Data Privacy Overview
Seer

Updated January 21, 2026

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 debugging agent that catches breaking changes before deployment, and debugs and fixes issues in production.

It consists of:

- **Code review:** Reviews your code changes against historical errors and performance issues to catch potential problems and offers suggestions on how to fix them before you merge your pull request.
- **Issue fix:** Uses LLM agents and issue context from Sentry—such as stack traces, distributed traces, replays, logs, profiling data, commit history, and tags—and your relevant code, to identify the root cause of errors and performance issues and suggest how to fix them.
    - ***Code changes***: Can also generate suggested code changes and open pull requests in GitHub to resolve issues faster. Or, if you prefer, you can take the analysis and context and use your favorite LLM to help draft the final fix.

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 pull requests, stay within Sentry production infrastructure for Seer analysis. The AI-generated outputs from your data inputs are shown only to you, and never shared with other customers.

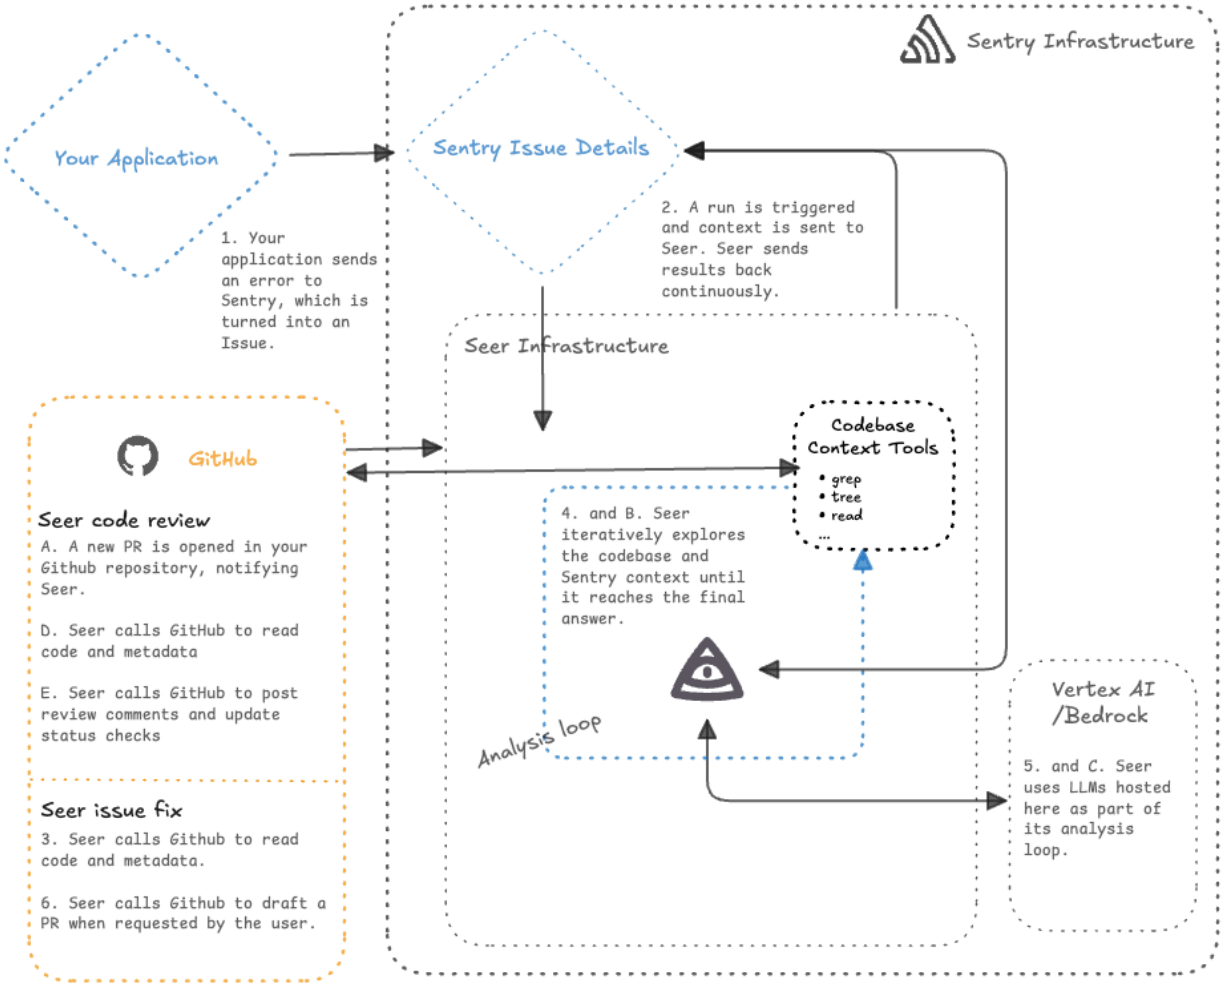

Seer code review involves flows A to E; Seer issue fix involves flows 1 to 5.

Note: The communication between Seer and GitHub depicted above is required to access Seer features. You can control which GitHub repositories Seer has access to.

How Seer Leverages Your Code

Seer requires direct access to your code to provide accurate root cause analysis and solution suggestions, and to find bugs in your code. To use Seer, you must connect Sentry to your Github repositories via our Sentry GitHub integration.

Sentry GitHub Integration

What does it do?

The Sentry GitHub integration allows Seer to access your code and your Github repository information, metadata, and workflow objects 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 and write access**. Once the integration is enabled, Sentry (including Seer) will have access to the GitHub repositories you select with write access to meta data and workflow objects–specifically, checks, commit statuses, issues, pull requests, and repository hooks–as well as repository contents to:
    - Post review comments and update status checks;
    - Create new development branches; and
    - Draft pull requests with Seer's suggested code changes to resolve an issue.

  Sentry 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.
- **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.

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.