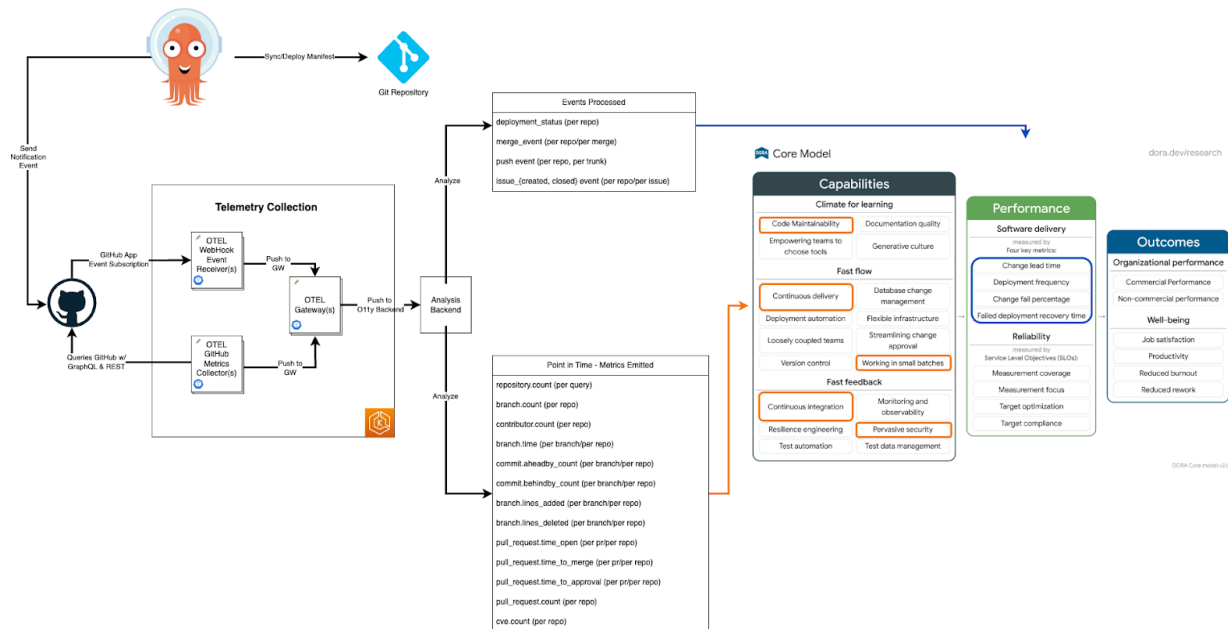


VCS Metrics and Visualization

Here's some information on visualizations created from the VCS metrics that are generated in the [GitHub Receiver](#). Largely, the metrics defined act as leading indicators that map to [DORA capabilities](#). Prior to the adjustments that have been made these metrics aligned closely with the [Engineering Effectiveness leading indicators](#). These are mentioned for historical context, providing information around how these metrics are interpreted.

The following diagram is a depiction of collecting telemetry from GitHub to populate DORA Capability and Four Key metrics.



These VCS metrics are point-in-time and scraped from the API. This gives a view of the current state. From there, events could be used to continue adding or removing from the metrics. When developing these metrics, we thought about how they would be used to tell a story. The story is human. Consider the following hypothetical scenario.

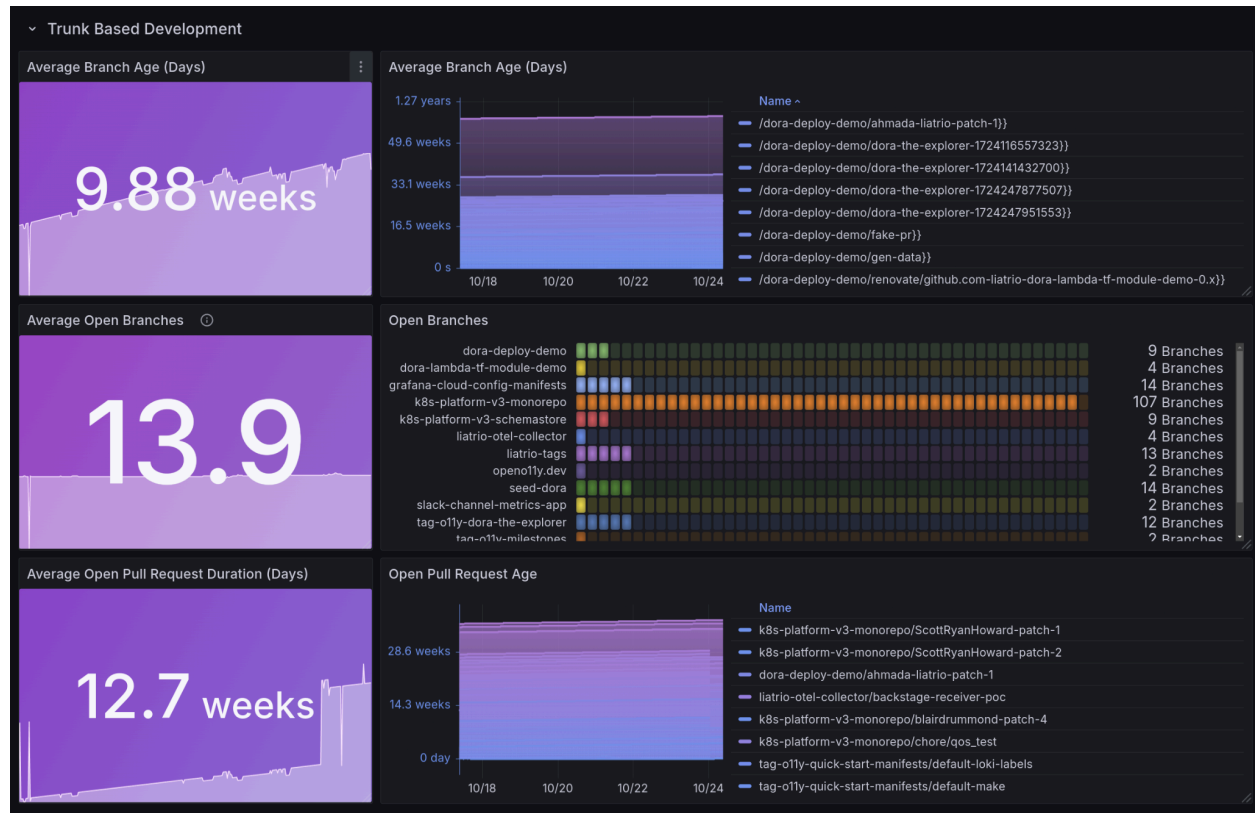
Technical leader Melissa noticed that her team had what felt like a high number of open pull requests (`vcs.change.count && vcs.change.state = open`) across services her team owns. She logically went to see what the average open time for pull requests was. From there she deduced that pull requests were sitting open for weeks at a time, contributing to the high open count. Why were pull requests open for so long? She knew that her team was small, consisting of only four engineers including herself. However, her team of four owned twenty-two repositories, consisting of six services. She realized this was a lot to maintain, but it didn't paint the entire picture. Her team certainly received inner-source contributions from other teams across the organization. Melissa then took a look at the average line addition and removal accounts for the various branches that existed. There she realized that each pull request included too many

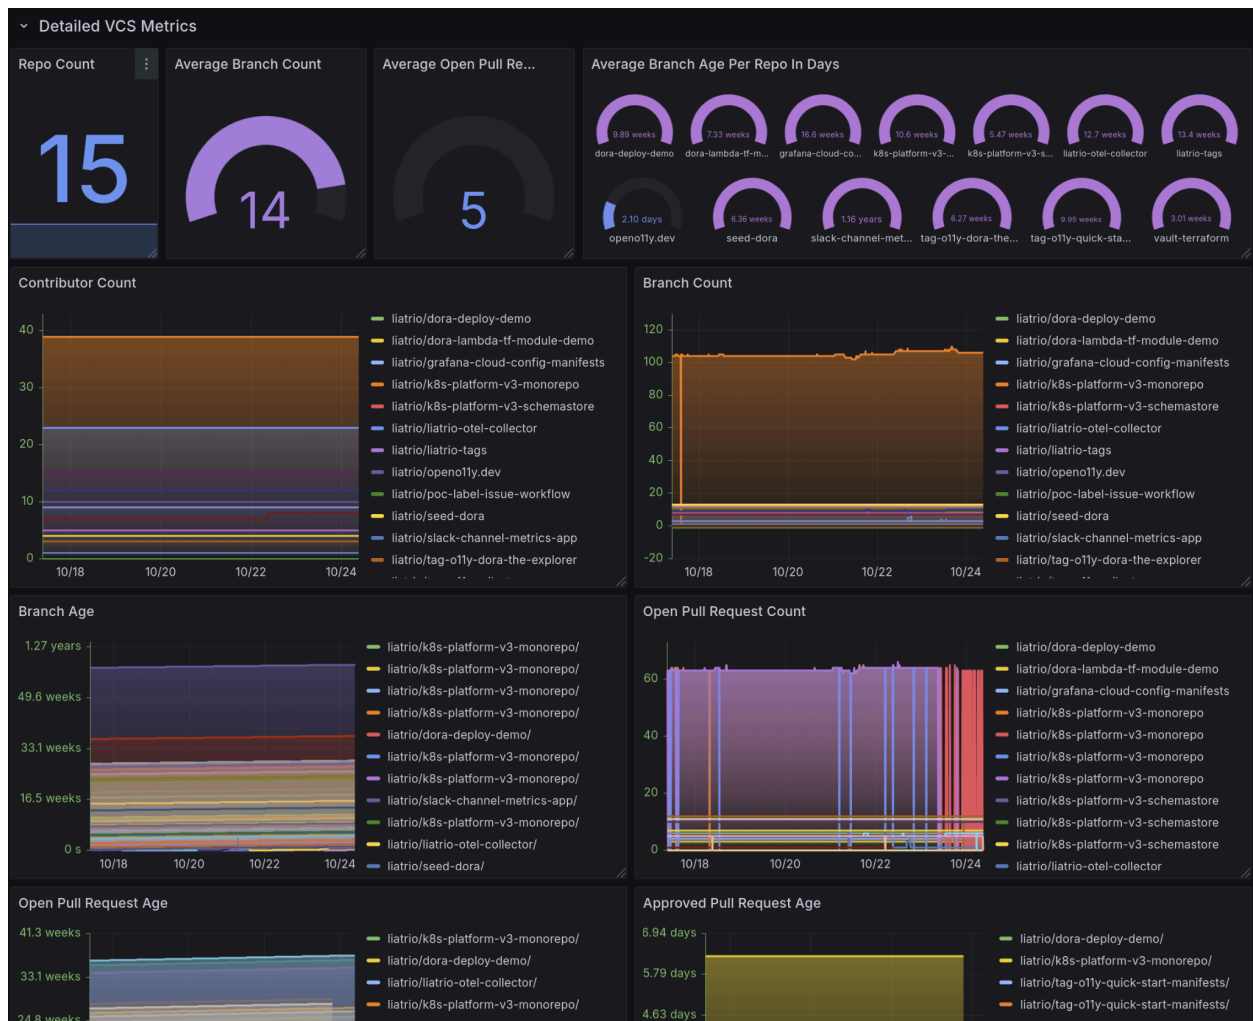
changes, making the pull requests hard to review. All these data points allowed her to reason that there was change needed to reduce cognitive load on the team, and accelerate delivery of value.

This was the original thinking behind the value of the metrics, measuring aspects of the DORA capabilities and providing the fundamental data points to create aggregations against. Each datapoint, often a simple gauge. This allows the user to create simple, but arguably effective aggregations through queries like (in prometheus as the example):

- This query gets the average branch age for repositories owned by a team.
 - `avg by(team_name) (vcs_repository_ref_time_seconds{organization_name=~"$organization", repository_name=~"$repo", team_name=~"$team"}) / 86400`
- This query gets the average branch count across repositories owned by a team.
 - `avg by(team_name) (vcs_repository_ref_count{organization_name=~"$organization", repository_name=~"$repo", team_name=~"$team"})`
- This query gets the average pull requests open time for repositories by team.
 - `avg by(team_name) (vcs_repository_change_time_open_seconds{organization_name=~"$organization", team_name=~"$team", repository_name=~"$repo"}) / 86400`
- This query gets the average branch time by repository.
 - `avg by(repository_name) (vcs_repository_ref_time_seconds{organization_name="$organization", repository_name=~"$repo", team_name=~"$team"}) / 86400`

The resulting dashboard would like as follows ([taken from the public DORA Capabilities Dashboard](#)):

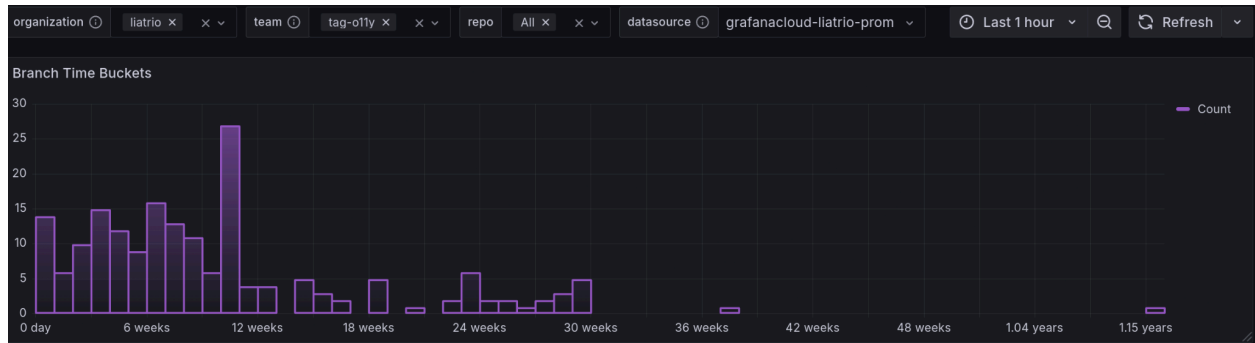





Histograms can be beneficial to detect outliers and distributions. From the fundamental gauge metric of time in seconds, one can create a simple histogram fairly easily. This may feel different, depending on the analysis backend, but in the examples above we're using Grafana Cloud, where the metrics are stored in Mimir (Prometheus). The query itself stays the same.

```
vcs_repository_ref_time_seconds{organization_name=~"$organization",
repository_name=~"$repo", team_name="tag-o11y"} / 86400
```

For the buckets, I created a bucket size of seven, and the unit is represented in days. Therefore each bucket is a week. I ran one transformation on the data which was a **reduce** function to reduce the series to rows, taking the max of the value. This way I'm not counting multiples of the same metrics. The rendering of the visualization is seen below.



Now, because this is all based on attributes, you can pair down to see distributions by repository simply changing the variable.



To get the resolution though and analyze further, one would want to dive deeper by looking at the fundamental telemetry prior to aggregation.

Summary

Based on the fundamental metric types, it's easy enough to query for the information you want, and still render distributions. The dashboards shown here are publicly available through the Grafana Dashboard marketplace and get updated as the metric conventions are updated.