

Counting Logs

Another exciting feature of LogQL is that it can envelop a query and allow the log stream selector to count entries per stream.

Range Vector Aggregation

LogQL range vector is unique from other log language queries because the selected range of samples usually includes a value of 1 for each entry. Hence, the aggregate can envelop a selected range into an instance vector. The range vector supported functions are:

- `rate`: calculate the number of entries per second
- `count_over_time`: counts the entries for each log stream within the given range.
- `count_over_time({job="MySQL"}[5m])`

For instance, the range vector counts all the log lines among the first 5 minutes for the MySQL query. This example showcases the per-second rate of all non timeout errors

rate:

```
(( {job="mysql"} |= "error" != "timeout")[10s] )
```

Aggregation operators

LogQL aggregation operators act similarly to PromQL, which supports a subset that can aggregate the elements of a single vector. This action results in a new vector of fewer elements with aggregated values.

- `sum`: Calculate the sum over labels
- `min`: Select minimum over labels
- `max`: Select maximum over labels
- `avg`: Calculate the average over labels
- `stddev`: Calculate the population standard deviation over labels
- `stdvar`: Calculate the population standard variance over labels
- `count`: Count the number of elements in the vector
- `bottomk`: Select the smallest k elements by sample value
- `topk`: Select largest k elements by sample value

Also, by using a *without* or *by* clause, whether a *without* or *by* clause, the log operators can use the aggregation operators to aggregate over all label values or a set of distinct values.

For instance

```
<aggr-op>([parameter,] <vector expression>)  
[without | by (<label list>)]
```

Note. Using *parameters* is only needed when using `topk`, `bottomk`, and `topk`, considered group vectors. Aggregators such as `bottomk` are different from other aggregators within the sample subset.

The *without* clause removes listed labels from results, while the *by* clause does the opposite, unlisting labels that are not part of the clause, regardless of label identity.

For instance: filter logs for the top 10 applications by highest log throughput;

```
topk(10,sum(rate({region="us-east1"}[5m])) by  
(name))
```

Get the count of logs during the last five minutes, grouping by level:

```
sum(count_over_time({job="mysql"}[5m])) by (level)
```

Get the rate of HTTP GET requests from NGINX logs:

```
avg(rate(({job="nginx"} |= "GET")[10s])) by (region)
```