

Product or feature is in a pre-release state and might change or have limited support. For more information, see the [product launch stages](#) (/products/#product-launch-stages).

This page describes how to modify your Node.js application to capture profiling data and have that data sent to your Google Cloud project. For general information about profiling, see [Profiling concepts](#) (/profiler/docs/concepts-profiling).

Profile types for Node.js:

- Heap
- Wall time

Supported Node.js language versions:

- 8.9.4 or higher on the 8.x version branch.
- 10.4.1 or higher on the 10.x version branch.

Supported operating systems:

- Linux. Profiling Node.js applications is supported for Linux kernels whose standard C library is implemented with `glibc` or with `musl`. For configuration information specific to Linux Alpine kernels, see [Running on Linux Alpine](#) (#running_with_linux_alpine).

Supported environments:

- Compute Engine
- Google Kubernetes Engine (GKE)
- App Engine flexible environment
- App Engine standard environment
- Outside of Google Cloud (For information on the additional configuration requirements, see [Profiling applications running outside of Google Cloud](#) (/profiler/docs/profiling-external).)

Before you use the profiling agent, ensure that the underlying Profiler API is enabled. You can check the status of the API and enable it if necessary by using either the Cloud SDK `gcloud` command-line tool or the Cloud Console:


In all of the supported environments, you use the Profiler by installing the package `@google-cloud/profiler`, adding a `require` statement to your application, and then deploying the application in the usual way.


The package `@google-cloud/profiler` depends on a native module. Pre-built binaries for this native module are available for all supported language and platform combinations. To determine which pre-built binary to install, `@google-cloud/profiler` uses `node-pre-gyp`.


If your build system doesn't match the Node version, libc version, architecture, or platform for the system where you deploy your application, then you must specify the correct value in the `npm install` command. For details, see [node-pre-gyp](https://www.npmjs.com/package/node-pre-gyp) (<https://www.npmjs.com/package/node-pre-gyp>).

To install the latest version of Stackdriver Profiler, do the following:

If you are also using the Trace agent, when you modify your application, import the Profiler package after the Trace agent package (`@google-cloud/trace-agent`). For more information, see [Setting up Stackdriver Trace for Node.js](#) (</trace/docs/setup/nodejs/>).







After Profiler has collected data, you can view and analyze this data using the Profiler interface. To get started using this interface, see [Opening the Profiler interface \(/profiler/docs/using-profiler\)](/profiler/docs/using-profiler).

When you load the Profiler agent, you specify a service-name argument and an optional service-version argument to configure it.

In the App Engine flexible environment, you do not have to specify these arguments; they are derived from the environment.

The **service name** lets Profiler collect profiling data for all replicas of that service. The profiler service ensures a collection rate of one profile per minute, on average, for each service name across each combination service versions and zones.

For example, if you have a service with two versions running across replicas in three zones, the profiler will create an average of 6 profiles per minute for that service.

If you use different service names for your replicas, then your service will be profiled more often than necessary, with a correspondingly higher overhead.

When selecting a service name:

- Choose a name that clearly represents the service in your application architecture. The choice of service name is less important if you only run a single service or application. It is more

important if your application runs as a set of micro-services, for example.

- Make sure to not use any process-specific values, like a process ID, in the service-name string.
- The service-name string must match this regular expression:

```
^[a-z]([-a-z0-9_]{0,253}[a-z0-9])?$
```

A good guideline is to use a static string like `imageproc-service` as the service name.

The **service version** is optional. If you specify the service version, Profiler can aggregate profiling information from multiple instances and display it correctly. It can be used to mark different versions of your services as they get deployed. The Profiler UI lets you filter the data by service version; this way, you can compare the performance of older and newer versions of the code.

The value of the service-version argument is a free-form string, but values for this argument typically look like version numbers, for example, `1.0.0` or `2.1.2`.

The profiling agent can report logging information. To enable logging, set the `logLevel` option when starting the agent. The supported `logLevel` values are:

- `0`: disables all agent logging.
- `1`: enables error logging.
- `2`: enables warning logging (default).
- `3`: enables info logging.
- `4`: enables debug logging.

Set the `logLevel` value in the same object that provides the service context:

If you use Docker images that run with [Linux Alpine](https://alpinelinux.org) (https://alpinelinux.org) (such as `golang:alpine` or just `alpine`), you might see the following authentication error:

Note that to see the error you must have agent logging enabled.

The error indicates that the Docker images with Linux Alpine don't have the root SSL certificates installed by default. Those certificates are necessary for the profiling agent to communicate with the profiler API. To resolve this error, add the following `apk` commands to your Dockerfile:

You then need to rebuild and redeploy your application.

The following are known issues in the beta release of Stackdriver Profiler for Node.js:

- The profiling agent for Node.js interferes with the normal exit of the program; it can take up to an hour for the program to terminate after all the tasks in the program have completed. Forcibly exiting the program, for example, by using `Ctrl-C`, causes the program to terminate immediately.

To learn about the Profiler graph and controls, go to [Using the Stackdriver Profiler Interface](/profiler/docs/using-profiler) (/profiler/docs/using-profiler). For advanced information, go to the following:

- [Filtering profiles](/profiler/docs/filtering-profiles) (/profiler/docs/filtering-profiles)
- [Focusing the graph](/profiler/docs/focusing-profiles) (/profiler/docs/focusing-profiles)
- [Compare profiles](/profiler/docs/comparing-profiles) (/profiler/docs/comparing-profiles)

