Data Privacy Overview
Seer

Updated June 11, 2025

Key Points

- AI generated output from your data is shown only to you, never other customers.
- Your data does not leave Sentry-controlled infrastructure for Seer analysis.
- No data of yours is used to train Seer generative AI models by default and without your permission.
- The data, security, and compliance commitments we make to you about the overall Sentry service also apply to Seer.

Seer Overview

Seer is Sentry's AI agent that automates the triaging, debugging, and fixing of application issues for our customers. It consists of:

- **Issue scan**: Seer runs a lightweight "fixability" model against errors and performance issues as they are reported to your Sentry instance. It uses the output of this model to help you identify and prioritize working on fixable issues, including letting Seer handle them for you.
- **Issue fix:** Automatically fixes issues, including bugs and performance problems, by performing a deep analysis of the problem with LLM agents—identifying the root cause, providing a suggested code fix, and opening a pull request.
    - **Root cause analysis:** Analyzes issues using contextual data from Sentry—such as stack traces, traces, profiling data, commit history, tags, and logs—and, optionally, your relevant code, to suggest the root cause of an issue.
    - **Solution and code changes:** Generates code fixes and can create pull requests in GitHub to resolve issues faster.

Data Architecture

Seer is powered by third-party generative AI models that are sourced and hosted within Sentry's production infrastructure. This is enabled by our existing infrastructure subprocessors via platforms such as GCP VertexAI and AWS Bedrock. The diagram below shows the data architecture for Seer. By design, your data inputs, including code, and any suggested code changes generated from your data inputs (i.e., outputs) stay within Sentry production infrastructure. The AI-generated outputs from your data inputs are shown only to you, and never shared with other customers.

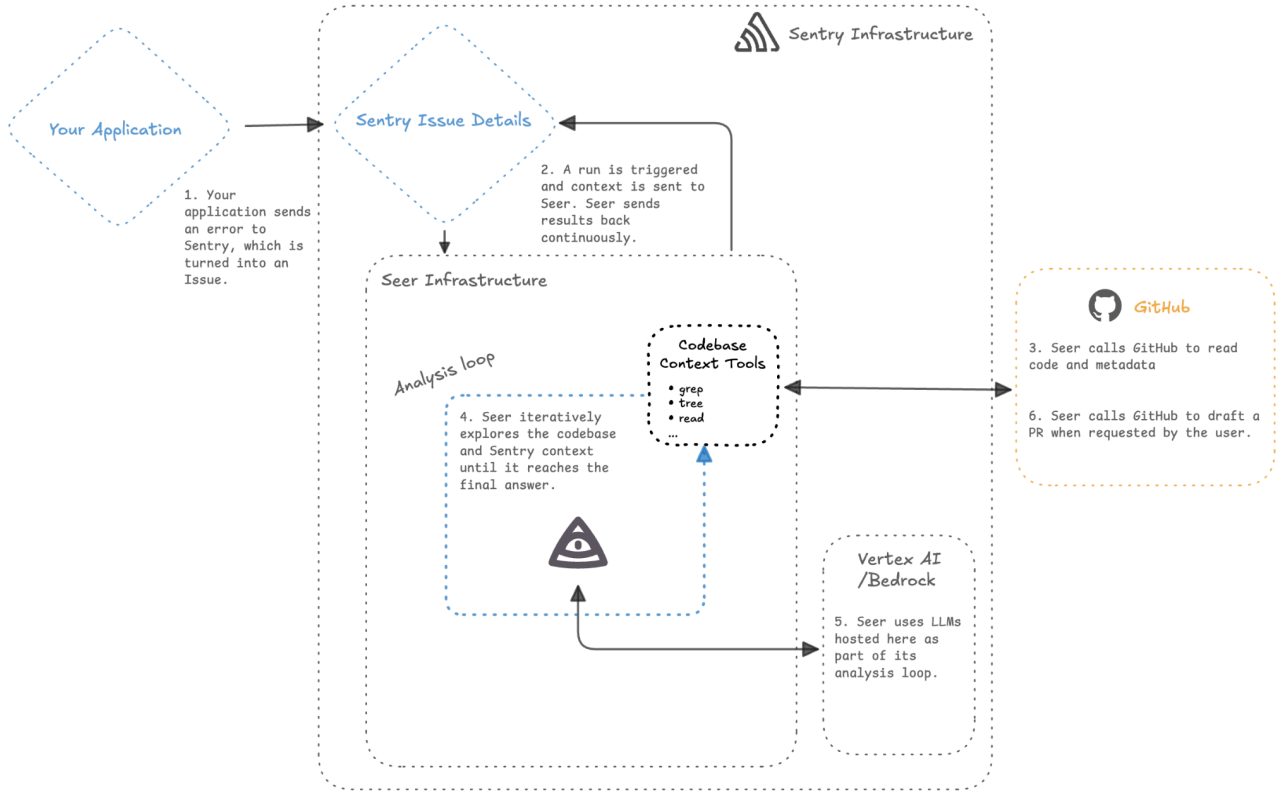

Note: The communication between Seer and GitHub depicted above is optional, based on your configuration, and you can control read and write access independently.

How Seer Leverages Your Code

By default, Seer performs analysis on the data already collected and available within your Sentry instance based on your configuration of Sentry SDKs. Seer does not require additional data or separate, direct access to your code. You can choose to enable such access to your code by connecting Sentry to your GitHub repositories via our optional Sentry GitHub integration.

Sentry GitHub Integration

What does it do?

If enabled, the Sentry GitHub integration allows Seer to provide deeper, more accurate root cause analyses and solution suggestions. Specifically, Seer can access sections of code to enrich your stack trace data, identify and read the specific files relevant to an issue, find relevant class or function definitions, and better understand your application.

How does it work?

The Sentry GitHub integration is designed with the following features.

- **Read-only access**. If the integration is enabled, Sentry (including Seer) will have read-only access to the code within your selected GitHub repositories.
- **Configuration**. You can configure the repositories that Sentry (including Seer) can access when you set up the integration, and your Seer settings allow you to designate the repositories, including specific branches, on which Seer analysis is run on a per-project basis.
- **Processing**. To provide its analyses, Seer may use the GitHub API to fetch your configured repositories into a containerized workload (Kubernetes pod) that's part of the same, trusted infrastructure that runs the rest of the Sentry service. Your code is only persisted while the Seer analysis is actively running, and is immediately and permanently deleted once the run completes. In the future, we may explore offering an opt-in, secure repository caching system for customers who want to improve Seer's speed.

Seer GitHub App

If you choose to install the Seer GitHub app, Seer can draft a pull request for your review directly in the GitHub repositories you select and configure.

This requires the app to have write-access in order to:

- Create new development branches; and
- Draft pull requests with Seer's suggested code changes to resolve an issue.

Seer is not designed to write to your main branch or other pre-existing branches, nor to merge any pull requests without your review and action on GitHub.

Terms and Data Rights

Sentry's Terms of Service (or other agreement with you for the Sentry service, including any Master Subscription Agreement) applies to Seer. These terms grant us limited rights to use your data to provide the Sentry service to you.

We will not train Seer generative AI models using your data by default and without your permission.

If you want to help improve the Sentry service, including Seer, you may opt-in to further use of your data for product improvement, including to train generative AI models.

Security and Compliance

SOC 2 Type 2 Report + HIPAA Attestation, ISO 27001:2022 Certification

Seer runs on the same production infrastructure as the broader Sentry service, and is secured to the same information security and compliance standards, including as set forth in our Security Policy. Seer is in-scope for Sentry's SOC 2 Type 2 Report and ISO 27001 certification, as well as our HIPAA attestation.

Data Commitments

For clarity, the commitments Sentry makes to you for the broader Sentry service apply to Seer, including for data retention and data storage location. If you are sending us personal data or PHI, this also includes our DPA and BAA commitments and flow-down of these commitments to the infrastructure subprocessors that host the models used for Seer.