



AZURE SENTINEL BEST PRACTICES

Strategies for success in data ingestion and
incident response

[Abstract](#)

This whitepaper details recommendations for configuring data sources for Microsoft Azure Sentinel and using Azure Sentinel during incident response and proactive threat hunting.

Azure Sentinel Best Practices

About this whitepaper

This whitepaper outlines best practice recommendations for configuring data sources for Microsoft Azure Sentinel, using Azure Sentinel during incident response, and proactively hunting for threats using Azure Sentinel.

Azure Sentinel makes it easy to collect security data across your entire hybrid organization from devices, users, apps, servers, and any cloud. Using the power of artificial intelligence, Sentinel ensures that real threats are identified quickly and unleashes you from the burden of traditional security incident and event management solutions (SIEMs) by automating setting up, maintaining, and scaling infrastructure.

Introduction

Overwhelming volumes of security data continue to prove a challenge for Security Operations Centers (SOCs) and the teams (SecOps) who operate them. Research shows that 76 percent of organizations reported increased security data¹. Combined with shortages of qualified professionals in the cybersecurity space (estimates suggest 3.5 million unfilled security jobs in 2021), this has resulted in 44 percent of an organization's security alerts never getting investigated. The issue is that successful security monitoring and response strategies *require* the collection and analysis of data at scale, and data fuels the machine learning models that power today's security solutions. This is a situation that will not improve in the near term.

For more than a decade SecOps has addressed collecting, analyzing, and responding to the deluge of alerts by deploying SIEMs to give their security analysts a "single pane of glass" to monitor. Results have been less than ideal. The scale, complexity, and rate of change in enterprise environments result in SIEM solutions that are unwieldy and expensive to build and run. They produce tremendous amounts of data which either overwhelm human analysts, or require locating and hiring data scientists to build, test, and deploy their own data analysis models. It's a lose-lose situation.

We created Microsoft Azure Sentinel to deal with these exact issues. Azure Sentinel is the first SIEM solution built into a major public cloud platform which delivers intelligent security analytics across enterprise environments and offers automatic scalability to

¹ *ESG: Security Analytics and Operations: Industry Trends in the Era of Cloud Computing 2019

meet changing needs. It features in-built artificial intelligence (AI) and machine learning (ML) and is built on top of Azure, which means it offers nearly limitless cloud speed and scale, has no infrastructure requirements, and can automate 80 percent of the most common tasks that SecOps analysts spend time on.

Since Azure Sentinel is designed to become a SOC's core technology, it is important to configure Azure Sentinel correctly, to connect the right sources of logs and data, and to ensure that your incident response processes are set before a breach occurs. This whitepaper will share Microsoft's best practices in these areas. For more information on Microsoft Azure Sentinel visit the product website at <https://aka.ms/azuresentinel>.

Enabling Azure Sentinel in an Azure tenant

To begin using Azure Sentinel, the service must be enabled in an Azure tenant, and then one or more data sources must be connected to the service. Azure Sentinel includes a number of pre-built data connectors for a broad range of Microsoft products and services and several built-in connectors for many additional non-Microsoft solutions. Additionally, Azure Sentinel can ingest data from Common Event Format (CEF), syslog, or REST-API sources by building new connectors.

There are three prerequisite steps for enabling Azure Sentinel:

- An active Azure subscription
- A Log Analytics workspace
- The correct permissions to deploy and use Azure Sentinel

For guidance on these steps visit <https://docs.microsoft.com/en-us/azure/sentinel/quickstart-onboard>.

Identifying data sources for Azure Sentinel

Today we no longer rely on signals from network security devices for the bulk of our security signals. The world of work has changed. No longer are our users, their devices, the data they access, and the applications and infrastructure they use to access that data under the direct control of organizations. They need access to sensitive data quickly and from any device. This puts a great deal of pressure on organizations. They still need to monitor network controls, but now they also must be much more reliant on identity signals to be sure the right users are accessing the right data on the right devices.

To help us make good security decisions, we recommend configuring Azure Sentinel to ingest security signal from a range of products, services, and locations.

Azure Sentinel can ingest data from a wide range of sources including Microsoft products and services, on-premises systems, leading SaaS applications, and non-Microsoft cloud environments including Amazon Web Services (AWS). Data sources can be connected to Azure Sentinel using one of these methods:

- Leverage the out-of-the-box data connectors included in Azure Sentinel to establish a connection in only a few clicks
- If a connector is not available, logs and alerts may be ingested using syslog, Common Event Format, or REST-API sources
- Some non-Microsoft solutions are connected via APIs provided by the connected data source

For more information on connecting data sources to Azure Sentinel see <https://docs.microsoft.com/en-us/azure/sentinel/connect-data-sources>.

Before connecting data sources to Azure Sentinel it is important to understand the potential costs of doing so. The following range of Microsoft generated logs and alerts can be ingested into both Azure Sentinel and Azure Monitor Log Analytics free of charge:

- Azure Activity Logs
- Office 365 Audit Logs including all SharePoint activity and Exchange admin activity
- Alerts from Microsoft Threat Protection products: Azure Security Center, Office 365 ATP, Azure ATP, Microsoft Defender ATP, Microsoft Cloud App Security, Azure Information Protection

Please note that Azure Active Directory (AAD) audit data is not free and is billed for ingestion into both Azure Sentinel, and Azure Monitor Log Analytics.

For full details of Azure Sentinel pricing including ingestion and storage costs, please visit <https://azure.microsoft.com/en-us/pricing/details/azure-sentinel/>.

To connect data sources to Azure Sentinel you will be working in the Data Connectors page inside Azure Sentinel:

The screenshot displays the Azure Sentinel 'Data connectors' page. At the top, it shows '39 Connectors', '18 Connected', and '1 Coming soon'. A search bar and filters for 'PROVIDERS: All' and 'DATA TYPES: All' are present. A list of connectors includes Amazon Web Services, Azure Active Directory, Azure Active Directory Identity Protection, Azure Activity, Azure Advanced Threat Protection (Preview), Azure Information Protection (Preview), Azure Security Center, Azure Security Center for IoT (Preview), and Barracuda CloudGen Firewall. The 'Amazon Web Services' connector is expanded, showing it is 'Connected' and 'Last Log Received' 20 minutes ago. It includes a description, last data received timestamp (05/05/20, 01:27 PM), related content (2 Workbooks, 2 Queries, 13 Analytic rules templates), and a line graph of data received from April 26 to May 3, totaling 1.02M.

Selecting which data sources to connect to your Azure Sentinel instance is an important choice. Microsoft recommends these sources as essential:

- Active Directory Federation Services (ADFS):** ADFS lets you securely share digital identity and entitlements rights across security and enterprise boundaries. Using a single sign-on within a single security or enterprise boundary to internet-facing applications, ADFS streamlines the user experience for customers, partners, and suppliers a streamlined user experience while they the web-based applications of an organization. A solution to allow Azure Sentinel to ingest ADFS sign-in logs is currently in private preview, but this document will be updated when it moves to public preview status.
- Azure Activity Directory (AD) activity logs:** To determine the “what, who, and when” for any action performed on resources in your subscription, we recommending setting Azure Sentinel to ingest AD activity logs like the [Azure AD audit logs activity report](#), the [Azure AD sign-in activity report](#), and [Azure activity logs](#). These logs can be connected with a single click using the pre-installed [Azure Activity connector in Azure Sentinel](#). There are separate instructions for ingesting Azure AD activity logs from [SumoLogic](#), [ArcSight](#), and [Log Analytics](#).
- Azure AD Identity Protection alerts:** [Azure AD Identity Protection](#) is a security control that lets organizations automate the detection and remediation of

identity-based risks, investigate risks using data in the portal, and export risk detection signals for further analysis and action. These alerts can be ingested using the pre-installed [Azure AD Identity Protection connector in Azure Sentinel](#).

- **Azure Advanced Threat Protection (ATP) alerts:** [Azure ATP](#) is a cloud-based security solution that leverages your on-premises Active Directory signals to identify, detect, and investigate advanced threats, compromised identities, and malicious insider actions directed at your organization. Azure ATP establishes a baseline of expected user behavior and will flag anomalous activities for your investigation. [Azure ATP can be connected to Azure Sentinel using the pre-installed connector \(currently in public preview\)](#).
- **Azure Information Protection (AIP) alerts:** [AIP](#) is a cloud-based solution that classifies and protects documents and emails by applying labels. These labels can be applied manually by users, automatically using admins-defined rules and conditions, or a combination of the two where users are given recommendations. AIP alerts can indicate suspicious activity such as unapproved attempts to access classified data, attempts to exfiltrate classified data, or attempts to reduce the classification labels on documents. [Azure Sentinel can ingest AIP alerts using the pre-installed connector \(currently in public preview\)](#).
- **Azure Key Vault logs:** [Azure Key Vault](#) is a tool for securely storing and accessing secrets, such as API keys, passwords, or certificates. [Azure Key Vault logs](#) can be [accessed and analyzed in Azure Monitor](#) and its logs and events from Azure Monitor can be ingested into Azure Sentinel.
- **Azure Security Center (ASC) alerts:** [ASC](#) provides security posture management for your cloud workloads, on-premises virtual machines, Linux and Windows servers, and Internet of Things solutions. Connecting ASC to Azure Sentinel allows it to ingest alerts, automatically create incidents, and trigger automated Azure Sentinel workbooks for investigation and remediation of the threat. [Azure Security Center alerts can be ingested by Azure Sentinel using the pre-installed connector](#).
- **Business critical applications:** If you have critical business applications that can export security alerts over syslog or CEF they can be ingested into Azure Sentinel. Instructions for doing this can be found [here](#).
- **Connect external solutions via agent:** Azure Sentinel can perform real-time log streaming of all other data sources using the Syslog protocol.
- Most appliances use the Syslog protocol to send event messages which include the log itself and data about the log. While the format of these logs may vary,

most appliances support the CEF based formatting for logs data. Solutions that can connect to Azure Sentinel using the agent include the following:

- Check Point
- Cisco ASA
- ExtraHop
- Reveal(x)
- F5
- Forcepoint products
- Fortinet
- Palo Alto Networks
- One Identity Safeguard
- Trend Micro Deep Security
- Zscaler
- Threat intelligence providers
- DNS machines
- Linux servers
- Non-Microsoft clouds such as Amazon Web Services
- Other solutions that support syslog or CEF.

If you are using these solutions in your environment, instructions for connecting them to Azure Sentinel can be found [here](#).

- **Connect external solutions via API:** External data sources can be connected to Azure Sentinel using APIs. Typically, this can be done using the APIs provided by the specific; these APIs connect to Azure Sentinel, gather specific data types, and send them to Azure Log Analytics. Appliances which can be connected to Azure Sentinel via API include:
 - Barracuda
 - Barracuda CloudGen Firewall
 - Citrix Analytics (Security)
 - F5 BIG-IP
 - Forcepoint DLP
 - Squadra Technologies secRMM
 - Symantec ICDX
 - Zimpreium

If you are using these solutions in your environment instructions for connecting them to Azure Sentinel can be found [here](#).

- **Microsoft Cloud App Security (MCAS) alerts:** [MCAS](#) is a cloud access security broker that supports various deployment modes including log collection, API connectors, and reverse proxy. It provides rich visibility, control over data travel, and sophisticated analytics to identify and combat cyberthreats across all your Microsoft and third-party cloud services. MCAS alerts can indicate unsanctioned attempts to download or export sensitive data, data access from non-corporate networks or risky IP addresses, and enforce your organization's policies to define user behavior in the cloud. [MCAS alerts can be brought into Azure Sentinel using the pre-installed connector.](#)
- **Microsoft Defender ATP alerts:** [Microsoft Defender ATP](#) is a platform designed to help enterprise networks prevent, detect, investigate, and respond to advanced threats. Microsoft Defender ATP is a full endpoint detection and response (EDR) available on a range of operating systems—Windows 10, macOS, Linux (in public preview), iOS and Android (both in private preview). The platform offers preventive protection, post-breach detection, and automated investigation and response. These alerts indicate attacks, compromises, and other threat indicators which can be automatically or manually remediated. [Microsoft Defender ATP alerts can be ingested into Azure Sentinel using the pre-installed connector.](#)
- **Microsoft Graph Security API:** Some security solutions can be connected to the [Microsoft Graph Security API](#) which can then send [threat indicators to Azure Sentinel](#).
- **Network Security Group (NSG) logs:** A typical NSG includes rules that allow or deny traffic to a virtual network subnet, network interface, or both. However, these logs can be large and noisy. We recommend you don't push all of the NSG logs into Azure Sentinel, just the alerts that flag NSG "deny rule" violations which indicate attempts to circumvent the security controls that we have implemented [These logs can be ingested into Azure Sentinel through Azure Monitor.](#)
- **Office 365 Advanced Threat Protection (ATP) alerts:** [Office 365 ATP](#) can be configured to protect cloud-based, on-premises, and hybrid deployments of Exchange, as it alerts whenever the service detects message-based threats targeting your users or organization such as malicious email attachments or phishing lures. At the time of writing there is no pre-installed connector for Office 365 ATP alerts but you can [ingest these alerts into Azure Sentinel using the Graph Security API.](#)
- **Office 365 Audit Logs:** Connecting your Office 365 Audit Logs to Azure Sentinel gives you visibility into a range of user activities, like various user, admin, system,

and policy actions and events. This data can be used to build dashboards, create custom alerts, and improve your investigation process. [Office 365 Audit Logs can be ingested into Azure Sentinel using the pre-installed connector.](#)

If you need to ingest additional logs Azure Monitor provides a mechanism:

- **Azure Monitor:** [Azure Monitor](#) lets you collect, analyze, and act on telemetry data from your Azure and on-premises environments. Azure Monitor helps you maximize performance and availability of your applications and proactively identify problems in seconds. [Azure Sentinel can ingest data from Azure Monitor using the REST-API.](#)

Using Azure Sentinel for incident response

Including Azure Sentinel in your incident response plan

Azure Sentinel can be extremely valuable in helping you build and run incident response plans. This section will cover high-level guidance on creating and implementing an incident response plan, and then using Azure Sentinel when responding to incidents.

If you do not have an already established security response process, we highly recommend you create one. [The Microsoft Security Response Center has published guidance on Microsoft best practices for building a security response process.](#)

Key concepts

Before we understand how Azure Sentinel can help with incident response it is important to understand three key concepts:

- **Alerts.** An alert is generated by a product or service when something suspicious happens. Examples of alerts include detection of suspicious files, detection of suspicious user activities, or attempted elevation of privilege.
- **Incidents.** An incident is an investigation package created by Azure Sentinel. A single incident can include evidence from multiple alerts. Azure Sentinel can automatically correlate the individual alerts and present a single package for an automated playbook to respond to or for a security analyst to investigate.
- **Entities.** An entity refers to a user, host, or IP address and act as entry points to explore all alerts and correlations associated with that entity. This can be extremely useful when investigating incidents. Instead of analyzing all the identity alerts, network alerts, and data access alerts in an investigation separately, entities allow us to view all of the alerts associated with a particular “thing” in our

environment. For instance, an entity would let you see all of the alerts associated with the CFO, their host machine, other hosts the CFO may have accessed, which IP addresses are associated with the CFO. This offers visibility into how seemingly unrelated events and alerts can be part of the same attack.

Investigating alerts

After you connect data sources to Azure Sentinel, you want to be notified when something suspicious happens. To help you to do this, Azure Sentinel lets you create advanced alert rules that generate incidents which you can assign and investigate.

Alerts triggered in Microsoft security solutions connected to Azure Sentinel, such as Microsoft Cloud App Security and Azure Advanced Threat Protection, do not automatically create incidents in Azure Sentinel. By default, when you connect a Microsoft solution to Azure Sentinel, any alert generated in that service will be stored as raw data in Azure Sentinel, in the Security Alert table in your Azure Sentinel workspace. You can then use that data like any other raw data you connect into Sentinel.

Using Microsoft Security incident creation analytic rules

Use the built-in rules available in Azure Sentinel to choose which connected Microsoft security solutions should create Azure Sentinel incidents automatically in real time. You can also edit the rules to define more specific options for filtering which of the alerts generated by the Microsoft security solution should create incidents in Azure Sentinel. For example, you can choose to create Azure Sentinel incidents automatically only from high-severity Azure Security Center alerts. [Full step-by-step instructions for creating incidents can be found here.](#)

Investigating incidents

Once you are generating incidents through Azure Sentinel the next step is to investigate those incidents. Azure Sentinel includes advanced investigation and analysis tools to help you understand what is happening and take remediation steps. The main techniques are [accessing the incidents in your environment through the Azure Sentinel Incidents page](#), [diving deeper using the investigation graph](#), and [responding to threats using automated playbooks](#).

Investigating incidents in multiple workspaces

Microsoft recommends using a single Log Analytics workspace to store and analyze data. However, managed security services providers (MSSPs) by definition will be working with multiple customers across multiple workspaces. Additionally, some organizations may have multiple workspaces in their environments either by design such as segmentation or through organic growth such as through mergers and

acquisitions. In these circumstances you can use the [Multiple Workspace View](#) to gain visibility across those workspaces. MSSPs working across multiple customer tenants can use [Azure Lighthouse](#) capabilities in their own tenant to manage Azure resources without connecting directly to the customer's tenant.

Threat hunting in Azure Sentinel

Azure Sentinel offers extensive proactive threat hunting capabilities for seeking out attacks before alerts have been raised. This can be a complex endeavor—the hunter is looking for new anomalies that were not detected by their organization's security products and solutions. To accomplish this, hunters need the ability to write, edit, and execute queries against the large volumes of data being logged and stored in their environment.

To aid this process, Azure Sentinel has built-in hunting queries and guidance to you ask the right questions about finding anomalies in that data. To get started with hunting in Azure Sentinel, including using the Hunting page; understanding the built-in Kusto query language and its operators; building, saving, and executing queries; and accessing the GitHub Azure Sentinel query repository see [here](#).


The screenshot displays the Azure Sentinel Hunting interface. The top navigation bar shows 'Home > Azure Sentinel workspaces > Azure Sentinel - Hunting'. The main header includes 'Azure Sentinel - Hunting' and 'Selected workspace: ADHunt0M0Worksp00014 - PREVIEW'. Below this, there are search and action buttons: 'Search (Ctrl-V)', '+ New Query', 'Bookmark Logs', 'Refresh', and 'Last 24 hours'. A summary row shows '20 Total Queries', '94 Total Results', '4 Total Bookmarks', and '0 My Bookmarks'. The 'Queries' section is active, displaying a table of hunting queries. The first query, 'Uncommon processes/files - bottom 5%', is selected, and its details are shown on the right. The query details include a description, provider (Microsoft), data source (SecurityEvent), and results (12). The 'Query Information' section shows the Kusto query:

```
let startdatetime = "2019-02-19T18:26:31.916Z";
let enddatetime = "2019-02-20T18:26:31.916Z";
let ProcessCreationEvents = (
    let processEvents = SecurityEvent
    | where TimeGenerated > start and TimeGenerated < end
);
```

 The 'Entities' section shows 'Timestamp' and 'Since'. The 'Tactics' section shows 'Execution', 'Initial Access', 'Persistence', and 'Privilege Escalation'. A 'Run Query' button is visible at the bottom right.

Many security organizations are interested in adopting machine learning as part of their threat hunting capability. However, finding security analysts with ML knowledge may prove difficult, and even if those analysts are lucky enough to get access to the correct resources, developing ML models from scratch is a steep hill to climb. Leveraging extensive knowledge and experience at Microsoft, Azure Sentinel provides a framework for you to grow your ML capability with us.

Azure Sentinel offers a spectrum of features allowing you to leverage the power of ML no matter where you are in terms of your ML investments.



Our built-in ML models offer the models we developed and have used for years. Using them only involves plugging your data into the model, then alerts are delivered to your case management page.

Azure Security also offers a build-your own ML platform for you to get all the flexibility to develop ML models for your unique business problem. We provide our algorithms and best practices in the form of templates to accelerate your ML projects.

Additional threat hunting techniques

In addition to the threat hunting techniques outlined above there are additional methods of conducting threat hunting in Azure Sentinel.

- **Hunting using notebooks.** The core of Azure Sentinel is the data store in Log Analytics. Here, high performance querying and dynamic scalability to massive volumes gives Azure Sentinel rich hunting opportunities. The Azure portal and all Azure Sentinel tools utilize a common API to access the data store. This API is also available to external tools such as Jupyter notebooks and Python. Jupyter notebooks provide full programmability and a huge collection of libraries for machine learning, visualization, and data analysis. This makes Jupyter notebooks a powerful hunting tool, and Azure Sentinel has integrated the Jupyter notebook experience into the Azure portal. To learn how to create, store, run, and share Jupyter notebooks with fellow hunters see [here](#).
- **Hunting using bookmarks.** Threat hunting typically requires reviewing mountains of log data looking for evidence of malicious behavior. During this process, investigators find events that they want to remember, revisit, and analyze as part of validating potential hypotheses and understanding the full story of a compromise. Hunting bookmarks in Azure Sentinel help you do this by preserving the queries you ran in *Azure Sentinel—Logs*, along with the query results that you deem relevant. You can also record your contextual observations and reference your findings by adding notes and tags. Bookmarked data is visible to you and your teammates for easy collaboration. You can revisit your bookmarked data at any time on the *Bookmarks* tab of the *Hunting* pane. You can use filtering and search options to quickly find specific data for your current investigation. Alternatively, you can view your bookmarked data directly in the *Hunting Bookmark* table in your *Log Analytics* workspace. To learn about hunting with notebooks in Azure Sentinel see [here](#).
- **Hunting using livestream (currently in public preview).** Hunting in Azure Sentinel using livestream lets you create interactive sessions where you can test newly created queries as events occur, get notifications from those sessions when matches occur, and launch investigations. Livestream sessions can be quickly created using any Log Analytics query. Livestream sessions give hunters the

ability to create queries and conduct them against live data in real time. To learn about hunting using livestream in Azure Sentinel see [here](#).

Conclusion

Traditional SIEMs have proven to be expensive to own and operate, often requiring you to commit upfront and incur high cost for infrastructure maintenance and data ingestion. Azure Sentinel provides you with SIEM-as-a-service and SOAR-as-a-service for the SOC: your birds-eye view across the enterprise; putting the cloud and large-scale intelligence from decades of Microsoft security experience to work. Following the best practices outlined within this white paper will help you eliminate security infrastructure setup and maintenance and provide you with scalability to meet your security needs—all while reducing costs and increasing visibility and control.