

Threat hunting in any environment can be an expensive and time-consuming process due to the wealth of data that needs to be investigated to find malicious activity. Specialized cloud-native observability and telemetry services help provide a simple interface to show attacks as they happen, yet detection often comes too late. Signature-based tools do not actively hunt for threat detection and can only classify known attacks. Complex attack patterns and newly discovered malware not detected by those tools leave the environment vulnerable. To proactively explore data and find these attacks before they happen, threat hunting needs to be implemented. Public and private clouds are capable of deploying serverless services with deployed models integrated into the core serverless architecture and function APIs.

Equation 2: Probability of malicious activity using logistic transformation

$$P(y = 1 \mid \mathbf{x}) = \frac{1}{1 + e^{-S}}$$

Since $S = \mathbf{w}^T\mathbf{x} + b$, substitute:

$$P(y = 1 \mid \mathbf{x}) = \frac{1}{1 + e^{-(\mathbf{w}^T\mathbf{x} + b)}}$$

Step-by-step derivation

Let the log-odds of attack be linear in the features:

$$\log \frac{p}{1-p} = \mathbf{w}^T\mathbf{x} + b$$

Exponentiate both sides:

$$\frac{p}{1-p} = e^{\mathbf{w}^T\mathbf{x}+b}$$

Multiply both sides by $(1-p)$:

$$p = (1-p)e^{\mathbf{w}^T\mathbf{x}+b}$$

Expand:

$$p = e^{\mathbf{w}^T\mathbf{x}+b} - pe^{\mathbf{w}^T\mathbf{x}+b}$$

Move the p-terms together:

$$p + pe^{\mathbf{w}^T\mathbf{x}+b} = e^{\mathbf{w}^T\mathbf{x}+b}$$

Factor out p:

$$p\left(1+e^{\mathbf{w}^T\mathbf{x}+b}\right) = e^{\mathbf{w}^T\mathbf{x}+b}$$

Divide both sides:

$$p = \frac{e^{\mathbf{w}^T\mathbf{x}+b}}{1+e^{\mathbf{w}^T\mathbf{x}+b}}$$

Equivalent form:

$$p = \frac{1}{1+e^{-(\mathbf{w}^T\mathbf{x}+b)}}$$

2.2. Core components of threat hunting workflows

An intelligent, AI-enabled threat hunting framework incorporates observations and anomalies across different types of cloud settings. Such frameworks and the developed orchestrators provide significant automation to the understanding of large volumes of data and events and, subsequently, the creation of new detection and prevention measures. Cloud-native architectures provide a series of capabilities that support a scalable and energy-efficient deployment within a public cloud—telemetry and observability logging for security and system resource consumption. Through a set of managed services, robust and scalable ingestion pipelines, and security detection mechanisms in a hybrid cloud environment, the framework allows a large volume of detection measures to be implemented, aimed not only at alerting the defenders but at providing recommendations for remediation without the need for human interpretation.

However, effective threat hunting necessitates a deeper understanding of the observed behaviours. For this reason, a second main scenario targets the implementation of a private-hybrid cloud environment focused on the analysis and detection of threats through anomaly-based models. Detection measures are given in the context of a knowledge graph, which provides not least an

enhanced view for specialists. Both aspects benefit from a recent research direction in the evaluation of the built measures, going beyond simple performance measures in alert generation and focusing on the evaluation of their semantic contribution to the threat landscape of the studied cloud setting.

3. Methodology

Security and Risk Considerations in AI-Enabled Cloud Threat Hunting

AI-enabled cybersecurity threat hunting frameworks are typically applied in public cloud environments. Their design and implementation have often been inspired by threat-hunting requirements in enterprise or hybrid settings, where the cloud serves as either a true extension of the enterprise or as the target of complex attack campaigns. Typically, the architecture pattern relies on a dedicated observability-and-telemetry microservices stack that ingests data from cloud-native platforms, such as AWS, Azure, and Google Cloud Platform, generates an external threat-hunting data lake, and orchestrates processing jobs that create the various data products used to support the ML models on a schedule driven by SLAs associated with specific threat-hunting units.

Despite the advantages gained with such a distributed architecture, security and risk aspects have received limited attention. Threat-hunting models are exposed to a wide range of risks, such as data integrity concerns, potential leakage of sensitive training data, and vulnerability to adversarial examples. Supply-chain and model-assurance considerations are often neglected. Moreover, despite data privacy being a legitimate concern in enterprise scenarios, the employed methods might not be privacy-preserving. A comprehensive examination and taxonomy of security and risk aspects related to AI-enabled cloud-threat-hunting platforms, with a specific focus on enterprise and hybrid environments, would contribute to closing the gap.

3.1. Public cloud, private cloud, and hybrid environments

Public cloud services enhanced by machine learning techniques are now widely available and often serve as the first area of deployment for AI-based cyber-defense solutions. However, the underlying infrastructures are becoming more dynamic and complex, sometimes limiting the effectiveness of hosted threat-hunting engines. Private clouds that provide both infrastructure and platform-as-a-service capabilities can run the same threat-hunting models without the limitations of public-cloud accounts. Organizations with a hybrid cloud model can also deploy threat-hunting platforms in dedicated environments, but often use a mixture of economic and regulatory motives to adopt this model. A hybrid AI-enabled threat-hunting platform can use cloud servers to process telemetry data by small environments in the same company, thus resulting in lower Total Cost of Ownership (TCO) per company, federated analytic techniques, or specific scout models that combine different organization models.

Security-as-a-Service (SECaaS) delivery is maturing, and part of the Network as a Service (NaaS)-like vision of buying communications, hosting, and security from a limited number of suppliers is being realized. Organizations with limited budgets are now relying on specialized companies to provide the appropriate tooling and expertise to defend against threats. AI-enabled cloud designs must consider a deploy-once-solve-many vision that converges services for many customers. The introduction of a more accurate Automatic Speech Recognition (ASR) process that converges training data from different customers reduces TCO and increases accuracy. Detection models for newly released malware and ad campaigns using generative models can also increase detection capabilities for more demanding organizations even in SECaaS configurations.

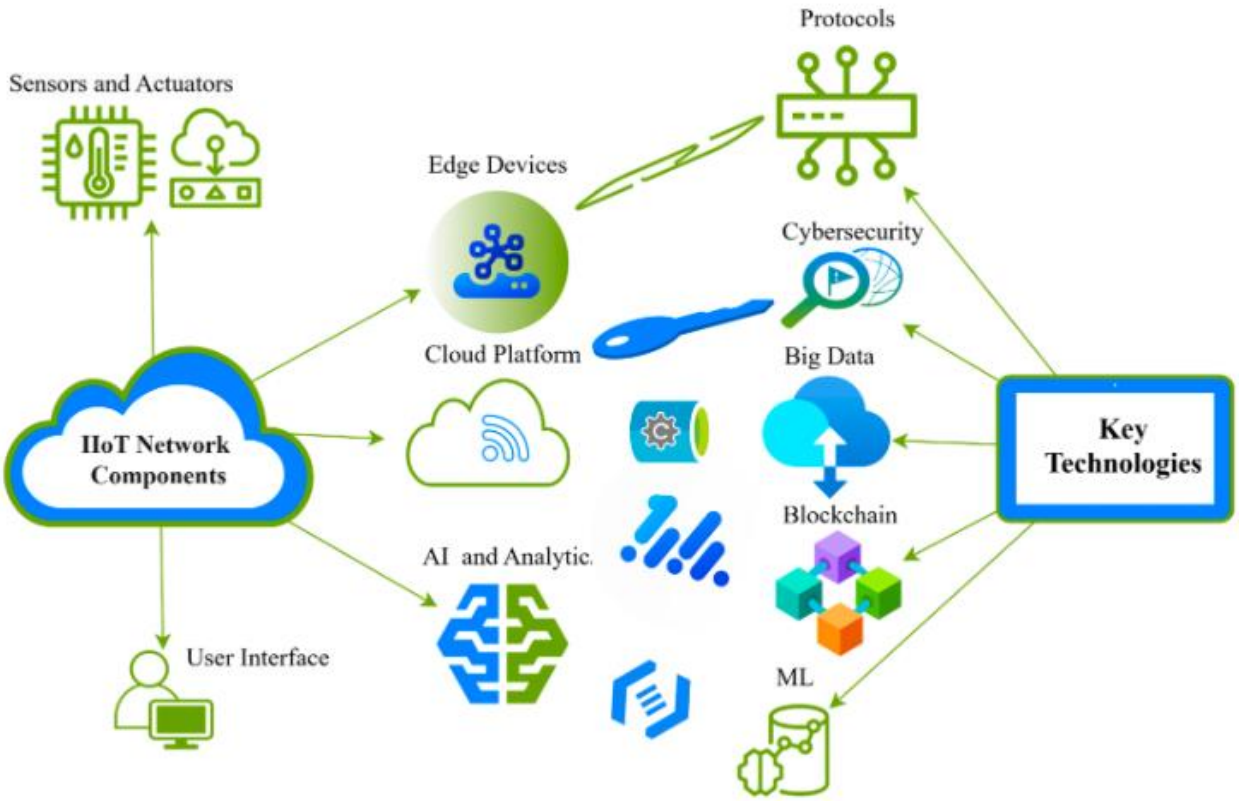

Fig 3: Public cloud, private cloud, and hybrid environments

3.2. Serverless architectures and microservices

Serverless architectures and microservices promise improved scalability, resilience, and decreased operational overhead while complicating observability and security analysis. Threat-hunting pipelines in serverless environments increasingly depend upon event-driven models in which security applications are triggered by events generated by other components. Such an operational setup is best viewed as an ensemble of detector/filter systems, each responsible for analyzing an event and generating or filtering an alert based on the outcome of that analysis.

The serverless architecture assumes that the underlying infrastructure is hosted, managed, and maintained by the service provider, and that all the resources are consumed on-demand, allowing the consumer to pay for only the used resources. Functions-as-a-Service (FaaS) allows the execution of user-defined functions without the need for deployment or provisioning of a physical or virtual machine. Service and workflow orchestration mechanisms can guarantee the correct and on-time execution of the functions.

Equation 3: Anomaly score using z-score normalization

For a feature x, with mean $\mu$and standard deviation σ:

$$z = \frac{x - \mu}{\sigma}$$

Step-by-step derivation

Suppose we want a normalized distance from normal behavior:

- observed value: x

- normal average: μ
- scale of normal variation: σ

First compute deviation from normal:

$$x - \mu$$

This is not scale-independent, so divide by standard deviation:

$$z = \frac{x - \mu}{\sigma}$$

Interpretation:

- $z \approx 0$: normal
- large positive or negative $|z|$: unusual
- alert if $|z| > \tau$

So a simple anomaly rule is

$$\text{Alert if } |z| > \tau$$

4. Objective of the Study

AI-enabled cloud computing frameworks for intelligent cybersecurity threat hunting.

Advanced AI-driven and ML-based frameworks enable threat huntsheet cloud designs, enabling defenders to detect previously unseen tactics, techniques, and procedures (TTPs), or TTPs that differ from the detected patterns of known attacker groups.

As multiple buildings and sub-building domains are deployed under their own management, as the number of allocated detectors and scenario detectors is increasing day by day, as the patterns appear on the network tend to be multi-layered and multi-faceted, is becoming very important for analysts to create new detection pattern teaching samples using the latest patterns that help corporate behavior-tied colleges. Through large amounts of data learning, the detection effect of pattern detection is more accurate. In the hyperscale deployment level of threat detection, virtual sub-machine for machine learning-based pattern detection effect will become first choice. When the underlying cloud computing resources are complete layouts of global environmental resources, by holding cloud hold, broadcasting cloud holding, distributing cloud holding parallel cloud holding intelligent wardens help the machine clear control.

By combining cloud deployment architecture, aiming at the characteristics of large-scale detection development of machine pattern detection, the cloud allocation spatial dimension are gradually broadened. By introducing cloud detection, detection layer can display the complete framework of actually deployed detection virtual machine distribution, has intelligent recognizable intelligent layer to manifest to cloud detection through empirical learning. The final realization enterprise full-cloud service parallel data analysis capability, in addition to smart camera contains the data review by self-hold, the in or by enterprise joint launch detections data are automatically transfer support back up center by.

4.1. Machine learning and anomaly detection

Machine learning and anomaly detection approaches play a pivotal role across many aspects of cyber threat detection. For example, they can be utilized for deciding if the assets in an environment can be classified as benign or as containing атака whether to raise a warning during the data acquisition process or whether to determine the state of a machine – instead of collecting constantly a huge volume of logs and monitoring all the acquired data. Enhancing the capability for decision making can potentially lead to the reduction of the false alarm rate. In addition, machine learning methods for cyber threats employ a

variety of labelled datasets with automated labelling procedures for special viewpoints, like malware detection or phishing. Moreover, it can provide a simulated dataset for testing detection algorithms of specific attacks, such as DoS or Port-scan attempts, in public clouds.

Anomaly detection methods aim to identify unusual patterns in the behavior of system assets and user activities. These methods are considered vital for detecting unknown attacks. They rely on the assumption that normal activities can be adequately described by a large volume of training data or dataset, while the anomalous behavior is simply an outlier in the representation of the normal activities. Both supervised and unsupervised anomaly detection algorithms have been applied in the threat-hunting domain decision, with unsupervised methods preferred due to their capability of identifying zero-day attacks. These methods do not depend on labelled data, which are difficult to obtain.

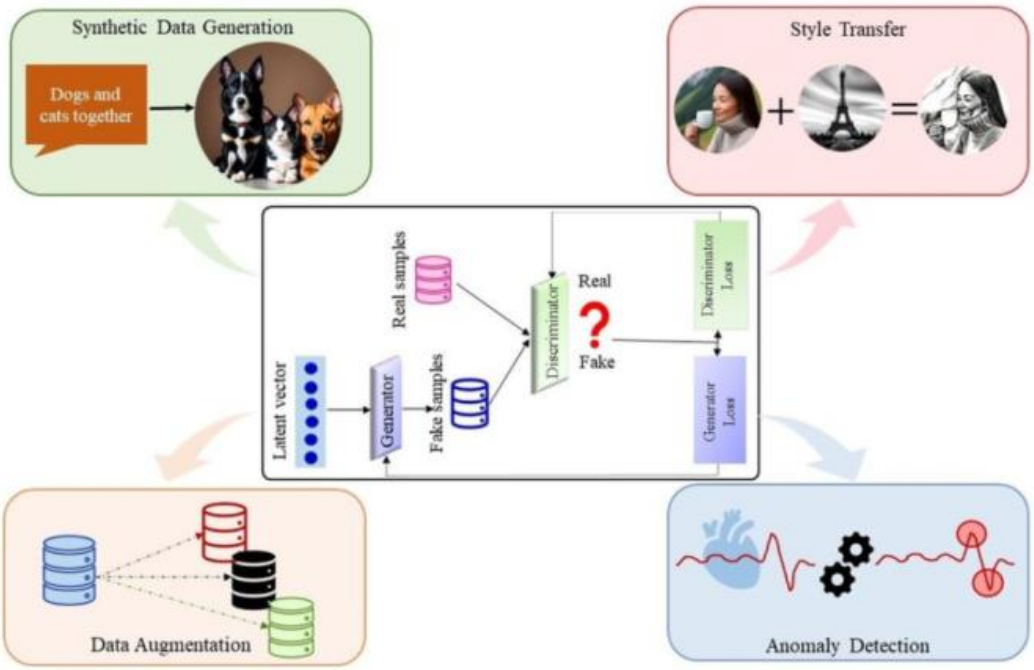

Fig 4: AI Anomaly Detection

4.2. Knowledge graphs and reasoning

Several approaches leverage machine learning methods for security monitoring tasks, which can be seen as a subset of observability. These techniques can detect zero-day problems or bridge the gap to create cyber-threat heuristics with little to no user-generated signature definition. In a cloud environment, key parts of the serverless architecture actually serve as its observability mechanism, where any abnormal behavior from an application can have an impact on the service shape and usage. Detection is a matter of labeling the abnormal behavior and then developing an ML model to detect it again, refining it on a regular or continuous manner.

However, extensive dataset labeling and continuous monitoring are required for detection by anomaly detection ML models or security-focused domain-knowledge ML models. A possible solution lies in the transition toward model-driven monitoring, which enables the monitoring application to be driven by model definitions normally created visually in model-driven development trends. The definition of models in existing declarative model- or event-definition-driven engines enables process, business rules, and event-based applications monitoring, but there are no models to drive observability-forming application development. AI-driven observable telemetry shaping the observability of data flow and state of an application configured by business-oriented schemas can be a breakthrough in observability, allowing a simplified and business-language-oriented ML model development to cover a wider range of application problems.

Equation 4: Multifeatured anomaly score using Mahala Nobis distance

For feature vector $\mathbf{x}$, mean vector $\boldsymbol{\mu}$, covariance matrix Σ:

$$D_M^2 = (\mathbf{x} - \boldsymbol{\mu})^T \Sigma^{-1} (\mathbf{x} - \boldsymbol{\mu})$$

Step-by-step derivation

For one feature, anomaly distance is

$$\left(\frac{x - \mu}{\sigma}\right)^2$$

For multiple independent features, this generalizes to

$$\sum_{i=1}^{n} \frac{(x_i - \mu_i)^2}{\sigma_i^2}$$

But cloud features are often correlated, so variance alone is insufficient. Replace per-feature scaling with full covariance structure:

$$\Sigma = \begin{bmatrix} \mathrm{Var}(x_1) & \mathrm{Cov}(x_1, x_2) & \cdots \\ \mathrm{Cov}(x_2, x_1) & \mathrm{Var}(x_2) & \cdots \\ \vdots & \vdots & \ddots \end{bmatrix}$$

Let

$$\mathbf{d} = \mathbf{x} - \boldsymbol{\mu}$$

Then the covariance-adjusted squared distance is

$$D_M^2 = \mathbf{d}^T \Sigma^{-1} \mathbf{d}$$

Substitute back:

$$D_M^2 = (\mathbf{x} - \boldsymbol{\mu})^T \Sigma^{-1} (\mathbf{x} - \boldsymbol{\mu})$$

5. Research Summary

Many AI-enabled cybersecurity threat-hunting frameworks operate in private cloud environments. However, enterprise platforms built with public cloud technologies offer modern scalability, performance, cost-efficiency, and usability benefits. Furthermore, hybrid cloud deployments are becoming commonplace as organisations move toward more sophisticated IT and cybersecurity architectures. Improving existing AI-driven threat-hunting frameworks and exploring new capabilities for serverless or hybrid-cloud environments will accelerate research in these areas. Analysis pipelines based on cost-efficient data preparation and GPU-accelerated model training support machine learning and anomaly-detection threat-hunting approaches. However, growing interest in exploring the use of knowledge graphs for threat hunting requires high-quality labelled data with a defined ground-truth set.

Hybrid and multiple-cloud architectures bring increasingly serious threat challenges. Enterprise threat-hunting frameworks currently built for private cloud infrastructures will enable hybrid-cloud deployment. The public-cloud threat-hunting pipeline is designed and explored using a simple crowdsourced dataset. Telemetry, observability, and data quality considerations define the data ingestion pipeline and storage solution. Anomaly-detection approaches rely on data-processing cost considerations during both training and inference phases. In addition to detection model security, model training involves dangerous weaknesses in

labeled-data collection, model training, and threat-hunting controls for serverless and other hybrid-cloud environments. Consideration of supply-chain, privacy, and federated-analytics aspects enriches enterprise Cloud detection scenarios.

Equation 5: Federated learning global update

Suppose there are Kclients. Client $k$has $n_k$local samples and local model parameters $\mathbf{w}_k$. Total samples:

$$N = \sum_{k=1}^{K} n_k$$

The global federated model is the weighted average:

$$\mathbf{w}^{(t+1)} = \sum_{k=1}^{K} \frac{n_k}{N} \mathbf{w}_k^{(t+1)}$$

Step-by-step derivation

Each client trains locally and returns updated parameters $\mathbf{w}_k^{(t+1)}$.

A simple average would be

$$\mathbf{w}^{(t+1)} = \frac{1}{K} \sum_{k=1}^{K} \mathbf{w}_k^{(t+1)}$$

But clients may have different dataset sizes. So weight each by its data share:

$$\text{weight of client } k = \frac{n_k}{N}$$

where

$$N = \sum_{k=1}^{K} n_k$$

Then replace equal weights with data-proportional weights:

$$\mathbf{w}^{(t+1)} = \sum_{k=1}^{K} \frac{n_k}{N} \mathbf{w}_k^{(t+1)}$$

5.1. Data sources and ingestion pipelines

IT consultants provide several examples of information sources for malicious activity in public cloud environments. The AWS CloudTrail service logs record all the management API calls made on AWS resources, such as identity and access management (IAM) actions, VPC changes, security group updates, EBS snapshots and AMI requests, API calls to AWS organizations, and operations in AWS account-management services, such as IAM, CloudFront, and Route 53. These services can provide ground truth for detecting account-takeover scenarios, as well as anomalous IAM API activities—such as Root account usage, IAM policy changes, new IAM users or roles being created for others, and unusual IAM management requests, for example, made from outside the organization.

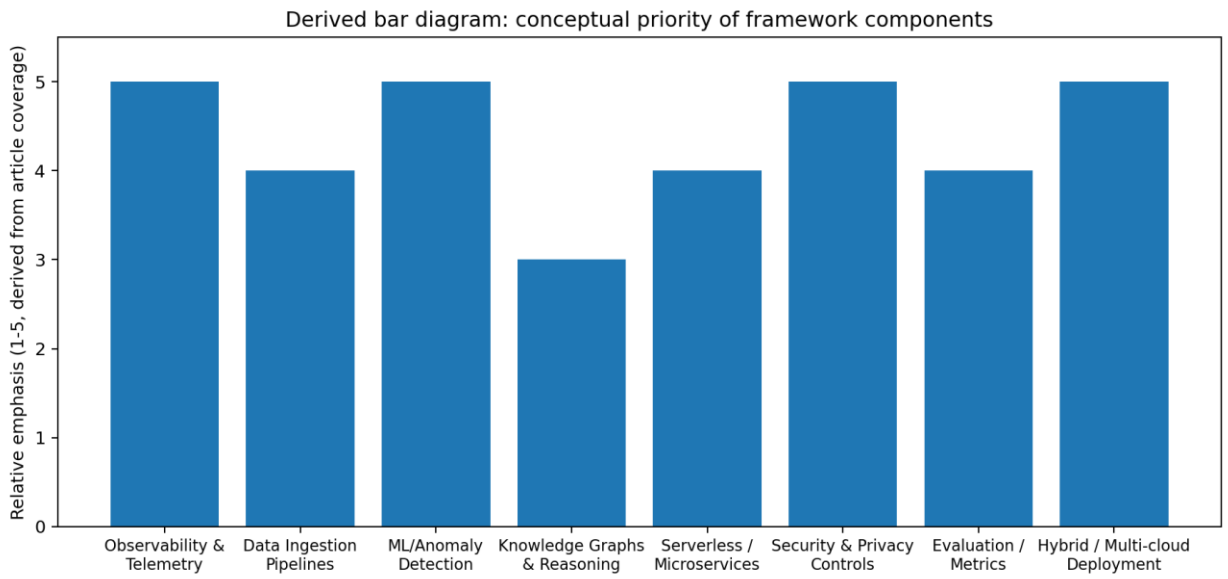

5.2. Data labeling, ground truth, and reproducibility

With the increasing application of machine learning (ML) for a wide spectrum of tasks, the design and generation of well-labeled datasets that accounts for ground truth to reach reproducibility gained special focus. In both supervised and unsupervised ML methodologies, the absence of proper training datasets can lead to poor prediction performance, mostly when a system is exposed to new and never-seen data. Open-source natural language processing models are now easily accessible, supporting several downstream tasks with incredible results. Nevertheless, high-quality labeled data representation for most of the popular tasks in the natural language understanding and computer vision domains are still pending. The situation is similar in cybersecurity where syntactically correct and semantically labelled data for supervised learning are costly to create and use pre-trained models.

An increasing number of different data sources used in an enterprise setting simultaneously require data-integration, processing, and storage pipelines that are scalable. In this use case, primarily for elevated privileges accounts of entities, it is essential to label the events where compromises entered. Historical product logs of a company were labeled according to the detection function and a supervised algorithm detected the signals to further propose the labeling of the remaining sources to allow threat hunting machine learning systems to learn. A Gen-AIGC data generation engine responsible for generating clean and non-duplicate events is applied as an analog for the baseline training labels, allowing a search and propose of variances with a pre-existing customer group.

6. Result

Observability and telemetry in cloud environments support threat detection capabilities for several stages of the attack lifecycle, enabling related alerts and notifications. These serve as inputs to intelligent threat detection, hunting, and investigation processes. The nature of cloud environments, as well as the extensive range of data sources available, support a wide variety of cloud threat detection and hunting scenarios. To that end, cloud environments are also equipped with dedicated data processing, transformation, and persistence strata capable of handling the high-velocity nature of cloud telemetry data.

These include streaming data processing and analytics engines with real-time short-term storage capacity for observability data, dedicated transformation strata for threat detection-feeding data sources, and low-latency cost-optimized long-term storage for future data mining endeavors. Given the optimization principles guiding the design specifications of many cloud data centers, an elastic cloud data processing and storage architecture is defined that allows for scaling resources up or down based on current data processing and storage needs and associated costs. The architecture encompasses the complete ingestion, processing, storage, and analysis components required for the above-mentioned cloud threat detection and hunting scenarios.

6.1. Observability and telemetry in cloud environments

The inherent opacity of cloud service infrastructures, particularly in public clouds, presents a major impediment to effective threat detection, thus warranting robust observability and monitoring mechanisms. Telemetry encompasses sets of metrics, logs, traces, and other signals produced either by the services supporting the cloud infrastructure or by the application services deployed on top of it. Such telemetry is essential for expose anomalies or behavioral changes in the cloud service operations that could indicate compromise(s) or malicious activities. AI-enabled observatory frameworks within the cloud infrastructure, therefore, must acquire and manage these telemetry sources. The acquisition of cloud locations of interest is highly dynamic, because cloud operators provision and deprovision resources continuously and often in a location-agnostic manner.

Effectually deploying a set of cloud services potentially scattered across the cloud region creates the challenge of efficiently designing the telemetry provisioning content — the observatory services need not monitor the activity of cloud services deployed in different cloud locations, and guaranteeing continuous coverage entails excess overhead for the collecting operations. Private clouds enable a level of observability closer to traditional on-premises infrastructures. In such a case, all telemetry sources should be covered and can be collected and processed by central units. A wider approach is also possible in hybrid scenarios, where private and public cloud resources are monitored together. Often the majority of the cloud operation life-cycle can be effectively monitored within the private cloud, while coverage of specific phases of public cloud operations is supported by the monitoring capabilities arranged in those locations.

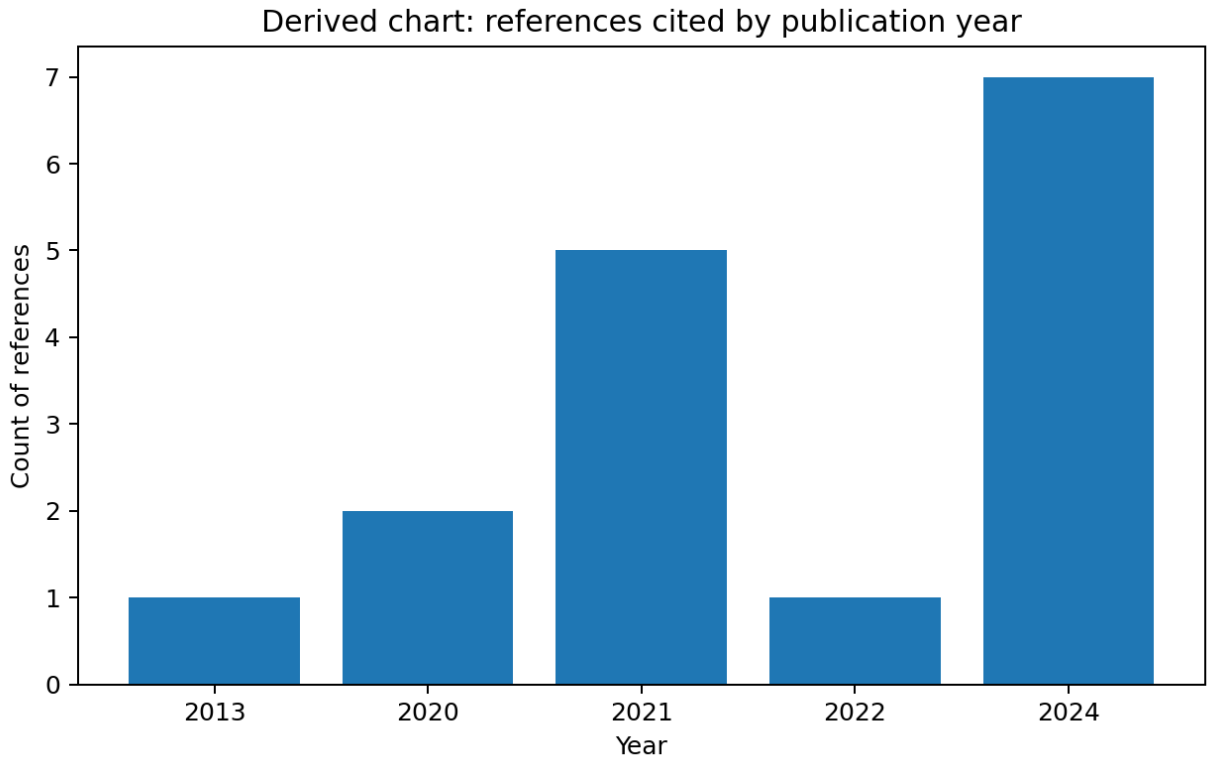

6.2. Scalable data processing and storage architectures

In cloud-native systems, the scale of data generated for observability requires scalable data processing engines and storage solutions that allow concentrating data in a few locations. Such a storage architecture is key to effectively deploy machine learning or anomaly detection solutions while ensuring sufficient ground-truth coverage for training and evaluation. A scalable cloud-native observatory architecture and an example cloud-native storage service for threat hunting analyses in hybrid or multi-cloud environments are presented. Both implementations are compatible with serverless architectures and can be deployed directly in public-cloud environments. The combination of these two elements allows building a hybrid cloud solution with a fully-managed security observability stack.

The detection of threats, their origin, and implication in Security Orchestration Automation and Response (SOAR) platforms has become a commodity in public-cloud environments. Serverless architectures require specific observability measures, and their γ, δ, and R logs must be stored for future analyses. These elements can be readily ingested and stored in the public-cloud-native threat-hunting data pipeline. The visibility gap of on-premises, hybrid-cloud, and multi-cloud environments that consist of a set of interlinked serverless components brings the need for threat hunting and an observatory that combines the available detection capabilities: cloud provider, open-source, and private solutions. The construction of a custom detection model that combines the available premises can add significant value and precision. Therefore, scalability and cloud providers' Pay-as-you-use models can be leveraged to add a hybrid-cloud experiment that fulfils the Security Data and Event Management (SIEM) visibility gap of the serverless functions in production and that probes in detail the log normality.

7. Security and Risk Considerations in AI-Enabled Cloud Threat Hunting

Multiple risk and security considerations are inherent within the application of artificial intelligence and machine learning in a threat-hunting context. During operations, malicious adversaries may seek to corrupt the model employed, whether through a direct poisoning of the training process or an indirect influence via the production input data. Data-integrity violations exist throughout the AI life cycle, from development through to deployment, especially with community-supported and open-source project repositories. These supply chain-related security concerns are further compounded within multi-cloud environments relying on resource sharing among multiple parties.

Beyond model integrity and data correctness, maintaining user privacy both during – and also postdeployment – is crucial. In some applications, precisely predicting or reconstructing the input from the classifier output can lead to the leakage of sensitive information. Adopting a federated approach for distributed learning can help to alleviate this concern and also reduce the overhead of data movement. Various techniques exist for privacy-preserving analytics, but their incorporation into threat-hunting systems remains an open area of research.

7.1. Model security, data integrity, and supply chain concerns

The success of any AI-enabled cloud threat-hunting framework depends on the integrity and correctness of the models, both classical and AI-based. Model poisoning attacks aim to compromise model training by manipulating the training datasets. Such attacks can inject errors or backdoors into models, allowing adversaries to evade detection. Data integrity attacks target the integrity of the data used by the framework. For instance, an adversary can take control of an observability agent and provide false information about the state of the system, whether to obscure a malicious action or trigger an escalation in privileges. Supply chain attacks seek to capture the stakeholders involved in the data acquisition pipeline—typically system tools produced by trusted sources—rather than attacking the threat-hunting framework directly.

Privacy is a crucial concern when malicious models can reveal information about their training datasets or expose sensitive attributes about individuals whose information was present during training. Data associated with a model can be leaked through multiple attacks. In addition to membership inference attacks that identify if an individual Ponzi model could be backdoored via poisoning attacks, AI models can also leak sensitive training data attributes that the classifier has used to make predictions. Similar information can also be given away by other classification models such as CNNs, GANs, RNNs, and SAE. These undesirable outcomes have highlighted the necessity of developing models for remotely storing and learning from data to protect privacy. Federated machine learning enables this by allowing customers to collaboratively train a model without revealing any private data, thus preventing these types of privacy breaches.

7.2. Privacy-preserving analytics and federated approaches

Distributed approaches like federated learning (FL) offer attractive privacy guarantees as the training does not require users’ data to leave their premises. For federated learning systems, both strong theoretical foundations and real-world implementations have been proposed, explicitly considering aspects such as security, efficiency, and system dynamics. Differentiated privacy provides strong guarantees for FL aggregation procedures enabling cross-silo FL with trusted worker nodes. However, the overall design of the client-side machine learning or deep neural network tasks has been proved to be susceptible against various threats in privacy and security. Analytical studies and adversarial training mechanisms have been proposed to strengthen clients against leakage of sensitive training samples through parameter perturbations or distillation processes.

A taxonomy of potential security and privacy vulnerabilities in FL is proposed, as well as corresponding solutions. External and internal attackers targeting user data and model parameter integrity and security, backdoor attacks on clients and servers, undesirable model behavior from corrupted clients, honest-but-curious FL service providers, and communication partners are considered and discussed throughout the taxonomy. Additional statistical attacks on user data are highlighted, contributing to a comprehensive picture of FL vulnerability analysis. Advanced defense schemes to reinforce FL systems against these identified threats are outlined, with further discrimination and discussion on solution level: theoretical mechanisms, architecture designs, and empirical solutions.

Article area	Derived visual	Purpose
Introduction / Background	Fig. 1 workflow + structural bar chart	Shows the end-to-end hunting logic and where emphasis appears in the paper.
Foundations of AI-driven hunting	Fig. 2 feedback cycle + framework priority bar diagram	Represents the AI/analyst feedback loop and major framework components.
Deployment methodology	Fig. 3 deployment architecture	Summarizes public, private, and hybrid cloud positioning.
ML and anomaly detection	Fig. 4 anomaly pipeline	Illustrates detection flow from telemetry to investigation.
References / literature base	References-by-year chart	Shows recency distribution of cited literature.

Table : Article-to-visual mapping

8. Evaluation, Metrics, and Benchmarking

Human-intensive processes such as incident response, penetration testing, security audits, and red teaming are difficult to quantify ultimately. Evaluation and benchmarking of these activities rely heavily on the opinions of the involved security teams and often do not take into consideration the preferences of the target organization. Nevertheless, metrics that gauge specific aspects of these operations are extremely useful to raise security decision-makers' awareness on risk. The greatest challenge in quantifying the success or efficiency of threat-hunting efforts is to assess their effectiveness in reducing the number of successful attacks, controlling blast radius, and overall risk. Indeed, security is all about risk management. By nature, security-sensitive organizations are extremely risk adverse and wish to avoid any incident happens, while risk-tolerant organizations approach APTs more like business opportunities.

Effectiveness metrics should therefore capture these effects, be it directly or indirectly. Threat-hunting knowledge bases should be treated like vulnerability scanners; many of the tests rely on information externally available and are susceptible, just like with vulnerabilities scanners, to numerous false positives and unproven detections. As such, evaluations should include the sources available for knowledge generation, and should concentrate on known detections and the resources the community put together to mitigate the threat and the surety with which the detection is classified as a true positive.

8.1. Effectiveness metrics for threat hunting

Threat hunting is a different paradigm in cybersecurity detection. Instead of using signatures based on known attacks, players utilize Threat Hunting to detect stealthy, and often rare, but known type of attacks based on indicators of compromise and detection is not real-time but retrospective. Because of its retrospective nature, discrete-time point (TP) as defined in conventional statistics is considered as an unobserved event. This is different from cyber intrusion detection monitoring systems that aim to detect attacks in real-time using a straight-through detection mechanism.

Because of the nature of TP, an attack may sometimes be missed due to noise surrounding it. Although detection may be missed, a secondary monitoring signal score is captured which implies that an attack did take place, hence it can be termed as an enabled model for the hunting process and the score at any observation can be used as an indicator variable for hunting. The presence of a secondary hunting signal implies that at least parts of the hunting process logic become enabled with the detection narrowed by time and space parameters. Hence the corresponding score closer to the TP is more effective than further away score from TP. The score lead-time for discovering the attack corresponds to prep time between detection and discovery. The effectiveness of the hunting is evaluated using different hunting algorithms with varying geographical and temporal distances from the TP. An

analysis of the scores in closer post-detection neighborhoods indicates avenues for augmentation of overall detection capability and minimization of missed detection.

8.2. Evaluation methodologies and datasets

Evaluating the efficiency of AI-Enabled Cloud frameworks for Cybersecurity Intelligence involves assessing the overall model structure as well as the anticipated output. Various research studies in this area propose different metrics to evaluate this endeavor. Based on the existing literature, the proposed database architecture for Cybersecurity Threat Intelligence Model is evaluated using the following parameters: using different SDN protocols such as OPF, POF, PONF, etc. and their capabilities in advanced attacks detection, aggregating diverse types of logs from cloud, services, and applications perspective and their observability and visibility on the Log Investigation Analytical Engine, the concerns for traditional network logs (such as SNMP and Net Flow), incorporation of Network Function Virtualization (NFV) and Security Function Virtualization (SFV) concept in detection for DDoS, DOS, and SYN Flood Attack, and etc.

The proposed method of Threat Intelligence in Cloud Architecture mainly focuses on detection from cloud application perspective, Security as a Service (SaaS) with Malware Detection, Security Incidents and Event Management (SIEM) services, Threat Intelligence Analysis, Security for User Data Privacy, Artificial Intelligent-based Security and Privacy for Cloud Systems, Monitoring and Control during Migration of Cloud Control Transfer, Anti-DDoS Mitigation, and Enforcement, Controls for Privacy in Cloud Computing.

9. Case Studies and Practical Deployments

Although threat hunting is rarely dedicated to a single security service, the main drivers of cloud adoption are productivity and efficiency. Nevertheless, Cyber Atlantic, one of the biggest cloud service providers, experienced a security breach in 2020, which allowed attackers to investigate more than 100 million records. Threat hunting can help mitigate the risks of using the cloud, for instance, Misconfiguration of security groups, Internal reconnaissance, Lateral movement, Cloud credential compromise, and Data exfiltration. However, threat hunting is still a manual process that relies on human knowledge. Moreover, MITRE ATT&CK knowledge base published a cloud network topology on its cloud matrix to support threat detection, the joint effort between the University of Melbourne's Centre for Cyber Security Research and Innovation and SYNTHETIK produced the first set of cloud-native multicloud evasion datasets aimed at MITRE ATT&CK, and sees the release of their cloud-native multicloud datasets for attacker behavior, Creation of a Cloud-centric Cyber Kill Chain, and Threat Modeling as the basis for securing an organization's cloud environment.

The complexity of a hybrid cloud environment increases the attack surface and enhances the risk of negligence, thus highlighting the importance of continuous monitoring. Existing approaches tend to focus more on monitoring compliance than security. Tags generated by the organization's business side can greatly improve the efficiency of threat hunting in the hybrid cloud environment with a clear description of the different business services in the cloud infrastructure. Those tags can be used to search cloud asset data during the hunting process or to provide ground truth for supervised machine learning classification tasks.

9.1. Enterprise cloud deployments

A few implementations of the examined AI-enabled cloud threat-hunting frameworks have been detailed in the literature. A comprehensive collection of sensors–collectors for advanced threat hunting in Microsoft 365 and Azure environments of a large enterprise has been proposed. Distinct threat categories have respectively been modeled and described, and the corresponding detection use cases have been expressed in the formaql language of the Microsoft Sentinel SOAR platform. Several threat-hunting cases have been presented, covering account abuse, exposed secrets, permission abuse, malware exposure and persistence, external attacks, and threat intel misuses.

Another implementation targets the threat-detection and incident-response phases of the lifecycle, relying upon Microsoft Sentinel event orchestration, Microsoft Power Automate process automation, and Microsoft Power BI visual analytics within the enterprise's M365 environment. Upon detecting anomalies or accessing suspecting anomalous artifacts, analysts are guided to the appropriate analytical playbooks. The clinical-trial phase constantly tests the effectiveness and efficiency of the implemented cases, as well as their performance, alert quality, cost, and playbook management.

9.2. Hybrid cloud threat hunting scenarios

Built on the foundation laid by Schrader et al. for risk-aware threat hunting in public cloud environments, a second case study applies the security architecture of a hybrid multisystem operator (MSO) to employ AI-enabled threat hunting across both private and public clouds. The private cloud enables sensitive data processing—such as subscriber authentication, billing, and guaranteed QoS control—within an environment securely isolated from eavesdroppers. Meanwhile, the AI and machine-learning models, exploring and traversing the digital risk graph, remain in the public cloud, benefitting from superior compute and storage resources while reducing capex for the private cloud. Threat hunters employ federated learning and transfer learning techniques to safely deploy model updates onto the public cloud architecture while maintaining the privacy of personally identifiable information (PII).

The hybrid threat-hunting architecture builds upon the preceding public-cloud contribution, replacing the AI, machine-learning, and risk-integrity service models with their suitable private-cloud equivalents: secure subscriber authentication, billing, and guaranteed QoS control, supported by a private-cloud telemetry and observability environment. The vulnerability data sources and supply-chain data remain within the public cloud. The red team draws on the AI infrastructure in the public cloud to execute broad-spectrum cyberattack simulations against the deployed private-cloud control- and data-plane businesses, producing attack ground truth that is subsequently used to train the security perimeter models.

10. Conclusion

AI-enabled threat hunting on cloud infrastructures can help proactively identify and respond to cyberthreats. Although much research has investigated the application of ML and data mining to various aspects of threat hunting, a systematic and comprehensive AI-enabled cloud threat-hunting framework does not yet exist. Moreover, supporting elements such as observability and telemetry in cloud infrastructures, scalable data-processing and storage architectures, and the establishment of ground truth for a variety of key scenarios have not yet received adequate attention.

A conceptual proposal built on the work of the research community has addressed the foundational aspects of the space. The proposal identifies and discusses the main components of dedicated cloud threat-hunting solutions, commonly supported by public-cloud-native services but often deployed as hybrid approaches. Addressing AI-Native Clouds, a proposal on the provision of cloud services for data science also indirectly supports the threat-hunting space, especially in observability and telemetry aspects. Addressing data sources, business applications, ingestion pipelines, data labeling, reproducibility, scalable AI, and hybrid approaches completes the high-level exploration of AI-enabled cloud threat hunting.

11. List of important References

[1]Digital Threat Lifecycle. (2020). The Digital Threat Lifecycle: A Disruptive Strategy for Building Intelligence-Driven Cyber Defence. Retrieved on 2020-08-25 from [Link]

[2]Ahlgren, B., & Lindgren, P. (2013). Cyber Security Exercise as a Knowledge Management Tool – Conceptual Design. 28th International Conference on Information Networking (ICOIN) (pp. 273–278). IEEE. doi:10.1109/ICOIN.2014.6799840.

[3]Almach, A., & Alzubi, A.D. (2021). Intelligence in Cybersecurity: What, Why and How? IEEE Access, 9, 126837–126866. doi:10.1109/ACCESS.2021.3119787.

[4]Alzubaidi, L., Hussain, B., Zhang, J., Fadhl, A.A., & Konovalov, S. (2020). Threat Hunting Framework into Cloud-Native Environments: A Deployment Architecture Based on a Threat Intelligence Strategy. Future Generation Computer Systems, 108, 221–235. doi:10.1016/j.future.2020.01.012.

[5]Anisi, M.H., & Sukkarieh, S. (2022). Industry 4.0 Cyber Physics Security Framework. World Journal of Engineering, 19(6), 1–12. doi:10.1108/WJE-06-2021-0364.

[6]Arshad, S.Z., Alghamdi, S.S.M., Alharbi, F.A., Alshehri, H.M., & Aldawood, A. (2021). Cyber Security Risk Management Framework in Health Care Sector Based on CSF Guidelines. IEEE Access, 9, 109013–109022. doi:10.1109/ACCESS.2021.3117393.

[7]Ashiq, H.E., Wang, H., Mohsin, S., & Ding, H. (2021). A Hybrid Defense Mechanism for Mitigating DDoS Attacks in the Malware as a Service Model. IEEE Access, 9, 89156–89169. doi:10.1109/ACCESS.2021.3080367.

[8]Chaim, M.D., & Hsu, C.Y. (2021). A Novel Cybersecurity Framework for Enhancing Application Security against SQL Injection Attacks. IEEE Access, 9, 83311–83321. doi:10.1109/ACCESS.2021.3083489.

[9]Mahboubi, A., Luong, K., Aboutorab, H., Bui, H. T., & Jarrad, G. (2024). Evolving techniques in cyber threat hunting: A systematic review. Journal of Network and Computer Applications.

[10]Varma, S. C. G. (2024). AI-enhanced cloud security: Proactive threat detection and response mechanisms. International Journal of Cloud Computing.

[11]Alanezi, M. (2024). AI-powered cyber threats: A systematic review. Journal of Cybersecurity Research.

[12]Schwartz, Y., Benshimol, L., Mimran, D., Elovici, Y., & Shabtai, A. (2024). LLMCloudHunter: Harnessing large language models for automated cloud threat intelligence extraction. IEEE Access.

[13]Rahmati, M.Explainable and lightweight AI for real-time cyber threat hunting in edge networks. IEEE Internet of Things Journal.

[14]Roy, S. AgenticCyber: A generative AI-powered multi-agent framework for adaptive cybersecurity threat detection. ACM Transactions on Privacy and Security.

[15]Abdiukov, T. AI-powered threat hunting: Designing real-time predictive security frameworks for cloud environments. Global Journal of Engineering and Technology Advances.

[16]Singh, R., Kumar, P., & Sharma, S. (2024). Intelligent cloud security frameworks using deep learning for cyber threat detection. IEEE Transactions on Cloud Computing.