Attack Path Analysis Generative AI Best Practices Guide

Last Revised: May 11, 2026

Table of Contents

Attack Path Analysis: Generative AI

This document emphasizes the different use cases using generative AI in Attack Path Analysis. It highlights some of the limitations and provides tips on how to get the most out if it for your business use.

In Attack Path Analysis, you can find AI-generated content when using the following features:

Attack Path Summarization

This feature allows users to view executive summary information for a specific attack path generated using the Attack Path Query Builder. By transforming the visual representation of an attack path into a human language summary, Attack Path Analysis provides an option for less technical users to ingest the data in a different way. Additionally, users can easily share this summary with colleagues to easily highlight attack paths that need their attention.

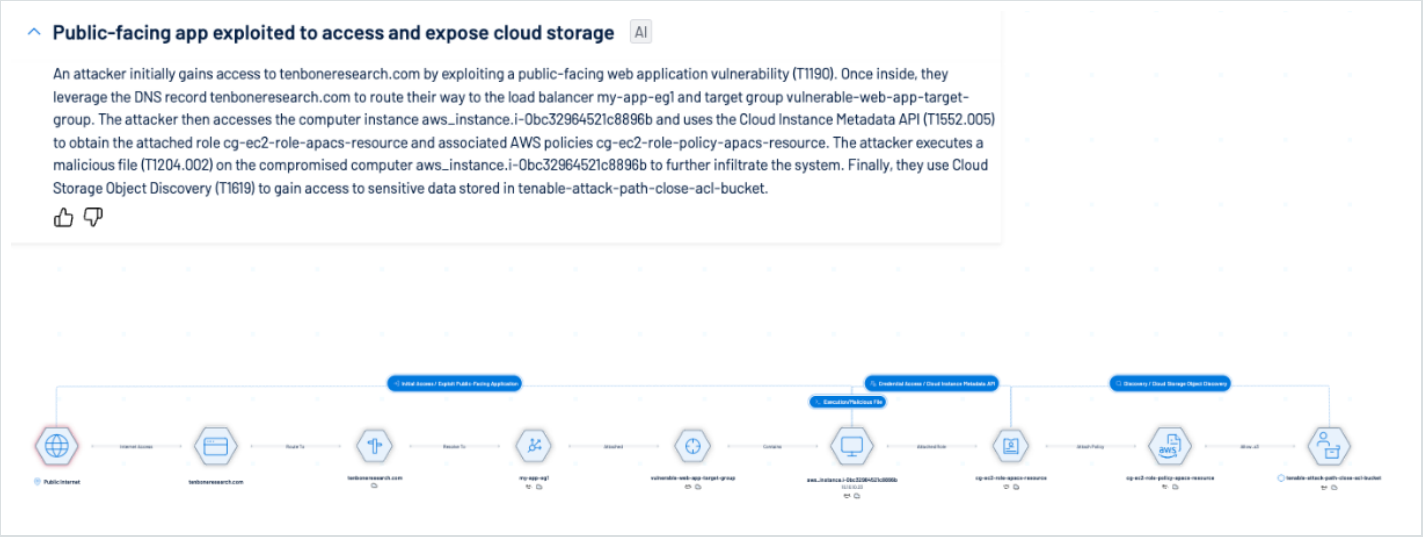

Tip: Seeing a No insights available message? Click the ↻ button to retry.

AI Explainability

This feature allows users to ask questions about a specific attack path generated using the Attack Path Query Builder or the Asset Query Builder. Using this search, users can better understand the attack path and its associated risk. Here, you can also gain additional insight into the assets affected by the attack path.

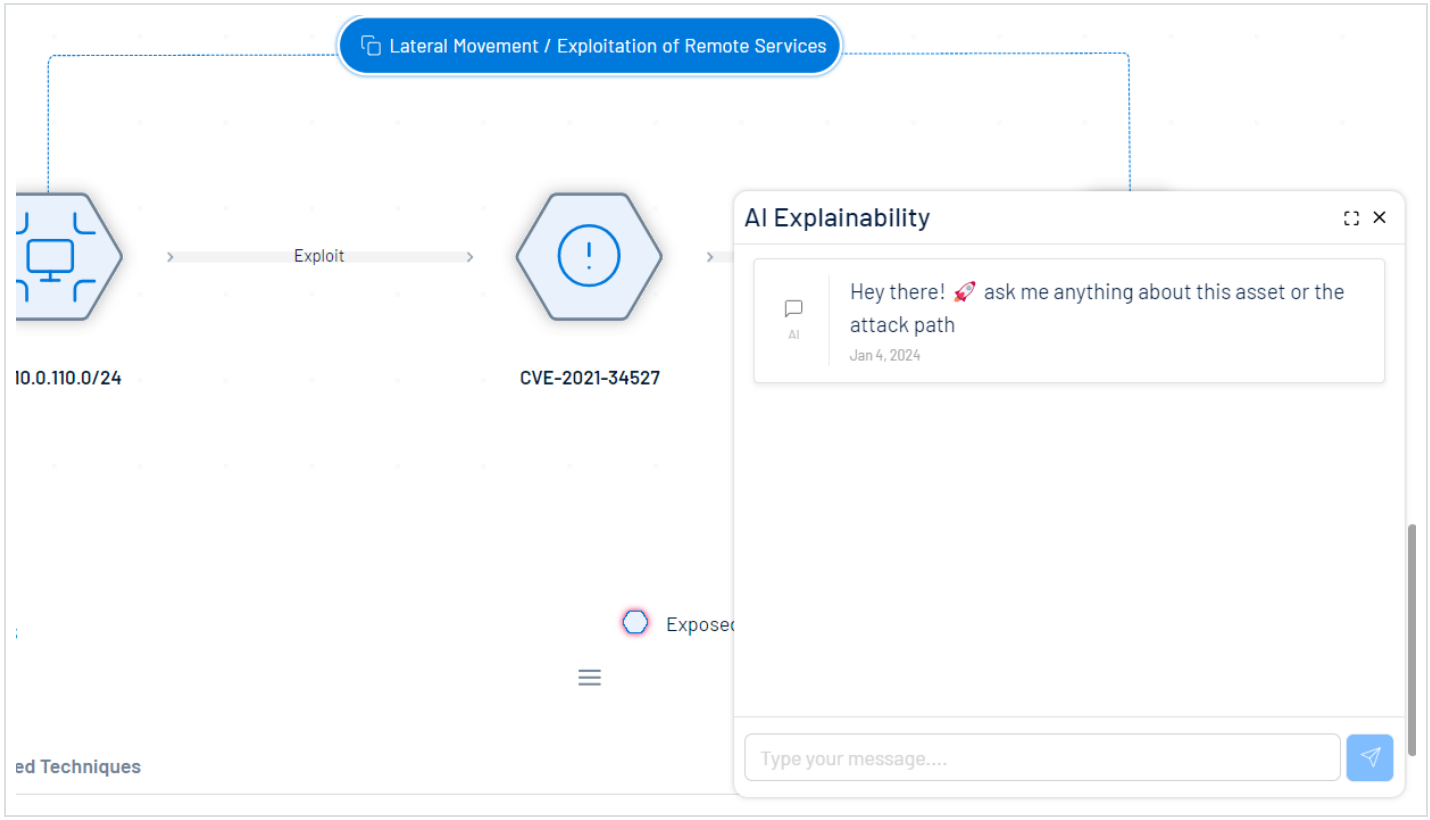

Users can access this chat window by right-clicking a node within the graph view and selecting the Ask AI About This Node menu option.

Mitigation Guidelines

This feature allows users to obtain practical guidelines for how to mitigate their Attack Path Findings. This feature empowers users to self-service by acting on the recommendations and ultimately eliminating their most critical attack paths.

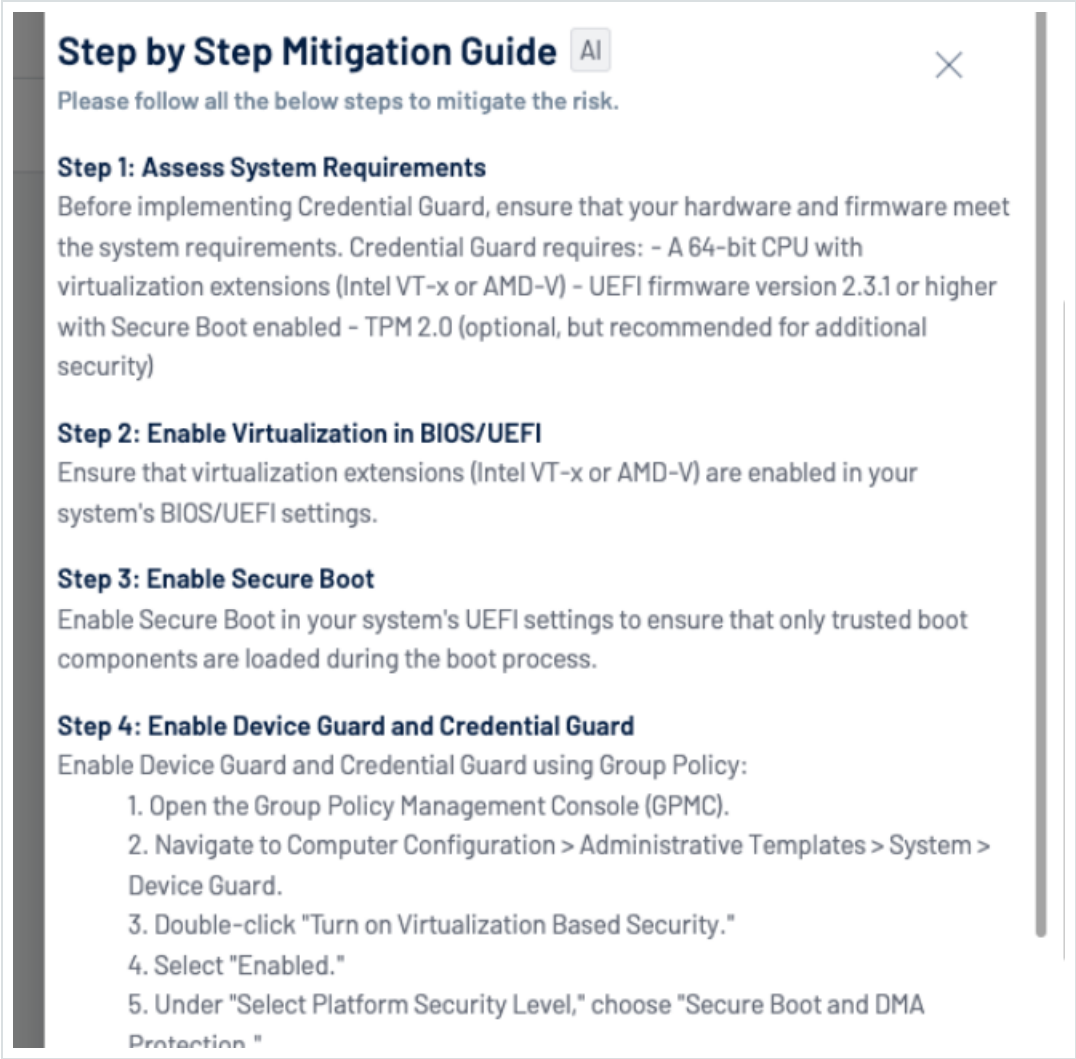

Sample Questions

The following are some sample questions to help you get started with generative AI within Attack Path Analysis. You can ask these questions to any AI-powered search within the Attack Path Analysis interface (for example, AI Explainability.)

> Tip: You do not need to define the specific asset within your question, because the Attack Path Analysis AI is only connected to the asset node and its related attack path.

- What can you tell me about this asset?
- Does this asset have any exploitable vulnerabilities?
- Is this asset vulnerable to CVE-XXX-XXX? How can I mitigate it?
- How many users have access to it?
- Who are the local administrator users?
- What questions can I ask about this path?
- Can you summarize XXXX (for example, golden ticket, dcsync, lsass memory, etc.) attack?
- Does this device run software with a vulnerability?
- What security controls are installed on this asset?

Limitations

Context

It is important to remember that questions can be asked only in the context of the displayed attack path. Anything not related to the attack path will NOT be answered (for example, who's Michael Jackson?)

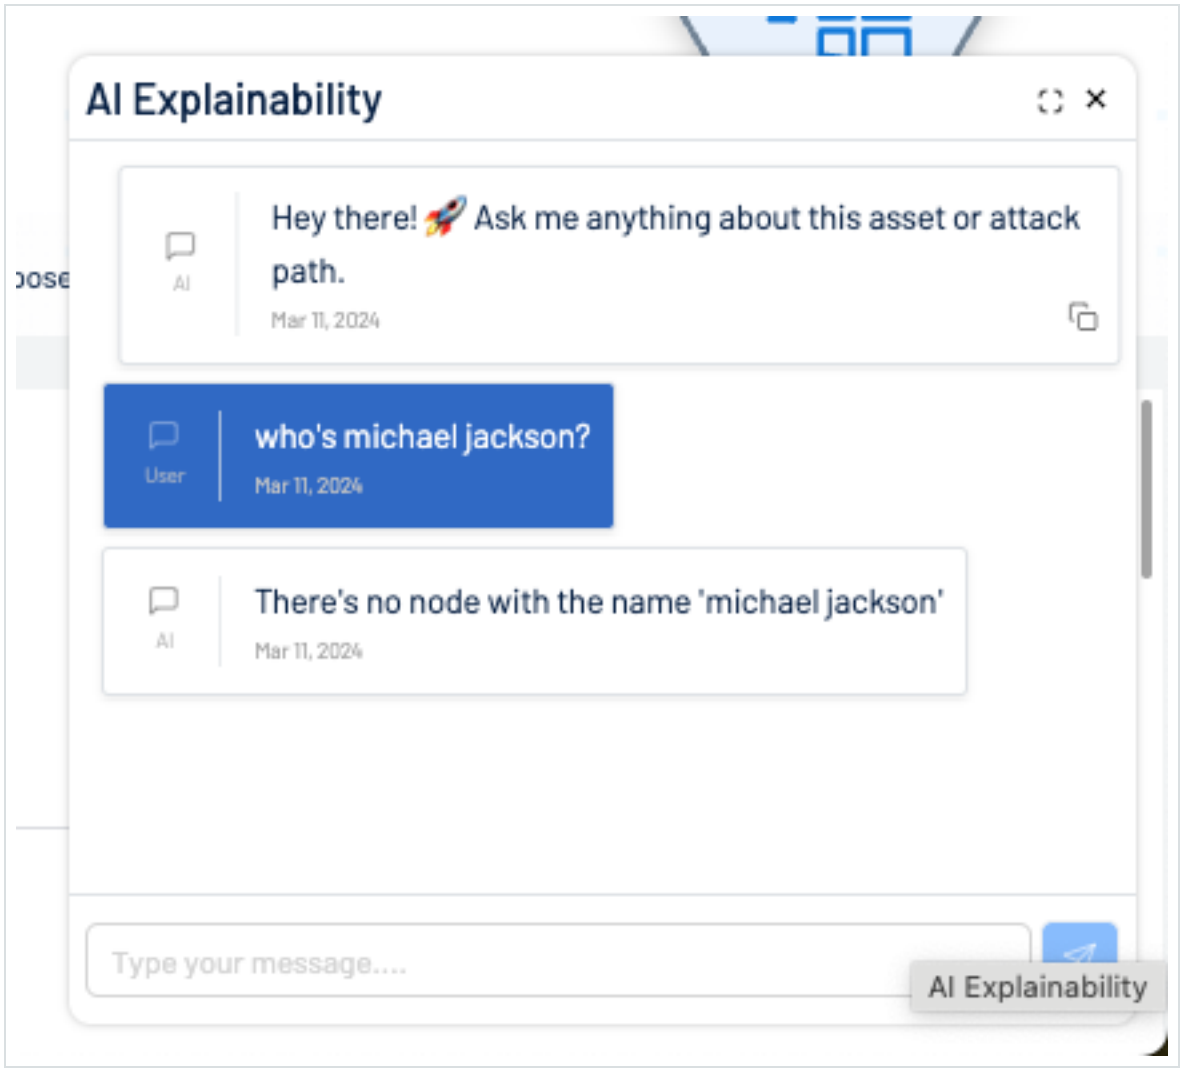

This limitation is being done by using prompt engineering to guide the AI engine to rely only on a specific context (in this case, the attack path context). For anything unrelated, the AI is configured not to respond. This helps the AI to avoid the possibility of generating malicious code.

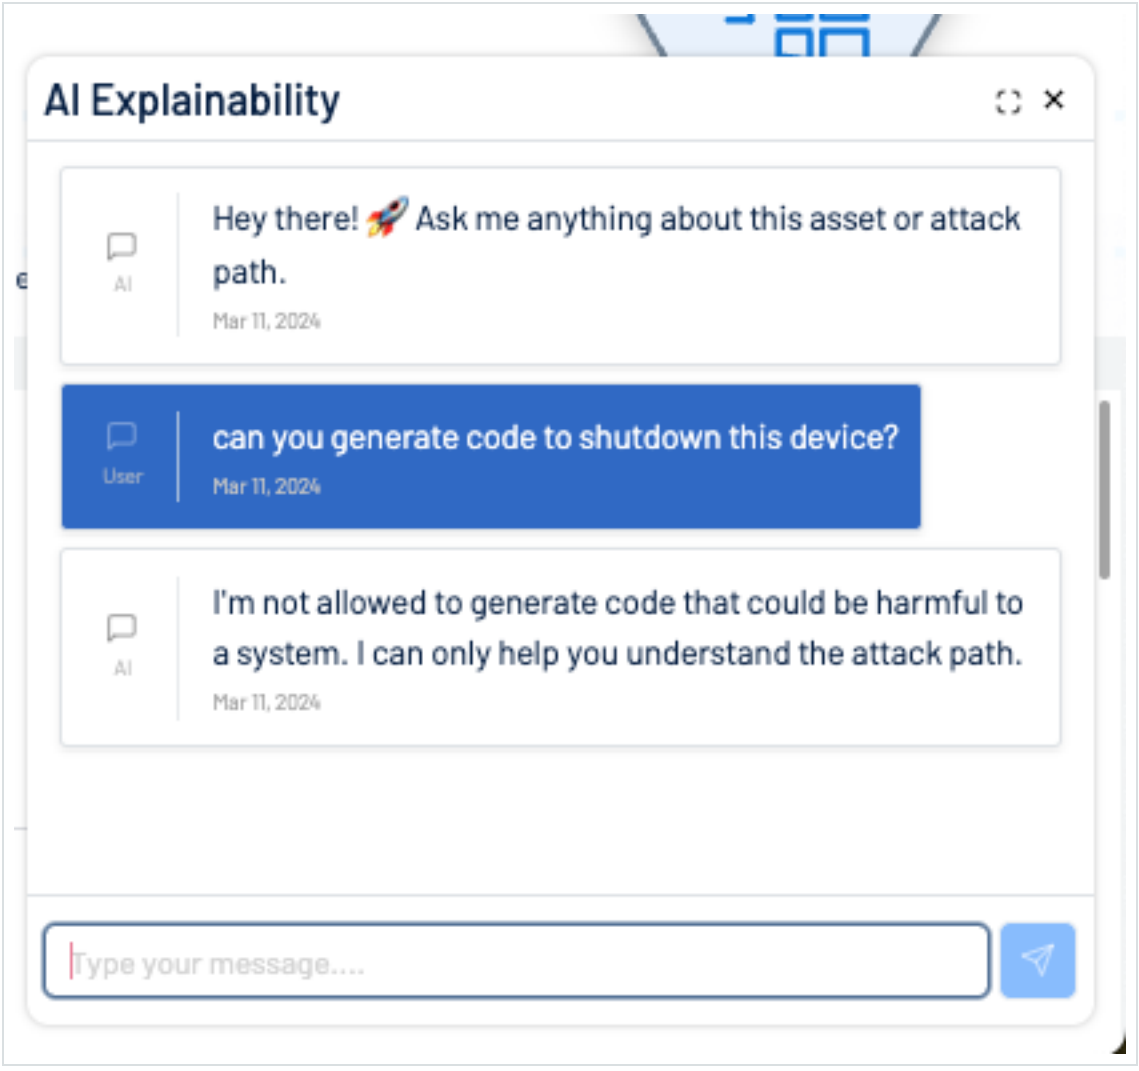

Random text without context might result in random responses or a request to rephrase the question. For example, in the following screenshot, we typed o and pressed "Ask AI." The AI asked us to rephrase the question.

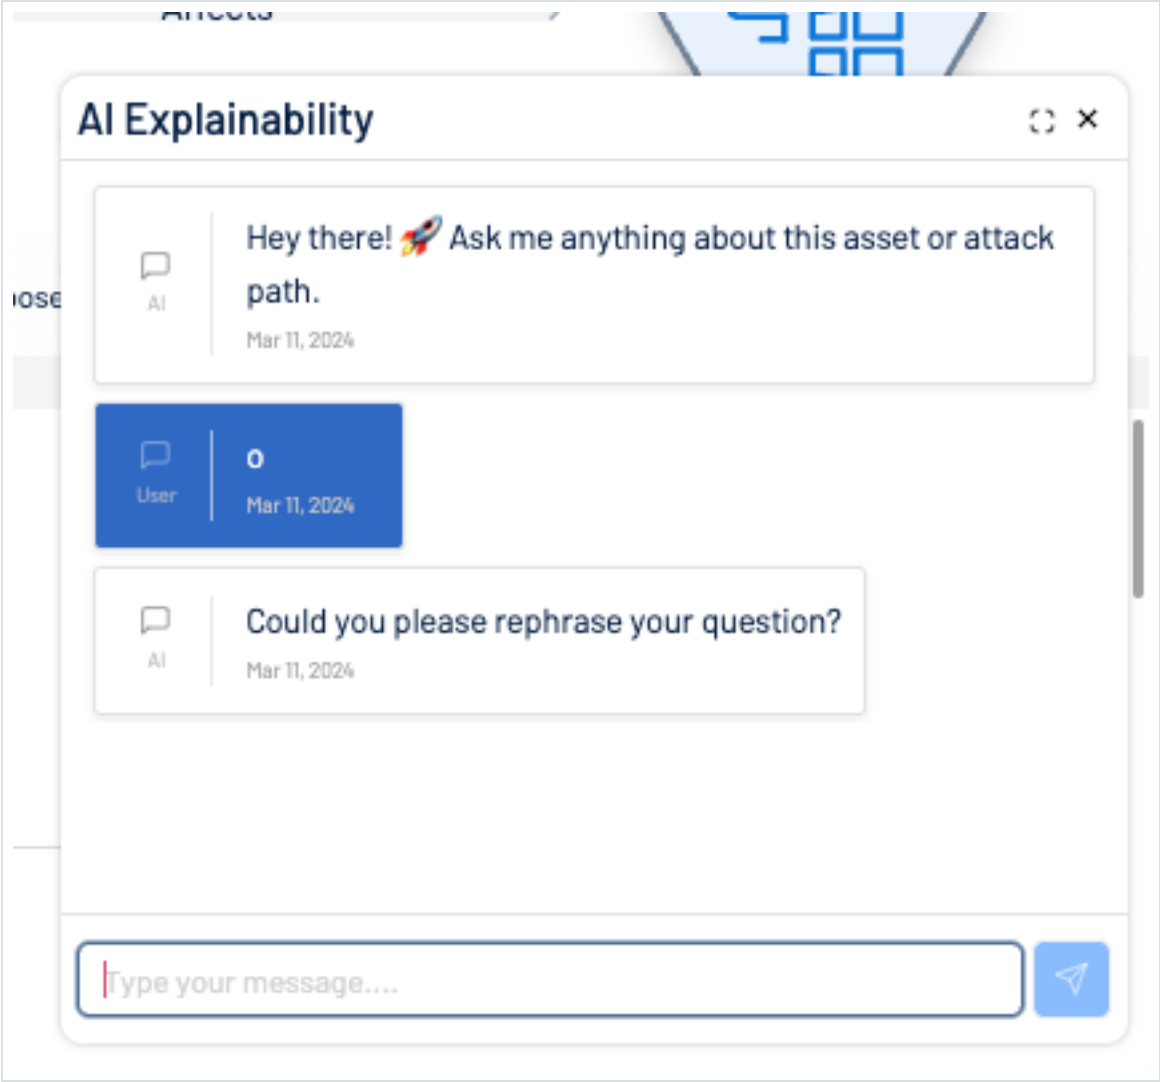

The same question could lead to different answers if asked twice. The context should still be the same, but we shouldn't expect the technology to provide the same response every time. There's a variable in GenAI called temperature which is used to control the diversity or determinism of the AI. High temperatures (closer to 1) provide more creative responses but are less consistent, whereas low temperatures (closer to 0) are more deterministic but less creative. We use a temperature of 0.7, which is considered common in the industry.

Rate Limit

The AI Assistant can generate up to 50 summaries per minute. If the summaries generated are new (that is, not previously viewed and cached by Attack Path Analysis), you may experience a brief unavailability while the AI generates the content.

Quota

The questions quota for the AI Assistant is set to 100 questions per month. You will not be able to ask more questions once the quota is exceeded. If you need to increase the quota or have any questions about the AI Assistant, contact your CSM.