Best Practices for Operationalization of Vectra Detect with Microsoft Sentinel

Version: May 24, 2023

Table of Contents

Introduction

This document intends to provide guidance to operationalize Vectra Detect in Microsoft Sentinel to get the most out of the Vectra Platform and its AI prioritization technology. We recommend following these guidelines as a starting point and then making adjustments that are specific to your environment and desired workflows.

Vectra Attack Signal Intelligence has 3 core components:

- AI Detection
- AI Triage
- Ai Prioritization

Those technologies work hand in hand to provide the best signal to the SOC and focus on what matters. In the scenario where Microsoft Sentinel is used as the primary tool by the SOC to manage security incidents across their Security solutions, it is key to bring this signal into Microsoft's Sentinel workflow. This guide relies on Analytic templates provided by Vectra to implement an effective approach.

This guide is for the Vectra Classic UX (syslog based using the Vectra appliance platform). An update will be released for the Vectra Respond UX (Vectra cloud platform) in the near future.

Available content in Microsoft Sentinel for Vectra Detect

The below table represents what is currently available and maintained by Vectra if you deploy our Detect solution from the content hub (or Marketplace):

MS Sentinel Data Point	Vectra Detect Solution
Package Version	2.0.5
Data Connector	1 (version 1.1.1)
Workbooks	1
Analytic Templates	7

- This document is an attachment to the Sending Vectra Detect Events to Microsoft Sentinel KB article on the Vectra support site.
    - That KB article details how to send data from Vectra Detect to Microsoft Sentinel while this document covers best practices around the operationalization of that data in Sentinel.

The most common mistake we are seeing is configuring and enabling all our Analytic templates. We strongly discourage doing this! Our analytic templates have been developed to cover different use cases and for implementing the best integration, we recommend following the guidelines shared in this document:

Template Name	Description	Version	Guideline	Incident	Severity
Vectra AI Detect - Suspected Compromised Account	Create an incident when an Account is suspected to be compromised. The higher the severity level is, the more immediate attention it requires as Vectra AI engine is more confident that this is a real threat. Level of severity are: Low, Medium, High, Critical). Recommended configuration is to trigger an alert for at least High and Critical.	1.0.5	Highly Recommended	Yes	Dynamic
Vectra Account's Behaviors	This analytic rule is looking for new attacker behaviors observed by the Vectra Platform. This rule is focused on account's detections.	1.0.1	Recommended	No	Dynamic
Vectra AI Detect - Suspected Compromised Host	Create an incident when a Host is suspected to be compromised. The higher the severity level is, the more immediate attention it requires as Vectra AI engine is more confident that this is a real threat. Level of severity are: Low, Medium, High, Critical). Recommended configuration is to trigger an alert for at least High and Critical.	1.0.5	Highly Recommended	Yes	Dynamic
Vectra Host's Behaviors	This analytic rule is looking for new attacker behaviors observed by the Vectra Platform. This rule is focused on host's detections.	1.0.1	Recommended	No	Dynamic
Vectra AI Detect - Detections with High Severity	Create an incident for high severity malicious behavior detected by Vectra AI (Threat score superior to 7.0). The Severity is a mapping with the Threat score assigned to a detection. It ranges between 0 and 10. The *severity_threshold* variable can be adjusted as desired.	1.0.6	Optional	Optional	High
Vectra AI Detect - New Campaign Detected	Identifies when a new Campaign has been detected. This occurs when multiple Detections across different Hosts are suspected to be part of the same Attack Campaign.	1.1.5	Optional	Optional	Medium
Vectra AI Detect - Suspicious Behaviors	Create an incident for each new malicious behavior detected by Vectra Detect. By default, it looks through all tactics. This can be modified to create incident only for a subset of tactics. Example of a common use case: "*Create an incident when a detection in the category Exfiltration is triggered*"	1.0.7	Optional	Optional	Dynamic

Recommended workflows

Overview

The Vectra philosophy is to focus on prioritized entities versus individual detections. The recommended configuration and workflow follow that same logic.

Incident creation for prioritized entities

Analytic templates to enable:

- Vectra AI Detect - Suspected Compromised Host
- Vectra AI Detect - Suspected Compromised Account

High level view:

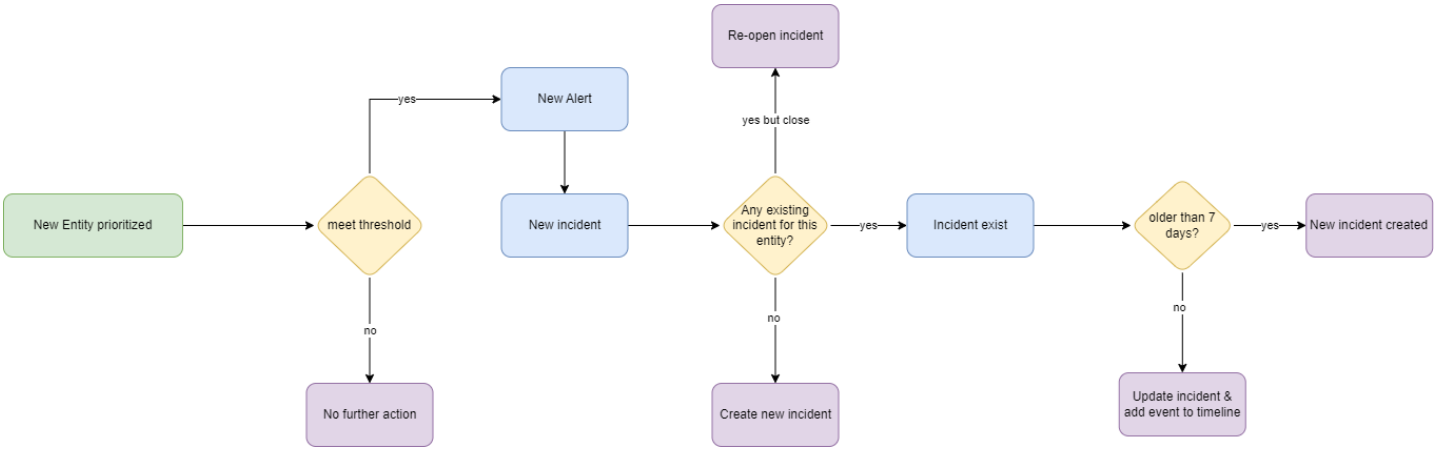

Alert creation for detections

Analytic templates to enable:

- Vectra Host's Behaviors
- Vectra Account's Behaviors

The recommended workflow to handle detections associated to entities would be:

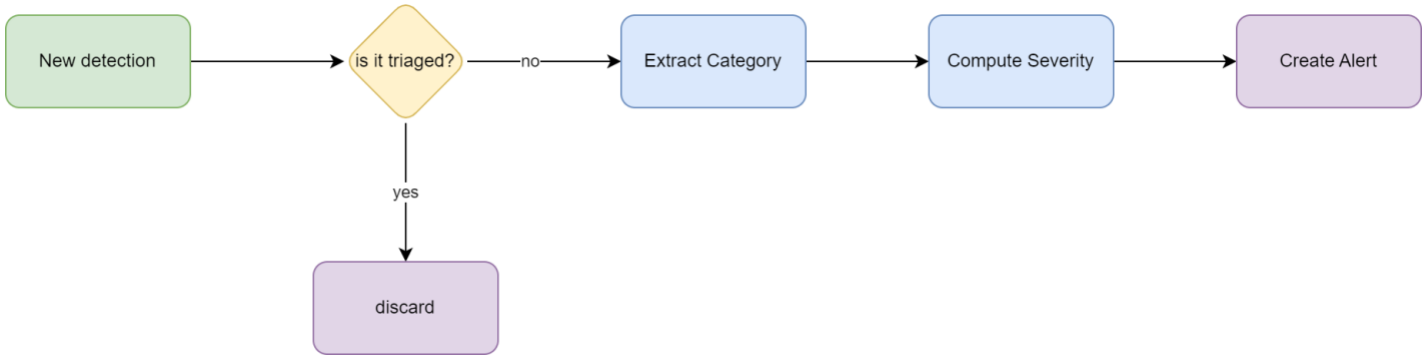

Anatomy of a Microsoft Sentinel incident for Vectra's signal

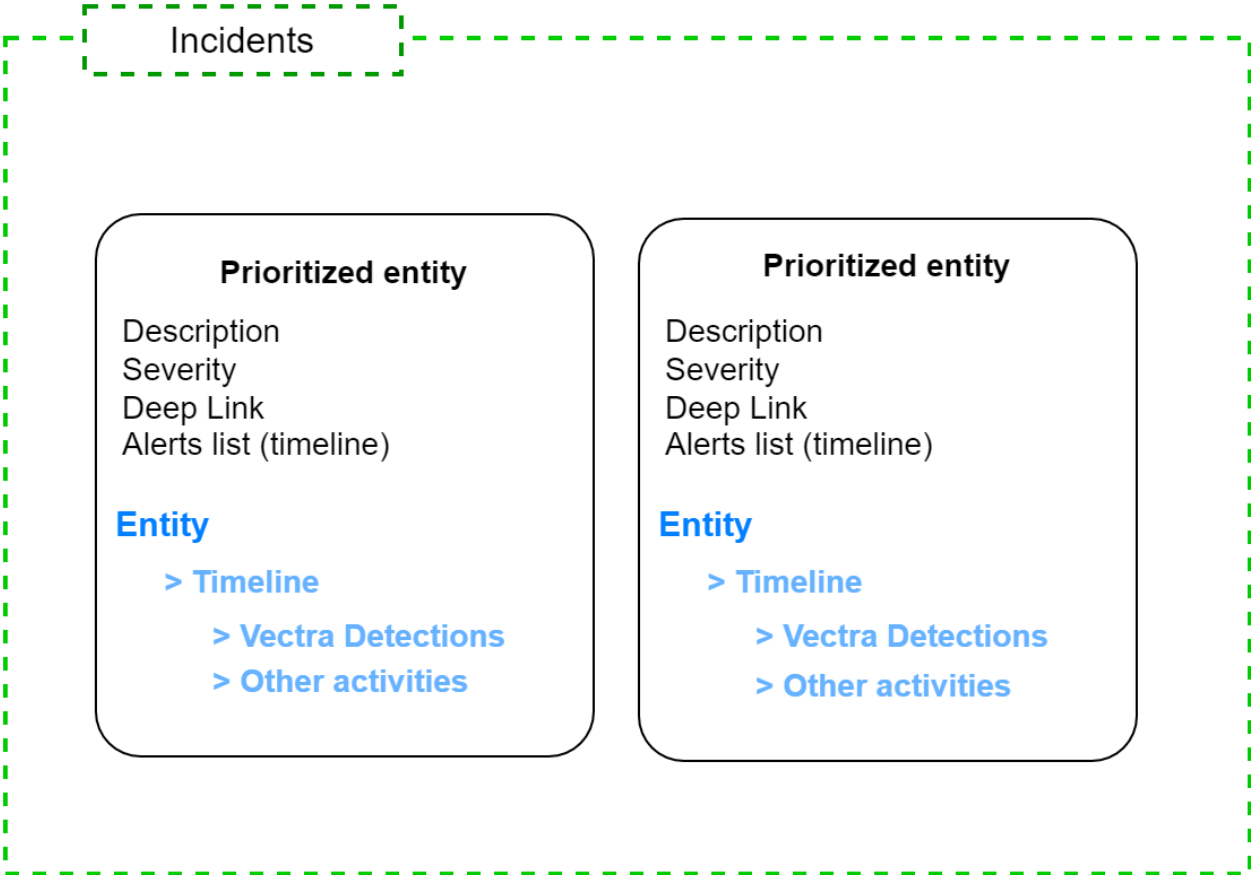

The creation of the incident is driven by Analytic rules (created from Vectra's provided Analytic templates) per entity type:

- Description shows the Threat and Certainty scores for this entity.
- Deep link is provided to pivot to the Vectra UI.
- Incident's severity is dynamically computed based on Vectra's severity score.
- Map the incident to one Entity (Host or account)
  - Note that it would be unlikely to be attached to multiple entities.
- The Entity's timeline also shows activities, including the behaviors (detections) generated by the Vectra Platform.
  - The analyst can decide to add them to the incident timeline.

!! Please Note: Vectra's detections are written to the *SecurityAlert* table but are not promoted to be incidents (created from Vectra's provided Analytic templates). The entity mapping allows correlation of detections to the entity. Alerts are visible in the timeline view of the entity. An alternative approach would be to use some automation to automatically add those alerts to the incident timeline. This is not done automatically today but is a potential for future development. Please contact your Vectra account team with feedback.

Analytic Templates logic

Priority mapping

The severity level in Vectra's classic UX is slightly different than the one in Sentinel. This is mapping that is applied in our Analytic Templates:

Vectra Severity	Sentinel Severity
Informational	Informational
Low	Low
Medium	Medium
High	High
Critical	High

Our analytic templates have a default severity of "Informational". The correct severity is set dynamically according to the above mapping. The severity of the Alert or the Incident is automatically mapped to Vectra's severity.

Frequency

This defines when a threat is identified and prioritized. Time to investigate is critical. Our analytic templates use the following parameters:

- Run every 5 minutes.
- Look at the events from last 5 minutes.

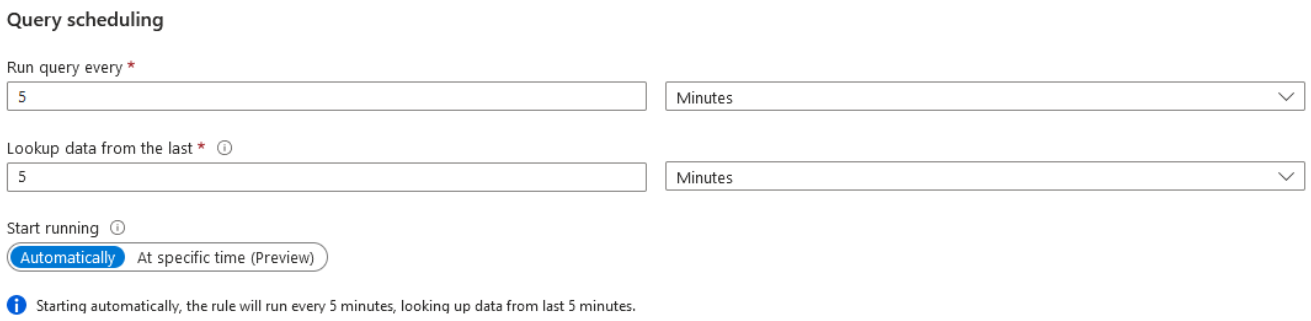

Alerts

Alert (Event) grouping must be disabled (this is set by default in our Analytics Templates). The Kusto query is already grouping by entities or detection and takes care of duplicates. A single alert or incident must be created for each entry of the result table.

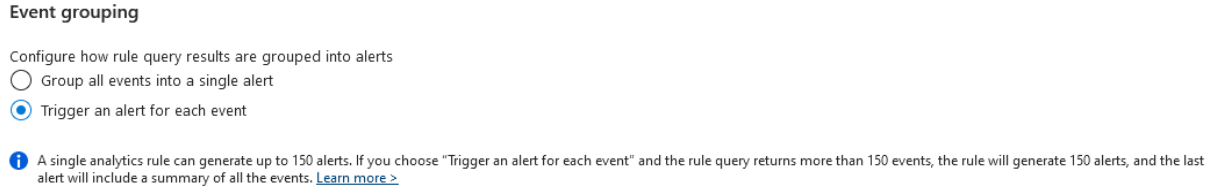

Incidents

There are few things we tried to avoid:

- Creating duplicate incidents.
- Creating too many incidents.
- Creating incidents every time a detection shows up.

The table below provides a mapping between our Analytics templates and what we recommend:

Analytic Template Name	Type	Grouping settings
Vectra AI Detect - Suspected Compromised Host	Incident	7 days (per entity)
Vectra Host Behaviors	Alert	None
Vectra AI Detect - Suspected Compromised Account	Incident	7 days (per entity)
Vectra Account Behaviors	Alert	None

We do not include other analytic templates as their configuration would depend on the use case that needs to be addressed. We do not recommend enabling them by default as they could generate too many incidents if not properly tuned.

Step by step configuration of recommended workflows

Incident creation for prioritized hosts

First, we want to create Incidents when Hosts are prioritized in the Vectra appliance platform. To do so, we are going to modify the analytic template named: Vectra AI Detect - Suspected Compromised Host.

- Menu: *Configuration* > *Analytics* > *Rule templates*.
- Search for Vectra and select the template name **Vectra AI Detect - Suspected Compromised Host.**
- Click "Create rule".

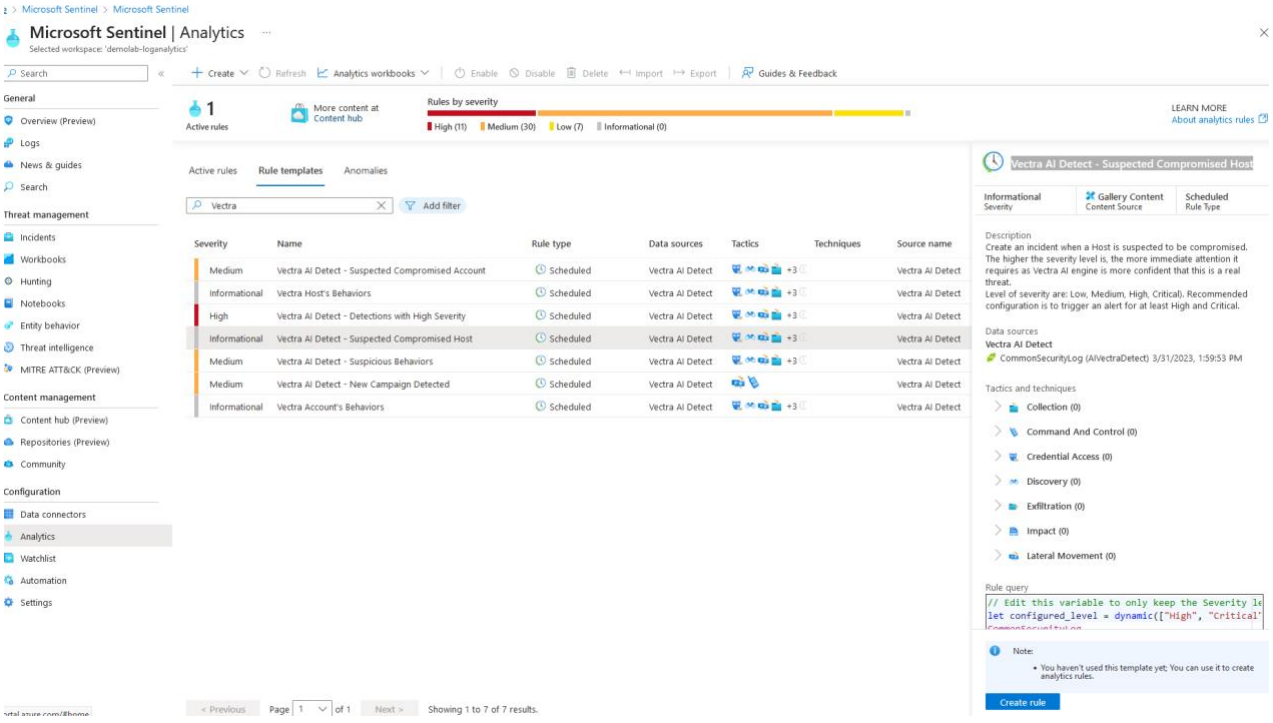

- Set the status to **Enabled** and make sure than Severity is set to **Informational** (it should be default).

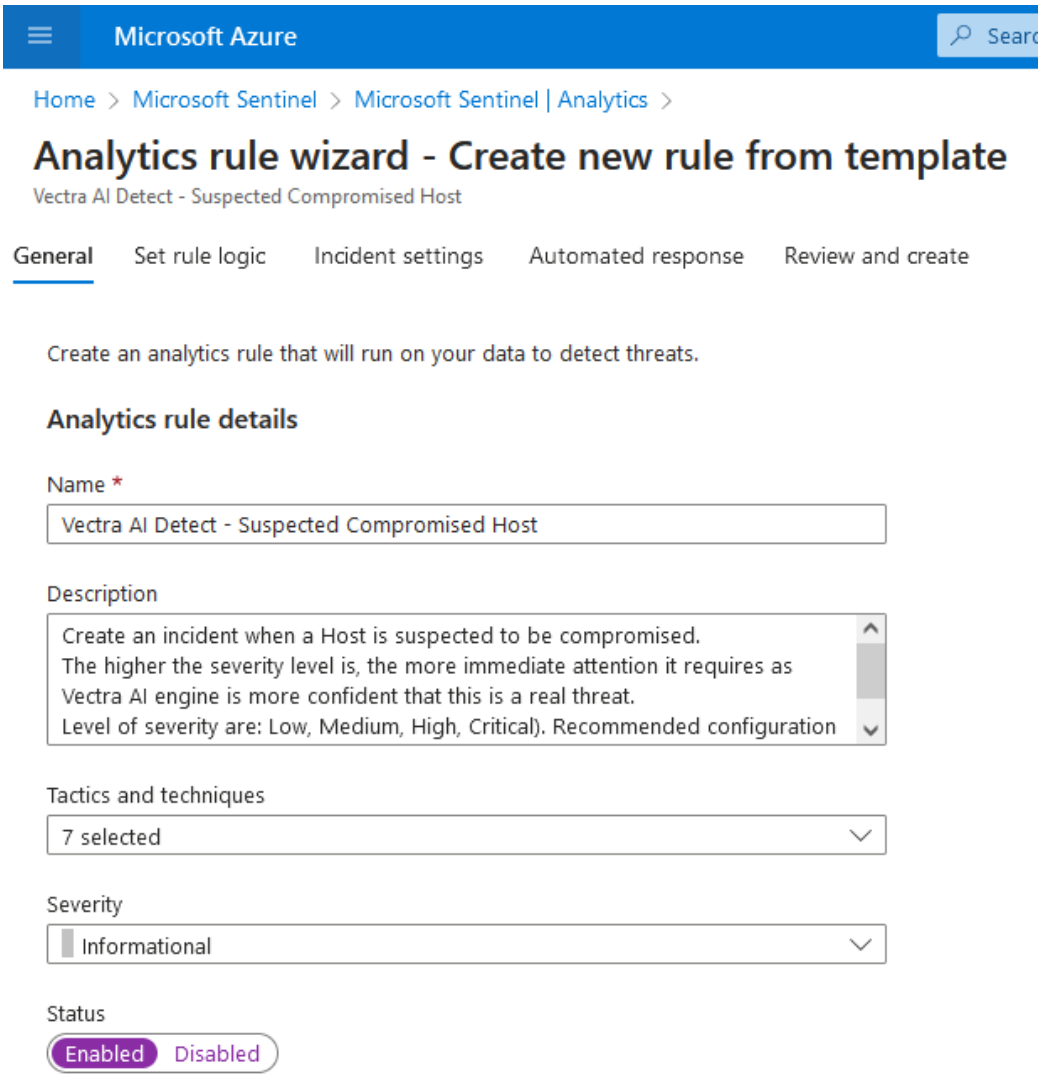

On the next page, the **rule logic** settings are all preconfigured:

- The query (use the view query results to validate if it returns any results).
- An incident is created when a host is in the High or Critical quadrant.
- The Alert enrichment:
  - Entity mapping.
  - Custom details.
  - Alert details.
- The scheduling of the query and the threshold.
- Event grouping.

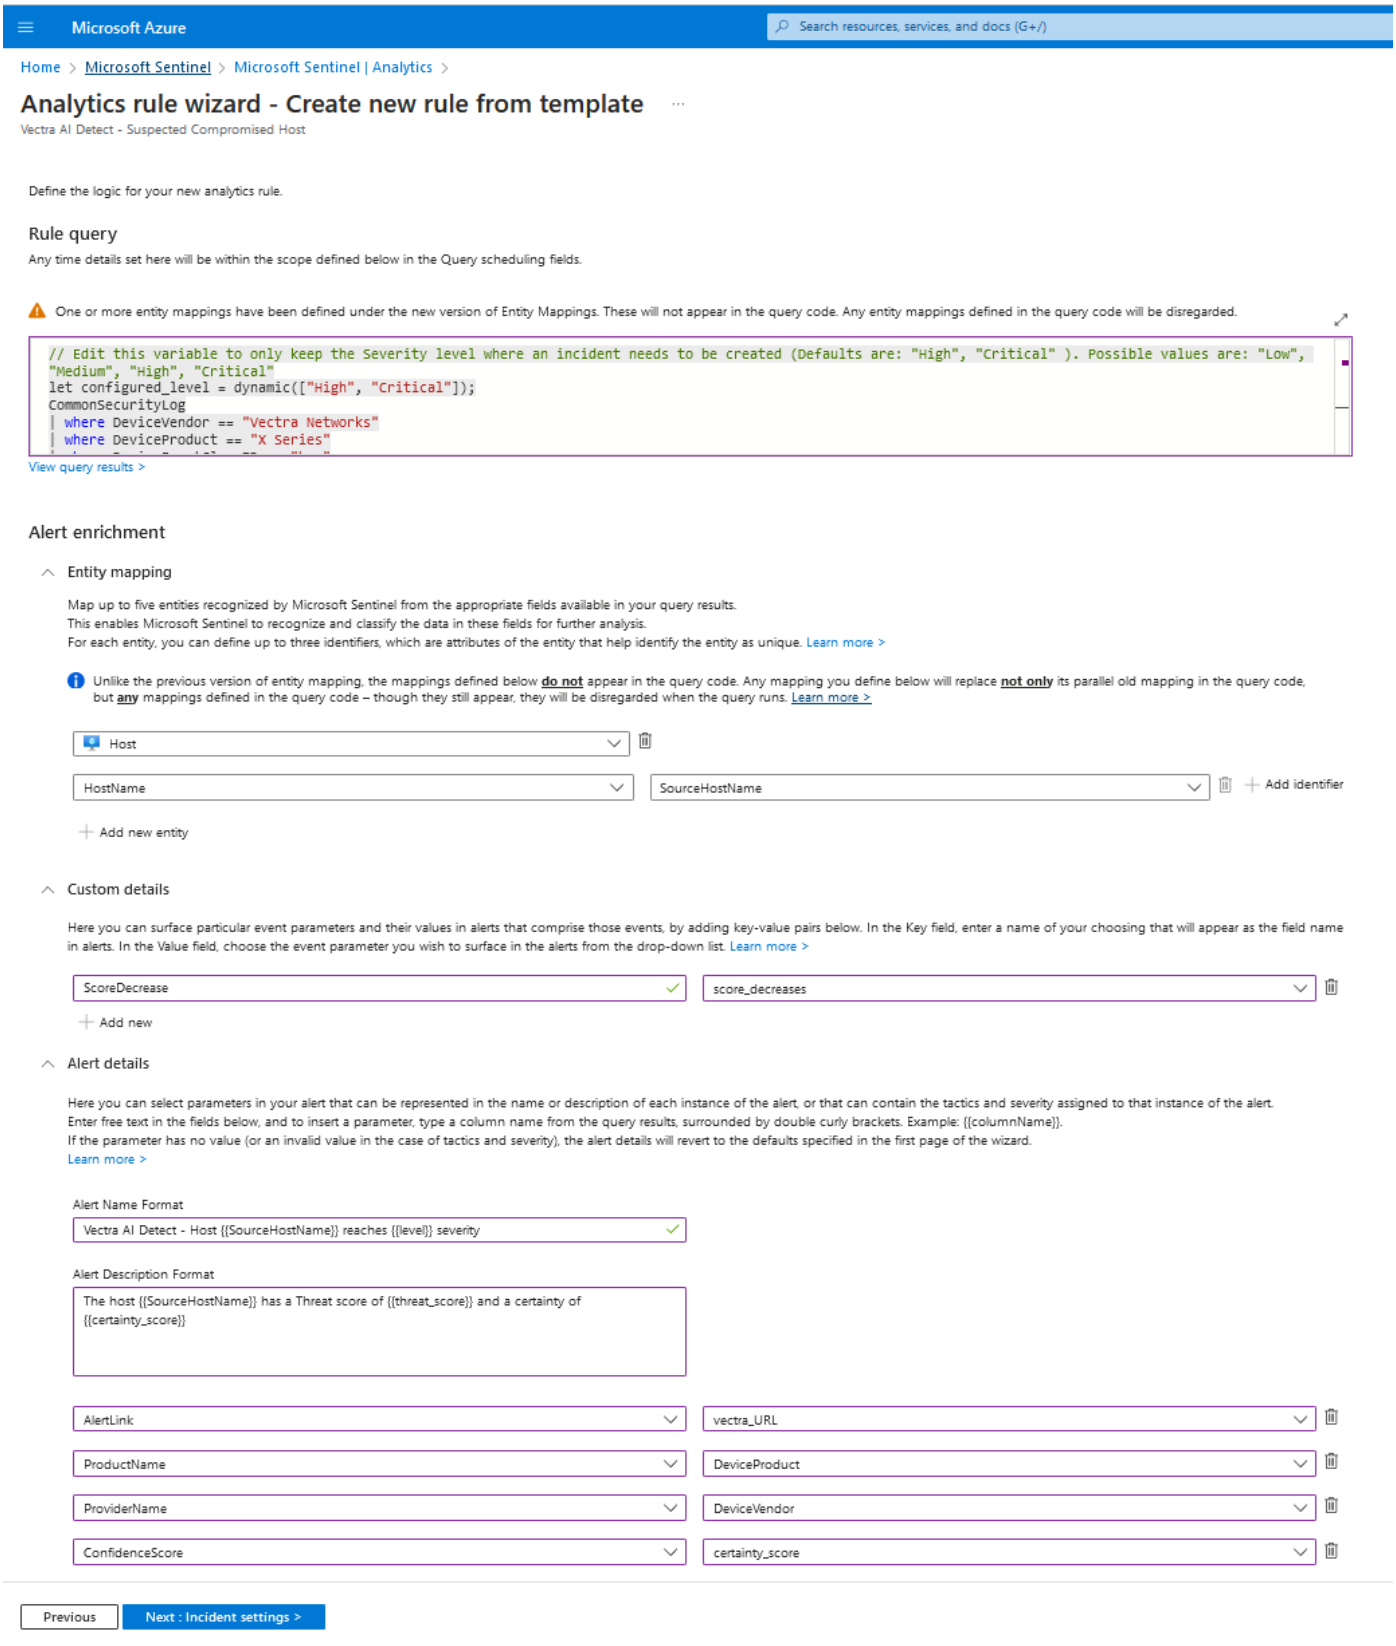

▼ Next, the Incident settings should all be preconfigured as well but please check them for incident creation as well as the grouping:

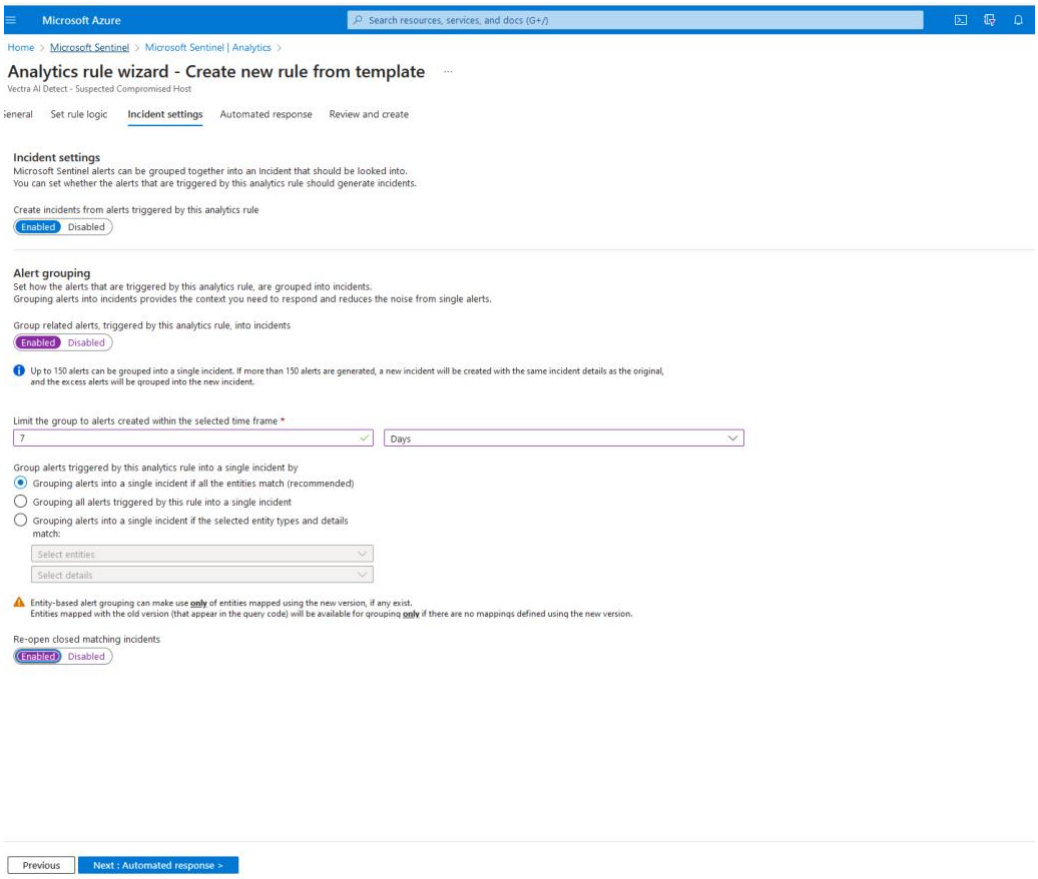

- The **automated response** depends on your environment and does not have anything configured by default.
- Click **Review** then **Create**.

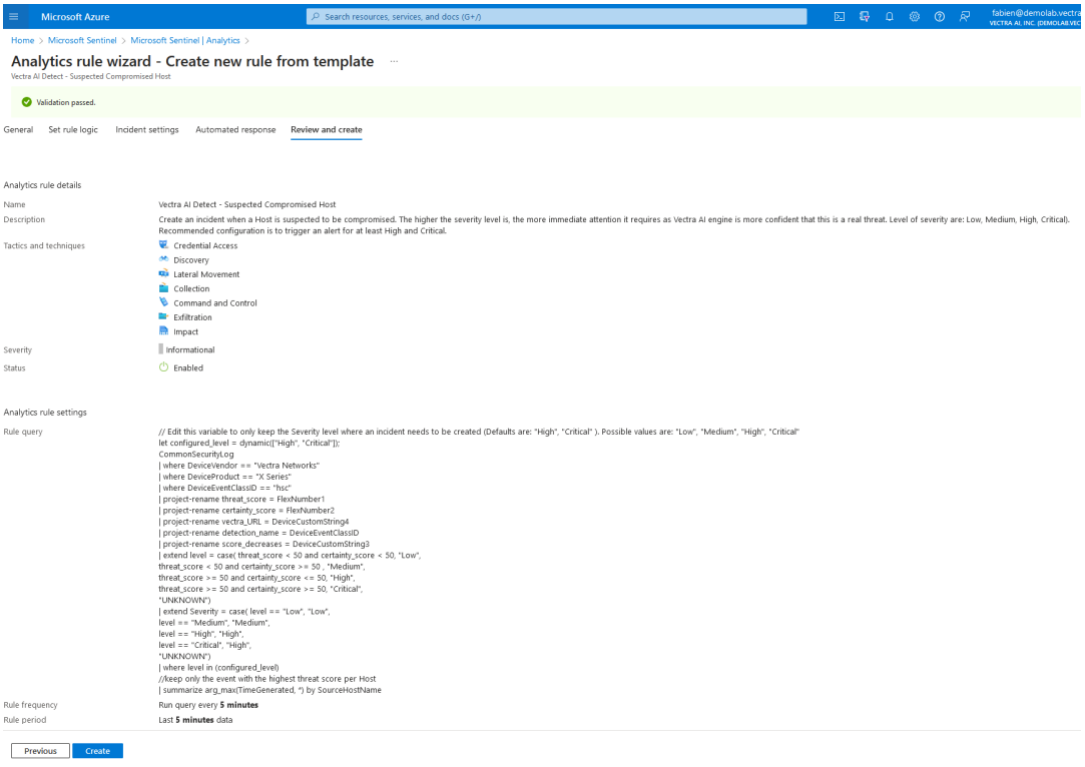

- Once deployed, you should see the analytic rule in the Active rules list:

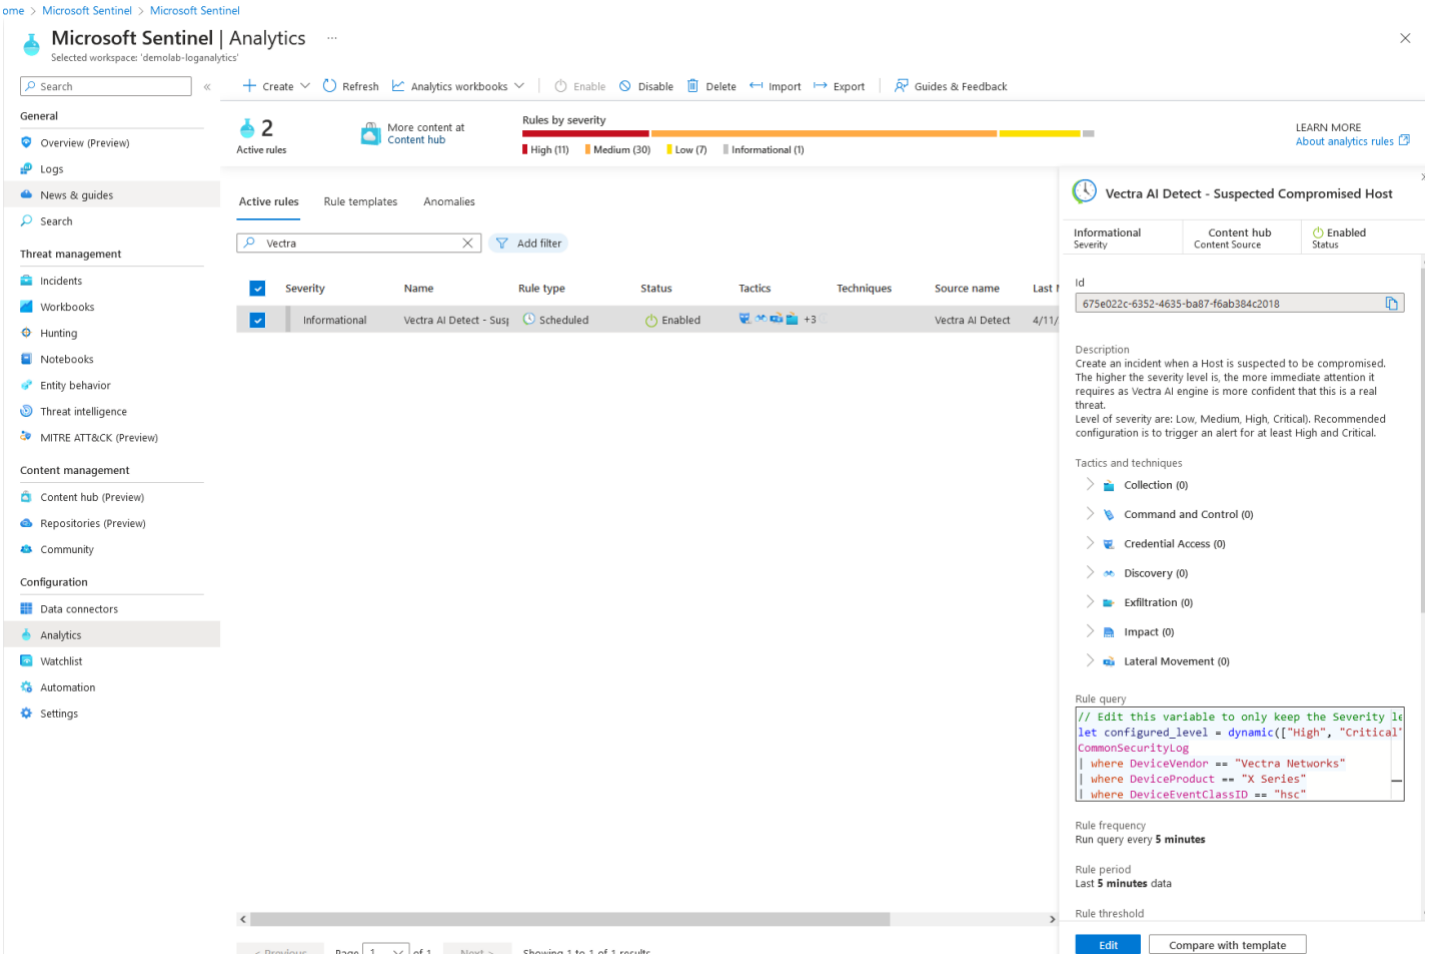

Alert creation for host detections

The intent is to create alerts for every host detection triggered by Vectra's Attack Signal Intelligence. Alerts go to the *SecurityAlert* table. To do so, we are going to modify the analytic template named: Vectra AI Detect - Vectra Host's Behaviors.

- Menu: *Configuration* > *Analytics* > *Rule templates*.
- Search for Vectra and select the template name **Vectra AI Detect - Vectra Host's Behaviors**.
- Click "Create rule"

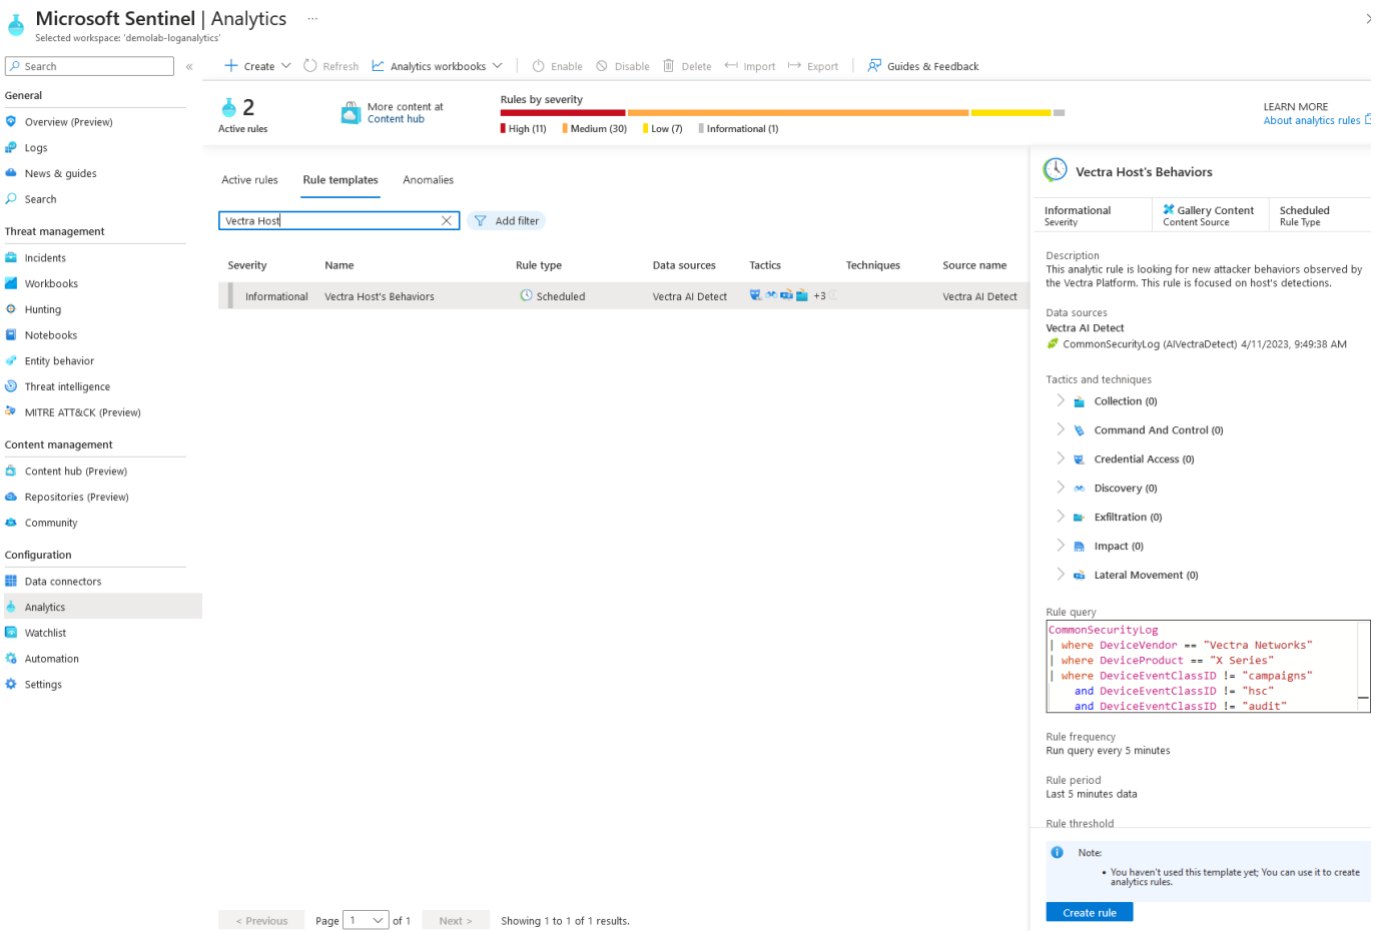

- Set the status to **Enabled** and make sure than Severity is set to **Informational** (it should be default).

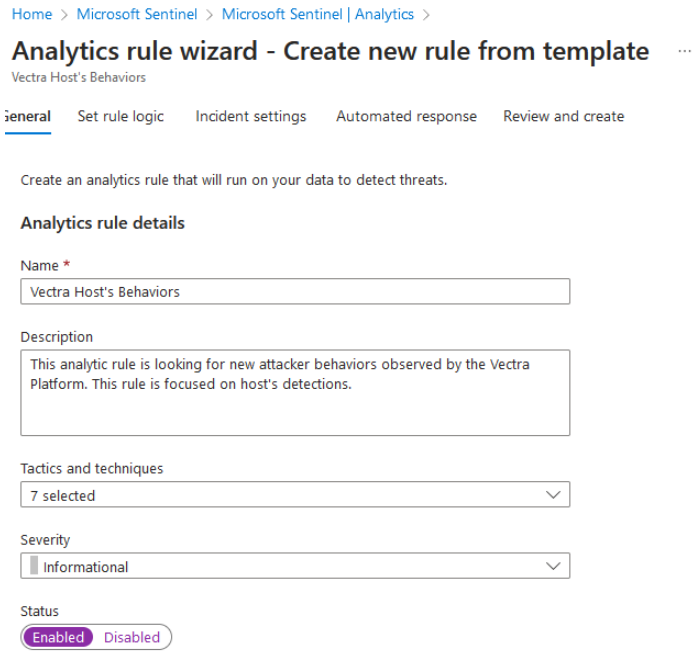

On the next page, the **rule logic** settings are all preconfigured:

- The query (use the view query results to validate if it returns any results)
  - An alert would be created for every detection that is not triaged.
- The Alert enrichment:
  - Entity mapping
  - Custom details
  - Alert details
- The scheduling of the query and the threshold
- Event grouping

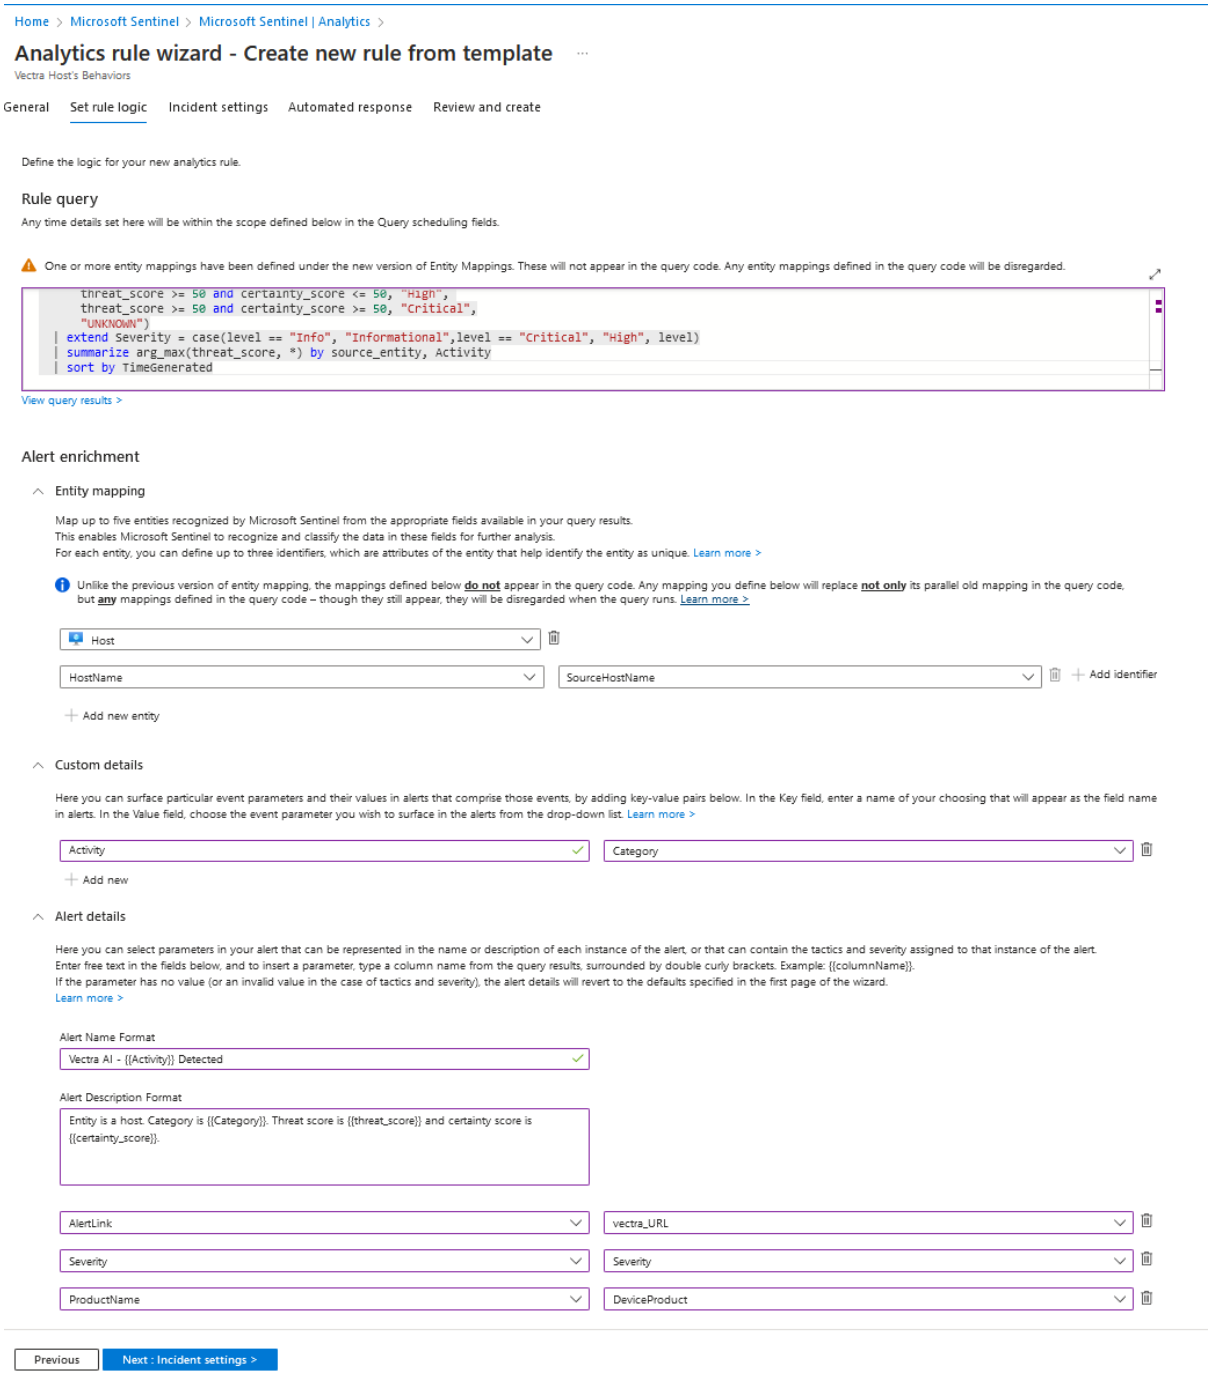

- Next, the **Incident settings** must **disabled**. We want to create only alerts.

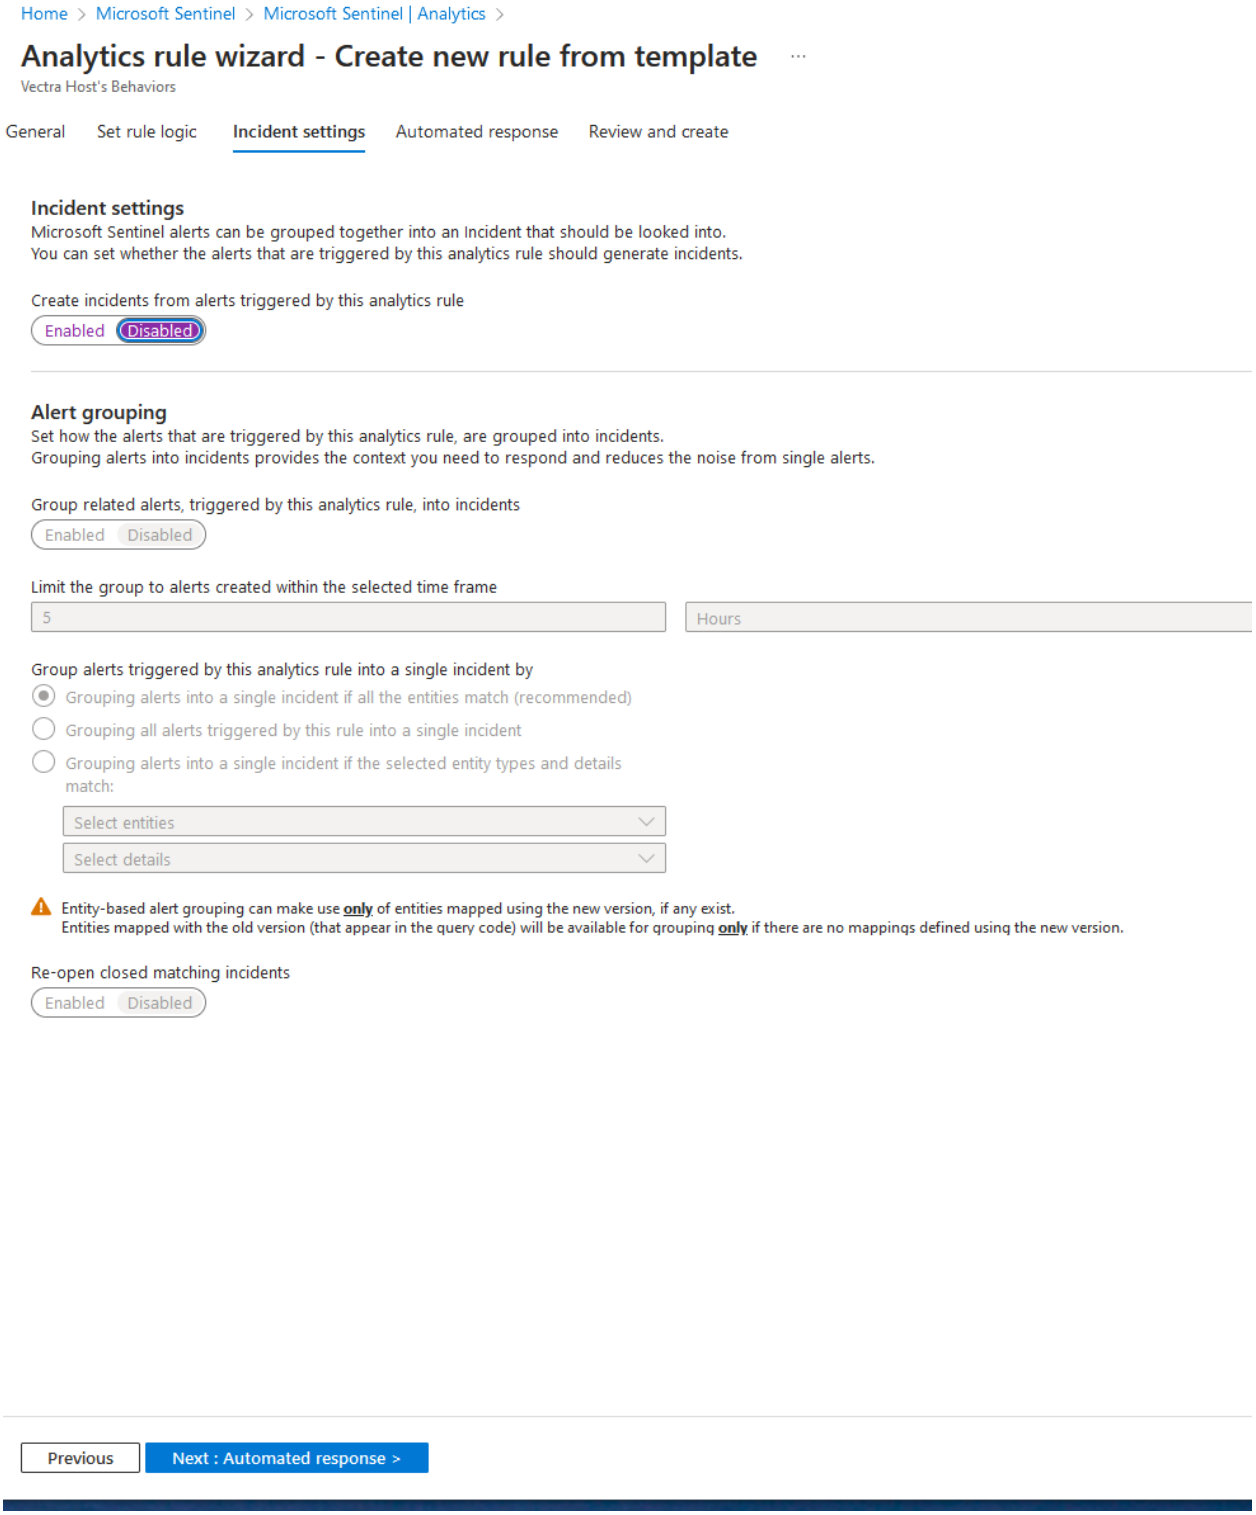

- The **automated response** depends on your environment and does not have anything configured by default.
- Click "**Review"** and then "**Create”**.

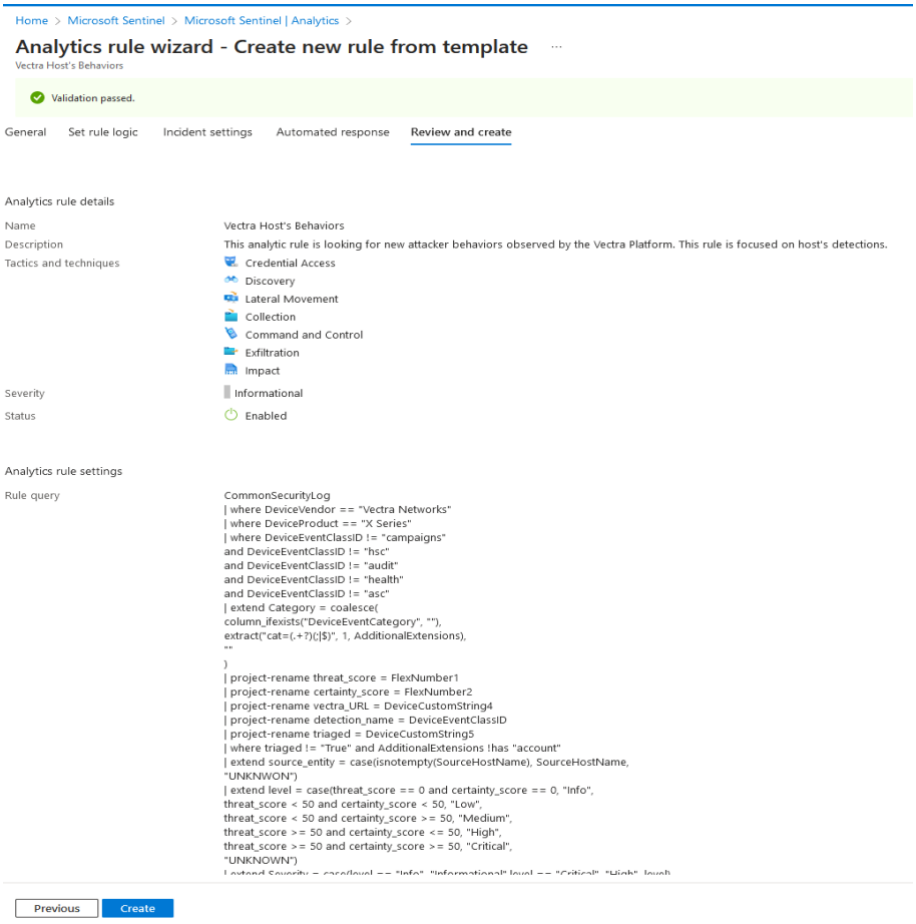

- Once deployed, you should see the analytic rule in the Active rules list:

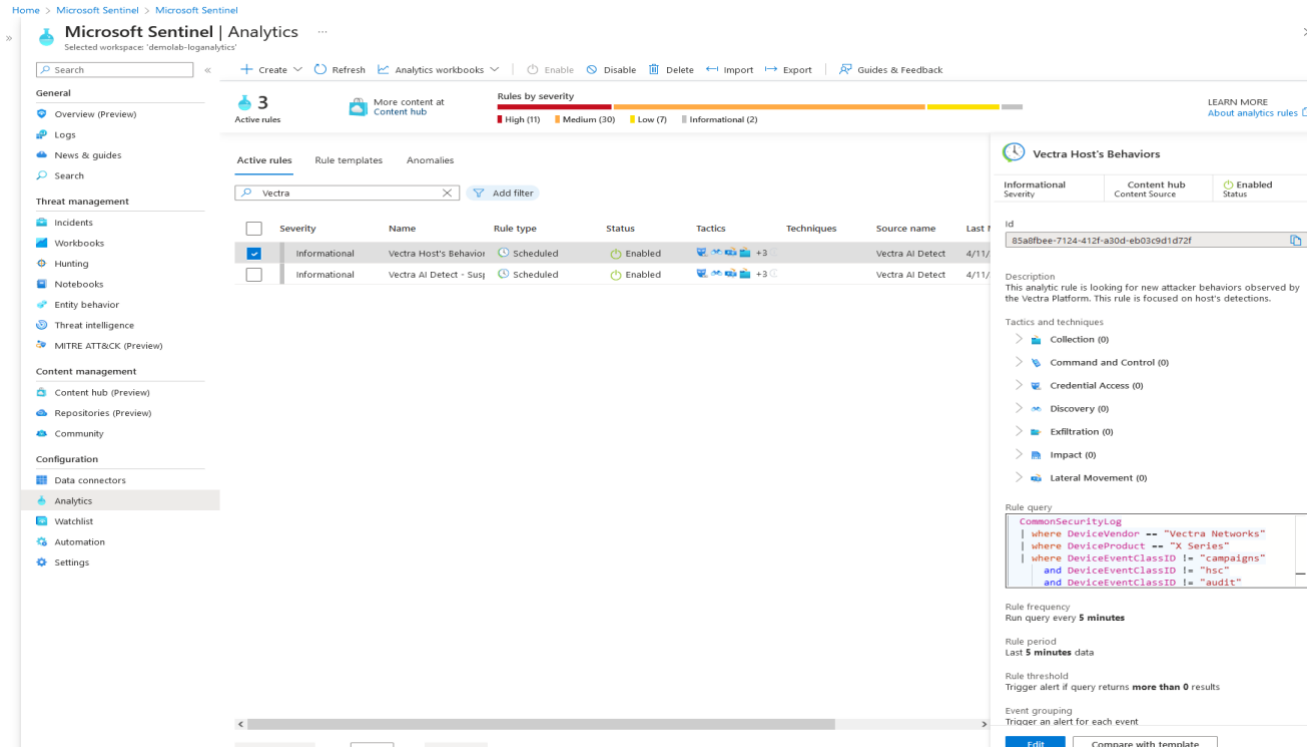

Incident creation for prioritized accounts

The same logic is being applied to accounts.

- Menu: *Configuration > Analytics > Rule templates*.
- Search for Vectra and select the template name **Vectra AI Detect - Suspected Compromised Account.**
- Click "Create rule".

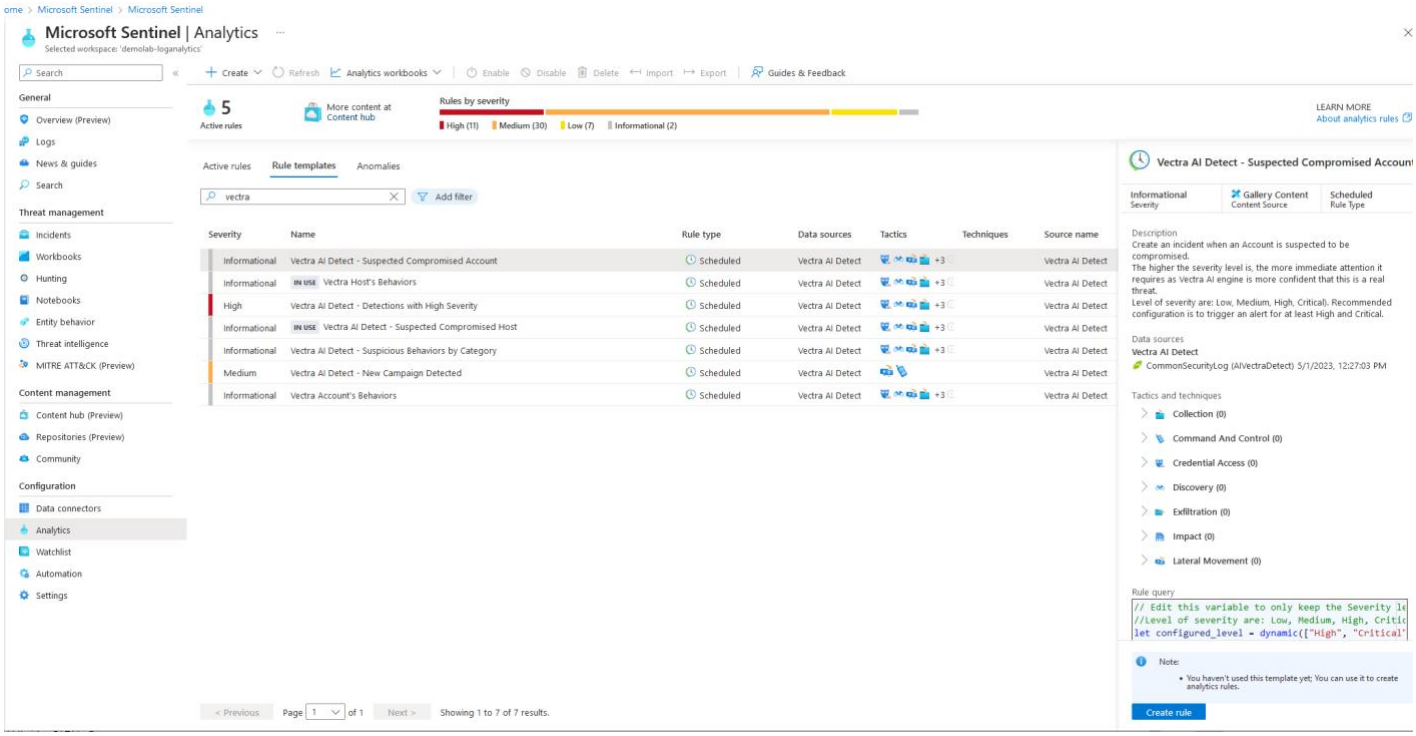

- Set the status to **Enabled** and make sure that Severity is set to **Informational** (it should be the default).

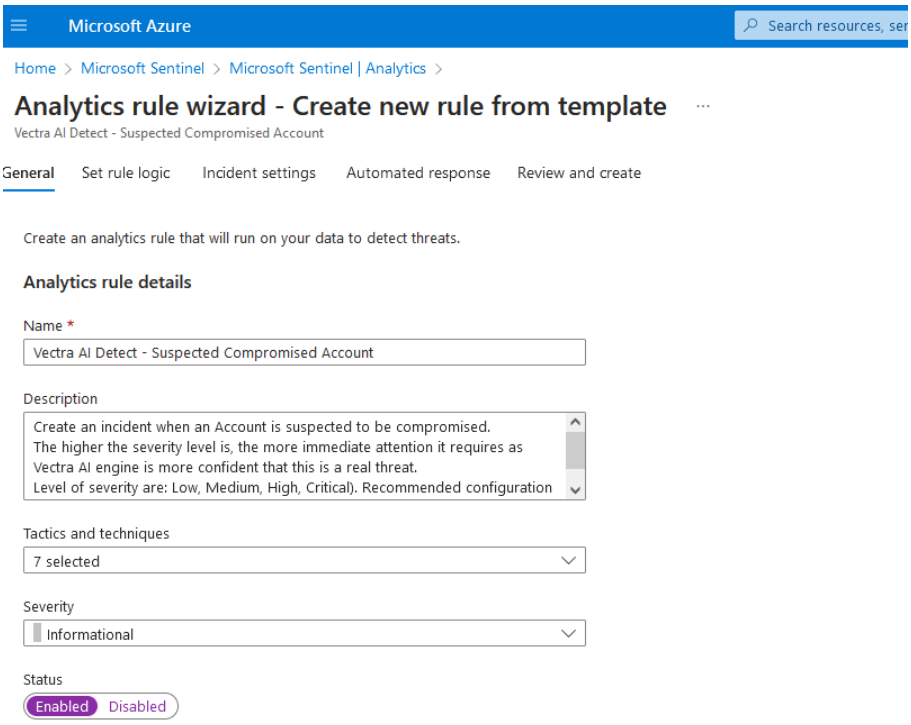

- In the next screen: *Set rule logic*, everything is preconfigured and should be left with default values (An incident would be created when an account is in the High or Critical quadrant).

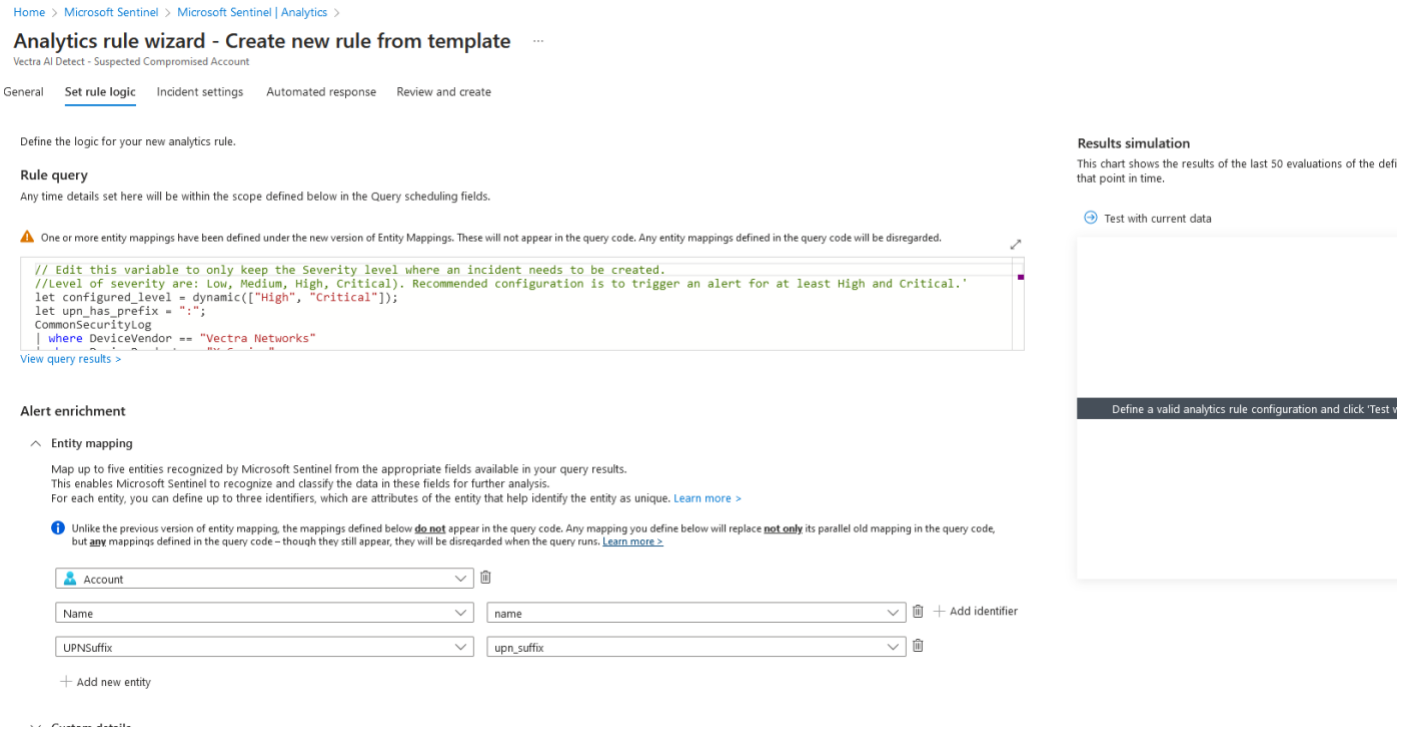

- In the Incident settings page, make sure that:
  - Incident settings is enabled.
  - Alert grouping is enabled with 7 days (Grouping alerts into a single incident if all the entities match).
  - Re-open closed matching incidents is enabled.

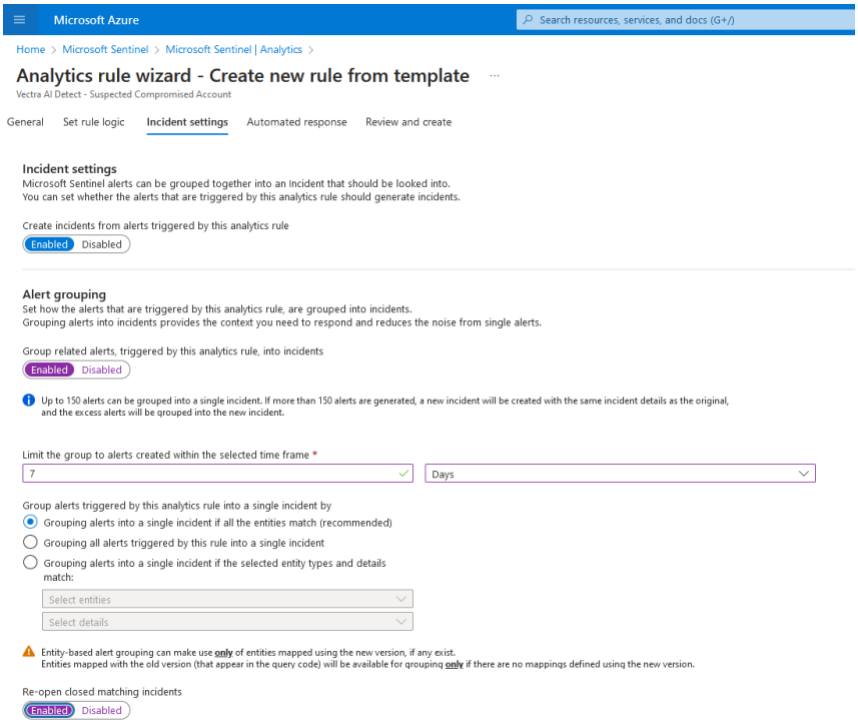

- Review and create. The rule is now listed as active:

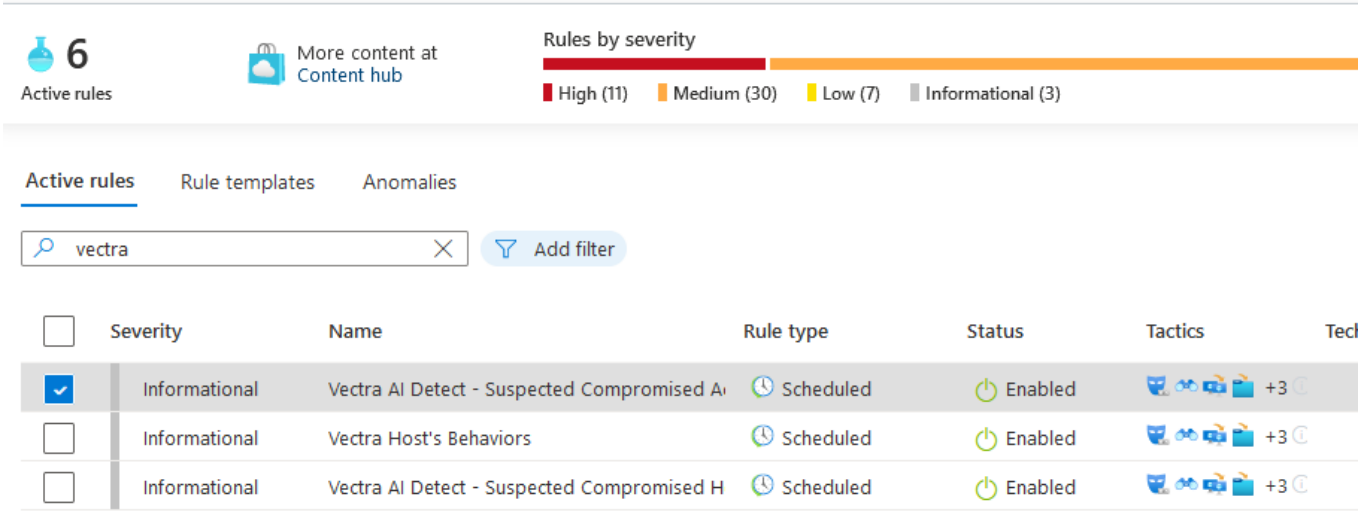

Alert creation for account's detections

The intent is to create alerts for every account's detection triggered by the Vectra Attack Signal Intelligence. Alerts go to the *SecurityAlert* table. To do so, we are going the analytic template named: Vectra AI Detect - Vectra Account's Behaviors.

- Menu: *Configuration* > *Analytics* > *Rule templates*.
- Search for Vectra and select the template name **Vectra AI Detect - Vectra Account's Behaviors**.
- Click "create rule"

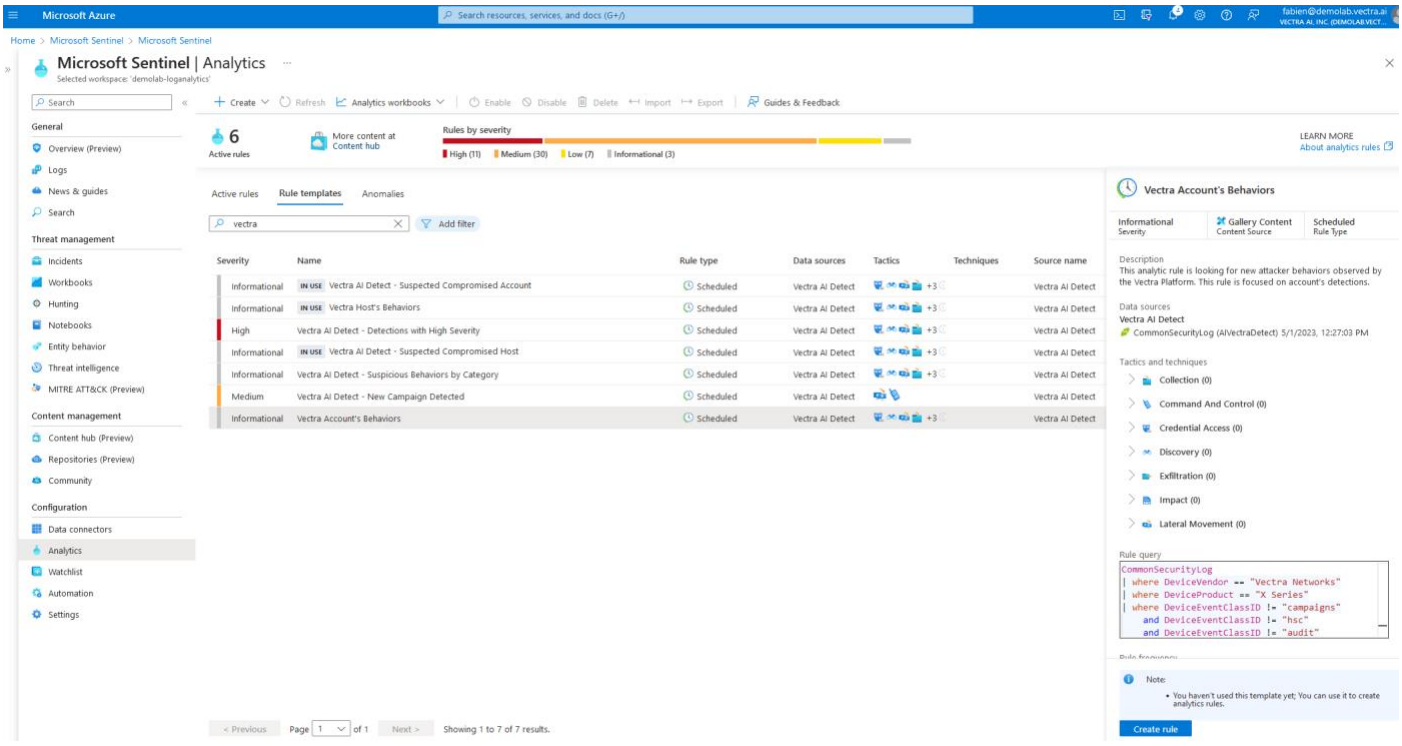

- Set the status to **Enabled** and make sure than Severity is set to **Informational** (it should be default).

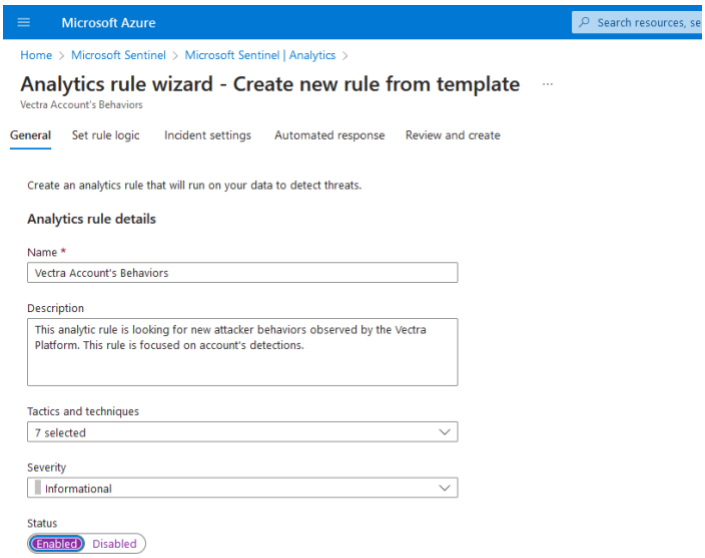

- In the next screen: *Set rule logic*, everything is preconfigured and should be left with default values (query, entity mapping, alert details, frequency, etc.)

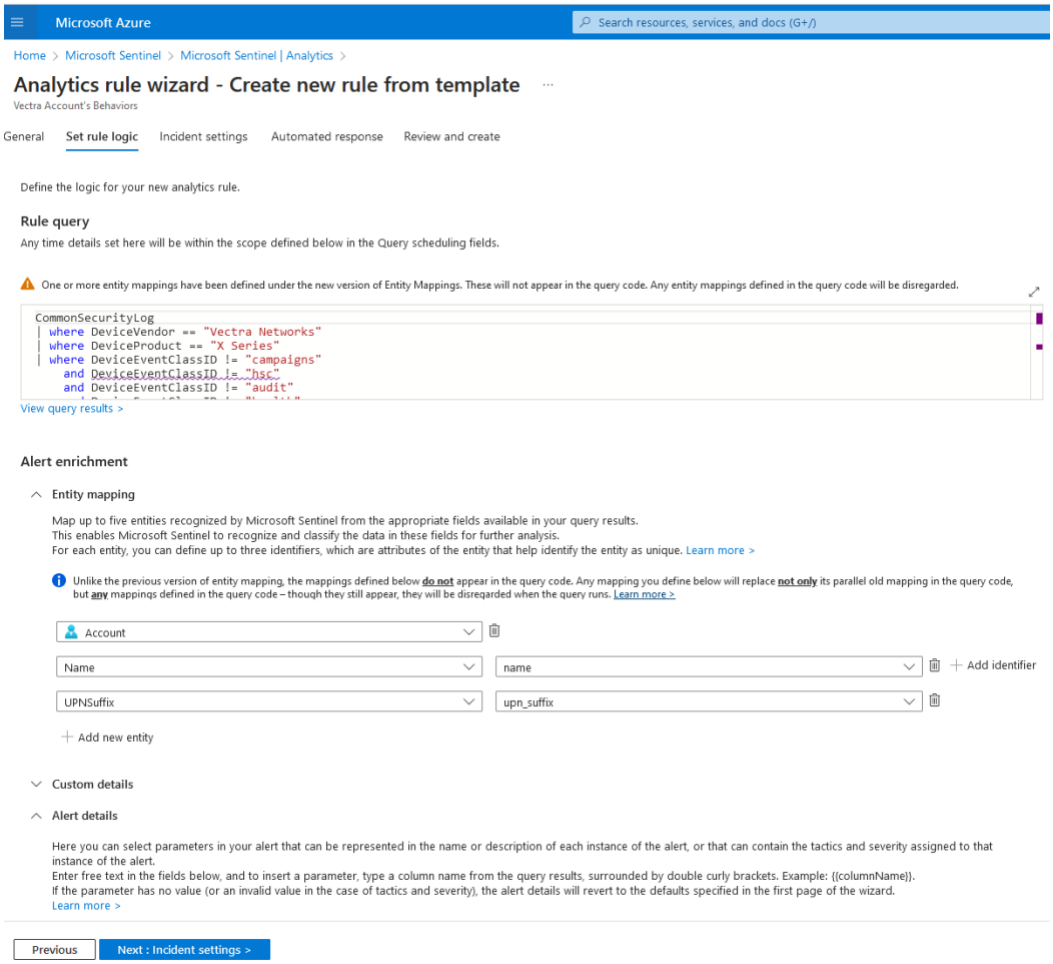

- In the next screen, *Incident settings* must be set to **Disabled.**

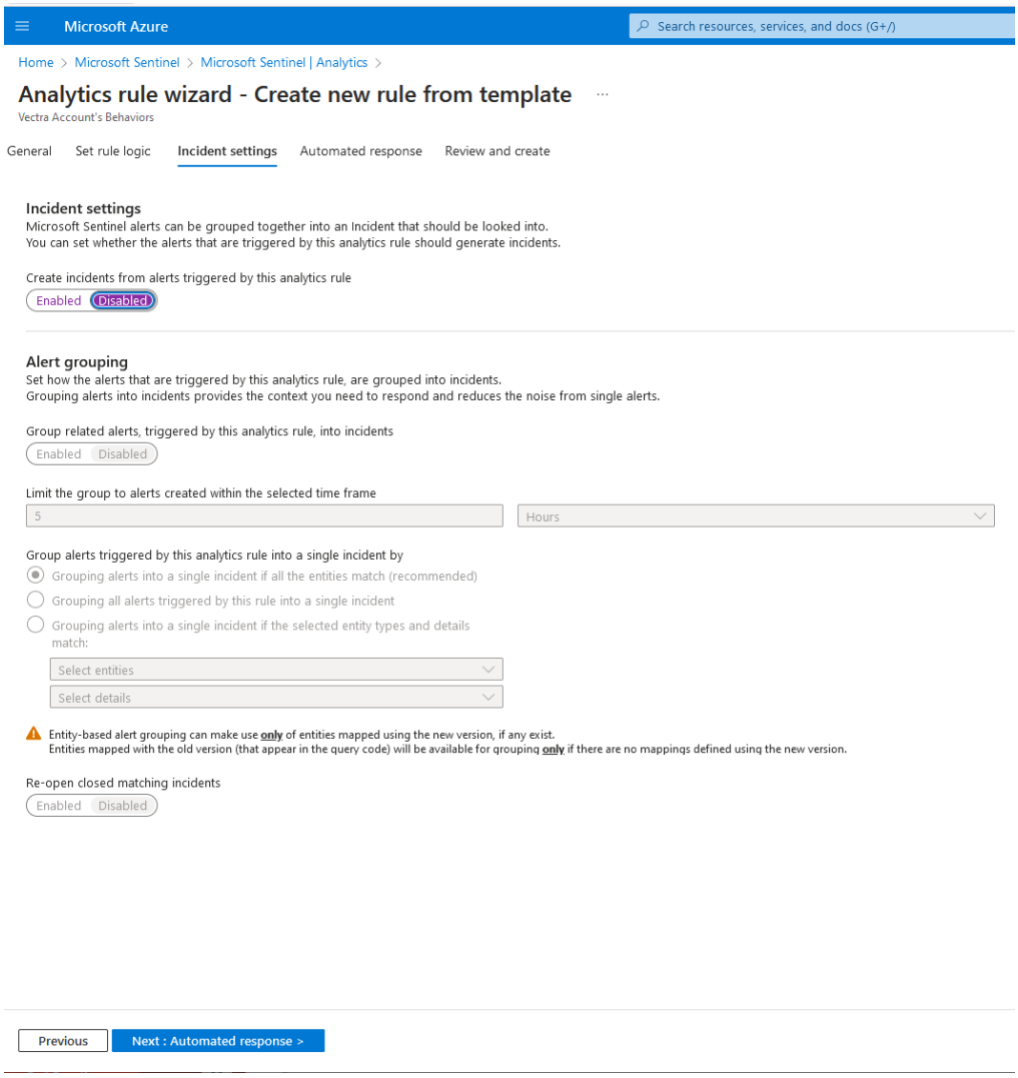

- Click **"Review"** and then **"Create"**.
- The rule is now listed as active:

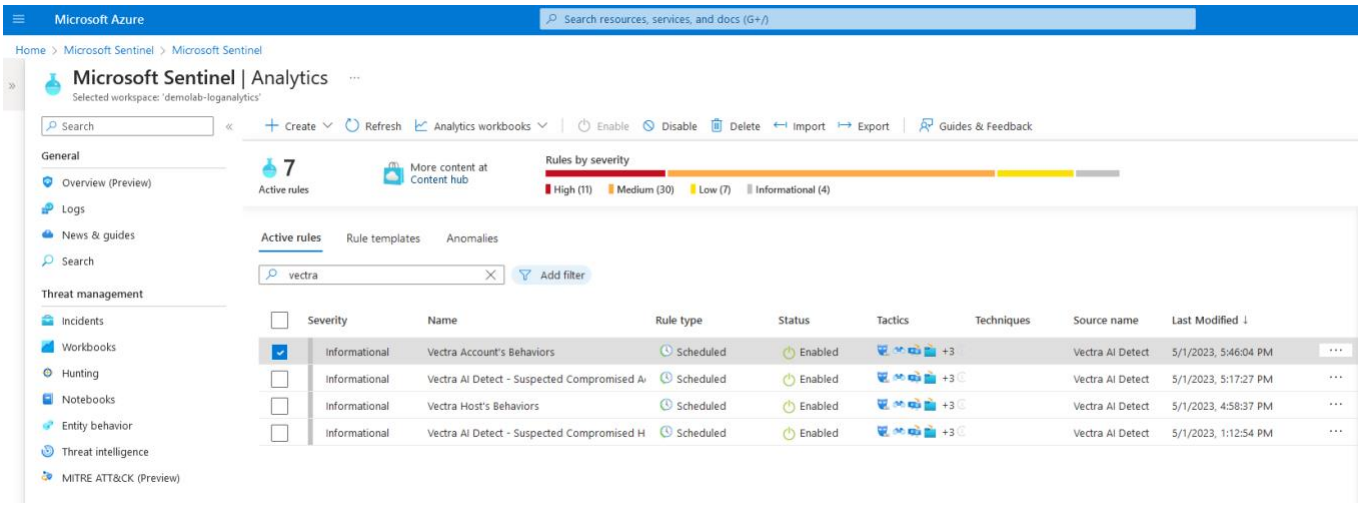

The SOC analyst journey

This section describes an example of the workflow for an analyst. It starts in the Incidents page where, as a SOC analyst, new and existing incidents would be picked up.

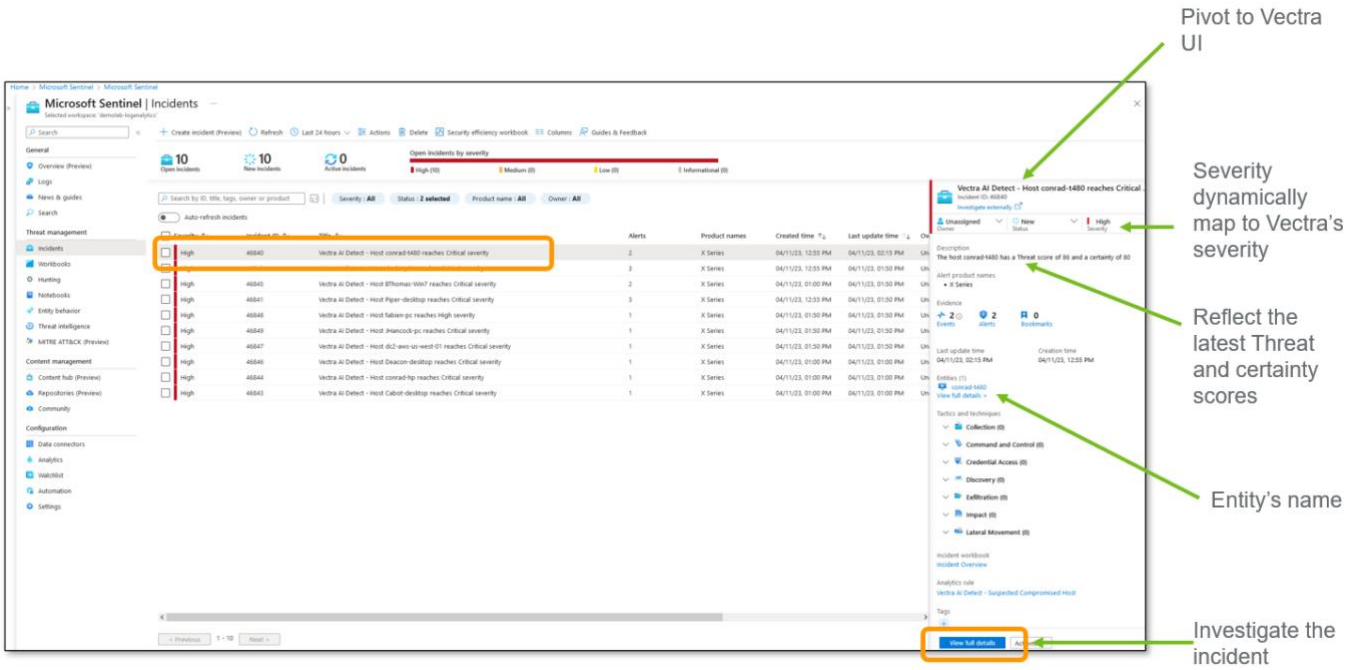

The full details provide additional detail such as the timeline where you can see incidents associated to this entity. The grouping configuration of events (7 day windows) allows a view of how the score has been evolving over time which will directly impact the severity associated with that incident.

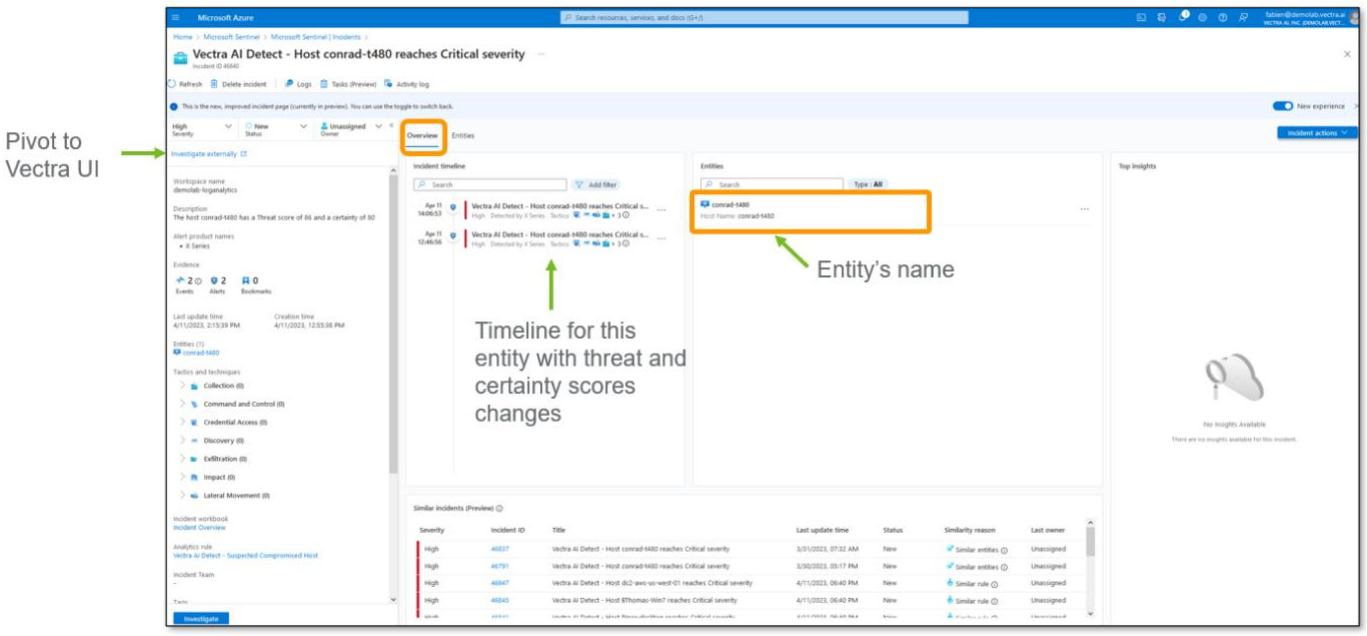

Now, as the next step, an analyst would like to see all the different attacker's behaviors that have been observed for that entity. This can be done right in Sentinel by clicking on the entity's name, and then the timeline tab:

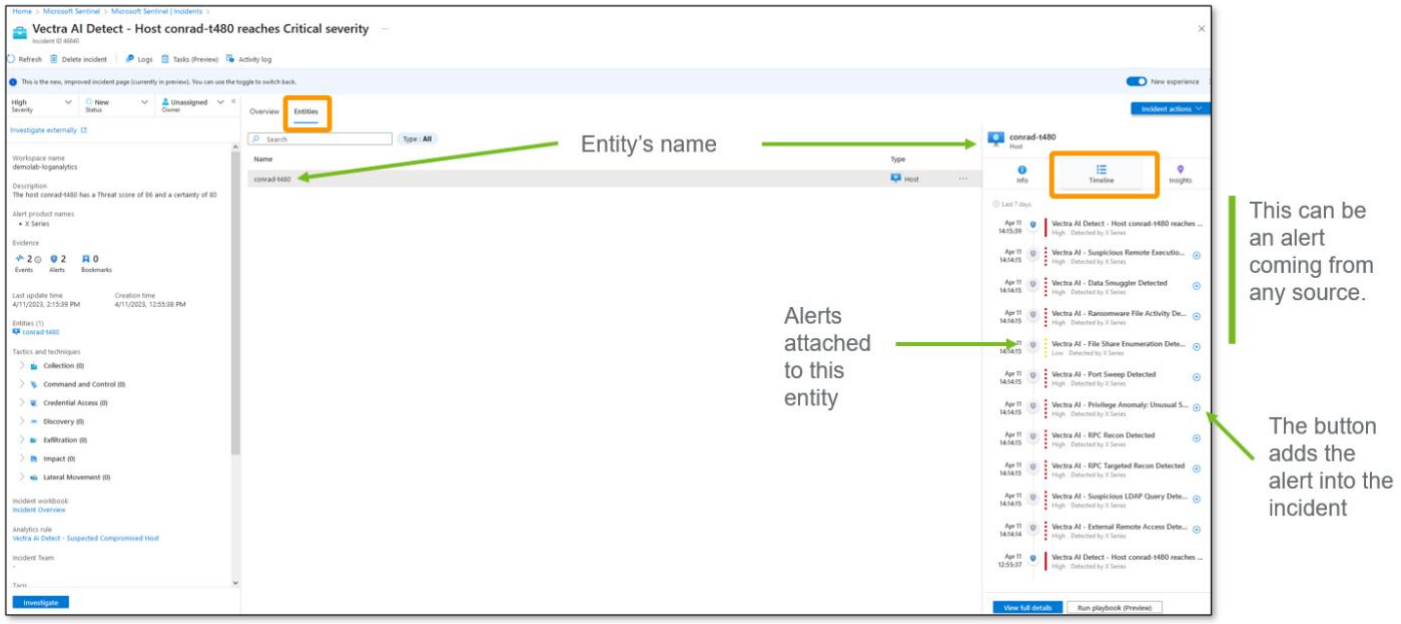

Alerts in the entity's timeline view can come from any source (EDR, FW, etc.). This is why the entity mapping configuration is key to successfully correlate them. Alerts of interest from this entity can be added to the incident by clicking the + button, then it shows up in the incident timeline view. Navigate back to the overview tab to see the timeline view. You can click on any new alert that you added in your incident to see additional details for each of them that allow you to continue your investigation:

- Detection name
- Category
- Severity
- Threat and Certainty Scores
- Deep link to Vectra UI

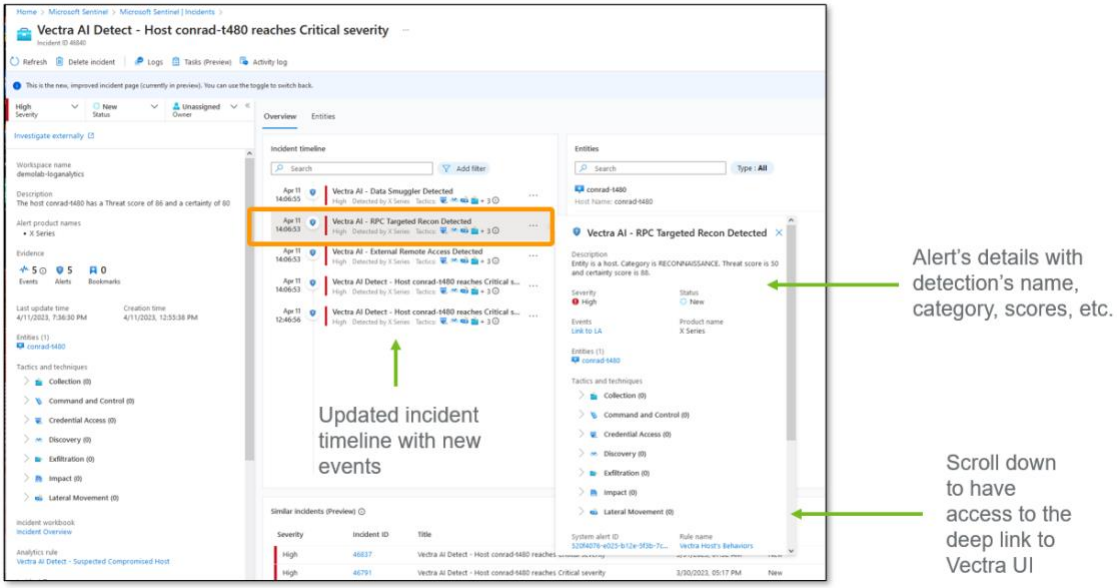

Best practices

Keep the content up to date!

As we continuously improve our content for Microsoft Sentinel, please make sure that you are using the latest version. New versions can include enhancements for new features as well as bug fixes. When updates are available, it can be seen in content hub:

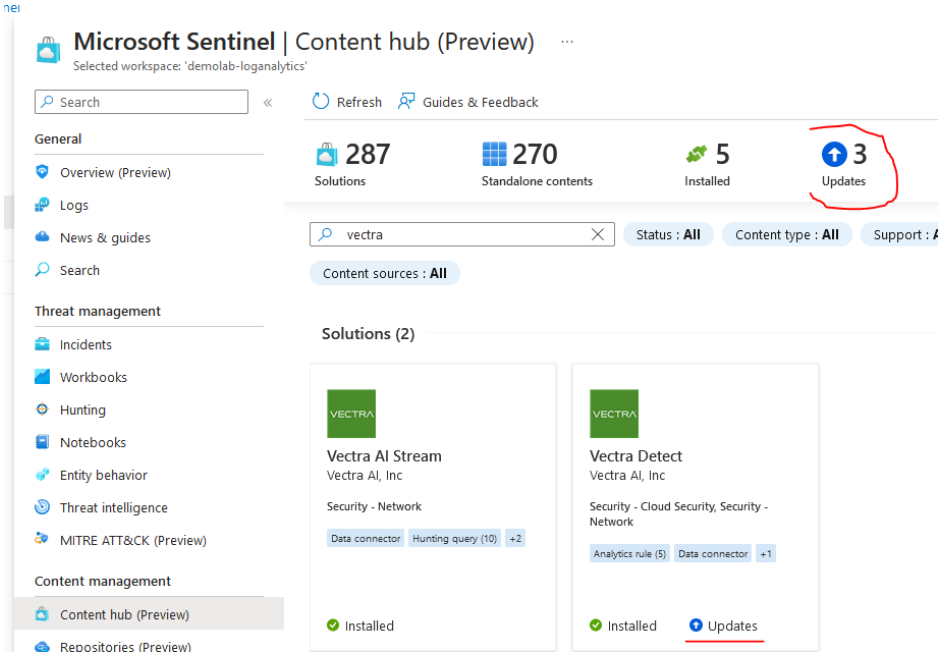

- Once you click on the tile, on the bottom right, click on **Actions > Update** to deploy latest version of the content and follow the wizard to deploy it in the right workspace.

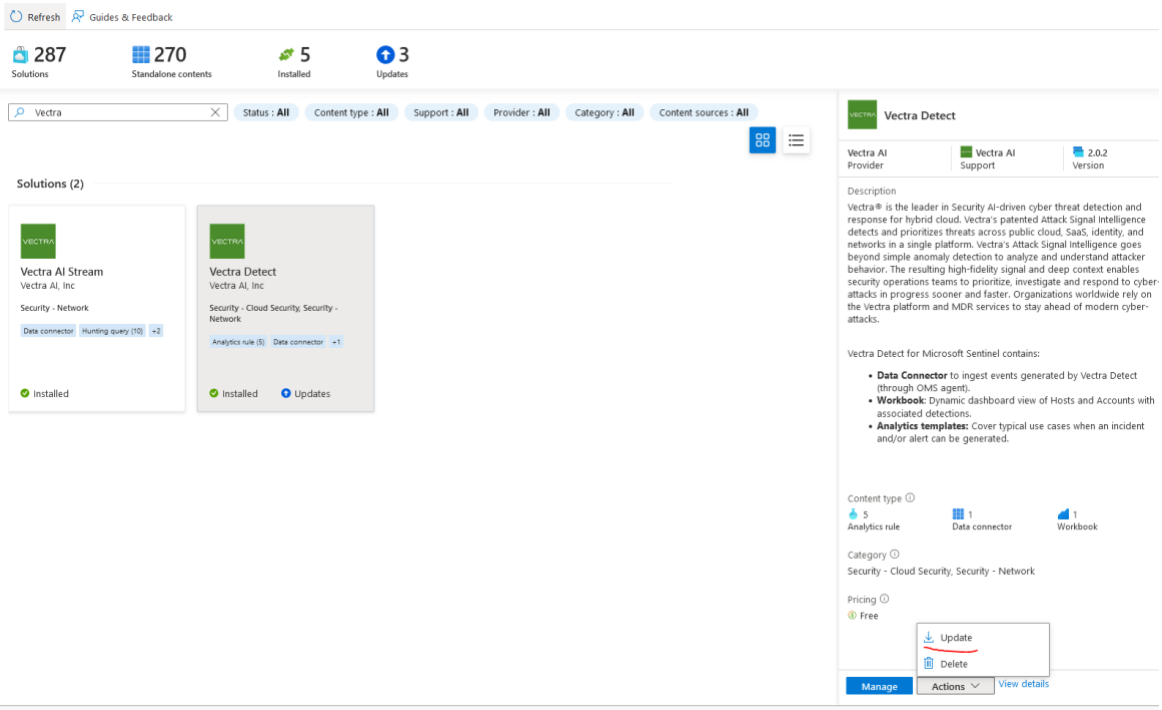

- Once it has been deployed, you would see that an update is available for the templates you already deployed. In the scenario of analytic templates, the tag "UPDATE AVAILABLE" is added to rules that have an update.
- To proceed, click the "Review and update" button at the bottom right of the page.

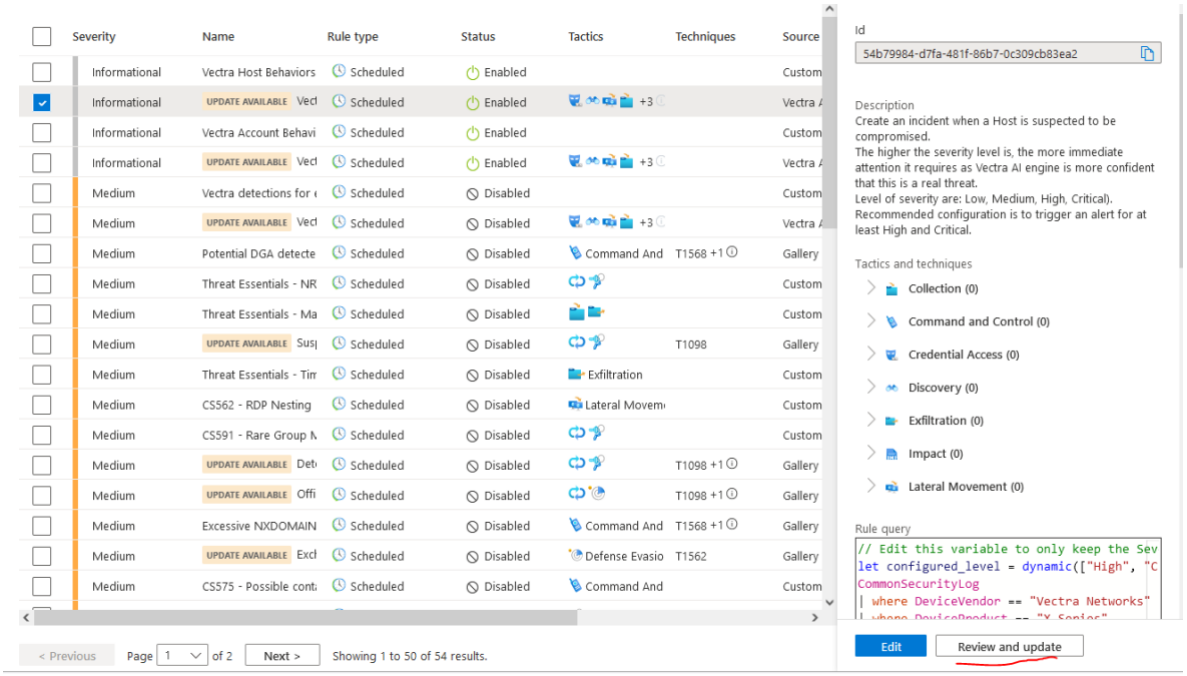

Create an alert when no event is received

As described in this KB, we rely on a Microsoft agent to send Vectra's events to Sentinel. This agent is receiving events via syslog and forwards them to the correct Log Analytics Workspace within Azure. To detect any issues with the middleware component, we recommend using Azure Monitor to setup an alert when no events are received for 24 hours. It is very unlikely that the Vectra platform does not generate any events for so long. Depending of your preferences, you can go as low as 1 hour.

- In Azure Monitor, create a new alert similar to the below:

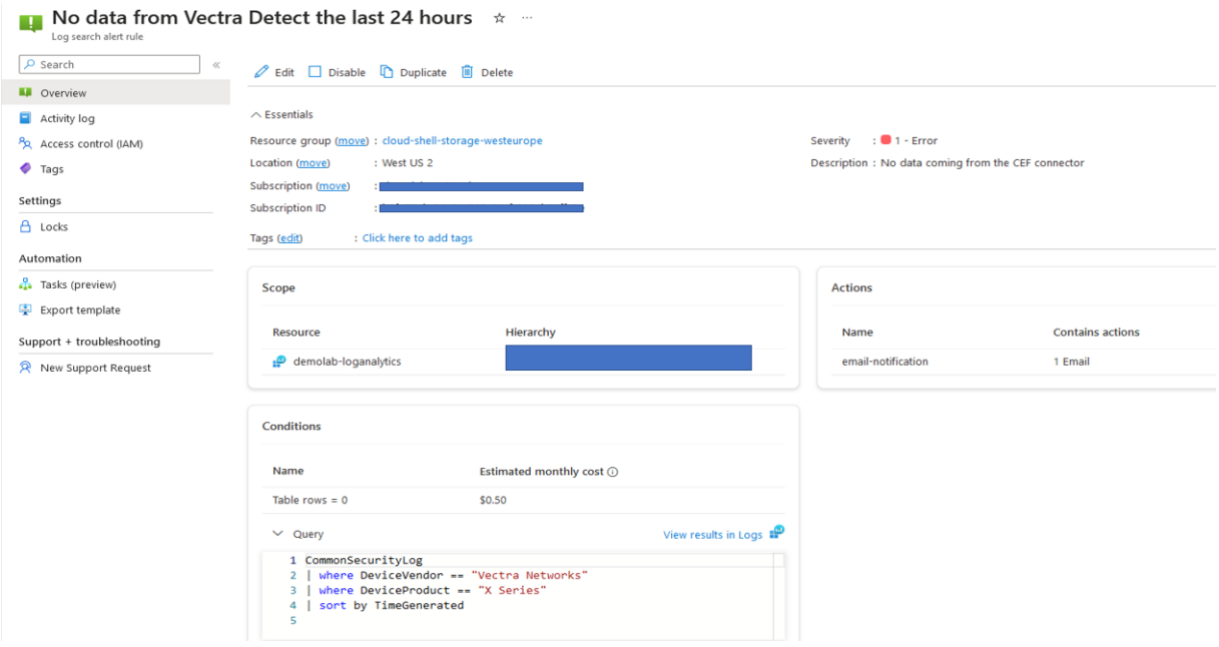

KQL query used:

```
CommonSecurityLog
| where DeviceVendor == "Vectra Networks"
| where DeviceProduct == "X Series"
| sort by TimeGenerated
```

In this case, we send out an email notification when the alert is triggered. You can change the action to whatever fits best in your environment.

Worldwide Support Contact Information

- Support portal: https://support.vectra.ai
- Email: support@vectra.ai (preferred contact method)
- Additional information: https://www.vectra.ai/support