CyberOne

Deploying Microsoft Sentinel Like A Pro – Your 8 Key Questions Answered

Congratulations, you've deployed Microsoft Sentinel! You've found it quite easy to set-up and are taking advantage of the out-of-the-box monitoring of all user activity in Microsoft 365. You're using the available rules to handle key incidents and the dashboard is giving you some basic data analysis. So, what next?

You now have two SIEMs in operation, which is time consuming and creating visibility challenges, not to mention the consumption costs. Your non-Azure cloud platforms aren't integrated, and Sentinel isn't connected into your 3rd party cybersecurity solutions. Your time is being swallowed handling thousands of alerts and your security appliances for syslog and CEF forwarding are log intensive and noisy. It's time to tune your Sentinel like a pro and leverage its game-changing capabilities to the max, to improve your security capabilities, lower your business cost, and make your day-to-day life easier.

Here we give you the 8 most frequently asked questions by Comtact clients who are busy bringing Microsoft Sentinel into their business, along with some initial advice to assist with your deployment.

1 How to transition from your previous SIEM to a fully rolled-out Sentinel solution?

Running two SIEM solutions in parallel is likely to be time-consuming and costly for consumption-based models. Monitoring across two solutions creates a fragmented environment and so also negatively affects visibility.

The first step in your migration journey is to fully map-out the monitoring controls in your original SIEM to enable you to export and translate into the Sentinel architecture.

Sentinel provides various capabilities to alleviate the workload and cost across dual SIEMs, such as the free ingestion of OfficeActivity, AzureActivity and Microsoft Defender events. New deployments are given 30 days of free log inges-tion to allow for initial integration set-up activity and log refinement before any charges are incurred. This provides time to convert to the Sentinel query language before analytical rules, SOAR automations, and dashboard views are tested.

To simplify migration tracking, leverage the Sentinel migration tracking workbook to comprehensive visualise and document the staged process of onboarding and migration.

2 How to integrate with your non-Microsoft infrastructure?

In order to achieve a comprehensive understanding of your organisation's full estate, it is important to scope the extent of your infrastructure footprint, both in the cloud and on-premise.

While Azure resources can be quickly integrated for visibility in Sentinel, it is important to extend this visibility across to other cloud platforms being utilised such as AWS and GCP. Sentinel provides native solutions to ingest log events from these platforms directly without the requirement for additional infrastructure. Similarly, Sentinel provides the ability to pull security events and service logs from on-premise devices via the Azure Monitor Agent.

After deploying this to devices, data collection rules should be defined for any custom log locations to allow for granular log ingestion from the systems. After expanding monitoring coverage to your additional cloud or on-premise environments, these should be incorporated into the current ruleset and have custom analytical rules defined to monitor these data sources.

3 How to select genuine alerts from a sea of false positives?

After the initial deployment of data connectors and default analytical rules into Sentinel, you are likely to face a high number of alerts which need to be refined.

To reduce the number of false positive triggers for analytical rules, we can leverage the tuning and testing capabilities built into Microsoft Sentinel. When identifying an analytical rule as being noisy and inaccurate, we can use the configuration page for this alert rule to see the number of triggers and identify the most common associated entities. Within this view, we can make changes to the query logic and test this updated logic by querying the last 50 occurrences of this rule, allowing us to see how the changes impact the number of alert triggers and the event count in each instance.

In addition to this, we can make use of the Analytics Efficiency workbook view to deep-dive into this activity for multiple analytical rules and visualise the most common closure reasons, highlighting any rules with a high false-positive count.

Using these capabilities, detection rules can be tuned and refined with exclusions applied for expected business operations to reduce the false positive count and increase the confidence factor of generated incidents.

4 How to ingest and filter out noisy data sources and appliances?

When configuring security appliances for syslog and CEF forwarding, these appliances can be log intensive and noisy. In some cases, these sources are critical for monitoring, but with a default syslog configuration these can be expensive to ingest.

There are several approaches to be considered for reducing the ingested data to limit this to only the necessary logs. Initially, the syslog configuration on these appliances should be evaluated to utilise any controls provided for granular filtering of forwarded event logs, but these are sometimes unavailable or limited to certain values such as event severity.

In conjunction with this, we can define additional specifications within the log analytics workspace to establish which log event types should be forwarded by the agent.

If more granular log filtering is still required, these filters should be applied on the syslog forwarder itself. By creating custom filtering configurations for the syslog agent, we can tune out log events based on the content of the event message itself allowing for the most granular level of log filtering.

5 How to connect Sentinel into 3rd party solutions for more actionable intelligence?

When integrating your Sentinel instance with 3rd party applications and products, there are a wealth of native connectors that can be leveraged to establish these connections.

If looking to ingest data from a 3rd party data source that doesn't have a native connector, there are a variety of options available to support this activity. It is important to evaluate the best possible ingestion method for these sources depending on their capabilities for log forwarding.

If there is a syslog capability this can be ingested using a syslog forwarder, or if the solution uses an API this can often be leveraged for data ingestion. When using APIs to pull log events from these 3rd party products, an Azure Logic App or Function App can be used to regularly query and ingest the data using a custom script.
It is crucial that once these data sources are ingested, the logs are parsed using KQL functions so that the events can be correlated and included in the ruleset.

Any further custom analytical rules and dashboard views for these data sources should be defined and deployed to expand the coverage for these solutions.

6 How to deploy at the next level – way beyond out-of-the box capabilities?

Once data connectors have been configured and log events are being ingested into Sentinel, it is important to use entity mapping to ensure that we can correlate activity more effectively. This capability allows events to be joined between data sources more effectively and enables the mapping of identified entities to threat intelligence and Sentinel watchlists.

To increase the effectiveness of analytical rules, the watchlist capability in Sentinel allows for definition and grouping of entities and assets. This can allow heightened alert sensitivity to be applied to specific asset groups and equally can be used to exclude asset groups from certain expected activities.

Equally, it is recommended to use the automation capabilities to define SOAR processes to apply immediate actions for certain incident cases. By scoping your required use cases for automation, these can be converted into automation rules within Sentinel using Azure Logic Apps. These can integrate with a vast number of external platforms to enable immediate response actions.

Using the watchlists in conjunction with the automation rules, it is possible to build comprehensive conditional automation workflows that will take different routes of action depending on the affected entity and its classification.

7 How to monitor and track the ongoing status of the environment?

Once your Sentinel instance is live and running with an actionable set-up and configuration, it is important to monitor and track the ongoing status of the environment. Using some of the pre-built workbooks within Sentinel, you can monitor many of these key areas. It is crucial to monitor the health of the Sentinel service, as well as the health of data sources that you are collecting logs from. These can be translated into analytical rules to provide immediate alerts on service health issues.

In addition to the health of the service, it is important to regularly review the statistics for log ingestion and the Workspace Usage Report provides insight into the expected costs of data ingestion and retention. By default, retention will be set to between 30 – 90 days for data sources, however these retention terms can be extended or reduced as required by operations and compliance.

Additionally, Microsoft Sentinel supports the usage of long-term log storage which should be evaluated and considered, either through exportation to Azure Data Explorer or via the Tables feature within the log analytics workspace.

8 How to maintain a comprehensive threat intelligence programme?

To leverage threat intelligence within Sentinel, it is important to consider your sources for ingestion. These can be configured directly with Threat Intelligence platforms or can be ingested via TAXII format.

It is crucial to ensure that a comprehensive threat intelligence programme is maintained and that these values can be identified and matched with log events in the ingested data tables. In addition to ingesting Threat Intelligence feeds, entities seen elsewhere in incidents and investigations that are determined to be malicious should be added as custom indicators. This allows repeated targeting via the same entities to be better detected and expands the context of the detected activity.

Additionally, all threat intelligence indicators should be managed with the appropriate lifecycles so that their expiration date is appropriate.

We are Microsoft Gold Partners specialized in deploying Microsoft Sentinel. Contact us to understand your specific deployment needs and receive tailored advice.

Get in touch →

About CyberOne

CyberOne strives to offer the most trusted cyber experience in the digital world. Founded in 2005 and based in Milton Keynes, our assessment, detection and response services empower you to stay one step ahead of the latest cyber threats. We are trusted by leading UK and global businesses and our people provide a service that is second to none.

www.cyberone.security
+44 (0)345 757575
enquiries@cyberone.security