

# How **OpenTelemetry** Builds a Robust Observability Practice

Unlock immediate benefits without vendor lock-in.

**splunk**>





# What is it?

## OpenTelemetry is an open-source approach to collecting observability data.

As a project supported by the Cloud Native Computing Foundation (CNCF), OpenTelemetry provides APIs and SDKs in over a dozen languages to help you instrument, collect, and export telemetry data:



metrics

0010  
01010  
0101

traces



logs

... along with their associated metadata. This includes support for profiling and other data types we haven't considered yet.



# Why should you care?

## Because it's your data.

We value flexible data solutions, flexible thinking, and flexible jeans (and, hopefully, at least one dad joke). OpenTelemetry flexes to your data needs: users can take complete control of their data since they are not locked into any one vendor or their proprietary agents.

Send your data anywhere. Collect and process it through as many inputs and outputs as you choose. Generate custom data fields through what are now industry-standard APIs. OpenTelemetry continues to evolve. [Profiling](#) is the newest advance — continuous profiling gives insights into resource utilization at a code-level and allows for this profiling data to be stored, queried, and analyzed over time and across different attributes.


Besides that, it's a great time to learn about OpenTelemetry. New career opportunities are popping up because:

- Companies everywhere are adopting observability.
- Data ownership is more important than ever.
- Delivering better customer experiences is critical.

Observability with vendor agents is already here. Advances in OpenTelemetry and open collection are trending for a reason. You're still early. Even if you don't want a new job, using your data however you want will never grow old. There's time to build your skills and become a top contributor to the project.

OpenTelemetry enables software developers and SREs to find and solve problems faster, making customers happier. It ensures that control of your data remains with you and that you can instrument anything at your own speed rather than relying on a vendor.

[It's the future of observability](#) and provides the most flexible solution for transferring observability data from an application into an observability system. It is the premier open-source project for collecting performance data such as metrics, traces, and logs from applications and infrastructure, processing that data, and exporting it to observability systems.



**Use data in whichever way makes sense to your business — without any strings attached.**



# Why don't you want to be locked into a proprietary data model?


To illustrate OpenTelemetry's importance, here's a simple analogy:

Raise your hand if you started gaming on an Atari. Later, you added Nintendo. Then, Sega. Each time you wanted to play the hottest video game, you had to consider buying an entirely new gaming system. Can you imagine if there were a universal gaming system? (Dare to dream with us; maybe you would imagine something close to OpenTelemetry.)

Your gaming system will work perfectly if you want to play the same video games forever. As soon as something with flashier graphics comes out, you may need an entirely new system. This is what it's like to be trapped in a proprietary model. We've all been there. The same can happen with many vendors who sell you the license and infrastructure to enable your observability practice. Even though vendors may claim to be faster, cheaper, or better, the long-term costs often outweigh whatever is gained in the short term.

You may write it off as the cost of doing business until the subsequent need for an upgrade arises. Costs may balloon over time. What a vendor presents as a "good deal" may change as your data needs expand. If you want to adopt a new technology and your vendor isn't interested in it or hasn't devoted resources to it, you're out of luck. Unfortunately, adding a new proprietary infrastructure to your organization will far exceed the costs of the newest gaming console.

If you have OpenTelemetry, you're in the clear. Access to performance and troubleshooting data is easy and will flex to new data applications and observability advances.





# Here are the top ways that OpenTelemetry will unlock your observability practice.

- **Improve performance:** Collect metrics, logs, traces, and more with one lightweight collector, process them, and emit data to observability systems to solve your application performance problems faster.
- **Reduce toil and avoid vendor lock-in:** Once set up, practically any vendor can use the OpenTelemetry instrumentation. Changing observability vendors no longer requires some reinstrumentation effort.
- **Deliver software faster:** OpenTelemetry supports various languages, frameworks, and libraries, making it easier to instrument code regardless of the stack used. Thanks to the project's open-source nature, support for practically anything can be added.
- **Own your data:** The OpenTelemetry ecosystem permits you to send data where you want, process it how you want, and maintain full ownership of your telemetry data without vendor involvement.






# Easy to deploy and start today

OpenTelemetry supports hundreds of integrations with the most popular applications, operating systems, language runtimes, and application frameworks, with more integrations being developed daily. It's the most flexible solution, freeing you from needing to plan instrumentation — OpenTelemetry works now and will work in the future, no matter what your observability vendor does.

If you've read this far, you already know that OpenTelemetry is more than just a way to avoid vendor lock-in. While that is a great reason to go all in, it's not the only reason. Engineering teams need telemetry data to triage and respond to the issues across the systems they monitor. As organizations adopt cloud-

native infrastructures, an observability practice becomes even more important due to the complexity of these infrastructures and applications.

OpenTelemetry offers a unified way to instrument code with multiple integration points. As a result, a single application can send data to multiple observability backends with minimal added performance impact.





OpenTelemetry benefits from an active community of developers and contributors. It's early enough to join the community and become an expert. It's an opportunity for organizations looking to improve visibility and monitoring of distributed systems without vendor lock-in.

OpenTelemetry is the default choice for instrumentation, and if you aren't on board, it's like committing your business to a gaming system from the 1980s. The benefits of ownership, performance, and data reliability make OpenTelemetry the logical choice. If your platform can't handle it, then that's a clear weakness in your platform.

Splunk is a top contributor to the [OpenTelemetry project](#). Splunk Observability Cloud uses OpenTelemetry natively, speeding time to value and eliminating vendor lock-in. Thanks to the efforts of the broader OpenTelemetry project, Splunk customers benefit from modern framework and language support being delivered more quickly.

Read [A Beginner's Guide to Observability](#) to learn how to build observability into your application environment.

[Get started for free.](#)

Keep the conversation going with Splunk.



Splunk, Splunk> and Turn Data Into Doing are trademarks and registered trademarks of Splunk Inc. in the United States and other countries. All other brand names, product names or trademarks belong to their respective owners. © 2024 Splunk Inc. All rights reserved.

24.CMPEbook.how-otel-builds-a-robust-ot1y-practice\_v6

**Observability is a practice and OpenTelemetry is fundamental to that practice.**

