Heroku dyno monitoring

HG Heroku Add-on for your Heroku applications

This article will outline the default metrics given to you by the default Heroku Dyno monitoring tools available on every account. However, the default monitoring tools just aren't enough. In order to get great observability of your systems, you need to be using a third party monitoring tool.

The HG Heroku Monitoring add-on is the best tool for monitoring the many different layers of an application. You can pull in metrics from your Heroku Dynos, your front end, your cloud provider as well as other systems such as your CI/CD pipelines.

When you log into HG Heroku Monitoring, your Heroku Dyno monitoring dashboards will be automatically generated. You should install the HG Heroku Monitoring app right now!

Check out how it looks here:

Default Heroku Dyno Monitoring

The default Heroku Dashboard view provides a 24-hour rundown of site StatsD and sparklines for each program. Summary metrics are the cumulative amount of errors in dyno and router and the current 95 percentile reaction time and a 10 min resolution throughput value. We view only applications with network dynos. Additional specific information on the app is presented, including dyno formation/location, current implementation, and language.

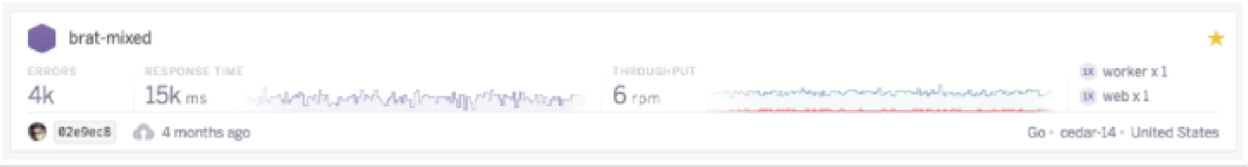

brat-mixed

ERRORS 4k

RESPONSE TIME 15k ms

THROUGHPUT 6 rpm

1X worker x 1

1X web x 1

02e9ec8 · 4 months ago

Go · cedar-14 · United States

If you need more information than this, consider using the HG Heroku Monitoring add-on to correlate your Dynos with the rest of the system. You can pull in metrics from anywhere and see them side by side with your Dynos.

Dyno metrics

The following measurements are obtained for all process types. The list of process types, and the average measurements for each process types, are shown below:

A single plot is seen as cumulative maximum memory usage, adding maximum RSS and maximum swap-reported memory for 10 or 1-minute intervals. A blue line shows the mean cumulative memory (RSS + Swap). The memory limit is seen as a slashed grey line of red-flagged quota violations. For the chosen timespan (i.e. 24 hours), the current, medium, and peak percent memory and the raw value are shown.

1. Memory quota: The maximum RAM of your dyno type to use, which will cause an R14 memory error.
2. Total memory: Mean total memory reflects the part of the memory to be optimized by users. In 10 or 1 minute RSS and swap are seen in an average of all dynos. Total memory.
3. RSS: The quantity of memory of a certain process type in RAM. Each 10 or 1-minute interval is recorded for Max RSS.
4. Swap: The dyno memory portion is stored on the disc in megabytes. A few megabytes of the swap are typically used per dyno for an operation. However, higher swap levels may suggest an unnecessary memory consumption in comparison with the dyno. This can lead and should be avoided for too long response times. For every 10 or 1-minute period, the max swap is recorded.

In this case, a shift in dyno size from 1X to 2X raises the memory quota.

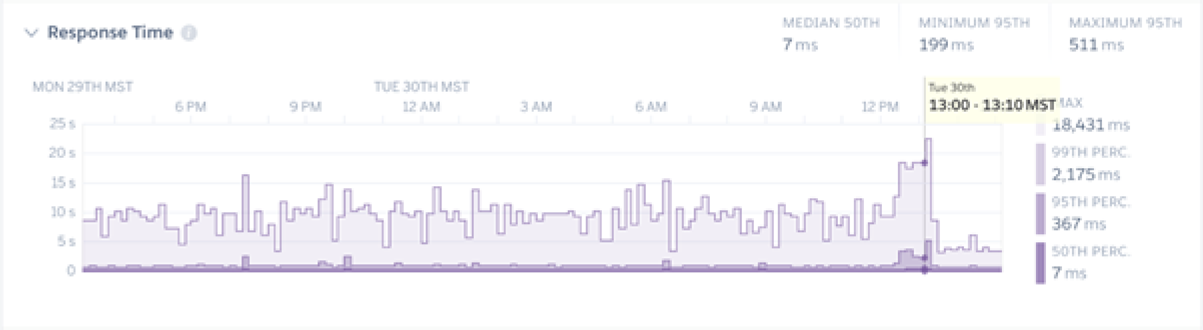

For further metrics, and metrics more specific to your services and application, check out the [HG Heroku Monitoring]() add-on. HG Heroku monitoring is immensely flexible, giving you both the default metrics, as well as the ability to expand beyond that and get custom metrics for your application.

Handling dyno loads

The load value means a work-on task (a loop or thread) that operates on a CPU or waits for a CPU to work, but has all its energy to work on otherwise. The load value does not contain IO-waiting tasks. Examples include:

1. 1 m Load average: The mean load average of 1 minute per every 10 minutes is seen for sample periods of 10 minutes. The average load is seen explicitly for 1-minute windows. The average load represents the number of Processor activities over the previous 30 minutes as an exponentially humidified average.
2. 1 m Load Limit: This is the sum of a 1-minute average for the period for 10-minute windows. 1 m Load Max The maximum load average for sampling periods of 20 seconds is shown over 1-minute intervals.

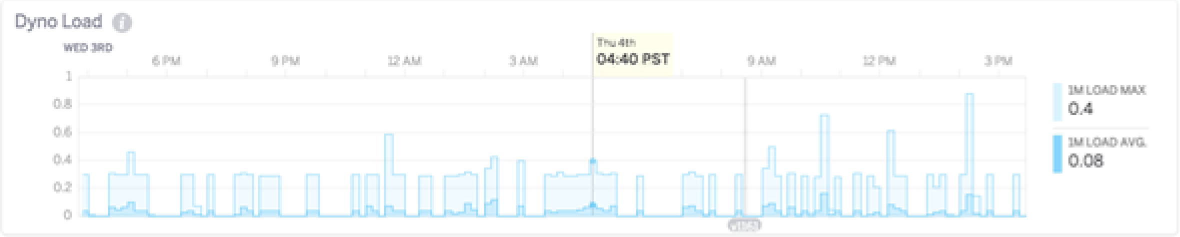

Dyno events

The Table of Events includes Heroku errors and user-initiated events that impact device health. User operations currently monitored include deployments, improvements in setup, and reboots. Blue is displayed for operation events (e.g. installation and setup changes). Color gradients represent the relative number of events of each type happening at any moment, as a limit of only an event type marker would be shown per interval of time. Further information can be sought by hovering over the actual case. These specifics include the explanation of errors, and for incidents triggered by the person who made the move.

There are crucial mistakes in red, orange alert levels, and grey detail.

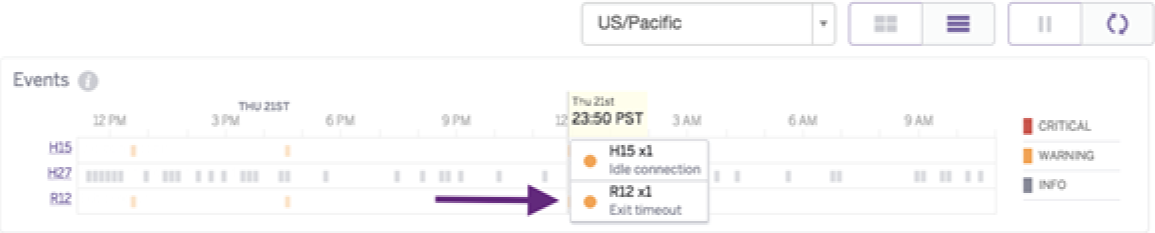

● H15 x1
Idle connection

● R12 x1
Exit timeout

There are two types of platform status events shown: (yellow) accidents and planned (grey). Just the activities for your area are available to you.

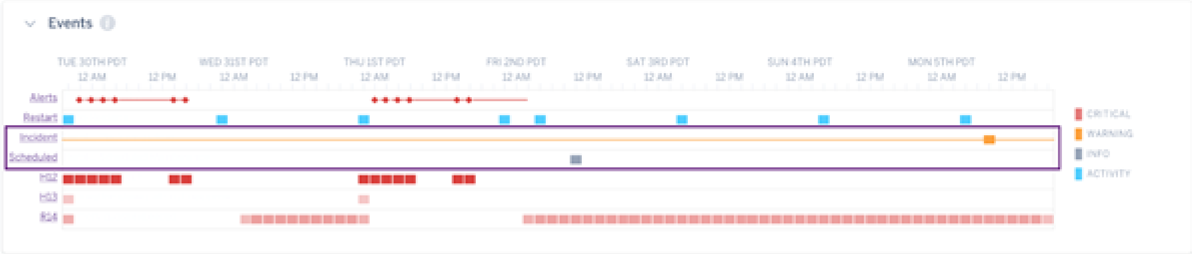

Changes in configuration variables are often reported as events with the shift in the event information shown in the variable.

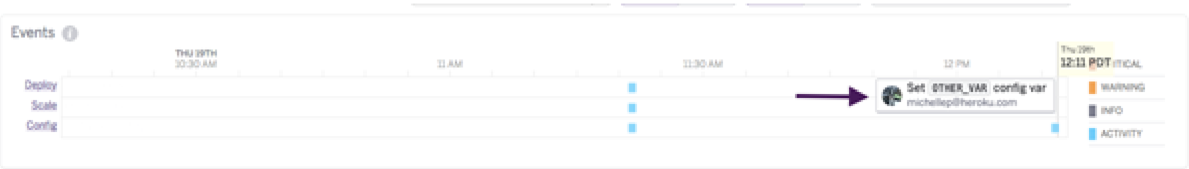

Set `OTHER_VAR` config var
michellep@heroku.com

Deploys will also be included in the list of events. The Metrics plots are expanded to cover deployment activities as markers which help users visualize the application health effect of deployments.

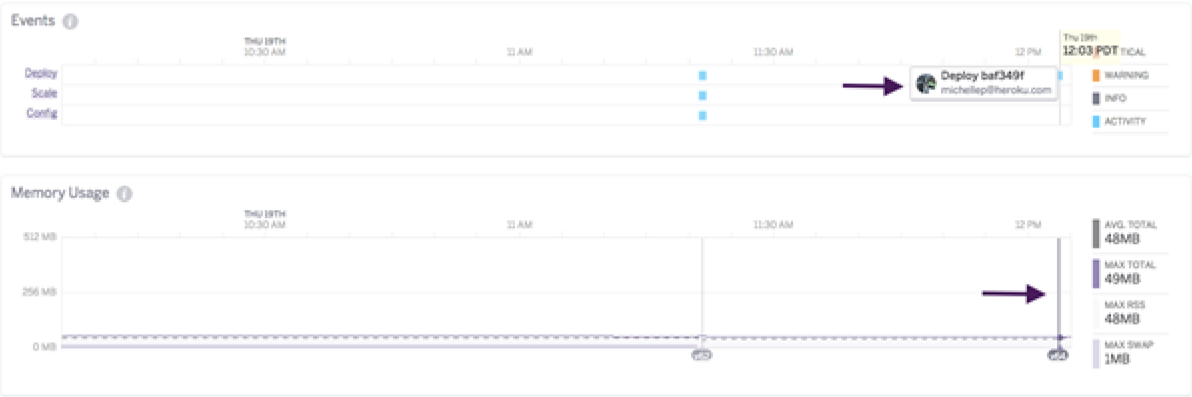

Deploy baf349f
michellep@heroku.com

Horizontal and/or vertical scaling events represent behaviors for dyno scaling.

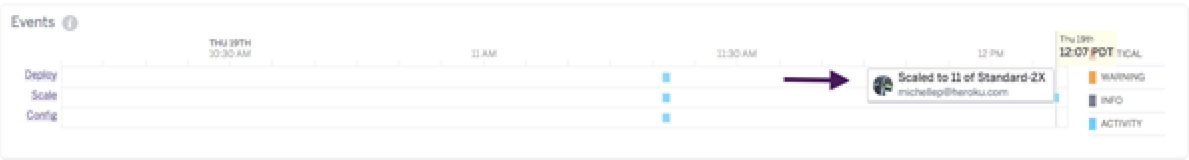

Scaled to 11 of Standard-2X
michellep@heroku.com

Heroku Dyno Monitoring with HG Heroku Add-on

As you can see, the depth of metrics that you can get with your default dyno monitoring metrics is not very significant. You need to go beyond that default, and dig deeper into your application in order to get useful monitoring.

The best tool for this is the HG Heroku add-on. Built on Hosted Graphite, Prometheus and Grafana, HG Heroku monitoring gives you great observability. You should check out the listing here, and install the add-on!