CyberProof®
A UST Company

Migrating to cloud-native threat detection and response

Contents

Security operations have moved to the cloud

More organizations are modernizing and rationalizing their IT estates. As part of this process, business services, infrastructure, and applications are being migrated to cloud-based environments. These changes in business IT processes open up new attack surfaces and cyber risks. They require security operations teams to quickly adapt their practices to what is needed to detect and respond to security threats in their cloud environments.

This eBook aims to guide security teams to extend their threat detection and response practices effectively from on-prem. to cloud environments when migrating to Microsoft's Azure Security stack.

But first, in order to understand the benefits of adopting Microsoft threat detection and response technologies, it's important to understand the concept of cloud-native architecture that Microsoft adopts.

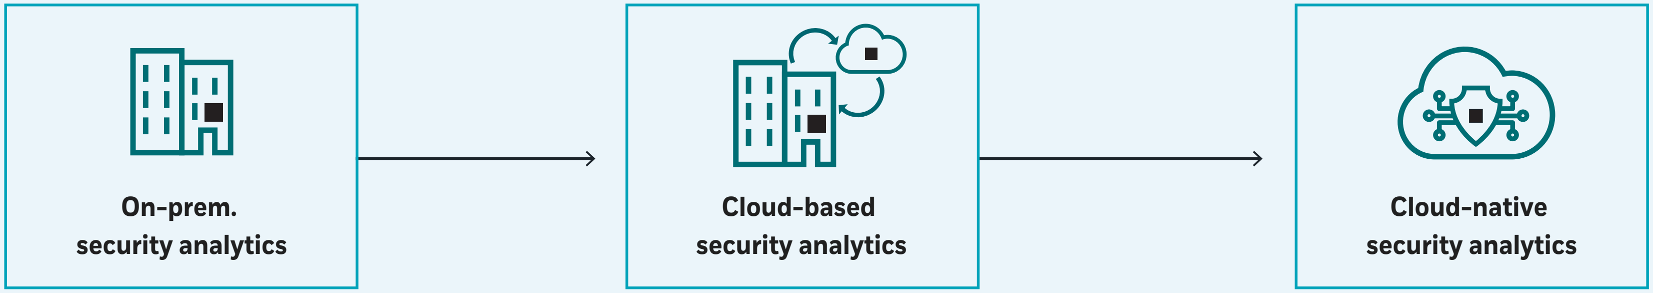

On-prem. security analytics

- Focused on on-prem. schema, not cloud
- Infrastructure costs:
  - Logging/hunting cloud logs is expensive
  - Initial purchase/licensing costs
  - Integration of data sources

Cloud-based security analytics

- Built for an on-prem. architecture but stored in a cloud-hosted environment
- SIEM does the analytics in the cloud but still needs additional compute/storage resources for enabling new functionalities
- Periodic updates to security features through software releases

Cloud-native security analytics

- Built for microservices-based architecture
- TCO Savings - server infrastructure, maintenance & storage on-prem.
- Continuous improvement of new use case content
- Elastic storage of security data adjusts as IT estate scales

EVOLUTION TO CLOUD-NATIVE SECURITY ANALYTICS

What is cloud-native security analytics?

Cloud-Based vs. Cloud-Native

As organizations migrate to the cloud, security operations teams are increasingly transitioning to Software-as-a-Service (SaaS) security analytics platforms to streamline operational costs and increase the speed of security delivery. This fast adoption is driving Security Information and Event Management (SIEM) solutions - primarily built for on-prem architectures - to transfer their codebase to a cloud-hosted environment, resulting in cloud-based security analytics.

There's an important distinction to make, however, between cloud-based and cloud-native security analytics. Solutions that are cloud-based are simply migrations from on-premises applications with a few modifications and often still require processed data to be sent back to a private, on-prem. datacenter, increasing the costs of data ingestion and slowing the threat detection process.

Cloud-based	Cloud-native
Built for on-prem. architecture but stored in a cloud-hosted environment	Built for a cloud-native architecture from the ground up
Integrates cloud/on-prem. systems that still require human interaction	Provides API-driven approach, making software/services integrations quick and easy to launch
SIEM does the analytics in the cloud but still needs to send information back to on-prem. data centers	Analytics sits on top of the cloud data lake enabling flexibility in scaling analysis, storage, and processing for big data
New security features and updates are made available in periodic software releases	DevOps approach enables continuous deployment of new detection rules, tuning, and enforcement as applications and services evolve
Restrictive storage limits and costs result in setup and maintenance of additional infrastructure to create new security content	Elastic storage allows storage and service processes to be scaled up and down on demand

When migrating applications, infrastructure, and services to the cloud, security monitoring should adopt cloud-native architecture to keep up with the pace of IT changes. Cloud-native architecture helps facilitate the collection and querying of data from a wide range of sources and supports the integration of third-party tools to carry out faster responses.

Microsoft Sentinel is a cloud-native SIEM that provides security data correlation, advanced analysis, and hunting of large volumes of events from hybrid environments to obtain high-context alerts. Microsoft Sentinel uses Machine Learning to proactively find anomalies hidden within acceptable user behavior and generate alerts.

It natively incorporates other foundational Microsoft services such as Azure Logic Apps to help build playbooks and connectors, enabling you to automate workflows and integrate with third-party services.

Built on top of a Log Analytics workspace, Microsoft Sentinel is able to store and query data collected from multiple sources at speed.

By using a cloud-native security analytics solution such as Microsoft Sentinel that is purpose-built for cloud scalability, you can identify threats faster as your digital estate changes.

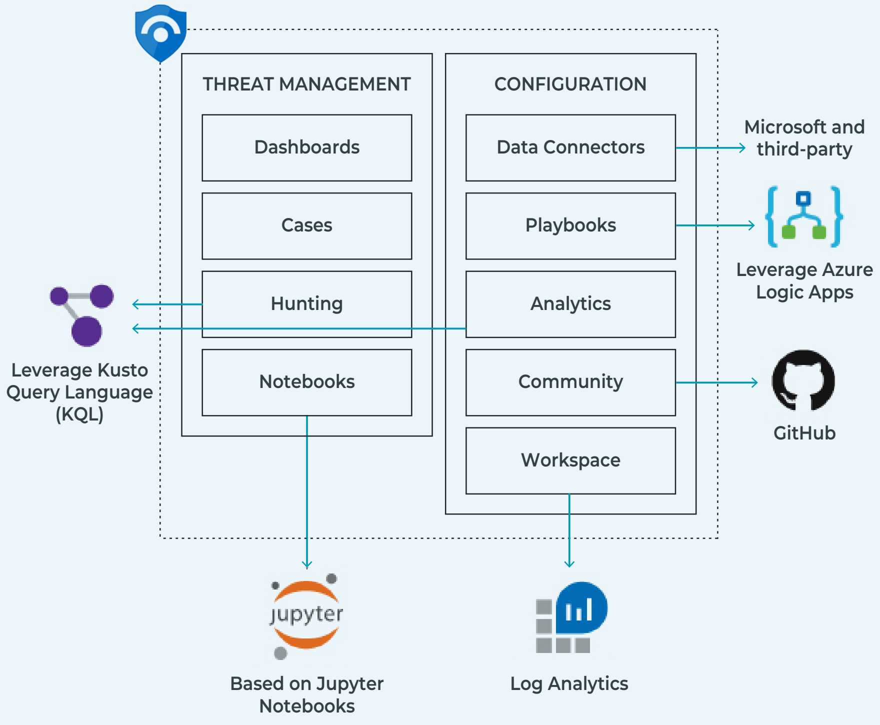

MICROSOFT SENTINEL ARCHITECTURE

Bringing in the right skills for for threat detection and response

Simply transferring the same security operations processes used for on-premises environments won't work with cloud-native architectures. There are various aspects of Microsoft-based Security Monitoring deployments that require up-to-date skills, experience, and knowledge of Microsoft technology as well as an understanding of the dynamics of cloud security policies and processes.

Acquiring these skills are all attainable if you spend the time and resources training your staff in these capabilities. However, just as important is the need to make sure you're bringing in practitioners with the right experience in applying these skills into real environments, or even better, into complex environments. This will help you to not only use best practices, but also focus on what really matters to your business. Backing up these skills with proven experience will ensure that deployment and migration is streamlined and doesn't introduce security risks.

Here are just a few examples of specific skills you should look for:

- **Program management –** highly organized managers with strong project management and stakeholder management experience. They should be able to put in place a governance and resource management plan when carrying out a large migration to cloud-based security monitoring infrastructures
- **Cloud forensics experts –** DFIR experts that specialize in cloud forensics; the ability to identify digital evidence and artifacts in cloud infrastructures
- **SIEM content engineering –** technically proficient engineers with experience in creating custom analytics rules, queries, and dashboards, and can leverage the automation, hunting and querying tools along with their own processes to reduce the risk of an incident
- **SC-200 certified –** Level 1 and Level 2 analysts that have been certified in the required skills to mitigate threats across, configure, and deploy Microsoft 365 Defender, and Microsoft Sentinel
- **Az-104 certified –** subject matter experts in implementing, managing, and monitoring Microsoft environments along with a strong understanding of core Microsoft services, Microsoft workloads, security, and governance
- **Az-500 certified security engineers –** subject matter experts in implementing security controls and threat protection, managing identity and access, and protecting data, applications, and networks
- A**z-400 DevOps Engineers –** designing and implementing Microsoft DevOps solutions using knowledge of Infrastructure as Code (IaC) to continuously improve services

If you don’t have these skills in place internally, consider working with a Managed Detection & Response (MDR) provider like CyberProof who has already invested in these multiple disciplines and can implement their resources alongside yours. MDR providers can act as an extension of your team both in terms of day-to-day operational support and for compliance and governance by leveraging a hybrid delivery model. This type of model uses tools and practices that enable transparency between internal teams and those of the service provider, augmenting your activities without requiring you to relinquish knowledge or control.

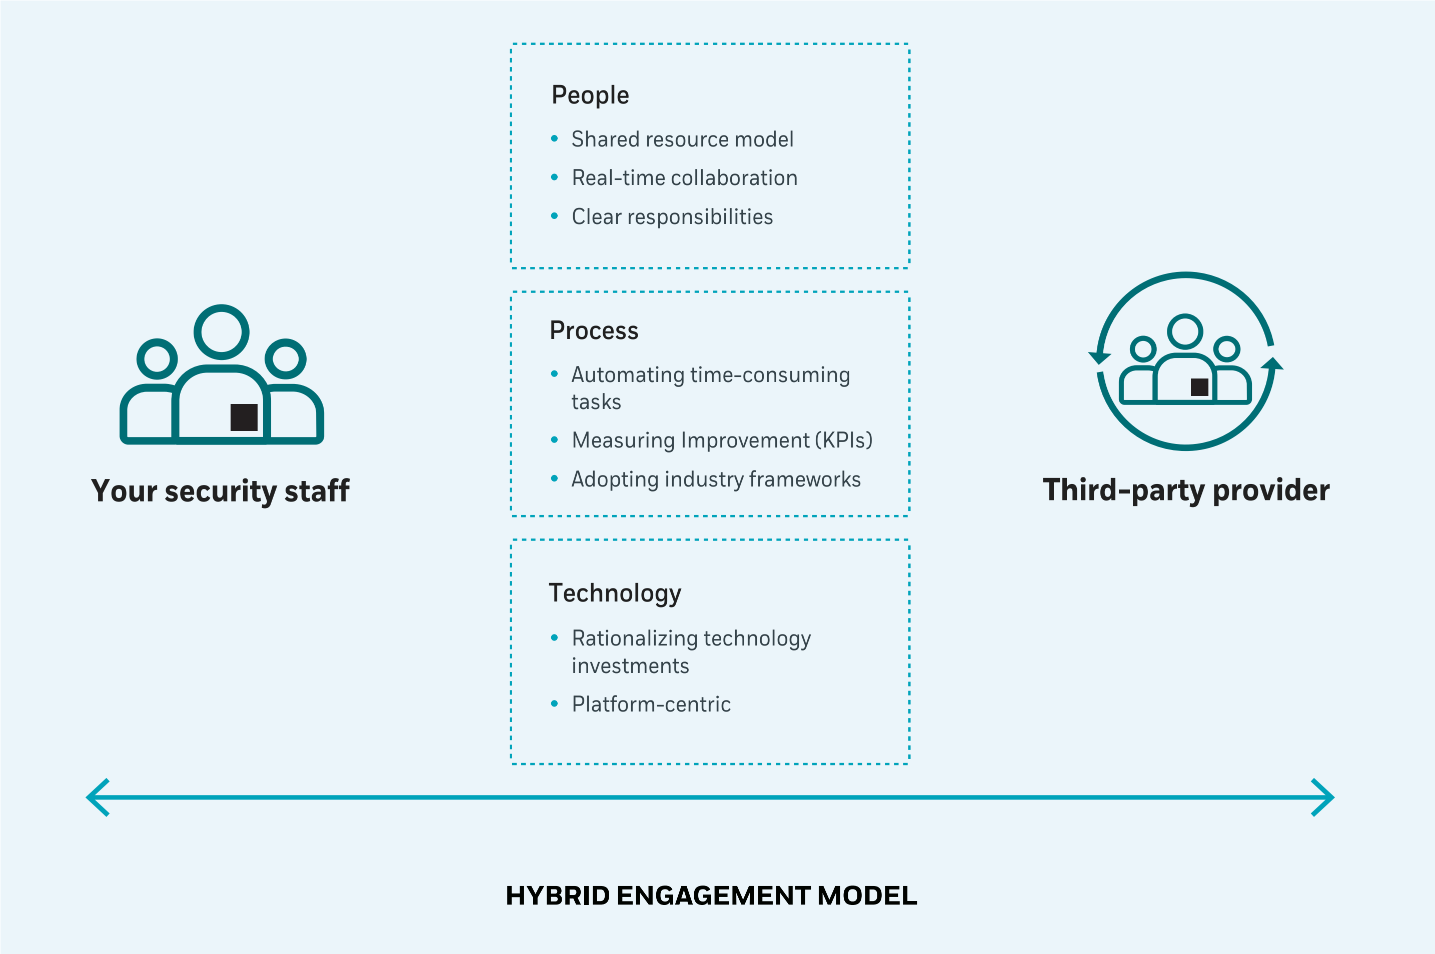

People

- Shared resource model
- Real-time collaboration
- Clear responsibilities

Process

- Automating time-consuming tasks
- Measuring Improvement (KPIs)
- Adopting industry frameworks

Technology

- Rationalizing technology investments
- Platform-centric

HYBRID ENGAGEMENT MODEL

Defining use cases to maintain accurate threat visibility during migration

You can't detect a threat that you can't define. Developing use cases is a way of defining not only which types of threats to monitor, but also how best to respond to them. This is particularly important when onboarding new cloud sources that provide new attack surfaces. Use cases will also help you evaluate which log sources from your legacy SIEM should be transitioned over to Microsoft Sentinel.

This is an important part of the migration, which is frequently missed when organizations evaluate log sources, resulting in a large increase of more complex and noisy alerts. Done correctly, however, the evaluation process makes threat detection and response processes much easier to manage. It also makes the system as a whole more accurate and allows for faster threat mitigation. But first, let's define what a use cases actually means.

Use cases are commonly associated with sources to collect relevant data and apply corresponding detection rules to a SIEM, which automatically raises an alert when a threat is identified. However, this definition of a use case is an outdated process that focuses on filling in detection gaps in the SIEM rather than reducing the risk of the primary threats to your business critical assets.

Using the traditional approach often results in noisy alerts that provide very limited context regarding the severity and the potential threat to a critical asset. This approach also lacks agility in detecting and responding to the latest threats, because of the difficulty in customizing the content as your infrastructure and threat landscape conditions change.

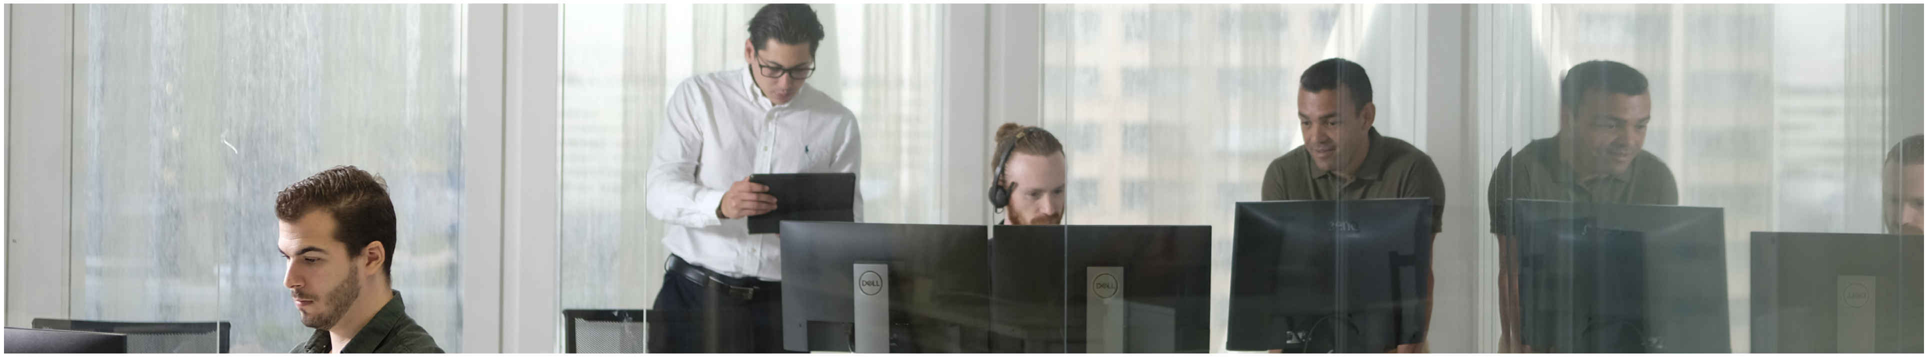

Traditional Process

The Result

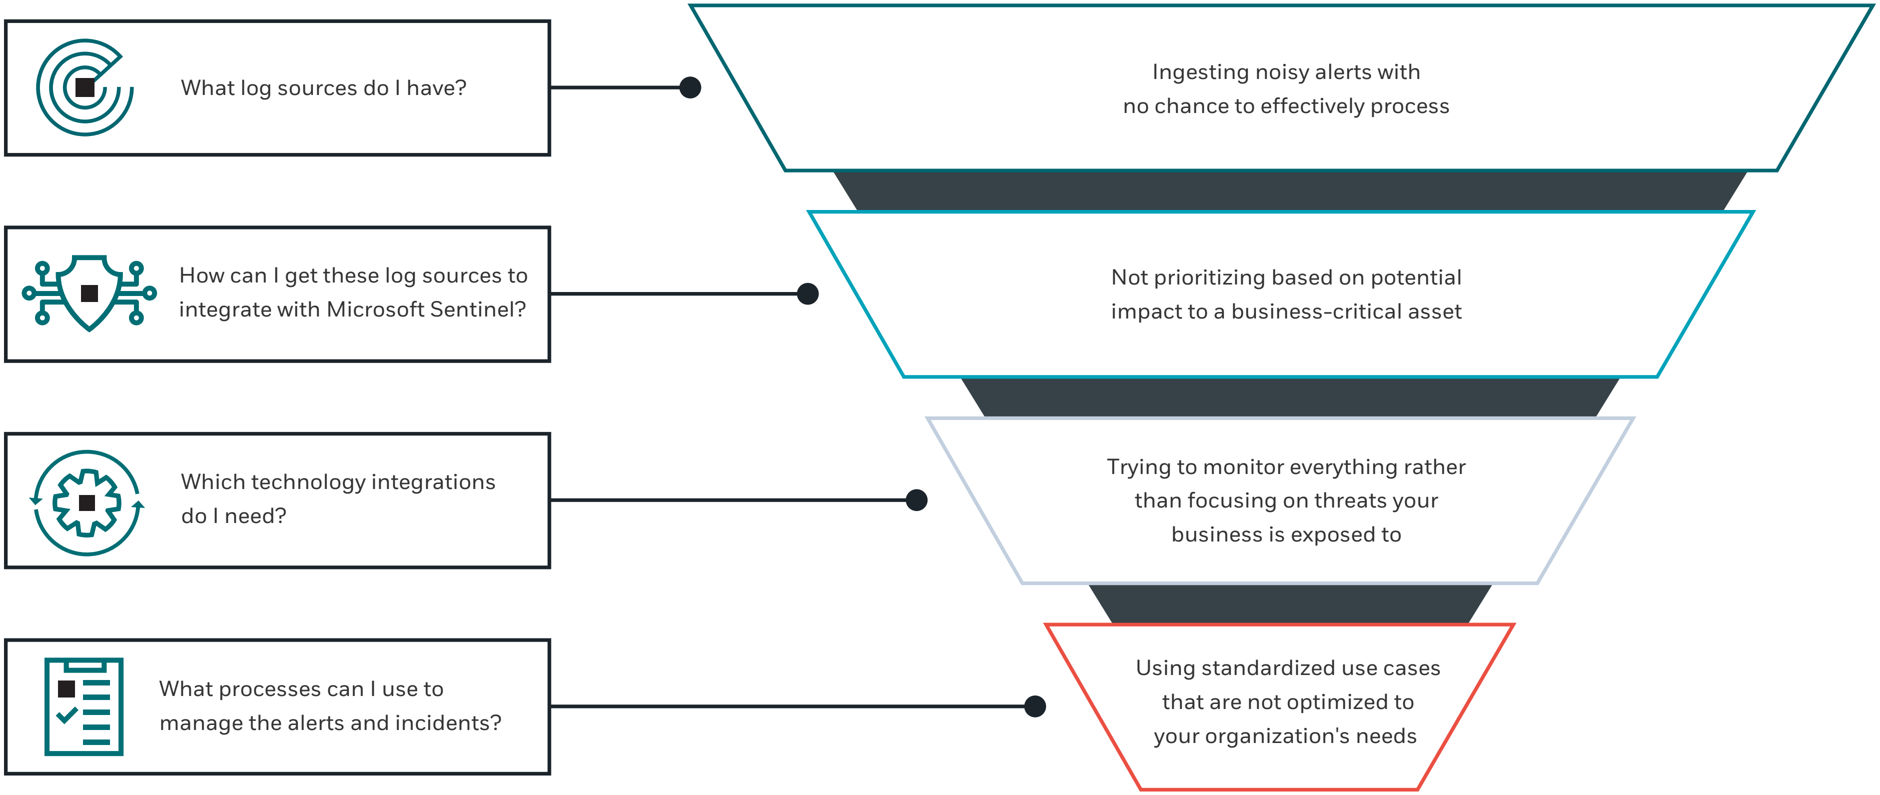

What log sources do I have?

How can I get these log sources to integrate with Microsoft Sentinel?

Which technology integrations do I need?

What processes can I use to manage the alerts and incidents?

Ingesting noisy alerts with no chance to effectively process

Not prioritizing based on potential impact to a business-critical asset

Trying to monitor everything rather than focusing on threats your business is exposed to

Using standardized use cases that are not optimized to your organization's needs

In reality, a use case – sometimes referred to as an attack scenario – should represent the outcome of an attack, or the attacker's desired outcome with targeted business assets. Defining this outcome requires you to identify which assets you want to protect, the systems that interact with those assets, and the techniques, tools, and procedures (TTPs) that would be required by an attacker to exploit those systems.

Once these use cases are defined, you can start evaluating your detection and response gaps that need to be filled, the required content – log sources, analytic rules, response playbooks, data connectors, automations, KQL queries, reporting, etc. –and response times that need to be met in order to mitigate the negative impact to the business.

Optimized Use Case Process

The Result

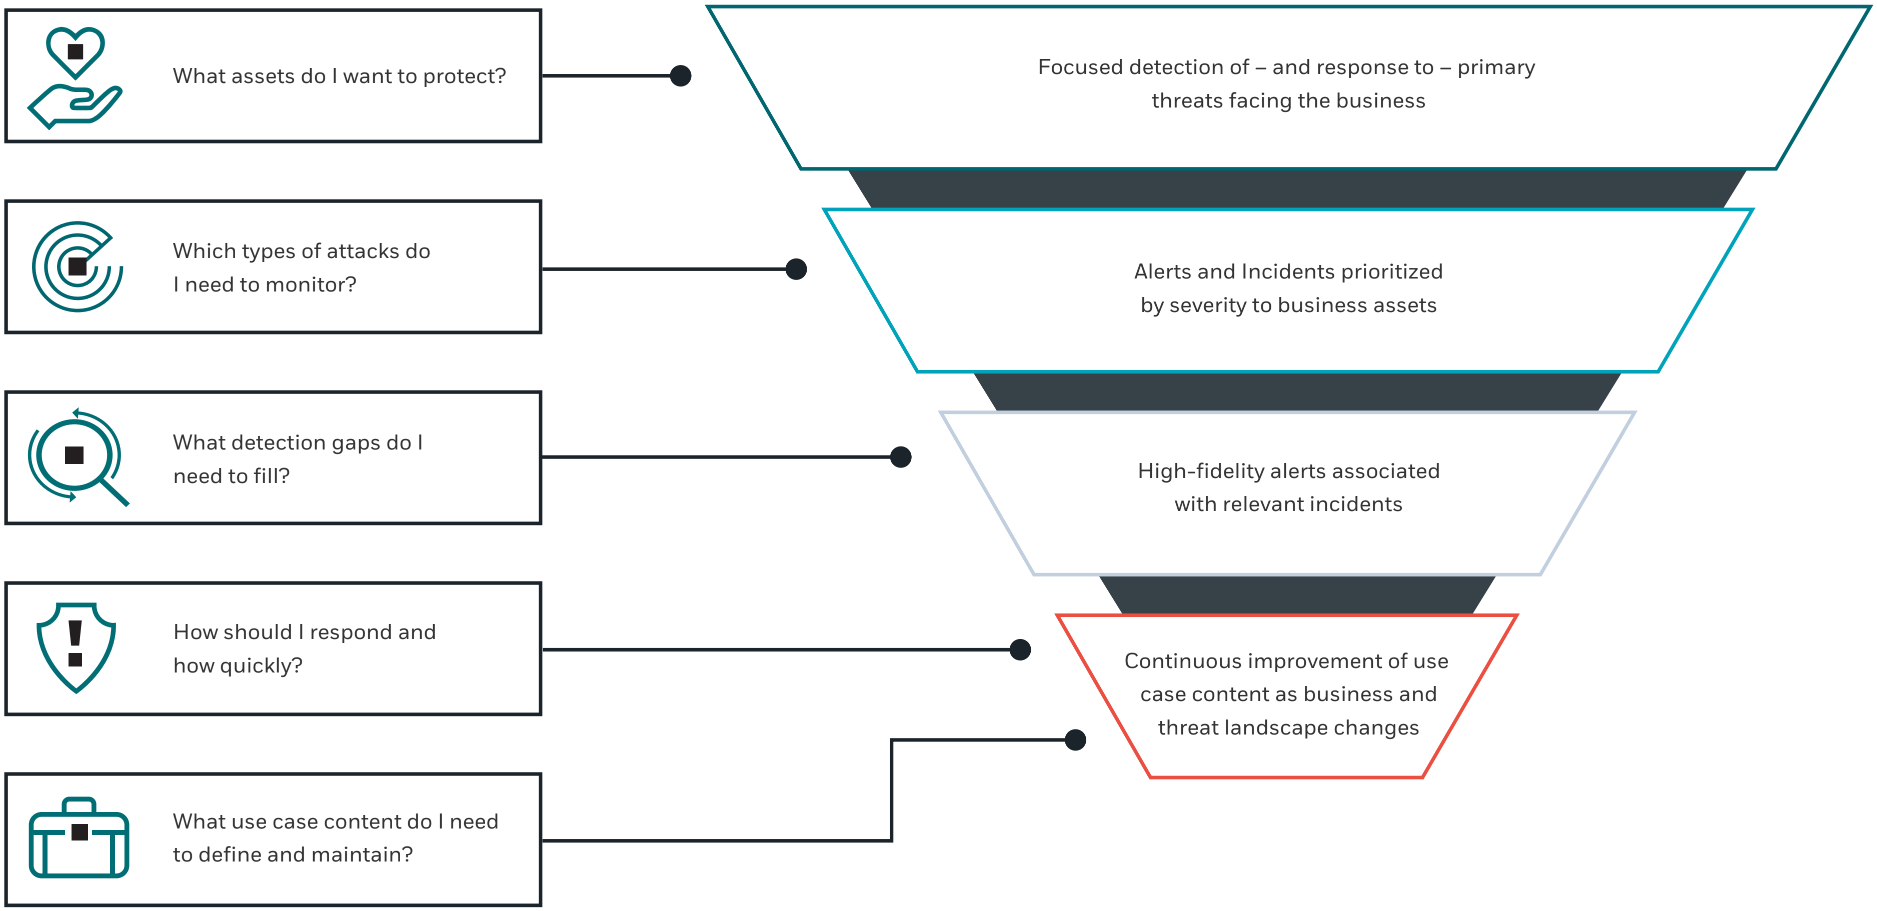

What assets do I want to protect?

Which types of attacks do I need to monitor?

What detection gaps do I need to fill?

How should I respond and how quickly?

What use case content do I need to define and maintain?

Focused detection of – and response to – primary threats facing the business

Alerts and Incidents prioritized by severity to business assets

High-fidelity alerts associated with relevant incidents

Continuous improvement of use case content as business and threat landscape changes

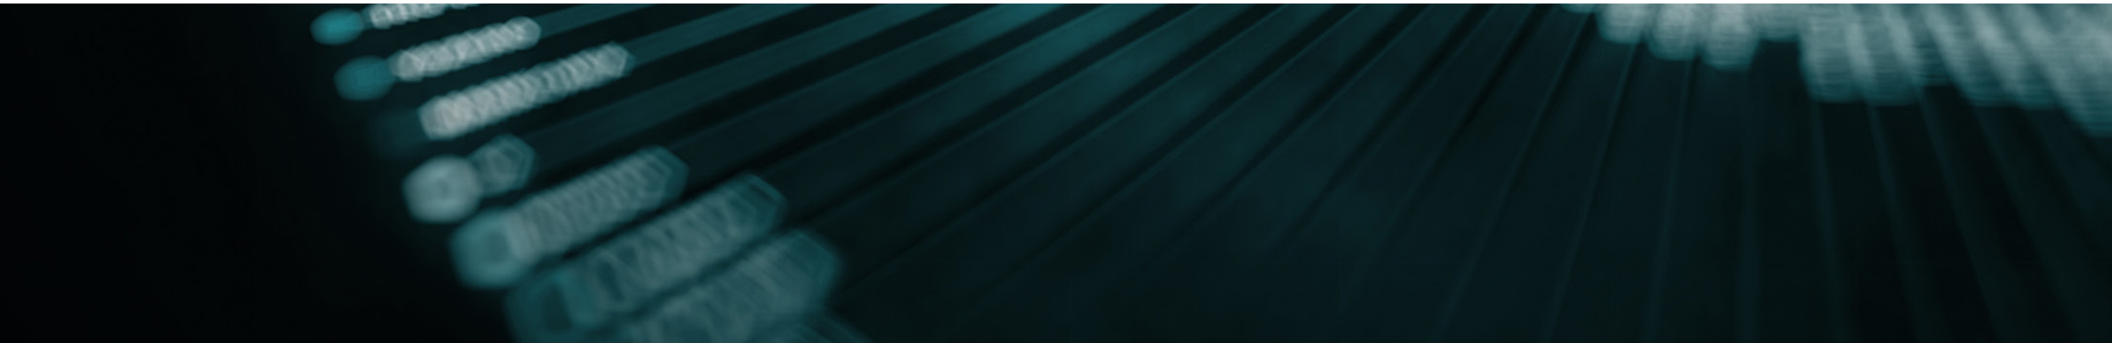

To help with fast onboarding of use cases, Microsoft Sentinel includes out-of-the-box detections covering 40+ connectors and 12+ MITRE tactics. Remember, however, that without combining with other components such as relevant response playbooks, queries, reports packaged together for a specific Use Case, responses will be slower, less accurate, and harder to optimize as attack scenarios change.

This is why at CyberProof, we've created a catalog of over 200 Microsoft Use Case Kits that package the content and enable fast onboarding and visibility of cloud threats.

This catalog is continuously updated in line with MITRE ATT&CK which means our managed detection and response customers are able to have new Use Case Kits deployed on a continuous basis.

Example: Cloud-related use case kit

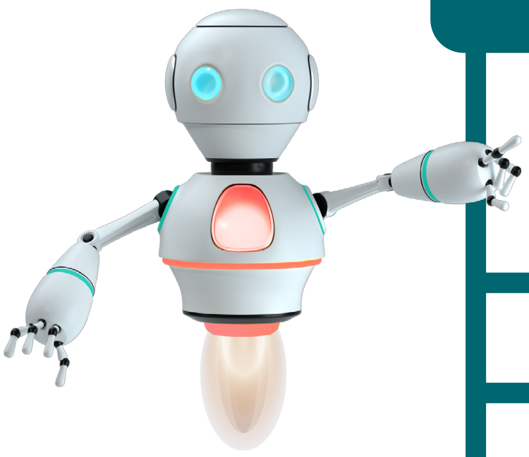

Suspicious number of resource creation or deployment activities in Microsoft

Description
Indicates when an anomalous number of VM creations or deployment activities occur in Microsoft Sentinel

Business Loss Outcome
Customer Portal Takedown

Techniques
Account Manipulation

Log Source Type
Cloud-Microsoft Sentinel events

Kill Chain
Actions on Objectives, Target Manipulation

Automation Maturity
Partially Automated

Enrichment
CyberProof Threat Intelligence Feeds, Office 365 Logs

Response Playbook

- Data Gathering – Gather data from SIEM including Resource ID, Operation Name, UserID, Timestamp, Activity Status, IP Address
- Investigation – Check if ClientIP is a known threat actor (WhoIS, Virustotal, Cymon.io). Look for action performed by userID in log sources AD, O365, firewalls, VPN, proxy for suspicious footprint
- Escalation – Classify alert severity, send email to soc@xyz.com

Transitioning use cases

During the initial stages of the migration process, it's best to continue running your legacy SIEM in parallel to a staged version of Microsoft Sentinel. This way, you can dual-feed use case content from your legacy SIEM, including detection rules and security events, into the staged environment. You can test and tune this data without impacting the production environment.

But how do you maintain visibility of this content on both platforms if they're running in parallel? Using a centralized service delivery platform with orchestration and automation capabilities will enable you to manage multiple SIEMs, while maintaining a single pane of glass for security monitoring. This will enable a phased transition without losing functionality, because no changes will be made to log sources in the initial stages - and rules and alerts can be tested early before you migrate fully to Microsoft Sentinel.

With this setup, your new cloud log sources go directly into Microsoft Sentinel, along with the use cases that are related to them. The use cases related to existing on-prem. log sources can be built on the legacy SIEM and then migrated to the dual-fed Microsoft Sentinel.

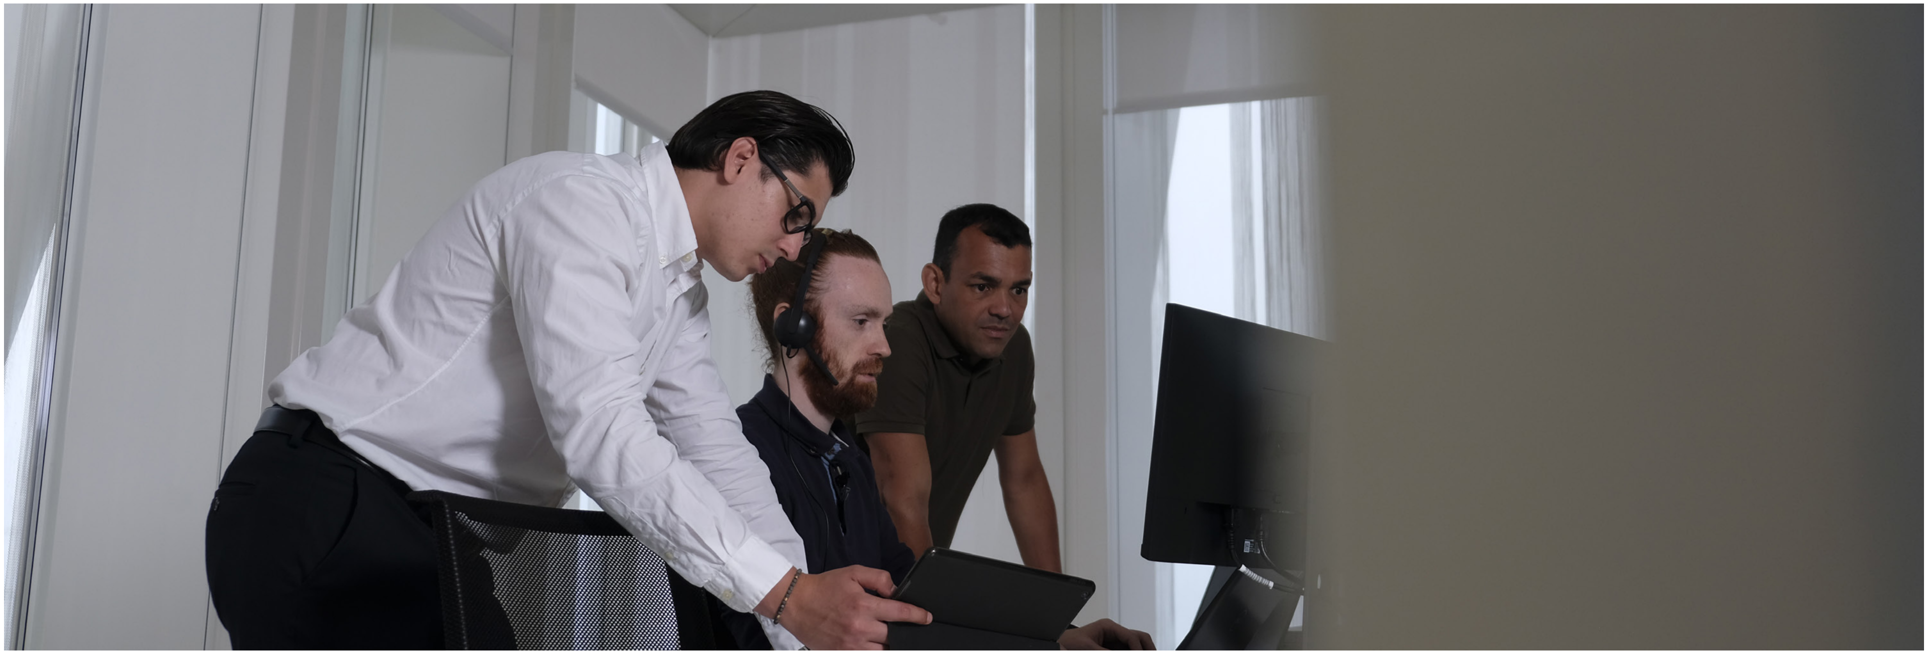

Working with an MDR provider like CyberProof that has a centralized platform pre-integrated with Microsoft Sentinel helps significantly. An MDR provider can support at this stage of the migration - but make sure that they give you full transparency, i.e., that you have as much visibility into activity as the provider's team has. For example, at CyberProof we leverage the CyberProof Defense Center (CDC) platform to help migrate from legacy tools to Microsoft Security platforms such as Sentinel and Defender. It acts as a single pane of glass before, during, and after migration, ingesting and triaging alerts from SIEM, EDR, Vulnerability Scanning, and Threat Intelligence Platforms to facilitating incident investigation and response.

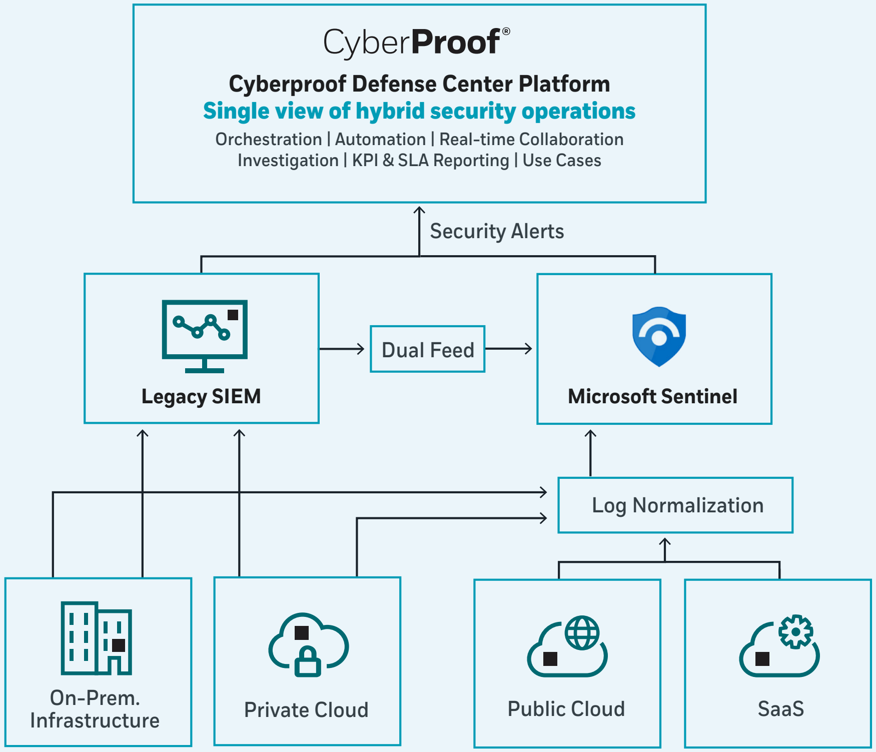
CyberProof®
Cyberproof Defense Center Platform
Single view of hybrid security operations
Orchestration | Automation | Real-time Collaboration
Investigation | KPI & SLA Reporting | Use Cases
Security Alerts
Legacy SIEM
Dual Feed
Microsoft Sentinel
Log Normalization
On-Prem. Infrastructure
Private Cloud
Public Cloud
SaaS

CENTRALIZED VIEW OF ALERTS FROM BOTH SIEMS

Optimizing log collection for faster and more cost-effective detection

An extensive log collection layer should cover on-prem. applications, cloud-hosted apps, and SaaS applications. It should extend to all of your organization's log sources – including those connected to regional offices, remote workers, and data centers (where relevant).

This means you're getting visibility into your cloud environments as quickly as possible. Microsoft Sentinel uses a cloud-native Log Analytics Data Lake which sits underneath its security analytics layer, enabling seamless data storage, management, and analytics at scale.

When adopting a cloud-native Detection and Response strategy, there are numerous out-of-the-box API data connectors supporting native and external sources that can stream data into the Microsoft Security Stack easily without custom configurations. The diagram here illustrates those that can be included by default.

Data Connectors

- Microsoft Activity
- AWS CloudTrail
- GCP Audit Logs
- Microsoft Security Center
- Azure Active Directory
- Azure Active Directory Identity Protection
- Microsoft 365
- Microsoft Defenders for Cloud Apps
- Microsoft Defender for Identity
- Microsoft Defender for Endpoint

If you are using components of Microsoft Defender, the range of existing security controls generating security alerts include identities, endpoints, data, email, collaboration, IoT, OT, cloud infrastructure, and cloud applications – all of these data sources must be included in the log collection layer.

Think of Microsoft Defender as a suite of tools that actually carries out the prevention, detection, and remediation of threats to end-user (365 Defender) and infrastructure (Microsoft Defender for Cloud) assets. In contrast, Microsoft Sentinel enables the overarching security visibility and management of the alerts and incidents coming in from Defender or third-party sources.

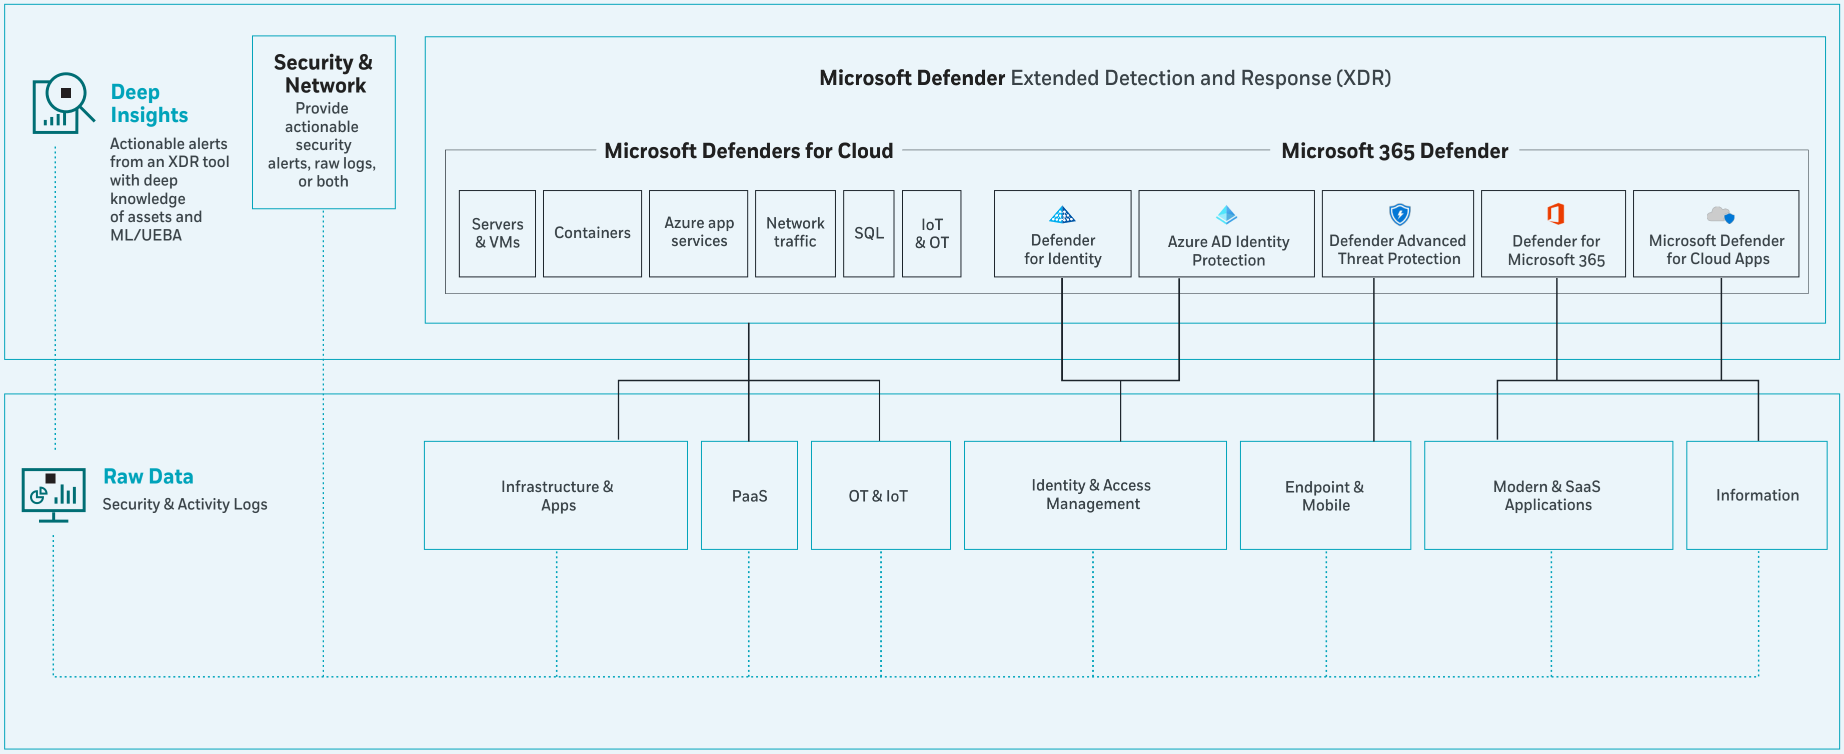
Deep Insights
Actionable alerts from an XDR tool with deep knowledge of assets and ML/UEBA
Security & Network
Provide actionable security alerts, raw logs, or both
Microsoft Defender Extended Detection and Response (XDR)
Microsoft Defenders for Cloud
Microsoft 365 Defender
Servers & VMs
Containers
Azure app services
Network traffic
SQL
IoT & OT
Defender for Identity
Azure AD Identity Protection
Defender Advanced Threat Protection
Defender for Microsoft 365
Microsoft Defender for Cloud Apps
Raw Data
Security & Activity Logs
Infrastructure & Apps
PaaS
OT & IoT
Identity & Access Management
Endpoint & Mobile
Modern & SaaS Applications
Information

MICROSOFT XDR ARCHITECTURE

There are several ways to ensure you're collecting and migrating relevant logs from hybrid environments into Microsoft Sentinel. What's important is to select methods that are:

- Fast
- Cost-efficient
- Secure
- Protect the integrity or availability of the data
- Support compliance requirements

For virtual machines and other compute resources, for example, Azure Monitor or Log Analytics agents will be required. On-prem. machines may require a Syslog or CEF forwarder. Security event data can be collected using Microsoft's native collectors supporting cloud and on-prem. infrastructure and services.

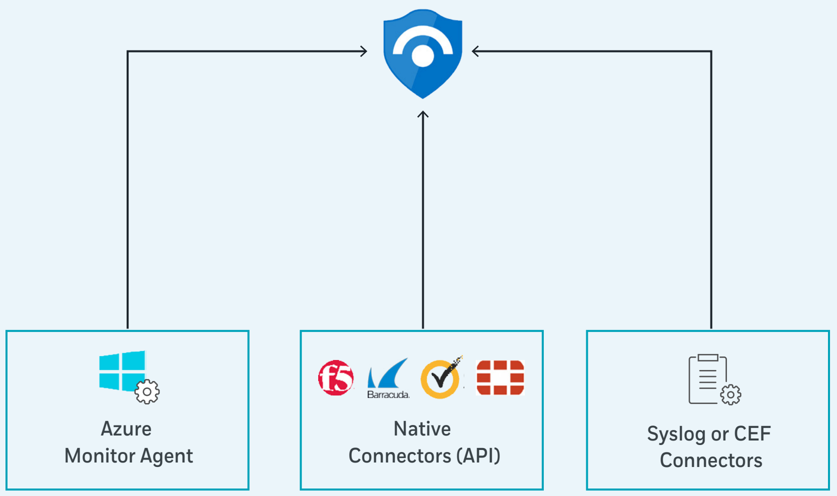

STANDARD MICROSOFT SENTINEL CONNECTORS

As new applications, services, and systems are adopted, some log sources may not be supported. This might happen at the outset or could come up in the future. When these new assets generate data and events in unsupported formats, this will slow down threat detection efforts.

Without finding a way to collect data from sources that aren't covered natively by Microsoft Sentinel, you're essentially leaving blind spots that attackers can use to hide in your network undetected.

To solve this, you should develop a continuous process for collecting new formats of data through non-native collection methodologies. If you have data scientists in your team, we recommend leveraging their experience in coding and data engineering to leverage PowerShell scripts and tools such as LogStash to create custom parsers that filter unstructured data, identify relevant fields, and converge on a common format.

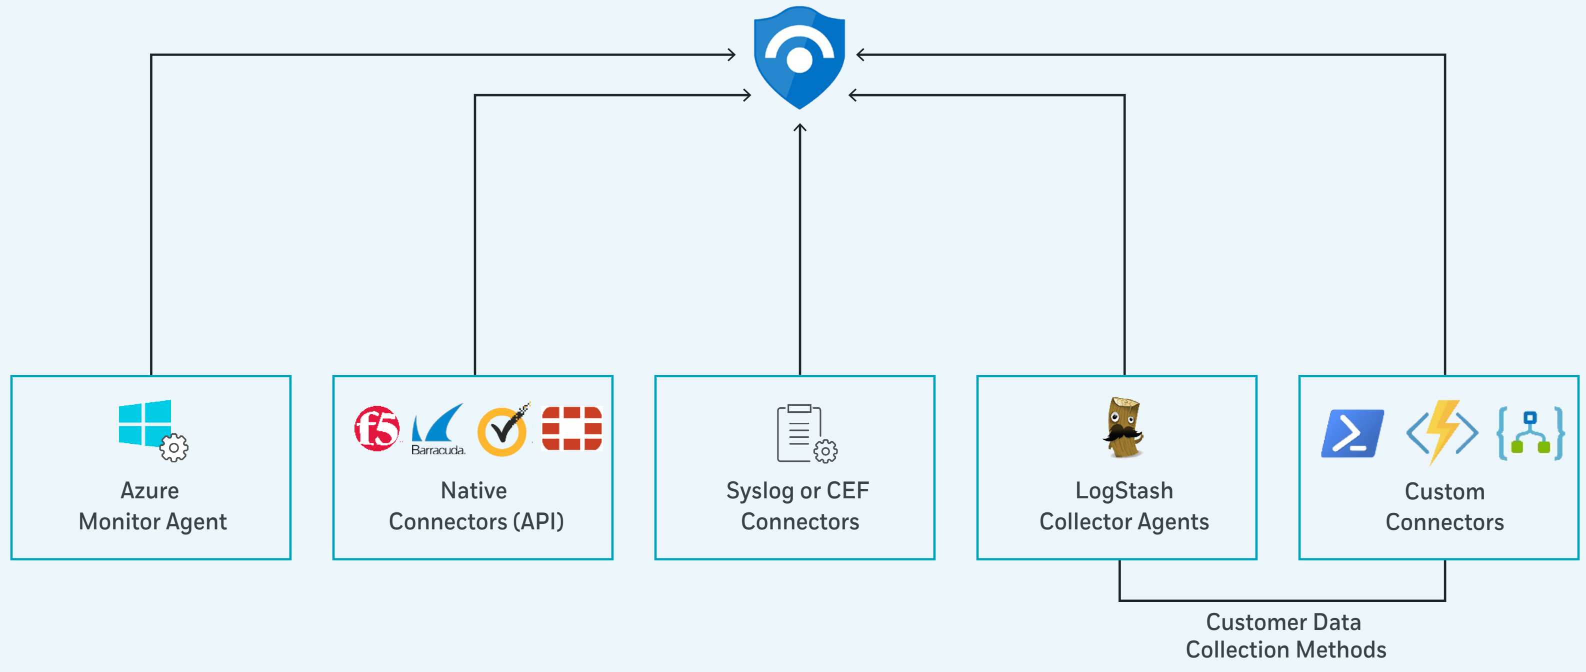

WIDER COVERAGE OF LOG SOURCES WITH BOTH STANDARD AND CUSTOMIZED CONNECTORS

If your business is a multinational enterprise or has multiple entities, you'll be needing to tag and filter a huge amount of data while providing Role-Based Access Control (RBAC). This will require extensive computing power and incur significant log ingestion and retention costs.

You can dramatically reduce these costs by leveraging a long-term data lake solution such as Azure Blob Storage or Azure Data Lake (ADX), a Platform-as-a-Service (PaaS) solution for big data analytics and data exploration by Microsoft. These solutions allow you to store low fidelity telemetry while routing parsed, cleaned vent data – specific for threat detection use cases – into your Azure Security Stack.

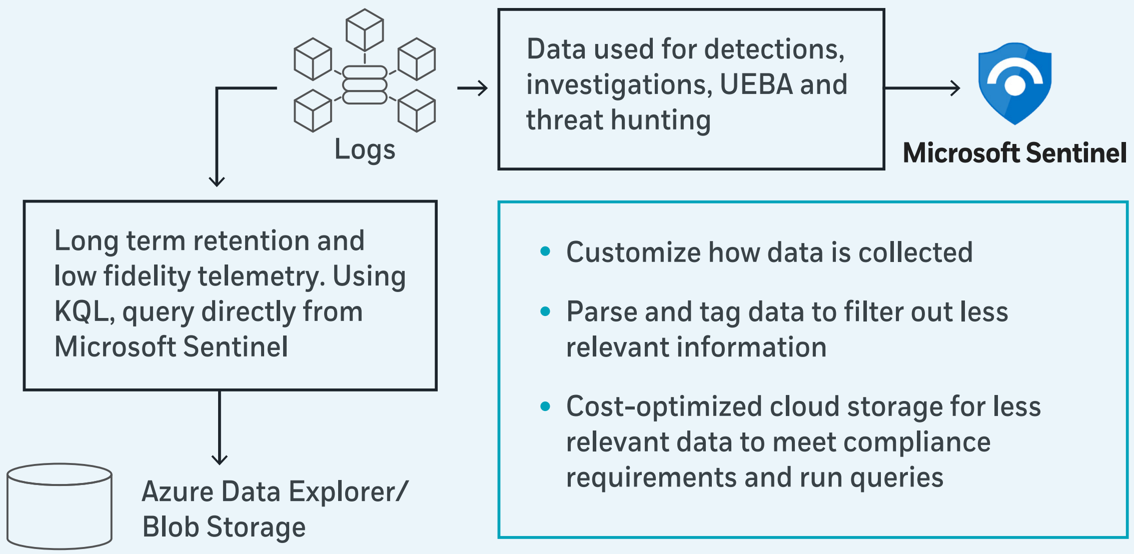

Data used for detections, investigations, UEBA and threat hunting

Long term retention and low fidelity telemetry. Using KQL, query directly from Microsoft Sentinel

Azure Data Explorer/ Blob Storage

- Customize how data is collected
- Parse and tag data to filter out less relevant information
- Cost-optimized cloud storage for less relevant data to meet compliance requirements and run queries

REDUCING COSTS OF LOG INGESTION AND RETENTION

How the CyberProof Log Collector can help

The CyberProof Log Collector (CLC) is designed to solve the issue of log ingestion and retention costs. The CLC is based on a microservices architecture and plays a key role in helping organizations transition to Cloud-native MDR. The CLC is purpose-built to collect all types of data from any source at scale using a container model. It can take any log and handle the parsing, tagging, cleaning, and aggregation of the data before it is ingested into Microsoft Sentinel.

The CLC also works with any programming language, including Python, PowerShell, .NET, etc.

The CLC improves the flow and handling of data, augmenting Microsoft Sentinel's predefined rules and capabilities to provide customers with automated and dynamically updated threat detection.

The CLC reduces costs by more than 40% due to the filtering of log data and routing of less relevant data into a cost-effective, cloud-native storage solution.

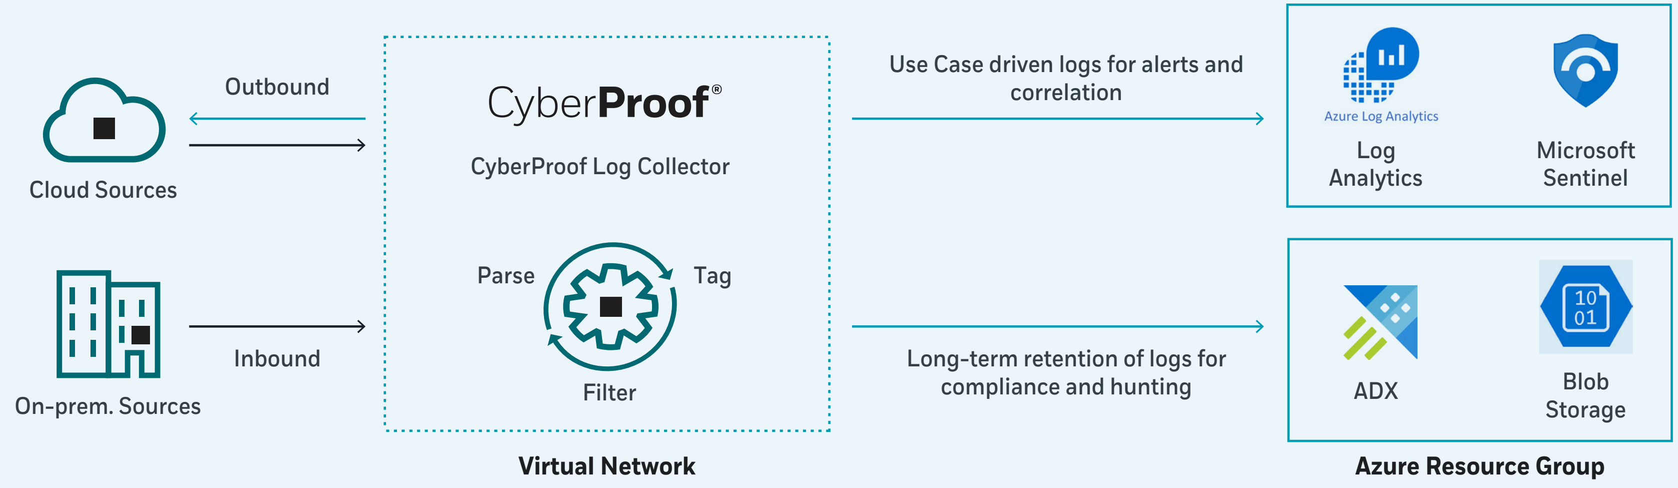

CYBERPROOF LOG COLLECTOR

Applying custom analytics, reporting, and queries

A key stage in cloud migration involves optimizing alert triage and response. This begins when an organization starts to receive a reliable stream of clean data in Log Analytics. Microsoft Sentinel's analytic rules and querying capabilities should be leveraged to optimize the process.

Analytic Rules are rules that can be created within Microsoft Sentinel using their native Kusto Query Language (KQL) to define the logic for automating the detection of threats. Within the Analytics Rule Wizard, you can set the parameters, the type of data that forms part of an alert (such as IP addresses or host names, etc.), and schedule queries to run at certain times. In a single query, you can search for Indicators of Compromise (IOCs) in the form of either malicious IP addresses, domains, URLs, or hashes.

Having a healthy threat intelligence collection and dissemination strategy is key here, because one of the applications of your threat intelligence gathering will be to apply relevant IOCs from these feeds to your Analytic rules in Microsoft Sentinel. These rules can then search for specific IOCs in your environment.

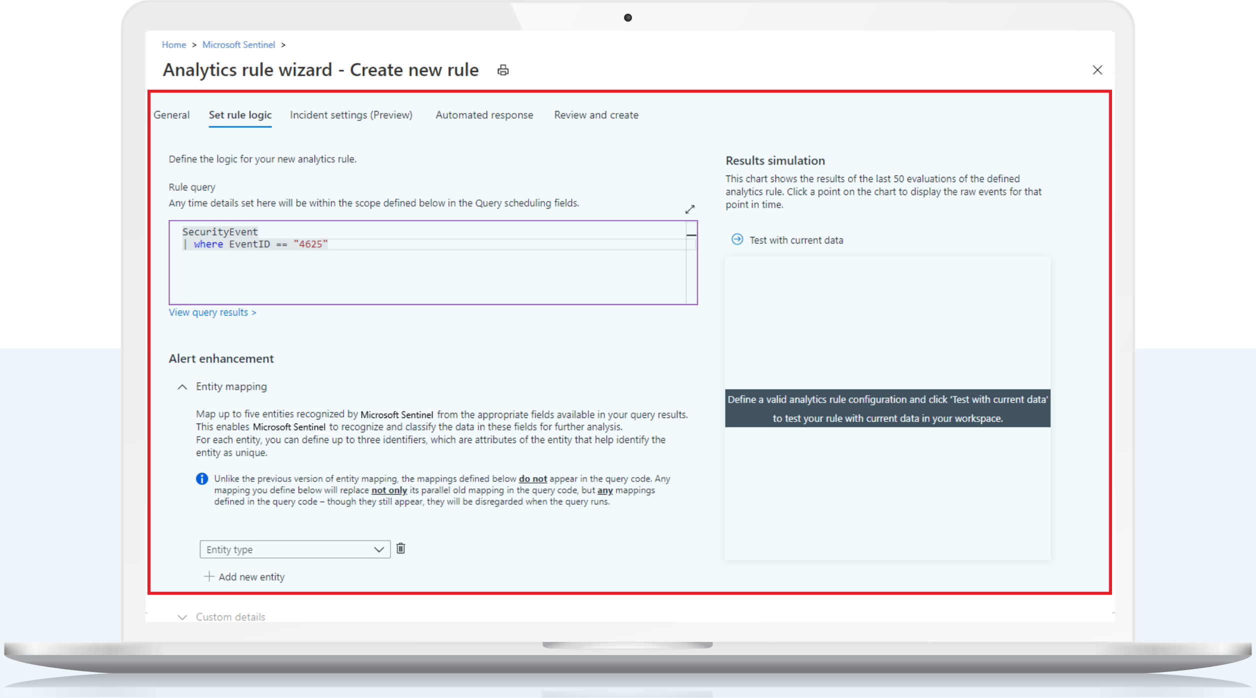

Analytics rule wizard - Create new rule

Define the logic for your new analytics rule.

Rule query

Any time details set here will be within the scope defined below in the Query scheduling fields.

```
SecurityEvent
| where EventID == "4625"
```

Alert enhancement

Entity mapping

Map up to five entities recognized by Microsoft Sentinel from the appropriate fields available in your query results. This enables Microsoft Sentinel to recognize and classify the data in these fields for further analysis. For each entity, you can define up to three identifiers, which are attributes of the entity that help identify the entity as unique.

Unlike the previous version of entity mapping, the mappings defined below **do not** appear in the query code. Any mapping you define below will replace **not only** its parallel old mapping in the query code, but **any** mappings defined in the query code – though they still appear, they will be disregarded when the query runs.

Results simulation

This chart shows the results of the last 50 evaluations of the defined analytics rule. Click a point on the chart to display the raw events for that point in time.

Define a valid analytics rule configuration and click 'Test with current data' to test your rule with current data in your workspace.

MICROSOFT DOCUMENTATION, "CREATE CUSTOM ANALYTICS RULES TO DETECT THREATS," 2021

This is where having an MDR partner like CyberProof is particularly helpful. An MDR partner can manage the creation and execution of relevant queries and leverage access to various sources of threat intelligence - open source, commercial feeds, dark web marketplaces, and forums - to integrate more tailored IOCs and TTPs into these queries, enabling you to find the right types of threats faster. At CyberProof, we're continuously updating custom queries based on the client's threat profile. When a query generates an alert in Microsoft Sentinel, this alert is shown in our CDC platform where our Virtual Analyst, SeeMo, automates the alert triage, enrichment, and containment actions as defined by the playbook. If this is linked to other alerts being generated by Sentinel, SeeMo will match these alerts under a single incident.

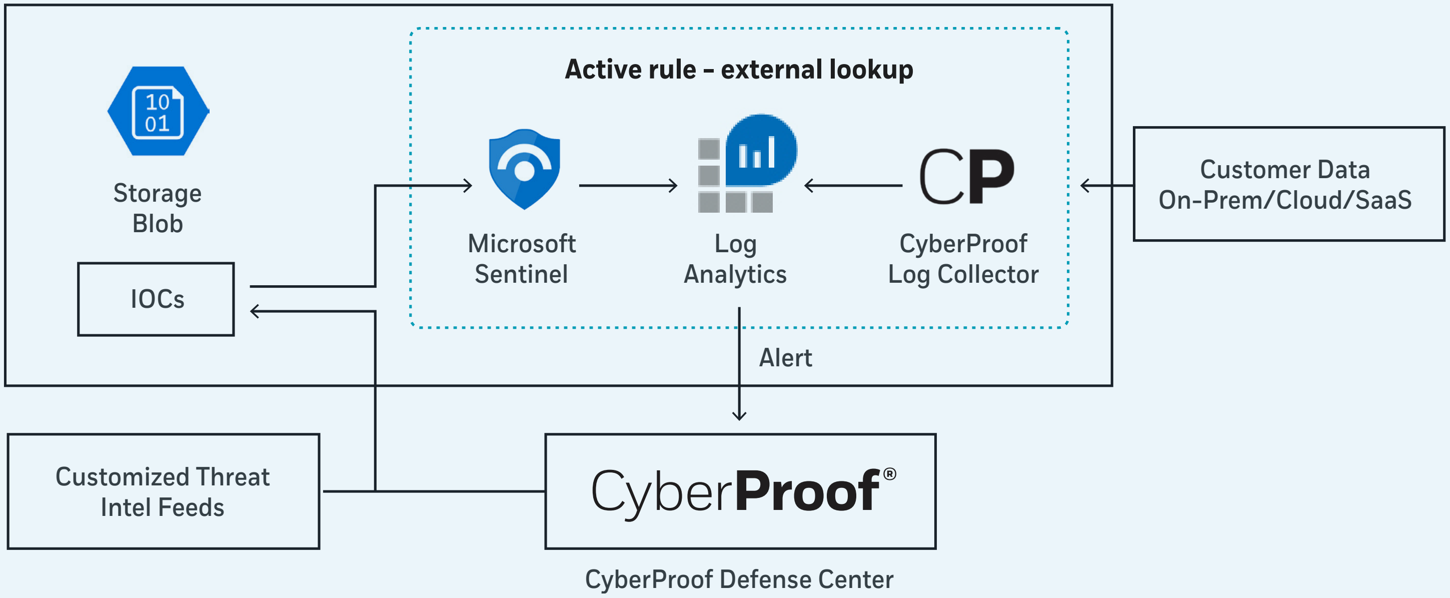
Active rule - external lookup
10 01
Storage Blob
IOCs
Microsoft Sentinel
Log Analytics
CP
CyberProof Log Collector
Customer Data On-Prem/Cloud/SaaS
Alert
Customized Threat Intel Feeds
CyberProof®
CyberProof Defense Center

INTEGRATING THREAT INTELLIGENCE-DRIVEN IOCS

Using Azure Data Explorer (ADX) as an Investigations Tool

We've established that log storage solutions such as ADX can be a viable option to reduce the costs of log ingestion and retention. But this platform can provide a lot more value to security operations teams.

Given that ADX uses the same query language (KQL) as Microsoft Sentinel, it's possible to scale hunting and reporting capabilities across the two platforms. This helps accelerate threat detection processes while optimizing costs of using Microsoft Sentinel. ADX can consume huge volumes of data for long periods of time in structured and unstructured formats. If optimized correctly, leveraging this information can create dynamic workbooks for visualizing data analysis results, as well as hunting rules to query large volumes of data for suspicious activitiy over longer time frames.

We therefore recommend leveraging ADX as part of your cloud-native detection and response practices, so you're able to scale your investigations and hunting activities as new cloud sources are added without significantly driving the costs up.

These tools will be essential parts of your use case development cycle, as you'll want a continuous and agile process for carrying out these queries and workbook customizations.

Summary

With more services, applications, and infrastructure transitioning to and being created in the cloud, migrating to a cloud-native detection and response strategy is highly recommended.

The Azure Security Stack is an advanced suite of technologies that will only become more advanced and convenient for businesses due to its native integrations with Microsoft's other services.

The costs of transitioning to Cloud-native detection and response are daunting if the processes is not managed in an optimized way. This is why partnering with an MDR provider like CyberProof that has deep knowledge and experience in Microsoft Security Migrations is key to a successful cloud migration.

CyberProof's security engineering and R&D teams work closely with Microsoft's product teams and have extensive experience deploying, managing, and optimizing the Microsoft security stack, having deployed some of the largest Microsoft-based Security Operations Infrastructures in the world. CyberProof has been recognized by Microsoft in these contexts:

- Microsoft Cloud Partner Program (MCPP)
- Microsoft Solutions Partner in the areas of Security and Specialist in Threat Protection and Cloud Security
- Azure Expert Managed Services Provider (MSP) status

CyberProof®
A UST Company

About CyberProof

CyberProof, a UST company, helps our clients transform their security to a cost-effective, cloud-native technology architecture. Our next-generation Managed Detection & Response (MDR) service is built to support large, complex enterprises by combining expert human and virtual analysts. Our services are enabled by our purpose-built platform, the CyberProof Defense Center – enabling us to be more agile, collaborate better, and deliver powerful analytics. Our integrated security services include Threat Intelligence, Threat Hunting, and Vulnerability Management. Our experts innovate to meet our clients' needs with custom use cases, integrations, and automations.

For more information, visit www.cyberproof.com.

Barcelona | California | London | Singapore | Tel Aviv | Trivandrum

cyberproof.com