

Automate, Group, and Get Alerted: A Best Practices Guide to Monitoring your Code



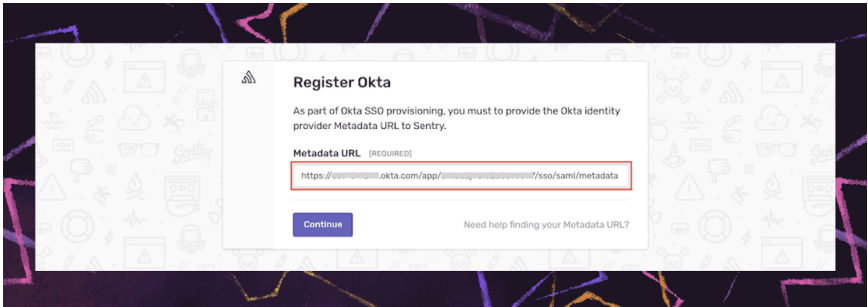
As companies grow, so do their products, teams, and the number of external tools. For engineers, that can mean code sprawl, data silos, notification fatigue, and some ‘WTF?!’ moments along the way as they try to make sense of it all.

To help combat the chaos, Sentry provides several ways to manage the volume, noise, and potential team disconnects that come with launching and scaling projects so you can see the issues that actually matter and solve them faster – without the expletives.

1. Automate Sentry Access with SCIM and SSO

Managers and IT departments can save time on user management tasks while reducing security risks by setting up [SSO](#) and [SCIM](#).

Automatically provision and deprovision users and teams directly in [SAML2](#) identity providers like Okta and Azure AD using SCIM. Whether employees change roles, leave the company, or teams experience structural changes, you can programmatically add or remove team members, update team attributes such as name, or replace an entire member set – without leaving your identity provider.



2. See Issues that matter with Review List, Grouping, and Filters

Events in Sentry tend to grow proportional to the organization – with more customers, more products, and more lines of code you’re likely to also see more events.

Filter and group events

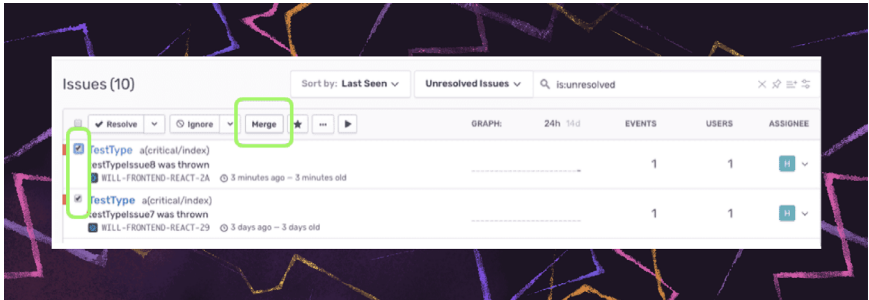
To help manage events, Sentry’s [grouping strategy](#) provides a built-in grouping algorithm to join related errors into a single issue, so you don’t receive a deluge of issue notifications.

Grouping works because all events have a fingerprint. Events with the same fingerprint are grouped based on information available within the event, such as *stacktrace*, *exception*, and *message*.

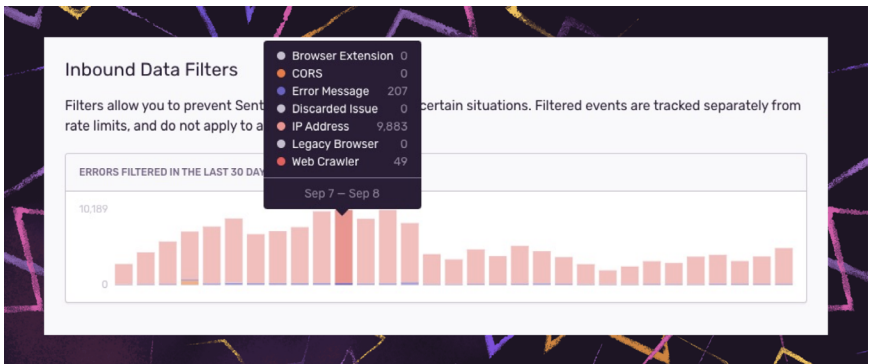
Effective issue management sometimes involves looking across the constellation of these events to fine-tune the signal from the noise with custom grouping.

When customizing [fingerprinting](#) rules, it can be helpful to use [Discover](#) to explore the projects that create the most Sentry issues relative to absolute error volume. Teams can

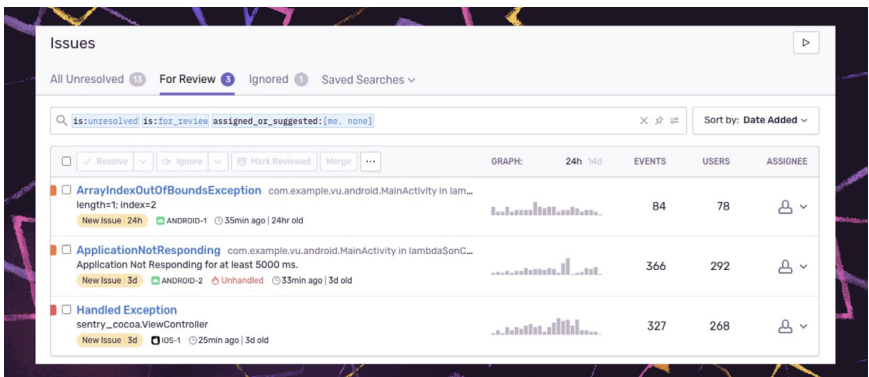
then combine patterns from event data alongside [intimate knowledge](#) of their codebase to inform custom fingerprinting or [stack trace](#) rules. These rules are all configured on a per-project basis.



Once you have grouped events, you can also filter events with custom [filter rules](#) using a similar approach above. You can apply filter rules at both the [SDK level](#) and via [inbound filters](#) within Sentry.



Access a hyper-focused issue view



To help you see the issues you care most about, [Review List](#) provides a centralized view of a sub-set of all your unresolved issues. By default, Review List filters to display issues assigned to you or suggested for you. This view offers a more manageable collection of events that team members can review and triage daily.

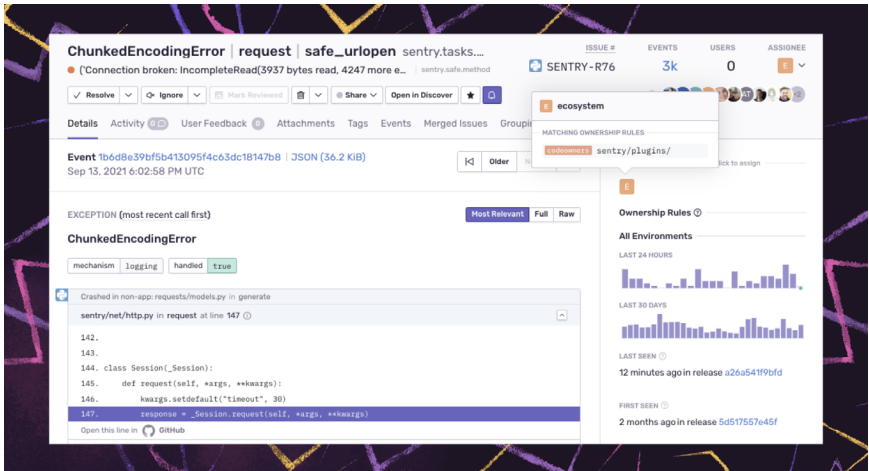
Sample and manage event quota

Understanding your events can also inform [sampling](#) decisions. For example, if a project has many unique issues, using rate limiting to control quota could bias your data toward noisier events, whereas sampling at the SDK level may provide a more representative sample of your errors. Our [user guide](#) is also a good starting point for exploring available options for sampling events across your Sentry projects.

3. Solve issues faster with ownership rules, alerts, and trace view

Teams can take action faster when issues are owned. Sentry provides [ownership rules](#) to facilitate finding the who behind the code so you can streamline time to resolution.

For those that currently use GitHub or GitLab CODEOWNERS files, you can now import your existing files in Sentry using [Code Owners](#), allowing you to automatically assign issues or route alerts to the responsible individuals or teams.



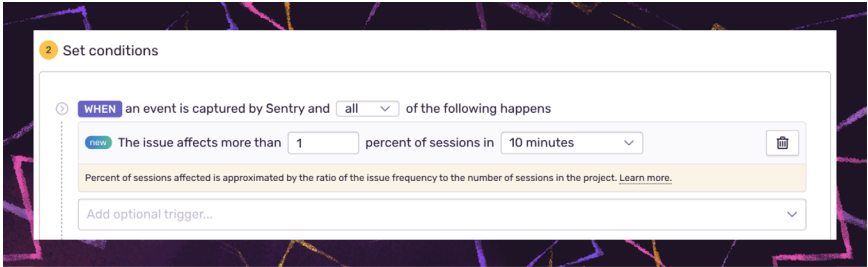
Sentry uses ownership rules to surface suggested assignees during the triage process and assign those issues to team members automatically, so developers only see the issues they care about, and product owners spend less time managing the queue.

We also have the concept of [Suspect Commits](#). Suspect Commits highlights the commit that introduced the error and the person who introduced it, so you can reproduce and solve the issues you own faster. In your release process, add a step to create a release in Sentry and associate it with commits from your linked repository to get started.

Set alerts and notifications

Coupled with Ownership Rules, you can set alerts and notifications to go directly to the corresponding individual or team who owns the code in the tools they use every day like Jira, Slack, email, and Pagerduty. Slack, for example, can send workflow notifications, deploy notifications, and alerts directly to [teams and individuals](#) where they can triage and resolve issues directly in chat.

However, it is not enough to alert the right people in the right place. The most effective alerts are also sent at the right time, frequency, and with the right context to take action quickly.



That's why we offer [issue](#) and [metric alerts](#), so you only receive alerts when an issue meets specific [trigger](#) criteria you set, like if a resolved issue is re-appearing or an issue is affecting a percentage of sessions.

Percent-based alerts, for example, adjust to the changes in app usage so you can set alerts to trigger when an issue exceeds a certain percentage of user sessions in a period, allowing you to prioritize issues based on user impact.

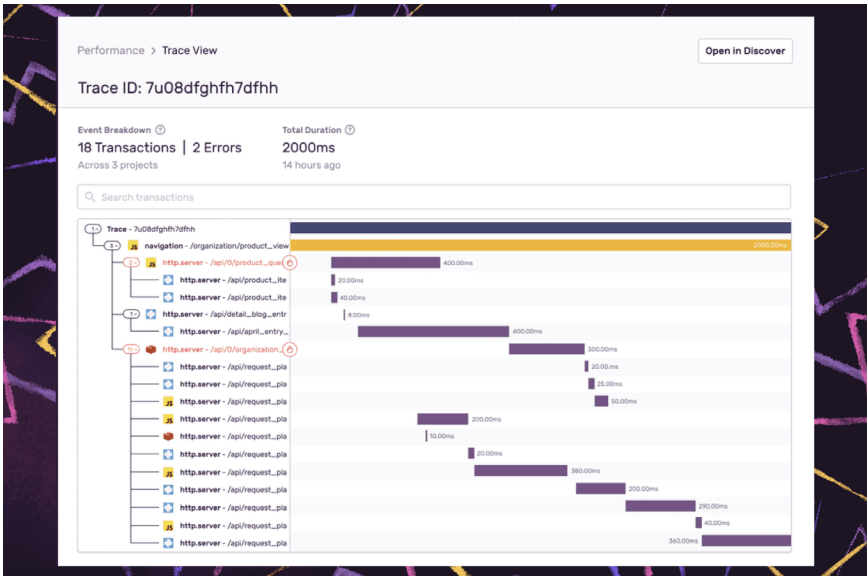
Performance-based [metric alerts](#) are another way to leverage aggregates (p50 duration, p75 LCP, etc.) to improve alerting triggers and behavior. Metric alerts tell you when a [metric](#) crosses a threshold, such as a spike in the number of errors in a project, or a change in a performance metric, like [latency](#), [Apdex](#), [failure rate](#), or [throughput](#).

To further control noise, Sentry not only evaluates the specified ["When" trigger criteria](#) each time it receives a new event but also checks the ["If" conditions filters](#), such as if the issue is older than a certain duration or if the event is from your latest release. Last, Sentry will only send a notification after checking the [Action Interval](#), which controls how often the alert can be triggered for a particular issue.

Trace performance issues with trace view and trace navigator

Leveraging transactions in addition to errors forms a more complete picture of the interconnectedness of your services as they grow. [Trace view](#) and [trace navigator](#) are tools to help accelerate triage at scale.

Trace view provides a broad visualization of the transactions comprising a single trace. This view can be helpful when debugging errors related to slow services or transactions. Trace navigator provides an abbreviated view of the entire trace giving you a relative context to reference as you follow the path of a request across services and into errors that occurred during the journey.

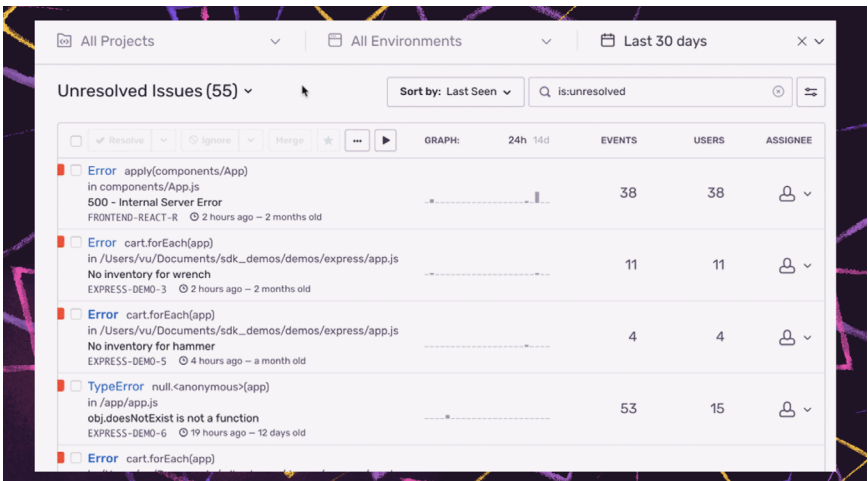


4. Collaborate across teams with project filters and activity tab

With a growing and more distributed codebase, issues can often cut across project concerns and time zones.

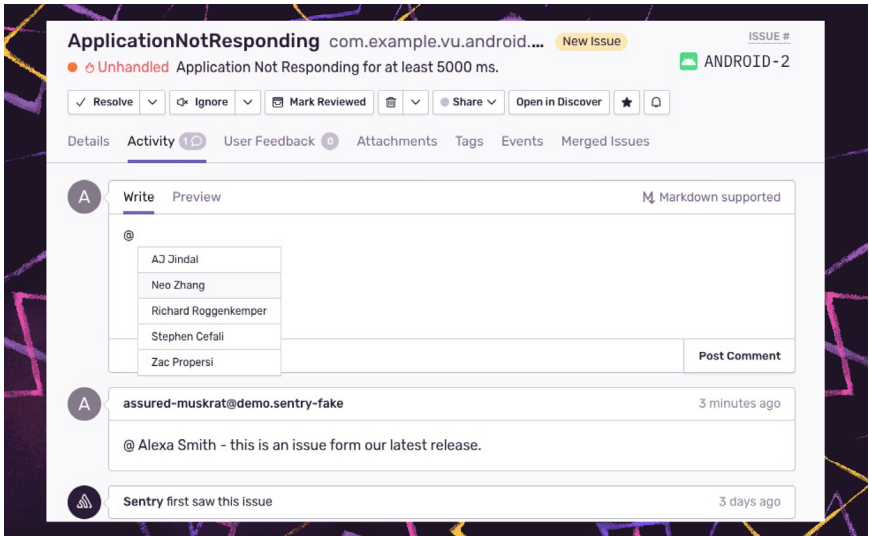
A universal issue view with cross-project issues

[Cross-project issues](#) within the Issue Stream offer stakeholders the ability to aggregate errors or performance issues across multiple Sentry projects into a single view.



With cross-project issues, you get a unified view of all your projects and environments so you can see what's impacting your service across any client. You can also drill down into specific projects or environments and sort issues by things like impact, frequency, or priority, so you can quickly see your top issues and triage accordingly.

Streamline issue chatter with activity tab



Last, consolidate communication either within Sentry using the [activity tab](#) where team members can see and comment on their event activity history directly, or sync context from Sentry into the [project management](#) tools your teams use already.

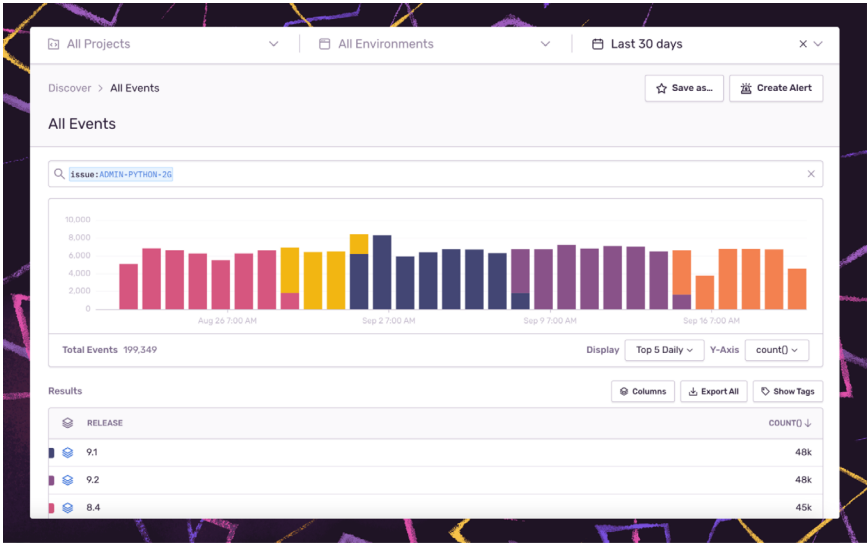
5. Learn from issues to improve your workflow with Discover and Releases

To understand and share how your application health and stability change over time as event and release volume increase, use [Discover Queries](#), [Dashboards](#), and [Release Health](#). Here are a few ways you can incorporate meaningful data into your workflow:

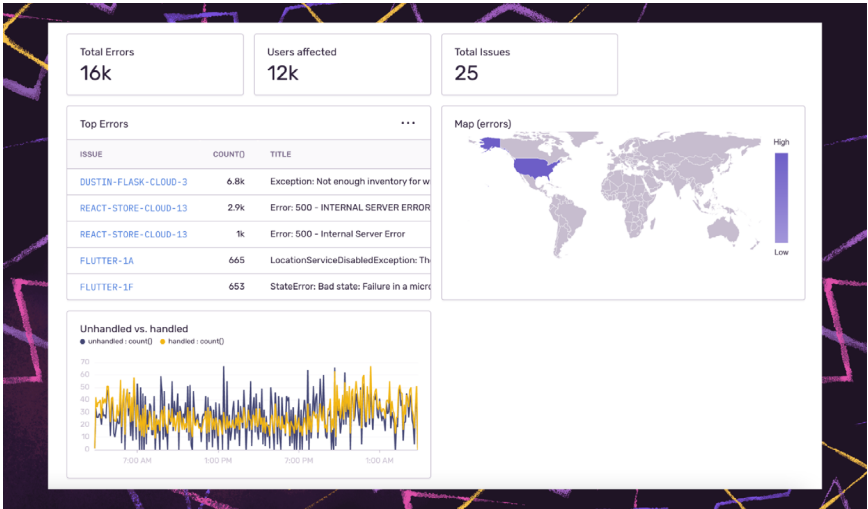
Answer critical business issues with Discover

Sentry offers a few ways to drill into your raw error, performance, and release data to answer questions like 'Which issues are persisting across releases?' or 'What are my noisiest projects and why?'

You can either write queries in Discover, using the same [syntax](#) used for Sentry search, or use [pre-built queries](#) to give you a starting point from where you can drill down to answer questions like 'What custom filters should I implement?' or 'How can I improve my triage strategy?' Of course, the pre-built queries only scratch the surface. To help you build custom queries, check out the [available fields](#) and [equations](#).

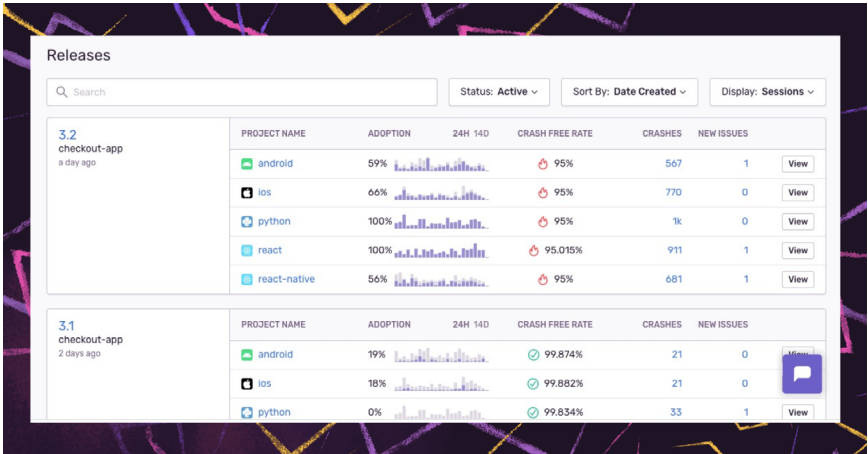


Share important team or organizational-level metrics using Dashboards



Dashboards are a visual way to capture and share what matters most across services, geographies, and stakeholders. Like Discover, there are [default dashboards](#) available out-of-the-box, or if you want something more customized to your business, you can either edit the default dashboards or build your own from scratch. A dashboard is shareable across the organization and can enhance weekly readouts, standups, or discussions around SLAs.

Get visibility into the stability of your application with Release Health



Sentry has added new options for better navigation of [release data](#) across your workflows. With [Release Health](#), you see core metrics like crash-free sessions, version adoption, and failure rate so you can quickly detect bad releases, investigate slowdowns, and prioritize the issues to solve across versions.

Plus, with a real-time view across all releases and [SemVer](#) (semantic versioning) support, you can compare active or past versions to understand release trends and learn how your team maintains release quality and version adoption as your platform grows.

Getting started with code observability at scale

Whether a Developer, Team Lead, or Product Owner, Sentry is continually working to help optimize your team's workflow so you can reduce time spent on troubleshooting issues and focus on more important things like building products that customers want.

If you're a current Sentry user, [log in](#) to your account to start implementing these tips today.

New to Sentry and want to learn more? [Request a demo](#) or [set up your free account](#) to get started.

