Got Logs? Get a PaperTrail: First thoughts

I stumbled upon Papertrail through a Twitter Ad (hey, those things work sometimes!) and figured that I should take a quick look. Given the amount of work I've been doing around compliance management and deployment of distributed systems, this seems like it may be an interesting fit. Luckily, they have a free tier as well which means it's easy to kick the tires on it before diving in with a paid commitment.

The concept seems fairly easy:

The signup process was pretty seamless. I went to the pricing page to see what the plan levels are which also has the Free Plan – Sign Up button nicely planted center of screen:

What I really like about this product is the potential to go by data ingestion rather than endpoints for licensing. Scalability is a concern with pricing for me, so knowing that the amount of aggregate data drives the price was rather comforting to me.

The free tier gets a first month with lots of data followed by a 100 MB per month follow on limit. That's probably not too difficult to cap out at, so you can easily see that people will be drawn to the 7$ first paid tier which ups the data to 1GB of storage and 1 year of retention. Clearly, at 7 days retention for the free tier, this is meant to just give you a taste and leave you looking for more if the usability is working for you.

First Steps and the User Experience

On completion of the first form, there is a confirmation email. You are also logged in immediately and ready to roll with the simple welcome screen:

Clicking the button to get started brings you to the instruction screen complete with my favorite (read: most despised) method of deploying which is pushing a script into a **sudo bash** pipe.

There is an option to run each script component which is much more preferred so you can see the details of what is happening.

Once you've done the initial setup process, you get a quick response showing you have active events being logged:

Basic logging is one thing for the system, so the next logical step is to up the game a bit and add some application level logging which is done using the remote-rsyslog2 collector. The docs and process to deploy are available inside the Papertrail site as well:

Now that I've got both by system and an application (I've picked the Apache error log as a source location) working, I'm redirected to see the live results in my Events screen (mildly censored to protect the innocent):

You can highlight some specific events and drill down into the different context views by highlighting and clicking anywhere in the events screen:

Searching the logs is pretty simple with a search bar that uses simple structured search commands to look for content. Searches are able to be saved and stored for reporting and repetitive use.

On the first pass, this looks like a great product and is especially important for you to think about as you look at how to aggregate logs for the purpose of search and retention for security and auditing.

The key will be making sure that you clearly define the firewall and VPC rules to ensure you have access to the remote server at Papertrail and then to make sure that you keep track of the data you need to retain. I've literally spent 15 minutes in the app and that was from first click to live viewing of system and application logs. All that and it's free too.

There is a referral link which you can use here if you want to try it out.

Give it a try if you're keen and let me know your experiences or other potential products that are freely available that could do the same thing. It's always good to share our learnings with the community!