Microsoft

Microsoft Sentinel Technical Playbook for MSSPs

How to deploy Microsoft Sentinel as a Managed Security Services Provider

Published: November-2022, Revision: V1.5.1

Authors:
Javier Soriano (Senior Program Manager, CxE Sentinel)
Ty Balascio (Senior Program Manager, CxE)
Yaniv Shasha (Senior Program Manager, CxE Sentinel)
Chris Boehm (Senior Program Manager, CxE Sentinel)
Chi Nguyen (Program Manager 2, CxE Sentinel)
Paul Cullimore (Senior Business Strategy Manager)
Edi Lahav (Principal PM Manager, CxE Sentinel)
Richard Diver (Senior Business Strategy Manager)
Waleed Bedair (Senior Program Manager, CxE Partners security)
Gary Bushey (Senior Program Manager, CxE Sentinel)
Margaret Mwaura (Program Manager, CxE Sentinel)
Didier Danloy (Program Manager, CxE Sentinel)
Jeremy Tan (Senior Program Manager, CxE Sentinel)

Disclaimer
This document is for informational purposes only. MICROSOFT MAKE NO WARRANTIES, EXPRESSED, IMPLIED, OR STATUTORY, AS TO THE INFORMATION IN THIS DOCUMENT.

This document is provided "as-is." Information and views expressed in this document, including URL and other internet website references, may change without notice. You bear the risk of using it.
This document does not provide you with any legal rights to any intellectual property in any Microsoft product. You may copy and use this document for your internal, reference purposes.
Some examples depicted herein are provided for illustration only and are fictitious. No real association or connection is intended or should be inferred.

NOTE: Certain recommendations in this white paper may result in increased data, network, or compute resource usage, and may increase your license, subscription, or ingestion costs.

Introduction

Thank you for considering Microsoft Sentinel as the heart of your Managed security service providers or SIEM integration practice. Microsoft Sentinel is a cloud-native security information and event manager (SIEM) platform that uses built-in AI, user and entity behavior analytics and threat intelligence from many sources to help analyze large volumes of data across an enterprise—fast. Microsoft Sentinel is also a platform that provides Security orchestration, automation, and response (SOAR) to help SOC analysts to automatically respond to known attacks and help them focus on new emerging threats.
Microsoft Sentinel aggregates data from all sources, including users, applications, servers, and devices running on-premises or in any cloud, letting you reason over millions of records in a few seconds. It includes built-in connectors for easy onboarding of popular security solutions. Collect data from any source with support for open standard formats like CEF and Syslog.

Updates in this Version Release

This section lists all the major features which were added in Microsoft Sentinel and related services, since this original document was released. It also contains list of items that have been added or updated since the last update of the document. You can click on the features/updates to reach directly to the section. If you're new to this architecture reference guide, we recommend you review the whole document.

Search & Archive

- Archive, Search and Restore page 24

Logs Ingestion

- AMA/MMA features parity + Migration path page 29
- Codeless Connector Platform (CCP) page 30
- Basic logs page 36

Pipeline Transformation

- Ingestion time transformation (filtering) page 35
- Updated Logstash connector page 36

Normalization

- Advanced Security Information Model (ASIM) page 34

Sentinel Health

- Analytics rules health page 47
- Data connectors health page 33

Security coverage

- New MITRE coverage blade page 43

Onboarding and Management

- Repositories for CI/CD page 53

New & updated technical content

- New Migration cookbook page 3
- Updating links in the training and community resources page 57

Appendix / Changelogs:

Type	Previous values	Updates values
Cross workspace Incident view	10	100
LA Cluster min daily ingestions	1 TB	500 GB
Number of alerts triggered by a single analytics rule in “Trigger an alert for each event” mode	20	150
Max file size to import watchlist items	Local storage: 3.8MB	Blob storage: 500MB
Comments maximum size	3000 characters	30 000 characters
Automation rule trigger	Incident creation	Incident creation, incident update, alert creation

Target audience

This document informs Microsoft partners researching how to integrate Microsoft Sentinel into their portfolio of services. It is written through the lens of Implementers & SOC architects who seek a distilled technical walkthrough of:

Microsoft Sentinel's capabilities
Technical dependencies
Data collection models
Multi-tenant management
Threat detection & analytics
Investigation processes
Strategies for automated response
Activity summaries and reports
Cost models and data storage

Beyond Microsoft Security Services Providers (MSSPs), this document aims to guide large organizations and institutions who operate security operations within environments requiring multi-tenant architectures.
This document is relevant both to MSSPs that are new to Sentinel and those who are already experienced with Sentinel. Experienced MSSP may want to focus on the following chapters:

- Architecture
- Sizing & Pricing/Cost - Long term storage options summary
- Automation/SOAR - MSSPs design considerations for automation rule and playbooks
- Analytic Rules
- Microsoft Sentinel Workbooks - Intellectual property protection
- DevOps / CI/CD automation

The Microsoft Sentinel value for MSSPs

Managed security service providers design business models based on scale and efficiency. The expense incurred by onboarding the first few clients define the templates for future deployments. Operating efficiently requires analyzing each use case required within a customer segment and applying a reusable process with a bias toward automation.

Your customers already understand the importance of employing rigorous security oversight. Introducing any change to tools or processes requires an examination of workload impact, system compatibility, and improved ROI. As compliance requirements, environmental complexity, and threat landscapes expand they rely on you for providing exhaustive security services as well as summaries and outcomes of the threats managed. Microsoft Sentinel’s architecture anticipates these requirements and scales to address them with extensive built-in capabilities complemented with rich extensibility for integrating your existing business processes.

Efficient customer onboarding

You can see value immediately with an expedited and automated customer onboarding of Microsoft Sentinel. Closing the time gap between customer commitment and fully deploy Microsoft Sentinel instance enables stronger security and lower costs for your customers and your team. It can be deployed in minutes using templates customized for your needs. Likewise, well tested data connectors cover a wide spectrum of popular resources, and a rich library of automation is available for nearly any scenario.

A new **migration guide** to help migrating from three major third-party SIEMs (ArcSight, Splunk and QRadar) to Microsoft Sentinel is now available in our official documentation.
This guide focuses on the following areas:

- Planning your migration
- Migrating detection rules
- Migrating SOAR
- Migrating historical data
- Converting dashboards to workbooks
- Upgrading SOC processes

The guide also provide a comparison and other considerations to help you select the correct Azure Platform to host historical data and how to ingest historical data to Azure.

There is also a dynamic Workbook to help you track the progress of your migration, including process and track different artifacts Microsoft Sentinel provides - data connectors, analytics rules, workbooks, automation and UEBA.

Optimizing system administration costs

Each customer operating environment contains unique characteristics. Tuning is necessary to ensure signals from all resources are collected, network topologies are understood, and monitors are in place. Moreover, resources are in a state of constant change. Microsoft Sentinel, in conjunction with Azure's governance capabilities, can adjust to these changing conditions programmatically.

Scaling SOC operations

As industries seek to improve efficiency, a larger dependence is placed on digitally transforming business processes. In addition to driving larger data estates, the ever-expanding variety of computing resources deployed result in exponential growth in threat signals requiring analysis. MSSPs differentiate themselves by filtering and distilling these signals down to the most important and most actionable. As a cloud native SIEM, Microsoft Sentinel's threat detection analytics capabilities, incident investigation and response and SOAR provides a strong foundation toward this need. Combined with the specialized tuning of your SOC team, your client's signal-to-noise ratio enriches.

Reutilization of intellectual property

From initial deployment through automated threat responses, all MSSPs seek opportunities for re-use across multiple client environments. Deployment automation, detection rule creation and tuning, visualizations, investigation workbooks, Threat Intelligence and client escalation workflows are but some of the many use cases ripe for scale. Not only will these efficiencies lead to more agility across your client base, but these investments also elevate your staff from repetitious, often mundane tasks.

MSSPs can use this guide as a reference for understanding Microsoft Sentinel's capabilities, planning, and budgeting considerations, and technical resources; formatted from a MSSP's point of view. This document complements official documentation, Certification learning paths, and community resources. Links to those resources provided in the last chapter.

Fundamentals

Microsoft Cloud fundamentals

Microsoft's Cloud offerings consist of Microsoft 365, Microsoft Azure and Microsoft Dynamics 365.

While Microsoft Sentinel supports signals from any workload, in this topic we'll focus mainly on Office 365 and Microsoft Azure cloud security offerings. It's important to understand the hierarchy of the major entities of these cloud offerings as it has a significant impact on the architecture of MSSP deployment.

Microsoft provides a hierarchy of Organizations, Tenants, Subscriptions, and User accounts for consistency of identities and billing across the cloud offerings.

- **Organization -** represents a business entity that is using Microsoft cloud offerings, typically identified by one or more public Domain Name System (DNS) domain names, such as contoso.com. The organization is a container for subscriptions.
- **Tenant -** a regional location that houses the servers providing cloud services. It is a single instance of Azure Active Directory (AAD)
- **Subscriptions –** A subscription is an agreement with Microsoft to use one or more Microsoft cloud platforms or services, for which charges accrue based on either a per-user license fee or on cloud-based resource consumption. Organizations can have multiple subscriptions for Microsoft's cloud offerings.
- **User accounts -** User accounts for all of Microsoft's cloud offerings are stored in an Azure Active Directory (Azure AD) tenant, which contains user accounts and groups. An Azure AD tenant can be synchronized with an existing on-prem Active Directory Domain Services (AD DS) accounts using Azure AD Connect, a Windows server-based service. This is known as directory synchronization.

Following is an example of an Azure AD Tenant that contains user accounts and groups for various Microsoft's Cloud Services:

- **Tenants –** For SaaS cloud offerings, the tenant is the regional location that houses the servers providing cloud services. For example, the Contoso Corporation chose the European region to host its Microsoft 365, EMS, and Dynamics 365 tenants for the 15,000 workers in their Paris headquarters. Azure PaaS services and virtual machine-based workloads hosted in Azure IaaS can have tenancy in any Azure datacenter across the world. You specify the Azure datacenter, known as the location, when you create the Azure PaaS app or service or element of an IaaS workload.
- An **Azure AD tenant** is a specific instance of Azure AD containing accounts and groups. Paid or trial subscriptions of Microsoft 365 or Dynamics 365 include a free Azure AD tenant. This Azure AD tenant does not include other Azure services and is not the same as an Azure trial or paid subscription.
- **Licenses -** For Microsoft's SaaS cloud offerings, a license allows a specific user account to use the services of the cloud offering. You are charged a fixed monthly fee as part of your subscription. Administrators assign licenses to individual user accounts in the subscription. In the example below, the Contoso Corporation has a Microsoft 365 E5 subscription with 100 licenses, which allows to up to 100 individual user accounts to use Microsoft 365 E5 features and services.

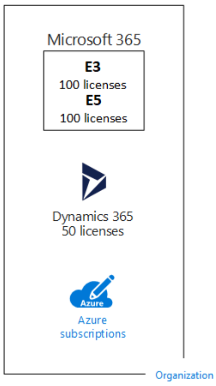

Connecting it all together

- An organization can have multiple subscriptions
- A subscription can have multiple licenses
- O365 licenses can be assigned to individual user accounts
- User accounts are stored in an Azure AD tenant

Multiple Microsoft cloud offering subscriptions can use the same Azure AD tenant that acts as a common identity provider. A central Azure AD tenant that contains the synchronized accounts of your on-premises AD DS provides cloud-based Identity as a Service (IDaaS) for your organization.

In the following diagram the common Azure AD tenant used by Microsoft's SaaS cloud offerings, Azure PaaS apps, and virtual machines in Azure IaaS that use Azure AD Domain Services. Azure AD Connect synchronizes the on-premises AD DS forest with the Azure AD tenant.

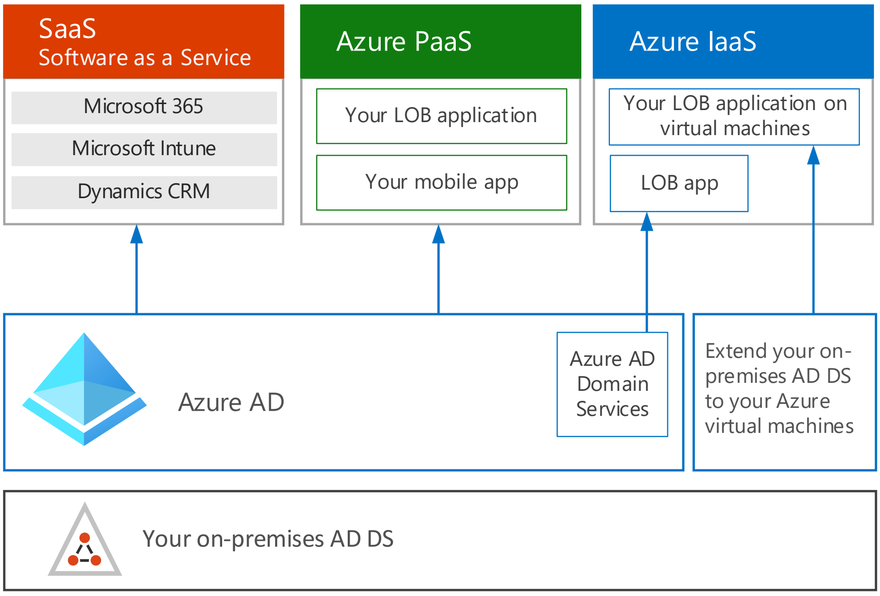

Relevant resources:

Microsoft’s hierarchy for Cloud Offerings

Microsoft Cloud Identity for Enterprise Architects

Microsoft 365 identity models and Azure Active Directory

Quickstart: Setup a tenant

Azure fundamentals

The Azure Cloud platform consists of more than 200 products and Cloud Services. It allows customers to build and run their solutions across multiple clouds, on-premises and at the edge with a variety of tools and frameworks.

Azure Services are available globally and partners can choose the best region for their customers' needs based on technical and regulatory considerations: service capabilities, data residency, compliance requirements, and latency.

A **Region** is a set of datacenters deployed within a latency-defined perimeter and connected through a dedicated regional low-latency network. Azure gives you the flexibility to deploy applications where you need to, including across multiple regions to deliver cross-region resiliency.

An **Availability Zone** is a high availability offering that protects your applications and data from datacenter failures. Availability Zones are unique physical locations within an Azure region. Each zone is made up of one or more datacenters equipped with independent power, cooling, and networking.

To ensure resiliency, there's a minimum of three separate zones in all enabled regions. The physical separation of Availability Zones within a region protects applications and data from datacenter failures. Zone-redundant services replicate your applications and data across Availability Zones to protect from single-points-of-failure.

To achieve comprehensive business continuity on Azure, build your application architecture using the combination of Availability Zones with Azure region pairs. You can synchronously replicate your applications and data using Availability Zones within an Azure region for high-availability and asynchronously replicate across Azure regions for disaster recovery protection.

Azure resources reside in **Subscriptions** and each resource is assigned to a **Resource Group**. For cases where many subscriptions exist a more efficient way is required manage access, policies, and compliance for those subscriptions.

A **Resource Group** is a container that holds related resources for an Azure solution. The resource group includes those resources that should be managed as a group. The decision of which resources belong in a resource group will be based on what makes the most sense for the customer.

Azure Resource Manager (ARM) is the deployment and management service for Azure. It provides a management layer that enables you to create, update, and delete resources in your Azure account. It includes management features, like access control, locks, and tags, to secure and organize the resources after deployment.

Azure management groups provide a level of scope above subscriptions. You organize subscriptions into containers called "management groups" and apply your governance conditions to the management groups. All subscriptions within a management group automatically inherit the conditions applied to the management group. Management groups give you enterprise-grade management at a large scale no matter what type of subscriptions you might have. All subscriptions within a single management group must trust the same Azure Active Directory tenant.

Azure Policies help to enforce organizational standards and to assess compliance at-scale. It also helps to bring your resources to compliance through bulk remediation for existing resources and automatic remediation for new resources. Common use cases for Azure Policy include implementing governance for resource consistency, regulatory compliance, security, cost, and management. Policy definitions for these common use cases are already available in your Azure environment as built-ins to help you get started.

The following diagram shows an example of creating a hierarchy for governance using management groups.

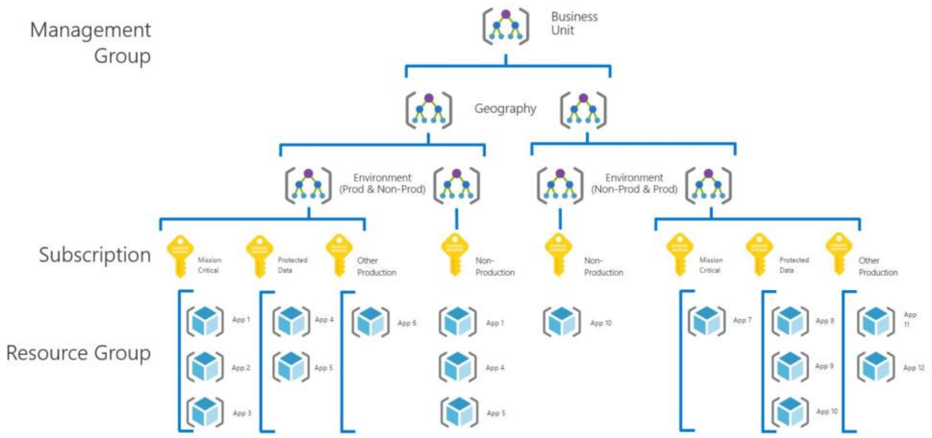

Relevant resources:

[Azure Security best practices]

[Microsoft Cloud Adoption Framework for Azure]

[Regions and Availability Zones in Azure]

[Overview of the reliability pillar]

[Products available by region]

[What are Management Groups?]

[What is Azure Policy?]

[What is Azure Resource Manager?]

[What is Azure role-based access control (RBAC)?]

Log Analytics fundamentals

Log Analytics is a tool in the Azure portal to edit and run log queries from data collected from various data sources and interactively analyze their results. Log Analytics queries can be used to retrieve records matching particular criteria, identify trends, analyze patterns, and provide a variety of insights into your data.

Azure Monitor, and its Log Analytics module, is the underlying log management platform powering Microsoft Sentinel. As such, any source that sends logs to Azure Monitor or Log Analytics inherently supports Microsoft Sentinel. As soon as Microsoft Sentinel is enabled for a Log Analytics workspace all collected sources are available in Microsoft Sentinel for further analysis, hunting, threat mitigation and additional actions taken by the SOC analyst.

Data collected by Azure Monitor Logs is stored in one or more **Log Analytics workspaces**. The workspace defines the geographic location of the data, access rights defining which users can access data, and configuration settings such as the pricing tier and data retention. At least one workspace must be created to use Azure Monitor Logs. A single workspace may be sufficient for all your monitoring data, or you may choose to create multiple workspaces depending on the requirements. For example, a requirement to split operational logs and security logs would result in 2 separate workspaces. The security logs workspace will be Microsoft Sentinel enabled.

Data is retrieved from a Log Analytics workspace using a log query which is a read-only request to process data and return results. Log queries are written in **Kusto Query Language (KQL)**, which is the same query language used by Azure Data Explorer.

The **Azure Log Analytics agent** or **Microsoft Monitoring agent (MMA)** collects telemetry from Windows and Linux virtual machines in any cloud or on-premises machines and sends the collected data to your Log Analytics workspace in Azure Monitor and/or Microsoft Sentinel.

MMA collects the following data sources: Windows Events, Syslog, Performance, IIS logs and Custom logs. Log Analytics workspace configuration allows to decide which data sources will be collected by all connected agents. The **Azure Monitor Agent (AMA)** provides new capabilities such as Windows Events filtering, multi-homing support for Linux, improved manageability of collection settings with DCR (Data Collection Rules) and more. DCR allows you to specify what data should be collected, how to transform that data, and where to send that data with different configurations for different groups of servers for example. For more information about AMA, please refer to the technical docs: https://docs.microsoft.com/en-us/azure/azure-monitor/agents/azure-monitor-agent-overview.

Data Destinations - With both the older MMA or the newer AMA, the Windows agent can be multihomed to send data to multiple workspaces and System Center Operations Manager management groups. With the older Linux agent also referred as OMS agent, logs can only be sent to a single destination, either a workspace or management group. Azure Monitor Agent for Linux also supports multihomed configuration.

In the following example various Connectors and log collection methods are listed to ingest logs from on-premises and/or Cloud Services such as: Office365, Azure, AWS, GCP and more

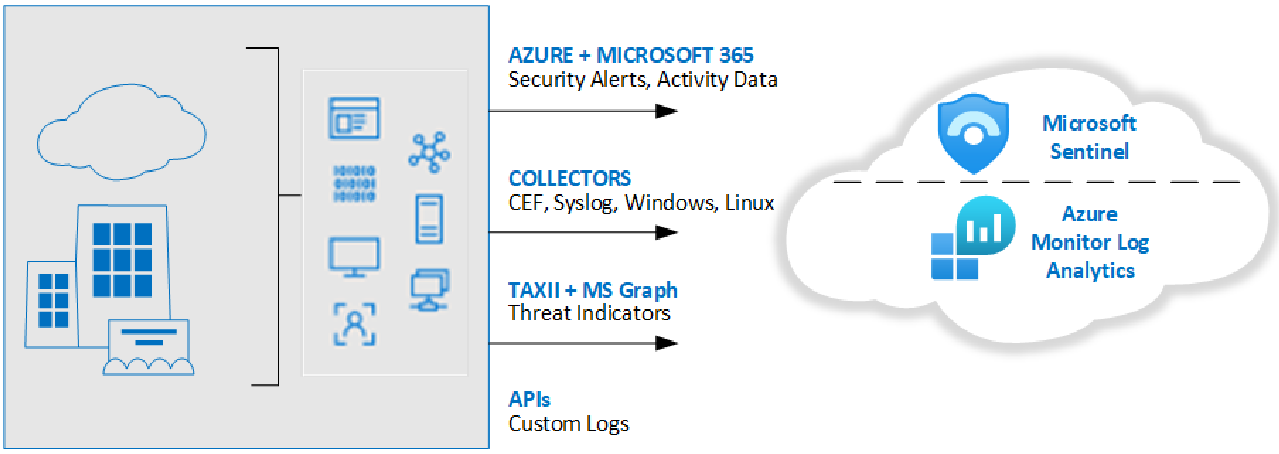

Data Collection in depth will be covered in the Data collection chapter.

Relevant resources:

Log Analytics Tutorial
Azure Monitor Logs overview
Overview of Log Analytics in Azure Monitor
Log Analytics agent overview
Quickstart: On-board Microsoft Sentinel
Kusto Query Language (KQL)
Managing and maintaining the Log Analytics agent

Architecture

This section focuses on the high-level design principles that need to be taken into account to provide a Microsoft Sentinel service as a Managed Security Services Provider (MSSP).

From now on in the document, we will describe the different concepts with one fictitious MSSP called Fabrikam, and two fictitious customers called Contoso and Wingtip.

Azure AD tenant topologies

In this section we will explain the different options available to structure your Azure AD tenants to provide a Microsoft Sentinel service to your end customers. There are fundamentally three options:

- Use a single identity for the MSSP internal services and applications and Azure management services.
- Use separate identities, one for MSSP internal services and applications, and a separate identity to manage your customers.
- Use identities on the client site. Not recommended due to the complexity of maintaining identities. A small portion of the MSSP' analysts may need to have identities on the client's site in order to provide services that Azure Lighthouse cannot handle.

Single identity model

In this approach, the MSSP has a single Azure AD that is used for internal operations and to serve customers. The following diagram describes this model:

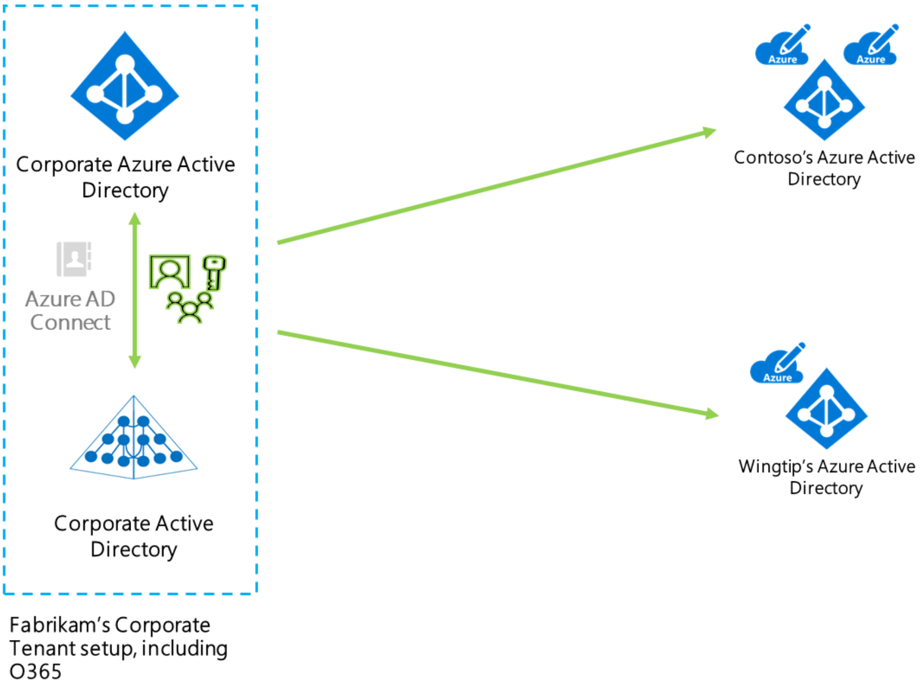
Corporate Azure Active Directory
Azure AD Connect
Corporate Active Directory
Fabrikam's Corporate Tenant setup, including O365
Azure
Azure
Contoso's Azure Active Directory
Azure
Wingtip's Azure Active Directory

For example, the user, User1, from Fabrikam's managed services team would have a single identity and set of credentials that he will use to access internal applications, like Office365, and also to access the customer environments that Fabrikam manages (we will discuss below what methods are available to access customer environments).

Pros: the great advantage of this model is that you only need to manage one identity for your users. That greatly simplifies your identity management processes and you also have a smaller identity footprint to secure. In our example, User1 will need to change his password just once whenever it expires and if he's terminated, Fabrikam will only have to disable one account.

Cons: the main disadvantage of this model is that you can't separate how your identities are secured by your internal IT to how you want to secure them for your end customer operations. For example, imagine that your internal IT policies require you to use MFA only when logging from outside your office, but your customer forces you to use MFA for each single login operation. This might force your internal IT team to implement new policies that they were not planning.

Multiple identities model

In this approach, the MSSP has two Azure AD tenants, one for its internal services and applications and a separate one to manage Azure customers. The following diagram describes this model:

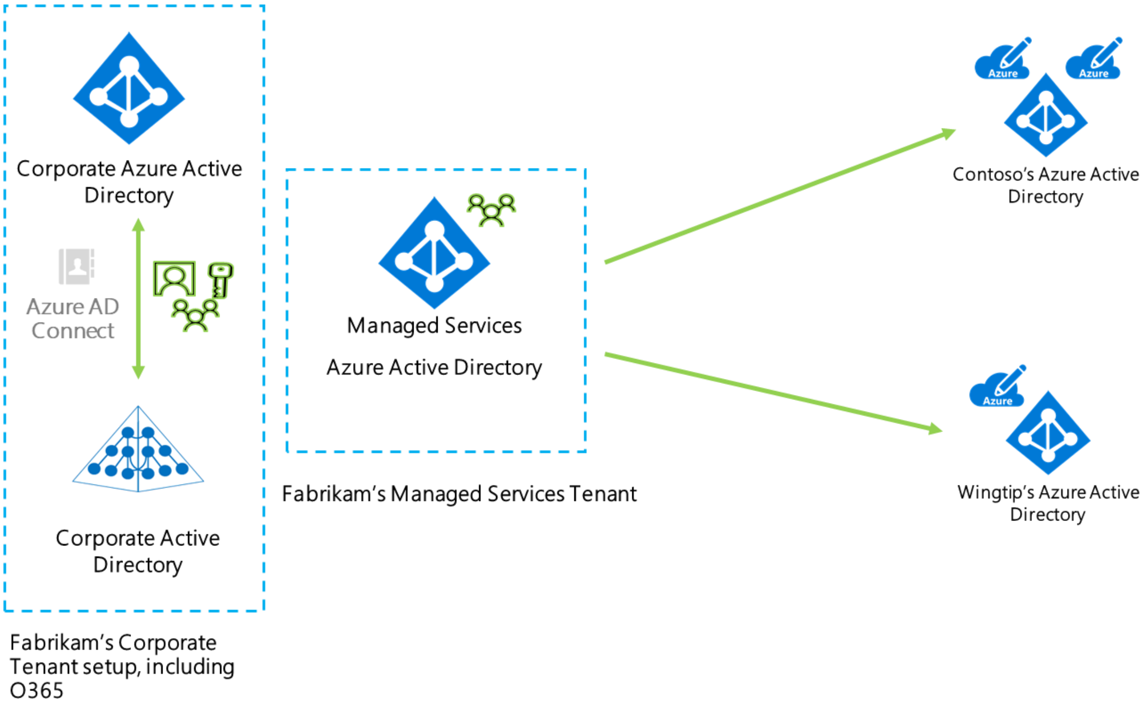
Corporate Azure Active Directory
Azure AD Connect
Corporate Active Directory
Fabrikam's Corporate Tenant setup, including O365
Managed Services Azure Active Directory
Fabrikam's Managed Services Tenant
Azure
Azure
Contoso's Azure Active Directory
Azure
Wingtip's Azure Active Directory

For example, the user, User1, from Fabrikam's managed service team, would have two identities, one to use his internal applications like email, and an additional identity that he will use when he needs to access customer environments (we will see in a later section how to access customer environments).

Pros: having two separate identities makes it easier for Fabrikam to adapt to different requirements from internal IT and from customers. For example, if Contoso requires Fabrikam to always use MFA when logging, they can just implement this policy in the managed services tenant and leave the corporate tenant unchanged.

Cons: Fabrikam needs to maintain multiple identities for users in the managed services team. Also, User1 will have to remember a separate set of credentials. If an employee is terminated, there needs to be a process in place to remove all related identities for the user.

Continue to the next section to understand how MSSP users can access customer environments.

Accessing the customer environment with Azure Lighthouse

Azure Lighthouse enables cross-tenant management, allowing for higher automation, scalability, and enhanced governance across resources and tenants.

This is the preferred method to access your customer environment because it allows you to manage customer resources as if they were in your own Azure AD tenant.

Azure Lighthouse is based on delegations. Each delegation contains three things: Identities, Roles and Scope.

- **Identities**: These are the identities (normally from the MSSP tenant) that will have access to customer resources. You can specify users, groups or service principals as the recipients of a delegation.
- **Roles**: these are the permissions that the identities will have when accessing customer resources. The roles that can be used here are all Azure Built-in roles with three exceptions captured here. Custom roles, roles with "DataActions" and Owner role are currently not supported. Also, you cannot grant Azure AD roles. See differences between Azure and Azure AD roles here.
- **Scope**: this indicates where the delegation will apply, valid scopes are *subscription* and *resource group*. Azure Lighthouse does not support delegations of subscriptions across a national cloud and the Azure public cloud, or across two separate national clouds.

An example delegation would be something like:

- **Identity**: User1 (user), *SecOps* (group) and *automation_spn* (service principal), all from Fabrikam Azure AD tenant
- **Roles**: *Microsoft Sentinel Contributor, Logic App Contributor*
- **Scope**: resource group *security_rg*

As you can see a single delegation can contain multiple identities and scopes. It can even contain multiple resource groups, but not multiple subscriptions.

There's two different ways to onboard a customer into your Lighthouse management: an ARM template or a Marketplace offer. Visit this article to get started.

In the context of Microsoft Sentinel, Azure Lighthouse can be used to manage the service across multiple customers. This is a high-level view of the setup:

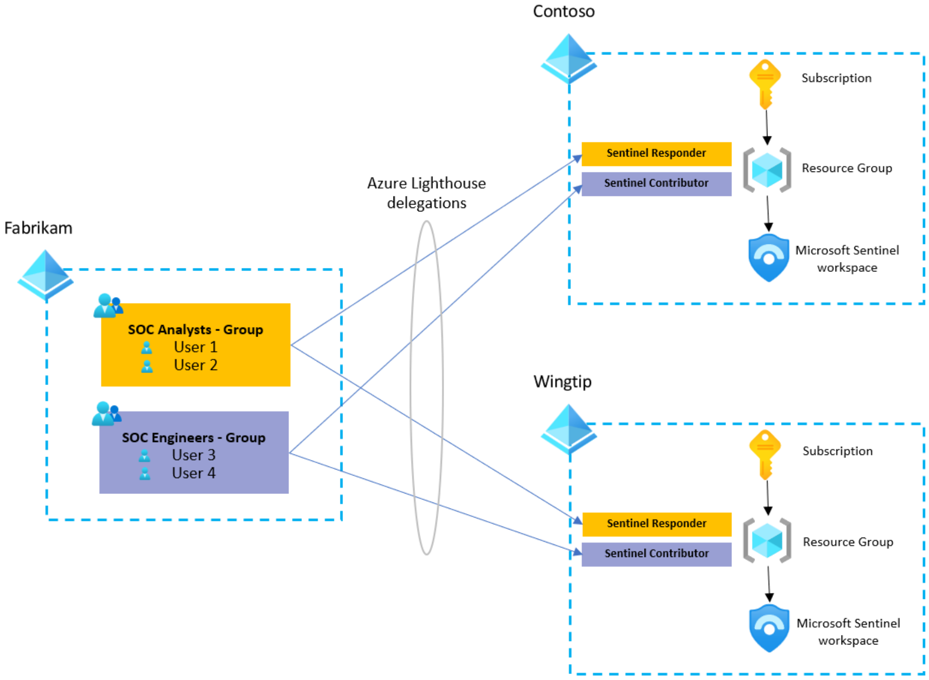

As you can see, in this case, the MSSP, Fabrikam, has two delegations for each customer, one for Engineers with the Microsoft Sentinel Contributor role and one for Analysts with the Microsoft Sentinel Responder role, all of them with delegated access at the resource group level where Microsoft Sentinel is located. This will effectively provide them with access to the whole resource group with the permissions included in the granted role.

Azure Lighthouse Onboarding

As already mentioned, there are two options to onboarding: ARM template or Azure Marketplace offer, being the latter preferred as it provides a very easy experience for end customers.

***NOTE** - there are some requirements before an MSSP can publish into the Azure Marketplace. The MSSP must have a silver or gold cloud platform competency level or be an Azure Expert MSP*.

Marketplace offers have an additional concept called **Plan**. A plan defines the service that you will provide to your customer. For example, you can have a marketplace offer to provide managed services for your customers, and within that offer several plans with different flavor: monitoring, backup and recovery, compliance, and full managed service.

In the Microsoft Sentinel context, you could have a Marketplace offer like this one:

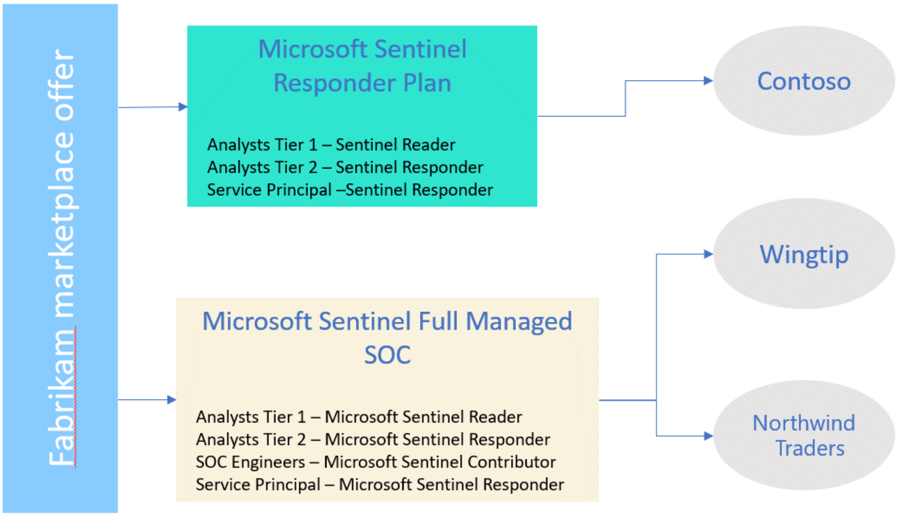
Fabrikam marketplace offer
Microsoft Sentinel Responder Plan
Analysts Tier 1 – Sentinel Reader
Analysts Tier 2 – Sentinel Responder
Service Principal –Sentinel Responder
Microsoft Sentinel Full Managed SOC
Analysts Tier 1 – Microsoft Sentinel Reader
Analysts Tier 2 – Microsoft Sentinel Responder
SOC Engineers – Microsoft Sentinel Contributor
Service Principal – Microsoft Sentinel Responder
Contoso
Wingtip
Northwind Traders

As you can see, inside each plan, you define the groups of users from your tenant that will have access to the customer environment and the permissions that will apply. The plan also includes the scope (resource group of subscription), but it is omitted for simplicity.

You can make these plans **public**, so everyone in Azure can see them, **or private**, if you only want a subset of customers to have access. This would allow for the option to create plans targeted just to specific audiences, like a particular customer or a vertical. Here is an example:

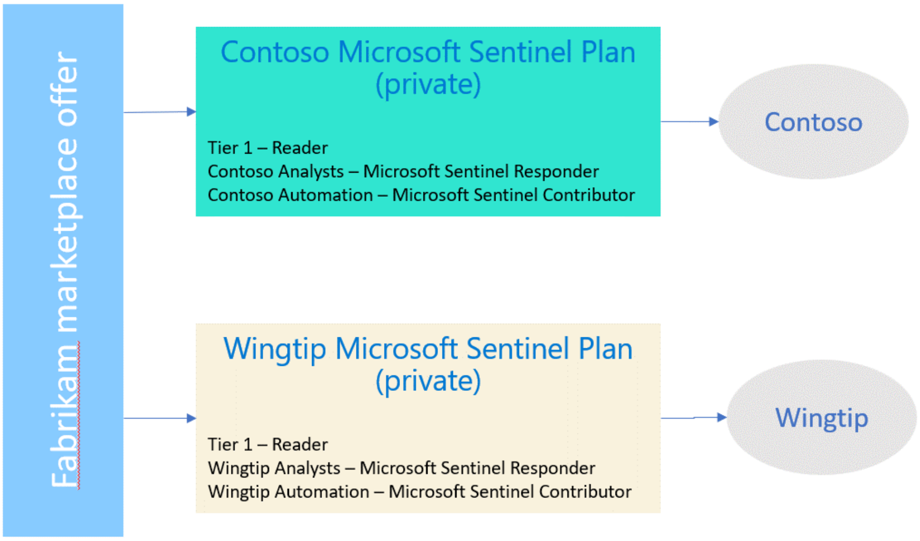
Fabrikam marketplace offer
Contoso Microsoft Sentinel Plan (private)
Tier 1 – Reader
Contoso Analysts – Microsoft Sentinel Responder
Contoso Automation – Microsoft Sentinel Contributor
Contoso
Wingtip Microsoft Sentinel Plan (private)
Tier 1 – Reader
Wingtip Analysts – Microsoft Sentinel Responder
Wingtip Automation – Microsoft Sentinel Contributor
Wingtip

Azure Lighthouse integration with Azure AD Privileged Identity Management (PIM)

Azure Lighthouse also has the possibility to integrate with Azure AD Privileged Identity Management (PIM) (preview). This lets you grant delegated permissions to customer tenants on a just-in-time basis so that users only have those permissions for a set duration at a time. For additional information, you can visit Create eligible authorizations - Azure Lighthouse | Microsoft Learn and samples.

Azure Lighthouse for Cloud Solution Providers

If you're a Cloud Solution Provider partner, a group of users from the MSSP tenant will automatically get Owner permissions for each customer Azure subscription. In this article, you will find specific guidance on how to combine this with Azure Lighthouse: Cloud Solution Provider program considerations - Azure Lighthouse | Microsoft Docs

What can't you do through Azure Lighthouse?

As previously explained, the MSSP preferred way to access a customer's Microsoft Sentinel environment will be utilizing Azure Lighthouse. However, there are operations that are not possible to do with just Azure Lighthouse:

- **Onboard some connectors** that require Security Admin or Global Admin permissions in the customer Azure AD tenant. Most of the Microsoft 1st party connectors require one of these permissions to enabled and these are not available through Azure Lighthouse.
- **Assign incidents to a user in the customer Azure Active Directory**. As incidents are managed in the customer workspaces, they can only be assigned to users in the MSSP's own tenant.

These capabilities can be implemented by using Azure B2B invites, so a user in the MSSP tenant can have a guest user in the customer AAD tenant with the appropriate permissions to perform these actions.

Azure Lighthouse is a very important service to get access to customer Azure resources, but it does not work for other workloads outside of Azure like Office 365 or Microsoft 365 security services. So, how do MSSPs manage customer services based on Office 365 or Microsoft 365? Two solutions exist:

Azure AD B2B

Azure Active Directory (Azure AD) business-to-business (B2B) collaboration is a feature within External Identities that lets you invite guest users to collaborate with your organization.

MSSP users can be "invited" to the customer tenant to perform management activities in that tenant; MSSP users will appear as *guest* users in the customer tenant. These guests can be then granted roles in the customer Azure AD tenant. The main difference with Azure Lighthouse, is that the guest users can be granted any Azure role (even custom ones) or Azure AD roles (remember that Azure Lighthouse can only grant Azure built-in roles). You can see a detailed discussion about the differences between Azure AD and Azure roles here.

The ability to grant Azure AD roles opens new possibilities like managing Office 365 and Microsoft 365 services.

Then, why do we need Azure Lighthouse?
The Azure B2B approach also has some disadvantages:

- **No cross-tenant management or visibility**. As you are invited into a customer tenant, you have to log into the customer tenant in order to see its resources. This blocks your cross-tenant visibility, as you cannot query multiple tenants at the same time. You would have to log in to one customer, manage that customer, log out, and then go to the next customer, which would be very cumbersome.
- **No ability to invite groups**. Azure B2B invitations are done on a user-by-user basis, you cannot invite an entire group. This brings challenges as you need to manage the lifecycle of each account in multiple places (this limitation can be removed by using Azure Entitlement Management, which we review below).

Taking these disadvantages into account, if an MSSP needs to manage both Azure and Office 365/Microsoft 365 workloads, **the best approach is to use a combination of Azure Lighthouse and Azure AD B2B invites** (see section below for instructions on how to automate)
This would look like this:

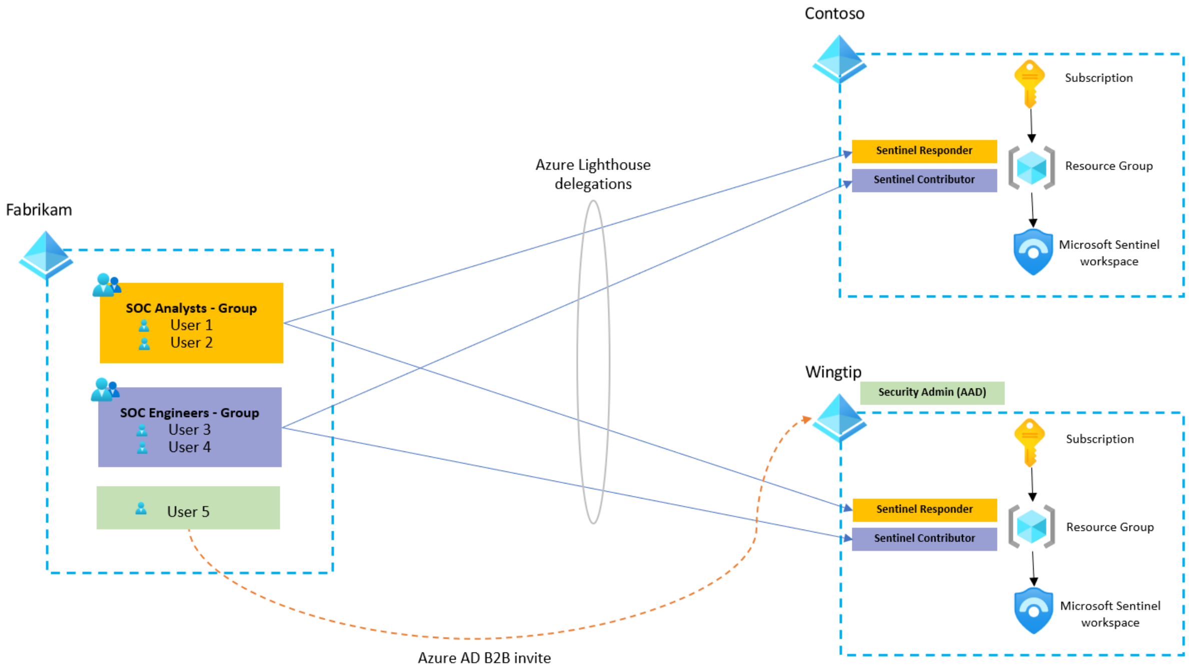

In this picture, there is a new user from Fabrikam (MSSP) that has been invited to Wingtip's Azure AD and is now a *guest* user in that AAD. They have also been granted the Security Admin role. Notice that Security Admin is an Azure AD role, so it can be granted to guest users, but it cannot be granted to users that have delegated access via Azure Lighthouse (remember that Azure Lighthouse can only grant Azure roles). Although not shown in the picture, **the same user can have Azure roles delegated through Azure Lighthouse and also be invited as a guest** and be granted Azure AD roles like Security Admin or Global Admin. It all depends on how you separate your teams internally.

Azure AD Entitlement Management

Azure Active Directory (Azure AD) entitlement management is an identity governance feature that enables organizations to manage identity and access lifecycle at scale, by automating access request workflows, access assignments, reviews, and expiration. **This feature requires an Azure AD P2 license**.

This feature can also be used to manage access from external Azure AD organizations, so it's a **perfect fit for the MSSP access needs when it comes to Microsoft 365 Defender workloads**. This is the high-level architecture for this access model:

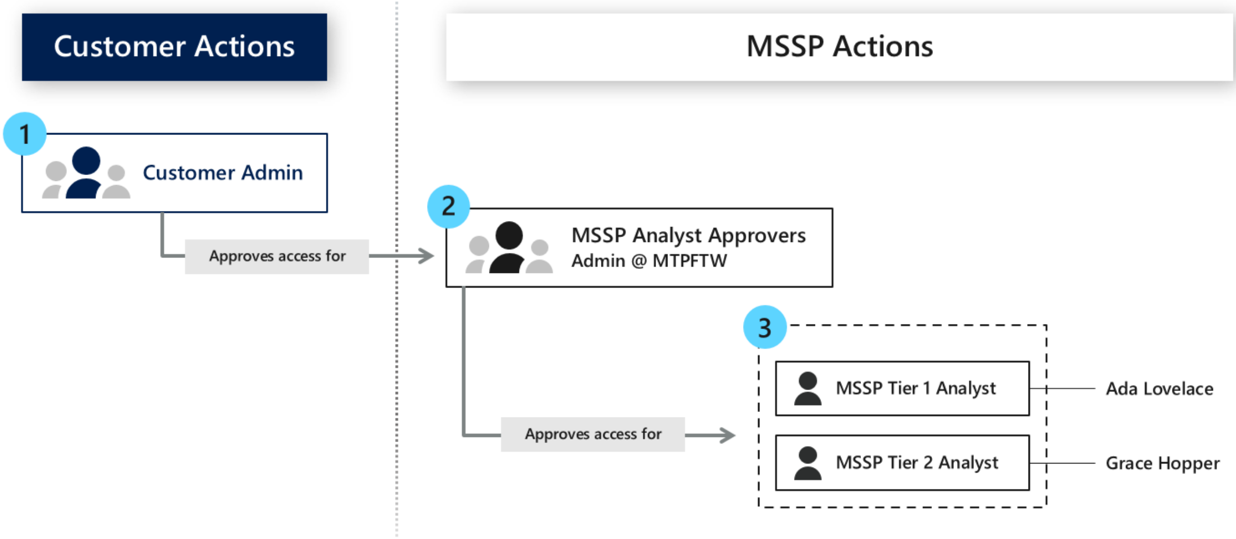

With this setup, MSSP users will be automatically invited into the customer tenant (after the appropriate approvals) to manage customer services. You can also assign which specific roles will be granted to those users; these roles should be specially crafted to manage Microsoft 365 Defender workloads.

For an in depth explanation on how to setup Entitlement Management for an MSSP to access customer environments, read this article. Remember, this is only needed to manage the M365 Defender part of the customer environment, the Microsoft Sentinel part will be managed through Azure Lighthouse as explained above.

Multi-Workspace design principles

A workspace is an element that makes up part of the Azure Monitor service. Azure Monitor allows for the collection and analysis of two types of data:

- Metric data – numerical values, typically performance data
- Log data – stores text values

Log data is viewed using Azure Log Analytics, with the actual data being stored within a Log Analytics Workspace. Once a workspace has been created within a Resource Group, within a Subscription (within an Azure Tenant) then a Microsoft Sentinel instance can be created on top of the chosen workspace.

The default workspace design decision is to keep it simple, with a 'new' single Log Analytics Workspace for Microsoft Sentinel. However, there are some valid reasons why more than one workspace would be required. (General documentation on workspace design can be found here: Design your Microsoft Sentinel workspace architecture | Microsoft Docs)

Following are a set considerations and best practice guidance for MSSPs that will assist in these design decisions:

New or reuse?

Many customers and partners already use Log Analytics to store and analyze data from a wide variety of data sources that have little to do with security. Many collect performance indicators and events from servers to monitor application use and operating system behavior. It might be tempting to simply reuse this workspace for Microsoft Sentinel as well. The Microsoft Sentinel solution applies to the entire workspace. If a workspace is ingesting 500GB daily, but only about 50GB of this is interesting to Microsoft Sentinel, Microsoft Sentinel will still charge for the full 500GB. Data that have less values in term of security can be sent directly to Basic Logs which will be automatically archived after 8 days and reduce the cost. There is no way of carving up a workspace so that Microsoft Sentinel focuses on specific tables only. Therefore, the advice is to create a new workspace for Microsoft Sentinel if a current workspace is already heavily used. In some cases, the non-security log data can be very small and there could be an argument to reuse a current workspace instance. The default approach remains, create a new workspace for Microsoft Sentinel.

Sovereignty + Regulatory Compliance

Workspaces are created within an Azure region. An Azure region is a collection (typically a pair) of datacenters grouped around a geographic area which could be a country, an area within a country or even a continent. The full list of where Log Analytics can be run, and therefore a workspace created, can be found here: Azure Products by Region | Microsoft Azure

Some customers or partners may require the data that Microsoft Sentinel collects to be stored within a specific country or Geo. Up to date information on Geographical availability and data residency is available here.

Larger or complex customers may require additional workspaces as different business groups are subject to different laws and compliance rules. Remember, when examining these areas that Microsoft Sentinel operates at two levels, Log Analytics for storing ingested data and houses the workspace, and there is a smaller data store which stores configuration data about the Microsoft Sentinel instance. This section is only concerned with workspaces. Regional versions of Microsoft Sentinel are covered under the Sovereignty section above.

MSSPs

The assumption is that MSSPs will almost certainly have to manage multiple workspaces. Microsoft cloud contracts effectively mean that customer's data must always be available to them and be exportable to other platforms. Therefore, every customer MUST run their own Microsoft Sentinel within their own Azure Tenant. There are technical reasons for this as well, as many of the more feature-rich data connectors only work within a tenant.

Azure Active Directory tenant boundary

A solid reason for additional workspaces is more technically focused. Many customers will have multiple Azure Tenants which are defined by an instance of Azure Active Directory (AAD). Many of the more feature rich direct connectors that ingest data into Microsoft Sentinel only work within an Azure Tenant boundary. These include the Microsoft 365 Defender (Preview) data connector, currently in preview, that can sync incidents between the two platforms as one example. There are approximately 20 other data connectors that must send logs to a workspace located within the same AAD tenant.

Access control

This can typically be achieved using RBAC as described in the next section. History has taught us that adding more complexity to enforce access control rarely works and is why many larger enterprises have so many on-premises Active Directory domains, many of which are poorly managed!

Split billing

Again, not normally a good reason for splitting out workspaces. Billing can be configured to be quite granular to show costs per workspace, resource group or subscription as needed.

Log Analytic clusters

Log Analytic Dedicated Clusters are collections of workspaces that must be CREATED within the same Azure region. An existing workspace can also be linked to a dedicated cluster and unlink it with no data loss or service interruption (only if the cluster is not enabled for CMK). When using dedicated cluster, a minimum of 500GB of

daily ingestion will be applied even if you only ingest 50GB / day. In addition, only five clusters can be created per subscription per region. though each cluster can hold up to 1000 workspaces.
Dedicated clusters support capabilities like Cross-query optimization, cost optimization and availability zones to MSSPs. However, clusters are not recommended option for MSSPs due to limitations like:

- MSSP cannot bill each customer per their individual usage and ingestion. The pricing tier of the workspace changes to cluster and ingestion is billed as per the cluster's commitment tier. Details here .
- MSSPs cannot allocate the 5MB per user per day data grant to their customers
- Other limits for dedicated clusters can be found here.

Limits

There are, however, some limits to managing multiple workspaces which are being addressed on the roadmap but exist today.

There is a limit of 100 workspaces than can be selected on the cross-incident view within the Microsoft Sentinel UI. Note Azure Lighthouse is needed for an MSSP to view multiple workspaces at once.

There is a limit of 100 workspaces that can be queried by a single KQL query. Though, in practice, the performance of a query running across 100 workspaces will be severely impacted by where these workspaces exist and any network latency which will slow down the query. Again, clustering will help with this area of performance.

A factor to consider when constructing queries across multiple workspaces is that common queries would need to be updated each time an additional workspace was added. To help avoid this admin complexity, the query that contains the workspace names to reference, can be saved as a KQL function and then just use an alias to reference that function (more details here).

An overview of the current limits is in the table below:

Type	Workspace Limit	Description
Cross workspace Incident view	100	Up to 100 workspace instances (e.g., customers) can be displayed at one time.
Cross workspace query	100	A KQL query can contain a workspace reference, well 100 of them though other performance bottlenecks (e.g., network latency) may restrict this max.
Cross workspace workbooks	100	Applies same limit as above
Cross workspace Analytic Rules	20	You can include up to 20 workspaces in a single analytics rule. See Analytics Rule section for more details.
Max analytic rules per workspace	512	We are considering increasing this limit for dedicated clusters in the future

Other Microsoft Sentinel limits are documented on this link, API limits here and LA limits here.

Moving a workspace

It is possible to move workspaces between different tenants within the same Azure region. And it's also possible to move workspaces between subscriptions within the same tenant and the connected ingestion agents will remain connected.

However, not all workspaces are created equal. Microsoft Sentinel workspaces are special as lots of different analytics including ML algorithms run across the data stored in their workspaces. Thus, **today a Microsoft Sentinel workspace CANNOT be moved**. Whilst technically Azure Log Analytics provides an interface allowing the move, **it is not a supported operation for Microsoft Sentinel**. Should a move be necessary, a new instance of Microsoft Sentinel must be created.

Role Based Access Control (RBAC)

This page in the Microsoft Sentinel official documentation contains detailed info about the different Azure built-in roles specific to Microsoft Sentinel. The same page, in the other roles and permissions paragraph, talks about other roles that are not part of Microsoft Sentinel (**SecurityInsights** provider), but that might be needed to operate your Microsoft Sentinel environment. Read the full article referenced above carefully to understand the roles that might be needed as part of your deployment.

Combining the Microsoft Sentinel roles and Azure Lighthouse, a MSSP should plan what permissions will be needed by the different MSSP teams operating in their customer's Microsoft Sentinel environment. Here is a sample table of permissions that can be leveraged as a (simplified) starting point:

Group	Role	Scope	Notes
Security Analysts	Microsoft Sentinel Responder	Microsoft Sentinel's Resource Group	View data, incidents, workbooks, and other Microsoft Sentinel resources. Manage incidents (assign, dismiss, etc.)
	Microsoft Sentinel Playbook Operator	Microsoft Sentinel's Resource Group (or the Resource Group where Playbooks are stored)	List, view and run playbooks. To attach playbooks to analytics rules, Microsoft Sentinel Contributor role is needed
Security Engineers	Microsoft Sentinel Contributor	Microsoft Sentinel's Resource Group	View data, incidents, workbooks, and other Microsoft Sentinel resources. Manage incidents (assign, dismiss, etc.). Create and edit workbooks, analytics rules, and other Microsoft Sentinel resources.
	Logic Apps Contributor	Microsoft Sentinel's Resource Group (or the Resource Group where Playbooks are stored)	Run and modify playbooks. Attach playbooks to analytics rules and automation rules.
	Monitoring Contributor	Subscription and/or Resource group and/or An existing data collection rule	Create or edit data collection rules
	Log Analytics Contributor	Microsoft Sentinel's Resource Group	Use the new Search feature *
	Virtual Machine Contributor Azure Connected Machine Resource Administrator	Virtual machines, virtual machine scale sets Arc-enabled servers	Deploy DCR associations (i.e. to assign rules to the machine)
	Template Spec Contributor	Microsoft Sentinel's Resource Group	Deploy v2.0 solutions from Content hub.
Service Principal	Microsoft Sentinel Contributor	Microsoft Sentinel's Resource Group	Automated configuration management tasks

* There is a cost associated with this feature and write permission on LA is required.

This approach can be further enhanced to account for MSSPs with additional complexity and/or bigger scale. For example, there can be separate groups of Security Analysts by customer, vertical (FSI, retail, health) or technical expertise (identity, network, etc.). Let's see an example where the accounts are separated by expertise:

Group	Role	Scope	Notes
Security Analysts	Microsoft Sentinel Responder	Microsoft Sentinel's Resource Group	View data, incidents, workbooks, and other Microsoft Sentinel resources. Manage incidents (assign, dismiss, etc)
	Microsoft Sentinel Playbook Operator	Microsoft Sentinel's Resource Group (or the Resource Group where Playbooks are stored)	List, view and run playbooks. To attach playbooks to analytics rules, Microsoft Sentinel Contributor role is needed
	Log Analytics Contributor	Microsoft Sentinel's Resource Group	Use the new Search Feature
Identity Engineers	Managed Identity Contributor	Subscription	Create and managed user assigned identities
Networking Engineers	Network Contributor	Subscription	Create and manage networks
Security Engineers	Security Admin	Subscription	View and manage Defender for Cloud (policies, rules, alerts) and view log analytics data
	Microsoft Sentinel Contributor	Microsoft Sentinel's Resource Group	View data, incidents, workbooks, and other Microsoft Sentinel resources. Manage incidents (assign, dismiss, etc). Create and edit workbooks, analytics rules, and other Microsoft Sentinel resources.
	Logic Apps Contributor	Microsoft Sentinel's Resource Group (or the Resource Group where Playbooks are stored)	Run and modify playbooks. Attach playbooks to analytics rules and automation rules.
	Log Analytics Contributor	Microsoft Sentinel's Resource Group	Use the new Search Feature
	Monitoring Contributor	Subscription and/or Resource group and/or An existing data collection rule	Create or edit data collection rules
	Virtual Machine Contributor	Virtual machines, virtual machine scale sets	Deploy DCR associations (i.e. to assign rules to the machine)

Group	Role	Scope	Notes
	Azure Connected Machine Resource Administrator	Arc-enabled servers	
	Template Spec Contributor	Microsoft Sentinel's Resource Group	Deploy v2.0 solutions from Content hub.

In addition to Azure roles, and as explained in the Azure Lighthouse section above, there might be a need to add:

- Azure Active Directory roles that might be required by security operations teams to manage Microsoft 365 workloads.
- Microsoft Sentinel Automation Contributor allows Sentinel to run playbooks. It isn't meant for user accounts.

Data RBAC in Hybrid model

There are two modes in which data can be accessed within a Microsoft Sentinel workspace: workspace context or resource context. The following table summarizes both:

	Workspace-context	Resource-context
Who is each model intended for?	Central administration. Administrators who need to configure data collection and users who need access to a wide variety of resources.	Application teams. Administrators of Azure resources being monitored.
What does a user require to view logs?	Permissions to the workspace.	Read access to the resource. Permissions can be inherited or directly assigned to the resource
What is the scope of permissions?	Workspace. Users with access to the workspace can query all logs in the workspace from tables that they have permissions to	Azure resource. User can query logs for specific resources, resource groups, or subscription they have access to from any workspace but can't query logs for other resources.
How can user access logs?	- Start Logs from Azure Monitor menu. - Start Logs from Log Analytics workspaces. - From Azure Monitor Workbooks.	- Start Logs from the menu for the Azure resource. - Start Logs from Azure Monitor menu. - Start Logs from Log Analytics workspaces. - From Azure Monitor Workbooks.

The recommendation is to use workspace context for teams that need the full Microsoft Sentinel experience and resource context for application teams.

For a full discussion on this topic, visit this article.

Azure Managed Applications

There might be some situations where the MSSP needs to keep the customer's Microsoft Sentinel environment locked down, so nobody from the customer organization can make changes to it. This would avoid situations where the customer accidentally deletes/disables analytics rules, removes playbooks or just close incidents. How can this be achieved?

In these situations, the MSSP can make use of Azure Managed Applications, which allow them to define infrastructure that will be deployed in a specific subscription but with the peculiarity that the resources can only be managed by the publisher, in other words: the MSSP. Not even a user with Owner role for the whole Management Group or Subscription will be able to modify the environment, they will only have **limited** access that is defined.

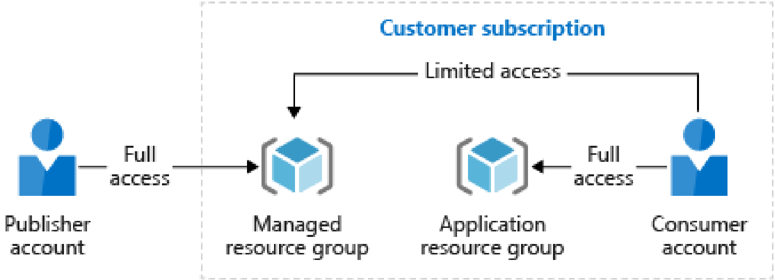

Refer to the documentation for additional details on how to use this deployment method. The good news is that the MSSP can reuse existing ARM template samples to define the Azure Managed Application, for example, Microsoft Sentinel All-in-One.

Sizing & Pricing / Cost

Cost components

Microsoft Sentinel is billed based on the volume of data ingested for analysis in Microsoft Sentinel and stored in the Azure Monitor Log Analytics workspace. The analytics enabled by Microsoft Sentinel do not include the related data ingestion charges for Log Analytics, and both costs must be considered within the estimation.

Once Microsoft Sentinel is enabled on your Azure Monitor Log Analytics workspace, every GB of data ingested into the workspace can be retained at no charge for the first 90 days. Retention beyond 90 days will be charged per the standard Azure Monitor Log Analytics retention prices.

Additional costs may be incurred based on usage of other services within Microsoft Sentinel. Utilizing Logic Apps for automation, or Machine Learning for analysis are two common examples.

Pricing details for Microsoft Sentinel

Pricing details for Azure Monitor Log Analytics

Microsoft Sentinel and Log Analytics offer ingestion & 90 day retention of some data at no cost. Those include:

- Azure Activity Logs
- Office 365 Audit Logs (all SharePoint, Exchange admin activity, Teams)
- Alerts from Microsoft Defender for Cloud, Microsoft 365 Defender, Microsoft Defender for Office 365, Microsoft Defender for Identity, Microsoft Defender for Endpoint and Microsoft Defender for Cloud Apps

Sizing and cost estimations

The key requirement in cost estimation is to identify the daily ingestion rate, usually in GB/day. In most cloud & hybrid environments, networking devices (firewalls, proxies, etc.), Windows servers, and Linux servers produce the most ingested data. An exhaustive inventory of data sources must be performed. The Microsoft Sentinel cost calculator includes tables useful to estimate footprints of data sources. These estimates are a starting point, but log verbosity settings and workloads will produce variances. Regular monitoring is recommended. To understand ingestion patterns real-time, refer to the Log Data Usage and Costs help article.

Using this data, tailor the customer's capacity reservation, which can save up to 60% compared to pay-as-you-go pricing. Capacity reservations provide the ability reserve a fixed amount of daily data ingestion capacity for Azure Monitor and Microsoft Sentinel for a fixed, predictable daily fee. The reservation capacity can be upgraded at any time. However, the minimum commitment period before it can opted out of or reduced is 31 days.

- Adding more capacity to the reservation – An upgrade can be requested at any time. The new capacity reservation will be effective at the start of the next UTC day.
- Reducing the selected capacity reservation – The capacity reservation plan can be reduced or opted out entirely from after the first 31 days. This 31-day clock resets every time any change (increase or decrease) is made to the selected capacity reservation. The new capacity reservation or business model choice will be effective at the start of the next UTC day.

For more details about capacity reservations a corresponding costs, visit the Microsoft Sentinel official pricing page.

For more detailed information on how to plan and understand costs, visit: Plan costs, understand Microsoft Sentinel pricing and billing | Microsoft Docs.

Customers that have Microsoft 365 E5, A5, F5 or G5 licenses are also entitled to special benefits with Microsoft Sentinel: Microsoft 365 E5 benefit offer with Azure Sentinel | Microsoft Azure

In June 2022, we also recorded a webinar to optimize cost in Sentinel using some of the new features.

Long term storage options summary

Customers normally need to keep data accessible for longer than three months. In some cases, this is just due to regulatory or audit requirements, but in some other cases they need to be able to run security investigations on data that is older than 3 months. Whether to keep that data in Microsoft Sentinel or not depends on a number of factors, just to name a few:

- How often does the customer need to access data that is older than the retention period set in Microsoft Sentinel?
- What kind of access is needed? Is it just to show it to an auditor or there a need to perform advanced hunting or investigations?
- How fast does the system need to respond? Is it ok if the system takes a few hours to return the data, or is the data needed in a few seconds?

When restoring data, the following limitations apply. See our documentation for the latest updates:

- Up to 4 restores per workspace per week
- One active restore at a time per table
- Minimum of 2 Days of log data to restore
- The upper limit of 60TB per single restore

Here is a table that summarizes the long-term retention options for Microsoft Sentinel data :

Storage Options	Workspace Retention	Archive	Azure Data Explorer	Azure Blob Storage
Performance	High	Medium (1)	High to Low (2)	Medium to Low
Maximum Retention	2 years	7 years	Unlimited	99 years
Cloud Model / Useability	SaaS / Great	SaaS/Great	PaaS / Good	IaaS / Fair
Cost	High	Low	Medium to Low (2)	Lower
Purpose	SecOps	Long term retention	Extended threat hunting, compliance, trend analysis, storage of non-security data or audit	Archive, Compliance, Auditing

(1) When using Search feature

(2) ADX performance and cost varies depending on cluster size, SKU, and other factors

Let’s now review each of them in detail.

Workspace Retention

This option offers the best performance and provides all the advanced features available in Microsoft Sentinel, like ML based detection rules, visual entity-based investigation, incident management, UEBA, advanced threat hunting, etc. This is the best place for your security data, as you will benefit from all the security-focused functionality mentioned above.

- **Performance**: Microsoft Sentinel/ Azure Log Analytics offers high performance consistently, with low latency and automated data management. There is no need to worry about scalability and performance, as this is handled in the backend.

- **Usability**: Microsoft Sentinel is a full SIEM+SOAR solution and, as mentioned above, offers many features that make it essential in a SOC. **These features are unique to Microsoft Sentinel** and are not part of any of the other long-term storage options.

- **Cost**: The retention costs in Microsoft Sentinel are high compared to the other two storage options. Retention is charged per GB/month after the default 90-day retention has expired. Retention is applied to all sources, including the ones that are ingested for free.

Archive

Archive lets you keep data in Log Analytics Workspace for 7 years in a low-cost archived state. Archive is the recommended long term storage solution for Sentinel. Configuring Archiving is very simple via API or the Microsoft Sentinel UI as described here. Each workspace has a default retention policy that's applied to all tables. You can set different retention policies on individual tables. Refer to this article for more information.

- **Performance:** Once logs are archived, you can use Search or Restore to access archived logs. Archived logs are accessible using search jobs from the Search blade. This lets you maintain access to the data, even when you don't need immediate access to the data. You can also use restore operation to query data in Archived Logs.
- **Usability:** The biggest advantage of Archive is ease of querying archived logs. This is done from Sentinel UI in the Search blade. However, the data isn't immediately available for queries. You need to perform a search to retrieve the data, which might take some time, depending on the amount of data being scanned and returned.
- **Cost:** Cost of data Archived is calculated per GB of data ingested per month. Search on archived logs incur a cost for the data scanned as part of the search, plus the cost of ingesting the search results (per the normal Log Data Ingestion prices). Restore for Archived data incurs a pro-rated cost per day and per GB based on the time the data is kept restored and the amount of data restored. Refer to this link, under the Log Data Archive and Restore section, for exact costs.

Search

The new search feature available in Sentinel queries all storage tiers including analytic tier, basic tier, and archive tier data. The new search feature is best suited when looking for a specific IOC in a specific table. A search job does not support full KQL operators, but the results will be saved in a new table in the same LA workspace, these tables end with the suffix _SRCH. Once the data is available in this new _SRCH table, it can be queried like any other table.

All the _SRCH tables use the retention value set for the workspace, but you can modify this value after the table is created. For MSSPs, it is good practice that only specific SOC engineers be allowed write permissions to perform Search operations as it re-ingests data and normal ingestion costs apply.

Below is a comparison between the Search feature and regular Azure Log Analytics query.

Search	Regular Azure Log Analytics Query
Timeout of 24 hours	Timeout of 10 minutes
Result capability: 1 million rows	Result capability: 30,000 rows
Can run concurrent searches without performance hit	Performance impacted when running concurrently
Can view partial results while job is running	Results can only be viewed once job is done

Restore

When you send data into archive you can use restore option which temporarily rehydrates/surface data into your workspace in the form of a result table. A Restore operation can be performed on both analytic tables and archived data. In Restore, specify the table you want restored and the time period. The Restore operation creates the restore table and allocates additional compute resources for querying the restored data using high-performance queries that support full KQL. Restore is great for investigations when looking for events related to multiple entities that occurred beyond the regular retention period.
The charge for maintaining restored logs is calculated based on the volume of data you restore, in GB, and the number or days for which you restore the data. Charges are prorated and subject to the minimum restore duration and data volume. There is no charge for querying against restored logs.

Azure Data Explorer (ADX)

This option allows you store great amounts of data at a cheaper cost than Log Analytics/ Microsoft Sentinel, while still retaining some advantages. On one hand, it can use the same KQL queries in ADX as in Log Analytics/Microsoft Sentinel, so the same hunting/exploration queries can be used in both places. On the other hand, ADX offers new options to store the data and make it more performant and cost efficient. These options require user attention that is not needed in Microsoft Sentinel. As an example, in ADX it can be defined what data is available in hot cache and move specific tables to it on demand, so queries run faster. In the context of long-term storage, ADX is a great option if the team needs to run light investigations on the data stored in ADX, but without all the security intelligence built on top of Microsoft Sentinel. This option, and its different architecture alternatives, are explained in detail in this blog post. It is also possible to use ADX to run cross-resources queries as explained here.

- **Performance**: several factors contribute to query performance of an ADX cluster: number of nodes in the cluster, cluster VM SKU, data partitioning. As you add more nodes to the cluster, the performance will increase, but also the price. You need to carefully size the cluster to find the sweet spot between performance and cost that works for your organization. This sweet spot will vary from one organization to another, highly dependent on how often the data is accessed and how quickly a response is expected.

- **Usability**: the usability of ADX in the context of security operations is good. You won't benefit from the many out of the box security features in Microsoft Sentinel (UEBA, visual investigation, incident management, etc.), but you can still explore the data using the same KQL queries you use in Microsoft Sentinel, and you can use visualization services like Azure Workbooks or Grafana on top of ADX data. Moreover, you can have an Azure Workbook that visualizes data spread across both Microsoft Sentinel and ADX. **You can also query ADX data right from within Microsoft Sentinel as explained here**, which also allows correlating data from both datastores. You can also do the opposite, query Microsoft Sentinel data from ADX (details here).

- **Cost**: As explained earlier, ADX can be tuned to offer the desired performance. It also offers autoscaling capabilities to adapt to workload on demand. On top of this, ADX can benefit from Reserved Instance pricing. You can run your own cost calculations in the ADX sizing tool. As the data needs to be moved from Microsoft Sentinel/Azure Log Analytics to ADX, there are several architectural options for this approach, and depending on which one you use, there might be additional costs to consider. For example, the Log Analytics data export feature has an associated cost that needs to be accounted for and more information can be found here. An additional cost component that you might need to consider is EventHub (used for queue management). EventHub is sized according to ingestion rate, so the faster data is sent into Azure Data Explorer, the higher the cost in this component.

Azure Blob Storage

This option offers a very low-cost alternative that might be suitable for purely audit/compliance purposes. The data is still secure and accessible, but querying the data requires a bigger effort as the KQL queries need to be modified to include individual SAS URLs pointing to multiple blobs. This option is ideal for scenarios where we need to keep the data to comply with regulation standards and frequent data access is not expected. There are two alternatives to send log data to Azure Blob Storage: using a Logic App or using the Data Export feature in Log Analytics. In order to decide between these two options, take into account the Logic App connector limits, if you expect to go beyond these limits, you should choose the Data Export option.

- **Performance**: Azure Blob Storage offers two performance tiers: Premium or Standard. Although both tiers are an option for long term storage, Standard is generally chosen due to greater cost savings.

- **Usability**: usability is the biggest concern when choosing Azure Blob as your long-term storage option. When we use Azure Blob, the data is saved in separate containers (folders) for each data type and individual blobs in a folder for each hour in the day. This will result in a huge number of files after several weeks of data retention. The easiest way to explore this data is using the *externaldata* operator in KQL. You can see more details about the usage of this operator in this blog post.

- **Cost**: Azure Blob Storage is the cheapest storage option of all three, but you still need to account for the cost of the data export feature (if that option is chosen), which is charged by export GBs, or the periodic execution of the Logic App.

ADX and Blob Storage combined

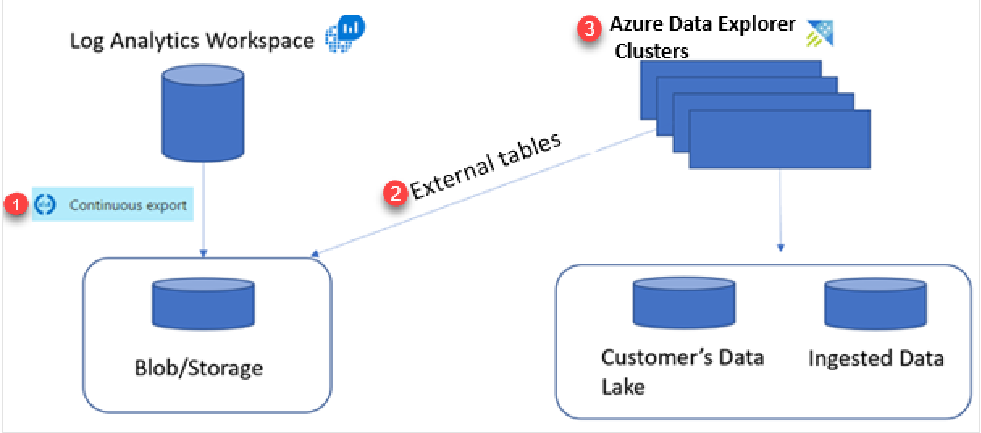

This option allows you to have the best of both worlds. You send the data to Azure Blob's cheap storage, but you can still run KQL queries on the data almost as if you had it local in the ADX cluster. In this article you can see how to implement this approach. As you can see in the aforementioned article, this approach uses Log Analytics data export feature to send the data to Azure Blob Storage, but instead of querying the data from Log Analytics with the *externaldata* KQL operator pointing to each blob, you create an **external table** in ADX that points to the whole container. This way you don't have to reference each individual blob as this is managed transparently for you.

- **Performance**: Azure Blob Storage offers two performance tiers: Premium or Standard. Although both tiers are an option for long term storage, Standard is generally chosen due to greater cost savings.

- **Usability**: usability is the biggest concern when choosing Azure Blob as your long-term storage option. This is mainly because using *externaldata* operator makes it very challenging when you have a big number of blobs to reference. With this approach we **completely eliminate that burden using external tables in ADX**. The external table definition understands the folder structure in the blob storage account and allows us to query the data contained in many different blobs and folders transparently.

- **Cost**: In this architecture, ADX plays a very important role, but the good news is that the size of the cluster doesn't matter because ADX only acts as a proxy. As you can see in the architecture diagram, we use the Log Analytics data export feature, so you will need to factor this into the cost. See below for a detailed cost simulation.

In summary, there are several options available to keep your data accessible while reducing costs. After capturing what are your needs in terms of access to the data and performance you can use this guide to decide the appropriate platform to store your logs.

Data Collection

Data Sources collection overview

After the relatively simple task of creating a workspace and enabling Microsoft Sentinel the real work begins with connecting data sources into Microsoft Sentinel. This is one of the areas where we see MSSPs adding value for their customer base with their skills and experience in connecting the correct log flows into Microsoft Sentinel without simply enabling all of them and the associated ingestion costs!

It is not an exaggeration to state that any security log can be ingested into Microsoft Sentinel by using one of the connector types below. Even a CSV file can be ingested if required, therefore supporting air gapped networks. To support this critical ingestion task there are five core categories of connectors that Microsoft Sentinel uses.

The five types of connectors are:

- **Direct** - native connection into Microsoft Sentinel, often with additional features beyond just log collection.
- **Syslog/CEF** - using the CEF format or even native Syslog, a collector server is needed for this connector type.
- **Agent** - Both Linux and Windows devices can install a monitoring agent to collect logs to send into Microsoft Sentinel. Note, the old MMA agent is slowly being replaced by the new Azure Monitor Agent. See below for more details.
- **Threat Intelligence** - used to ingest commercial IOCs from security partners.
- **Custom** - involve more work to create but ultimately allows for more control and scale as well.

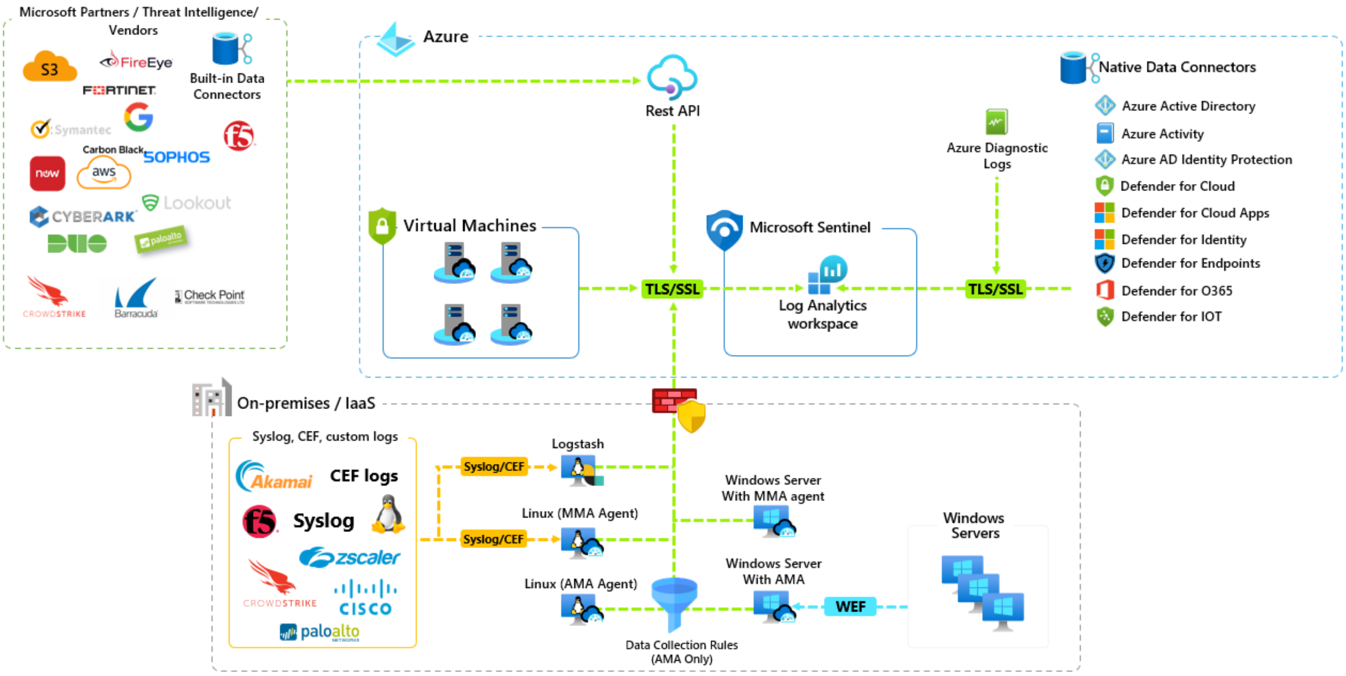

Data collection methods

Connector types

Direct connectors

These connector types are available by default from Microsoft Sentinel, they are essentially built in. These data sources have native connectivity which implies a cloud type of service from Azure, AWS and GCP as well. Some common connectors used when Microsoft Sentinel is first built are:

- **Azure Active Directory** – Audit, Sign-in logs, and additional logs that are in preview
- **Azure Activity** – logs subscription level events
- **Microsoft 365 Defender** — used to collect alerts from Microsoft Defender for Cloud, Microsoft 365 Defender, Microsoft Defender for Office 365, Microsoft Defender for Identity, Microsoft Defender for Endpoint and Microsoft Defender for Cloud Apps. This is a feature rich connector which will also allow for raw data (as well as events or alerts) to be ingested if required and configured.
  This option enables additional threat hunting, correlation, applied threat intelligence, and advanced machine learning algorithms. The data can also be sent to the security data warehouse (ADX) for long-term data retentions. This could occur at the same time as sending to the SIEM, or it can be sent there after the SIEM has expired the data
- **Office 365** - User and admin activities within Exchange, SharePoint and Teams. This connector uses the Office 365 Management Activity API but ingests only alert level information NOT the raw data from this API.
- **Microsoft Defender for Cloud** - ingest alerts from Microsoft Defender for cloud's various products

There is another category of direct connectors which reach out and connect directly to devices that can be located within the cloud but more typically on customer sites. These connectors use the source device API (when available) and are normally built by the device manufacturers. Some common direct API connectors are shown below:

- Barracuda WAF and Firewall
- Citrix Analytics
- F5 Big IP
- Forcepoint DLP
- Proofpoint
- Qualys VM (Azure Function based)
- Salesforce Service Cloud (Azure Function based)

The screenshot below shows the most commonly connected first security data source, Azure Active Directory sign-in and audit logs. Note the permissions needed to deploy the connector. Note, the license requirement is being retired imminently. You will also notice additional Azure AD sources currently in preview.

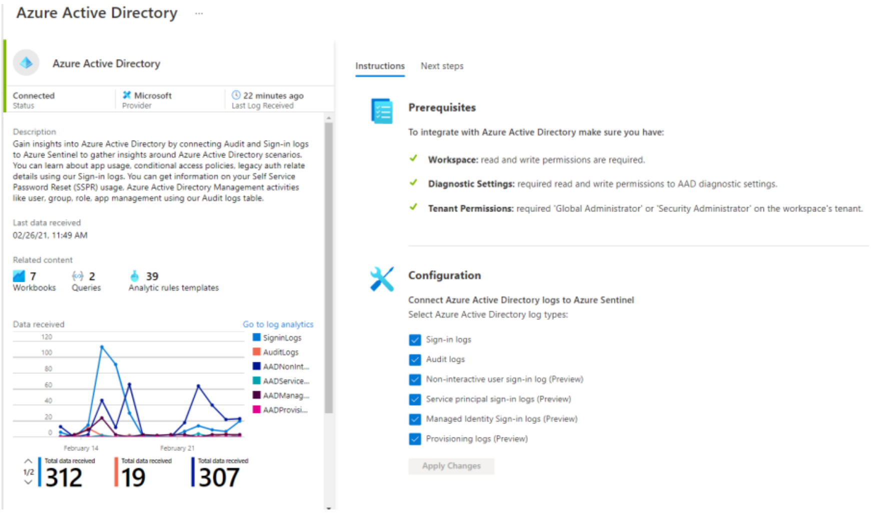

Azure Active Directory ···

Azure Active Directory

Connected
Status

Microsoft
Provider

22 minutes ago
Last Log Received

Description
Gain insights into Azure Active Directory by connecting Audit and Sign-in logs to Azure 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
02/26/21, 11:49 AM

Related content
7 Workbooks
2 Queries
39 Analytic rules templates

Prerequisites

To integrate with Azure Active Directory make sure you have:

✓ **Workspace:** read and write permissions are required.

✓ **Diagnostic Settings:** required read and write permissions to AAD diagnostic settings.

✓ **Tenant Permissions:** required 'Global Administrator' or 'Security Administrator' on the workspace's tenant.

Configuration

Connect Azure Active Directory logs to Azure Sentinel
Select Azure Active Directory log types:

☑ Sign-in logs
☑ Audit logs
☑ Non-interactive user sign-in log (Preview)
☑ Service principal sign-in logs (Preview)
☑ Managed Identity Sign-in logs (Preview)
☑ Provisioning logs (Preview)

Syslog connectors (Forwarders)

Many traditional security products (and many others) will support sending security logs using syslog, a protocol used for sending security event information. Microsoft Sentinel does have the ability to ingest raw syslog messages (via a collector) but the preferred approach is to use CEF (Common Event Format) formatted events transported over the syslog protocol.

CEF data is normalized before being stored in Log Analytics and so is easier to query. Syslog data is not normalized and therefore more complicated to query. Though Microsoft Sentinel, unlike many other SIEM products, can still store and query syslog events.

Either way both types are supported via a collector. This collector uses a Linux machine as a log forwarder, to forward events into Microsoft Sentinel. There is no direct connection from a syslog (or CEF) source to Microsoft Sentinel. For this to happen the device would need a well-documented API that allows security events to be read. See previous section on Direct connectors.

This collector can be installed on-premises or indeed in the cloud as a VM depending on the infrastructure available to the customer and MSSP. There are various guides on how to scale syslog collection, and at truly high ingestion volumes it may be more appropriate to use Azure Functions to allow ingestion to scale higher. Typically, daily ingestion amounts of over 1TB could need this investment.

There are potentially thousands of CEF and Syslog connected devices, many are included within Microsoft Sentinel and include a link to a manufacturers page with help on configuring CEF / Syslog. In this article, there is a list of different data sources and the preferred collection method.

Agent based connector

There are currently two agent types to choose from. There is the 'classic' and its soon to be seen replacement.

Azure Log Analytics agent (classic)
Also historically called the OMS agent or MMA (Microsoft Monitoring Agent) which gives an idea of its long history. The agent runs on Windows and Linux machines.

This agent is typically used by Microsoft Sentinel to collect OS Events. It also collects logs from Microsoft DNS Servers, Windows Firewalls and Windows Security Events.

There are many other types of data that can be collected by this agent depending on the role of the machine the agent finds itself on. This even includes importing files. More details can be found here.

NOTE - The Log Analytics agent will be retired on 31 August, 2024

Azure Monitor Agent (new)
The incoming replacement for the Log Analytics agent is the Azure Monitor Agent..

When used with Sentinel, the new agent supports the following:

- Windows Security Events
- Windows Forwarding Event (WEF) – Public Preview
- Windows DNS logs – Public Preview
- Linux Syslog CEF – Public Preview

This updated agent supports:

- An ability to filter events before ingesting into Microsoft Sentinel
- Support for multi homed solutions on Linux platforms (already supported on Windows)
- Support for Windows event collection and filtering

Features parity between Azure Log Analytics agent (MMA) and Azure Monitor Agent (AMA):

- Supported services and features of the AMA are available here
- Comparison with the legacy agent is available here

Migration guide from MMA to AMA is available in our official documentation using this link.

How to deploy the agent:
In the Azure portal you can create a new Data Collection Rule in Azure Monitor and specify the virtual machines you wish to include. This will enable the system-assigned managed identity, install the Azure Monitor agent extension and create and deploy the data collection rule associations. DCR must be created in the same region where the LA workspace resides.

More information on limits, configuration and deployment of Data Collection Rules can be found here.

Threat Intelligence connector

NOTE - There is a separate section focused on Threat Intelligence later in this document.

There are two methods of connecting commercial Threat Intelligence Providers (TIPs), or in fact any source of threat intelligence, into Microsoft Sentinel. Threat Intelligence can be any source of threat information though is often represented by a file hash, a URL, or an IP address.

Microsoft Graph API
A rich interface that allows TIPs to feed a wealth of information into Microsoft Sentinel (and also Defender for Endpoint if needed!). More details can be found here. A number of commercial providers provide a suitable Graph API connector including:

- MISP
- Palo Alto MineMeld
- ThreatConnect
- ThreatQuotient

TAXII connector
Microsoft Sentinel comes complete with a STIX/TAXII v2 (only, no v1 support) connector which enables a built-in TAXII client in Microsoft Sentinel to import threat intelligence from TAXII 2.x servers. No need to create an Azure AD enterprise application just select the feeds required and any needed account information and the threat data will start to flow into Microsoft Sentinel.

As a final note it is possible to add custom (via the dashboard) indicators directly into Microsoft Sentinel as well.

Custom connectors

There are several methods available to connect data into Microsoft Sentinel if the above four categories cannot connect the source required.

Codeless Connectors Platform (CCP)
Used to create a connector for any data source that exposes a public REST API endpoint. This is a full SaaS solution that does not require any other service such as Azure Functions. This means you don't need to scale it or pay for the additional services.

These connectors can be automatically deployed with API or with ARM template. More information about the Codeless Connectors Platform can be found here.

Logs ingestion API (public preview)
Log Analytics also has a log ingestion REST API which uses HTTP to stream data into Log Analytics. This means any popular programming language can be used to send data into Log Analytics and therefore Microsoft Sentinel. Data sent using this REST API requires that a Data Collection rule (DCR) be used which can transform the data before it gets written to the Log Analytics workspace and a Data Collection Endpoint (DCE) where you send the data. In addition, data can be written to a custom table as well as the CommonSecurityLog, SecurityEvents, Syslog, and the WindowsEvents tables. This new API also allows adding custom fields to native tables.

More information on this API can be found here.

Log Analytics HTTP Data Collector REST API (legacy)

This REST API is another way to send data into a Log Analytics table. However, when using this REST API, there is no ability to perform any data transformation and the data can only be written to a custom table.

More information on this API can be found here

PowerShell cmdlet

There is a cmdlet that can be found here which will allow upload of data into a custom table within Log Analytics, using the legacy HTTP Data Collector REST API. This even includes flat files such as a CSV.

Azure Logic Apps

Data collection can be orchestrated! A scheduled task or triggers can be used to connect to data sources and ingest this data into Microsoft Sentinel. Although this can be a relatively simple method for automating the ingestion of data is not recommended for larger data sets due to the underlying costs. Therefore, use for low volume data sources only or for enriching other data ingestions.

Some high-level examples of using Logic Apps to ingest data could be gathering data from:

- Any REST based API
- A SQL server
- A file system

NOTE - Logic Apps can be used with on-premises equipment using gateways such as the On-Premises Data Gateway.

Azure Functions

As the name implies, a built-in Azure feature that can be used for Microsoft Sentinel purposes. It is a way of running serverless code and is generally used as a Microsoft Sentinel data connector when ingesting large amounts of data or data with unique collection requirements.

These types of connectors need to be built and a wide range of programming languages such as PowerShell, Python, .NET C#, and Java are included.

Most of the newer connectors being built for Microsoft Sentinel use this collection method. You can see some examples in our official documentation, like Okta or Proofpoint TAP.

There are other connectors that are built by the community, like the Office 365 Management API connector, which collects not just the alerts that the standard built in connector handles, but additional audit logs. This can be found here.

The Microsoft Sentinel product team and the community often share new examples of Custom Data connectors. These are available on GitHub.

Multi-Cloud Environment

Customers might have a multi-cloud environment with the need to ingest logs from multiple cloud service providers into Microsoft Sentinel. Below are the native data connectors available in Microsoft Sentinel for Amazon and Google:

Amazon

Microsoft Sentinel provides native Amazon Web Services (AWS) connectors to pull AWS service logs into Microsoft Sentinel. There are two versions of AWS connectors available – AWS S3 connector (new) and AWS CloudTrail connector (legacy). It is recommended to use the S3 connector even for the CloudTrail logs.
The new S3 connector works using the following architecture and workflow. Review the architecture overview and configurations for setup information.

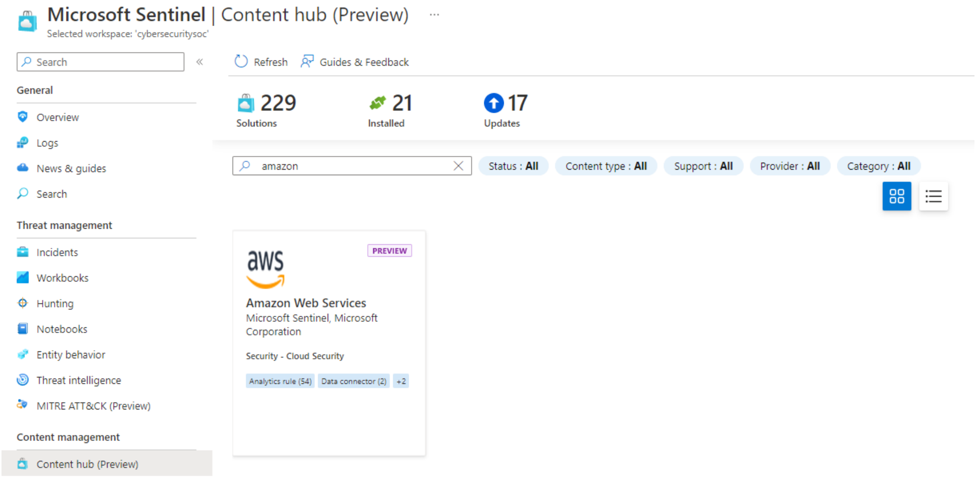

At the time of writing, the S3 connector supports AWS logs from VPC Flow Logs, GuardDuty and CloudTrail. But the connector will be extended to support custom logs from S3 bucket soon.

Google

Microsoft Sentinel has Google Cloud Platform (GCP) and Google ApigeeX solution from Content Hub which include data connectors. Besides that, there is a built-in connector for Google Workspace. Below is a list of data connectors or solutions available for Google.

- Google Cloud Platform DNS
- Google Cloud Platform IAM
- Google Cloud Platform Cloud Monitoring
- Google ApigeeX
- Google Workspace (G-Suite)

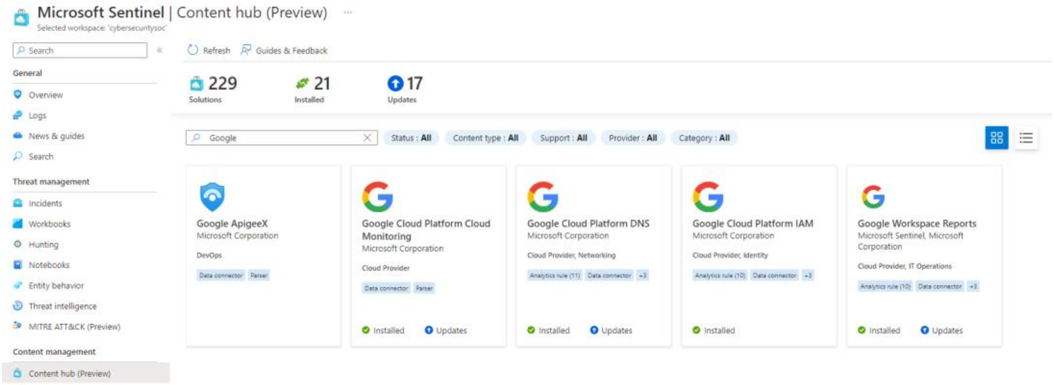

Data Connectors Health

Once you have enabled data connectors in your Microsoft Sentinel workspace, you can monitor the health of the connectors by enabling health monitoring for supported data connectors. This provides insights on your connector health such as authentication, throttling, and other service or data source issues.

Health data from Data Connector Health are stored in SentinelHealth data table. This allows you to run queries to detect health drifts or even configure alerts and automated actions for health issues.

As a Microsoft Security Services Provider, you can leverage Azure Lighthouse and Cross-workspace query to run the health queries, configure alerts and automated response for your customers from multiple tenants.

Besides Data Connector Health or SentinelHealth table, you can also use data collection health monitoring workbook to monitor your data ingestion status. The workbook provides general status of data ingestion in the selected workspace by showing data volume, EPS rates, log received time. It also detects ingestion anomalies and provides agent health information on the installed agents (MMA/AMA).

You can leverage Azure Lighthouse to select the subscription and workspace of your customer tenant in the workbook to gain insights of the data collection health.

Normalization

Advanced Security Information Model / ASIM

Microsoft Sentinel ingests data from many different sources and each source uses different column / field names to designate the same object or they have different fields values. Working with various data types and tables together requires you to understand each of them and write and use unique data sets for analytics rules, workbooks, and hunting queries for each type or schema. As an example, when working with authentication, the user ID has different names depending on where the data come from.

User unique identifier	Description	Example
SID	A Windows user ID.	S-1-5-21-1377283216-344919071-3415362939-500
UID	A Linux user ID.	4578
AADID	An Azure Active Directory user ID.	9267d02c-5f76-40a9-a9eb-b686f3ca47aa
OktaId	An Okta user ID.	00urjk4znu3BcncfY0h7
AWSId	An AWS user ID.	72643944673

To make normalization easier and more effective, a framework called **Advanced Security Information Model** or **ASIM** was created**.**

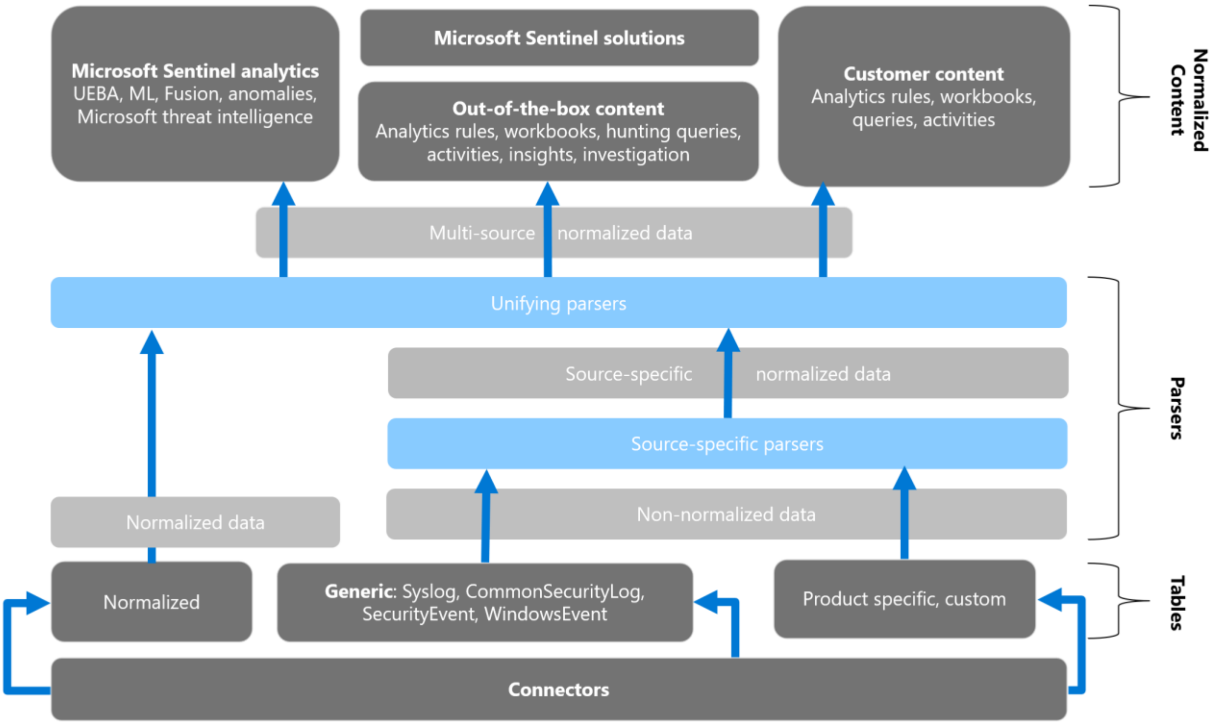
Microsoft Sentinel analytics
UEBA, ML, Fusion, anomalies, Microsoft threat intelligence
Microsoft Sentinel solutions
Out-of-the-box content
Analytics rules, workbooks, hunting queries, activities, insights, investigation
Customer content
Analytics rules, workbooks, queries, activities
Normalized Content
Multi-source normalized data
Unifying parsers
Source-specific normalized data
Source-specific parsers
Non-normalized data
Normalized data
Parsers
Normalized
Generic: Syslog, CommonSecurityLog, SecurityEvent, WindowsEvent
Product specific, custom
Tables
Connectors

This framework provides normalization to Microsoft Sentinel. It includes the following components:

Normalized schemas

Cover standard sets of predictable event types that are easy to work with and build unified capabilities. The schema defines which fields should represent an event, a normalized column naming convention, and a standard format for the field values. Query writers and analysts save time by learning only an intuitive standard schema and creating queries that support multiple built-in and custom sources based on this schema.

More information around Normalization for Parsing can be found here.

Examples of ASIM normalized schemas already available today:

- DNS schema: Parsers available include Cisco Umbrella, Infoblox NIOS, BIND, BlucCat, Microsoft DNS Server, Zscaler ZIA, Azure Firewall, Sysmon for Windows. More specific parsers are available here.

- Authentication schema: Parsers available for Windows Sign-ins, Linux sign-ins, Azure AD Sign-ins, AWS sign-ins, Okta authentication, PostgreSQL sign-ins. More information can be found here.

- Network session schema: Parsers available for Checkpoint Firewall, Cisco Meraki, FortiGate FortiOS, Azure Firewall, Azure Network Security Group, AWS VPC logs, Microsoft 365 Defender for Endpoint, Windows Firewall logs, Sysmon for Linux and more available here

The full list of ASIM normalized schemas is available here.

Parsing
Microsoft Sentinel supports parsing at query time which means security data can be ingested in its original format. This means connectors do not necessarily need to be updated if the data source structure changes. Perhaps due to a firmware update on the source system, an operating system update, new product range etc. Instead, parsing is done at time of query or even during an active investigation.

However, the better formatted the original data source is, the easier it is to query. Hence, CEF is preferred over the raw Syslog data sources.
You can find KQL parsers for major vendors here.

As we have seen earlier, ASIM brings many benefits for the SOC Analysts and it also help creates better detections rules, hunting queries or Workbooks. Beside these benefits, MSSP will also now be able to easily deploy substantial number of parsers using ARM templates as documented here.

Pipeline Transformation

Data ingestion and transformation

Ingestion time transformation allows manipulation and control over the data (or logs) before it's ingested into the Analytics workspace. It gives you more control over the data ingested into your workspace.
Ingestion time transformation uses Data Collection Rules (DCRs) to filter and enrich the workspace tables. You can then use KQL queries to specify the transformation criteria. Ingestion time transformation improves analytics, performance and reduces ingestion costs. The use cases include:

- Filtering out unwanted data
- Enrichment and tagging
- Masking or Obfuscation of sensitive data

The list of tables that support ingestion time transformation is included here. Custom logs must be ingested using the Logs ingestion API to support ingestion time transformation. The Custom Log API supports ingesting logs to Custom tables or supported Standard tables.

Logstash output plugin

Instead of using the Log Analytics agent, Logstash can be used instead. It gives some additional features over the default agent such as aggregation or enrichment with external sources. There is an output plugin for Logstash that will connect it directly into Log Analytics where the logs are stored within a custom table. The diagram below shows the different elements of a Logstash solution.

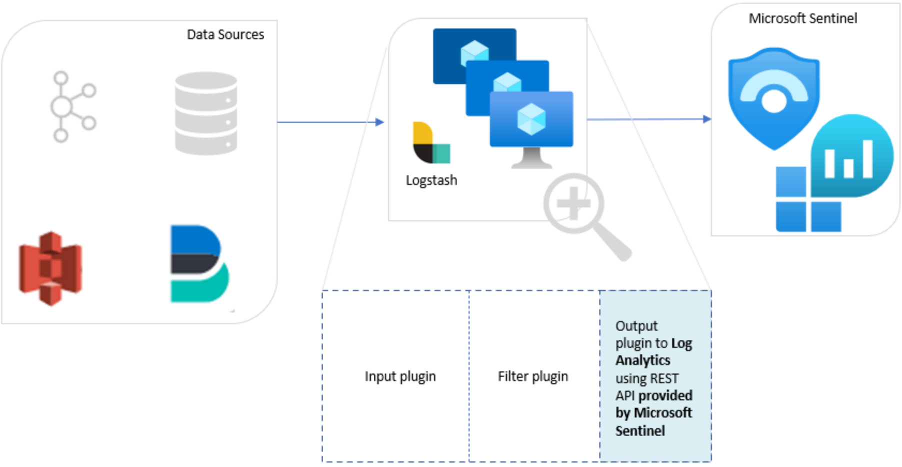
Data Sources
Logstash
Microsoft Sentinel
Input plugin
Filter plugin
Output plugin to Log Analytics using REST API provided by Microsoft Sentinel

The input plugin allows for customization of the collection of data from different sources.
The filter plugin allows for normalization of the data before being ingested into Microsoft Sentinel.
The output plugin is the connector provided by Microsoft Sentinel to connect the data.

Note that the current version of the Logstash plugin uses custom tables, not the built-in tables that first party connectors for Microsoft Sentinel write to. This will be updated in the future, but for now certain Microsoft Sentinel features such as UEBA and Fusion can only reason over the built-in tables, not custom tables.

More information about Logstash and its connection into Microsoft Sentinel can be found here.

NOTE - a newer version of the Logstash output plugin is currently being developed. It will offer the following benefits over the initial release:

- Supports DCR-based custom logs
- Performance improvements & optimizations (data compression)
- Co-existence with custom logs v1
- Better error handling & telemetry

Enrichment

Generally, involves tying together data from different sources to provide a richer stream of data, and to enhance incidents when they occur. This is typified by adding areas such as GeoIP information, data from WhoIs, username and device name enrichment as well.

A common example is enriching suspicious file information from VirusTotal which can be automated using Logic Apps. More details on this approach here.

Basic logs (Preview)

You can configure some log sources to be ingested as tables in Basic Logs Tier to reduce the cost of ingesting high-volume verbose logs you use for debugging, troubleshooting, auditing and are required for ingestion. Tables in Basic logs tier are accessible via limited KQL for 8 days after ingestion for interactive queries, the query API, and the new Search UX. Data from Basic Logs can be used for investigation, IOC search, ad hoc queries, and as part of Logic App playbook automation. Beyond the initial 8 days, Microsoft Sentinel archives Basic Logs and are accessible via the new Search experience. A table can be converted from Basic Logs Tier to Analytics Logs Tier or back and forth. This conversion can only be done once per week. Check here to compare the Basic Tier and Analytics Tier. To see tables that support Basic logs click here. Tables created using Data Collector API are not supported.

Automation/SOAR

Built on the foundation of Azure Logic Apps, Microsoft Sentinel's automation and orchestration solution provides a highly extensible architecture that enables scalable automation as modern technologies and threats emerge. You can use these features to automate your common tasks and simplify security orchestration with playbooks and automation rules (Preview) that integrate with Azure services as well as with your existing tools.

Automation Rules

What is it?
Automation rules allow you to centrally manage all the automation when it comes to incident handling. Automation rules streamline automation use in Microsoft Sentinel and enable you to simplify complex workflows for your incident orchestration processes.

How does it work?
Automation rules are triggered by the creation or update (in preview) of incidents. You can set conditions to govern when actions will run, based on the incident properties, entity details and/or analytics rules that triggered them. You can also set the **order** of actions and the rule’s expiration time.

Automated incident response
Explicitly set incident status or severity, assign an owner, or add a tag when an incident is created or updated (in preview), **without** the need to run a playbook**.** Some of the benefits include:

1. **Running playbooks on incidents**
   You can still run playbooks from your automation rules to integrate with other services or create complex workflows. These automation rules, pass all incidents details to the playbooks, including alerts and entities.
2. **Trigger playbooks for Microsoft Providers**
   Automate the handling of Microsoft security alerts by applying automation rules to incidents.
   created from alerts.
3. **Do more with your playbooks**
   Automation rules allow you to Apply a single playbook to any of your analytics rules once, attach multiple playbooks to a single automation rule, and control the order of playbook execution. This will allow you to create simpler playbooks, easier to maintain and test, and arrange them in combinations as needed.
4. **Apply incident suppression**
   You can use rules to automatically resolve incidents that are known as false positives. For example, when running penetration tests, during scheduled maintenance or upgrades, or testing automation procedures.

Creating and managing automation rules

Automation rules are centrally managed in the new Automation tab under the Automation Rules sub tab (see screenshot below). From there, users can create new automation rules and edit the existing ones. They can also drag & drop automation rules to change the order of execution and enable or disable them.

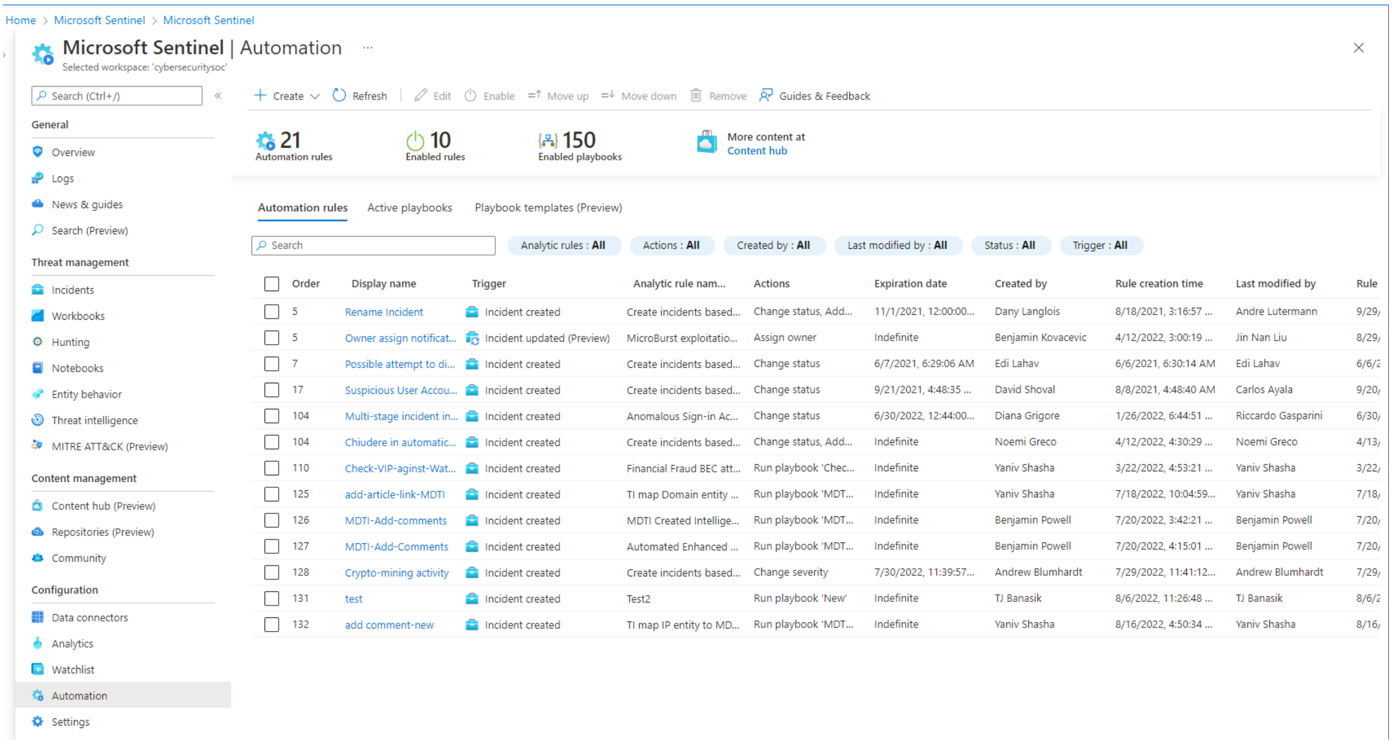

☐	Order	Display name	Trigger	Analytic rule nam...	Actions	Expiration date	Created by	Rule creation time	Last modified by	Rule
☐	5	Rename Incident	Incident created	Create incidents based...	Change status, Add...	11/1/2021, 12:00:00...	Dany Langlois	8/18/2021, 3:16:57 ...	Andre Lutermann	9/29/
☐	5	Owner assign notificat...	Incident updated (Preview)	MicroBurst exploitatio...	Assign owner	Indefinite	Benjamin Kovacevic	4/12/2022, 3:00:19 ...	Jin Nan Liu	8/29/
☐	7	Possible attempt to di...	Incident created	Create incidents based...	Change status	6/7/2021, 6:29:06 AM	Edi Lahav	6/6/2021, 6:30:14 AM	Edi Lahav	6/6/2
☐	17	Suspicious User Accou...	Incident created	Create incidents based...	Change status	9/21/2021, 4:48:35 ...	David Shoval	8/8/2021, 4:48:40 AM	Carlos Ayala	9/20/
☐	104	Multi-stage incident in...	Incident created	Anomalous Sign-in Ac...	Change status	6/30/2022, 12:44:00...	Diana Grigore	1/26/2022, 6:44:51 ...	Riccardo Gasparini	6/30/
☐	104	Chiudere in automatic...	Incident created	Create incidents based...	Change status, Add...	Indefinite	Noemi Greco	4/12/2022, 4:30:29 ...	Noemi Greco	4/13/
☐	110	Check-VIP-aginst-Wat...	Incident created	Financial Fraud BEC att...	Run playbook 'Chec...	Indefinite	Yaniv Shasha	3/22/2022, 4:53:21 ...	Yaniv Shasha	3/22/
☐	125	add-article-link-MDTI	Incident created	TI map Domain entity ...	Run playbook 'MDT...	Indefinite	Yaniv Shasha	7/18/2022, 10:04:59...	Yaniv Shasha	7/18/
☐	126	MDTI-Add-comments	Incident created	MDTI Created Intellige...	Run playbook 'MDT...	Indefinite	Benjamin Powell	7/20/2022, 3:42:21 ...	Benjamin Powell	7/20/
☐	127	MDTI-Add-Comments	Incident created	Automated Enhanced ...	Run playbook 'MDT...	Indefinite	Benjamin Powell	7/20/2022, 4:15:01 ...	Benjamin Powell	7/20/
☐	128	Crypto-mining activity	Incident created	Create incidents based...	Change severity	7/30/2022, 11:39:57...	Andrew Blumhardt	7/29/2022, 11:41:12...	Andrew Blumhardt	7/29/
☐	131	test	Incident created	Test2	Run playbook 'New'	Indefinite	TJ Banasik	8/6/2022, 11:26:48 ...	TJ Banasik	8/6/2
☐	132	add comment-new	Incident created	TI map IP entity to MD...	Run playbook 'MDT...	Indefinite	Yaniv Shasha	8/16/2022, 4:50:34 ...	Yaniv Shasha	8/16/

Automation rules can also be created from within the Incidents view. Here is a screenshot that shows you how to create an automation rule from a given incident:

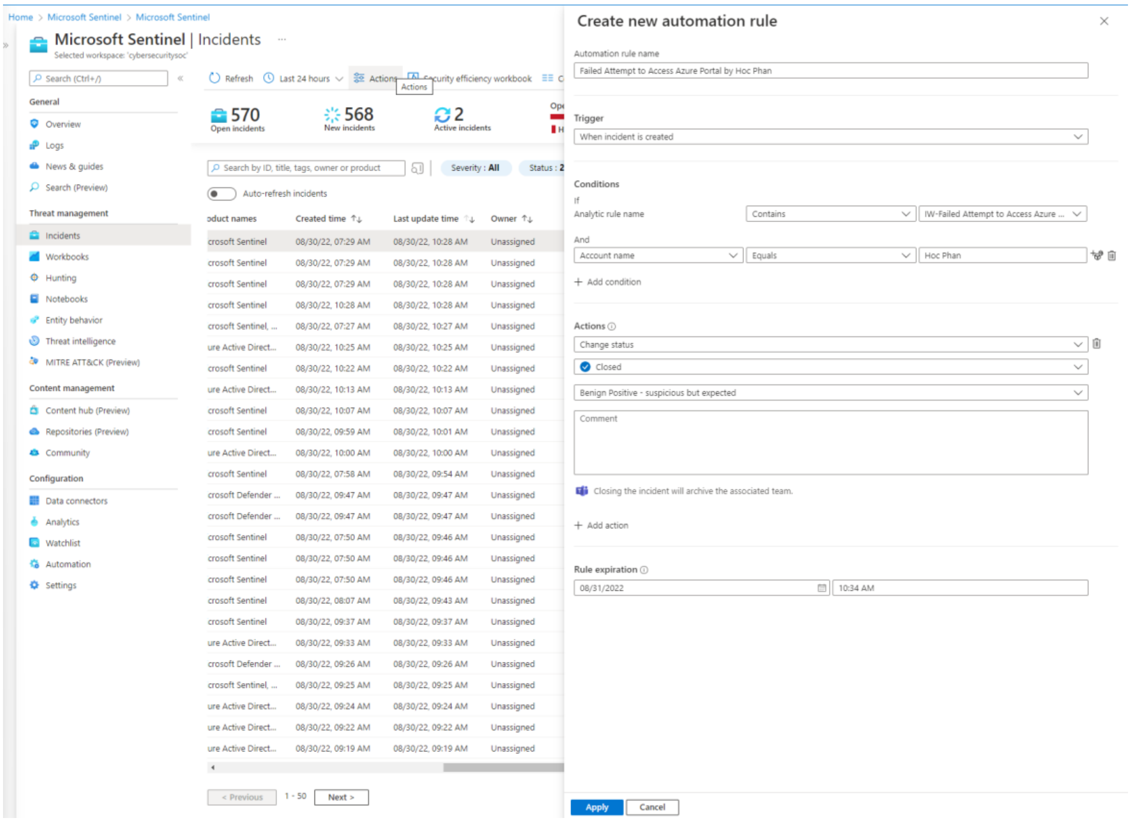

oduct names	Created time ↑↓	Last update time ↑↓	Owner ↑↓
crosoft Sentinel	08/30/22, 07:29 AM	08/30/22, 10:28 AM	Unassigned
crosoft Sentinel	08/30/22, 07:29 AM	08/30/22, 10:28 AM	Unassigned
crosoft Sentinel	08/30/22, 07:29 AM	08/30/22, 10:28 AM	Unassigned
crosoft Sentinel	08/30/22, 10:28 AM	08/30/22, 10:28 AM	Unassigned
crosoft Sentinel, ...	08/30/22, 07:27 AM	08/30/22, 10:27 AM	Unassigned
ure Active Direct...	08/30/22, 10:25 AM	08/30/22, 10:25 AM	Unassigned
crosoft Sentinel	08/30/22, 10:22 AM	08/30/22, 10:22 AM	Unassigned
ure Active Direct...	08/30/22, 10:13 AM	08/30/22, 10:13 AM	Unassigned
crosoft Sentinel	08/30/22, 10:07 AM	08/30/22, 10:07 AM	Unassigned
crosoft Sentinel	08/30/22, 09:59 AM	08/30/22, 10:01 AM	Unassigned
ure Active Direct...	08/30/22, 10:00 AM	08/30/22, 10:00 AM	Unassigned
crosoft Sentinel	08/30/22, 07:58 AM	08/30/22, 09:54 AM	Unassigned
crosoft Defender ...	08/30/22, 09:47 AM	08/30/22, 09:47 AM	Unassigned
crosoft Defender ...	08/30/22, 09:47 AM	08/30/22, 09:47 AM	Unassigned
crosoft Sentinel	08/30/22, 07:50 AM	08/30/22, 09:46 AM	Unassigned
crosoft Sentinel	08/30/22, 07:50 AM	08/30/22, 09:46 AM	Unassigned
crosoft Sentinel	08/30/22, 07:50 AM	08/30/22, 09:46 AM	Unassigned
crosoft Sentinel	08/30/22, 08:07 AM	08/30/22, 09:43 AM	Unassigned
crosoft Sentinel	08/30/22, 09:37 AM	08/30/22, 09:37 AM	Unassigned
ure Active Direct...	08/30/22, 09:33 AM	08/30/22, 09:33 AM	Unassigned
crosoft Defender ...	08/30/22, 09:26 AM	08/30/22, 09:26 AM	Unassigned
crosoft Sentinel, ...	08/30/22, 09:25 AM	08/30/22, 09:25 AM	Unassigned
ure Active Direct...	08/30/22, 09:24 AM	08/30/22, 09:24 AM	Unassigned
ure Active Direct...	08/30/22, 09:22 AM	08/30/22, 09:22 AM	Unassigned
ure Active Direct...	08/30/22, 09:19 AM	08/30/22, 09:19 AM	Unassigned

Create new automation rule

Automation rule name: Failed Attempt to Access Azure Portal by Hoc Phan

Trigger: When incident is created

Conditions

If Analytic rule name | Contains | IW-Failed Attempt to Access Azure ...

And Account name | Equals | Hoc Phan

+ Add condition

Actions

Change status

Closed

Benign Positive - suspicious but expected

Comment

Closing the incident will archive the associated team.

+ Add action

Rule expiration: 08/31/2022 | 10:34 AM

Apply | Cancel

Microsoft Sentinel Playbooks

A security playbook is a collection of procedures that can be run from Microsoft Sentinel in response to an alert. A security playbook can help automate and orchestrate your response and can be run manually or set to run automatically when specific alerts are triggered. Security playbooks in Microsoft Sentinel are based on Azure Logic Apps.

Playbook Templates (preview)

When needing to recreate the same playbook in multiple customers, using the Playbook Templates gallery can reduce the time it will take to do so (also see the CI/CD section below to see about deploying custom playbooks). This provides a listing of pre-built playbook templates that can be deployed easily to multiple customers.

For more information go here

Playbook Health monitoring

If your customers are using playbooks to perform automated tasks, it is a requirement to understand how well those playbooks are working. While there is no functionality within Microsoft Sentinel to perform this activity, the “Playbook health monitoring” workbook will allow you to view how selected playbooks/logic apps have performed for a specified time period.

In the image below, most of the logic apps have run successfully, but the “Get-SecurityScoreData” for the “CyberSecSOC” subscription should be looked at to determine why it is failing so much

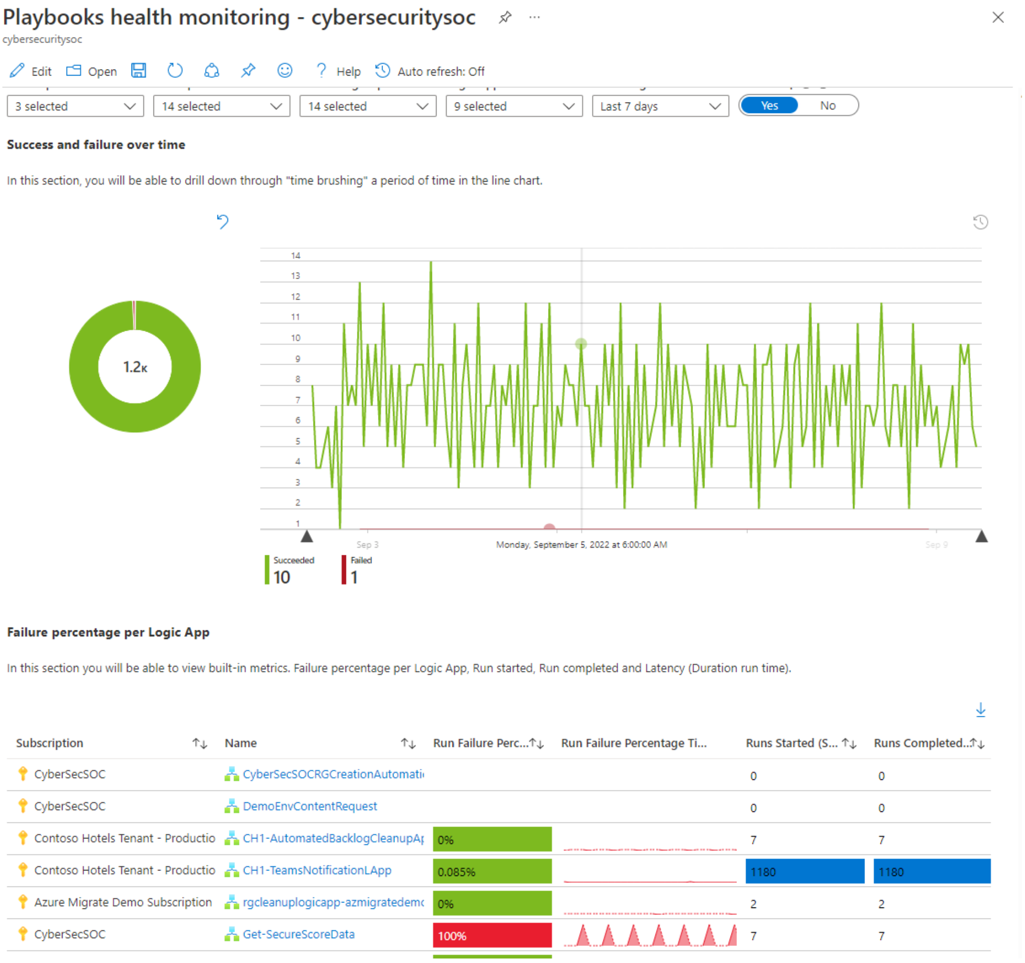

Playbooks health monitoring - cybersecuritysoc

cybersecuritysoc

3 selected | 14 selected | 14 selected | 9 selected | Last 7 days | [✓ Yes] No

Success and failure over time

In this section, you will be able to drill down through "time brushing" a period of time in the line chart.

Failure percentage per Logic App

In this section you will be able to view built-in metrics. Failure percentage per Logic App, Run started, Run completed and Latency (Duration run time).

Subscription	Name	Run Failure Perc...	Run Failure Percentage Ti...	Runs Started (S...	Runs Completed...
CyberSecSOC	CyberSecSOCRGCreationAutomati			0	0
CyberSecSOC	DemoEnvContentRequest			0	0
Contoso Hotels Tenant - Productio	CH1-AutomatedBacklogCleanupAp	0%		7	7
Contoso Hotels Tenant - Productio	CH1-TeamsNotificationLApp	0.085%		1180	1180
Azure Migrate Demo Subscription	rgcleanuplogicapp-azmigratedem	0%		2	2
CyberSecSOC	Get-SecureScoreData	100%		7	7

Microsoft Sentinel and Logic Apps integration

Security playbooks can be run either manually or automatically. Running them manually means that when you get an incident, you can choose to run a playbook on-demand as a response to the selected incident. Running them automatically means that while authoring the correlation rule, you set it to automatically run one or more automation rules or playbooks that use the alert trigger when the alert is triggered. Here are the two types of triggers that we support:

Alert Trigger

- Can only be triggered when an alert is created

Incident Trigger for Automated Response

- Used with Scheduled Analytics Rule
- Multiple Playbooks can be triggered for Incidents based on the automation rule(s) selected
- Can be triggered when an incident is created or modified (in preview)

The product team and the community shared many examples for Logic App templates that available on Github

MSSPs design considerations for automation rule and playbooks

As explained above, when dealing with SOAR in Microsoft Sentinel, we have two main components.

- Automation rule
- Playbook (Azure logic App)

Depending on the MSSP needs, we have two main models on how to deploy both playbooks and automation rules:

Model 1: Automation rules and Logic app components stored on the customer tenant.

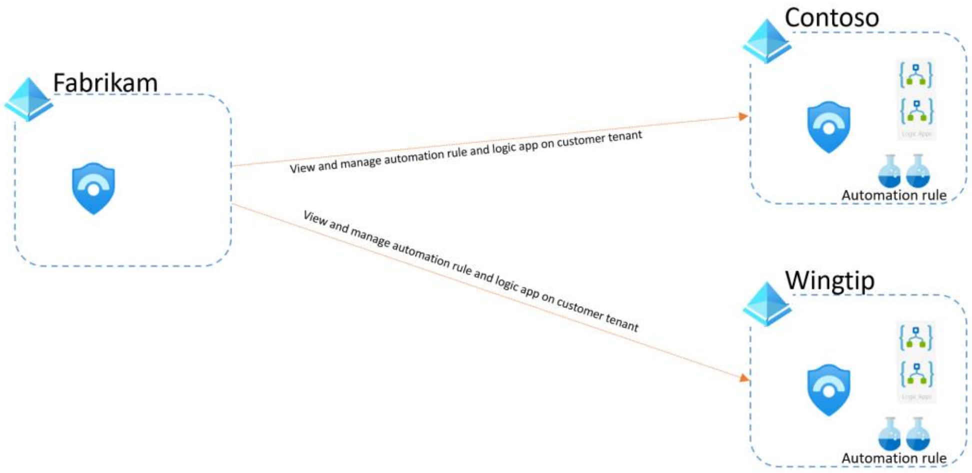
Contoso
Fabrikam
View and manage automation rule and logic app on customer tenant
View and manage automation rule and logic app on customer tenant
Automation rule
Wingtip
Automation rule

In this above example, the partner manages two customers over Azure lighthouse, the SOAR artifacts like the automation rules and playbooks are stored and run on the customer's tenant.

<u>Considerations for this model:</u>

- Logic app execution price is charged to the customer
- Monitoring for success or failure needs to happen against the customer environment

Model 2: Automation rule run on customer tenant and logic app stored and run on the partner side

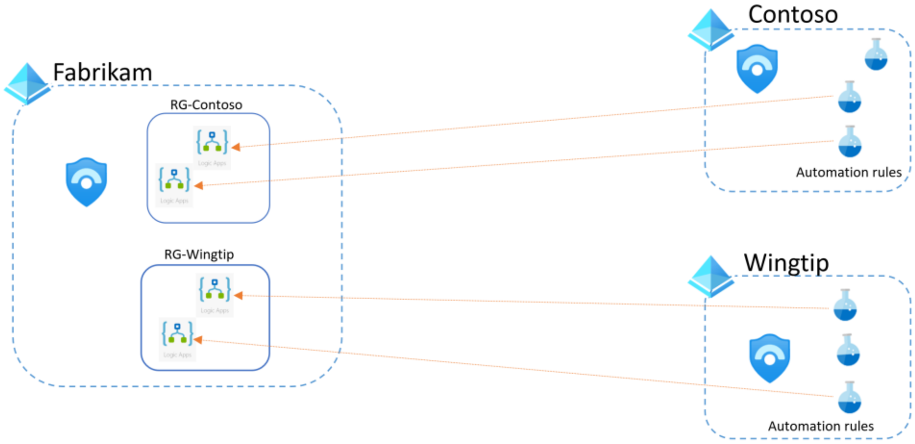

In this setup, the automation rule that runs on the customer side calls a Playbook on the MSSP side. This approach is **recommended when the MSSP needs to protect the Intellectual Property** built into the playbooks.

Considerations for this model:

- Logic app execution price is charged to the MSSP.
- Recommendation is to place the playbooks in a separate resource group per customer.

In model number 2, if the playbook needs to perform any actions in the customer tenant, we will need to create an identity object (SPN) **in that customer's tenant**.

After creating the SPN, we need to **make sure it has all the appropriate permissions** to perform all the actions required by the playbook. For example, if the playbook needs to block a user in Azure AD, the SPN will need to have the corresponding permissions that allow this task.

In the MSSP tenant, configure the logic app actions to use the identity object (SPN) that we created in the customer tenant:

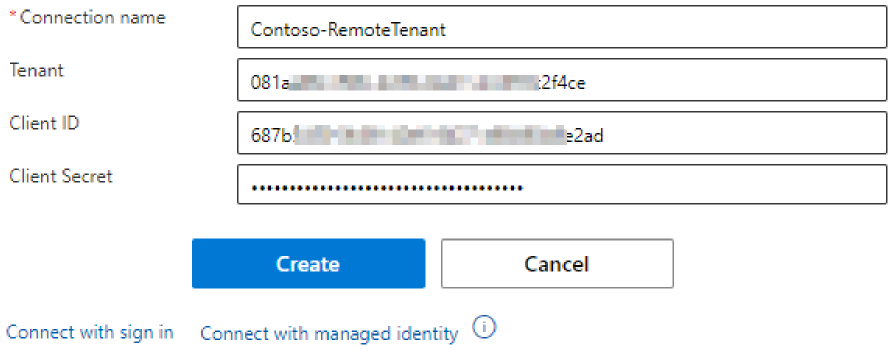

* Connection name	Contoso-RemoteTenant
Tenant	081a[redacted]2f4ce
Client ID	687b[redacted]e2ad
Client Secret	••••••••••••••••••••••••••••••••••••••

You will need to update other connectors to use the SPN as needed.

In summary, the SOAR capabilities in Microsoft Sentinel can be a huge help for MSSPs, greatly reducing the amount of time spent triaging and investigating incidents.

Threat Intelligence

What is it?

Cyber threat intelligence is information describing known existing or potential threats to systems and users. This type of information takes many forms, from written reports detailing a particular threat actor's motivations, infrastructure, and techniques, to specific observations of IP addresses, domains, and file hashes associated with cyber threats.
Threat intelligence is a primary data set used by SOC analysts to aid in detection of potential threats, prioritizing incidents, and contextualizing malicious activity. Our goal with Microsoft Sentinel is to ensure that you can easily incorporate and effectively use threat intelligence to enhance and enrich all areas of the SIEM.

Threat Intelligence can be sourced from many places, such as open-source data feeds, threat intelligence-sharing communities, commercial intelligence feeds, and local intelligence gathered in security investigations within an organization. Microsoft Sentinel lets you import the threat indicators your organization is using, which can enhance your security analysts' ability to detect and prioritize known threats.

How do you stream threat indicators to Microsoft Sentinel?

You can stream threat indicators to Microsoft Sentinel via the following built-in connectors and custom methods:

1. Using one of the integrated threat intelligence platform (TIP) products such as MISP Open Source Threat Intelligence Platform, Anomali Threat Stream, Palo Alto Networks MineMeld, ThreatConnect Platform, ElecticIQ Platform, ThreatQ Threat Intelligence Platform.
2. Connecting to TAXII servers
3. Using direct integration with the Microsoft Graph Security tiIndicators API
4. Create new threat indicators via Threat Intelligence blade on the Microsoft Sentinel portal
5. Bulk import of threat indicators from CSV or JSON file

Our recommendation is to have the same set of threat indicators imported to your customers' Microsoft Sentinel workspaces. Then in your MSSP workspace, you can leverage cross-workspace queries to aggregate the threat indicators from the customer's workspaces and correlate them with your incident detection, investigation, and hunting experience.

The imported threat indicators can then be viewed and managed from the **Threat Intelligence blade** and queried from the **ThreatIntelligenceIndicators** table in the Logs. Threat indicators will not be deleted from Log Analytics even when they expire. To delete them in bulk, you can use the following PowerShell script: Microsoft-Sentinel-Bulk-Delete-Threat-Indicators.

To access only your active threat indicators, add the following filter clause in your queries:

```
| where Active == true
```

Several features from Microsoft Sentinel then become available or are enhanced and should be utilized:

- **Analytics** includes a set of out-of-the-box scheduled rule templates you can enable to generate alerts and incidents based on matches of log events from your threat indicators. A typical rule using threat data starts with "TI map" in its name.
- **Threat Intelligence workbook** provides summarized information about the threat indicators imported into Microsoft Sentinel and any alerts generated from analytics rules that match your threat indicators.

- **Hunting** queries allow security investigators to use threat indicators within the context of common hunting scenarios.
- **Notebooks** can use threat indicators when you investigate anomalies and hunt for malicious behaviors.
- **Playbooks** can be used to enrich alerts and remediate threats with threat data. [Microsoft Sentinel GitHub repo]has a collection of playbook templates built by both Microsoft and our partners (i.e., RiskIQ, Recorded Future, HYAS, etc.)

MITRE ATT&CK

What is it?

The MITRE ATT&CK page will provide an overview of how well a Microsoft Sentinel instance is covered according to the MITRE framework. It will show the MITRE tactics that Microsoft Sentinel uses and the techniques that are used within those tactics in an easy-to-read table view as shown below:

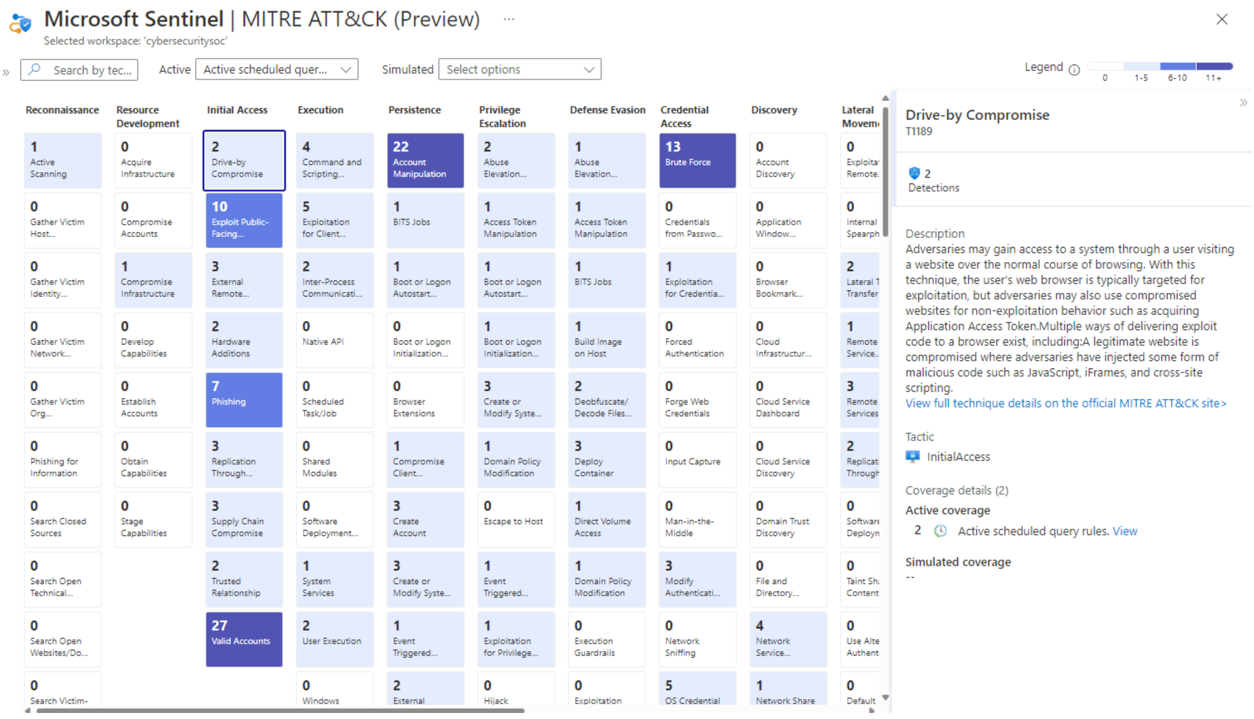

Microsoft Sentinel | MITRE ATT&CK (Preview)

Selected workspace: 'cybersecuritysoc'

Active: Active scheduled quer... Simulated: Select options

Legend: 0, 1-5, 6-10, 11+

Reconnaissance	Resource Development	Initial Access	Execution	Persistence	Privilege Escalation	Defense Evasion	Credential Access	Discovery	Lateral Movem
1 Active Scanning	0 Acquire Infrastructure	2 Drive-by Compromise	4 Command and Scripting...	22 Account Manipulation	2 Abuse Elevation...	1 Abuse Elevation...	13 Brute Force	0 Account Discovery	0 Exploita Remote.
0 Gather Victim Host...	0 Compromise Accounts	10 Exploit Public-Facing...	5 Exploitation for Client...	1 BITS Jobs	1 Access Token Manipulation	1 Access Token Manipulation	0 Credentials from Passwo...	0 Application Window...	0 Internal Spearph
0 Gather Victim Identity...	1 Compromise Infrastructure	3 External Remote...	2 Inter-Process Communicati...	1 Boot or Logon Autostart...	1 Boot or Logon Autostart...	1 BITS Jobs	1 Exploitation for Credentia...	0 Browser Bookmark...	2 Lateral T Transfer
0 Gather Victim Network...	0 Develop Capabilities	2 Hardware Additions	0 Native API	0 Boot or Logon Initialization...	1 Boot or Logon Initialization...	1 Build Image on Host	0 Forced Authentication	0 Cloud Infrastructur...	1 Remote Service...
0 Gather Victim Org...	0 Establish Accounts	7 Phishing	0 Scheduled Task/Job	0 Browser Extensions	3 Create or Modify Syste...	2 Deobfuscate/ Decode Files...	0 Forge Web Credentials	0 Cloud Service Dashboard	3 Remote Services
0 Phishing for Information	0 Obtain Capabilities	3 Replication Through...	0 Shared Modules	1 Compromise Client...	1 Domain Policy Modification	3 Deploy Container	0 Input Capture	0 Cloud Service Discovery	2 Replicati Through
0 Search Closed Sources	0 Stage Capabilities	3 Supply Chain Compromise	0 Software Deployment...	3 Create Account	0 Escape to Host	1 Direct Volume Access	0 Man-in-the-Middle	0 Domain Trust Discovery	0 Software Deployn
0 Search Open Technical...		2 Trusted Relationship	1 System Services	3 Create or Modify Syste...	1 Event Triggered...	1 Domain Policy Modification	3 Modify Authenticati...	0 File and Directory...	0 Taint Sh Content
0 Search Open Websites/Do...		27 Valid Accounts	2 User Execution	1 Event Triggered...	1 Exploitation for Privilege...	0 Execution Guardrails	0 Network Sniffing	4 Network Service...	0 Use Alte Authent
0 Search Victim-			0 Windows	2 External	0 Hijack	0 Exploitation	5 OS Credential	1 Network Share	0 Default

Drive-by Compromise
T1189

2
Detections

Description
Adversaries may gain access to a system through a user visiting a website over the normal course of browsing. With this technique, the user's web browser is typically targeted for exploitation, but adversaries may also use compromised websites for non-exploitation behavior such as acquiring Application Access Token.Multiple ways of delivering exploit code to a browser exist, including:A legitimate website is compromised where adversaries have injected some form of malicious code such as JavaScript, iFrames, and cross-site scripting.
View full technique details on the official MITRE ATT&CK site>

Tactic
InitialAccess

Coverage details (2)

Active coverage
2 Active scheduled query rules. View

Simulated coverage
--

The numbers show how many active rules use the specific technique. Clicking on the individual technique will open a blade on the right side to show more detailed information as shown above.

In addition, there are options under the "Simulated" drop down to see how enabling certain types of queries will assist with the coverage.

For more information go [here]. If you want to download the data into a comma separated value (CSV) file, you can download the "Export-AzSentinelMITREToCSV" PowerShell file from [here].

Analytics Rules

Once you have connected your data sources to Microsoft Sentinel, you'll need to create queries that will execute on a schedule and detect suspicious activities. Microsoft Sentinel comes with hundreds of built-in templates written by Microsoft's security experts and analysts based on known threats, common attack vectors, and suspicious activity escalation chains. New analytics rules also come with some of the new solutions deployed from the content hub. You have more information about analytics rules in general here, and specific to scheduled rules here.

But what is specific to MSSPs when it comes to Analytics Rules?

Cross-workspace Analytics Rules

As with other resource types, the key feature is to make Analytics Rules work across workspaces and ultimately across Azure AD tenants, these are called cross-workspace Analytics Rules. Here is an example:

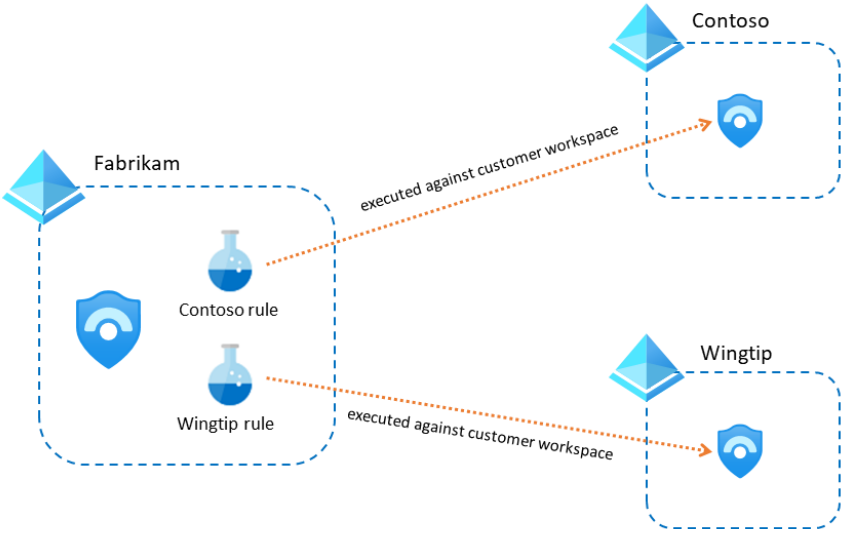
Contoso
Fabrikam
executed against customer workspace
Contoso rule
Wingtip rule
Wingtip
executed against customer workspace

In this setup, we call ***local*** the workspace created in the MSSP tenant and ***remote*** the one created in the customer tenant.

The KQL query of *Contoso rule* would be like this:

```
workspace('contoso_workspace').SecurityEvent

| where EventID == '4625'
```

If required, we can also add multiple workspaces (up to **20** workspaces) to the query, so the rule can aggregate or correlate data from multiple workspaces:

```
workspace('contoso_workspace').SecurityEvent

| union workspace('wingtip_workspace').SecurityEvent

| where EventID == '4625'
```

As you can imagine, if you add 20 workspaces to the query, it can become easily unmanageable. In order to make these queries easier to use and read, you can create a KQL function that serves as an **alias** for a query that contains multiple customer workspaces. For example, saving this query:

```
workspace('contoso').SecurityEvent
| union workspace('wingtip').SecurityEvent
```

as a function named *ContosoWingtip_SecEvents,* that way, you can write you analytics rules and hunting queries like this:

```
ContosoWingtip_SecEvents
| where EventID == '4625'
```

Also, take into account that KQL functions can be automated via PowerShell, ARM or API, so you could easily setup automation to update your KQL functions whenever you add a new customer.

Some important facts about cross-workspace analytics rules:

- The **Microsoft Sentinel solution needs to be installed on both source and target Log Analytics workspaces**.
- You can only add **up to 20 workspaces** to a given cross-workspace analytics rule, although we recommend keeping it under 5 for good performance
- Microsoft Sentinel **incidents and alerts** raised by a cross-workspace analytics rule, **will only be created in the workspace where the rule was defined** (they will not show up in the "remote" workspaces)
- Entities from the remote tenant would be available in the source workspace, but visual investigation won't fully work
- They are **only available for Scheduled analytics rules**. Fusion, Microsoft, and ML Behavior Analytics will only look at the data in the workspace where these rules are enabled
- **Use caution**. As you add more workspaces to the query, the performance may decrease, and rules might end up failing. If possible, avoid adding multiple workspaces into a rule if there's not a strong reason to do so. Further information about query optimization can be found here.

When should we use cross-workspace analytics rules?

- When the analytics rule needs to correlate data stored in multiple workspaces (rarely needed in MSSP setup)
- To protect the Intellectual Property created as part of an analytics rule

Apart from these two scenarios, the scheduled analytics rule should be created in the customer workspace. Here is a sample on how this could look like:

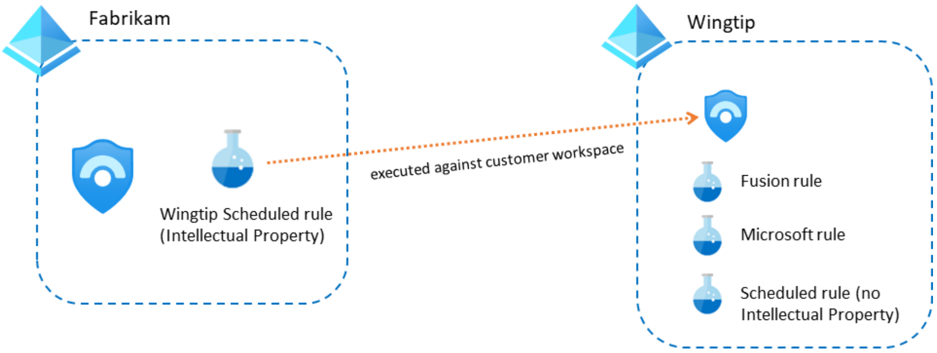
Fabrikam
Wingtip Scheduled rule (Intellectual Property)
executed against customer workspace
Wingtip
Fusion rule
Microsoft rule
Scheduled rule (no Intellectual Property)

As you can see above, we only create in our tenant the analytics rules that contain MSSP intellectual property. The rest should be created in the customer workspace. Also notice that in this situation you will have alerts and incidents in both workspaces (local and remote).

Managing analytics rules in the MSSP tenant

Just as a reminder, you should only **create rules in the MSSP tenant if you need to protect your intellectual property**. If that's the case, you would need to manage those rules appropriately.

As you may have noticed, in the examples above we always create a new analytics rule for each customer. For example, if we have a rule named "Suspicious logins in Azure AD", we will create one rule for Contoso named "Contoso – Suspicious logins in Azure AD" and another one for Wingtip name "Wingtip – Suspicious logins in Azure AD". Why not create a single rule that looks at multiple customer workspaces in parallel? There are several reasons:

- Query performance can be greatly impacted if customer workspaces are located in multiple regions.
- If an alert is triggered, it might be difficult to identity which customer is the alert coming from (there are methods to work around this, but they add complexity).
- The limit of workspaces in the alert rule is 20, so if you have more than 20 customers, you will still need to breakdown by groups of customers.

Because of the above reasons, we recommend having a single rule per customer as in the pictures above. As you can imagine, this can result in a high number of alert rules created in the MSSP tenant…imagine 50 rules per customer and having 100 customers, it would result in 5000 rules.

If you expect having a big number of proprietary rules and customers, you need to keep in mind that **a Microsoft Sentinel workspace has a limit of 512 analytics rules**. To work around this limitation, you could use an architecture where you create one Microsoft Sentinel workspace in the MSSP tenant for each customer that you manage. This would look like this:

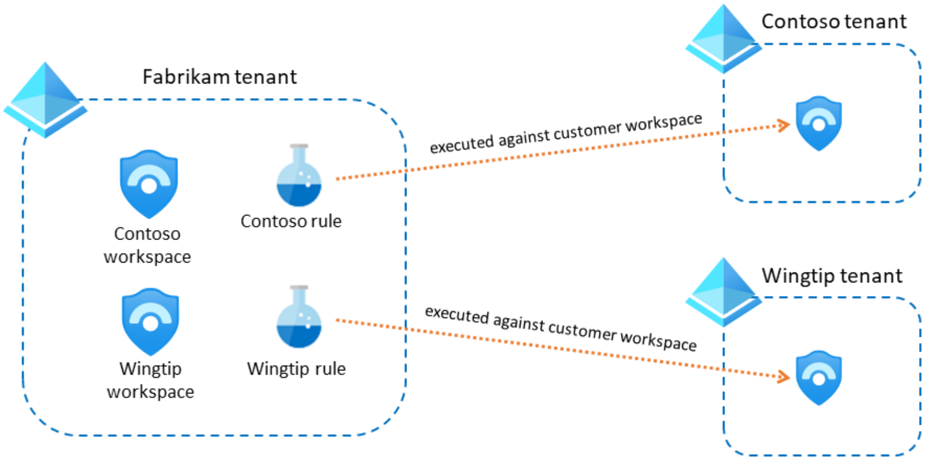

In this architecture, **we also recommend hosting each customer workspace in a separate resource group**. This resource group can also be the container of other customer-related artifacts like playbooks and workbooks (see Automation chapter for similar architecture pattern).

Cross-workspace incident view

As explained in this article, you can see and manage incidents being raised across workspaces. This feature is documented here and currently **has a limit of 100 workspaces**. This limit is expected to increase in the upcoming months, but we need to keep in mind how many customers can a group of analysts monitor in parallel without being overloaded and/or inefficient. If we assign a big number of customers to a team of analysts, it might not be able to cope with the number of incidents and the quality of the managed service offered to the customer will decrease.

As you may have noticed by now, **you may have alerts and incidents showing up in two or more workspaces for a given customer**. This can lead to situations where you can only monitor 5 customers (in the best-case scenario) with the current limit of workspaces in the incident view, because you have to select the local (where alerts/incidents coming from Fusion, ML Behavior Analytics, Microsoft and non-proprietary Scheduled rules are) and remote (where alerts/incidents coming from your Scheduled proprietary rules are) workspaces for each customer.

To alleviate this situation, you can use a workbook that shows full list of incidents regardless of the 100 workspaces limitation in the cross-workspace incident view. An example of this workbooks is the Microsoft Sentinel Central workbook shown here (see more in next section about Workbooks)

Rules Migration

When migrating from existing SIEM, the migration of Detection Rules usually becomes a requirement as part of the migration effort.

Please refer to the following content to learn more **about our recommendations for ways to migrate analytic rules from Splunk/QRadar/ArcSight:**
Webinar: Best practices converting detections rules from existing SIEM to Microsoft Sentinel
How do you export QRadar offenses to Microsoft Sentinel
Splunk to Kusto Query Language map
Plan your migration to Microsoft Sentinel | Microsoft Learn

Analytics rules health

Once the analytic rules have been setup in all your customers, it is imperative to make sure they are working correctly. A workbook called "Sentinel Central" has been created that will show information on the rules, hunting queries, saved searched, and retention policies for each of the selected Microsoft Sentinel instances.

As shown below, the "Resource Health" tab will allow for 1 or more resource types including Analytic rules, Playbooks, Data Connectors, and Automation rules to be selected (or all of them), to get an overview of the health of those resources.

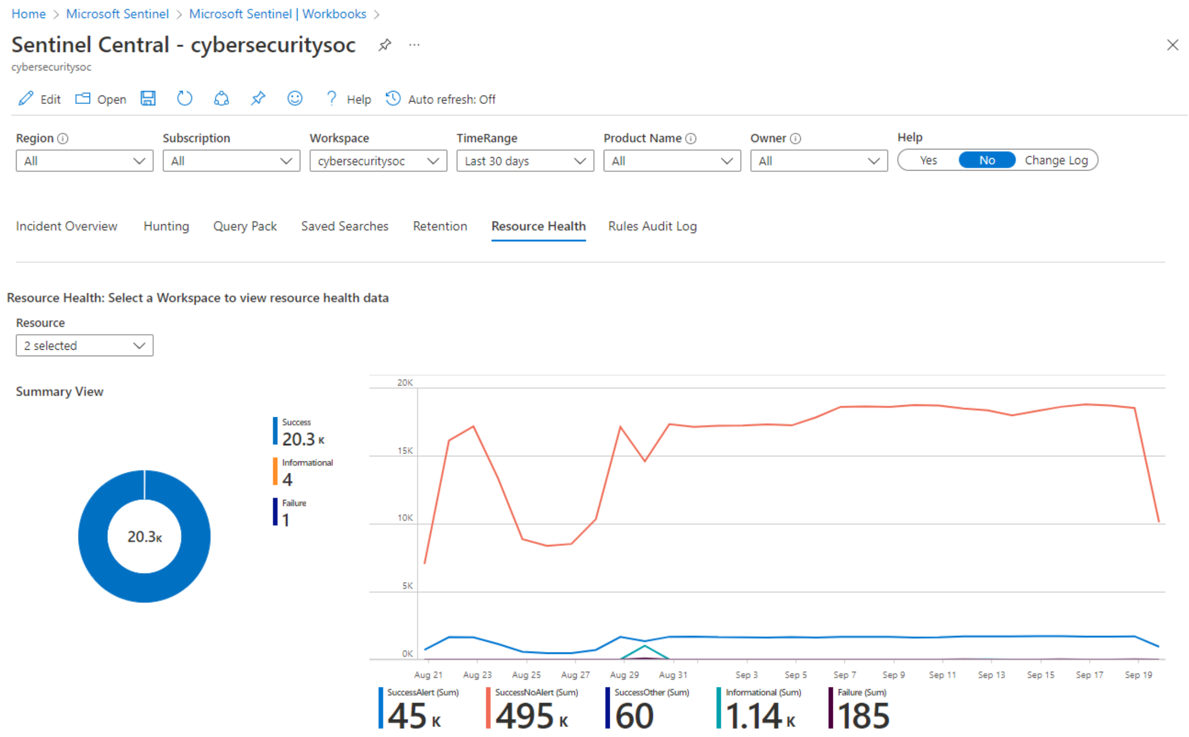

Home > Microsoft Sentinel > Microsoft Sentinel | Workbooks >

Sentinel Central - cybersecuritysoc

cybersecuritysoc

Region: All
Subscription: All
Workspace: cybersecuritysoc
TimeRange: Last 30 days
Product Name: All
Owner: All
Help: Yes [✓ No] Change Log

Resource Health: Select a Workspace to view resource health data

Resource: 2 selected

Below that, the run history of each Resource will be shown, followed by the run history of the selected resource.

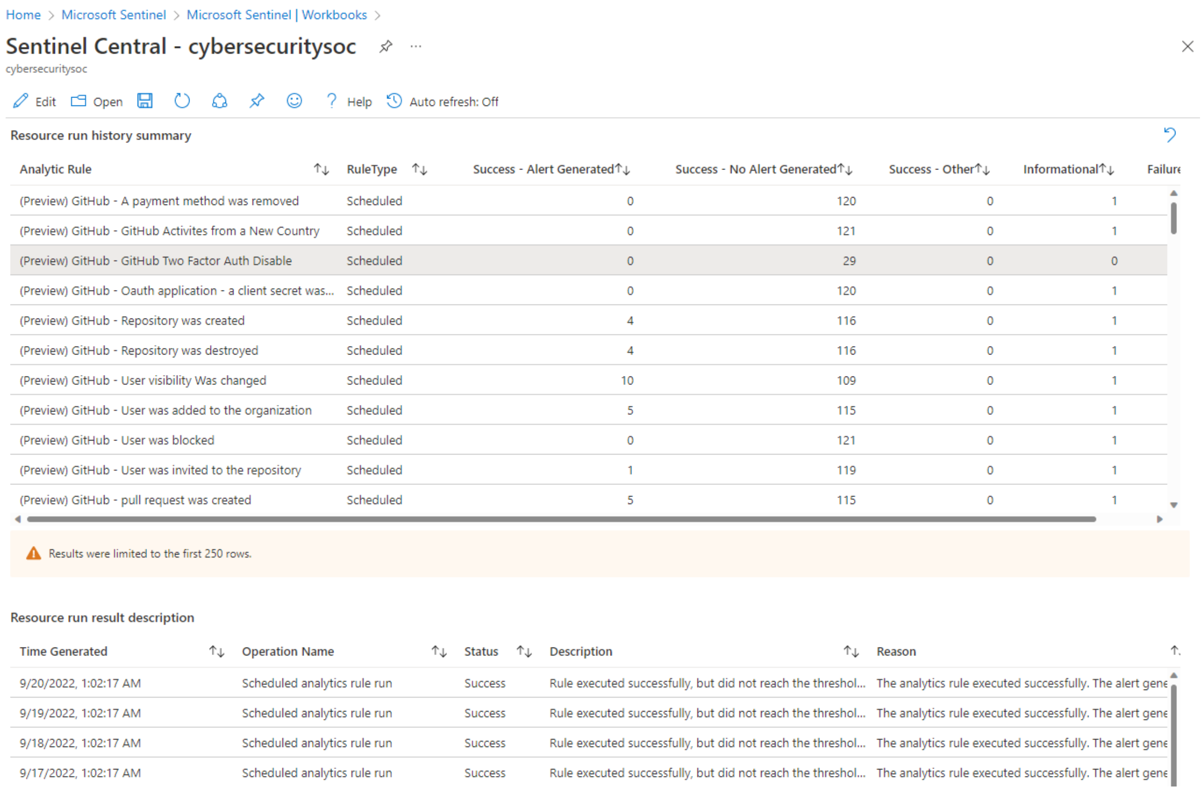
Home > Microsoft Sentinel > Microsoft Sentinel | Workbooks >
Sentinel Central - cybersecuritysoc
cybersecuritysoc
Edit
Open
Help
Auto refresh: Off

Resource run history summary

Analytic Rule	RuleType	Success - Alert Generated	Success - No Alert Generated	Success - Other	Informational	Failure
(Preview) GitHub - A payment method was removed	Scheduled	0	120	0	1	
(Preview) GitHub - GitHub Activites from a New Country	Scheduled	0	121	0	1	
(Preview) GitHub - GitHub Two Factor Auth Disable	Scheduled	0	29	0	0	
(Preview) GitHub - Oauth application - a client secret was...	Scheduled	0	120	0	1	
(Preview) GitHub - Repository was created	Scheduled	4	116	0	1	
(Preview) GitHub - Repository was destroyed	Scheduled	4	116	0	1	
(Preview) GitHub - User visibility Was changed	Scheduled	10	109	0	1	
(Preview) GitHub - User was added to the organization	Scheduled	5	115	0	1	
(Preview) GitHub - User was blocked	Scheduled	0	121	0	1	
(Preview) GitHub - User was invited to the repository	Scheduled	1	119	0	1	
(Preview) GitHub - pull request was created	Scheduled	5	115	0	1	

Results were limited to the first 250 rows.

Resource run result description

Time Generated	Operation Name	Status	Description	Reason
9/20/2022, 1:02:17 AM	Scheduled analytics rule run	Success	Rule executed successfully, but did not reach the threshol...	The analytics rule executed successfully. The alert gene
9/19/2022, 1:02:17 AM	Scheduled analytics rule run	Success	Rule executed successfully, but did not reach the threshol...	The analytics rule executed successfully. The alert gene
9/18/2022, 1:02:17 AM	Scheduled analytics rule run	Success	Rule executed successfully, but did not reach the threshol...	The analytics rule executed successfully. The alert gene
9/17/2022, 1:02:17 AM	Scheduled analytics rule run	Success	Rule executed successfully, but did not reach the threshol...	The analytics rule executed successfully. The alert gene

This will be very useful when making sure that your customer's Analytic Rules, Playbooks, Data Connectors, and Automation rules are working correctly.

The "Rules Audit Log" tab will show all the Analytic rules for the selected workspace and then the history of the selected rule as shown below. Scroll to the right to see additional fields that may have changed.

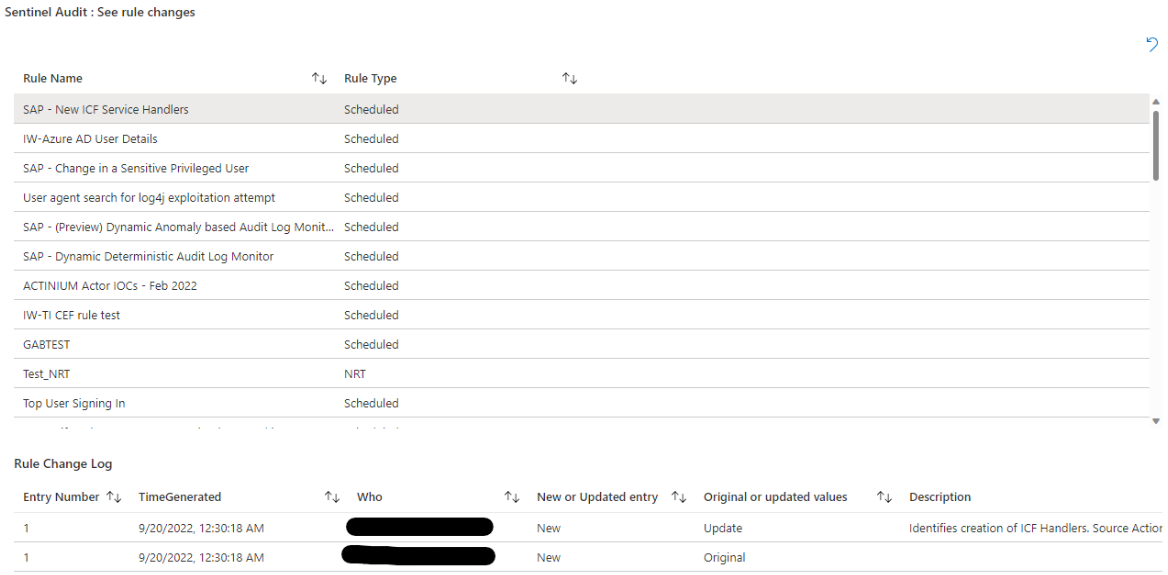
Sentinel Audit : See rule changes

Rule Name	Rule Type
SAP - New ICF Service Handlers	Scheduled
IW-Azure AD User Details	Scheduled
SAP - Change in a Sensitive Privileged User	Scheduled
User agent search for log4j exploitation attempt	Scheduled
SAP - (Preview) Dynamic Anomaly based Audit Log Monit...	Scheduled
SAP - Dynamic Deterministic Audit Log Monitor	Scheduled
ACTINIUM Actor IOCs - Feb 2022	Scheduled
IW-TI CEF rule test	Scheduled
GABTEST	Scheduled
Test_NRT	NRT
Top User Signing In	Scheduled

Rule Change Log

Entry Number	TimeGenerated	Who	New or Updated entry	Original or updated values	Description
1	9/20/2022, 12:30:18 AM	[redacted]	New	Update	Identifies creation of ICF Handlers. Source Actior
1	9/20/2022, 12:30:18 AM	[redacted]	New	Original	

This will be useful to determine if a rule has stopped working, to see if it has been modified and what has changed.

Microsoft Sentinel Workbooks

What is it?

Azure Workbooks provide a flexible canvas for data analysis and the creation of rich visual reports within the Azure portal. They allow you to tap into multiple data sources from across Azure and combine them into unified interactive experiences.

How does it work?

Workbooks can query data from multiple sources within Azure. Authors of workbooks can transform this data to provide insights into the availability, performance, usage, and overall health of the underlying components. For instance, security logs from Azure Active Directory and Azure Security Center can display the results as a grid in an interactive report allowing you to easily combine multiple data sources.

What does it do for you?

The real power of workbooks is the ability to combine data from disparate sources within a single report. This allows for the creation of composite resource views or joins across resources enabling richer data and insights that would otherwise be impossible.

Workbooks are currently compatible with the following data sources:

- Logs
- Metrics
- Azure Resource Graph
- Alerts (Preview)
- Workload Health
- Azure Resource Health
- Azure Data Explorer

Workbooks provide a rich set of capabilities for visualizing your data. For detailed examples of each visualization type you can consult the example links below:

- Text
- Charts
- Grids
- Tiles
- Trees
- Graphs
- Composite bar

Microsoft Sentinel provides workbooks templates around each data connector, these templates can be used as references to add your own modifications, or you can enjoy the updates of the templates provided by the Microsoft Sentinel community and the Microsoft Sentinel development teams.

Example of template management through Microsoft Sentinel:

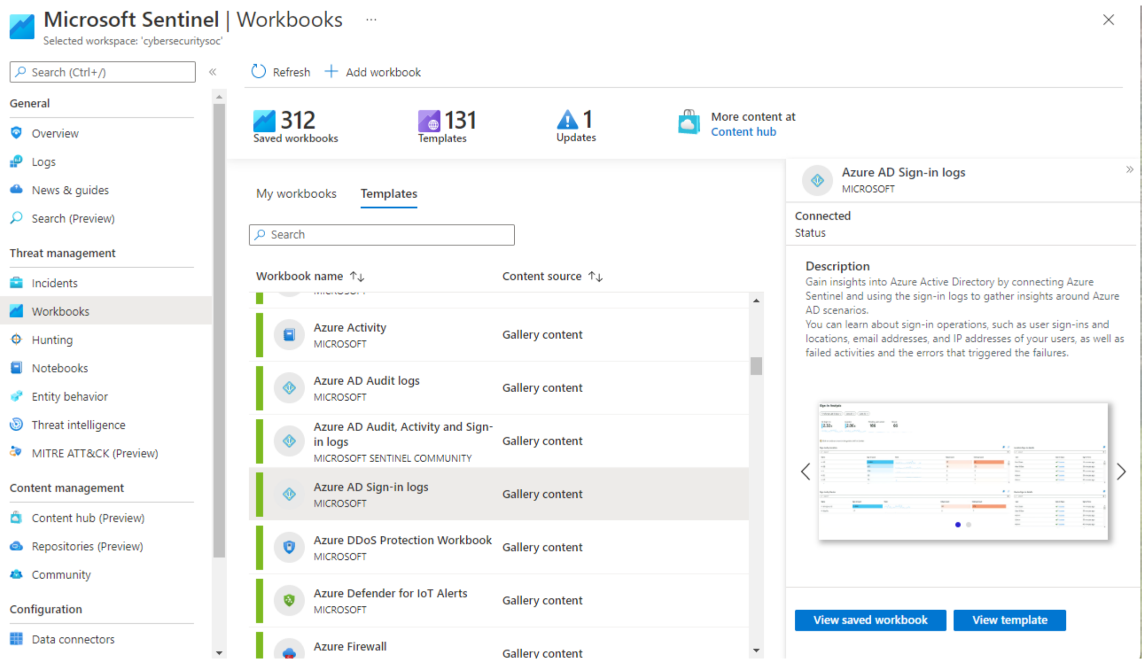

Microsoft Sentinel | Workbooks ···

Selected workspace: 'cybersecuritysoc'

312 Saved workbooks

131 Templates

1 Updates

More content at Content hub

Workbook name ↑↓	Content source ↑↓
Azure Activity MICROSOFT	Gallery content
Azure AD Audit logs MICROSOFT	Gallery content
Azure AD Audit, Activity and Sign-in logs MICROSOFT SENTINEL COMMUNITY	Gallery content
Azure AD Sign-in logs MICROSOFT	Gallery content
Azure DDoS Protection Workbook MICROSOFT	Gallery content
Azure Defender for IoT Alerts MICROSOFT	Gallery content
Azure Firewall	Gallery content

Azure AD Sign-in logs

MICROSOFT

Connected

Status

Description

Gain insights into Azure Active Directory by connecting Azure Sentinel and using the sign-in logs to gather insights around Azure AD scenarios.

You can learn about sign-in operations, such as user sign-ins and locations, email addresses, and IP addresses of your users, as well as failed activities and the errors that triggered the failures.

Workbooks have a top-down approach to data manipulation, for example I'd suggest using the Azure Firewall Workbook template. This Workbook offers Time range, Workspaces (allowing cross-workspace queries) and specific resource filtering. You can add this to any Workbook to increase the range of queries, currently there are limitations on log analytics workspace queries which are listed here: [Query across resources with Azure Monitor - Azure Monitor | Microsoft Docs](limit is currently 100 workspaces in a single query)

Example of a workbook with the multi-workspace filter:

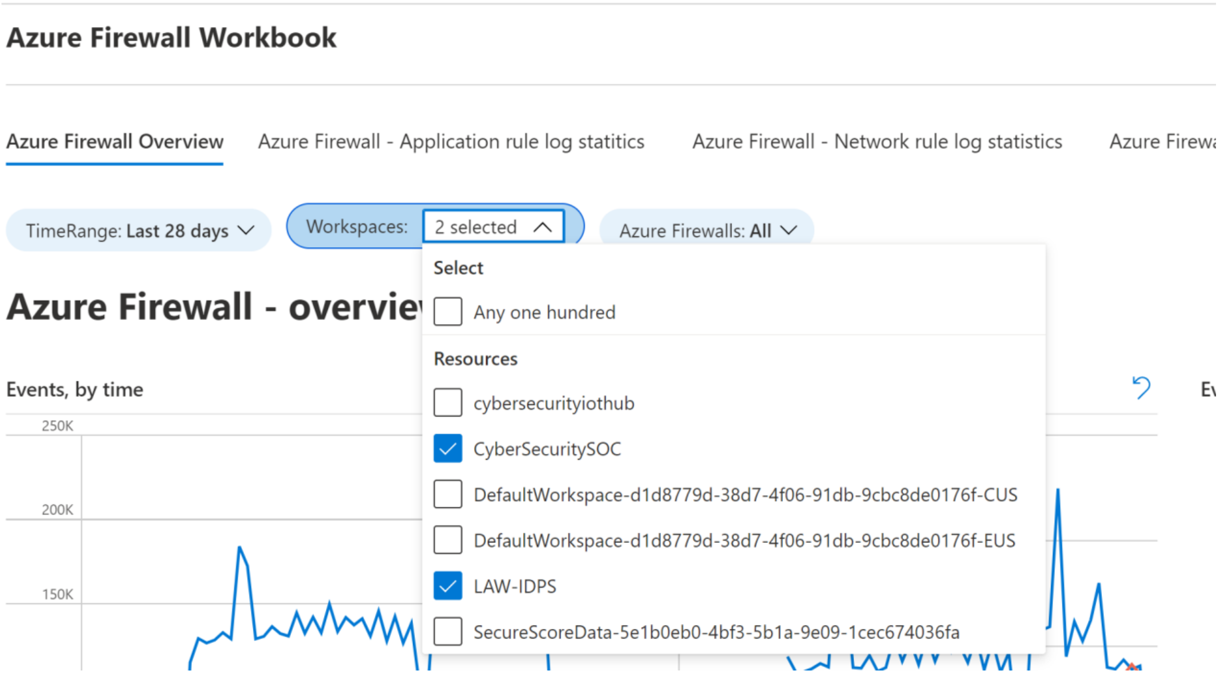

Azure Firewall Workbook

Azure Firewall Overview
Azure Firewall - Application rule log statitics
Azure Firewall - Network rule log statistics
Azure Firewa

TimeRange: Last 28 days

Workspaces: 2 selected

Azure Firewalls: All

Azure Firewall - overvie

Select

□ Any one hundred

Resources

□ cybersecurityiothub

☑ CyberSecuritySOC

□ DefaultWorkspace-d1d8779d-38d7-4f06-91db-9cbc8de0176f-CUS

□ DefaultWorkspace-d1d8779d-38d7-4f06-91db-9cbc8de0176f-EUS

☑ LAW-IDPS

□ SecureScoreData-5e1b0eb0-4bf3-5b1a-9e09-1cec674036fa

Microsoft Sentinel Central Workbook

There is one specific workbook that can be particularly important for MSSPs, its name is [Microsoft Sentinel Central]() and is authored by Clive Watson.

This workbook provides you with a cross-tenant view of the different subscriptions and workspaces managed via Lighthouse. At the top, it includes filters so you can focus on specific subscriptions or workspaces.

In the first table, it will give you a view with the number of incidents (grouped by severity) per workspace and a couple of very useful statistics to measure your service performance: mean time to triage and mean time to closure.

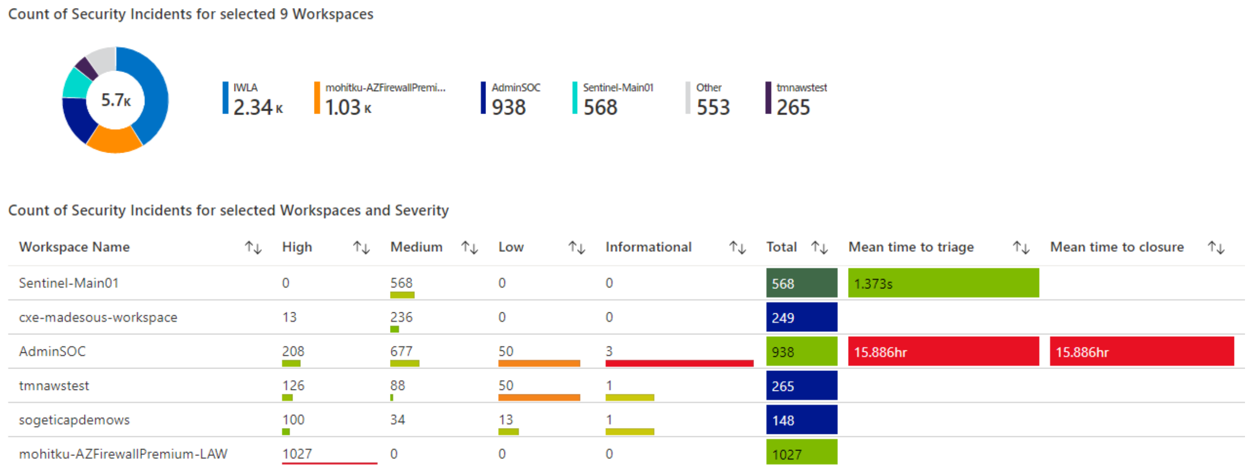

Count of Security Incidents for selected 9 Workspaces

Count of Security Incidents for selected Workspaces and Severity

Workspace Name ↑↓	High ↑↓	Medium ↑↓	Low ↑↓	Informational ↑↓	Total ↑↓	Mean time to triage ↑↓	Mean time to closure ↑↓
Sentinel-Main01	0	568	0	0	568	1.373s	
cxe-madesous-workspace	13	236	0	0	249		
AdminSOC	208	677	50	3	938	15.886hr	15.886hr
tmnawstest	126	88	50	1	265		
sogeticapdemows	100	34	13	1	148		
mohitku-AZFirewallPremium-LAW	1027	0	0	0	1027		

Below in the workbook you have a very useful view that has the list of incidents per workspace in the last 24 hrs, and links to jump to each specific incident. You also have the option to open the Investigation Insights workbook.

At the bottom you can see the full list of incidents for a given workspace.

Intellectual property protection

If you have developed your own intellectual property into a Workbook, you might need to hide it from your customers.

In order to do that, you can host the workbook in the MSSP tenant and make it multi-tenant as shown in the picture below:

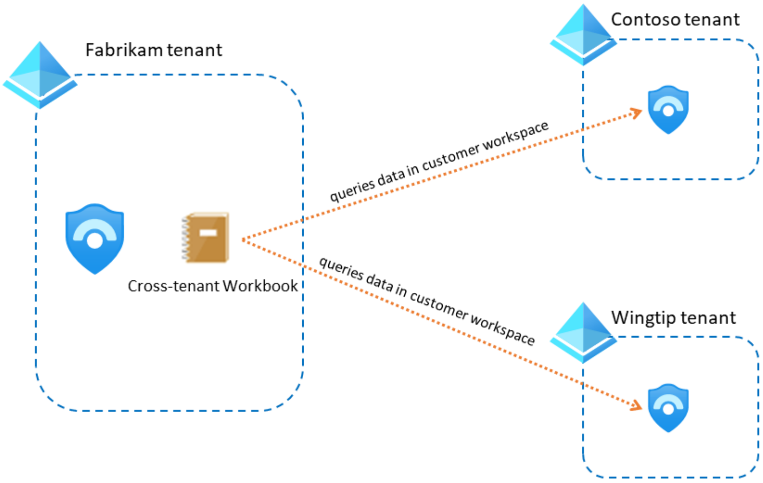

Fabrikam tenant
Cross-tenant Workbook
queries data in customer workspace
Contoso tenant
queries data in customer workspace
Wingtip tenant

This article explains how to do this step-by-step: [Making your Microsoft Sentinel Workbooks multi-tenant (or multi-workspace) - Microsoft Tech Community]()

In this previous scenario, the MSSP will have cross-customer visibility, but the **customer won't be able to access the workbook**. What if the customer needs to see the workbook visualizations but we want to keep the code underneath secret? In this case, the recommended approach is to export the workbook to **PowerBI**, as explained here. This provides several **additional benefits**, to name a few:

- Easier to share. You can just send a link to the PowerBI dashboard and the user will be able to see the report. No need to have Azure access permissions.
- Scheduling. You can configure a PowerBI to send an email on a given schedule, that will contain a snapshot of the dashboard.

Other interesting workbook resources:

- Microsoft Sentinel Central workbook by Clive Watson
- Azure Monitoring Workbooks Video - https://www.microsoft.com/en-us/videoplayer/embed/RE4B4Ap
- Microsoft Sentinel Workbooks 101 - Microsoft Sentinel Workbooks 101 (with sample Workbook) - Microsoft Tech Community
- Microsoft Sentinel Cross-Workspace - Extend Microsoft Sentinel across workspaces and tenants | Microsoft Docs

DevOps - CI/CD automation

Automation and DevOps practices are crucial components for a successful managed security practice. These are some of the key benefits of good automation:

- Reduction of human error.
- Much faster deployment and configurations.
- Improved change management as changes are tracked in source code control.
- Enhanced security as consistency is guaranteed.
- Time savings to allow employees to focus on adding value to our customers.

In this section we will try to review all the options and best practices to build a successful automation framework on top of Microsoft Sentinel.

At a high level, there are two different mechanisms to automate your Microsoft Sentinel deployment: using Microsoft Sentinel Repositories or using custom deployment methods (ARM/Bicep templates, PowerShell, Az CLI, API). Let's discuss them separately:

Microsoft Sentinel Repositories

Microsoft Sentinel Repositories allow you to create and manage your custom content from an external source control repository for continuous integration / continuous delivery (CI/CD). Currently Microsoft Sentinel supports Azure DevOps and GitHub. Repositories thus provide a way to automate Microsoft Sentinel deployment and operations of your custom content to your new and existing customers. Each Microsoft Sentinel workspace is currently limited to five repository connections. More information on how to connect your Repositories here.

Repositories content needs to be stored as ARM templates. The repositories deployment pipeline doesn't validate the content except to confirm it's in the correct JSON format. The following Microsoft Sentinel content types are supported:

- Analytic rules
- Automation rules
- Hunting queries
- Parsers
- Playbooks
- Workbooks

IMPORTANT: Microsoft Sentinel Repositories expect ARM templates to have a specific parameter name and resource type format; visit our RepositoriesSampleContent repository to find samples for each content type and useful script to convert from YAML to ARM format.

Updates made to the content in your external repositories are synced to your Microsoft Sentinel workspace and will overwrite any changes you make to that content through the Microsoft Sentinel portal. The repository that you connect to will now be your "single source of truth" for custom content in the connected workspaces. If you make changes outside of the repository, is important to ensure that the content in your repository is updated accordingly and that your deployments are happening as you expect them to.

Repositories can be especially useful for MSSPs serving multiple Sentinel customers and workspaces for content deployment. Repositories connections offer various ways for you to manage and customize the deployment experience across your customers. The first is choosing the content type(s) you'd like to deploy through a connection, which can be easily specified in the connection creation wizard. Upon creating your connection, you can easily modify the connection workflow (GitHub) or pipeline definition (Azure DevOps) to customize your trigger and deployment paths amongst other things. Learn more about customizing your connections in this article.

One of the most interesting exercises MSSPs run into when working with CICD pipelines for multiple customers is how to best structure their set of pipelines to service all their customers. There's certainly no one-size-fits-all structure, but below are three patterns to consider:

1. One repository for all, and one repository per customer

There are times when all customers need the same set of Sentinel content, whether it's the newest set of detections related to a popular cybersecurity incident, or it's a set of popular hunting queries. Having one generic content repository for all customers allows MSSPs to connect each customer's workspace to a centrally managed workspace. This allows the MSSP to add any generic content to that repository without having to deploy the content manually into each customer's repository. To complement this central repository, MSSPs can create a specific repository for each customer that needs tailored content and include their specific content there. Those per-customer repositories only need to be modified when customers need modification to their specific content. This results in each customer's workspace having two connections, one to the centrally managed repository with generic content, and another the one that is managed specifically for them. Here is a high-level diagram for this option:

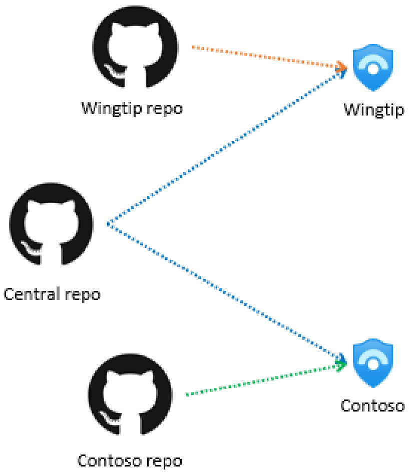

This structure is best for MSSPs that have a balance of content-for-all, and content tailored to specific customers. This same structure can be applied at a branch level if preferable to having multiple repositories (One repository with a main branch for all, and a specific branch per customer).

2. One repository for all with custom folders per customer

Some MSSPs prefer to not manage multiple repositories and/or prefer to group their content based on the shared data sources as opposed to splitting it by customer. In those cases, one structure to consider is having all your content in one repository, and connecting all your customers to this repository. You can then specify what folders each customer's connection deploys from by customizing each connections workflow or pipeline definitions as shown in this article. Here is a high-level diagram for this option:

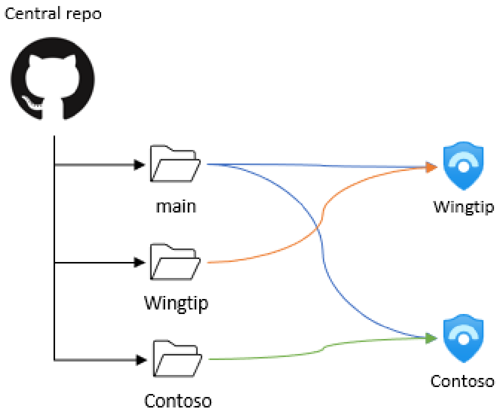

This structure offers flexibility and maintenance of a single repository but requires more upfront work during the onboarding stage to ensure that each connection is properly mapped to the folder(s) it should be deploying content from. You can opt to structure your folders based on the content they have (e.g. 'AAD Analytics folder') and map this folder to any connections that need AAD content. Another approach is naming your folders based on the customer(s) that this folder serves (e.g. 'Customer X folder') which has all of the customer's needed content and have it be mapped to their connection(s).

*Please note that anytime you update an existing connection (to include more or less content types) from the portal wizard, any workflow/pipeline customization will be reset.

3. One repository per customer

If your customers don't have any overlap in their Sentinel content needs, or the idea of a repository where multiple customers are sharing the same content is not appealing, simply creating a repository per customer might be the best method to use for managing your CICD pipelines. This allows for full separation content across your customers and can best serve customers with very specific needs across their workspaces. Here is a high-level diagram for this option:

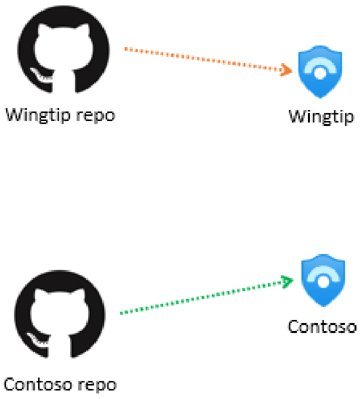

Of course, these are just example folder structures that have worked for many of our partners, and there are certainly many more that can be used to best optimize your scenarios. To further customize your CICD pipelines with any structure you use, you can utilize the newly supported configuration files with each repository branch to prioritize the deployment of high-priority content, exclude any content you want to avoid deploying, or map parameter files to their respective content files. Learn more repositories configuration files in this article.

Custom deployment methods

Microsoft Sentinel Repositories are the best way to manage your security content as code, but there might be situations where you need to use other methods, for example to deploy content types still not supported by Repositories, like Data Connectors or Watchlists. Another situation where you will need to use custom deployment methods is if you need to automate the deployment of the Log Analytics workspace and Microsoft Sentinel themselves.

Available custom deployment methods are ARM/Bicep templates, CLI (PowerShell or Az CLI) and API.

ARM/Bicep templates

ARM/Bicep templates are the default way to deploy Azure resource, and Microsoft Sentinel is no exception. At the time of writing these lines, all Sentinel content type support ARM template deployment and you can see all details about the expected format in the Microsoft Sentinel ARM/Bicep reference.

As you can see, multiple API versions are available for each resource type. Be aware that not all the latest features might be available in the latest stable version, and you might have to use one of the versions tagged as *preview*. Take into account that preview versions are subject to change. For some resource types you might also find associated quick start templates like Automation Rules. For data connectors, you can find some samples here.

Looking at the various templates, you might have noticed that there's two ways to specify the resource type in the ARM/Bicep template: one using the *scope* property and one without it. Both are equally valid. This is because all Sentinel resources are **extension resources** that are deployed on top of a Log Analytics workspace, i.e., they can't exist without an underlying workspace.

When using the *scope* property, you specify in that property the parent resource for the extension resource you're deploying, like this:

```
"type": "Microsoft.SecurityInsights/automationRules",

"name": "[<resource name>]",

"scope": "[concat('Microsoft.OperationalInsights/workspaces/', <workspace name>)]"
```

When not using *scope,* you create the Sentinel resource as follows:

```
"type": "Microsoft.OperationalInsights/workspaces/providers/alertRules",

"name": "[concat(parameters('workspace'),'/Microsoft.SecurityInsights/', <analytic rule id>)]"
```

You can use whichever method you want, but keep in mind that Repositories only supports the second one.

Take into account that **deploying Microsoft Sentinel on top of a workspace is now done using the *onboardingStates*** endpoint, you can find a sample ARM template here.

CLI (PowerShell and Az CLI)

As with any other Azure resource, Microsoft Sentinel also has support for CLI tools like PowerShell and Az CLI. In the following table you can find details about them:

CLI tool	Built-by	Framework	API coverage	Technology	Documentation
Az.SecurityInsights	Microsoft	PowerShell	SecurityInsights 2021-10-01-preview	Azure SDK for .NET	Blog and samples
az sentinel	Microsoft	Azure CLI	SecurityInsights 2021-10-01-preview	Azure SDK for .NET	Documentation

API

Microsoft Sentinel has its own RESTful API, which is called **SecurityInsights** (name of the Microsoft Sentinel resource provider). For details about the different endpoints and operations, visit Microsoft Sentinel API reference.

NOTE - As you may have noticed in the API documentation, the API version is 2021-10-01, which is the current stable API version. There are also preview versions that are documented here and are subject to change. Be careful when using the preview version, as operations and schemas might change.

Review this article if you want to get started with the Microsoft Sentinel API.

How to handle other resource providers

As you may know, Microsoft Sentinel uses resources that are not part of Microsoft.SecurityInsights, but are part of other Azure resource providers. Here is a list of those and where to find ways to deploy them:

Artifact	ARM/Bicep/Terraform	PowerShell	Azure CLI	REST API
Hunting Queries	ARM/Bicep/Terraform reference + samples	PowerShell	Az CLI	API
Parsers	ARM/Bicep/Terraform reference + samples	PowerShell	Az CLI	API
Playbooks	ARM/Bicep/Terraform reference + samples	PowerShell	Az CLI	API
Workbooks	ARM/Bicep/Terraform reference samples	N/A	N/A	N/A

Training and community resources

- Ninja Training –
  - Become a Microsoft Sentinel Ninja: The complete level 400 training - Microsoft Tech Community
  - Become a Microsoft Sentinel Automation Ninja! - Microsoft Tech Community
  - Azure Sentinel notebook ninja - the series! (microsoft.com)
- Microsoft Sentinel Learning Path-https://docs.microsoft.com/en-us/learn/paths/security-ops-sentinel/
- SC-200 – Security Operations Analyst Associate- https://docs.microsoft.com/en-us/learn/certifications/security-operations-analyst/
- Microsoft Sentinel Tech community – Blogs – https://techcommunity.microsoft.com/t5/microsoft-sentinel-blog/bg-p/MicrosoftSentinelBlog
- What's new in Microsoft Sentinel - https://docs.microsoft.com/en-us/azure/sentinel/whats-new
- Top 10 Best Practices for Azure Security – https://aka.ms/azuresecuritytop10
- Microsoft Sentinel Documentation site - https://docs.microsoft.com/en-us/azure/sentinel/
- Azure Security Podcast – https://aka.ms/azsecpod
- Microsoft Sentinel in the Field - Microsoft Sentinel in the Field - YouTube
- Azure Security Community Webinars - https://aka.ms/SecurityWebinars Azure Security Community Webinars - https://aka.ms/SecurityWebinars
- Microsoft Sentinel GitHub - https://github.com/Azure/Azure-Sentinel