Microsoft Security

BlueVoyant

Microsoft Sentinel Deployment Best Practices

2nd Edition

Authors:

Adrian Grigorof, CISSP, CRISC, CCSK

Marius Mocanu, CISSP, SABSA, CISM, CEH

Jordan Shaw-Young, CISSP, CISM

Preface to the 2nd Edition

BlueVoyant and Microsoft Security collaborated to produce the first edition of *Azure Sentinel Deployment Best Practices* in July 2021 to provide enterprise and public sector organizations with a practical field guide to deploying Microsoft's then-new cloud SIEM platform. Since 2021 Microsoft Sentinel has seen rapid development, releasing many new and improved features, and has gained broad adoption from security teams around the world. The solution has even seen a name change over this time, rebranding from "Azure Sentinel" to "Microsoft Sentinel" to better reflect its capabilities as a full enterprise SIEM solution rather than a tool exclusively for Azure workloads.

Our initial guide filled a gap for security practitioners and leaders who needed a view of real-world considerations that come with deploying Microsoft Sentinel from an experienced security team. In the intervening months since its publication, both Microsoft and the global Partner community have produced a growing body of high-quality documentation on the deployment and migration of the Microsoft Sentinel platform. Thousands of global enterprises now have project experience in operationalizing and integrating the tool.

The 2nd edition of this guide will serve two purposes. First, we will continue to provide practical, experience-derived deployment recommendations covering the latest features and capabilities of the Microsoft Sentinel platform. Second, we will push further into more challenging examples and use cases that we have encountered in the field through our project work and how enterprise organizations were able to solve them.

Thank you for reading.

The BlueVoyant Team

Table of Contents

Introduction

The purpose of this whitepaper is to provide security organizations with a practical field guide to assist in developing a deployment strategy for Microsoft Sentinel. It will employ best practices to support a stable, cost-effective, and operationally effective implementation of Microsoft's cloud-native security information and event management (SIEM) platform. This document is written from a security practitioner perspective, based on experience deploying and managing Microsoft Sentinel in a wide range of organizations.

We intend for this guide to serve as a reference and planning document primarily for chief information security officers, security architects, and enterprise architecture and project management leaders. It defines adoption and migration strategies, budgeting, project planning, and resourcing requirements for a successful implementation of Microsoft Sentinel. It can be read as a companion document to other Microsoft Sentinel technical whitepapers, such as the [Microsoft Sentinel Technical Playbook for MSSPs.]

Microsoft Sentinel cloud-native SIEM architecture

Microsoft Sentinel is Microsoft's cloud-native SIEM solution and the first cloud-native SIEM from a major public cloud provider. Microsoft Sentinel is deployed in an organization's Azure tenant and accessed via the Microsoft Azure portal, ensuring alignment with pre-existing organizational access control policies.

Leveraging native integrations with Microsoft Defender tools and Azure services such as Log Analytics and Logic Apps for analysis and automation capabilities, Microsoft Sentinel allows organizations to ingest, correlate, and analyze security signals from across the enterprise.

The ability to leverage elastic compute and storage capabilities inherent in Azure for data-intensive applications such as SIEM, is a significant advantage over on-premises based log analysis solutions.

Additionally, Microsoft Sentinel can make use of infrastructure-as-a-service (IaaS) and platform-as-a-service (PaaS) capabilities available in Azure to deliver workflow automation, analysis, and long-term log retention that are typically only available as add-on services from other SIEM providers.

Microsoft Sentinel for Security Operations

In addition to architectural advantages like elastic scalability and a rapid software development lifecycle (SDLC) that a cloud-native SIEM platform provides to Security Architecture and Engineering teams, experience has shown us the advantages provided to Security Operations and Content Engineering teams can be even more impactful.

The programmability and API-first nature of Azure opens new capabilities that have traditionally been the exclusive domain of application development operations (DevOps) teams and makes these available to Security Operations. New threat detection content developed in KQL can be rapidly deployed at scale through a variety of methods, including tools such as Azure DevOps for internal content teams, or by leveraging Azure Lighthouse to enhance and extend content development through a partner or MSSP.

In many cases, we see security teams doing both concurrently – developing internal SecOps capabilities while partnering with an external organization that can provide fast and threat-centric detection content.

This guide will expand on Microsoft Sentinel best practices for both Security Architecture as well as Security Operations, based on our experience in over 400 enterprise deployments of the platform and as a Microsoft Verified Managed XDR Solution partner. Use cases and solutions described in this document are from real organizations, with all identifying details having been anonymized.

Case Studies

Where applicable, we have added anonymized case studies based on actual Microsoft Sentinel deployments to provide practical examples of situations and outcomes in enterprise environments where we have been engaged.

Case Study: Leading Technology US-based company

Company ABC Inc. has 3,000 employees and has been running Microsoft Sentinel for more than one year in their primary Azure tenant for both corporate and customer networks. Ingesting a variety of log sources such as Windows security events, firewall threat and traffic logs, VPN events, authentication logs, and several PaaS resources running in Azure and AWS cloud instances.

Initially, the average log ingestion volume was 110GB/day, using a 100GB/day Commitment Tier and 12 months for online data retention. ABC Inc. spend related to Microsoft Sentinel was approximately US$12,700/month in the US East Azure Region. Over 200 alert rules and 15 active playbooks were deployed in the Microsoft Sentinel instance.

Through acquisition, ABC Inc. acquired a company running an on-premises LogRhythm SIEM. As part of the post-acquisition integration activities, ABC Inc. planned to quickly migrate all log sources from the LogRhythm SIEM platform into the Microsoft Sentinel instance to reduce costs and avoid redundancies.

The onboarding activities, including the build of a local Syslog collector, reconfiguration of the MMA agents for Windows and Linux devices, development of 2 new custom data connectors, alert tuning, and use case development was completed within 3 weeks.

The log ingestion volume in the ABC Inc. Microsoft Sentinel instance following the cutover was 180GB/day, running 219 alert rules and 16 playbooks. Post migration, ABC Inc. the Azure bill for Microsoft Sentinel was US$20,900/month. Our experience has shown that integrating SIEM environments generally will result in only a small increase in detection use cases, provided an initial strategy is appropriately developed in the primary Microsoft Sentinel and followed through integration activities.

Core Microsoft Sentinel Solution Components

In this section, we provide guidance on the deployment of the core Microsoft Sentinel solution components to be deployed in your Azure subscription.

Azure Log Analytics Workspace

The first deployment prerequisite of Microsoft Sentinel is a Log Analytics workspace where all ingested data will be analyzed and initially stored. A Log Analytics workspace is created within a specific Azure region. It has a configurable log retention period, defining how long data will be stored within the Log Analytics workspace (database). The default is 30 days of hot storage, but this can be configured to as long as 730 days (2 years). When Microsoft Sentinel is enabled for a log analytics workspace, the free log storage is extended to 90 days. A log archiving option ("warm" storage) is available as well and allows the retention of logs for up to 7 years.

Log Analytics is a component of overall Microsoft Sentinel cost and is calculated based on the volume of ingested data and the data retention period. Special consideration should be paid to the extended retention period. Certain event tables may only contain system performance metrics or verbose logging of services, which may not be ideally suited for security analysis within a SIEM solution. Data unrelated to security monitoring may not be worth storing over a long period of time when balanced against ingestion costs. Conducting a thorough analysis of the incoming data and aligning it to organizational compliance policies will determine if raw data must be kept online in Log Analytics or if other storage options should be considered.

Microsoft Sentinel benefits from the inherent elastic storage capabilities of Azure Cloud. As such, it can dynamically scale on demand to meet even the most demanding data ingest requirements. For larger enterprises and organizations that see more than 500 GB/day, Microsoft offers an optional dedicated cluster for Microsoft Sentinel within Azure's infrastructure. This can improve search performance and, depending on your configuration of Microsoft Sentinel workspaces, can provide cost savings and efficiencies.

Case Study: Large Manufacturing Organization

ABC Inc. is currently using Microsoft Sentinel as their primary security monitoring tool. The current log ingestion volume is 450GB/day with an upward trend of +10% every 90 days. The company has a log retention policy for 2 years. Based on this volume of data the Azure monthly bill for Microsoft Sentinel would be around US$60,000, using the Log Analytics workspace as a hot storage.

The company has decided to use 3 months included storage available in Microsoft Sentinel and 21 months within the data archiving option at the cost of US$0.02/GB and using a data restore approach of 20GB/restore for 2 days as a restoration interval. Using this option, the company is able to save around US$20,000 in Azure costs per month.

Log retention options are configurable within Log Analytics at any time, however cost implications can be significant, and it is recommended that log storage requirements are documented at the outset of a project and initially configured within the SIEM solution.

For organizations that need to keep data available for longer than 90 days in a cost-effective storage repository while still being able to perform real-time queries or Kusto Query Language (KQL), there is the option to use Azure Data Explorer (ADX). It is a big data analytics platform that is highly optimized for all types of logs and telemetry data analytics. It provides low latency, high throughput ingestions with fast queries over extremely large volumes of data. It is feature-rich in time series analytics, log analytics, full-text search, advanced analytics visualization, scheduling, orchestration, automation, and many more native capabilities.

Learn more about Microsoft ADX here: https://docs.microsoft.com/en-us/azure/data-explorer/

There are a few initial best practices to follow when configuring Azure Log Analytics for use with Microsoft Sentinel:

- In multi-region architectures, deploy your Log Analytics workspace in an Azure region that will minimize the egress cost of data transfer between regions. In complex architectures with multiple Microsoft Sentinel instances, initial consideration should be paid to the region where most data is produced and consumed to avoid data export charges when providing Microsoft Sentinel with data from disparate Azure regions. In most cases, data export charges between Azure regions are usually lower than the price difference for Log Analytics between regions. Export charges between regions for Azure resources are only applicable to IaaS services (virtual machines [VMs]) and not to Azure PaaS services.
- For global organizations with a business presence in many countries, data sovereignty compliance may require the creation of multiple log analytics workspaces, each with its own settings for log retention and access to the stored data.
- Limit the number of Log Analytics workspaces, where possible. Understanding the relationship between security and operational data early in the project, and how each will be ingested, can save data ingestion charges at later dates.

- Implement a comprehensive role-based access control (RBAC) strategy for Log Analytics access early in the project.
- Configure Microsoft Sentinel analytic rules for monitoring various parameters related to data ingestion and costs. Often analytic rule requirements are built based purely on security operations needs; however, analytic rules are powerful and can be configured to perform monitoring on operational aspects of Microsoft Sentinel itself.
- Data requiring longer retention periods can be stored in alternative solutions, such as Azure Data Explorer (ADX) or Azure Blob Storage.
- Decide as early as possible on the data retention and restoration available options in Microsoft Sentinel, if the built-in Microsoft Sentinel data archiving solution is chosen. Search job limitations around the number of concurrent queries and time intervals may be an overly restrictive option for large enterprise organizations.

Azure Logic Apps

Combined with Microsoft Sentinel Automation Rules, Azure Logic Apps provides security orchestration and automated response (SOAR) capabilities in Microsoft Sentinel. Azure Logic Apps power “playbooks” and are, effectively, a sequence of procedures that can be run in response to a security alert. Playbooks can help automate and orchestrate response actions that would typically be undertaken by security analysts. These can be triggered manually or set to run automatically when specific alerts are triggered.

Azure Logic Apps is a great beneficiary of the capabilities of elastic compute and uses the power of the Azure Cloud platform to automatically scale and meet demand—you do not have to worry about the complexity of infrastructure capacity, hosting, maintenance, or availability for your workflows. It is highly likely that if an organization has workloads in Azure Cloud, Logic Apps are already used in automations for other services.

Azure Logic Apps comes with many different out-of-the-box connectors that enable organizations to easily create Microsoft Sentinel playbooks for automated workflows. Azure Logic Apps’ pricing structure is based on the number of transactions (or executions) and the type of connector.

As a Microsoft Azure Cloud service, Azure Logic Apps can run under a consumption-based pricing and metering model. This means that the fees are related only to how many workflow actions Azure Logic Apps execute.

The monthly price for deploying and using Logic Apps to orchestrate security event response is generally not a significant factor in the total cost of running Microsoft Sentinel. The versatility provides organizations with a wide range of options for reporting, alerting, and orchestration involving Microsoft Sentinel alerts. Other enterprise-grade connectors for non-security applications can come with higher cost price tags, so evaluation on the volume of playbook runs should be undertaken before deployment.

Case Study: Insurance Company

Insure Inc., an insurance provider with 4,000+ employees, is currently using Microsoft Sentinel for security monitoring and automated security incident response. The following Microsoft Sentinel playbooks were configured using Azure Logic Apps as part of Insure Inc. Sentinel alert rules response:

- **Azure Usage** playbook runs a daily query in the Log Analytics usage table. It sends an email notification to the service owner with aggregate daily Azure ingest costs per log source.
- **Incident Notification** playbook runs when a Microsoft Sentinel incident is created and automatically opens a ServiceNow ticket with incident details. The playbook is executed through an automation rule that allows a wide range of filtering criteria to limit the playbook run to specific types of incidents.
- **Data Enrichment** playbook runs when a Microsoft Sentinel incident is created, enriching users and entities based on data available from some Insure Inc. repositories. It sends email notifications to the SOC team when any suspicious activities are seen on a VIP user account.
- **Send DLP Notification** playbook runs when a Microsoft Sentinel incident is created for a USB storage device alert and adds DLP tags, and marks Microsoft Sentinel incident as critical.
- **Health Monitoring** playbook runs when a Microsoft Sentinel playbook fails and sends an email and SMS notifications to the company SOC team.

Running in the West U.S. Azure Region, in 2022, Insure Inc paid a total monthly aggregated fee for Azure Logic Apps of $112.00, which includes standard connector actions, built-in actions, and data retention fees. Relative to the value produced by automating response actions, the cost of building security workflow automations in Logic Apps is often negligible and a powerful way to drive business value from a Microsoft Sentinel deployment.

Data Sources

We regularly encounter a common misconception among security executives and practitioners that Microsoft Sentinel can only be used for Azure Cloud resources. In fact, Microsoft Sentinel is successfully used to ingest and correlate data from a wide range of log sources located in a variety of cloud platforms (Azure, Amazon Web Service (AWS), Google Cloud Platform (GCP) and Oracle Cloud Infrastructure (OCI)), on-premises network and compute infrastructure, 3rd party security tools, and software as a service (SaaS) applications. In addition to the growing set of out-of-the-box data connectors, the Microsoft Sentinel public community is regularly demonstrating new use cases and data connectors that expand the capabilities of the solution.

Case Study: Transportation company

Wheels Transportation is a transportation and logistics company that has been in business for over 100 years. At the time of initial engagement, Wheels did not have any workloads in Azure, including standard productivity tools such as Microsoft Office 365. Infrastructure was primarily on-premise, including servers and network infrastructure.

The security operations team at Wheels had been using an on-premises SIEM solution that required regular capital expenditure on storage and server hardware, which had to be aligned to forecasted events per second (EPS) ingestion requirements several years into the future. This EPS forecast was often missed, resulting in unbudgeted purchases of additional hardware and log ingestion licenses.

Wheels chose to deploy Microsoft Sentinel as its first Azure resource and primary SIEM solution because of the ability to operationalize log ingestion costs using a public cloud. Although Wheels did not have any existing Azure Cloud infrastructure, the flexibility of the Microsoft Sentinel SIEM solution was the right fit.

The wide variety of potential data types as log sources means that the consideration paid to each different data type is important at the outset of a Microsoft Sentinel project. Microsoft Sentinel includes hundreds of connectors out of the box, with the ability to create custom connectors to meet custom requirements. We have collected a summary table of some of the more common data collection methods, with experiential commentary relevant for deployment teams configuring new data ingest sources.

Log ingestion method	Typical log sources	Experience
Microsoft Sentinel built-in data connectors	Built-in Microsoft Sentinel data connectors for Microsoft native log sources, such as Azure Active Directory, Microsoft Defender, Office 365 and others.v	● Several of these log sources are free to ingest. Example: Microsoft Office 365, Azure subscriptions activity logs and alerts from Microsoft Defender products. ● In many cases, the configuration is simply a click on a "Connect" button. ● As of January 2023, there were 125 Microsoft Sentinel connectors available out-of-the-box, with 273 additional solutions available through Microsoft Sentinel Content Hub

Microsoft Monitoring Agent (MMA) and Azure Monitor Agent (AMA)	Windows and Linux machines deployed on-premises or in any other cloud environments. Microsoft Internet Information Server (IIS) Web Servers and Microsoft DNS logs can be collected via these agents. Any logs from other applications running on the same machine where MMA or AMA agents are running. MMA will be end-of-life as of August 2024. We recommend that any new Microsoft Sentinel deployments should use AMA for all endpoints. We also recommend that the endpoints be servers, not workstations. Workstation logs usually offer limited value or duplicate existing EDR signals.	● Windows events are typically very noisy, and the MMA/AMA filtering capabilities should be used to select the type of data that is sent to Microsoft Sentinel. AMA provides advanced filtering capabilities and can be used to specify individual Windows events that should be collected. For organizations that have strict log retention policies, data collection polices (DCRs) can be configured to split the logs and store them in alternate locations for a more cost effective log retention. ● Windows Performance Counters can be collected via this agent, too. Depending on the sample rate interval, this can increase the volume of collected events, increasing the overall Azure consumption.
Syslog log collector	Windows and Linux machines deployed on-premises or in any other cloud environments. Microsoft Internet Information Server (IIS) Web Servers and Microsoft DNS logs can be collected via these agents. Any logs from other applications running on the same machine where MMA or AMA agents are running. MMA will be end-of-life as of August 2024. We recommend that any new Microsoft Sentinel deployments should use AMA for all endpoints. We also recommend that the endpoints be servers, not workstations. Workstation logs usually offer limited value or duplicate existing EDR signals.	● Windows events are typically very noisy, and the MMA/AMA filtering capabilities should be used to select the type of data that is sent to Microsoft Sentinel. AMA provides advanced filtering capabilities and can be used to specify individual Windows events that should be collected. For organizations that have strict log retention policies, data collection polices (DCRs) can be configured to split the logs and store them in alternate locations for a more cost effective log retention. ● Windows Performance Counters can be collected via this agent, too. Depending on the sample rate interval, this can increase the volume of collected events, increasing the overall Azure consumption.

Syslog log collector	Firewalls, intrusion prevention systems (IPSs), L2/L3 network devices, and others.	● Syslog data is collected in Syslog format and may require creation of log parsers in Microsoft Sentinel. ● All Syslog messages are stored in a single Log Analytics table (Syslog table). ● As Microsoft Sentinel syslog collector runs on top of standard syslog servers such as rsyslog, their configuration can be used to filter the type of log data that is sent to Microsoft Sentinel. Microsoft Sentinel itself provides options to filter the log data based on Syslog facility and severity.
Common Event Format (CEF) log collector	● Firewalls, SD WAN, and various security solutions with the ability to log using the Common Event Format. ● Microsoft provides the installation scripts and documentation for a Linux agent that can be deployed on the same machine where a Syslog agent runs.	● CEF has a standard schema used by many security vendors, allowing interoperability among different platforms. ● Data is stored in the CommonSecurityLog table and typically does not require additional log parsing. ● Many platforms, like firewalls, allow customization of CEF templates, which is a great tool to optimize the volume of ingested logs at the source point.

Logic App playbooks	An Azure Logic Apps playbook can use various Logic Apps connectors and actions in order to achieve the ability to perform complex tasks (i.e. using a REST API call can be used to pull events from an application or tool). An HTTP connector typically retrieves the data through a REST API call while a “Send Data” connector is pushing the retrieved data to Log Analytics. This method is typically used for SaaS applications. Data is ingested in Microsoft Sentinel Log Analytics workspace and is placed in a custom table.	● Using remote application REST API calls, the data connectors can be set up to extract specific events only, which is an effective tool for log optimization. ● Data enrichment can be performed once the raw logs are retrieved from a log source. ● Log ingestion is based on a “pull,” within a predefined time interval (no real-time capabilities unless the log source itself can push new data when it becomes available). ● Relies on availability of Microsoft Sentinel playbooks; therefore, additional monitoring is required for playbook health and status. ● The customer has control over Log Analytics table schema definition.
REST API calls from an Azure function app	This method requires custom development using remote application REST APIs and Azure functions. Multiple tutorials available on building Azure function apps with a large number of programming languages supported (the most common being Python and PowerShell). While this type of data connector may require more advanced programming capabilities, it compensates through versatility. A large number of built-in Microsoft Sentinel connectors and solutions deploy Azure function app connectors.	● Customer does not need to run a separate machine/VM to host the connector code ● Logs are ingested in a Log Analytics custom table. ● Log optimization is dependent on remote application REST API though the function app may perform additional filtering and data enrichment. ● This is the Microsoft recommended method for log collection from custom log sources.

Codeless Connector Platform	The Codeless Connector Platform (CPP) allows the creation of log ingestion connectors for SaaS platforms. By defining the log ingestion parameters (REST API interface, authentication, pagination, polling frequency, etc) in a JSON file, it provides options to integrate with any log sources that provides an API for log retrieval. One can deploy these connectors as ARM templates or through Azure APIs.	● Typical challenges are around obtaining the right credentials ● Some SaaS logging APIs are not mature enough and may cause problems. The API should be tested prior to any efforts to configured it via CPP.
Logstash collector	Firewalls, IPS, network devices, and others. A Logstash collector needs to be deployed on-premises or in a cloud environment on a VM.	● Data enrichment (for example geo-location) can be done on collection point. ● The raw logs can be parsed using various Logstash filters such as Grok patterns ● This allows log optimization, by collecting only required log fields. ● Once ingested in Microsoft Sentinel Log Analytics workspace, data will be stored in a custom table. Special Data Collection Rules (DCRs) can be configured in order to send the logs to the standard Microsoft Sentinel tables such as Syslog and CommonSecurityLog. ● Many parsers and data enrichment tools are available in the open-source Elasticsearch-Logstash-Kibana (ELK) community. ● Microsoft maintains and supports the Log Analytics Logstash output plugin.

Azure diagnostics	Azure PaaS resources. Not always considered a separate log ingestion method, collecting events via Azure Diagnostics is applicable to Azure PaaS resources only. Turning Azure Diagnostics for Azure PaaS resources, the audit logs and metrics events can be collected and stored in an Azure Diagnostics Log Analytics table.	● Data ingested via Azure Diagnostics has to be reviewed for each log source type as it is very large and not always relevant from a security perspective. ● No data optimization can be done via this method. ● Microsoft provides in Microsoft Sentinel a set of standard data connectors (e.g., Azure Key Vault, Azure Kubernetes) for PaaS resources and in most cases, these are instructions on configuring DCRs (Data Collection Rules) that can be used to configure the resources automatically without a need to manually configure each of them. ● If there are no specific compliance requirements for log retention, using Microsoft Defender for Cloud for monitoring and protecting Azure PaaS resources, in general, is a more cost- effective solution for this situation.
File-based ingestion	Used for ingestion of data from files located on the same machines where MMA or AMA agents are running. Examples of such log sources are Microsoft DHCP Server or on-premises Exchange servers.	● Logs are collected in a Log Analytics custom table (_CL).

Amazon Web Services	AWS PaaS log sources (CloudTrail, GuardDuty, VPC Flow, CloudWatch, etc.).	● Microsoft provides an AWS CloudTrail Microsoft Sentinel data connector and an AWS S3 data connector out of the box. The AWS S3 connector allows the collection of logs from AWS CloudTrail, AWS GuardDuty and AWS VPC Flow logs via an AWS S3 bucket. ● Collecting events from other AWS resources, a REST API function can be developed. ● Data is placed in dedicated AWS tables such as AWSCloudTrail, AWSGuardDuty and AWSVPCFlowLogs, and it is already parsed.
Observability pipelines	Applicable to organizations with complex logging data routing requirements and needs to centralize log collection across a large number of locations. Products such as Cribl, can receive, filter, enrich and redirect the log data to various log analytics or storage solutions. They stream data to Microsoft Sentinel through the Azure Monitor HTTP Data Collector REST API.	● In some cases, observability pipeline solutions may provide additional capabilities (log data routing, advanced filtering, centralized monitoring) that can complement Microsoft Sentinel deployments.
Threat intelligence platforms: Structured Threat Information Expression (STIX) and Trusted Automated eXchange of Indicator Information (TAXII)	Available from external Threat Intelligence Platforms and from STIX/TAXII feeds, both open source and premium, these can enrich Microsoft Sentinel with additional indicators of compromise (IoC) such as known malicious internet protocols (IPs) or domain name service (DNS) names	● Employ either the Microsoft Sentinel Threat Intelligence Platforms data connector or Microsoft Sentinel TAXII data connector. ● Microsoft Threat Intelligence and several 3rd party solutions are available in the Microsoft Sentinel Content Hub and it links to a wide range of TI sources.

Microsoft provides a set of guidelines for vendors and Microsoft Sentinel community for development of new data connectors.

Microsoft Sentinel uses Azure Log Analytics (itself a component of Azure Monitor) as the backend for the log storage and querying capabilities through KQL. A wealth of information is available from various log sources stored in Log Analytics as tables. Different resources can create a variety of tables (see Azure Monitor Logs table reference for details) default tables, though not always populated with data while others created as specific Microsoft Sentinel data connectors are enabled and configured. Another range of tables, not covered in the list following, are represented by custom logs that can be used to ingest logs from custom applications that fall outside the scope of standard SIEM log sources.

Below is a list of the most common tables used in Microsoft Sentinel (the full list is over 500 tables with active logging), with the "Billable" column indicating if the logs ingested are subject to the Microsoft Sentinel ingestion price or are free of charge.

Table name	Description	Log sources	Relevant data	Billable
AuditLogs	Azure Active Directory (Azure AD) activities audit such as creation or modification of users, groups, applications	Azure AD	Account, location, activity	Yes
AWSCloudTrail	Audit of account activity across an AWS infrastructure	AWS CloudTrail	Account, location, activity	Yes
AWSVPCFlow	Information on network traffic through AWS VPCs	AWS VPC Flow	Source, destination, protocol, action	Yes
AWSGuardDuty	Monitoring of AWS for potential malicious activities	AWS GuardDuty	Threat, resource	Yes
AzureActivity	Azure activity such as creation/ modification/ deletion of Azure resources, policy updates	Azure Subscriptions	Account, activity	No
AzureDiagnostics	Diagnostic logs for Azure resources	Azure IaaS/PaaS resources	PaaS diagnostic data	Yes
AzureMetrics	Metrics recorded by various Azure resources	Azure IaaS/PaaS resources	PaaS metrics	Yes

CommonSecurityLog	Logs from security devices logging via syslog using CEF	Network and security devices	Source, destination, protocol, action	Yes
DeviceEvents DeviceFileCertificateInfo DeviceFileEvents DeviceImageLoadEvents DeviceInfo DeviceLogonEvents DeviceNetworkEvents DeviceNetworkInfo DeviceProcessEvents DeviceRegistryEvents	Detailed activity from devices with the Defender for Endpoint installed	Microsoft Defender for Endpoint	Account, file, registry and network activities	Yes
DnsEvents	Microsoft DNS events (registrations, configuration changes). Note: DNS queries outside the authoritative zone are not recorded	Microsoft DNS Server	DNS registrations, failures, queries	Yes
DnsInventory	Logs of DNS records created on the DNS zone	Microsoft DNS Server	DNS records	Yes
Dynamics365Activity	Activity from users in Microsoft Dynamics CRM platform	Microsoft Dynamics	User activities	Yes
EmailAttachmentInfo EmailEvents EmailPostDeliveryEvents EmailUrlInfo	Email related activity	Microsoft Defender for Office 365	Various email-related stats	Yes
Event	Windows event log entries (excluding Security event log)	Windows events	System and application-specific events	Yes

Heartbeat	MMA/AMA heartbeats	MMA and AMA agents	MMA/AMA health	No
McasShadowItReporting	Microsoft Defender for Cloud Apps IT information: records of access to applications typically used in ”shadow IT” (file sharing, meetings)	MMA and AMA agents	MMA/AMA health	No
NetworkMonitoring	Azure network performance monitoring	Azure network monitoring solution	Stats on network performance	No
OfficeActivity	Office 365 activity: Exchange, SharePoint, Teams, data loss prevention (DLP), OneDrive	Office 365	Office 365 user and admin activities	No
Operation	Records related the functionality of monitoring agent logs (data collection, availability, issues)	MMA/AMA	Status of MMA/AMA agents	Yes
Perf	Windows and Linux performance counters collected by MMA and AMA	Windows and Linux performance counters	Performance counter (i.e. CPU%, hard disk $)	Yes
ProtectionStatus	Microsoft Defender for Cloud records related to the status of endpoint protection solution on monitored endpoints	Microsoft Defender for Cloud	Status of endpoint protection	Yes
SecurityAlert	Alert details (Microsoft Sentinel, Microsoft Defender for Cloud, Microsoft Defender for Endpoint, Active Directory Identity Protection	Microsoft Sentinel, Microsoft Defender solutions	Alert metadata (as captured by the triggered alert)	No

SecurityBaseline	Microsoft Defender for Cloud records related status of monitored endpoints versus configured policies for security baseline (e.g., levels of patching)	Microsoft Defender for Cloud	Status of updates versus security baseline	Yes
SecurityBaseline- Summary	Microsoft Defender for Cloud records with statistics for the monitored endpoints related to compliance with configured policies	Microsoft Defender for Cloud	Policy compliance stats	Yes
SecurityDetection	Microsoft Defender for Cloud Apps logs for potential security issues detected on the monitored endpoints	Microsoft Defender for Endpoint	Alert name, alert type, severity, description, remediation suggestions	Yes
SecurityEvent	Window Security event logs entries	Windows endpoints	Account, domain, activity, process	Yes
SecurityIncident	High-level details of security incidents created in Microsoft Sentinel or other Microsoft 365 security tools	Microsoft Sentinel, Microsoft 365 Security tools (if "create incidents" is selected)	Incident description, severity, source app	No
SentinelHealth	Stats on Microsoft Sentinel functionality	Microsoft Sentinel	Data connectors health, playbook runs	No
SigninLogs	Azure AD interactive sign-ins	Azure Active Directory	Account, source, location, application, device details	Yes

StorageBlobLogs StorageFileLogs StorageQueueLogs StorageTableLogs	Records on Azure storage activity	Azure storage	Resource, activity, IP address, agent, bytes, uri, status, protocol	Yes
Syslog	Logs from syslog devices (non CEF)	Syslog devices (network/ security)	Raw message, severity, syslog facility	Yes
ThreatIntelligenceIndicator	Indicator of compromise ingested via Microsoft Graph Security REST API	Microsoft Graph	IP, domain, URL, hash	Yes
Update	Microsoft Defender for Cloud missing/required updates (Windows, Linux)	Microsoft Defender for Cloud	Computer, update	Yes
UpdateSummary	Microsoft Defender for Cloud records with the status of current updates for the monitored endpoints	Microsoft Defender for Cloud	Computer, update	Yes
W3CIISLog	Microsoft IIS access logs	Microsoft IIS	Source, destination, URL, protocol, status, bytes	Yes
WindowsEvent	Windows event logs	Windows events collected through WEF	Account logins/ logouts, user creation, group activities, processes, etc.	Yes
WindowsFirewall	Microsoft Windows Firewall log entries (firewall running on endpoints)	Microsoft Firewall	Traffic allowed and traffic dropped on endpoints	Yes

Project Resourcing

With an initial view of the key components of Microsoft Sentinel, including solution components and data sources, we will provide an overview of recommended approaches to deploying your new Microsoft Sentinel environment.

Project Planning

Duration and complexity of a Microsoft Sentinel deployment project will vary depending on a variety of factors. Some key variables that should be captured at the project planning stage, which will affect project durations, include:

- Access to log sources, system, and data owners
- Types of log sources (i.e., standard data connectors versus custom development)
- Complexity of Azure architecture (e.g., multiple tenants, cross-tenant searching)
- Requirement for custom SOAR automation playbooks, and interaction with other connected systems
- Azure cost assessment and optimization strategy
- Customer change management processes

Key roles suggested for a successful Microsoft Sentinel deployment include:

1. Project manager
2. Security architect
3. Cloud engineer
4. Engineering–systems owner
5. SIEM Engineer
6. Network Engineer
7. Business analyst
8. SOC Analyst
9. Developer (languages vary, but C#, Java and Python are often beneficial)
10. Compliance manager

Project Manager

Experienced project management staff with Project Management Professional (PMP) and Information Technology Infrastructure Library (ITIL) backgrounds are recommended, as stakeholder management requirements can be quite broad. Microsoft Sentinel projects will involve input and work effort from teams supporting both cloud and on-premises infrastructure, end-user-facing services such as SaaS applications and workstations, as well as mission-critical server infrastructure.

Cost impacts from each of the various log sources must be quantified prior to the project, with variances tracked as new log ingestion begins. Change management from existing security tools is a critical factor to ensure business continuity and cyber defenses are not compromised due to poorly planned or executed changes.

Security Architect

Initial scope questions concerning what log sources are to be ingested and the specific purpose for ingestion are important to analyze at the early stages of a Microsoft Sentinel project, taking a clear risk-based approach. There are numerous methods to gain security visibility to assets in the organization's information technology environment, however log ingestion from these sources must always be accompanied by analysis of the cost impact of ingestion and analysis of the data.

Design of the data collection methods, including the Syslog collectors sizing, location in the company network and network-level design such as load balancing, encryption and routing, is also the security architect role in Microsoft Sentinel deployments.

The Microsoft Sentinel environment will contain highly sensitive data, and appropriate role-based access control must be applied to the Azure resources in scope of the project. The security architect will have responsibility for the security design of the Microsoft Sentinel solution.

Cloud Engineer

Microsoft Sentinel will likely be one of many services running in your organization's Azure tenant, and determining resiliency requirements, Azure regions, data residency, and required tagging or templates applicable to Microsoft Sentinel will be the domain of the organization's Azure Cloud engineer/ administrator. The cloud engineer may also be responsible for deploying various Azure policies that may help the configuration of logging for multiple Azure resources.

Engineering – Systems Owner

Configuring log sources to send data to Microsoft Sentinel is often one of the more time-consuming activities in a Microsoft Sentinel deployment, particularly in complex organizational structures. Log structuring and format will vary from source to source, and organizational owners of assets such as SaaS applications, workstations, servers, cloud endpoints, and security infrastructure are often widely distributed.

Subject matter experts (SMEs) and asset owners with the administrative ability to provide samples of logs and configure log-forwarding parameters on each asset will be required to dedicate effort to working with the project team to ensure data is sent to Microsoft Sentinel.

For most data connectors available in Microsoft Sentinel, the prerequisite permissions are specified in the connector details. For example, for the Azure Active Directory data connector, these are Global Administrator or Security Administrator for the Azure tenant, log analytics workspace read/write and AAD diagnostics read/write.

Engineering–SIEM

The SIEM engineer(s) are responsible for configuring Microsoft Sentinel, including Log Analytics, Logic Apps, workbooks, and playbooks. Working with the security architect, systems owner, and project manager, the SIEM engineer will be responsible for the following high-level tasks:

- Initial configuration of the Azure tenant, including provisioning required resources, assigning access roles, and configuring workspace parameters such as log retention, resource tagging, and blueprints.
- Deployment and configuration of syslog/CEF log collection agents in appropriate locations to collect logs from on-premises devices. This step will also include provisioning appropriate servers to run syslog or other log collection solutions. In most cases, this will be a joint effort between the system engineer and the SIEM engineer, involving configuration and potential troubleshooting.
- Working with system owners to enable log forwarding and configuring any required parsing of log data in Log Analytics.
- Working with security operations to create and deploy KQL analytic rules to provide detections for SOC/computer security incident response team (CSIRT) use.
- Tuning of alert rule parameters, including thresholds, detection logic, and assigned criticality levels to minimize false positives and appropriately identify potential attacker behavior.
- Working with project management for stakeholder engagement, creating workbooks for data visualization, and dashboarding of Microsoft Sentinel content. Security operations is likely to be the primary consumer of workbooks; however, data visualizations in Microsoft Sentinel may be created and customized for a wide audience of stakeholders; therefore, appropriate requirements gathering through project governance is advised. The security architect may provide valuable input on the KPIs to be captured in dashboards.
- Creating automated workflows using Azure Logic Apps is recommended at the project phase. Working with security operations to document response workflows for various incident types and provisioning playbooks to automate response actions per incident type are effective ways to provide immediate value from the Microsoft Sentinel implementation.
- Working with systems owners for other IT systems, such as IT service management (ITSM) and helpdesk tools to build integrated ticket workflows, are recommended at the project phase. Microsoft Sentinel has capabilities to integrate with platforms such as ServiceNow or other ITSM API-enabled tooling to provide workflow automation for incident handling.

It is highly recommended to have the SIEM engineers involved in Microsoft Sentinel deployments to be trained as much in advance as possible, ideally along with a good introduction to related products such as the Microsoft Defender family of products.

Microsoft has developed online training courses for an introduction to Microsoft Sentinel platform. This provides a solid start for becoming familiar with the platform. It is recommended especially for SOC Engineers or Analysts to pursue the training and certification for Microsoft Security Operation Analyst. A good certification to obtain is SC-200: Microsoft Security Operations Analyst as it covers Microsoft Defender and Microsoft Sentinel. Another good training resource specifically for KQL would be following online course from Pluralsight.

Network Engineer

Network engineering resources will be required on demand to apply changes to firewalls or network infrastructure to facilitate log forwarding from data sources to Azure.

Business Analyst

Capturing and evaluating the budget and resource impact of Azure data ingestion and various data source ingestions are important aspects of a Microsoft Sentinel project. As a cloud-native SIEM, organizations are shifting costs from capital expenditure to operational expenditure, and cost forecasting for the Microsoft Sentinel solution is recommended at the project stage.

As part of the Microsoft Sentinel project, the business analyst (BA) should be able to provide an Azure cost analysis of each technical requirement. In conjunction with the SIEM engineer, the BA should model this expected Azure cost impact over time as changes to the IT environment are seen. An effective risk-based security program will be able to quantify the risk mitigation effects of security controls as related to the mitigation cost for a specific control.

Security Operations (SOC Analyst)

Security operations stakeholders in the Microsoft Sentinel project are primarily assigned to document the detection, alerting, and threat-hunting requirements of the solution. While the security architect and SIEM engineer are able to provide access to security-relevant data and present this back to security operations in the form of alerts, incidents, dashboards, or reports, it is ultimately the responsibility of security operations as the end consumer of the service to articulate the build requirements.

Developer

Developer resources are often the most overlooked requirements for a Microsoft Sentinel project. Programming languages such as C# and Python and developer effort are often required to obtain data from log sources such as some SaaS applications and can be leveraged to great effect by Azure functions. Developers can also be leveraged to build automation playbooks using Logic Apps, which can make use of a wide range of code to automate security operations tasks.

Compliance Manager

If your organization has legal, regulatory, or industry-specific compliance requirements that will need to be satisfied by Microsoft Sentinel, the interaction between the core Microsoft Sentinel team and compliance manager is mandatory.

Decisions such as log retention period, custom workbooks, and compliance reporting mandates are overseen by this resource.

Benchmark Project Effort and Duration

Below are high-level benchmarks for typical effort per resource type for a sample 5,000-employee organization with typical on-premises and cloud/SaaS log sources. Actual effort is highly variable depending on organization-specific factors and, in this example, we estimated the project duration to 6 calendar weeks.

Resource type/ function	Benchmark effort (days)	Key tasks
Project manager	5	● Project planning ● Stakeholder engagement ● Resource planning ● Change management ● Project governance
Security architect	1.5	● Data ingestion strategy and methods ● RBAC controls ● Compliance requirements ● Access control ● Identify existing security controls ● Identify log sources to be ingested into Microsoft Sentinel ● Provide expertise in use-case creation Identify gaps in data ingestion ● Provide consulting on governance/ risk/ compliance
Azure Cloud engineer	2.5	● Managing administrative permissions/RBAC in Azure tenant ● Service resiliency ● Provision Azure resources ● Configure Azure AD service accounts ● Configure Azure AD groups and assign membership

Engineering– systems owner	0.2 (per log source)	• Configuring log forwarding on assets required to send data to Microsoft Sentinel • IT sysadmin: deploying monitoring agents on endpoints • Deploy on-premises Microsoft Sentinel log collector • Provide expertise around application logging capabilities and potential use cases • Configure application logging • Assist in validation of alert rules and SOAR playbooks
Engineering–SIEM	10	• Developing KQL detection rules Developing workbooks for data visualization • Creating Azure functions for data retrieval, moves • Developing KQL detection rules • Developing workbooks for data visualization • Creating Azure functions for data retrieval, moves
Network engineer	0.5	• Assist with required connectivity between Microsoft Sentinel and logging sources
Business analyst	0.2	• Azure cost analysis • Documentation of business workflows to be automated
Developer	0.5 per log source	• Development of new custom data connectors to ingest data to Microsoft Sentinel
SOC Analyst	10	• Documentation of detection use cases for detection rules • Documentation of detection parameters (e.g., alert thresholds, exclusions, threat intelligence sources)
Compliance manager	0.1	• Provide recommendations on the compliance requirements applicable to Microsoft Sentinel SIEM • Review and provide feedback on Microsoft Sentinel components designed and built for compliance purposes

Case Study: International Consumer Packaged Goods Manufacturer

ABC Corporation, a large international consumer packaged goods manufacturer with offices on three continents and more than 75,000 employees, decided to centralize all security monitoring to the Microsoft Sentinel platform. ABC Corporation had a variety of tools and applications deployed in three regional datacenters, as well as in AWS, Google Cloud, and Azure. In addition, the organization was using a large number of SaaS applications across a variety of business units.

During initial planning, the ABC Corporation security team created a list of 49 log source types to be included in the scope of work for the Microsoft Sentinel implementation. Seven log source types required development work for the creation of custom data connectors.

The Microsoft Sentinel deployment project was led by a project manager, a security architect, and a Microsoft Sentinel engineer. The team obtained executive support and guidance from the ABC Corporation Chief Information Security Officer (CISO). During each phase of the project, the implementation team interacted with various teams and resources from different geographical regions and business units, with the project completed in 4 months.

The following diagram presents the milestones and deliverables completed in each project phase:

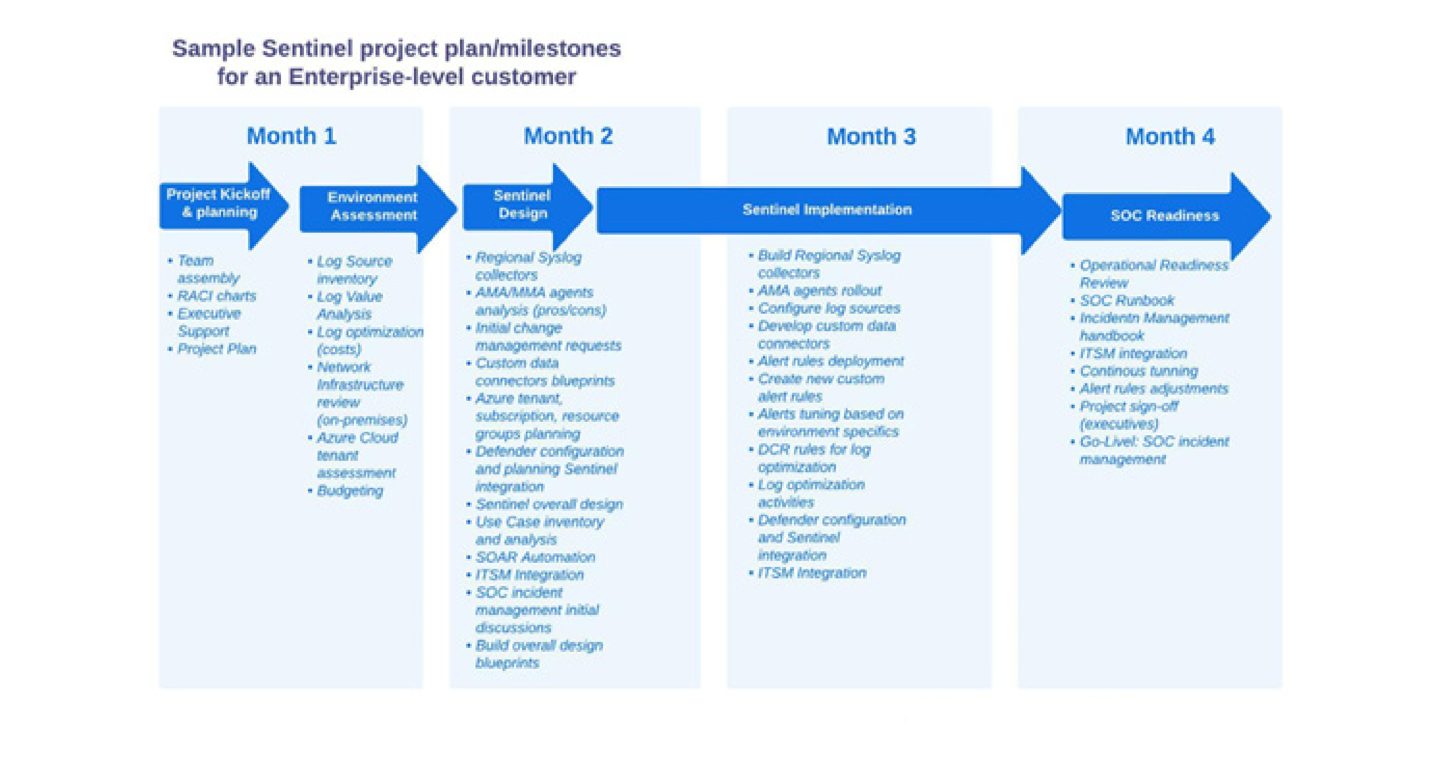

Sample Sentinel project plan/milestones for an Enterprise-level customer

Month 1

Project Kickoff & planning

- Team assembly
- RACI charts
- Executive Support
- Project Plan

Environment Assessment

- Log Source inventory
- Log Value Analysis
- Log optimization (costs)
- Network Infrastructure review (on-premises)
- Azure Cloud tenant assessment
- Budgeting

Month 2

Sentinel Design

- Regional Syslog collectors
- AMA/MMA agents analysis (pros/cons)
- Initial change management requests
- Custom data connectors blueprints
- Azure tenant, subscription, resource groups planning
- Defender configuration and planning Sentinel integration
- Sentinel overall design
- Use Case inventory and analysis
- SOAR Automation
- ITSM Integration
- SOC incident management initial discussions
- Build overall design blueprints

Month 3

Sentinel Implementation

- Build Regional Syslog collectors
- AMA agents rollout
- Configure log sources
- Develop custom data connectors
- Alert rules deployment
- Create new custom alert rules
- Alerts tuning based on environment specifics
- DCR rules for log optimization
- Log optimization activities
- Defender configuration and Sentinel integration
- ITSM Integration

Month 4

SOC Readiness

- Operational Readiness Review
- SOC Runbook
- Incidentn Management handbook
- ITSM integration
- Continous tunning
- Alert rules adjustments
- Project sign-off (executives)
- Go-Live: SOC incident management

At the end of the implementation, the Microsoft Sentinel environment included a list of 38 unique log source types, including the seven new custom data connectors for ABC Corporation SaaS applications. In total were deployed a list of 257 alert rules, including some custom alert rules and eight playbooks. The ongoing log ingestion volume in day-2 was 720GB/day with 500GB/day Commitment Tier on both Log Analytics and Microsoft Sentinel.

Design Planning

Architecture Planning and Considerations

The following section provides key factors affecting the initial architecture for deployments of new Microsoft Sentinel instances or the migration from existing SIEM platforms.

Data residency requirements

There are over 60 Microsoft Azure regions spanning 140 countries, with over 30 of them supporting Log Analytics workspaces. The geographical regions include North America, South America, Europe, Middle East, South-East Asia, and Africa. While additional support is added on a regular basis, Microsoft Sentinel is not available in all regions.

Microsoft Sentinel is also available in Azure Government cloud and in special Azure cloud instances such as Azure China. Some of the data connectors in these environments may have certain restrictions, so each must be analyzed case by case.

Additional Resources:

Quickstart: Onboard in Microsoft Sentinel | Microsoft Docs for the current list of Azure regions supported by Microsoft Sentinel.

Support for data types in Microsoft Sentinel across different clouds | Microsoft Learn

Depending on the type of business and customer residency, organizations may have compliance restrictions related to the logged data. Compliance regulations are not always clearly defined with respect to logging requirements and subject to change over time. Organizations may choose to use a local region to avoid complications due to changes in legislation or auditing processes.

The selection of the region also carries implications for Microsoft Sentinel and Log Analytics costs as well as the availability of resources for the specific region. Regions such as the U.S. can offer a significant cost advantage versus other regions. For example, East U.S. offers a 17% discount compared with Canada Central.

Depending on the volume of log ingestion, the discount can be significant. Therefore, it is recommended that the project team obtains organizational requirements relating to data residency prior to deploying.

Case Study: International Grocery Chain in South America.

Company Grocery Inc., a major chain of grocery stores in Chile with stores in Columbia, Venezuela, Argentina, and Brazil, decided to use Microsoft Sentinel SIEM as their core security analytics platform across the entire network infrastructure, mainly logging on-premises resources. Having stores and datacenters in five different countries but not having any corporate data residency or compliance requirements to force the log analytic data in South America, Grocery Inc. chose US East Azure Region.

This decision saved Grocery Inc. approximately 50% in Azure costs for both Microsoft Sentinel and Azure Log Analytics when compared to resources deployed in a local Azure region.

Number of Azure AD Tenants

An Azure AD tenant provides identity and access management (IAM) capabilities for applications and resources used within an organization. An identity is a directory object that can be authenticated and authorized for access to a resource. Identity objects exist for human identities (e.g., employees) and non-human identities (e.g., computing devices, applications, and service principals).

While most organizations have a single Azure AD tenant, some may have one or more through mergers and acquisitions or the need to segregate environments such as corporate versus production infrastructure.

Each Azure tenant requires a dedicated Microsoft Sentinel instance, as some of the data connectors only work within the current Azure tenant. For example, the Azure AD and Office 365 logs can only be connected within the local tenant.

Solutions based on custom connectors and Azure REST APIs can be developed to aggregate remote tenant logs into a single Microsoft Sentinel instance, but such solutions can introduce an unnecessary level of complexity and increased maintenance that unless are a must, should be avoided.

Azure Lighthouse can provide cross-tenant management experience for unified platform tooling, management at scale, and increased visibility.

Additional resources on Microsoft Sentinel integration with Azure Lighthouse:

Microsoft Sentinel and Azure Lighthouse (microsoft.com)

Build a scalable security practice with Azure Lighthouse and Microsoft Sentinel

Multi-tenant access for Managed Security service providers - Microsoft Tech Community

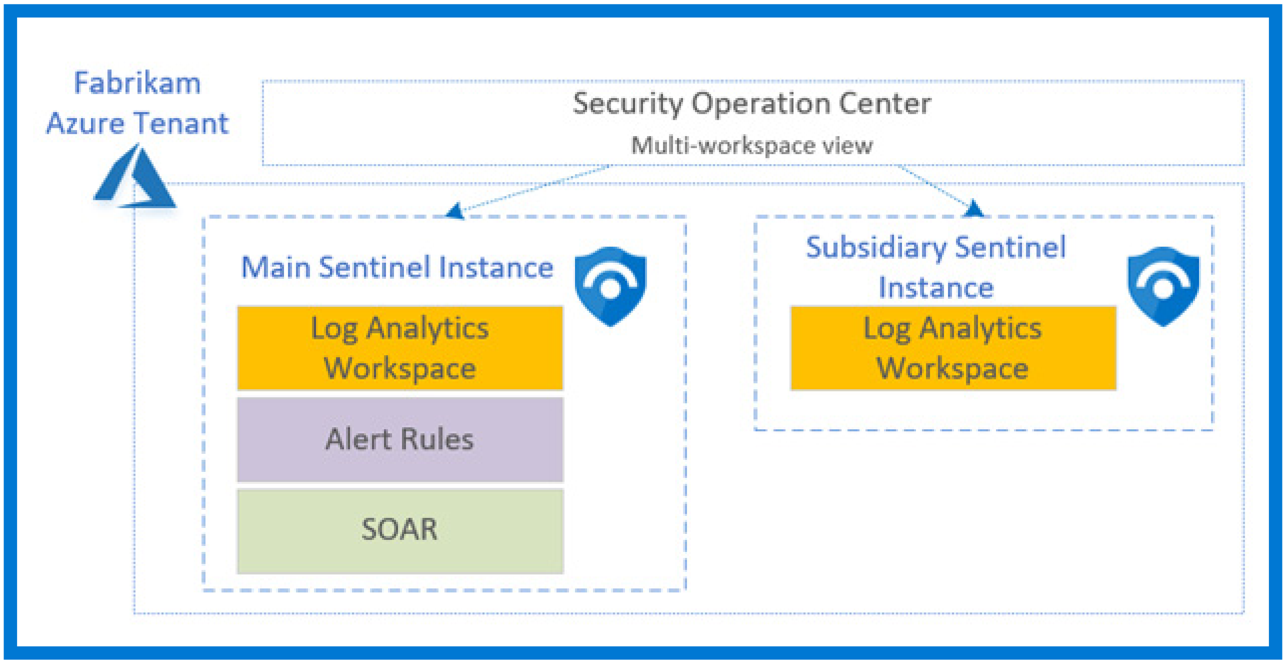

Fig. 1. A single Azure Microsoft Sentinel instance can be used to monitor subsidiary Azure Microsoft Sentinel instances across multiple Azure Active Directory tenants

Number of Azure Subscriptions

A single Microsoft Sentinel instance can integrate data from multiple Azure subscriptions; however, some security or operational solutions may have restrictions on their logging capabilities that limit sending logging/ diagnostics data to different subscriptions. For example, Microsoft Endpoint Manager can only send audit logs to a log analytics workspace within its subscription. That is a limitation of that solution, not of Microsoft Sentinel. Such solutions require a custom connector that can be configured using a variety of methods.

Microsoft Sentinel can be deployed in existing subscriptions or in its own subscription without any implications for its functionality.

A dedicated subscription is recommended in the following situations:

- There is a need to clearly identify or segregate any costs associated with Microsoft Sentinel.
- Permissions need to be assigned at the subscription level to allow the creation and management of various resources required for a full Microsoft Sentinel configuration. In a complex environment, this could be VMs, function apps, automation accounts, storage accounts, data explorer clusters, machine learning, key vaults, and databases.

If Microsoft Sentinel is deployed in multiple subscriptions, access can be managed centrally through regular assignment of Azure AD roles. Azure Lighthouse is not required in this case, as it is designed to provide cross-tenant access.

Case Study: K-12 school board

ABC School is a large school board in North America with more than 100,000 students and 7,000 staff members. ABC School has organizational separation between network operations, security operations, and server management teams, where each team has specific access rights to view and access resources (segregation of duties) and different cost centers.

ABC School has decided to deploy Microsoft Sentinel in a new Azure subscription under the abcschool.com tenant, as Microsoft Sentinel was to be used.

Number of Azure Resource Groups

As is the case with most Azure resources, a Microsoft Sentinel Log Analytics workspace resides in a resource group. A resource group is a container that holds related resources for an Azure solution; in this case, it would be Microsoft Sentinel. Resource groups allow for granularity in assigning permissions and logical grouping of resources based on their purpose.

As a solution, Microsoft Sentinel will use multiple types of resources—some mandatory, some optional—such as Log Analytics workspaces, workbooks, Logic Apps, API connections, function apps, automation accounts, storage accounts, key vaults, application insights, VMs, and many others. In most cases, a single resource group is sufficient, but in certain instances, such as those where different types of Azure Function Apps need to be used (some based on Linux, some on Windows), the full solution may span multiple resource groups.

If a dedicated subscription is not practical, it is highly recommended to maintain all Microsoft Sentinel–related resources in a dedicated resource group.

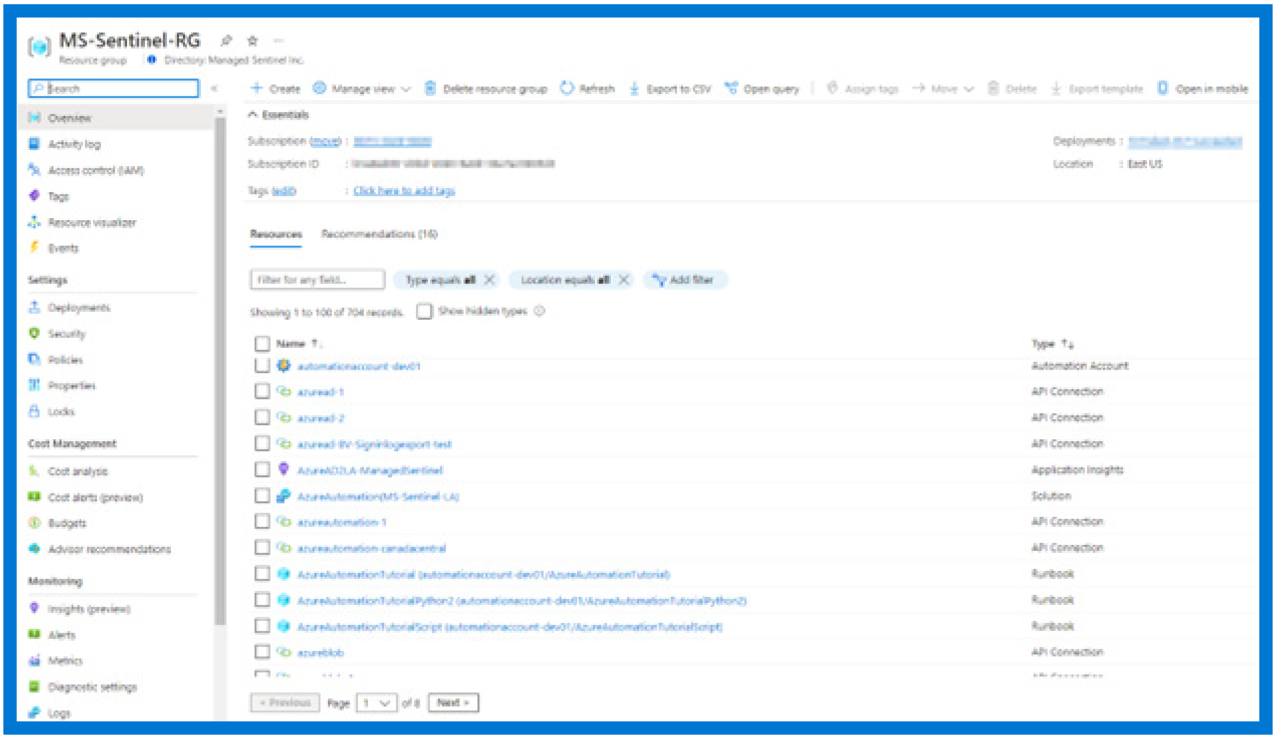

MS-Sentinel-RG

Resource group | Directory: Managed Sentinel Inc.

Essentials

Subscription (move) : [redacted]

Subscription ID : [redacted]

Tags (edit) : Click here to add tags

Deployments : [redacted]

Location : East US

Showing 1 to 100 of 704 records.

Name	Type
automationaccount-dev01	Automation Account
azuread-1	API Connection
azuread-2	API Connection
azuread-BV-Signinlogexport-test	API Connection
AzureAD2LA-ManagedSentinel	Application Insights
AzureAutomation(MS-Sentinel-LA)	Solution
azureautomation-1	API Connection
azureautomation-canadacentral	API Connection
AzureAutomationTutorial (automationaccount-dev01/AzureAutomationTutorial)	Runbook
AzureAutomationTutorialPython2 (automationaccount-dev01/AzureAutomationTutorialPython2)	Runbook
AzureAutomationTutorialScript (automationaccount-dev01/AzureAutomationTutorialScript)	Runbook
azureblob	API Connection

Fig. 2. Example of resources used in a Microsoft Sentinel solution

Distribution of Azure PaaS Resources

There is no cost for traffic that spans between Azure PaaS regions, but traffic egressed to non-Azure environments such as internet and on-premises virtual private network (VPN) incurs a bandwidth cost. When considering an Azure region for the Log Analytics workspace used for Microsoft Sentinel, the cost of transferring data out of other Azure regions or other cloud providers should be understood and taken into consideration as an additional cost.

For this reason, the preferred location should be in the region with the majority of log-generating Azure resources. Multiple Microsoft Sentinel instances can be deployed, but the reduced bandwidth costs should be weighed against the additional management complexity.

Additional resources for Microsoft Azure bandwidth pricing:

Pricing - Bandwidth | Microsoft Azure

Data Segregation Requirements

Some organizations have strict requirements on the accessibility of logging data between different business units due to legislative or regulatory compliance requirements or internal dictates. Permissions can be applied to specific types of logging data within a single Microsoft Sentinel instance, but for full, clear isolation, a dedicated Microsoft Sentinel instance should be considered.

A common scenario is a central SOC that requires visibility across the entire environment but also needs organizational units to be able to access the logging data from only their own resources. For example, a manufacturing business may have a dedicated Microsoft Sentinel instance that collects logging data from operational technologies (OT) devices and must provide access to analysts specialized in OT. These can include application developers, dedicated security analysts, and other specialized roles. By creating separate Microsoft Sentinel instances, it is possible to provide the OT teams full visibility to the logs from devices under their purview while maintaining full visibility for the central SOC.

Because the overall volume of logs will remain the same, there are no additional Azure ingestion costs for multiple Microsoft Sentinel instances.

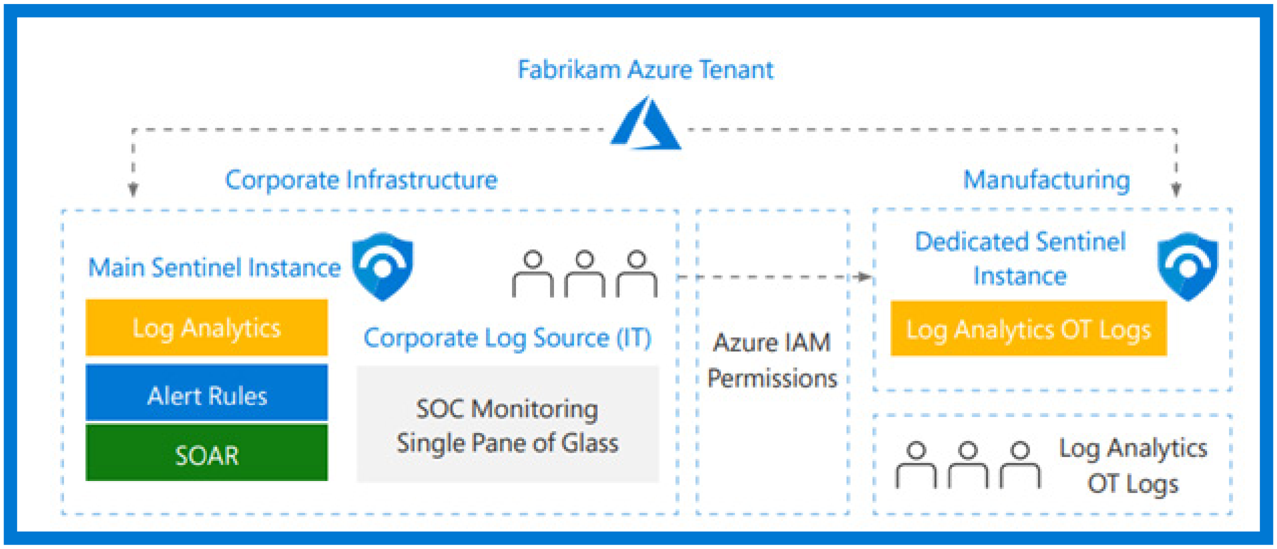

Fig. 3. Separation of Microsoft Sentinel instances for corporate versus manufacturing units.

Complex Organizational Structures

SIEM deployments are typically driven by the IT security department to address specific security needs. A SIEM can be an expensive security control; therefore, it is common for multiple business units to contribute to the overall expense.

Based on the IT security tenet of "separation of duties," access to the analytical data and the security events generated should generally stay within the Chief Information Security Officer's control.

Various organizational units, such as human resources (HR), may require a level of access to specific dashboards or sets of data. Microsoft Sentinel provides the ability to assign table-level permissions and limit the level of access to the minimum required to perform requisite job functions.

Role-based Access Control (RBAC) Requirements

- Microsoft Sentinel provides an extensive list of Azure built-in roles that can be used to provide granular access based on job requirements and permitted level of access. Part of these roles are several dedicated Microsoft Sentinel roles:

- Microsoft Sentinel Contributor - Can perform all engineering-related configuration, such as creating alert rules, configuring data connectors, and additional similar tasks.

- Microsoft Sentinel Reader - Can query the log data stored in Microsoft Sentinel but cannot modify any settings.

- Microsoft Sentinel Responder - Can query the log data, can view and update incidents raised by security rules, but cannot modify alert rules.

- Microsoft Sentinel Playbook Operator - Can view and run automation playbooks.

- Microsoft Sentinel Automation Contributor - Not applicable to users, is used to grant the ability to associate playbooks with automation rules.

In addition to Microsoft Sentinel-specific roles, there are additional roles required to fully configure and use Microsoft Sentinel:

- Logic App Contributor - For creation and management of SOAR playbooks.

- Workbook Contributor - For creation and management of workbooks (Microsoft Sentinel dashboards).

- Workbook Reader - For accessing / reading workbooks.

Depending on the requirements, custom RBAC roles can be created and associated with the Microsoft Sentinel instance. While custom roles are supported within the same Azure tenant, they cannot be used for access via Azure Lighthouse.

The roles can be applied at the subscription or resource group level with the recommendation being to provide the minimum permissions required to perform the job.

Assigning the Microsoft Sentinel Reader role provides access to all the logs by default, but custom roles can be assigned on a "per table" basis to limit visibility to just specific Microsoft Sentinel tables.

For Azure resources, resource-specific permissions can be applied using a resource-context RBAC that can provide multiple options to identify specific resources and grant permissions to the appropriate users or groups.

If access to data is needed for only a subset of the data in Microsoft Sentinel tables, there are options to provide read-only dashboards or present the data based on custom queries via Microsoft Power BI.

Case Study: SLG Organization in the US

An SLG (State Local Government) organization in the US with more than 75,000 employees across 50+ agencies decided to build its centralized security monitoring platform with the Microsoft Sentinel platform. Having a single Azure tenant with some log sources being centralized posed several challenges to the SLG security team regarding RBAC, data segregation between different agencies, and various data retention policies.

The team decided to build a multi-workspace Sentinel infrastructure with each agency having its own Microsoft Sentinel instance where agency-related logs are collected, and only agency personnel are authorized to see and manage data. Azure AD, Azure Activity, and Office Activity are tenant-level logs that were ingested in a centralized, single-purpose Sentinel instance to which only the SLG overall SOC team would have access.

A custom Azure function was built to extract specific agency logs from these three tables based on agency labels stored in the IdentityInfo table created by the Microsoft Sentinel UEBA engine. All logs related to specific agency personnel get moved every 5 minutes to a custom table in the agency Sentinel workspace, where custom parsers have been created.

This method allowed each agency to see only data related to their environment and build advanced correlation rules between agency infrastructure device logs and user authentication events from Azure AD.

Additional resources about roles and permissions:

Azure built-in roles - Azure RBAC | Microsoft Docs

Azure custom roles - Azure RBAC | Microsoft Docs

Roles and permissions for working in Microsoft Sentinel | Microsoft Docs

Table-level RBAC in Microsoft Sentinel - Microsoft Tech Community

Permissions in Microsoft Sentinel | Microsoft Docs Controlling access to Microsoft Sentinel data: Resource RBAC - Microsoft Tech Community

Ingestion of Operational Logs Versus Security Logs

Organizations may also collect performance and operational logs for monitoring purposes and day-to-day operations analytics, such as various performance counters and diagnostics logs.

These types of logs fall under the "availability" umbrella of the confidentiality–integrity–availability (CIA) triad.

This type of data sometimes overlaps with the security-related logging data, and from a logistics perspective, all the logging data may be aggregated into a single log analytics workspace. Enabling Microsoft Sentinel will apply the additional costs to the full set of data, so if the operational data represents a significant proportion of the overall logs, then a separation of operational and security logs is recommended. Some organizations consider that the additional visibility or correlation across the entire data set justifies the additional expense.

In Our Experience,

- Existing Log Analytics workspaces that store both security and operational logs can increase Microsoft Sentinel costs. Often the costs associated with Microsoft Sentinel are intended to be borne by a security organization, but operational logs in Log Analytics may result in IT operations costs being hidden within security spending.
- After cost versus benefit analysis, organizations often decide to maintain the aggregated operational and security logs together. That takes advantage of the potential correlation capabilities, with the operations department contributing to the operational costs through internal cost allocation mechanisms.

Case Study: Manufacturing Company

Fabrikam Inc. is a large international manufacturing company with offices and data centers on three continents. Fabrikam is in the process of migrating all internal workloads to Azure Cloud, with substantial Windows and Linux server estate located on-premises. Fabrikam Inc. has decided to use Azure Monitor for server monitoring as part of its Azure migration strategy.

The legacy environment included a total of 780 servers, both Unix and Windows, located in Azure and within the three data centers on-premises, logging to a QRadar SIEM solution. The Fabrikam server operations team deployed MMA and Operations Management Suite (OMS) agents to all servers. The team has determined that the total daily Azure Monitoring logging was approximately 360 GB/day.

Fabrikam corporate security decided to migrate the on-premises QRadar SIEM to Microsoft Sentinel. During the initial setup of Microsoft Sentinel, the team enabled security events collection for all remote MMA agents using common-level event streaming. This configuration increased the log collection by adding another 290 GB/day in the SecurityEvent Log Analytics table.

Operational logs can satisfy key business and technical requirements; however, in our experience, projects should account for operational logging requirements as a separate cost.

Estimation of Log Ingestion Volume and Pricing Model

The estimation of future log ingestion is a challenging exercise. During a Microsoft Sentinel deployment, organizations are typically trying to include many log sources that are new to them and have not been included in a prior SIEM solution. They are also trying to include the full range of logs from existing log sources that were previously only logging partially (or not at all).

The additional difficulty arises from the fact that each organization has a different distribution of log source types that are based on their industry, number of users, internet presence, distribution of their assets (on-premises versus cloud versus SaaS), and compliance requirements.

Selecting a sample of log sources and configuring them to send full logs for a typical day (or a typical week) is often the most precise way to estimate the log ingestion volume.

With all the caveats mentioned previously and if a testing sample is not possible, our own statistics run against 400+ Microsoft Sentinel instances show the following 30-day log ingestion averages based on three types of number of active users (Azure AD, Office 365 or Windows AD).

Customer size stats using the number of Azure AD Users as customer size

Customer Size	Average Data Size (30 days)	Estimated monthly cost (USD)
Under 1,000	0.70 TB	$3,007.10
Between 1,001 and 3,000	2.30 TB	$8,278.73
Between 3,001 - 5,000	3.30 TB	$9,768.00
Between 5,001 - 10,000	3.34 TB	$10,063.70
Over 10,000	8.02 TB	$21,974.53

Customer size stats using the number of Office 365 Mailboxes as customer size

Customer Size	Average Data Size (30 days)	Estimated monthly cost (USD)
Under 1,000	0.64 TB	$3,944.93
Between 1,001 and 3,000	1.24 TB	$5,320.07
Between 3,001 - 5,000	3.44 TB	$10,182.70
Between 5,001 - 10,000	2.47 TB	$8,669.27
Over 10,000	6.63 TB	$18,166.20

Note: The apparent discrepancy in the log volume between 3001-5000 and 5001-10000 (larger organizations have lower log volume) is the result of the relatively low from a statistical perspective (several hundred Sentinel instances) number of data points. We did not want to adjust the data to fit the expected volume.

Customer size stats using the number of Windows AD users as customer size

Customer Size	Average Data Size (30 days)	Estimated monthly cost (USD)
Under 1,000	0.92 TB	$2,714.27
Between 1,001 and 3,000	2.77 TB	$8,880.00
Between 3,001 - 5,000	3.56 TB	$10,537.90
Between 5,001 - 10,000	4.34 TB	$12,846.70
Over 10,000	11.31 TB	$30,160.00

Important notes about the estimates:

- The costs are for the specified log ingestion volumes, with the corresponding discounts (i.e. if a commitment tier such as 300 GB/day is used), for the East US Azure region.
- The estimates are based on 500+ Microsoft Sentinel deployments and this may not be a number sufficiently large to be statistically significant.
- The standard deviation for the data sets is very high, meaning a wide variation in the values collected. The values used are for the last 30 days from the current day, but values may vary if a longer time interval is used.
- The "per user" estimates are typically the most common request from our customers. Depending on the technology used by each organization, the most relevant set of data (i.e., Azure AD users for an Azure-heavy organization or Windows AD users with a significant on-prem infrastructure) should be used.

Based on the expected logging volume, a pricing model can be selected to take advantage of the commitment tier discounts offered by Microsoft, especially when logging volume is estimated to be consistently over 100 GB/day. (Note: The figures above are for example only and may vary over time.)

Architecture Design Output

At the conclusion of the high-level design phase, the following items should be decided and documented:

- Azure region used for the Microsoft Sentinel Log Analytics workspace (or workspaces, if more than one region is to be used).
- Azure subscription, resource group, and log analytics workspace, including naming convention and tags.
- Azure AD groups and the RBAC to be applied to each.
- Log sources in scope (on-premises, cloud, SaaS).
- Microsoft Sentinel data connectors required to ingest in-scope log sources.
- Custom data connectors to be developed (if applicable).
- On-premises syslog collectors (quantity, location, operating system [OS] type, any additional configuration).
- Initial list of use cases to be implemented.
- Internet/VPN/LAN/WAN connectivity between log sources and Microsoft Sentinel.
- Estimated log ingestion volume (GB/day). Data retention policy (in days or months).
- Pricing model (pay-as-you-go) versus reserved capacity, depends on the estimated log ingestion volume.

Deployment

Once the high-level design is completed, the provisioning of Microsoft Sentinel and the related resources can be initiated. The following sections highlight the key considerations for the deployment phase.

Azure Resources

As an analytical solution built around Log Analytics workspace and Logic Apps, Microsoft Sentinel requires the following resources to be created:

- Subscriptions (if a dedicated subscription(s) will be used)
- Resource group(s)
- Log Analytics workspace(s) Automation rules/playbooks Alert rules
- Workbooks

Azure Global Admin, Subscription Owner, or Azure Security Administrator permissions are typically required to create these resources and enable Microsoft Sentinel for the selected Log Analytics workspace.

During the deployment, in addition to the Azure region, some basic configuration will be required, such as log retention (the default is 90 days) and the selection of a pricing model.

Automation rules, playbooks, alert rules, and workbooks are typically created gradually following the onboarding of various log sources. These resources are part of the ongoing SIEM tuning and maintenance and should follow the typical change control procedures.

Microsoft Sentinel provides hundreds of alert rules, workbooks and automation playbook templates along with hunting scripts (threat-hunting scripts that are typically used ad-hoc and not as alerting rules). The templates can be used to activate/deploy schedule alerts, create customized dashboards, create automation playbooks and perform threat-hunting activities. In most cases, once deployed, the resources created have to be adjusted to match the existing environment, configure local credentials, etc.

Methods of deployment:

- **Manual** - Using the Azure portal, the administrator manually configures the Microsoft Sentinel resources. Any manual process has the inherent risks of human operator error, lack of compliance with potential change control procedures, and undocumented changes.
- **Automation tools** - Microsoft Sentinel resources support several infrastructure-as-code tools, such as Hashicorp Terraform, that can provide consistency to processes. One of the native Microsoft Sentinel features, allows the use of a GitHub or Azure DevOps repository to automate the deployment of detection rules, Microsoft Sentinel functions and workbooks. The repositories can be used to deploy content across multiple Microsoft Sentinel instances. An upcoming Microsoft Sentinel feature, provisionally called Workspace Manager, allows for the configuration of various

groups of Microsoft Sentinel instances with subsets of detection rules pushed automatically from a centralized Microsoft Sentinel instance.

- **Custom scripts/programs** - Microsoft provides a PowerShell automation library called Az.SecurityInsights that can be used to script a wide range of Microsoft Sentinel–related deployment tasks. In a similar fashion, the SecurityInsights REST API is accessible through Python as part of the azure.mgmt.securityinsight module. The Microsoft Sentinel community provides additional resources, such as the AzSentinel PowerShell library and a wide range of Azure Resource Manager (ARM) templates for a variety of Microsoft Sentinel playbooks, alert rules, and workbooks. You can also use the All-In-One project from GitHub that will deploy a basic Sentinel instance. More information can be found at: [Azure-Sentinel/Tools/Sentinel-All-In-One at master · Azure/Azure-Sentinel (github.com)](https://github.com)

Optional Azure resources:

- Service principal names (SPNs) - SPNs are typically used in automation playbooks for authentication when accessing various resources for retrieval of log data or execution of automation tasks. SPNs should be provided with the minimal permissions required to perform their tasks.

- Storage accounts - Storage accounts can be used for temporary storage of analysis data, long-term log retention, and other tasks that require basic storage outside Log Analytics.

- Function apps - As serverless compute resources, function apps can perform a wide variety of tasks related to log collection and automation tasks. Only one type of compute platform can be used for one resource group. If, for example, both Python and .Net function apps are required, they need to be deployed in different resource groups. Some of the built-in Microsoft Sentinel data connectors require the deployment of function apps.

- Key vaults - Typically used for secure storage of secrets used by various log collectors.

- Event hubs - In certain designs that require integration with legacy SIEMs or third-party analytical platforms, Azure Event Hubs can be used to share analysis data.

None of the optional resources are required for the initial Microsoft Sentinel configuration, but they may end up being used as the complexity of the deployment increases and new automation tasks are required.

In Our Experience,

Microsoft Sentinel deployment projects can run into foreseeable challenges that can be avoided with appropriate project management and advance planning. We have included a few key items to consider:

- Understand your data as it pertains to potential data residency requirements; there may be a need to choose one region over another.
- Clearly identify and coordinate with Azure administrators at the inception of your project.
- Keep the Azure administrators briefed on the scope and status of the project.
- Consider naming conventions, tagging, etc. Where none exist, plan to future-proof the deployment.
- Run a risk assessment with your stakeholders about whether to create new log analytics workspaces (starting fresh) versus using the pre-existing legacy log analytics workspaces.
- Understand your planned automation tasks and their touchpoints to ensure that required SPN permissions are pre-staged. For example, what permissions are needed to enable or disable an Azure AD user account?

Additional Resources:

Azure Sentinel PowerShell Module Az.SecurityInsights has been released to GA - Microsoft Tech Community

Enable Continuous Deployment Natively with Microsoft Sentinel Repositories! - Microsoft Tech Community

Log Source Onboarding

Microsoft Sentinel includes many data connectors for a wide range of log sources such as Azure (Azure AD, PaaS), Microsoft 365 (Defender) solutions, non-Azure cloud (e.g., AWS, GCP), on-premises sources (e.g., firewalls, Network Access Control, VPN, LAN/WAN, Windows Active Directory, DNS), SaaS (multiple solutions), and threat intelligence feeds.

Depending on the type of log source, the onboarding process varies from just a few clicks of the mouse, deployment of AMA, or to more complex configurations involving the deployment of additional log collection resources such as Microsoft Sentinel playbooks based on Azure Logic Apps, Azure Function Apps, and vendor-provided log retrieval tools.

For each supported data connector, Microsoft Sentinel provides full instructions on how to onboard a particular log source and, where applicable, with options to automate the deployment of the required data collectors and relevant log parsers.

Case Study: European Pharmaceutical Company

Pharma Inc. is a multinational company with offices in 7 countries on three continents. Prior to the migration to Microsoft Sentinel, Pharma Inc. was running a LogRhythm SIEM on-premises with a limited number of log sources. Pharma Inc. SOC and Compliance teams requested to have visibility into the entire environment, including cloud resources, all SaaS applications and all healthcare-related applications used by the production teams. During the initial discovery sessions, the Pharma Inc. security team identified a list of 105 unique log source types that will be required to be onboarded in Microsoft Sentinel. Out of those, only 35 were out of the box data connectors from the Microsoft Sentinel platform and the rest required development work.

During the initial onboarding in Microsoft Sentinel, analysis was completed based on sample data from several log sources. The Pharma Inc. security team decided to not pursue the log ingestion activities for 40 log sources, which were considered to have very little value from a security analytics perspective.

The team developed custom data connectors for 40 log source types using various methods described in this whitepaper. The total effort required for the development of these new custom data connectors was 4 weeks. This case study highlights the importance of performing an initial analysis of the security value of data being analyzed in a SIEM solution. The project duration and overall cost were reduced substantially with appropriate advanced planning.

Built-in Data Connectors

Microsoft Sentinel includes many connectors that can be deployed with a few clicks via the Microsoft Sentinel portal and the requisite RBAC permissions. That includes Azure AD, Azure subscription activity, Office 365, and the whole family of Microsoft Defender products. New data connectors for other products are added on a regular basis. Consider the built-in data connectors over custom ones, where feasible, as they are fully supported by Microsoft and the Microsoft Sentinel community.

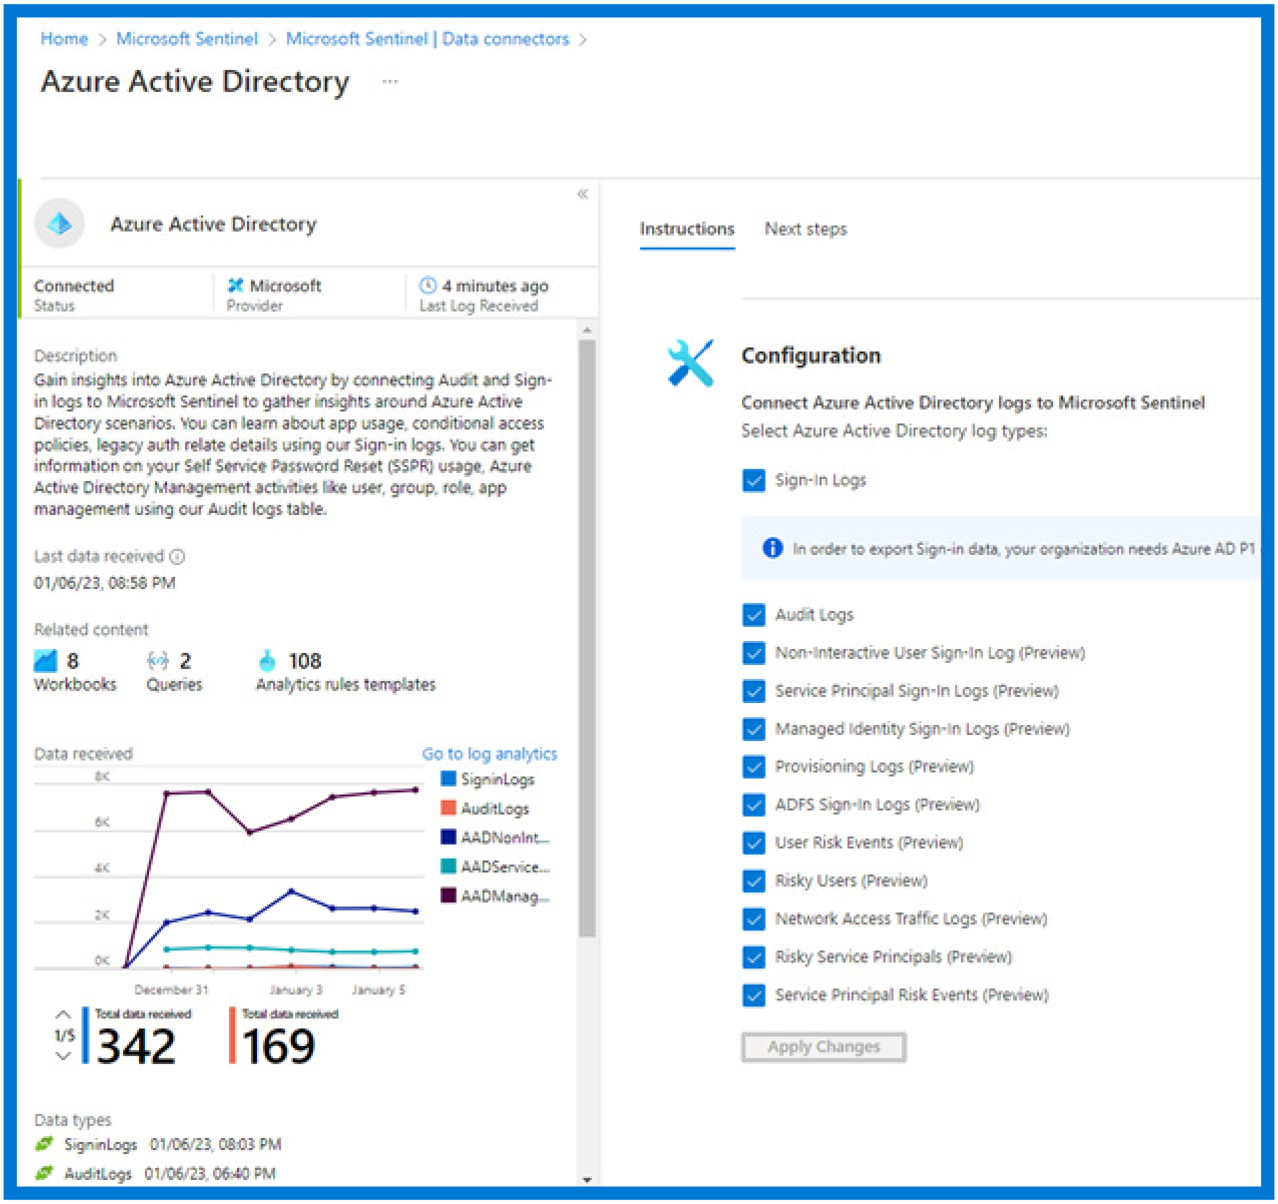

Azure Active Directory ···

Azure Active Directory

Connected
Status

Microsoft
Provider

4 minutes ago
Last Log Received

Description

Gain insights into Azure Active Directory by connecting Audit and Sign-in logs to Microsoft Sentinel to gather insights around Azure Active Directory scenarios. You can learn about app usage, conditional access policies, legacy auth relate details using our Sign-in logs. You can get information on your Self Service Password Reset (SSPR) usage, Azure Active Directory Management activities like user, group, role, app management using our Audit logs table.

Last data received

01/06/23, 08:58 PM

Related content

8 Workbooks
2 Queries
108 Analytics rules templates

Data received

Total data received 342

Total data received 169

Data types

SigninLogs 01/06/23, 08:03 PM

AuditLogs 01/06/23, 06:40 PM

Configuration

Connect Azure Active Directory logs to Microsoft Sentinel

Select Azure Active Directory log types:

- ☑ Sign-In Logs

In order to export Sign-in data, your organization needs Azure AD P1

- ☑ Audit Logs
- ☑ Non-Interactive User Sign-In Log (Preview)
- ☑ Service Principal Sign-In Logs (Preview)
- ☑ Managed Identity Sign-In Logs (Preview)
- ☑ Provisioning Logs (Preview)
- ☑ ADFS Sign-In Logs (Preview)
- ☑ User Risk Events (Preview)
- ☑ Risky Users (Preview)
- ☑ Network Access Traffic Logs (Preview)
- ☑ Risky Service Principals (Preview)
- ☑ Service Principal Risk Events (Preview)

Fig. 4. Azure AD built-in Microsoft Sentinel data connector

Currently estimated to take place in Q2, 2023, an upcoming change in Microsoft Sentinel will include all the data connectors, detections (rule templates), hunting queries, log parsers, automation playbooks and workbooks available through the Microsoft Sentinel Content Hub, allowing for the deployment of a particular solution in a single step.

Microsoft Sentinel Content Hub

Microsoft Sentinel content is Security Information and Event Management (SIEM) solution components that enable customers to ingest data, monitor, alert, hunt, investigate, respond, and connect with different products, platforms, and services.

Content in Microsoft Sentinel includes any of the following types:

- Data connectors provide log ingestion from different sources into Microsoft Sentinel
- Parsers provide log formatting/transformation into ASIM formats, supporting usage across various Microsoft Sentinel content types and scenarios
- Workbooks provide monitoring, visualization, and interactivity with data in Microsoft Sentinel, highlighting meaningful insights for users
- Analytics rules provide alerts that point to relevant SOC actions via incidents
- Hunting queries are used by SOC teams to proactively hunt for threats in Microsoft Sentinel
- Notebooks help SOC teams use advanced hunting features in Jupyter and Azure Notebooks
- Watchlists support the ingestion of specific data for enhanced threat detection and reduced alert fatigue
- Playbooks and Azure Logic Apps custom connectors provide features for automated investigation, remediation, and response scenarios in Microsoft Sentinel

Microsoft Sentinel offers these content types as solutions and standalone items. Solutions are packages of Microsoft Sentinel content or Microsoft Sentinel API integrations, which fulfill an end-to-end product, domain, or industry vertical scenario in Microsoft Sentinel. Both solutions and standalone items are discoverable and managed from the Content hub.

You can either customize out-of-the-box (OOTB) content for your own needs, or you can create your own solution with content to share with others in the community. For more information, see the Microsoft Sentinel Solutions Build Guide for solutions' authoring and publishing.

Important: The Microsoft Sentinel Content hub and solutions are currently in PREVIEW, as are all individual solution packages. See the Supplemental Terms of Use for Microsoft Azure Previews for additional legal terms that apply to Azure features that are in beta, preview, or otherwise not yet released into general availability.

Discover and manage Microsoft Sentinel content

Use the Microsoft Sentinel Content hub to centrally discover and install out-of-the-box (OOTB) content.

The Microsoft Sentinel Content hub provides in-product discoverability, single-step deployment, and enablement of end-to-end product, domain, and/or vertical OOTB solutions and content in Microsoft Sentinel.

- In the Content hub, filter by categories and other parameters, or use the powerful text search, to find the content that works best for your organization's needs. The Content hub also indicates the support model applied to each piece of content, as some content is maintained by Microsoft and others are maintained by partners or the community.

Manage updates for out-of-the-box content via the Microsoft Sentinel Content hub, and for custom content via the Repositories page.

- Customize out-of-the-box content for your own needs, or create custom content, including analytics rules, hunting queries, notebooks, workbooks, and more. Manage your custom content directly in your Microsoft Sentinel workspace, via the Microsoft Sentinel API, or in your own source control repository, via the Microsoft Sentinel Repositories page.

Why content hub solutions?

Microsoft Sentinel solutions are packaged integrations that deliver end-to-end product value for one or more domain or vertical scenarios in the content hub.

The solutions experience, powered by Azure Marketplace, helps you discover and deploy the content you want. For more information on authoring and publishing solutions in the Azure Marketplace, see the Microsoft Sentinel Solutions Build Guide.

- **Packaged content** are collections of one or more components of Microsoft Sentinel content, such as data connectors, workbooks, analytics rules, playbooks, hunting queries, watchlists, parsers, and more.

- **Integrations** include services or tools built using Microsoft Sentinel or Azure Log Analytics APIs that support integrations between Azure and existing customer applications, or migrate data, queries, and more, from those applications into Microsoft Sentinel.

You can also use solutions to install packages of out-of-the-box (OOTB) content in a single step, where the content is often ready to use immediately. Providers and partners use Sentinel solutions to add value to their customers' investments by delivering combined product, domain, or vertical value.

Use the Content hub to centrally discover and deploy solutions and OOTB content in a scenario-driven manner.

For more information, see:

- Centrally discover and deploy Microsoft Sentinel out-of-the-box content and solutions

- Microsoft Sentinel solutions catalog in the Azure Marketplace

- Microsoft Sentinel catalog

Categories for Microsoft Sentinel out-of-the-box content and solutions

Microsoft Sentinel out-of-the-box content can be applied with one or more of the following categories. In the Content hub, select the categories you want to view to change the content displayed. You can discover community delivered items centrally in Content hub as standalone content or solutions.

Category name	Description
Application	Web, server-based, SaaS, database, communications, or productivity workload
Cloud Provider	Cloud service
Compliance	Compliance product, services, and protocols
DevOps	Development operations tools and services
Identity	Identity service providers and integrations
Internet of Things (IoT)	IoT, OT devices and infrastructure, industrial control services
IT Operations	Products and services managing IT
Migration	Migration enablement products, services, and
Networking	Network products, services, and tools
Platform	Microsoft Sentinel generic or framework components, Cloud infrastructure, and platform
Security - Others	Other security products and services with no other clear category
Security - Threat Intelligence	Threat intelligence platforms, feeds, products, and services
Security - Threat Protection	Threat protection, email protection, and XDR and endpoint protection products and services
Security - 0-day Vulnerability	Specialized solutions for zero-day vulnerability attacks like Nobelium
Security - Automation (SOAR)	Security automations, SOAR (Security Operations and Automated Responses), security operations, and incident response products and services.
Security - Cloud Security	CASB (Cloud Access Service Broker), CWPP (Cloud workload protection platforms), CSPM (Cloud security posture management and other Cloud Security products and services

Security - Information Protection	Information protection and document protection products and services
Security - Insider Threat	Insider threat and user and entity behavioral analytics (UEBA) for security products and services
Security - Network	Security network devices, firewall, NDR (network detection and response), NIDP (network intrusion and detection prevention), and network packet capture
Security - Vulnerability Management	Vulnerability management products and services
Storage	File stores and file sharing products and services
Training and Tutorials	Training, tutorials, and onboarding assets
User Behavior (UEBA)	User behavior analytics products and services

Industry vertical categories

Microsoft Sentinel out-of-the-box content can be applied with one or more of the following categories. In the Content hub, select the categories you want to view to change the content displayed. You can discover community delivered items centrally in Content hub as standalone content or solutions.

Category name	Description
Aeronautics	Products, services, and content specific for the aeronautics industry
Education	Products, services, and content specific for the education industry
Finance	Products, services, and content specific for the finance industry
Healthcare	Products, services, and content specific for the healthcare industry
Manufacturing	Products, services, and content specific for the manufacturing industry
Retail	Products, services, and content specific for the retail industry
User Behavior (UEBA)	User behavior analytics products and services

Support models for Microsoft Sentinel out-of-the-box content and solutions

Both Microsoft and other organizations author Microsoft Sentinel out-of-the-box content and solutions. Each piece of out-of-the-box content or solution has one of the following support types:

Support model	Description
Microsoft-supported	Applies to: ● Content/solutions where Microsoft is the data provider, where relevant, and author. ● Some Microsoft-authored content/solutions for non-Microsoft data sources. Microsoft supports and maintains content/solutions in this support model in accordance with Microsoft Azure Support Plans. Partners or the Community support content/solutions that are authored by any party other than Microsoft.
Partner-supported	Applies to content/solutions authored by parties other than Microsoft. The partner company provides support or maintenance for these pieces of content/solutions. The partner company can be an Independent Software Vendor, a Managed Service Provider (MSP/MSSP), a Systems Integrator (SI), or any organization whose contact information is provided on the Microsoft Sentinel page for the selected content/solutions. For any issues with a partner-supported solution, contact the specified support contact.
Community-supported	Applies to content/solutions authored by Microsoft or partner developers that don’t have listed contacts for support and maintenance in Microsoft Sentinel. For questions or issues with these solutions, file an issue in the Microsoft Sentinel GitHub community.

Content sources for Microsoft Sentinel content and solutions

Each piece of content or solution has one of the following content sources:

Content source	Description
Content hub	Solutions deployed by the Content hub that support lifecycle management
Standalone	Standalone content deployed by the Content hub that is automatically kept up-to-date
Custom	Content or solutions you've customized in your workspace
Gallery content	Content from the feature galleries that don't support lifecycle management. This content source is retiring soon. For more information see OOTB content centralization changes.
Repositories	Content or solutions from a repository connected to your workspace

Microsoft Monitoring Agent (MMA)

The MMA is used across multiple Microsoft solutions and has undergone a few name changes as its features and functionality have evolved. Currently (January 2023), MMA is considered the "legacy" agent and its end-of-life is August 31st, 2024. The recommended agent for new Microsoft Sentinel deployments is the Azure Monitor Agent (AMA).

If you are still considering using MMA due to various factors, here are areas to ensure that have been covered during a Microsoft Sentinel deployment project:

- Consider the timelines and the strategy for a migration to AMA before the MMA end-of-life
- Consider the central deployment and management methodology for MMA
- Determine which endpoints are in scope for deployment (Windows domain controllers, critical application servers, standalone servers or workstations) as these affect the log ingestion volume significantly. It is very uncommon to collect logs from workstations through MMA.

Azure Monitor Agent (AMA)

While AMA has been available for a while, it has only recently caught up with all the features offered by the Microsoft Monitoring Agent, and in many cases, this delayed the adoption of AMA. Organizations that have MMA as the main Microsoft Sentinel agent have to start planning the migration as soon as possible in order to avoid a rushed migration in Q3 2024.

One of the main advantages of AMA is the ability to control the log collection policy centrally and apply different policies against different groups of computers, regardless if they are Azure VMs, on-premises servers or VMs in third-party cloud platforms like AWS and GCP. Through the deployment of Arc agent (an extension of Azure Resource Management that gives support to resources running outside of Azure), non-Azure compute infrastructure becomes manageable through Azure policies, including deployment of AMA and enforcement of data collection rules (DCRs).

AMA is supported on both Windows and Linux OS (MMA was used only on Windows, with the OMS agent acting as its counterpart on Linux).

AMA allows for increased granularity of the types of log collected. For example, for Windows event logs, individual event id / log sources combinations can be specified, allowing for the optimization of the log ingestion volume.

Multi-homing - AMA provides the ability to implement multiple log collection policies, i.e., logs can be sent to different log analytics workspaces, with different log collection options for each of them. This is available for both Windows and Linux (OMS on Linux did not support multi-homing).

The multi-homing option combined with the granular filtering enables the ability to split the stream of log data, with the critical one sent to Microsoft Sentinel for real-time monitoring, while less important (from a security perspective) log can be sent to a more cost effective storage option.

In Our Experience,

- Various Microsoft Sentinel data collectors vs. basis Azure Monitor DCRs - As AMA is quite versatile and can be used independently from Microsoft Sentinel (just as the basic Azure Monitor agent), some of the logs collected through Microsoft Sentinel have their own connectors to ensure that the logs arrive in the standard Microsoft Sentinel tables instead of the generic Event table used by Azure Monitor. For example, for Windows security event logs, the log collection has to be configured through the Microsoft Sentinel "Windows Security Events via AMA" data connector in order to have the events sent to the SecurityEvents table. The data connector using the MMA agent is called "Security Events via Legacy Agent." Configuring both the AMA and the legacy connectors can lead to data duplication so special attention has to be paid to ensure that only one connector is configured to collect the Windows security event logs (one of the log sources that generate large volumes of logs in most environments).

- DNS logs - Windows DNS logs can be collected through AMA through its own dedicated collector ('' Windows DNS Events via AMA (Preview)" and the logs are saved in the ASimDnsActivityLogs table, a different one from the DNS collected by the MMA agent – DNSEvents). Any detection rules written for the DNSEvents table should be adjusted to match the name and schema used in ASimDnsActivityLogs).

- WEF log collection – Some organizations have deployed WEF (Windows Event Forwarding) as a centralized Windows event log collection, mostly in order to avoid polling individual log servers for events of interest or deploying various analytical platform agents on just one server. AMA supports the logs collected through WEF (a dedicated data connector is available in Microsoft Sentinel, called "Windows Forwarded Events (Preview)". One caveat about this connector is that it stores the logs in the WindowsEvents table instead of the SecurityEvents used by the normal AMA connector. The schema used by these tables is different, and in order to analyze them using the same alert rule templates, their output has to be normalized using Microsoft Sentinel parsers (Microsoft Sentinel functions).

- Collection of local log files – One caveat for the collection of local text-based logs on endpoints with AMA installed is that the format of the files has to be consistent with the specified schema specified in the data collection rule and files with the inconsistent format will not be collected. In our experience, MMA is more tolerant to an inconsistent schema, discarding the entries that don't match the expected input while ingesting those that did match. An example of such files is the logs generated by the Microsoft DHCP Server, with its logs recording a header with some generic metadata, followed by the actual log entries for the DHCP activities. For such situations, MMA may be considered as a temporary log ingestion method until AMA provides a workaround.

- Defender for Cloud implications – AMA is used by other Microsoft solutions, one of the most important being Defender for Cloud. The engineers working on implementing Microsoft Sentinel and Defender for Cloud have to consider the implications of AMA configuration changes that can affect the other solution.

- Some of the AMA features are still in Preview mode (as of January 2023), and if there are formal restrictions in place regarding such "beta" or "preview" functionalities, these have to be reviewed accordingly. One example is the collection of WEF or DNS logs, still in preview mode.

- Compared to MMA, data collection rules (DCRs) in AMA may require additional expertise due to their increased versatility, especially when configuring custom DCRs for Windows event logs.

- While for certain uses, the Azure Arc agent requires a license, this is not the case when used just to deploy AMA.

Migration from MMA to AMA:

- Initiate the project as soon as possible in order to avoid a rushed migration in 2024

- Use tools such as the AMA Migration Helper workbook (available in Azure Monitor) to keep track of the migration progress

- Fully understand the differences on achieving various outcomes on AMA vs MMA or OMS (see AMA Supported services and features)

- Avoid duplication of data by turning on/off log collection policies are migration is progressing

Case Study: Large US Public School Board

A Public School organization is a large school board in the US with more than 100,000 students and 2,500 staff. Part of the Microsoft Sentinel onboarding activities, the school board wanted to reduce log consumption for the Windows security events and take advantage of the AMA agent functionality in terms of log collection policies.

That was applied to all 95 AD domain controllers to filter out some Windows Events collected from these servers and only keep those used for security detections. The team has created an XPATH query to collect only 21 Windows event IDs from AD domain controllers and forward them to the Log Analytics workspace where Microsoft Sentinel is running.

This approach allowed the school board to reduce the Microsoft Sentinel and Log Analytics consumption by 34%.

Additional Resources:

Azure Monitor Agent overview – Microsoft Docs

Migrate to Azure Monitor Agent from Log Analytics agent – Microsoft Docs

Tools for migrating from Log Analytics Agent to Azure Monitor Agent – Microsoft Docs

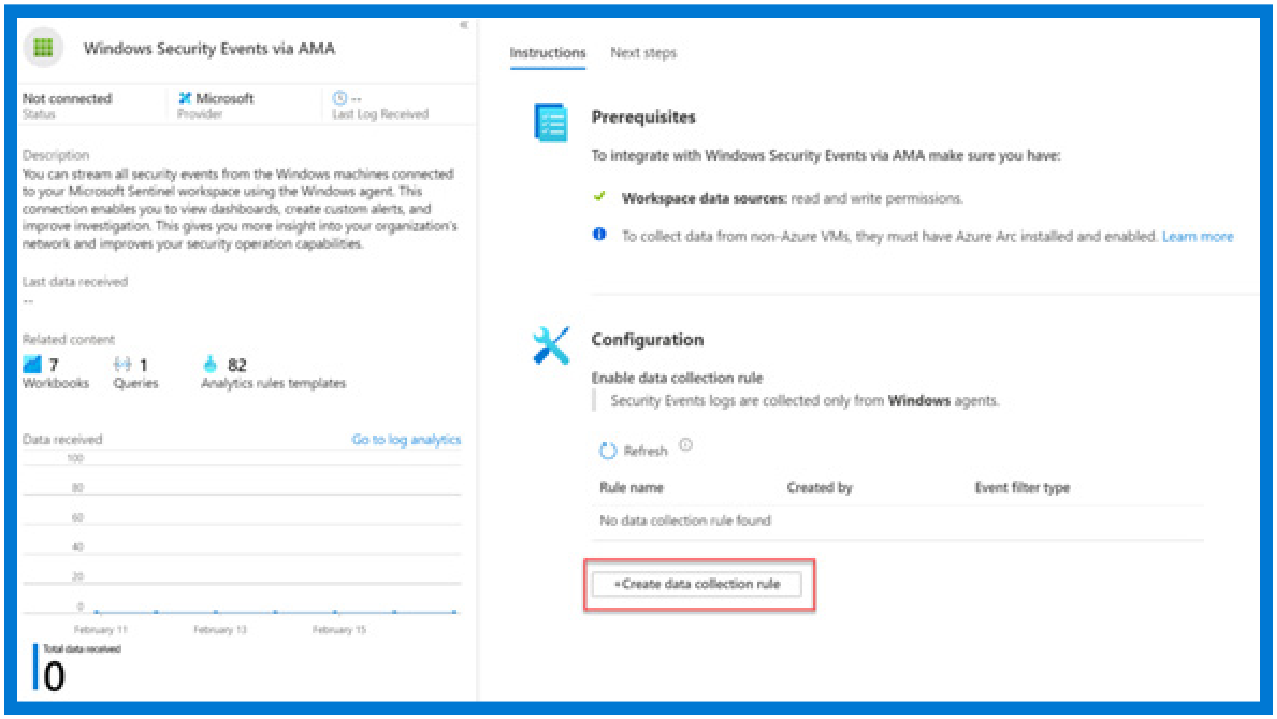

Description

You can stream all security events from the Windows machines connected to your Microsoft Sentinel workspace using the Windows agent. This connection enables you to view dashboards, create custom alerts, and improve investigation. This gives you more insight into your organization's network and improves your security operation capabilities.

Prerequisites

To integrate with Windows Security Events via AMA make sure you have:

✓ **Workspace data sources:** read and write permissions.

To collect data from non-Azure VMs, they must have Azure Arc installed and enabled. Learn more

Configuration

Enable data collection rule

Security Events logs are collected only from **Windows** agents.

Rule name	Created by	Event filter type
No data collection rule found		

Fig. 5. Configuration of the Windows Security Events via AMA data connector

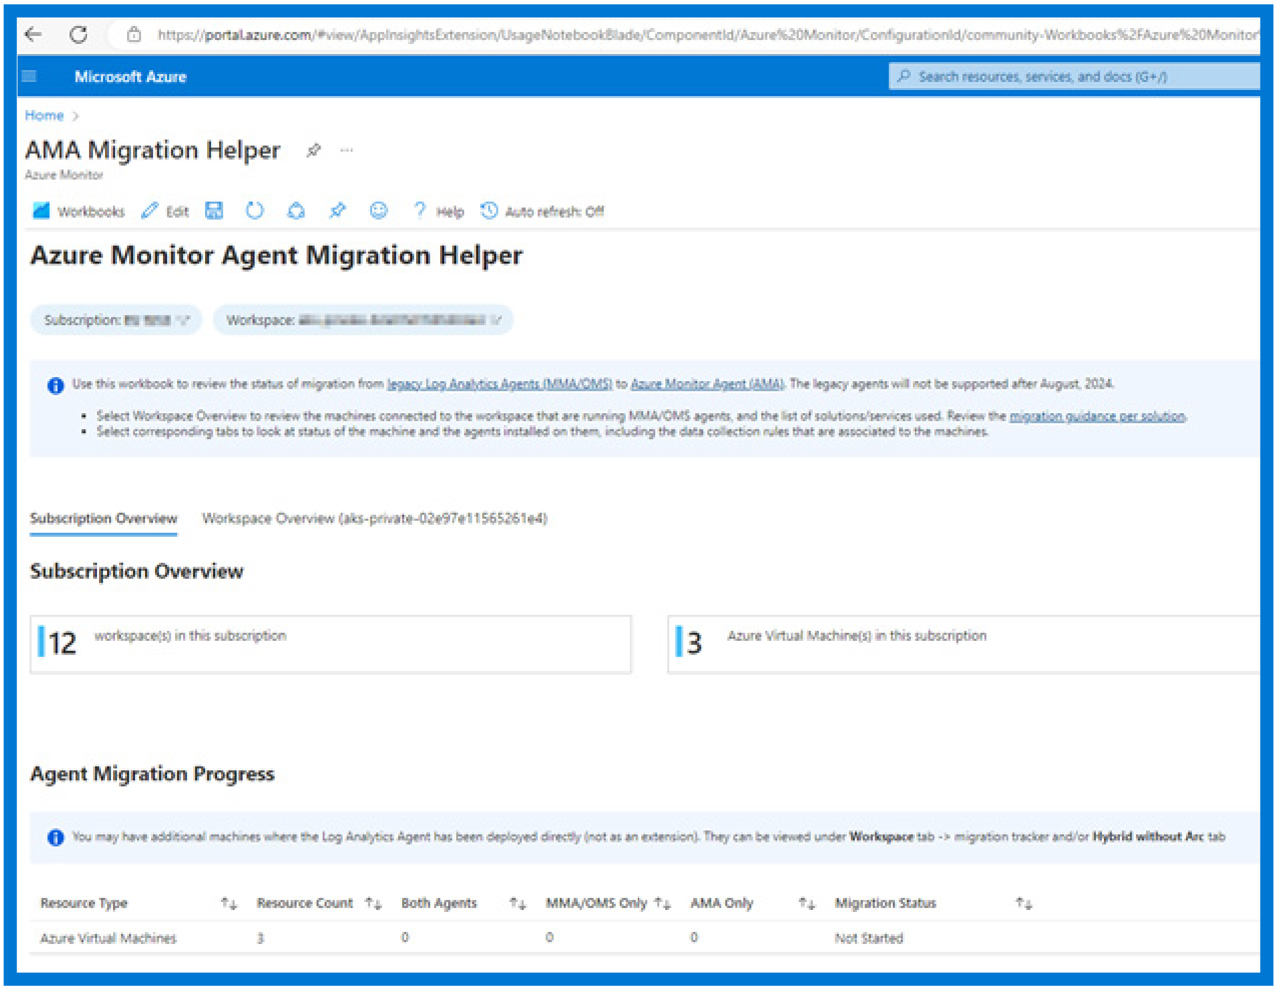
https://portal.azure.com/#view/AppInsightsExtension/UsageNotebookBlade/ComponentId/Azure%20Monitor/ConfigurationId/community-Workbooks%2FAzure%20Monitor
Microsoft Azure
Search resources, services, and docs (G+/)
Home >
AMA Migration Helper
Azure Monitor
Workbooks
Edit
Help
Auto refresh: Off

Azure Monitor Agent Migration Helper

Subscription: [redacted]

Workspace: [redacted]

Use this workbook to review the status of migration from legacy Log Analytics Agents (MMA/OMS) to Azure Monitor Agent (AMA). The legacy agents will not be supported after August, 2024.

- Select Workspace Overview to review the machines connected to the workspace that are running MMA/OMS agents, and the list of solutions/services used. Review the migration guidance per solution.
- Select corresponding tabs to look at status of the machine and the agents installed on them, including the data collection rules that are associated to the machines.

Subscription Overview

12 workspace(s) in this subscription

3 Azure Virtual Machine(s) in this subscription

Agent Migration Progress

You may have additional machines where the Log Analytics Agent has been deployed directly (not as an extension). They can be viewed under **Workspace** tab -> migration tracker and/or **Hybrid without Arc** tab

Resource Type ↑↓	Resource Count ↑↓	Both Agents ↑↓	MMA/OMS Only ↑↓	AMA Only ↑↓	Migration Status ↑↓
Azure Virtual Machines	3	0	0	0	Not Started

Fig. 5.1 – Azure Monitor Agent Migration Helper Workbook

Deploying the AMA on Linux allows for the collection of any syslog message or local logs that follow a consistent naming convention. The log collection can be filtered by both syslog facility and syslog message severity. Any Linux agent with AMA installed can act as a syslog collector for remote syslog log sources.

Deploying a Syslog Collector

For remote syslog log collection, Microsoft Sentinel requires a syslog server with Linux rsyslog or syslog-ng syslog servers, with rsyslog as the most common choice.

The server can be deployed on-premises, as a VM or physical server - or as a VM in Azure or other cloud environments. The main requirement for a VM is to provide routing and connectivity from the log sources that need to send syslog data.

For isolated locations with limited WAN or VPN connectivity to the syslog collector location, depending on the log sources capabilities, a TLS-based syslog server can be deployed in an Internet accessible location (such as an Azure VM) to receive encrypted logs via the Internet. Most TLS-aware log sources support self-signed certificates, but if a Public Key Infrastructure (PKI) solution is available, those encryption/authentication certificates can be employed.

Access to the TLS syslog can be restricted to selected sources as applicable.

Some SaaS solutions use TLS-enabled syslog as the sole logging option offered to their customers.

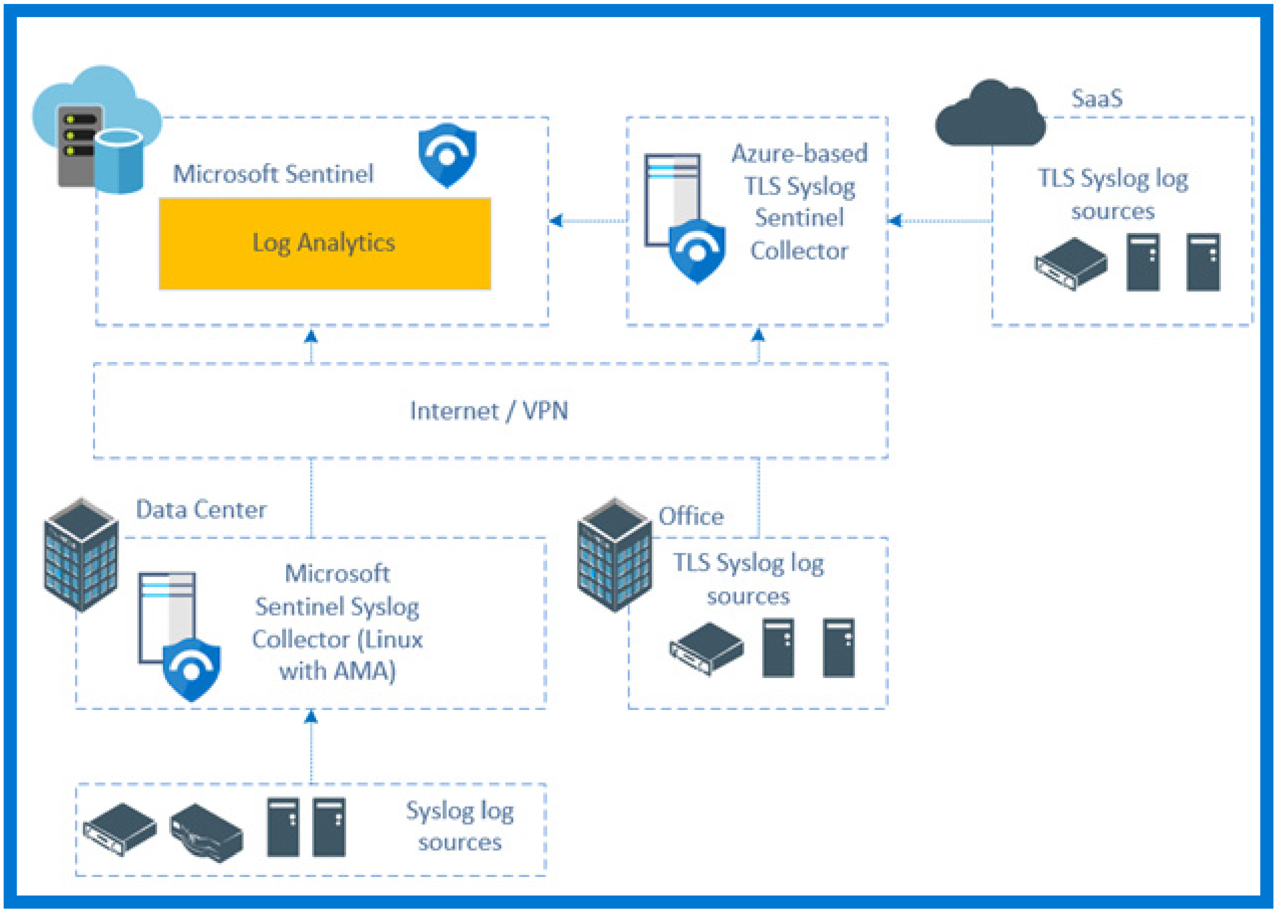

Fig. 6. Typical Microsoft Sentinel syslog collector deployment configurations

The hardware requirements for the syslog collector depend on the expected logging volume. Ingestion rate limits depend on several factors, such as the type of log source, size of a typical log entry, internet connectivity, protocol used (regular UDP syslog, TCP syslog, TLS syslog) and others. Based on our tests with a range of typical logs, such a mix of firewalls and LAN devices, a single collector using 2 CPUs, 8 GB of RAM and 50 GB of local storage can handle up to 6,000 EPS (peak value).

In most situations, the syslog collector will simply receive the logs and forward them to Microsoft Sentinel. If the intention is to also keep a copy of the raw logs on the syslog server as an offline queue, log backup, or for long-term retention, then required disk space must be provisioned.

The syslog collector is used to receive both plain syslog logs and CEF. If that is the case, pay special attention that the same facility is not used by both CEF and non-CEF log sources.

In Our Experience,

Some considerations, based on our experience deploying syslog collectors:

- Coordinate with the team to ensure that a technical resource is available to help configure the syslog collector. That may require personnel that have not previously been engaged in the Microsoft Sentinel project.
- Be sure to provision access to the syslog collector, especially if Microsoft Sentinel is deployed by an external consultant.
- Try to avoid commingling existing syslog collectors to minimize potential ingest or parsing issues.
- Architect to your risk tolerance in terms of syslog collector availability. Determine potential impacts to log collection if a network outage occurs. If applicable, load balancers can be deployed to ensure fail-over/high-availability for the syslog collector. If DevOps resources are available, the collector can be containerized.
- Deploy your syslog collector on a locally common Linux distribution. If there is none, this may be a good opportunity to help introduce a standard.
- In our experience, a number of log source types do not adhere to syslog standard (i.e. RFC 5424: The Syslog Protocol) or their logging options have limited configuration options. These may require additional steps to properly ingest and parse data into Microsoft Sentinel. For example, Cisco Meraki and Firepower devices logging formats can be challenging, depending on version.
- The default message size in some syslog implementations is 2 kb, leading to some messages being truncated. By increasing the default size, the truncation can be avoided.
- IP addresses to computer names reverse DNS resolution performed by the syslog server requires a properly configured DNS server on the syslog collector.
- Monitor the ingestion rate (aggregate the Syslog and CommonSecurityLog tables) and ensure that log entries are not dropped due to lack of resources (a time chart graph with consistent peaks around 6000 EPS might indicate that).
- A common cause of syslog collector failures is to enable the local collections of the logs that are intended for Microsoft Sentinel. Unless the intent is to keep a local copy of logs (and this may require significant local storage and log rotation/compression), ensure that the syslog server is not configured to save incoming logs locally.

Example KQL query for EPS monitoring:

```
union (Syslog),(CommonSecurityLog)
| where TimeGenerated > ago(24h)
| summarize count() by bin(TimeGenerated,1h)
| extend EPS = count_/3600
| project-away count_
```

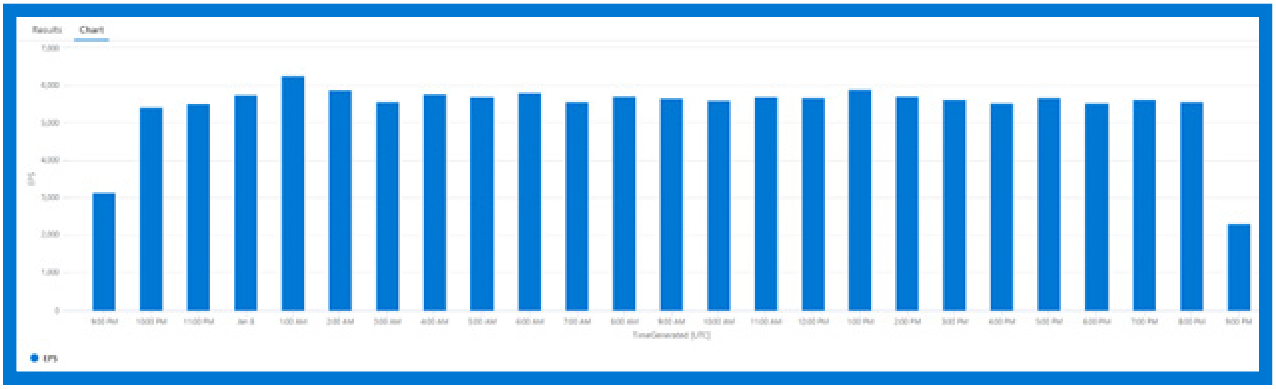

Fig 6.1 Example of potential syslog collector at maximum capacity

Case Study: Healthcare Organization in Canada

Health Inc., a healthcare provider in Canada with 3,500 employees, recently migrated their SIEM to the Microsoft Sentinel platform. The company is running Checkpoint firewalls for perimeter security, as well for protecting internal on-premises data centers. Firewall logs were identified as the source type generating the most ingest cost, based on average volume of data ingested daily in Microsoft Sentinel.

A corrective action was required to filter out the data not directly useful for security detection or compliance requirements. The Health Inc. team decided to use schema transformation (DCR) on the CommonSecurityEvents Log Analytics table to filter out all firewall deny events and ingest only specific fields from each firewall and VPN event log.

The overall log consumption for Checkpoint firewall was reduced by 40% from the initial ingestion volume with no substantial impact to detection quality.

Microsoft Sentinel Automation Playbooks

Based on Azure Logic Apps connectors, a Microsoft Sentinel automation playbook can be used to retrieve logs from various sources using Logic App connectors and actions. Providing that the log source offers a log query/collection capability via API, an HTTP request using a REST API connection to the log source interface is a good example of this type of scenario. Microsoft Sentinel playbooks should be used just for low volume/ log complexity log sources, as they are not designed to perform large data transfers. For additional security, the playbooks can retrieve authentication secrets from Azure Key Vault or authenticate using a managed security identity (when supported).

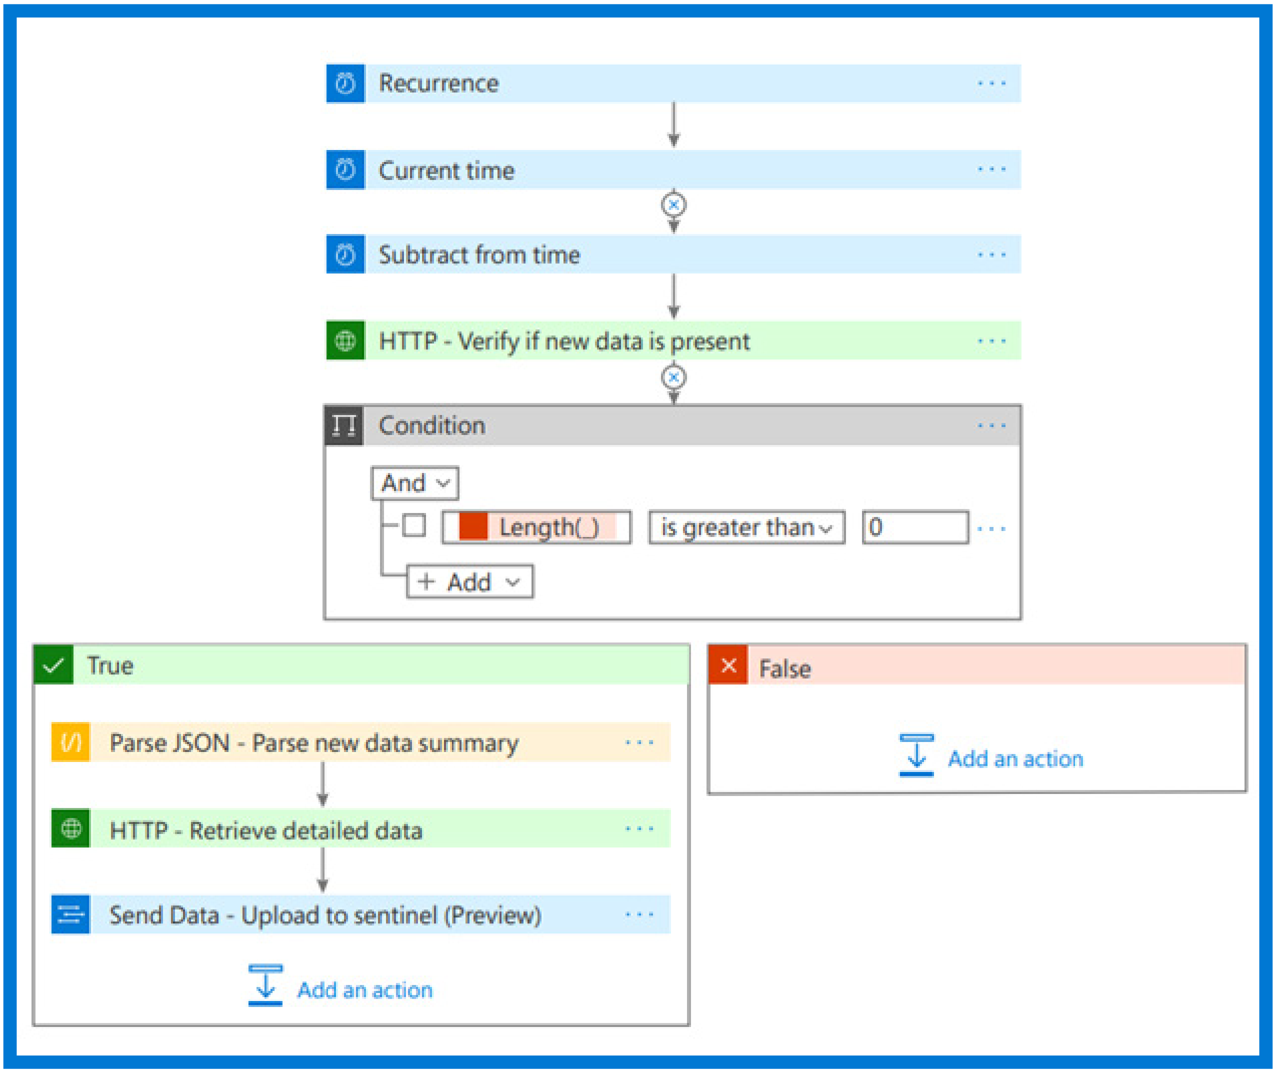
Recurrence
Current time
Subtract from time
HTTP - Verify if new data is present
Condition
And
Length()
is greater than
0
+ Add
True
False
Parse JSON - Parse new data summary
HTTP - Retrieve detailed data
Send Data - Upload to sentinel (Preview)
Add an action
Add an action

Fig. 7. Example of log ingestion via Microsoft Sentinel playbook–retrieval of Office 365 workloads not ingested via Office 365 Data Connector (e.g., DLP, Power BI)

In Our Experience,

Our recommendations on creating and deploying playbook automations in Logic Apps:

- Plan for the acceptance requirements of Microsoft Sentinel playbooks automation. Automated actions can be high-impact so understanding risk tolerances is important.
- Monitor playbook runs and plan to alert if issues are encountered. Therefore, the problem can be quickly remediated versus discovered too late. When configured to log diagnostics, all playbook activities are being recorded in the AzureActivity table.
- Consider any third-party tools involved and how they may work with Microsoft Sentinel playbooks. You may need to consult the third parties for details about their API or other connectivity options.
- Create and maintain API connections across various Logic App connectors. For example, consider the use of managed identity or SPNs with secrets stored in key vaults versus individual service accounts. Not all connectors support managed identity.

Azure Function Apps

As a low-footprint and relatively inexpensive resource, Azure Function Apps is one of the most stable and performant log ingestion methods. Functions apps provide the full capabilities of .Net, Python, PowerShell. Recently, Node.js programming runtimes can be used to perform a wide range of log ingestion tasks, including but not limited to log retrieval via REST APIs, pagination, filtering, parsing, and enrichment of data. Azure Function Apps require more advanced programming capabilities.

Many built-in Microsoft Sentinel connectors rely on the deployment of function apps, and typically they are developed by the log source vendor working in collaboration with Microsoft.

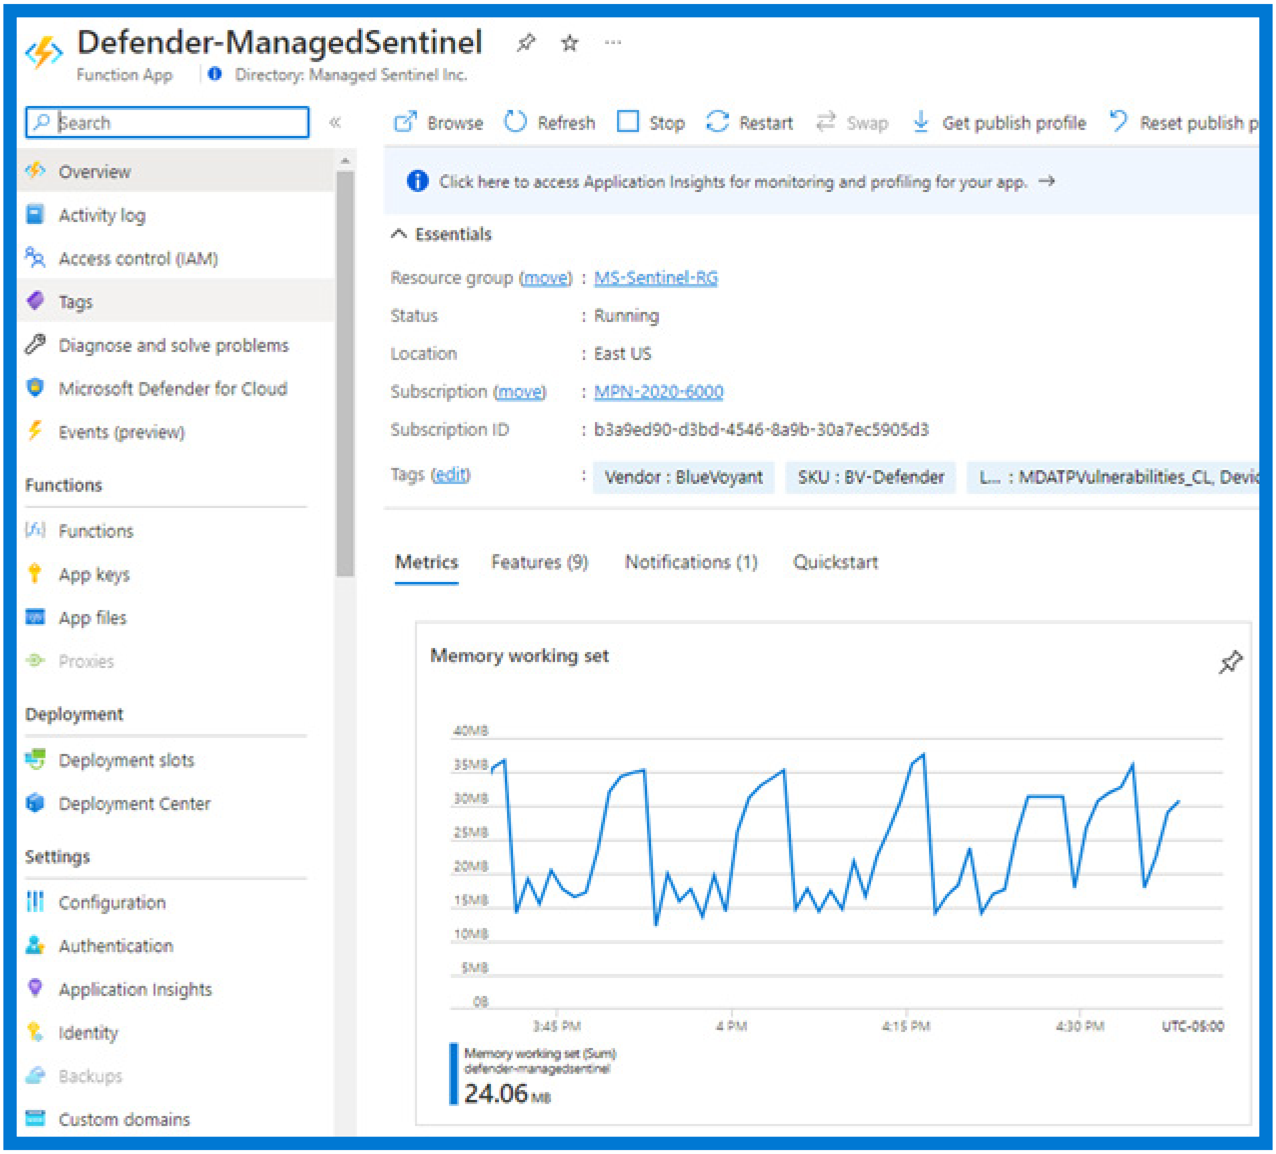

Click here to access Application Insights for monitoring and profiling for your app. →

Resource group (move) : MS-Sentinel-RG
Status : Running
Location : East US
Subscription (move) : MPN-2020-6000
Subscription ID : b3a9ed90-d3bd-4546-8a9b-30a7ec5905d3
Tags (edit) : Vendor : BlueVoyant SKU : BV-Defender L... : MDATPVulnerabilities_CL, Devic

Memory working set (Sum)
defender-managedsentinel
24.06 MB

Fig. 8. Azure Function App retrieving custom data from Microsoft Defender for Endpoint

In Our Experience,

In developing Azure Function Apps, we suggest the following considerations be addressed when deploying:

- Leverage members of the team who have programming expertise because deploying Azure Function Apps requires programming skills.
- Gain a good understanding of log source interface capabilities (i.e., REST API) before deploying. Vendor support (the developer of the 3rd party REST API) may be required for troubleshooting or to provide documentation. In our experience, many of the 3rd party REST APIs are not mature enough to provide stable performance or even cover all the advertised features.
- Deploy only one compute platform per resource group (i.e., Python versus .Net).
- Implement continuous monitoring of function apps health (workbook or alert rules). Many data connectors using function apps use API credentials that may need to be renewed from time-to-time so their expiration has to be monitored (this can be achieved indirectly by monitoring the flow of log data imported by the data connector).
- Avoid modifying Azure functions apps deployed as part of an out-of-the box data connector. These are typically tied to a Microsoft GitHub repository and any modifications might be overwritten when an update takes place.

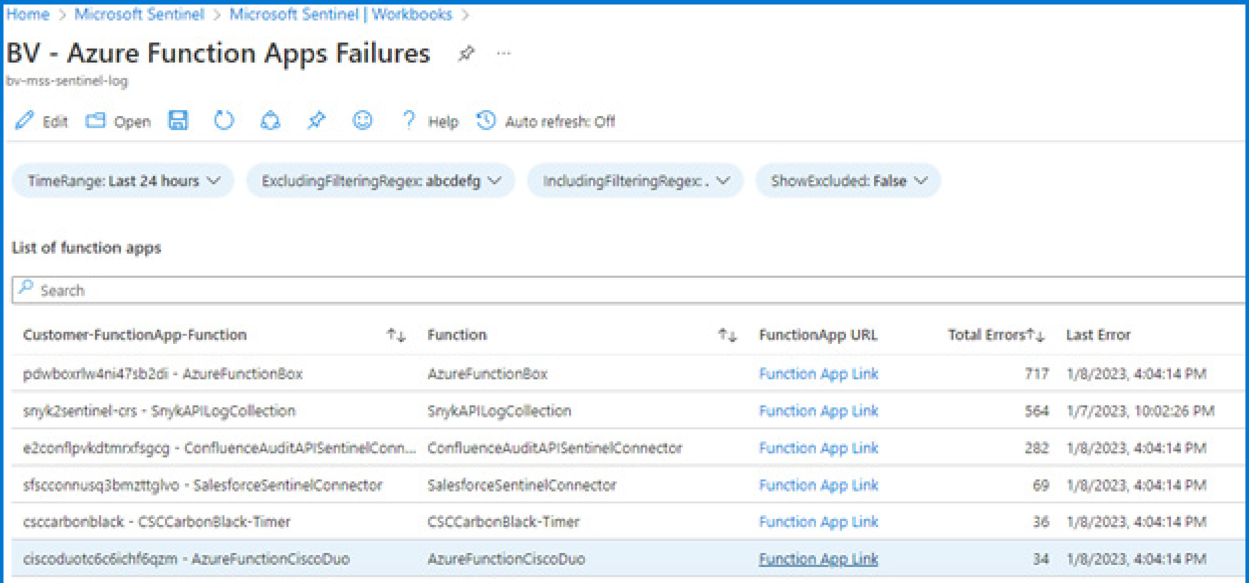
Home > Microsoft Sentinel > Microsoft Sentinel | Workbooks >

BV - Azure Function Apps Failures

bv-mss-sentinel-log
Edit
Open
Help
Auto refresh: Off

TimeRange: Last 24 hours

ExcludingFilteringRegex: abcdefg

IncludingFilteringRegex: .

ShowExcluded: False

List of function apps

Customer-FunctionApp-Function	Function	FunctionApp URL	Total Errors	Last Error
pdwboxrlw4ni47sb2di - AzureFunctionBox	AzureFunctionBox	Function App Link	717	1/8/2023, 4:04:14 PM
snyk2sentinel-crs - SnykAPILogCollection	SnykAPILogCollection	Function App Link	564	1/7/2023, 10:02:26 PM
e2conflpvkdtmrxfsgcg - ConfluenceAuditAPISentinelConn...	ConfluenceAuditAPISentinelConnector	Function App Link	282	1/8/2023, 4:04:14 PM
sfscconnusq3bmzttglvo - SalesforceSentinelConnector	SalesforceSentinelConnector	Function App Link	69	1/8/2023, 4:04:14 PM
csccarbonblack - CSCCarbonBlack-Timer	CSCCarbonBlack-Timer	Function App Link	36	1/8/2023, 4:04:14 PM
ciscoduotc6c6ichf6qzm - AzureFunctionCiscoDuo	AzureFunctionCiscoDuo	Function App Link	34	1/8/2023, 4:04:14 PM

Fig. 8.1 – Sample Azure Function App health monitoring workbook

Case Study: Engineering Company

CMEng Inc. is an engineering company located in Canada with customers worldwide. Currently, CMEng is running a hybrid infrastructure with workloads in Azure, AWS, and Oracle Cloud Infrastructure (OCI) as well on-premises. Also, CMEng is using several SaaS applications for various business lines.

CMEng decided to migrate to Microsoft Sentinel as its main security monitoring tool. Onboarding all log sources in Microsoft Sentinel would require the development of several custom data connectors. Some of the log sources selected for custom data connector development were OCI, Duo multi-factor authentication (MFA), and Cloudflare WAF. None of these data connectors existed within out-of-the-box Microsoft Sentinel data connectors or in easily accessible community forums.

The data connectors were developed based on vendor REST APIs using Azure Function App, and application data was ingested in Log Analytics custom tables (_CL). The effort required for this project was around 40 hours, completed by a senior cloud security developer. Access to team members with programming expertise can broaden the options for data ingest significantly.

Third-party and Vendor-provided Log Retrieval – Log Ingestion Tools

Depending on their specific design and maturity level, some SaaS platforms may offer specialized utilities that can be run at scheduled intervals to retrieve activity logs. Those could be scripts (such as PowerShell and Python) or executables. Once those scripts are configured according to the vendor's instructions, the MMA/AMA can be used to monitor the logs' location and update new data as it is downloaded by the log retrieval script.

For more advanced processing/filtering capabilities, tools such as Logstash or Cribl can be used to collect the local logs, process them, and upload them to Microsoft Sentinel using the Microsoft Log Analytics Logstash plugin.

A number of SaaS vendors offer the option to configure the log ingestions straight from their platform. Typically, this is achieved by providing a set of Azure Log Analytics Workspace IDs and keys that are used by the vendor to push their logs through the Azure HTTP Data Collection API.

For certain solutions, log collection tools are provided by third-party developers not associated with the vendor itself. Most of these tools are available in free-to-use GitHub repositories.

In Our Experience,

- Third-party tools such as Logstash can provide very powerful additional capabilities that include inputs for a variety of platforms (i.e. Kafka), filters for data parsing and enrichment, and a Microsoft-supported output for Azure Log Analytics Workspace.
- Compared with the standard Microsoft Sentinel Syslog Collector, third-party tools may require additional expertise, and the collectors need more resources (RAM, CPU)
- Implementing a third-party tool comes with the risk of moving away from Microsoft-supported log ingestion methods so support has to be obtained from other sources. Licensing a third-party tool may incur additional capital and operational costs.

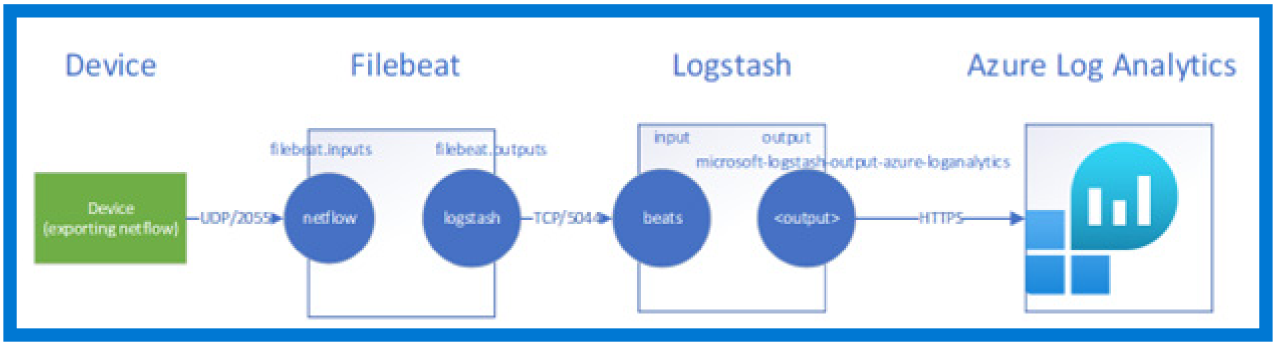

Fig. 9. Example of NetFlow log ingestion via Filebeat and Logstash

Additional Resources:

Use Logstash to stream logs with pipeline transformations via DCR-based API – Microsoft Docs

Integrating Cribl Stream with Azure Sentinel – Third party doc

Case Study: US-based Mortgage Corporation

Mortgage Inc. has recently migrated from an on-premises SIEM to the Microsoft Sentinel platform. Due to cost constraints, the Mortgage Inc. CISO instructed the team to ingest only data that can be used for security detections into the Microsoft Sentinel platform.

Windows events were identified as being the second largest log source based on the volume. Therefore a pragmatic approach was taken to ensure only Windows event IDs matching the Microsoft Sentinel alert rules KQL logic will be collected and stored in Log Analytics. The team decided to use Data Collection Rules (DCR) in Azure Monitor in the security events table. Based on the KQL query that was written with “where EventID !=” statement, the team was able to eliminate 122 Windows EventIDs and ingest only 21 eventIDs that match the 37 alert rules running based on this table.

Doing this resulted in the volume of logs related to Windows security events being reduced by 56% without substantial impact to detection efficacy.

Automation Playbooks

The main engine behind the Microsoft Sentinel automation capability is Azure Logic Apps. Logic Apps were released as general availability in 2016 as part of Microsoft Sentinel SOAR capabilities. They are a proven technology used across the Azure ecosystem.

A playbook is initiated by a trigger and performs its functionality through a variety of connectors using Microsoft Flow. Microsoft Flows allows for transfer of data between connectors and a wide range of processing functions. In Microsoft Sentinel, the typical trigger is “When a response to a Microsoft Sentinel incident is triggered,” and a Logic App connector allows the collection of data by a Microsoft Sentinel incident that can be passed to various connectors to perform tasks based on the desired outcome of the security orchestration/automation scenario.

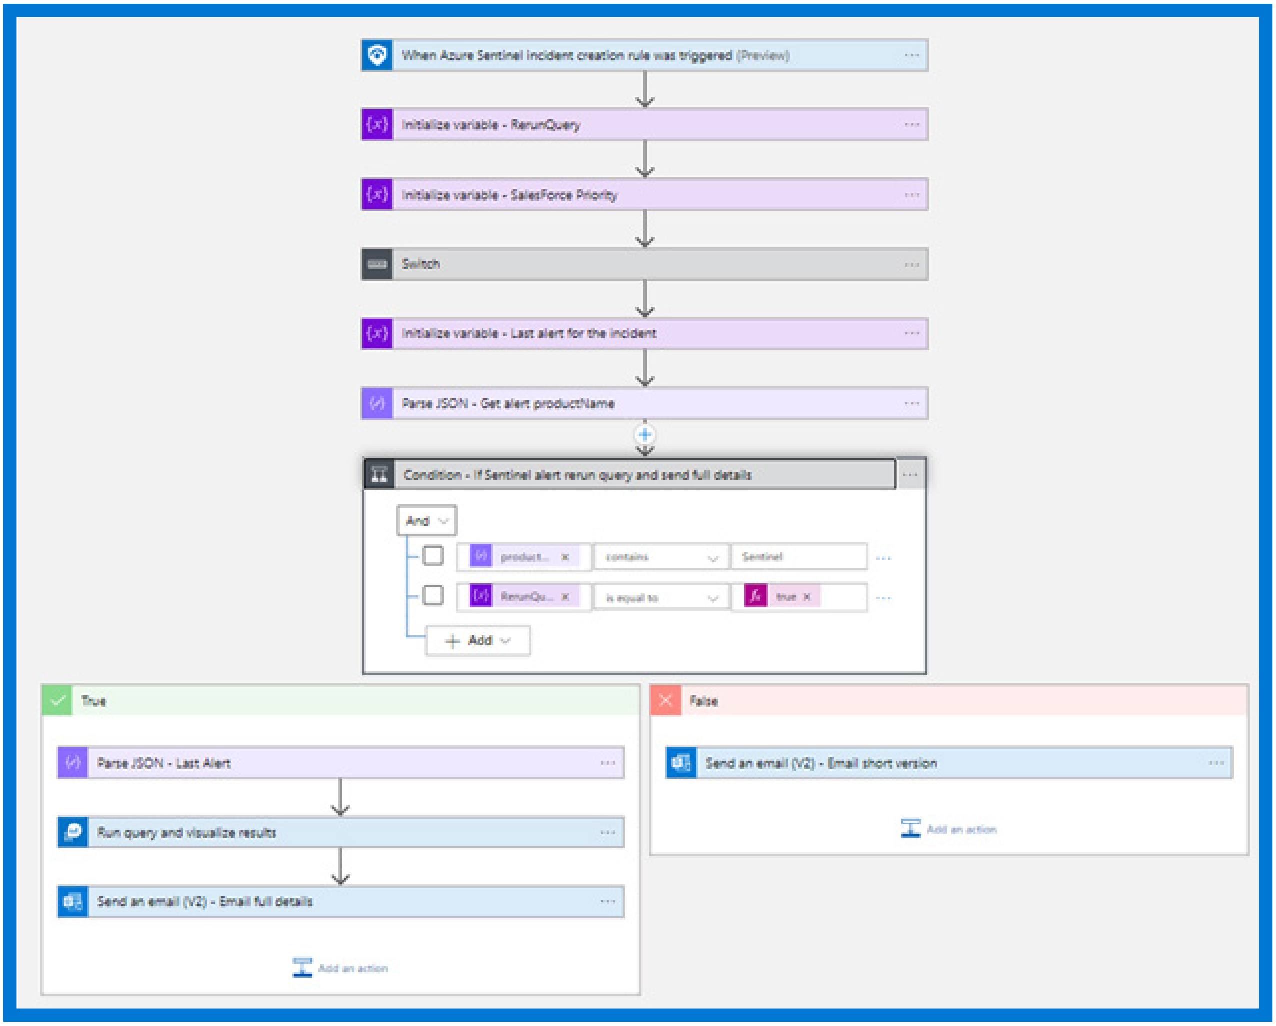
When Azure Sentinel incident creation rule was triggered (Preview)
Initialize variable - RerunQuery
Initialize variable - SalesForce Priority
Switch
Initialize variable - Last alert for the incident
Parse JSON - Get alert productName
Condition - If Sentinel alert rerun query and send full details
And
product... contains Sentinel
RerunQu... is equal to true
Add
True
Parse JSON - Last Alert
Run query and visualize results
Send an email (V2) - Email full details
Add an action
False
Send an email (V2) - Email short version
Add an action

Fig. 10. Sample Microsoft Sentinel playbook with a Microsoft Sentinel Incident trigger and connectors for sending alert details by email

Microsoft Sentinel playbooks can perform very advanced tasks with activities branching based on criteria identified in alerts or data retrieved using indicators collected in alerts. The full processing can be performed within the Microsoft Sentinel environment. It is limited only by the automation capabilities provided by the third-party security controls required to provide information or perform additional tasks.

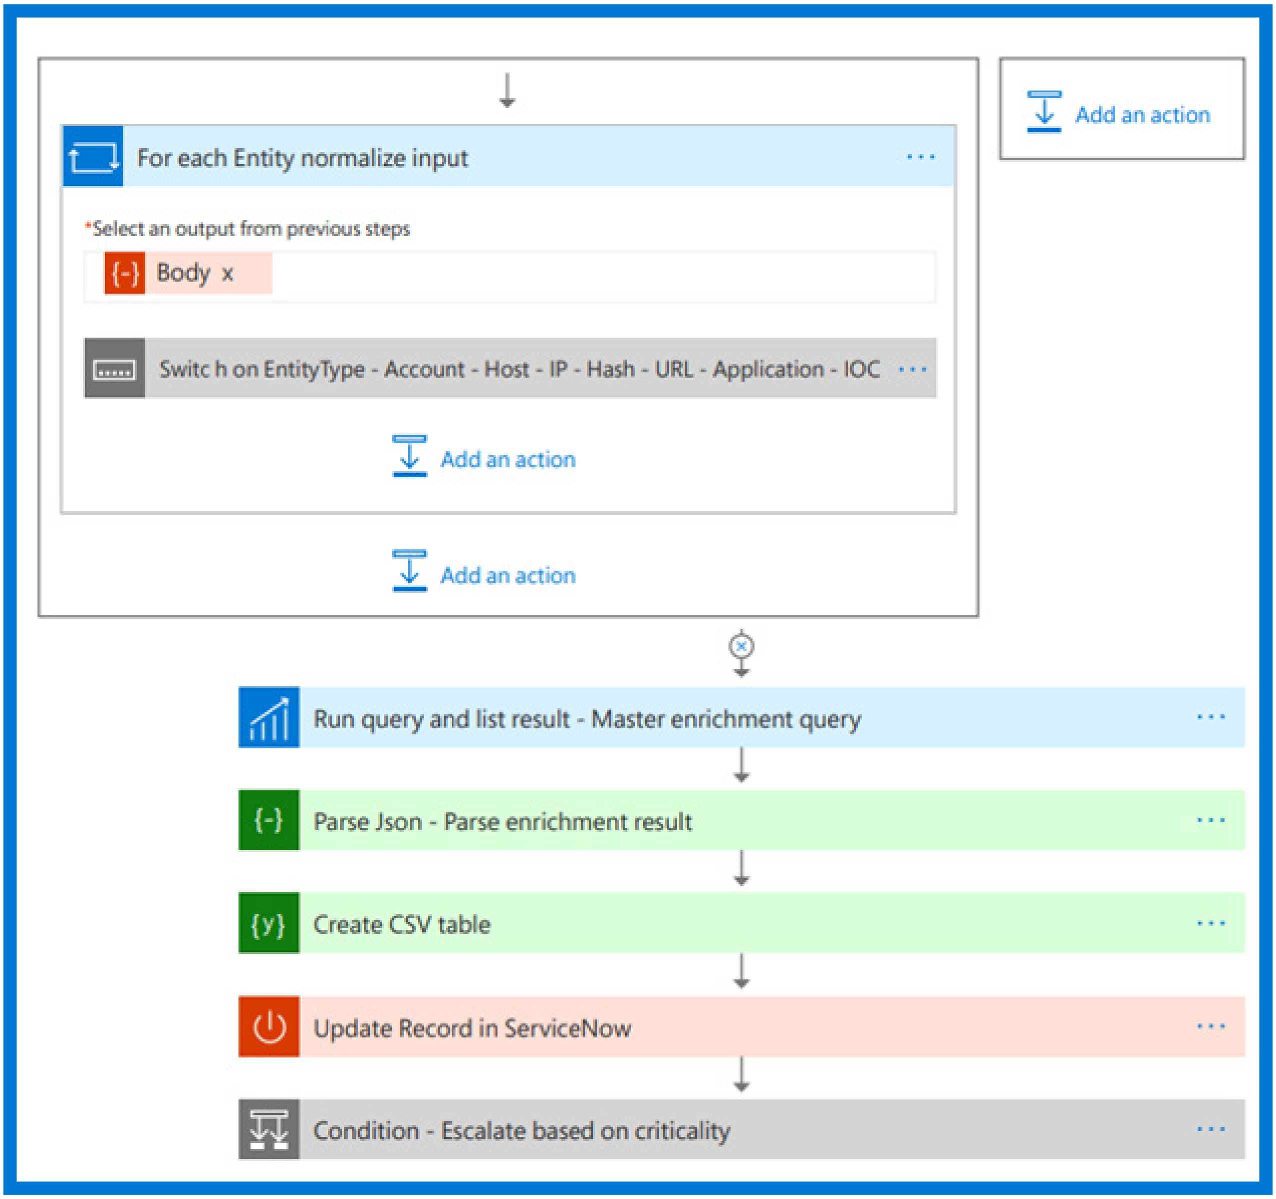
For each Entity normalize input
*Select an output from previous steps
Body x
Switc h on EntityType - Account - Host - IP - Hash - URL - Application - IOC
Add an action
Add an action
Add an action
Run query and list result - Master enrichment query
Parse Json - Parse enrichment result
Create CSV table
Update Record in ServiceNow
Condition - Escalate based on criticality

Fig. 11. Alert data enrichment playbook with integration with ServiceNow

The most used Logic App connectors in Microsoft Sentinel are:

- Occurrence triggers (running the playbook at regular intervals)
- HTTP request. To retrieve data from third-party log sources via REST API
- Azure Monitor. Run KQL queries to retrieve or visualize data
- Data Operations. Parse JSON, Compose, Convert to CSV
- Controls. Conditions (if ... else), For each, Switch
- Variables. Placeholders for data used during alert processing
- Send Data. To upload data into a Microsoft Sentinel table for further use Microsoft Sentinel
- Retrieve/Update incidents
- Office 365. Send emails
- Notification controls. PagerDuty, Signl4, Twilio
- ITS tools connectors. ServiceNow, Freshservice, Freshdesk

The Automation blade in the Microsoft Sentinel interface also provides access to a large number of playbook templates that can be deployed with a click of the mouse. For specific log sources, additional playbooks are available when the solution is deployed through the Content Hub.

Automation Rules

The automation rules allow for a more intuitive construction of SOAR activities. That provides the ability to build combinations of playbook runs and incident updates (severity, ownership, status, tagging) to match the required output.

An automation rule can be applied using a combination of filtering criteria, such as alert names, description, severities, or type of entities identified in the incident created by the alert rule. Multiple playbooks can be applied in succession, allowing for practically unlimited possibilities on how to automate the response for different types of incidents.

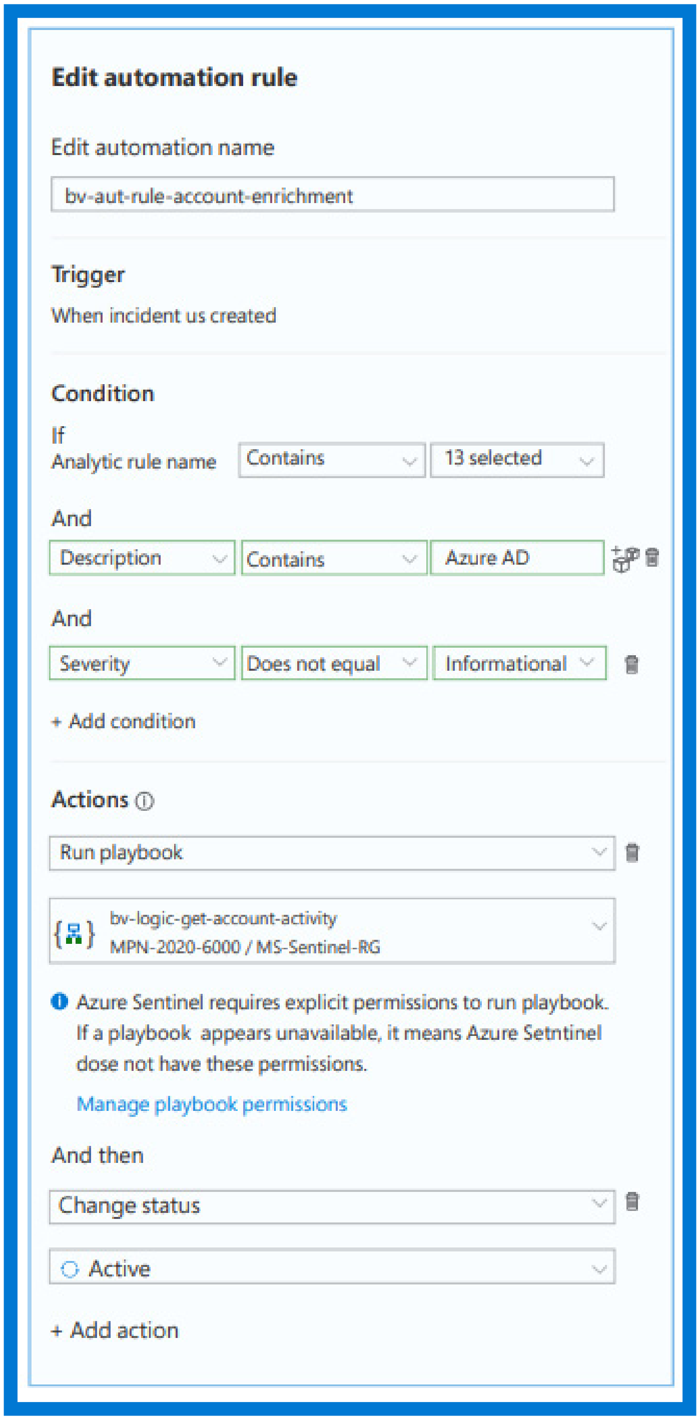

Edit automation rule

Edit automation name

bv-aut-rule-account-enrichment

Trigger

When incident us created

Condition

If
Analytic rule name | Contains | 13 selected

And
Description | Contains | Azure AD

And
Severity | Does not equal | Informational

+ Add condition

Actions ⓘ

Run playbook

bv-logic-get-account-activity
MPN-2020-6000 / MS-Sentinel-RG

Azure Sentinel requires explicit permissions to run playbook. If a playbook appears unavailable, it means Azure Setntinel dose not have these permissions.

Manage playbook permissions

And then

Change status

○ Active

+ Add action

Fig. 12 Sample automation rule performing enrichment of data for the captured alert entities

The playbook used in the sample automation rule above ("bv-logic-get-account-activity") will retrieve the list of account entities from the alert, run a summary of their activities in Azure AD SigninLogs, and post the results as a comment in the incident, allowing the security analyst to have a snapshot of the suspicious account activities without leaving the incident page.

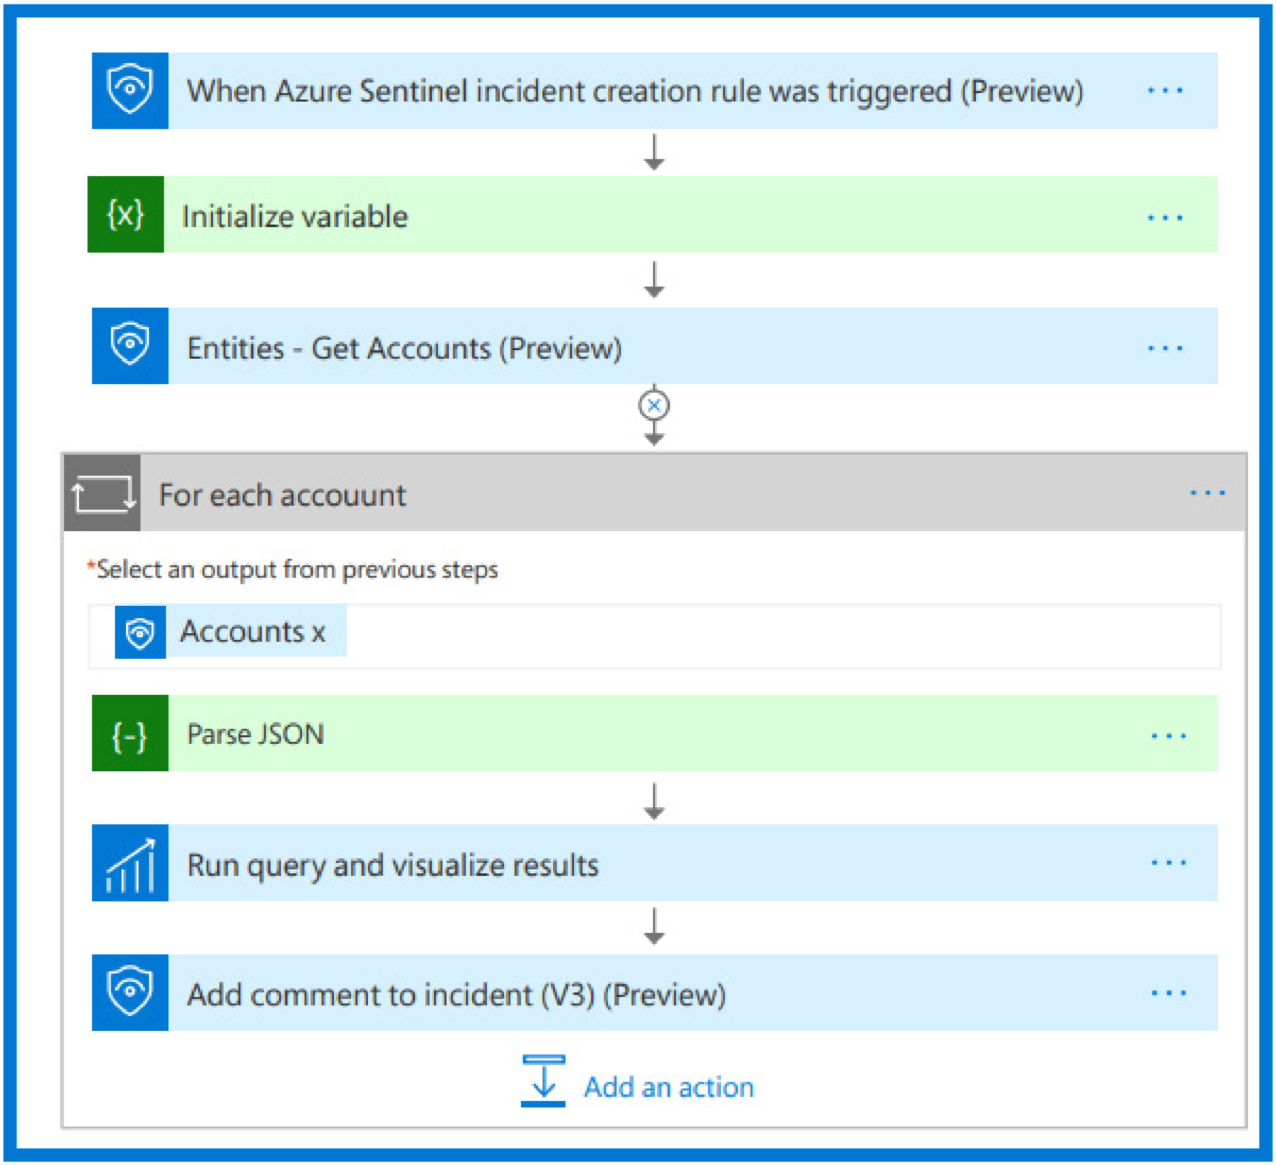
When Azure Sentinel incident creation rule was triggered (Preview)
Initialize variable
Entities - Get Accounts (Preview)
For each accouunt
*Select an output from previous steps
Accounts x
Parse JSON
Run query and visualize results
Add comment to incident (V3) (Preview)
Add an action

Fig. 13. Enrichment playbook

In Our Experience,

- Automation rules allow for quick tuning of the incident flow. That allows for the creation of customized actions to close false positives. An automation rule can be created from the incident details page, with the incident metadata already populated in the automation rule filtering conditions.
- The ease of configuration of automation rules may lead to a large number of rules implemented. That may be difficult to manage so they should not be abused.
- When configured to run multiple automation playbooks, they are executed in parallel so the playbooks logic has to be carefully reviewed to accommodate for such parallel runs.

Additional Resources:

Playbooks & Watchlists Part 1: Inform the subscription owner - Microsoft Tech Community

Playbooks & Watchlists Part 2: Automate incident response for Deny-list/Allow-list - Microsoft Tech Community

Deploying Workbooks

The Microsoft Sentinel workbooks provide a wide range of data visualization based on KQL queries and integration with additional Microsoft resources (via REST APIs). Over 135 workbook templates are provided for the typical log sources such as Azure Active Directory, Office 365, Windows Active Directory, and third-party log sources (e.g., firewalls, SaaS).

Workbooks provide several visualization controls (e.g., bar, pie, area, time charts), conditional formatting, and several other features commonly found in analytical platforms.

Workbooks can retrieve data from multiple sources, allowing for complex integration with various Microsoft services: Azure Log Analytics Workspace, Microsoft Graph, Azure Data Explorer, Azure Resource Manager, and many other sources. The existing templates that ship with Microsoft Sentinel can be reused for new workbooks tailored for customer-specific requirements.

In Our Experience,

- The top required custom workbook is for key performance indicators (KPIs) that would allow executives to make decisions around cybersecurity governance. Defining such KPIs and extracting them from the logs can be challenging. That is because most KPIs rely on measuring the efficiency of a processes-tools-people combination rather than log data provided by security controls. Through regular review and feedback from the consumers of reports, workbooks can become very effective tools.
- Low fidelity alerts can be captured in workbooks for regular review and avoid the cluttering of the incident management interface.

One of the most used workbooks is Security Operations Efficiency:

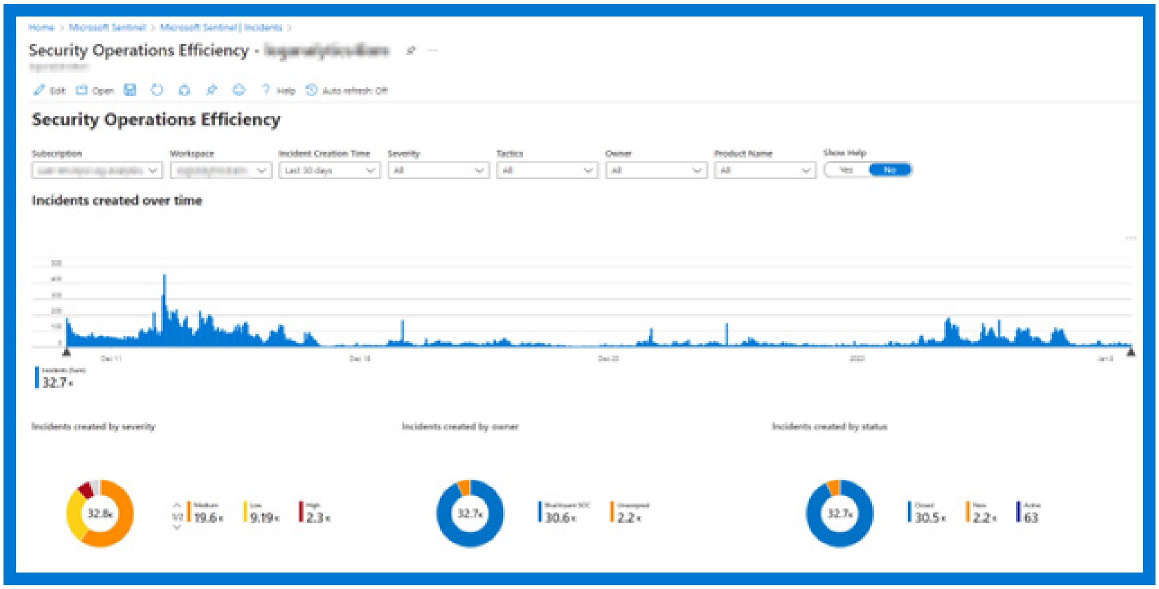
Home > Microsoft Sentinel > Microsoft Sentinel | Incidents >
Edit
Open
Help
Auto refresh: Off

Security Operations Efficiency -

Security Operations Efficiency

Subscription:
Workspace:
Incident Creation Time: Last 30 days
Severity: All
Tactics: All
Owner: All
Product Name: All
Show Help: Yes [✓ No]

Incidents created over time
Dec 11
Dec 18
Dec 25
2023
Jan 8
Incidents (Sum)
32.7K
Incidents created by severity
32.8K
Medium
19.6K
Low
9.19K
High
2.3K
Incidents created by owner
32.7K
30.6K
Unassigned
2.2K
Incidents created by status
32.7K
Closed
30.5K
New
2.2K
Active
63

Fig. 15. Security Operations Efficiency Workbook

Workbooks can also be used to convey complex information in a more visual and intuitive manner, without a need to extract data from Microsoft Sentinel tables. An example of such a workbook is the SOC Process Framework Workbook developed by Rin Ure, that contains a wide range of information on the processes required for a SOC development.

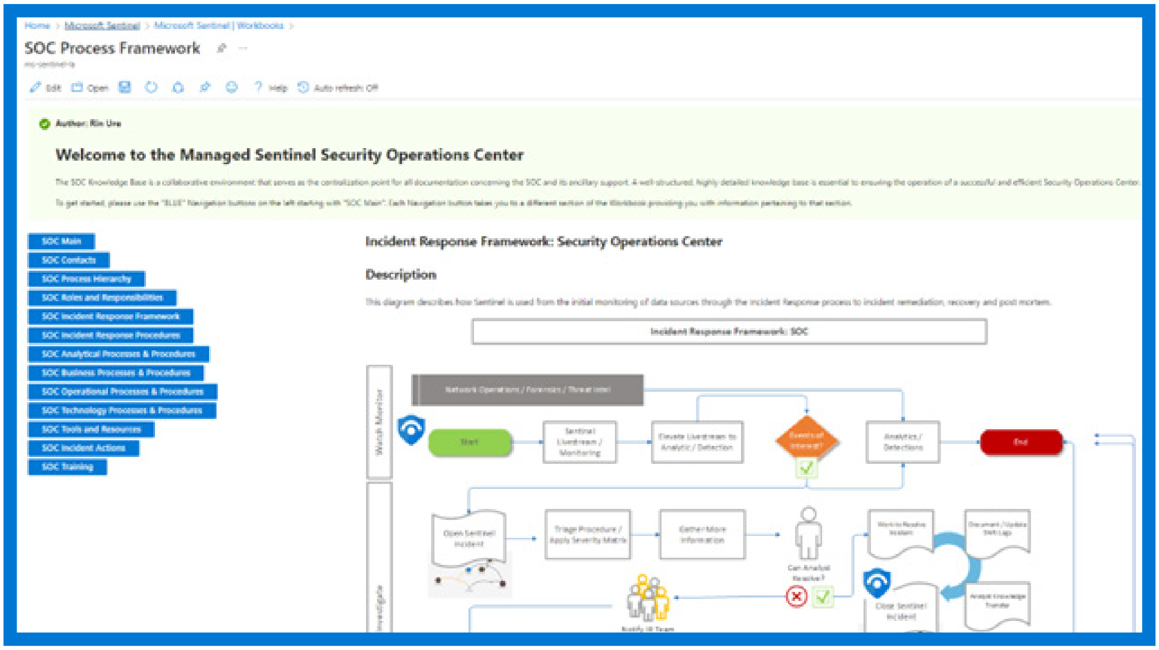
Home > Microsoft Sentinel > Microsoft Sentinel | Workbooks >
SOC Process Framework
Edit
Open
Help
Auto refresh: Off
SOC Main
SOC Contacts
SOC Process Hierarchy
SOC Roles and Responsibilities
SOC Incident Response Framework
SOC Incident Response Procedures
SOC Analytical Processes & Procedures
SOC Business Processes & Procedures
SOC Operational Processes & Procedures
SOC Technology Processes & Procedures
SOC Tools and Resources
SOC Incident Actions
SOC Training
Incident Response Framework: SOC
Watch Monitor
Network Operations / Forensics / Threat Intel
Start
Sentinel Livestream / Monitoring
Elevate Livestream to Analytic / Detection
Events of Interest?
Analytics / Detections
End
Open Sentinel Incident
Triage Procedure / Apply Severity Matrix
Gather More Information
Can Analyst Resolve?
Work to Resolve Incident
Document / Update Shift Logs
Close Sentinel Incident
Analyst Knowledge Transfer
Notify IR Team
Investigate

Author: Rin Ure

Welcome to the Managed Sentinel Security Operations Center

The SOC Knowledge Base is a collaborative environment that serves as the centralization point for all documentation concerning the SOC and its ancillary support. A well-structured, highly detailed knowledge base is essential to ensuring the operation of a successful and efficient Security Operations Center.

To get started, please use the "BLUE" Navigation buttons on the left starting with "SOC Main". Each Navigation button takes you to a different section of the Workbook providing you with information pertaining to that section.

Incident Response Framework: Security Operations Center

Description

This diagram describes how Sentinel is used from the initial monitoring of data sources through the Incident Response process to incident remediation, recovery and post mortem.

Fig. 16. The SOC Process Framework workbook

Case Study: Large International Manufacturing Company

Manufacture Inc. has over 12,000 employees located in 5 countries. Recently, Manufacturing Inc. migrated their SIEM from QRadar to Microsoft Sentinel.

The QRadar platform was used by the SOC team for more than 5 years. That led the migration team to request migration of all content from the legacy SIEM to Microsoft Sentinel using a like-for-like approach. The scope was intended to cover use cases and SOAR playbooks and dashboards. Providing a familiar user experience for their SOC team was considered a critical success criteria of the migration project.

The SOC team identified a list of 4 dashboards with various KPIs and metrics used on a daily basis that were required to be moved into Microsoft Sentinel. While some of the out of the box Workbooks existed in Microsoft Sentinel, the Manufacturing Inc. SOC team had specific requirements in terms of content and look and feel that had to be built in Microsoft Sentinel. For example, KPIs metrics presented in a Microsoft Sentinel workbook were required for case history and management of specific manufacturing floor application metrics.

Specialized customization is fully available within Microsoft Sentinel; however, replicating specific user interfaces from other SIEM tools can be complex and costly. In many cases, a rationalization of required functionality in the SIEM combined with additional Analyst training can allow for the migration to serve as an exercise in technology simplification. The migration effort from legacy SIEM to Microsoft Sentinel was 285 Professional Services hours completed over a period of 8 weeks.

Additional Resources:

[Microsoft Sentinel Workbooks 101 (with sample Workbook) - Microsoft Tech Community]

[Commonly used Microsoft Sentinel workbooks | Microsoft Docs]

[How to use Azure Monitor Workbooks to map Sentinel data - Microsoft Tech Community]

Deploying User and Entity Behavior Analytics

User and Entity Behavior Analytics (UEBA), shown as "Entity Behavior" in the Microsoft Sentinel management interface, represents one of the relatively new technologies provided by SIEM solutions and relies, in most cases, on machine learning capabilities to track the behavior of users and entities such, as hosts and IP addresses, and detect deviations from the expected patterns.

Microsoft Sentinel relies on a range of logs sources such as Azure AD, Windows Active Directory, Office 365, and others to build models around the activity of users and computers and build "insights" around their observed behavior. Microsoft Sentinel builds several baselines using 10, 30, 90, or 180 days of observed behavior, based on the type of log source. Through UEBA, multiple low fidelity signals can be aggregated to build a timeline of events that can be used to better understand the entity behavior and make more informed decisions during further investigations. The Microsoft Sentinel Entity Behavior component tracks not only users, hosts, and IP addresses but also IoT devices and Azure resources.

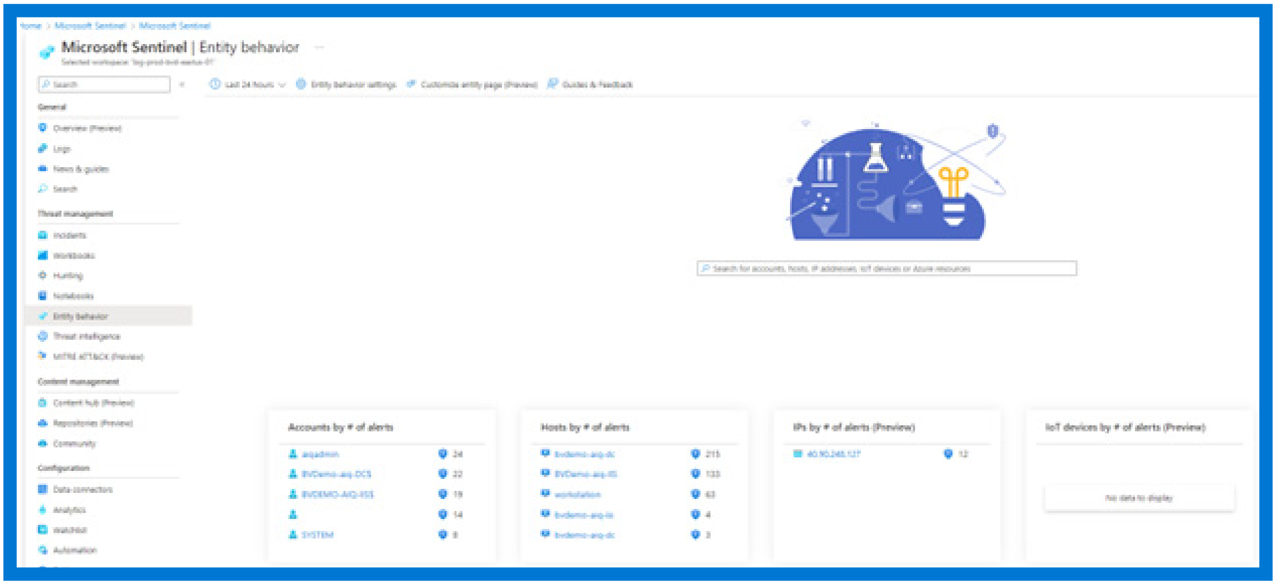

Accounts by # of alerts	
aiqadmin	24
BVDemo-aiq-DCS	22
BVDEMO-AIQ-IIS1	19
	14
SYSTEM	8

Hosts by # of alerts	
bvdemo-aiq-dc	215
BVDemo-aiq-IIS	133
workstation	63
bvdemo-aiq-iis	4
bvdemo-aiq-dc	3

IPs by # of alerts (Preview)	
40.90.248.127	12

IoT devices by # of alerts (Preview)
No data to display

Fig. 17. User entity Microsoft Sentinel Entity Behavior interface

The entity information interface allows for quick pivoting into a Microsoft Sentinel investigation graph for further details on the relation between the entity and the alerts raised by Microsoft Sentinel.

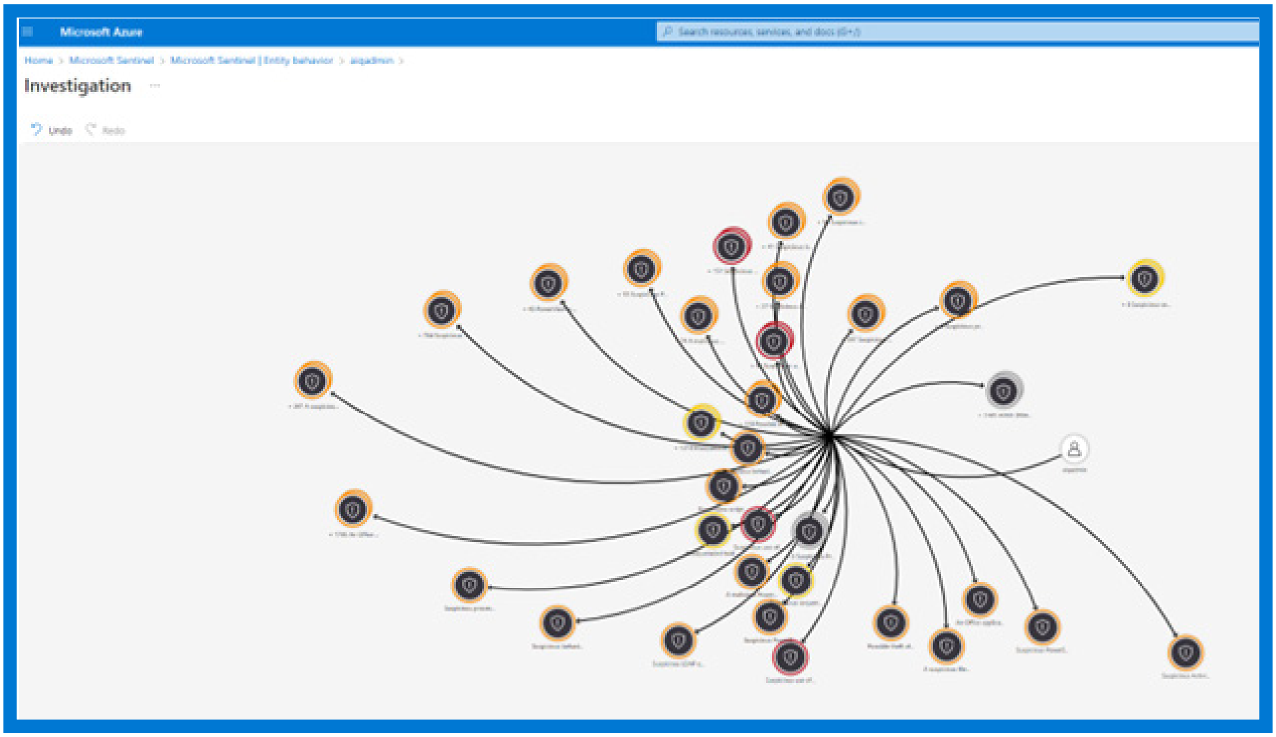
Microsoft Azure
Search resources, services, and docs (G+/)
Home > Microsoft Sentinel > Microsoft Sentinel | Entity behavior > aiqadmin >
Investigation
Undo
Redo

Fig. 18. Microsoft Sentinel investigation graph

The UEBA data is available as Microsoft Sentinel tables for full integration with alert rules and SOAR playbooks. That allows for the enrichment of alert metadata and adjustment of incident severity based on the behavior retrieved from UEBA.

The UEBA component monitors for a range of outliers in the common activities and records them in the Anomalies table. These can be used to trigger incidents or enrich incidents from other detection rules. Examples of monitored activities include account creation and removal, data destructions, failed sign-ins, code exectution and many others. See UEBA Activities for a full list.

In Our Experience,

- Substantial data is compiled by UEBA engines, which are often underutilized by security organizations in Microsoft Sentinel deployments. We have found these detections to be highly valuable and recommend utilizing them in security operations use cases.

Using the MITRE ATT&CK Dashboard

Based on the MITRE ATT&CK knowledgebase, the dashboard available in Microsoft Sentinel provides a mapping between existing alert rules deployed and the various attack techniques described in the MITRE ATT&CK knowledge base of adversary tactics and techniques.

The intent of the visualization is to provide a heatmap of the coverage of MITRE ATT&CK techniques and the detection rules deployed in the specific Microsoft Sentinel instance to help identify areas that are still not covered.

In Our Experience,

- MITRE ATT&CK coverage appears to be very important for many CISOs and other senior managers involved in the cybersecurity monitoring infrastructure.
- While an effective way to visualize various attack vectors and review their details, the chart provided should be used more for researching the type of detections that need to be implemented and the potential log sources. It should be less used for trying to achieve 100% coverage. The reason is that even when a specific technique appears to have several alert rules deployed that cover it, it does not mean that they offer visibility over all the potential uses of that specific attack vector. For example, even if technique T1190, **Exploit Public-Facing Application**, appears to be covered by 20 detection rules, there are potentially a large number of other aspects of the existing infrastructure vulnerable to this particular technique that are not covered by those rules.
- A large number of MITRE ATT&CK techniques are covered by other types of security controls, i.e., EDR solutions like Defender for Endpoint. For this reason, it may not always be practical to identify the true full coverage of the attack vectors just based on the alert rules configured in Microsoft Sentinel.

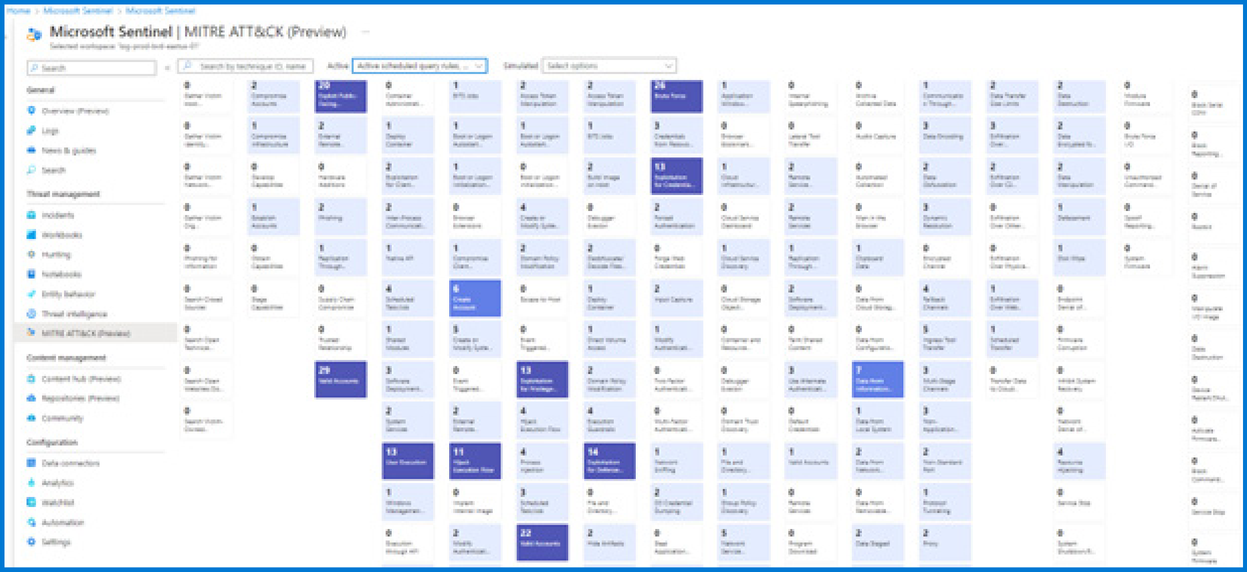

Additional Resources:

Understand security coverage by the MITRE ATT&CK® framework

Deploying Notebooks

Based on Jupyter Notebooks, Microsoft Sentinel Notebooks allow for advanced threat-hunting capabilities, using the data collected by Microsoft Sentinel and the processing capabilities available in multiple programming languages. Those include but are not limited to Python and C#/.Net. Any library, such as those related to threat intelligence enrichment and ML/AI available to the selected programming language, is available to use toward Microsoft Sentinel log data.

Microsoft Sentinel provides the ability to integrate notebooks with Azure Synapse Analytics for big data processing.

Deploying Cyber Threat Intelligence Functionality

Cyber Threat Intelligence (CTI) is available from a wide array of sources. Those can include open- source data feeds, threat intelligence sharing communities, premium curated feeds, and your own security investigations. CTI can be provided as a formal write up about a given threat actor's tactics, techniques, and procedures (TTP), underlying goals or motivations, or specific lists of observed domain names, IP addresses, email addresses, and file hashes—the latter are collectively known as indicators of compromise (IOCs) but also known within Microsoft Sentinel as threat indicator data. CTI can provide valuable contextual information when combined with your own data, helping to speed the time to detect, identify, and triage malicious or anomalous activity.

In Our Experience,

- Notebooks are very powerful and require knowledge of at least one programming language SME and, due to their technical complexity, can have a steep learning curve.
- MSTICPy Python library, developed and maintained by Microsoft, is extensively used to perform a wide number of threat intel enrichment tasks within Notebooks. There are many articles with examples of use for threat-hunting scenarios. (See Additional Resources)
- Notebooks are commonly used for proactive threat-hunting and post-mortem analysis.

Additional Resources:

Use notebooks with Microsoft Sentinel for security hunting | Microsoft Docs

MSTICPy and Jupyter Notebooks in Azure Sentinel - Microsoft Tech Community

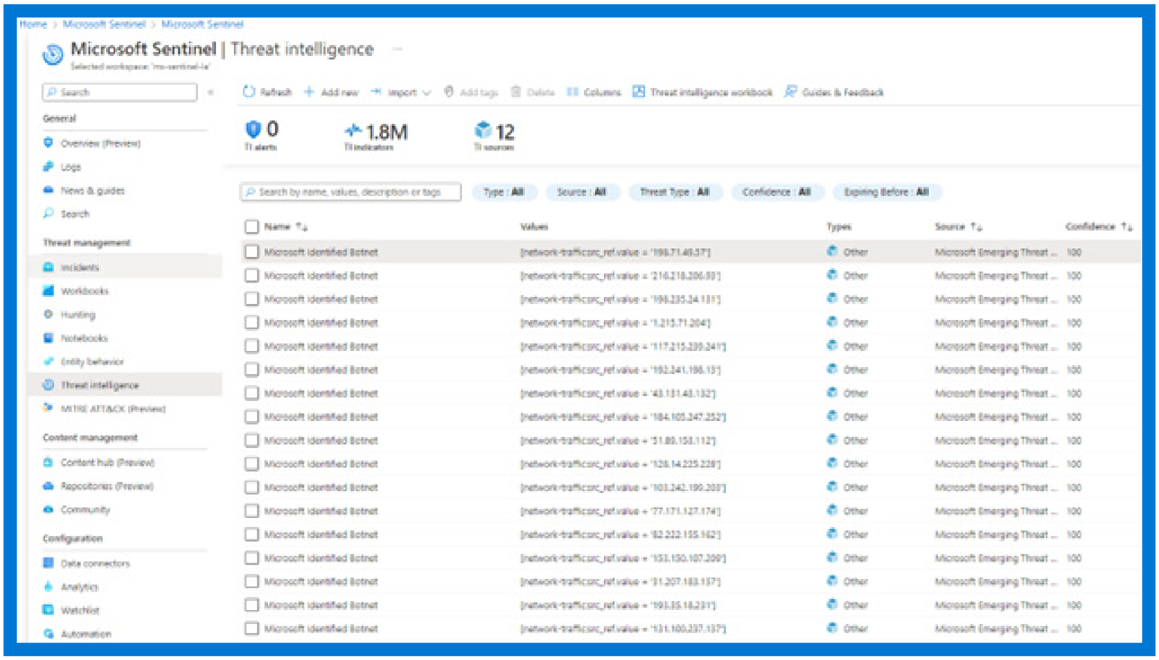

Name	Values	Types	Source	Confidence
Microsoft Identified Botnet	[network-traffic:src_ref.value = '198.71.40.57']	Other	Microsoft Emerging Threat ...	100
Microsoft Identified Botnet	[network-traffic:src_ref.value = '216.218.206.93']	Other	Microsoft Emerging Threat ...	100
Microsoft Identified Botnet	[network-traffic:src_ref.value = '198.235.24.131']	Other	Microsoft Emerging Threat ...	100
Microsoft Identified Botnet	[network-traffic:src_ref.value = '1.215.71.204']	Other	Microsoft Emerging Threat ...	100
Microsoft Identified Botnet	[network-traffic:src_ref.value = '117.215.220.241']	Other	Microsoft Emerging Threat ...	100
Microsoft Identified Botnet	[network-traffic:src_ref.value = '192.241.198.13']	Other	Microsoft Emerging Threat ...	100
Microsoft Identified Botnet	[network-traffic:src_ref.value = '43.131.43.132']	Other	Microsoft Emerging Threat ...	100
Microsoft Identified Botnet	[network-traffic:src_ref.value = '184.105.247.252']	Other	Microsoft Emerging Threat ...	100
Microsoft Identified Botnet	[network-traffic:src_ref.value = '51.89.153.112']	Other	Microsoft Emerging Threat ...	100
Microsoft Identified Botnet	[network-traffic:src_ref.value = '128.14.225.228']	Other	Microsoft Emerging Threat ...	100
Microsoft Identified Botnet	[network-traffic:src_ref.value = '103.242.190.203']	Other	Microsoft Emerging Threat ...	100
Microsoft Identified Botnet	[network-traffic:src_ref.value = '77.171.127.174']	Other	Microsoft Emerging Threat ...	100
Microsoft Identified Botnet	[network-traffic:src_ref.value = '82.222.155.162']	Other	Microsoft Emerging Threat ...	100
Microsoft Identified Botnet	[network-traffic:src_ref.value = '153.150.107.209']	Other	Microsoft Emerging Threat ...	100
Microsoft Identified Botnet	[network-traffic:src_ref.value = '31.207.183.157']	Other	Microsoft Emerging Threat ...	100
Microsoft Identified Botnet	[network-traffic:src_ref.value = '193.35.18.231']	Other	Microsoft Emerging Threat ...	100
Microsoft Identified Botnet	[network-traffic:src_ref.value = '131.100.237.137']	Other	Microsoft Emerging Threat ...	100

Fig. 19. Example IOCs

CTI can be imported to Microsoft Sentinel through many methods, but the most common is via the Microsoft Graph security API (most of the available Microsoft Sentinel connectors rely on this method).

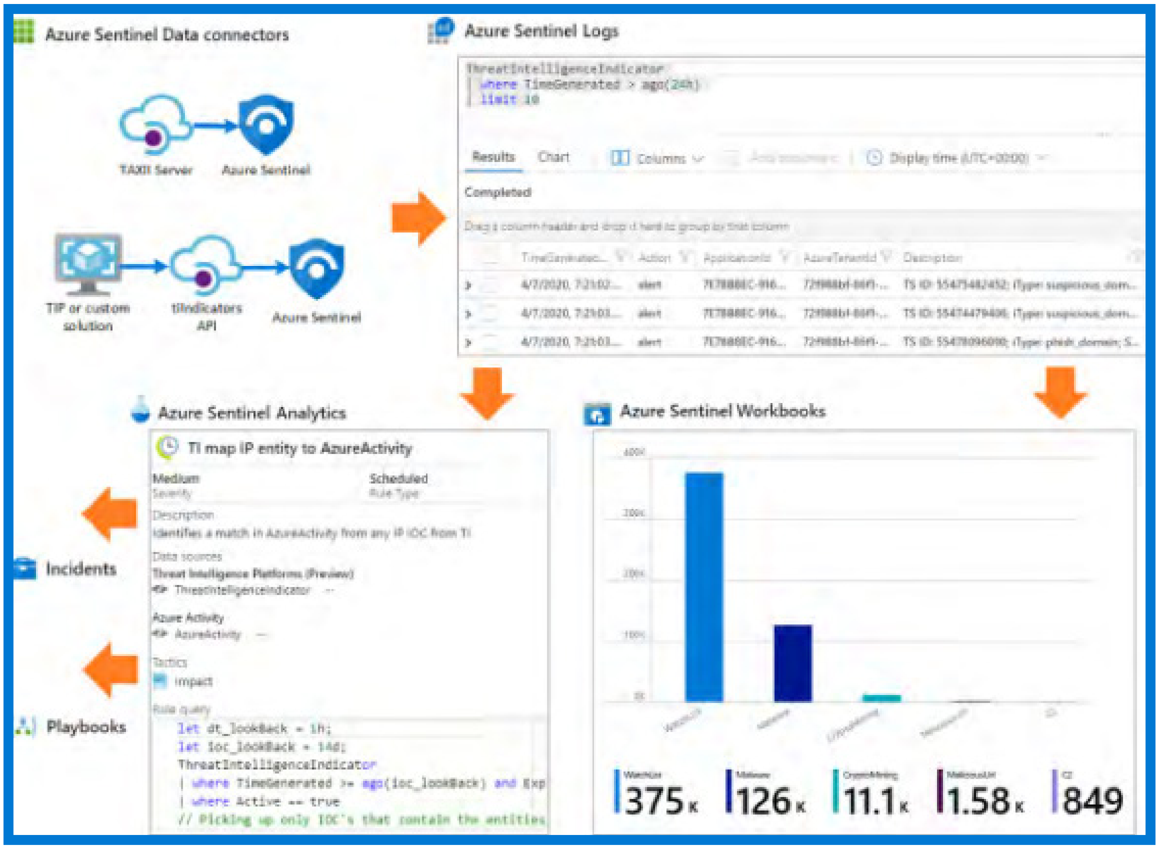

```
ThreatIntelligenceIndicator
| where TimeGenerated > ago(24h)
| limit 10
```

Completed

Drag a column header and drop it here to group by that column

	TimeGenerated...	Action	ApplicationId	AzureTenantId	Description
>	4/7/2020, 7:21:02...	alert	7E7888EC-916...	72f988bf-86f1-...	TS ID: 55475482452; iType: suspicious_dom...
>	4/7/2020, 7:21:03...	alert	7E7888EC-916...	72f988bf-86f1-...	TS ID: 55474479406; iType: suspicious_dom...
>	4/7/2020, 7:21:03...	alert	7E7888EC-916...	72f988bf-86f1-...	TS ID: 55478096090; iType: phish_domain; S...

TI map IP entity to AzureActivity

Severity: Medium

Rule Type: Scheduled

Description

Identifies a match in AzureActivity from any IP IOC from TI

Data sources

Threat Intelligence Platforms (Preview)
ThreatIntelligenceIndicator

Azure Activity
AzureActivity

Tactics
Impact

Rule query

```
let dt_lookBack = 1h;
let ioc_lookBack = 14d;
ThreatIntelligenceIndicator
| where TimeGenerated >= ago(ioc_lookBack) and Exp
| where Active == true
// Picking up only IOC's that contain the entities
```

Fig. 20. Microsoft Sentinel threat intelligence flow

To import using the Microsoft Sentinel TAXII data connector, you will need to know the advertised TAXII API root and collection ID. These are generally published in the vendor documentation and often include a discovery endpoint to help identify all available collection IDs.

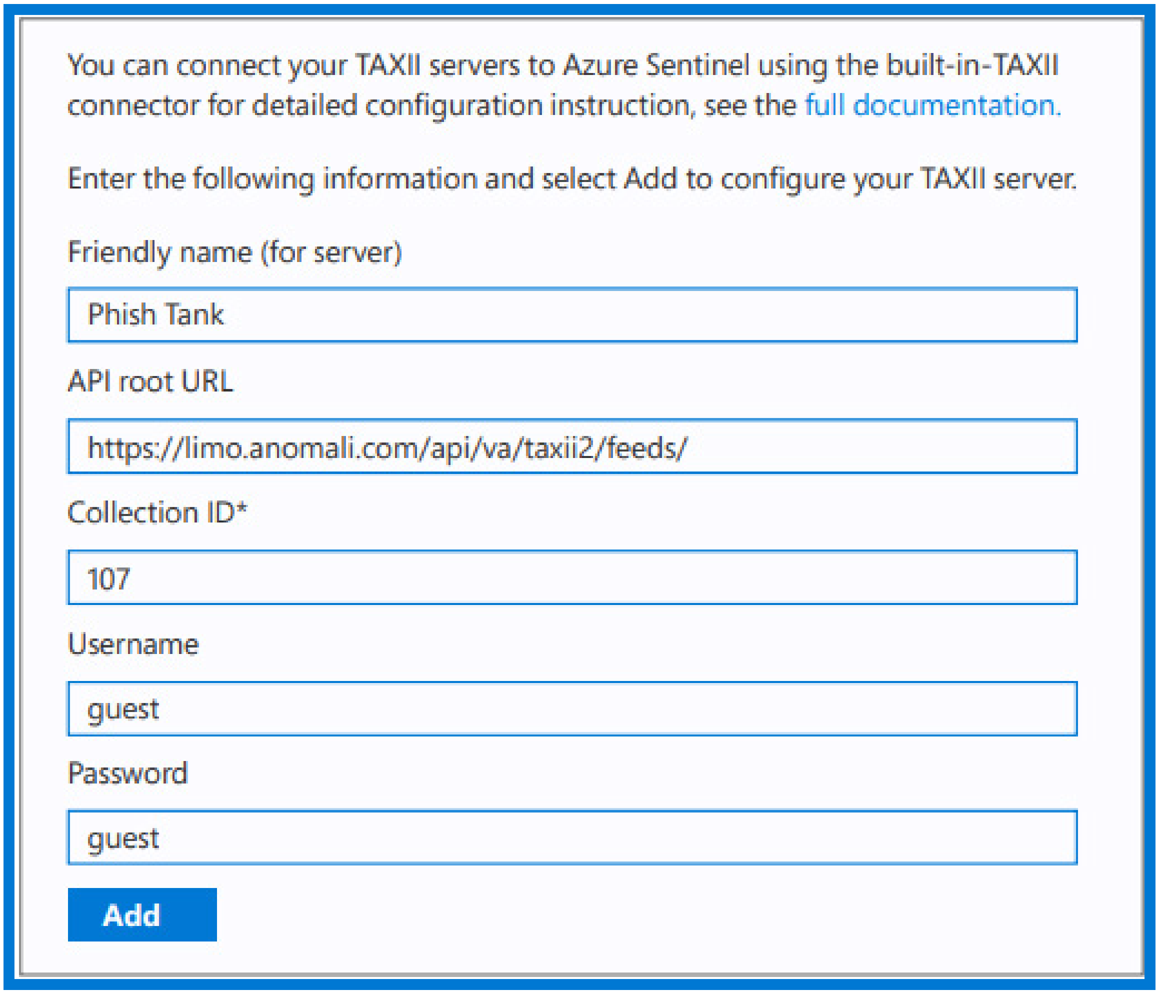

You can connect your TAXII servers to Azure Sentinel using the built-in-TAXII connector for detailed configuration instruction, see the full documentation.

Enter the following information and select Add to configure your TAXII server.

Friendly name (for server)

Phish Tank

API root URL

https://limo.anomali.com/api/va/taxii2/feeds/

Collection ID*

107

Username

guest

Password

guest

Fig. 21. Configuring a TAXII data connector

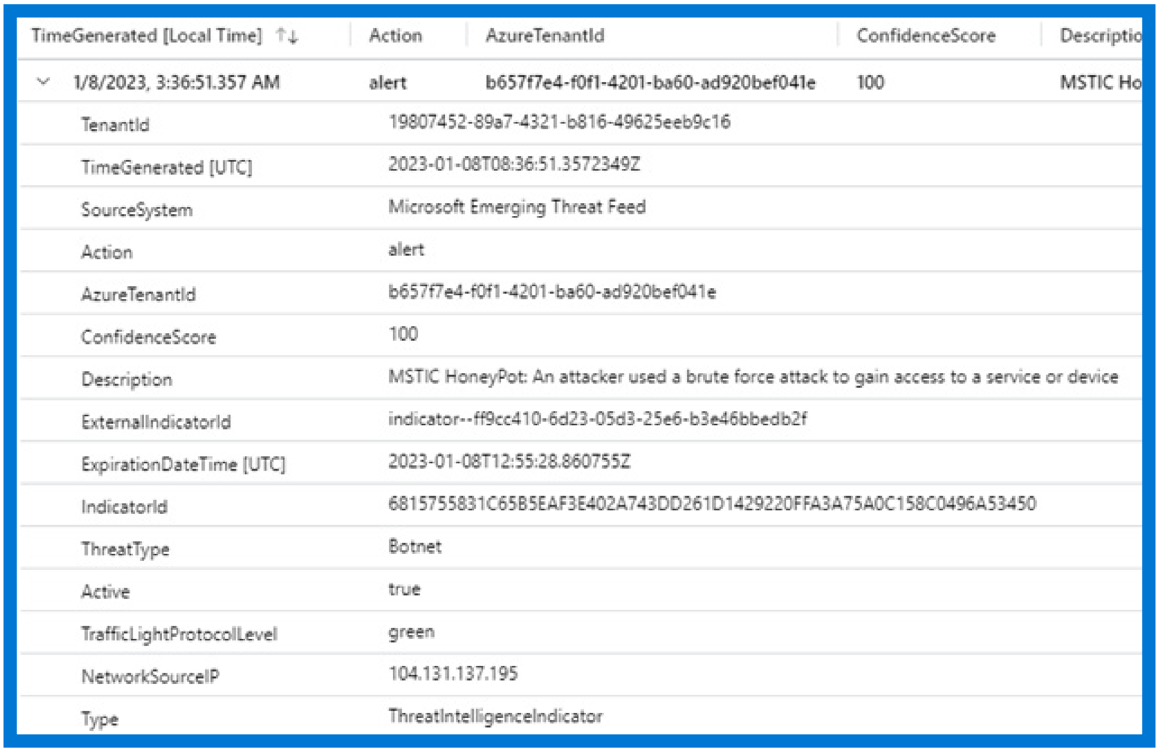

TimeGenerated [Local Time] ↑↓	Action	AzureTenantId	ConfidenceScore	Descriptio
˅ 1/8/2023, 3:36:51.357 AM	alert	b657f7e4-f0f1-4201-ba60-ad920bef041e	100	MSTIC Ho

TenantId	19807452-89a7-4321-b816-49625eeb9c16
TimeGenerated [UTC]	2023-01-08T08:36:51.3572349Z
SourceSystem	Microsoft Emerging Threat Feed
Action	alert
AzureTenantId	b657f7e4-f0f1-4201-ba60-ad920bef041e
ConfidenceScore	100
Description	MSTIC HoneyPot: An attacker used a brute force attack to gain access to a service or device
ExternalIndicatorId	indicator--ff9cc410-6d23-05d3-25e6-b3e46bbedb2f
ExpirationDateTime [UTC]	2023-01-08T12:55:28.860755Z
IndicatorId	6815755831C65B5EAF3E402A743DD261D1429220FFA3A75A0C158C0496A53450
ThreatType	Botnet
Active	true
TrafficLightProtocolLevel	green
NetworkSourceIP	104.131.137.195
Type	ThreatIntelligenceIndicator

Fig. 22. Sample threat intelligence indicator record (ingested via the Microsoft Threat Intelligence solution available in the Microsoft Sentinel Content Hub)

TIPs (Threat Intelligence Platforms) are external platforms that allow organizations to aggregate feeds from a variety of sources and curate the collected data to deduplicate redundancies. These platforms generally include mechanisms to apply these IOCs to security solutions, such as blocking known bad websites at your firewall or stopping malicious file hashes at your endpoints. Microsoft Sentinel can ingest this same data using the TIPs data connector. That uses an API and Application (client) ID, Directory (tenant) ID, and client secret from your TIP to connect and send threat indicators for use in Microsoft Sentinel.

The most important use case for threat indicators in Microsoft Sentinel is to drive the analytics that correlate events with IOCs to generate security alerts, incidents, and automated responses. Azure Playbooks, based on workflows built into Azure Logic Apps, can help automate and orchestrate a response when an IOC is identified in an alert. For example, contextual information can be added to an incident with the enriched data recorded as comments. That includes rich markup and HTML formatting to aid your analysts and incident responders even before they open the incident.

Microsoft Sentinel also provides several additional ways to ingest third-party TI:

- **Custom logs** - If the TI is available as a local file (i.e., a comma separated variable [CSV] file), MMA/AMA can be used to collect the updated data and send it to a Microsoft Sentinel custom log. In addition to MMA/AMA functionality, any custom log can be brought into Microsoft Sentinel using other log management solutions that integrate with Azure Log Analytics, such as Logstash and Fluentd.
- **Watchlists** - As one of the Microsoft Sentinel features, watchlists allow for the importing of CSV data and its use in a Microsoft Sentinel table to integrate with alert rules and queries.
- **External Data** - If the TI data is available as a URL (e.g., downloadable CSV files, Azure blobs), the KQL function externaldata can be used to download it on demand and use it as a temporary Microsoft Sentinel table.

In Our Experience,

- Very few organizations take advantage of the threat intelligence features available in Microsoft Sentinel (including the free Microsoft Threat Intelligence solution available in the Content Hub)
- Automation playbooks can be developed to integrate most of the available TI data, even if they don't have a standard data connector available.
- Through automation playbooks, indicators of compromise can be shared between Microsoft Sentinel and other tools, such as Defender for Endpoint

Additional Resources:

Threat indicators for cyber threat intelligence in Microsoft Sentinel – Microsoft Docs

Bring your threat intelligence to Microsoft Sentinel - Microsoft Tech Community

Connect threat intelligence data to Microsoft Sentinel - Microsoft Docs

Tutorial: Set up automated threat responses in Microsoft Sentinel – Microsoft Docs

Deploying Alert Rules

As the core functionality of a SIEM, the configuration of detection rules is a critical component of any Microsoft Sentinel deployment. Microsoft Sentinel includes a large number of built-in alert rules templates (over 500 as of January 2023), covering the array of typical log sources, with new alert rules added regularly. The solutions available through the Content Hub also include a wide variety of alert rules for each solution.

Additional rule templates can be obtained through the Microsoft Sentinel Community, where both Microsoft and third-party contributors publish new content. The community reviews the proposed alert rules, and those found valuable are published in Microsoft Sentinel.

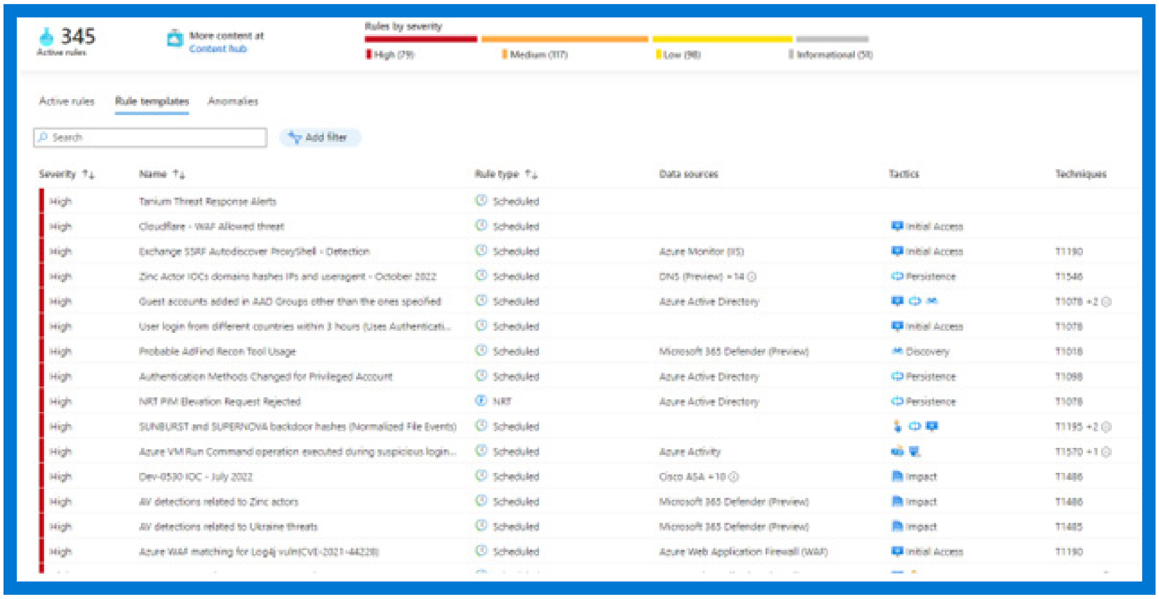

345
Active rules

More content at
Content hub

Rules by severity
High (79) Medium (117) Low (98) Informational (51)

Active rules | Rule templates | Anomalies

Search | Add filter

Severity ↑↓	Name ↑↓	Rule type ↑↓	Data sources	Tactics	Techniques
High	Tanium Threat Response Alerts	Scheduled			
High	Cloudflare - WAF Allowed threat	Scheduled		Initial Access	
High	Exchange SSRF Autodiscover ProxyShell - Detection	Scheduled	Azure Monitor (IIS)	Initial Access	T1190
High	Zinc Actor IOCs domains hashes IPs and useragent - October 2022	Scheduled	DNS (Preview) +14	Persistence	T1546
High	Guest accounts added in AAD Groups other than the ones specified	Scheduled	Azure Active Directory		T1078 +2
High	User login from different countries within 3 hours (Uses Authenticati...	Scheduled		Initial Access	T1078
High	Probable AdFind Recon Tool Usage	Scheduled	Microsoft 365 Defender (Preview)	Discovery	T1018
High	Authentication Methods Changed for Privileged Account	Scheduled	Azure Active Directory	Persistence	T1098
High	NRT PIM Elevation Request Rejected	NRT	Azure Active Directory	Persistence	T1078
High	SUNBURST and SUPERNOVA backdoor hashes (Normalized File Events)	Scheduled			T1195 +2
High	Azure VM Run Command operation executed during suspicious login...	Scheduled	Azure Activity		T1570 +1
High	Dev-0530 IOC - July 2022	Scheduled	Cisco ASA +10	Impact	T1486
High	AV detections related to Zinc actors	Scheduled	Microsoft 365 Defender (Preview)	Impact	T1486
High	AV detections related to Ukraine threats	Scheduled	Microsoft 365 Defender (Preview)	Impact	T1485
High	Azure WAF matching for Log4j vuln(CVE-2021-44228)	Scheduled	Azure Web Application Firewall (WAF)	Initial Access	T1190

Fig. 27. Microsoft Sentinel alert rule templates

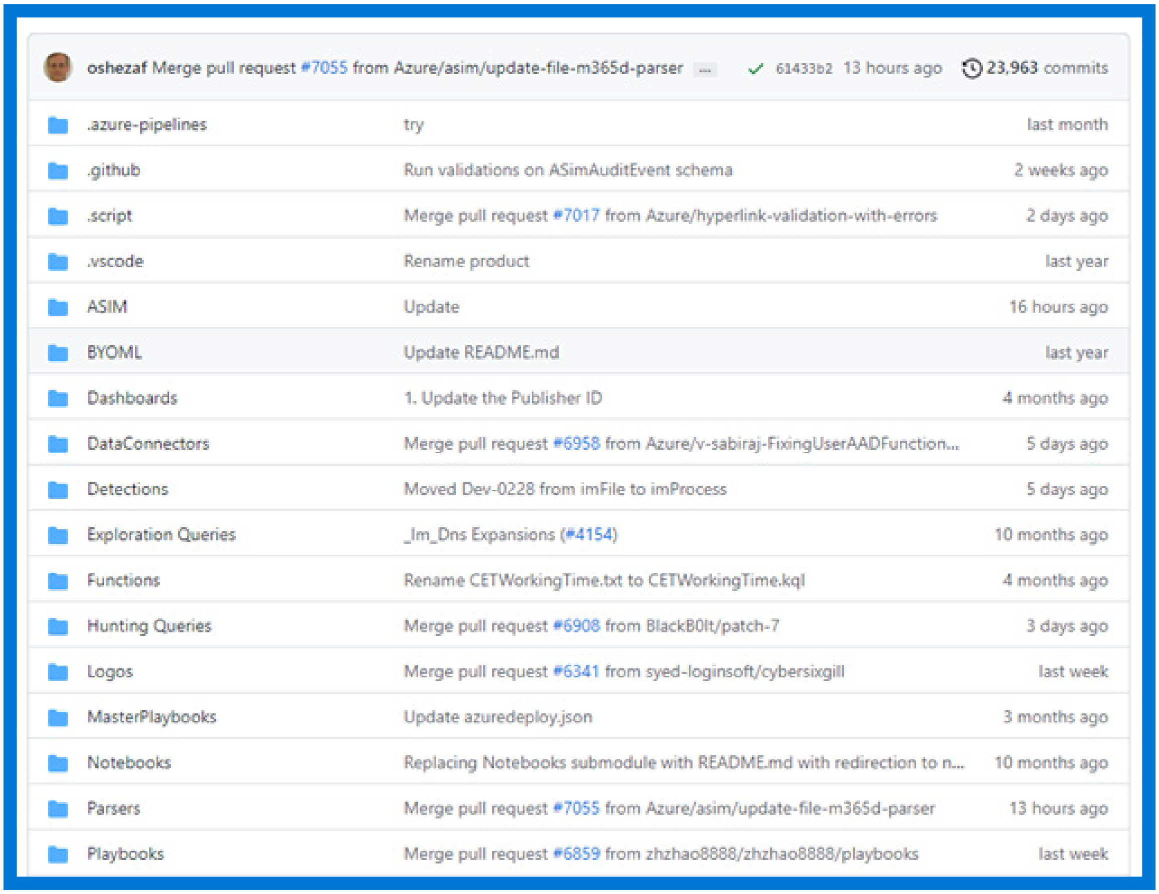

oshezaf Merge pull request #7055 from Azure/asim/update-file-m365d-parser ... ✓ 61433b2 13 hours ago 23,963 commits

.azure-pipelines	try	last month
.github	Run validations on ASimAuditEvent schema	2 weeks ago
.script	Merge pull request #7017 from Azure/hyperlink-validation-with-errors	2 days ago
.vscode	Rename product	last year
ASIM	Update	16 hours ago
BYOML	Update README.md	last year
Dashboards	1. Update the Publisher ID	4 months ago
DataConnectors	Merge pull request #6958 from Azure/v-sabiraj-FixingUserAADFunction...	5 days ago
Detections	Moved Dev-0228 from imFile to imProcess	5 days ago
Exploration Queries	_Im_Dns Expansions (#4154)	10 months ago
Functions	Rename CETWorkingTime.txt to CETWorkingTime.kql	4 months ago
Hunting Queries	Merge pull request #6908 from BlackB0lt/patch-7	3 days ago
Logos	Merge pull request #6341 from syed-loginsoft/cybersixgill	last week
MasterPlaybooks	Update azuredeploy.json	3 months ago
Notebooks	Replacing Notebooks submodule with README.md with redirection to n...	10 months ago
Parsers	Merge pull request #7055 from Azure/asim/update-file-m365d-parser	13 hours ago
Playbooks	Merge pull request #6859 from zhzhao8888/zhzhao8888/playbooks	last week

Fig. 28. Microsoft Sentinel Community GitHub Repository

Rule templates can be deployed as Microsoft Sentinel scheduled alerts in the Microsoft Sentinel portal. These are fully customizable for use cases that extend beyond the out-of-the-box detections.

The alert rule templates can be filtered based on their data source, e.g. "Azure Firewall" so that the relevant alert rules can be identified and deployed.

The out-of-the-box alert rule templates are designed to cover a wide variety of environments. In most cases, after the initial deployment, the alert rules may have to be tuned to match the existing infrastructure. Such tuning may involve adjustments of thresholds, exclusion of certain hosts or IP addresses, adjustment of frequency or correlation with an additional log source for increased accuracy.

Important: The tuning of the alert rules is critical for a mature SIEM deployment.

All alert rules deployments should be included in regular reviews for their value and adjustments made to make them more relevant to the organization. As a cloud-based solution, Microsoft Sentinel provides frequent access to new content. Alert rule deployment and tuning lifecycle should match the dynamics of your threat landscape. A weekly review is recommended during the initial deployment, followed by monthly reviews as the environment enters a more stable state.

While Microsoft Sentinel provides hundreds of detections, each organization will, at some point, require the development of new detections that may be specific to their own environment. In order to manage the detection rules and the additional Microsoft Sentinel resources (workbooks, automation playbooks, etc.) a content development and deployment team may have to be established. Such a team will be responsible for the creation, modification, and deployment of detection rules.

The alert rules development should be driven from the top:

- The assets involved in the scope of monitoring and the level of monitoring have to be identified based on the organization's security policies. Those can be driven by the cybersecurity framework adopted by the organization (i.e., NIST's Cybersecurity Framework) or other compliance requirements. That type of process may end up with a revision of the security policies. In many cases, existing policies do not keep pace with the realities of the hybrid infrastructure used by most organizations.

- The log data required for the monitoring ingested into Microsoft Sentinel. As that can be an ongoing process, the onboarding of log sources should be prioritized based on their importance and value for security monitoring.

- Required detection rules developed and deployed and their efficiency/quality monitored and measured on a regular basis. As with the log sources, the development of the detection rules has to be prioritized based on the gaps identified and the monitored systems in scope.

- Automation rules / SOAR playbooks developed

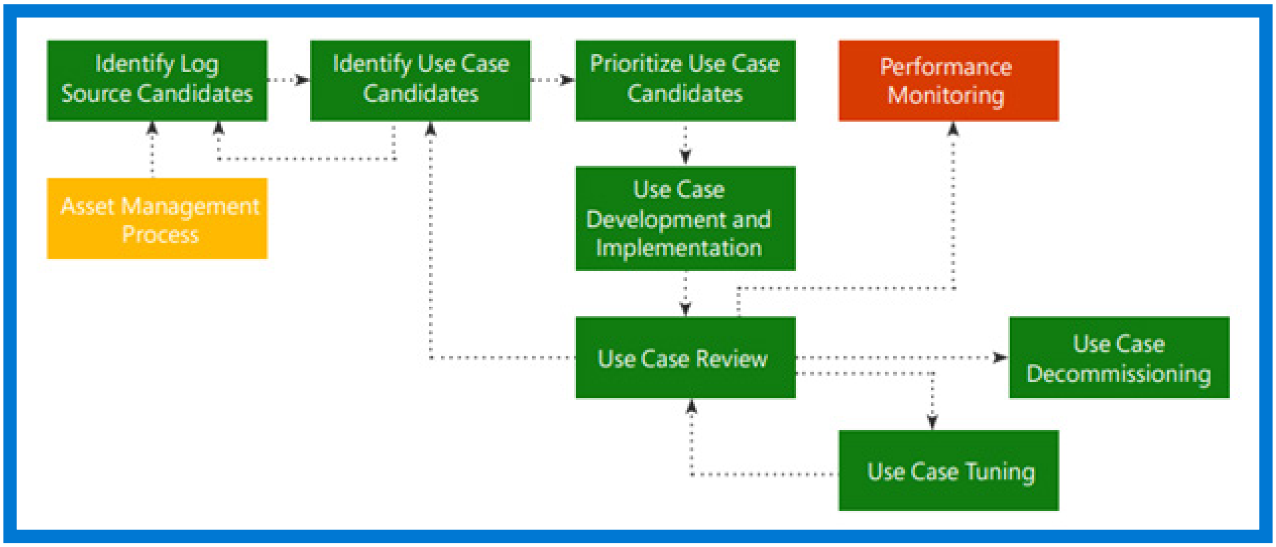
Identify Log Source Candidates
Identify Use Case Candidates
Prioritize Use Case Candidates
Performance Monitoring
Asset Management Process
Use Case Development and Implementation
Use Case Review
Use Case Decommissioning
Use Case Tuning

Fig. 29. Sample SIEM use-case lifecycle

If the Microsoft Sentinel infrastructure involves more than one log analytics workspace, one important aspect to consider is the distribution of content to all Microsoft Sentinel instances. Even for just one instance, it may help to use CI/CD pipelines for managing content such as detection rules, Microsoft Sentinel functions (used by parsers), workbooks, and even automation playbooks.

Microsoft Sentinel provides the “Repositories” feature that allows the use of GitHub or Azure DevOps repositories to automatically push content to Microsoft Sentinel instances that subscribe to that particular repository.

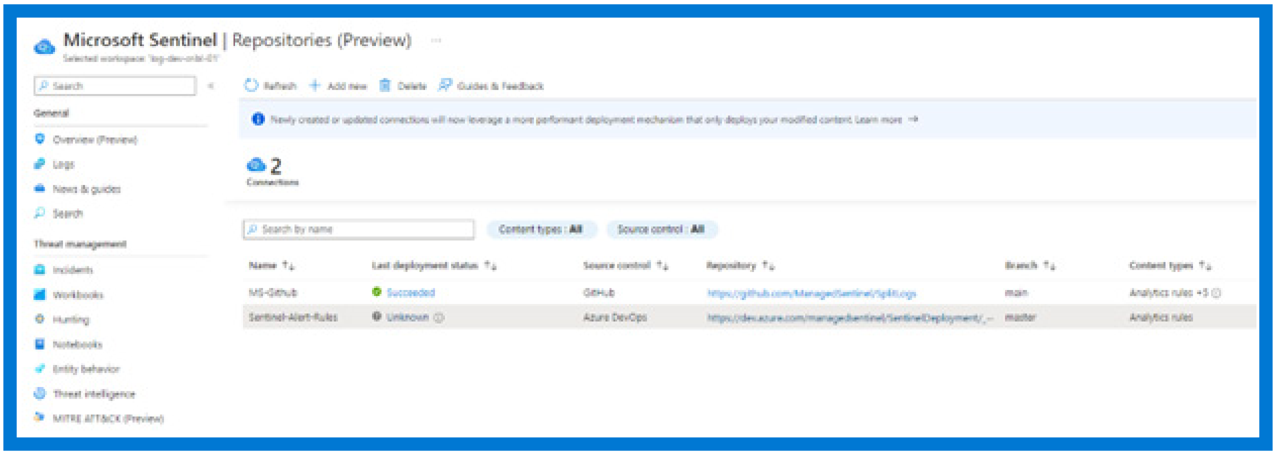

Newly created or updated connections will now leverage a more performant deployment mechanism that only deploys your modified content. Learn more →

Name ↑↓	Last deployment status ↑↓	Source control ↑↓	Repository ↑↓	Branch ↑↓	Content types ↑↓
MS-Github	Succeeded	GitHub	https://github.com/ManagedSentinel/SplitLogs	main	Analytics rules +5 ⓘ
Sentinel-Alert-Rules	Unknown ⓘ	Azure DevOps	https://dev.azure.com/managedsentinel/SentinelDeployment/_...	master	Analytics rules

Fig. 29.1. Sample configuration of Microsoft Sentinel content repository

Another related feature, in private preview as of January 2023, is the Workspace Manager. The new feature allows for the creation of groups of Microsoft Sentinel instances and deployment of sets of detections from a “golden image” (a Sentinel instance will have all the available alert rules).

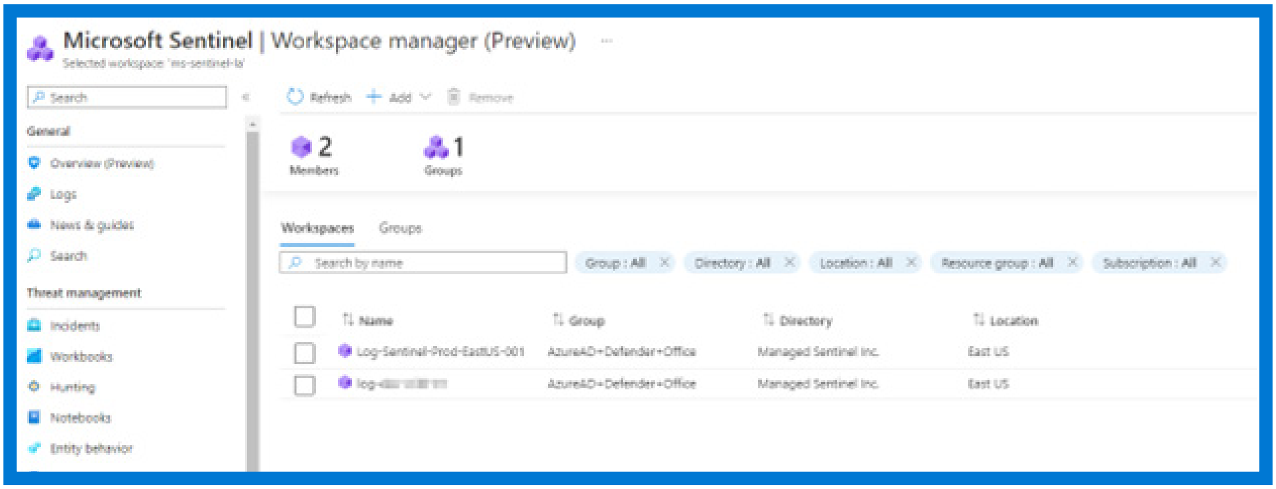

Name	Group	Directory	Location
Log-Sentinel-Prod-EastUS-001	AzureAD+Defender+Office	Managed Sentinel Inc.	East US
log-[redacted]	AzureAD+Defender+Office	Managed Sentinel Inc.	East US

Fig. 29.2. Sample configuration of Microsoft Sentinel Workspace Manager

The centralized management of Microsoft Sentinel content is critical for larger organizations that intend to have a unified approach to maintaining a standard set of detection rules or have compliance requirements for deploying by code only. The existing Microsoft Sentinel features provide a good start in implementing such a process. However, they may have to be augmented with additional tools to cover requirements specific to each organization. The fully featured REST API offered by Microsoft Sentinel and Log Analytics Workspace can be used to automate almost any aspect of Microsoft Sentinel content.

If Microsoft Sentinel is deployed or managed by a third-party, such as a Managed Security Services Provider (MSSP) providing Microsoft Sentinel management services, additional alert rules might be available from their catalog along with more automated deployment methods of monitoring and tuning advice. Most MSSPs build automation tools that allow consistent deployment and management of content across the Microsoft Sentinel instances that they manage.

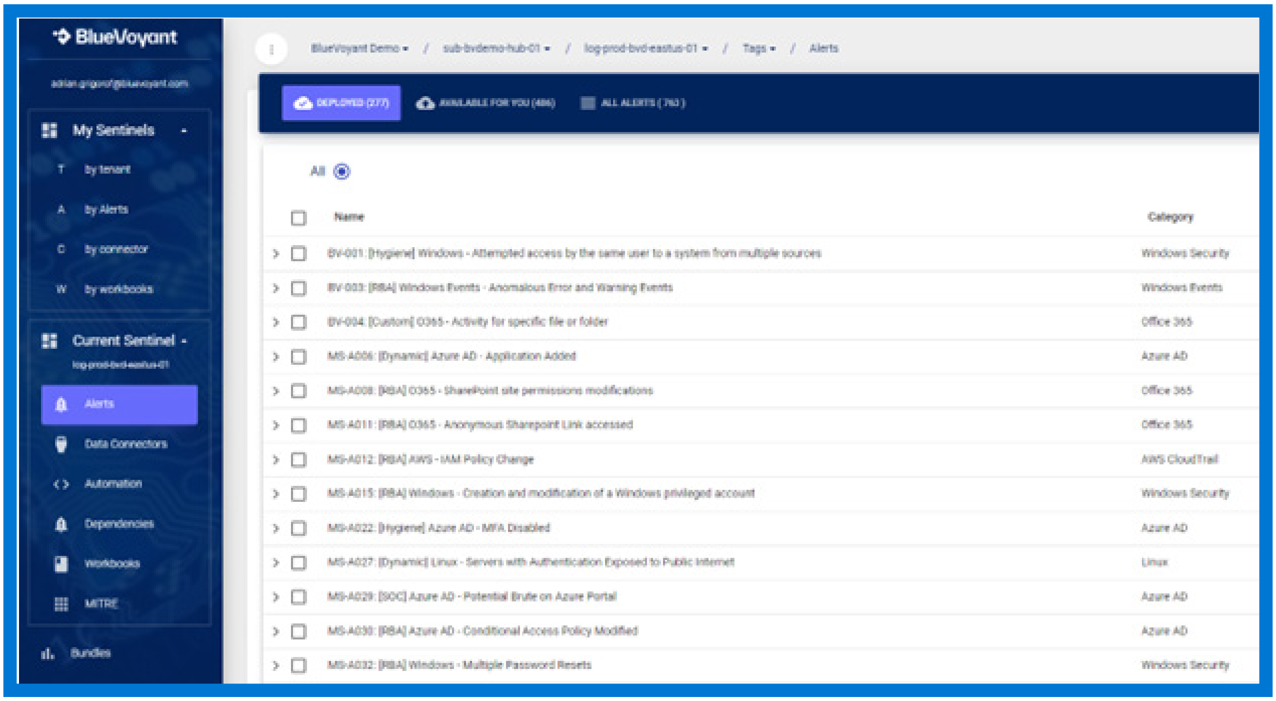

Name	Category
BV-001: [Hygiene] Windows - Attempted access by the same user to a system from multiple sources	Windows Security
BV-003: [RBA] Windows Events - Anomalous Error and Warning Events	Windows Events
BV-004: [Custom] O365 - Activity for specific file or folder	Office 365
MS-A006: [Dynamic] Azure AD - Application Added	Azure AD
MS-A008: [RBA] O365 - SharePoint site permissions modifications	Office 365
MS-A011: [RBA] O365 - Anonymous Sharepoint Link accessed	Office 365
MS-A012: [RBA] AWS - IAM Policy Change	AWS CloudTrail
MS-A015: [RBA] Windows - Creation and modification of a Windows privileged account	Windows Security
MS-A022: [Hygiene] Azure AD - MFA Disabled	Azure AD
MS-A027: [Dynamic] Linux - Servers with Authentication Exposed to Public Internet	Linux
MS-A029: [SOC] Azure AD - Potential Brute on Azure Portal	Azure AD
MS-A030: [RBA] Azure AD - Conditional Access Policy Modified	Azure AD
MS-A032: [RBA] Windows - Multiple Password Resets	Windows Security

Fig. 29.3. BlueVoyant's Microsoft Sentinel content management platform

Given the criticality of alert rules to core SIEM functionality, seeking consulting advice from third-party SIEM specialists may be advisable. This could be an on-going engagement, or an initial deployment of Microsoft Sentinel based on best practices, including knowledge transfer for its management.

In Our Experience,

- Most organizations lack a use-case development process. They rely on deploying what is available out-of-the-box from their SIEM platform. Microsoft Sentinel provides hundreds of detections, but without understanding what they cover and mapping them against the organization's requirements, it is difficult to assess the current security posture and how the detections match potential threats. That is probably the number 1 reason for failures to successfully adopt a SIEM platform and see the expected return of investment.
- Avoid the temptation to enable too many alert rules at once to avoid an influx of false positive alerts. If left alone, those can cause "the boy who cried wolf" alert fatigue among your analysts. Start with a handful, tune the associated alert rule logic to match the expectations of your organization, and then enable more.
- Alert tuning may feel a bit like a dark art at first. Take one alert rule at a time, step through the logic in the flow, and consider the trigger against the associated fields. You will find the field (or even an entire table) that should be added or removed. That will eliminate false positives and have the alert rule reflect your specific use case better. Remove (or add) that section and test. The more times you run through this exercise, the better you will learn your own data and the faster this will become.
- Custom tables and custom data are generally not included in default alert rules. Those may need to be added or have custom rules created to maximize the effectiveness of the newly sourced data.
- Professional services deployments are a good way to quickly deploy your instance. Include knowledge transfer, which is essential for your staff to learn the "what, when, and how" about alert rule tuning for future use cases.
- Centralized management of Microsoft Sentinel content is critical for organizations with multiple Microsoft Sentinel instances.

Some helpful links about learning Microsoft Sentinel, including alert rule tuning:

Microsoft Learning Path: Cloud-native security operations with Microsoft Sentinel – Microsoft Docs

Become a Microsoft Sentinel Ninja: The complete level 400 training – Microsoft Tech Community

Sentinel Community - Join discussions about Microsoft Sentinel

Deploy custom content from your repository – Microsoft Docs

Implement Microsoft Sentinel and Microsoft 365 Defender for Zero Trust

Migration from Existing SIEM Solutions

Microsoft Sentinel is a relatively new entrant to the SIEM market. Most projects involve the deployment of Microsoft Sentinel with the migration and cutover from another legacy SIEM. Based on our experience leading numerous migration projects from a wide variety of SIEMs to Microsoft Sentinel, we have compiled some recommended approaches presented in the following three scenarios.

Scenario 1 - There is a legacy SIEM solution

That involves migrating from an existing, older, and possibly obsolete SIEM. The vendor may not have kept up with the latest developments or the architecture is difficult or expensive to scale with the growing log volume. In many cases, the existing SIEMs are already overloaded with too many logs or a significant licensing expense. They are not providing expected or satisfactory security visibility. Quite often, these scenarios involve a need for additional and potentially significant costs in licensing or a need to upgrade aging hardware.

Use case migration

Legacy use cases will have to be inventoried and the ones that are deemed to provide value should be documented and converted in Microsoft Sentinel alert rules or playbooks. Fortunately, for the most common ones, Microsoft Sentinel already provides support through a large array of built-in alert rule templates.

Due to differences in the overall SIEM design and the significant shift from on-premises to cloud-based technology, the effort to convert the existing SIEM use cases may vary. It is common to become focused on reproducing a specific legacy alert in Microsoft Sentinel verbatim rather than target the intent of the original use case. The requirements should be documented as abstractly as possible to allow for the differences between analysis/alerting processes.

Particular attention must be provided to any log collection agents present in the legacy SIEM solution. The typical agents are deployed on endpoints such as Windows or Linux servers. Each SIEM solution has its own design and may require significant changes in how such agents are deployed and monitored. For example, a project may include pull versus push agents, agentless solutions, centralized collection of logs using Windows Event Log Forwarder, or Linux syslog versus native SIEM log collection method.

The differences in agent deployment and log collection must be very well understood and the relevant infrastructure staged to avoid last moment major redesigns.

A small pilot or development environment should be deployed to address challenges that are not easily visible on paper.

One very important difference between legacy on-premises SIEMs and Microsoft Sentinel is the need to send the log data to the cloud. This requires the proper network infrastructure in place with internet or VPN connections that can provide sufficient bandwidth as well as opening the necessary firewall ports for outbound access to Azure. Microsoft Sentinel provides multiple options for aggregating the logging data using a single collector or gateway before sending to the cloud repository for analysis.

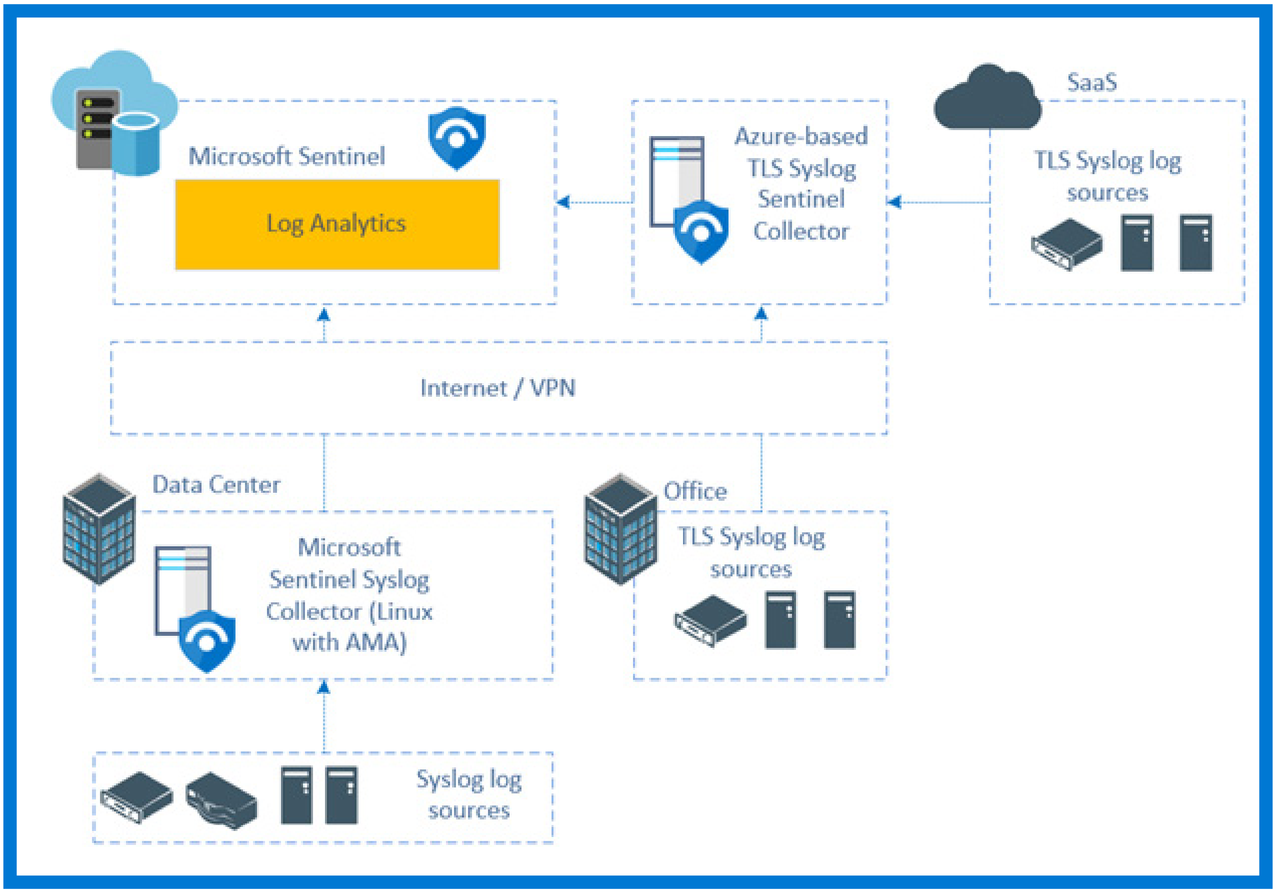

Fig. 30. Microsoft Sentinel collectors for on-premises log sources

A testing plan should be established for the removal of existing agents, configuration of AMA, and the migration of syslog-based log sources from the legacy collector to Microsoft Sentinel. In some situations, the collectors can be recycled, but it is seldom recommended. It is recommended to choose fresh installs wherever feasible.

The existing SIEM can be kept in place until all resources have been migrated and the use cases fully tested in Microsoft Sentinel. Special consideration must be given to the data stored in the legacy SIEM, with provisions made to allow access to historical data based on log retention policies or compliance needs. For example, the Payment Card Industry Data Security Standard (PCI DSS) requires access to the prior year of in-scope log data.

In Our Experience,

- Care should be taken to manage risk in cases where the legacy SIEM is approaching license/ support renewal, the organization has an imminent audit, or the organization has insufficient resources to devote to properly managing the migration. Decisions under pressure or duress can result in risk to the enterprise.
- If extended log retention was used with the legacy SIEM, a strategy should be developed on maintaining access to the retained data for the required period (based on the organization's policies). For most SIEMs, it is difficult and expensive to extract large volumes of logs that are stored for long-term retention.
- Many existing use cases in production SIEM solutions have an equivalent in the existing rule base for Microsoft Sentinel or can be re-created simply with a KQL query.
- Many organizations migrating to Microsoft Sentinel choose to completely discard existing SIEM use cases and start fresh in Microsoft Sentinel. That is due to lack of maintenance and relevance with the current log sources.
- SIEM analysts experienced with a specific legacy platform may desire training and assistance when transitioning concepts from legacy technology to cloud-native solutions. Partnering the analysts with professional services while implementing the new Microsoft Sentinel deployment may help streamline the adoption of the new platform.
- Evaluate the true nature of logs ingested by the legacy SIEM. For example, it may ingest 20 GB/ day due to some limitation while the true log volume may be much greater. Miscalculation of these estimates can lead to unexpected costs during the Microsoft Sentinel deployment.
- If a legacy SIEM has been maintained by a third party, such as an MSSP, be sure to check into any applicable contractual agreements around service exit provisions.
- If a large number of automation playbooks are in place, resources to validate their relevance and test the newly created ones might not be available. Such automations typically involve escalation/notification of various teams (i.e., networking, systems, databases) for specific events. As such content was developed for many years, it is difficult to convert and test them all in just a few months. When this is the case, the critical ones should be identified, and a roadmap built for the migration of the rest.

Scenario 2 - Legacy SIEM in place and a third-party analytics/SIEM platform deployment is still in progress

Another migration scenario is when the legacy SIEM is already in the process of being replaced by a new solution, but there are limitations around connecting Azure Cloud resources to the legacy SIEM, and there are critical Azure-based projects going live that need full SIEM coverage. Microsoft Sentinel can be deployed quickly to provide security analytics.

Most Azure-based resources have the capability to stream logs directly into a Log Analytics workspace and, through that, make the logs readily available to Microsoft Sentinel. This allows a reduction in time (from weeks/months to days) required to onboard raw logs from Azure resources into a SIEM. Depending on the volume of logging data, additional bandwidth costs related to cloud egress traffic can be avoided.

For certain Microsoft/Azure log sources, the ingestion of logging data into Microsoft Sentinel is free of charge so additional savings can be obtained by reducing the load on the on-premises SIEM. One such example is the onboarding of Office 365 logs, a non-billable table in Microsoft Sentinel.

Microsoft Sentinel has the capability to export new incidents into Events Hub and make the data available for third-party analytical platforms. Modern SIEM solutions typically can retrieve events from a data streaming platform like Azure Events Hub.

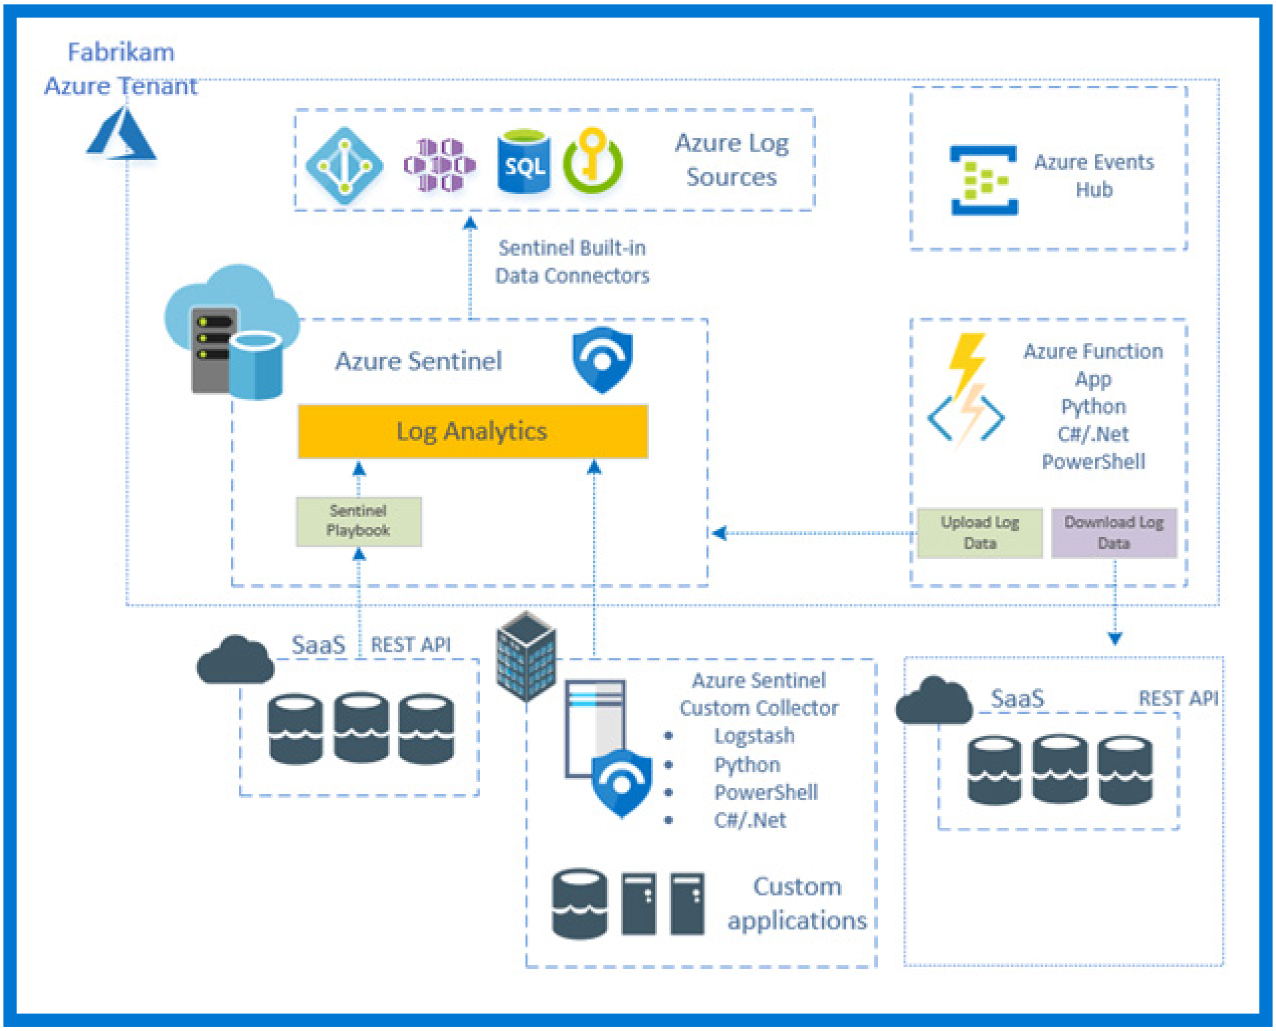

Fig. 31. Deployment of Microsoft Sentinel SIEM for fast-tracked analytics for Azure-based resources

A long-term strategy must be developed to guide day-to-day decisions about onboarding log sources in Microsoft Sentinel versus the third-party solution. In many cases, the argument for aggregation to a single repository is the ability to correlate data. Although in quite a few instances, especially for isolated application environments, the aggregation does not bring enough additional value to justify the increase in cost and complexity.

Scenario 3 - Migration from a limited logging/analytics platform such as syslog to Microsoft Sentinel

Organizations with basic log collection/analytics capabilities usually have a limited IT security budget or have an established raw log collection infrastructure and need to add advanced security analytics capabilities.

For new deployments, the value of each type of log source (from a security perspective) must be analyzed, and the ingestion of raw logs into the new SIEM must be prioritized on a cost/value basis. Keep in mind that Microsoft Sentinel log ingestion is free of charge for Office 365 logs, for Microsoft 365 security incidents, and several other log sources (see the full details earlier).

It is critical to onboard log sources gradually, understand their value from a security analysis perspective, and decide on their ingestion and retention strategy. Some common log sources that need increased scrutiny from a cost/value perspective are:

- Firewall traffic logs (allowed and denied traffic)
- Windows Active Directory logs (Windows Security event logs)
- Azure Diagnostics (raw logs from Azure resources)
- Web and WAF logs

The risk appetite, compliance requirements, and available budget vary from organization to organization. These are the typical factors that affect the decisions related to log ingestion and retention.

If a raw log collection and retention capability exists, it can be preserved for long-term retention or to meet compliance requirements. If necessary, older logs can be ingested in Microsoft Sentinel on an impromptu basis to perform more advanced analytics through KQL or Microsoft Sentinel Notebooks.

Additional Resources:

Webinar: Best practices converting detections rules from existing SIEM to Microsoft Sentinel

How do you export QRadar offenses to Microsoft Sentinel? – Microsoft Tech Community

Best practices for migrating detection rules from ArcSight, Splunk and QRadar to Microsoft Sentinel - Microsoft Tech Community – Microsoft Tech Community

Splunk to Kusto Query Language map – Microsoft Docs

Cost Management

Cost management for Microsoft Sentinel is inextricably tied to organizational risk management. Ingesting incremental data into Microsoft Sentinel from various network components allows for creating analytic rules to detect attacker behavior; however, every byte of data ingested into Log Analytics carries a cost. This section provides some practical guidance for building and evaluating a business case for adopting Microsoft Sentinel.

Evaluating Your Data Ingestion Against Use Cases

The expense of a given log source ingested into Microsoft Sentinel should be evaluated against a commensurate benefit. For example, the data aids in visibility for cyberattacks, can detect or prevent a data breach, or is simply required due to regulatory compliance. The ingestion of many types of log sources can incur additional expense for your Microsoft Sentinel deployment; therefore, a degree of cost/benefit analysis is required. For example, ingesting all data from the APIs of the top 100 employee-accessed SaaS applications could aid in certain investigations, but cost would likely be prohibitive. That would outweigh the potential benefit of increased security visibility, threat mitigation, or fulfilling a compliance requirement.

We advise identifying the applications that align with significant degrees of business risk. Project teams will want to survey all possible data sources and analyze variables such as log volume or anticipated risk mitigation. A log source with an unusually high volume of log data weighed against a relatively small number of potential security use cases could be an example to exclude from the project scope and Microsoft Sentinel budget.

Ultimately, this decision lies in the hands of the stakeholders and decision makers. It should require formal sign off for decisions by the enterprise risk management owners.

Log Ingestion Strategies

Any network device, endpoint, or application has the potential to generate log data, from basic information to verbose debugging-level details. In an ideal world, all logs will be recorded and stored indefinitely. The value of the information captured in the logs depends on the type of log source; its association with users, computers, and applications; and the level of detail built into the logging mechanism by the device or application vendor.

The requirements around logging levels, the type of information captured in the logs, and their retention requirements are driven by the organization's information security policies resulting from the cybersecurity governance process. In most cases, the policies are based on industry best practices and compliance requirements with standards (e.g., PCI, International Organization for Standardization [ISO] 27001, National Institute of Standards and Technology [NIST] Cybersecurity Framework, etc.). For many such standards, log retention requirements are vague and sometimes subject to an auditor's interpretation. As a principle, any logs related to the CIA (Confidentiality-Integrity-Availability) triad are in the scope of information security policies for logging. Still, most of the "must-have" requirements are related to auditing user activities and the integrity of IT systems and their processed data. Log collection must be treated as any security control driven by feasibility and its overall contribution to the organization's security stance. A common logging policy is to have the logs available for online analysis for 90 days with one year offline/archived to slower or less expensive storage.

In Our Experience,

- The majority of the organizations consider the current free 90 days, online storage offered by Microsoft Sentinel sufficient for their analytical needs
- Health organizations typically take advantage of the 7-year archiving option for most of their logs stored in Microsoft Sentinel
- Building a custom log retention infrastructure using solutions such as Azure Data Explorer (ADX) may be feasible for specific requirements (the larger the log volume, the better from a cost per GB stored perspective)
- Aggregating logs in Microsoft Sentinel may allow for complying with the log retention policy for solutions that don't offer extensive log retention options. For example, Microsoft Defender for Endpoint provides access to its raw logs for 30 days. If onboarded in Microsoft Sentinel they can be retained for up to seven years.

From a security perspective, the logs historically captured were those from perimeter security controls, such as firewalls and proxy servers, authentication/authorization servers, or Windows Active Directory Security event logs. Relatively few connections to untrusted locations, bandwidth limitations, plain text, unencrypted traffic, limited internet resources, and virtually nonexistent logging from endpoint security controls resulted in a log analysis strategy centered around ingestion and processing of perimeter security controls.

The exponential increase in internet usage has introduced many new devices with internet connectivity: mobile, Bring-Your-Own-Device (BYOD), Internet of Things (IoT) and SaaS. The commensurate adoption of encryption gradually reduced the visibility of traditional security controls, with the endpoints increasingly becoming the single location where the activities of potentially malicious actors are detectable.

Combined with the increase in logging data, the legacy approach to log analytics is no longer feasible. The "log everything" approach leads to either unmanageable costs or poor performance via quickly overwhelmed SIEM solutions. Our suggested approach is to analyze each type of log source in detail and weigh the costs versus benefits of ingestion for each type of log identified.

The analysis should consider not only the log source but also the logged field entries and their value from a threat detection perspective. A specific log source may generate a large number of event types, so they should not be considered under the same category. For example, a Fortinet firewall may log over 60 types of log entries from its various components: stateful firewall, IPS/IDS, configuration changes, VPN, authentication, wireless security, etc.

Consider the level of effort to detect use cases in relation to compensating controls in the dynamic of the threat landscape. That may provide equal or better visibility for the same cost. The following is a sample high-level analysis of log source volume versus threat detection value.

Log source	Log volume	Value for threat detection
Firewalls allowed traffic	High	Medium-low
Firewalls denied traffic	High	Low
Firewalls VPN	Medium	Medium
Intrusion prevention/detection system (IPS/IDS)	Low	High
URL filtering / URL access	High	Medium
Email security	Low	High
Windows Security events	High	High
AAA (Radius, terminal access controller access control system [TACACS])	Low	High
Cloud IAM	Medium	High
LAN/WAN	Low	Low
Cloud PaaS	High	Medium
Website access	High	Low
Database audit tools	Low	High
Endpoint Detection and Response (EDR) – alerts/incidents	Low	High
Endpoint Detection and Response (EDR) – Raw endpoint activity logs	High	High
Cloud security controls	Low	High
Vulnerability scanning	Low	Medium
File integrity	Low	High
Privileged Access Management (PAM)	Low	Medium
SD-WAN	High	Low
Multi-Factor Authentication (MFA)	Medium-Low	High

Detailed Analysis Examples

Firewall-allowed Traffic

- Volume: High. An organization with 1,000 end users may generate 20 GB–30 GB of firewall log per day.

- Log data collected: Source IP, destination IP, protocol/application, traffic (bytes/packets), firewall action (allow), firewall interfaces, firewall rule, and user.

- Sample alert rules: Matching with known malicious IPs, geolocation, volume of traffic (anomalies), number of concurrent connections, potential command and control (C&C) beaconing, and applications/ protocols used.

- Value for threat detection: Low. Malicious actors rattle the cages of the internet, endlessly scanning for vulnerabilities. This, combined with dynamic IP address assignment, creates a scenario where one can expect several hit matches from known malicious IPs inbound. Matches outbound are noteworthy but, without other correlating log sources, are difficult to investigate.

- Optimization options. Firewall logging rules adjustments, filtering of types of log entries recorded by firewalls, filtering at syslog collector level, and adjustment of firewall logging format.

EDR (alerts/incidents)

- Volume: Low. Only potentially malicious activities are logged and should not create a significant volume of logs.

- Log data collected: Endpoint host name, username, process (child/ parent), hashes (SHA256, MD5), file names, operations, network connections (destination IP, host name, URL), remediation status, severity, confidence, threat name, threat description URL, and recommended remediation.

- Sample alert rules: Incidents matching high and critical severities, incidents with no remediation, repetitive remediations, and incidents affecting multiple users/hosts.

- Optimization options: Tuning of EDR solution, exclusion lists of known or benign incidents.

Windows Security Events

- Volume: Medium–High. Depending on the organization's requirements around collecting data for user identity governance, the required logs may vary.

- Log data collected: User sign in/sign out, user creation, user disabled/enabled, password changes, group creation, group membership, process creation, and file access audit.

- Sample alert rules: Anomalous user logins, patterns matching brute force attacks/ password spraying, user creation, activation of disabled users, addition to high-sensitive or privileged groups, and suspicious processes.

- Optimization options: Filtering of security events collected, configuration of logging policies (via group policies), filtering at the connector level, and filtering at the SIEM solution-level.

If certain log sources have the potential to generate large volumes of data (and associated analytics costs), additional compensating controls can be considered to provide the same or similar visibility at a lower cost. For example, detections around the processes collected from Windows Security Event logs are typically overlapping with detections provided by EDR solutions. An organization can decide to disable the logging of process creation on Windows servers based on Microsoft Defender for Servers being deployed and covering potential malicious activities based on monitoring of processes created. Disabling the process creation audit may result in 50%-70% Windows Security event log volume. In many cases, the compensating controls offer superior detection alerts and are more frequently updated by the solution vendor.

Microsoft Sentinel and the log collectors it provides offers several levels where the raw logs can be filtered. These include the log source itself, the on-premises syslog collector, the Microsoft Sentinel data connector or Microsoft Sentinel itself. Each of these levels can be used to optimize the log ingestion for a good balance between detection value and ingestion costs. One of the most versatile, powerful, and yet easy to configure is the use of the data ingestion and transformation option, available as a feature of any Log Analytics Workspace. It provides a high level of control over the incoming data and can be configured centrally and applied in a very timely manner without any change in the logging configuration for the monitored devices. See Custom data ingestion and transformation in Microsoft Sentinel for more details.

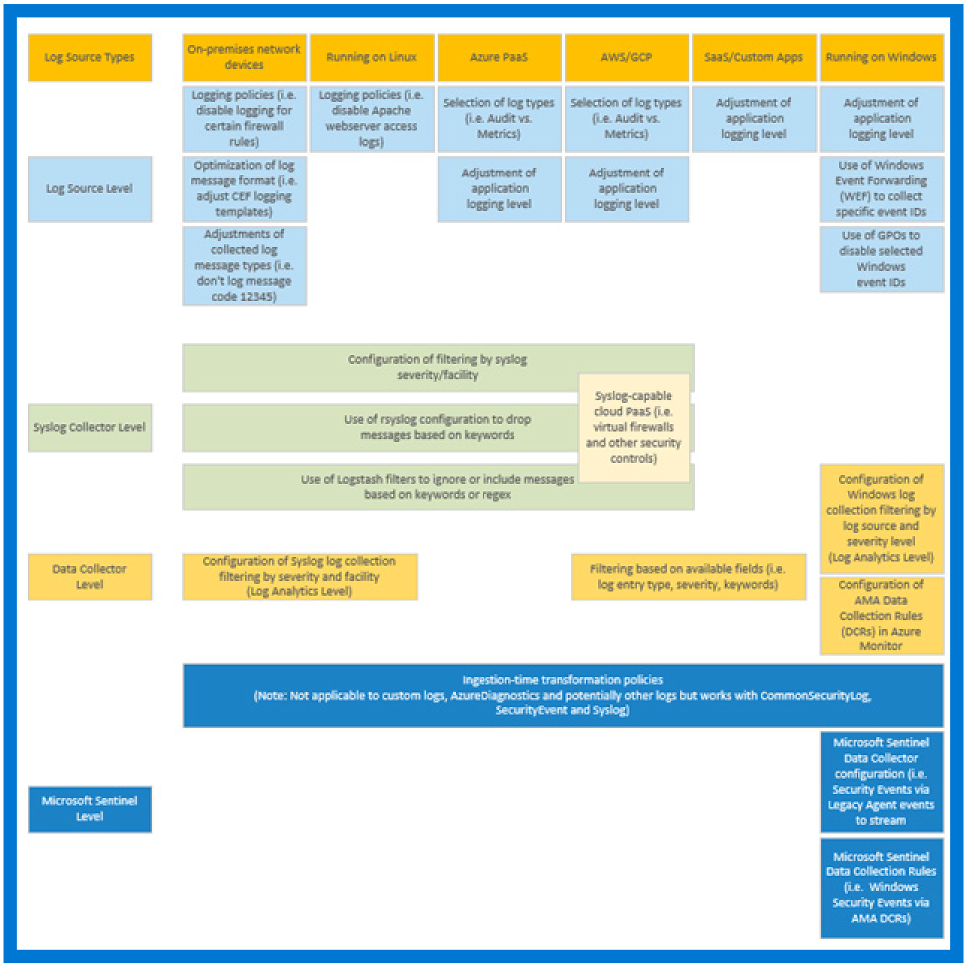

Log Source Types
On-premises network devices
Running on Linux
Azure PaaS
AWS/GCP
SaaS/Custom Apps
Running on Windows
Logging policies (i.e. disable logging for certain firewall rules)
Logging policies (i.e. disable Apache webserver access logs)
Selection of log types (i.e. Audit vs. Metrics)
Selection of log types (i.e. Audit vs. Metrics)
Adjustment of application logging level
Adjustment of application logging level
Log Source Level
Optimization of log message format (i.e. adjust CEF logging templates)
Adjustment of application logging level
Adjustment of application logging level
Use of Windows Event Forwarding (WEF) to collect specific event IDs
Adjustments of collected log message types (i.e. don't log message code 12345)
Use of GPOs to disable selected Windows event IDs
Configuration of filtering by syslog severity/facility
Syslog Collector Level
Use of rsyslog configuration to drop messages based on keywords
Syslog-capable cloud PaaS (i.e. virtual firewalls and other security controls)
Use of Logstash filters to ignore or include messages based on keywords or regex
Configuration of Windows log collection filtering by log source and severity level (Log Analytics Level)
Data Collector Level
Configuration of Syslog log collection filtering by severity and facility (Log Analytics Level)
Filtering based on available fields (i.e. log entry type, severity, keywords)
Configuration of AMA Data Collection Rules (DCRs) in Azure Monitor
Ingestion-time transformation policies (Note: Not applicable to custom logs, AzureDiagnostics and potentially other logs but works with CommonSecurityLog, SecurityEvent and Syslog)
Microsoft Sentinel Data Collector configuration (i.e. Security Events via Legacy Agent events to stream
Microsoft Sentinel Level
Microsoft Sentinel Data Collection Rules (i.e. Windows Security Events via AMA DCRs)

Fig. 31.1. Methods of log filtering available for Microsoft Sentinel

In Our Experience,

- Legacy log collection/retention policies enforce the collection of high-volume, low- value logs. Re-evaluate and consider updating corporate security policy.
- Misunderstanding logged data quality. Review both your log sources and inherent log fields for value in threat detections.
- Compliance may require long log retention volumes (e.g., 1 year for PCI logs). Consider reviewing your architecture for ways to reduce compliance scope.
- Log collection and retention technologies, especially syslog products, have a variety of filtering and parsing options. Assess your architecture to ensure your current solution meets your needs.
- Obtain senior management support to honestly review the value of your current logged sources. Staff may otherwise be unwilling to assume the risk of offering tuning recommendations.
- Custom data ingestion and transformation in Microsoft Sentinel provides a very effective way to manage log filtering centrally

Budgeting for Microsoft Sentinel Costs

With appropriate attention to cost management, Microsoft Sentinel is a highly cost-effective SIEM solution and provides substantial benefits over physical and premises-based virtualized solutions. However, as with the move of all IT infrastructure to the cloud, assumptions and modes of operation that held true for the on-premises world need to be re-examined and potentially adjusted.

Azure data ingestion costs can be difficult to project, particularly for data sources that produce logs at variable rates based on factors external to the organization. For example, log data from internet-facing web application firewalls may see a significant spike in volume based on the sudden popularity of a company website, causing Azure ingestion costs to spike in parallel. This may not be an event that the IT organization could have foreseen but will need to be reckoned with for subsequent budget cycles.

Case Study: Software company

During the deployment of Microsoft Sentinel and against recommendations, an organization decided to log inbound denies on an internet-facing firewall logging into Microsoft Sentinel. Shortly after deployment, the organization was targeted by a DDoS attack, causing a significant and sudden increase in log volumes. The cost for data ingestion quickly spiked, and adjustments were made rapidly to adjust logging levels. While there is value in a record of attacker IP addresses, ingesting large volumes of log data to a cloud service like Microsoft Sentinel is likely not the most cost-effective method of obtaining this information.

For more on “Economic Denial of Sustainability” attacks, visit:
https://www.managedsentinel.com/edos-attack-azure-sentinel/

Enumerating In-scope Log Sources and Phasing Deployment Projects Over Time

A complete and comprehensive view of all organizational assets is often an unattainable goal for security teams. However, having a clear view of the in-scope and out-of-scope log sources at the outset of a project will be critical to managing costs for Microsoft Sentinel over multiple budget cycles. Phasing in onboarding new log sources over time is also advised to help understand and control ingestion costs rather than experiencing a significant cost spike and inevitable pull-back.

An example of project phasing focusing on users, endpoints, and cloud infrastructure:

Phase	Log sources onboarded
1	Identity (Azure AD, Windows AD), EDR detections/incidents (Microsoft Defender products and other similar solutions)
2	Security infrastructure (firewalls, IDS/IPS, NAC)
3	Hybrid infrastructure (servers, L2/L3 devices, storage, PaaS)
4	SaaS applications (I.e. SalesForce, Workday)

Collecting Log Samples

Collecting log samples from in-scope log sources is an important and often overlooked step in preparing scope documentation and preparing multi-year budgetary requirements for Microsoft Sentinel deployments. Analysis of log types, logging behavior, and data volumes for in-scope sources will serve several purposes beyond cost analysis, including baselining and tuning analytic rules that may be set up for detecting anomalous logging behavior or developing log parsing in Log Analytics. Log samples will likely need to be collected from a variety of sources, which should be included in the pre-engagement schedule by the project team.

Common log sources that may be required to sample for analysis include:

- OS security and event logs requiring AMA
- Syslog or CEF logs from on-premises infrastructure
- Azure Diagnostics logs
- Logs from SaaS applications available via API calls

Ongoing Cost Monitoring and Evaluation

Monitoring costs for log ingestion on an ongoing basis is a critical task in ensuring the ongoing viability of the Microsoft Sentinel platform for any organization. Azure cost management, at a general level, is a broad and well-covered topic. We will provide more targeted recommendations that can be implemented by security teams within the Microsoft Sentinel solution.

Using KQL Queries

In addition to using KQL queries for analytic and threat-hunting rules, they can also be developed to monitor unusual deviations in log ingestion volumes or types and provide alerts for further investigation. Many log source types produce data volumes that follow predictable patterns, such as following business hours or seasonal activity. Creating statistical models and configuring analytic rules to provide notifications for deviations from normal log volumes is essential. It is a simple way to provide immediate notice to administrators or analysts to review spikes in data ingestion before they make a meaningful impact on a monthly Azure invoice.

Trending and visualizations in workbooks or in external tools such as PowerBI can provide an effective way to spot upward trends in data ingestion by source and allow system owners to take action to keep costs in line with budgets.

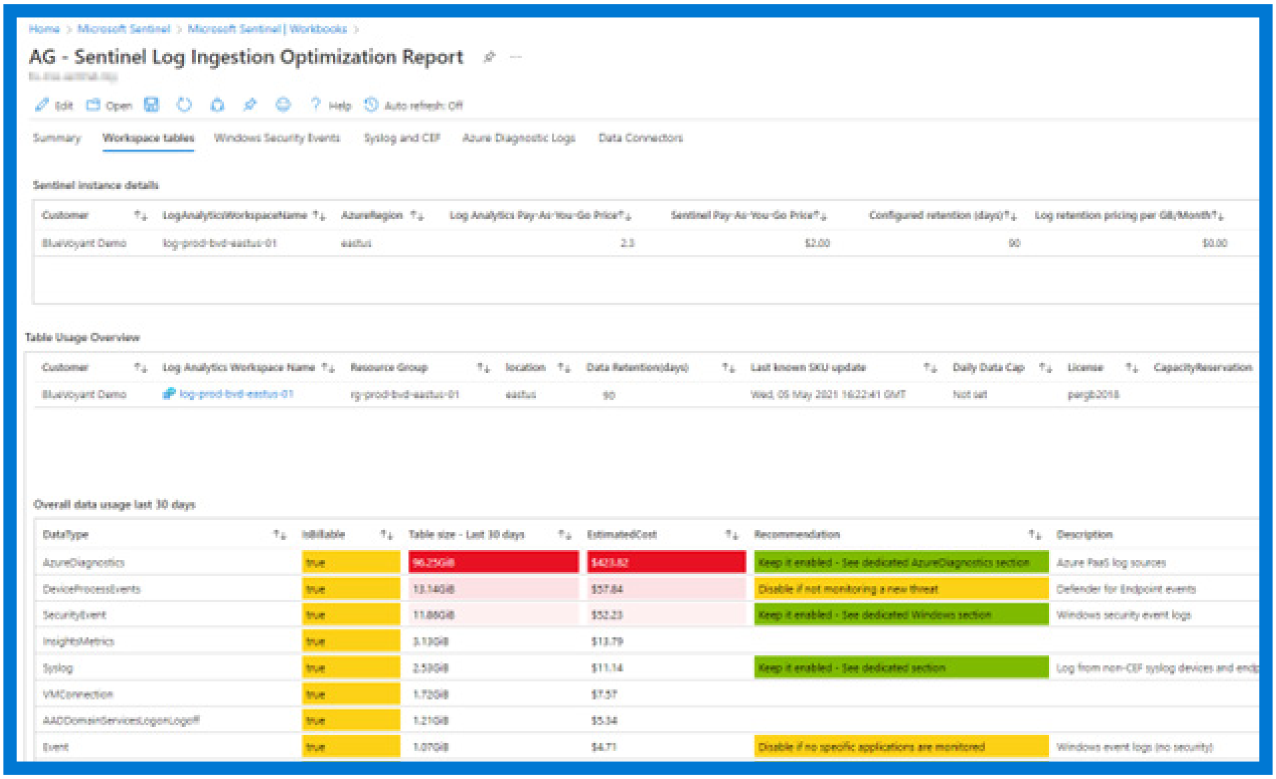

AG - Sentinel Log Ingestion Optimization Report

Sentinel instance details

Customer	LogAnalyticsWorkspaceName	AzureRegion	Log Analytics Pay-As-You-Go Price*	Sentinel Pay-As-You-Go Price*	Configured retention (days)*	Log retention pricing per GB/Month*
BlueVoyant Demo	log-prod-bvd-eastus-01	eastus	2.3	$2.00	90	$0.00

Table Usage Overview

Customer	Log Analytics Workspace Name	Resource Group	location	Data Retention(days)	Last known SKU update	Daily Data Cap	License	CapacityReservation
BlueVoyant Demo	log-prod-bvd-eastus-01	rg-prod-bvd-eastus-01	eastus	90	Wed, 05 May 2021 16:22:41 GMT	Not set	pergb2018	

Overall data usage last 30 days

DataType	IsBillable	Table size - Last 30 days	EstimatedCost	Recommendation	Description
AzureDiagnostics	true	96.25GiB	$423.82	Keep it enabled - See dedicated AzureDiagnostics section	Azure PaaS log sources
DeviceProcessEvents	true	13.14GiB	$57.84	Disable if not monitoring a new threat	Defender for Endpoint events
SecurityEvent	true	11.86GiB	$52.23	Keep it enabled - See dedicated Windows section	Windows security event logs
InsightsMetrics	true	3.13GiB	$13.79		
Syslog	true	2.53GiB	$11.14	Keep it enabled - See dedicated section	Log from non-CEF syslog devices and endp
VMConnection	true	1.72GiB	$7.57		
AADDomainServicesLogonLogoff	true	1.21GiB	$5.34		
Event	true	1.07GiB	$4.71	Disable if no specific applications are monitored	Windows event logs (no security)

Fig. 31.1. Methods of log filtering available for Microsoft Sentinel

Conclusion and Resources

We hope that this whitepaper has been informative and helpful for security practitioners and CISOs pursuing deployments of Microsoft Sentinel. The contents and recommendations provided have been developed by our team with hundreds of Microsoft Sentinel deployments around the globe in a variety of industries, and as a Microsoft Intelligent Security Association (MISA) MSSP Member and Threat Protection and Cloud Security Advanced Specialization partner.

Many thanks to the teams at Microsoft that have supported us, and the inspiring and forward-thinking customers we have the privilege to work with every day.

NOTE: Case studies are based on customer engagements, The names and company details have been altered to ensure their anonymity.

Additional resources

Microsoft training paths

Microsoft Sentinel in the Azure Marketplace

Produced in partnership with BlueVoyant LLC

See More

BlueVoyant