This page describes how to use the metric-selection tool to specify a target metric for an alerting policy. The chart next to the **Target** region gives you visual feedback on the data being captured by the target.

The **Target** region uses the same metric selector that is used in Metrics Explorer and for creating charts. If you are already familiar with it, you can skip this page.

To select a metric, use the **Find resource type and metric** field to choose one resource type and one metric type. You can specify them in either order. To begin, click in the field. This brings up one or two lists, based on any prior selections. The lists are indicated by headers, **Resource types** and **Metrics**, as seen in the following screenshot:

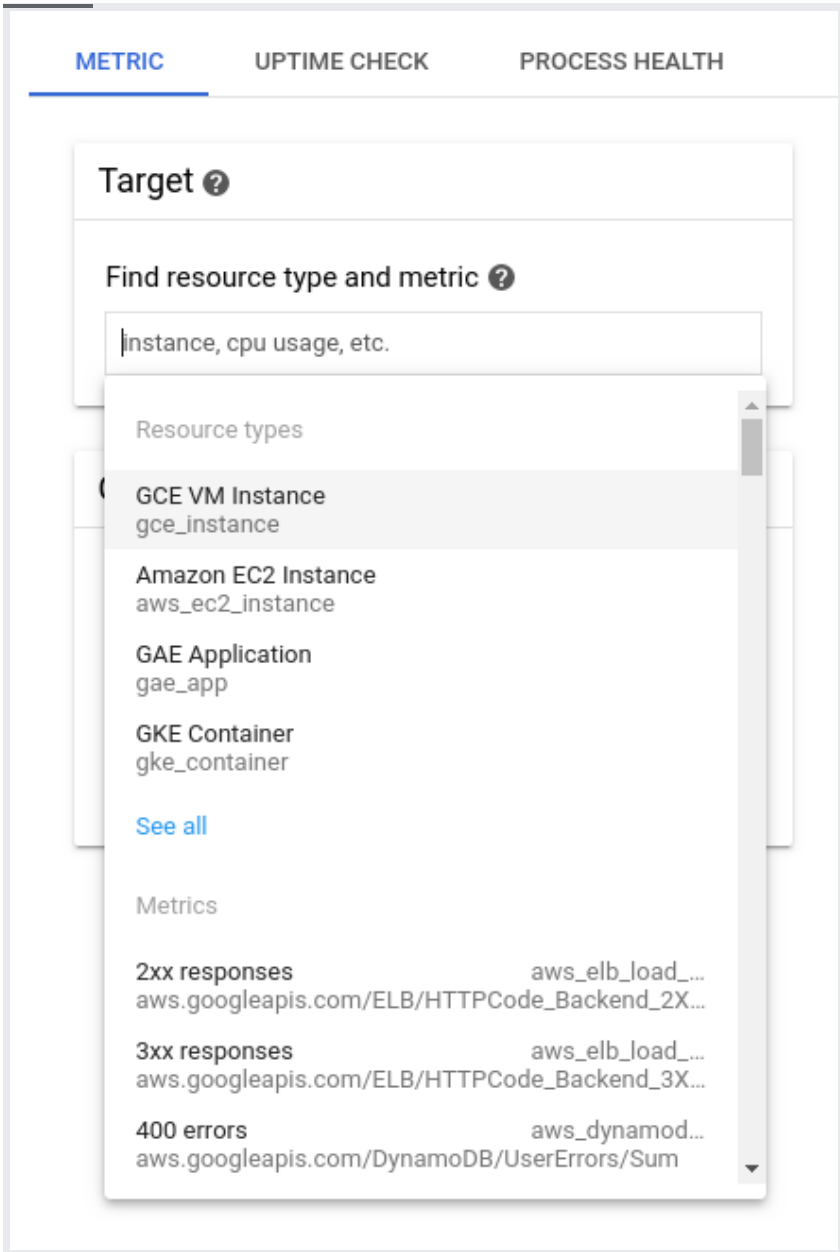

You cannot select metric types with values of type **`String`** from the console. String-valued data can't be displayed or or used in alerting policies. To retrieve data for these metric types, you must use a method in the API, like **`eries.list`** (/monitoring/api/ref_v3/rest/v3/projects.timeSeries/list).

You can select an entry in two ways:

- By selecting entries from the lists.

Because there fewer resource types than metrics, select the **Resource type** first. If your resource type isn't displayed in the short list, click **See all**. After you select the resource type, the metric list is automatically

filtered to display only relevant metrics.

- By entering a metric filter. To enter a metric filter, do the following:

  1. Next to **Find resource type and metric**, click **Help** ⓘ
  2. Click **Direct filter mode** in the help pane.

     When **Direct filter mode** is enabled, the **Find resource type and metric** option is replaced with an editable text box labeled **Resource type, metric, and filter**:

     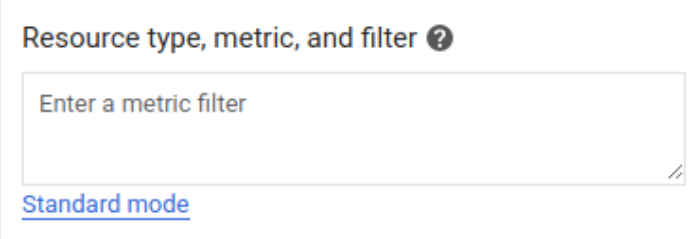

     If you made selections for a resource type, a metric, or a filter prior to selecting **Direct filter mode**, then those settings are used to prepopulate the **Resource type, metric, and filter** text box.
  3. Enter a metric filter in the **Resource type, metric, and filter** text box. Your filter must include a metric type and a resource type. You can also include label filters. For the filter grammar, see Monitoring filters (/monitoring/api/v3/filters).

     For example, to display the log entries for all Google Cloud VM instances in the `us-east1-b` zone, enter the following:

     If you've used direct filter mode to configure charts or alerting policies and no data is available, then the chart displays an error message. The exact error message is dependent on the filter you entered. For example, a typical message is `Chart definition invalid`. You might also see the message `No data is available for the selected time frame.`

Hovering over an item on either list brings up a tooltip that displays the information in the item's descriptor. For information on descriptors for metric types or monitored resources, see the metrics list (/monitoring/api/metrics) or the monitored resource list (/monitoring/api/resources).

When at least one resource type and metric pair is selected, the chart shows all the available time series, and additional items appear below the specifed metric on the **Metric** tab. The following screenshot shows the **Metric** tab after a metric has been specified:

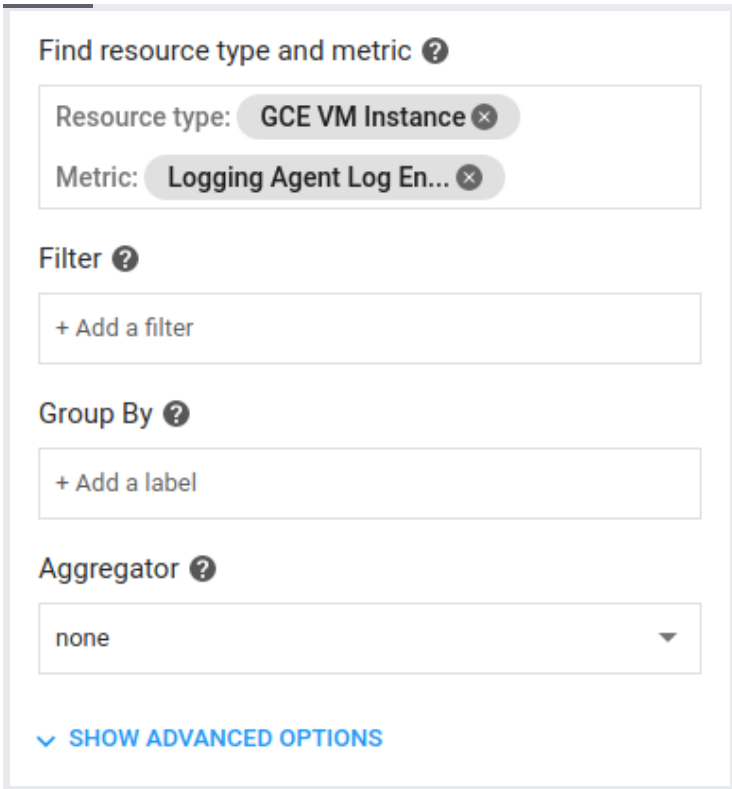

You can reduce the amount of data returned for a metric by specifying filter criteria, so that only time series that meet some set of criteria are used. Filtering removes some data from the chart by excluding time series that don't meet the filtering criteria. This results in fewer lines on the chart that displays this metric, which can improve the performance of the chart.

You can supply multiple filtering criteria. The corresponding chart shows only the time series that meet *all* of the criteria, a logical `AND`.

When you click in the **Filter** field, a panel containing lists of criteria by which you can filter appears. In broad strokes, you can filter by resource group (/monitoring/groups), by name, by resource label, and by metric label.

The following screenshot shows the known filter-by labels for a project:

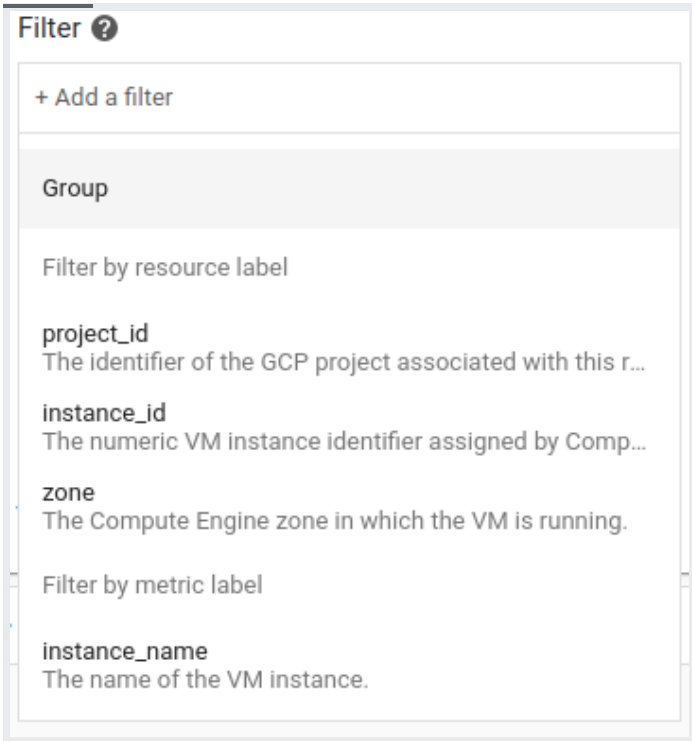

You can select from the lists or type to find matches. Additionally, you can create filters for data that has not yet appeared; such filter criteria won't appear on the selection list, but you can manually specify filters that you know will be valid in the future.

After you choose a label on which to filter, you have to specify the rest of the filter: a value or range of values and a comparison.

For example, the following screenshot shows a filter on the `zone` resource label. The **Filter** field supports a pair of comparison operators for equality, `=` and `=~`, and a pair for inequality, `!=` and `!=~`. The second item in each pair takes a regular-expression as a value. Simple equality, `=`, is the default.

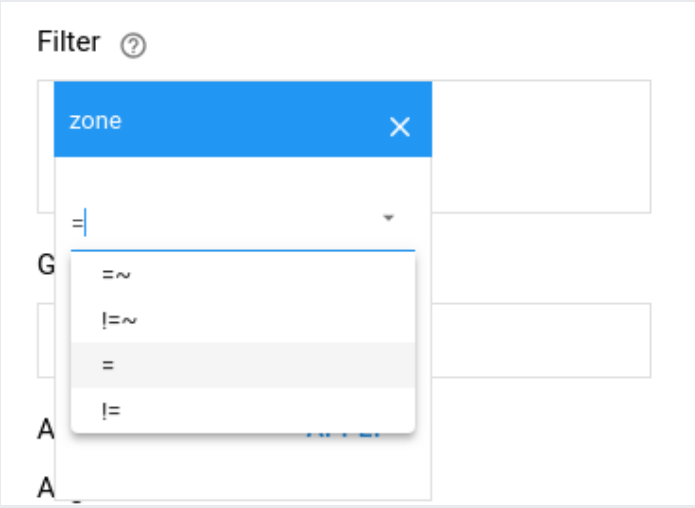

Below the list of comparison operators is a list of the available values. The following screenshot shows the names of zones in the project:

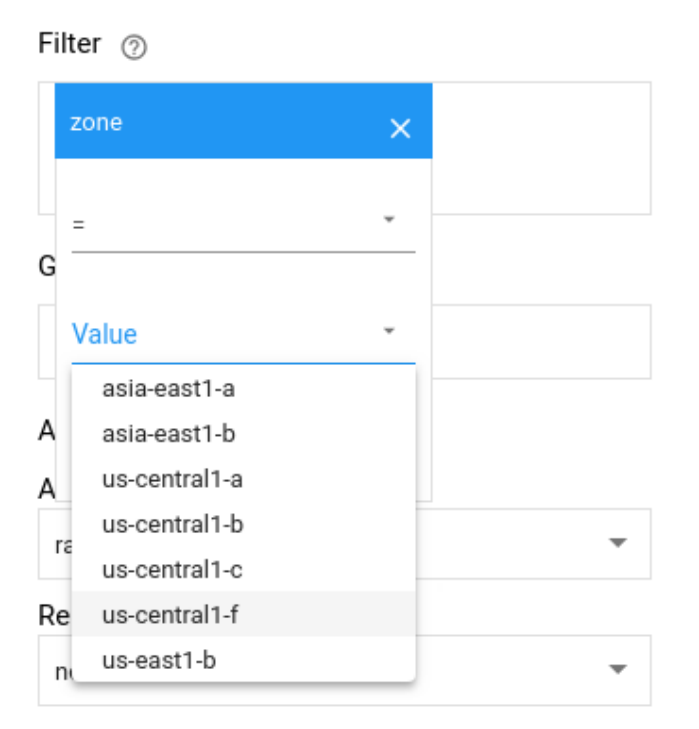

For the **Value** field, you can select one of the items on the drop-down list, or you can enter an expression that matches multiple items:

- If you use a direct comparison, `=` or `!=`, you can create a filter string like `starts_with`. For example, the filter string `starts_with("us-central")` matches any `us-central` zone:

  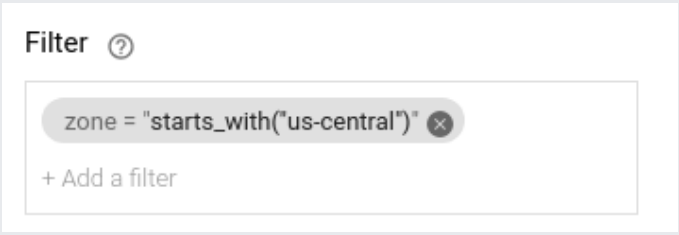

  See Monitoring filters (/monitoring/api/v3/filters#filter_syntax) for more information on filter strings.

- If you select `=~` or `!=~` you can use a regular expression in the value. For example, the regular expression `us-central1-.*` matches any `us-central1` zones:

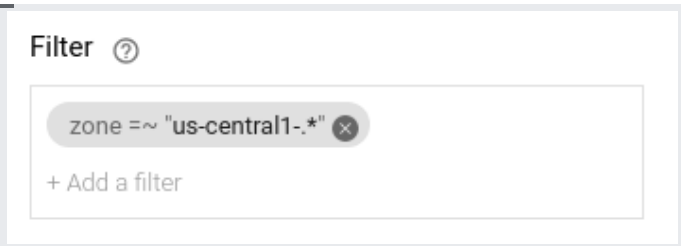

The regular expression `^us.*.a$` matches any US zone that ends with “a”:

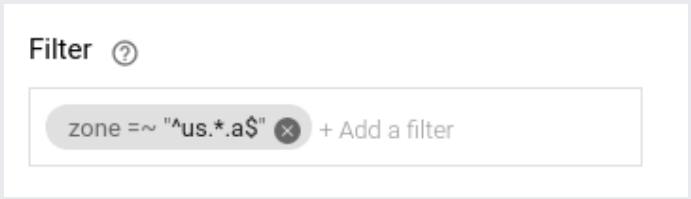

You can specify multiple filter criteria, and you can use the same label multiple times. This lets you specify a filter for a range of values. To add additional filters, click **Add a filter** near the bottom of the filter field. Currently, all of the filter criteria must be met; they constitute a logical `AND`. For example, you can use both `starts_with` and `ends_with` filter strings to show only “a” zones in the US:

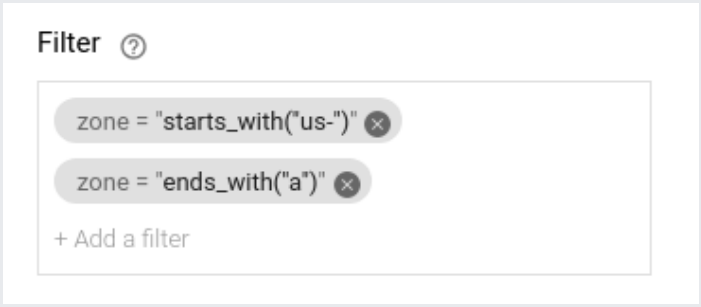

You can reduce the amount of data returned for a metric by combining data from similar time series.

Like filtering, grouping is way to reduce the amount of data you are manipulating. Filtering works by excluding some time series, and grouping works by identifying sets of time series that all meet some criterion and then combining, or *aggregating*, the members of the sets together.

Grouping is done on the basis of labels. When you click in the **Group By** field you get the same labels that are available for filtering.

The **Group By** option lets you group time series by resource and metric labels, and then combine the data within those groups. This creates a single new time series for each combination of group-by values, and the new time series represents all the members of the group.

Suppose, for example, you want to examine logging load but only at a regional level; that is, you don't care about the number of entries from specific servers. You can group time series by zone, and then chart the total, average, or other measure for each zone.

Group By chooses a default aggregation function based on the type of data being grouped, but you can change this selection. See Aggregation (#aggregation-options) for information about other choices.

The following screenshot shows a grouping by zone. The choice of `sum` under **Aggregator** is the default.

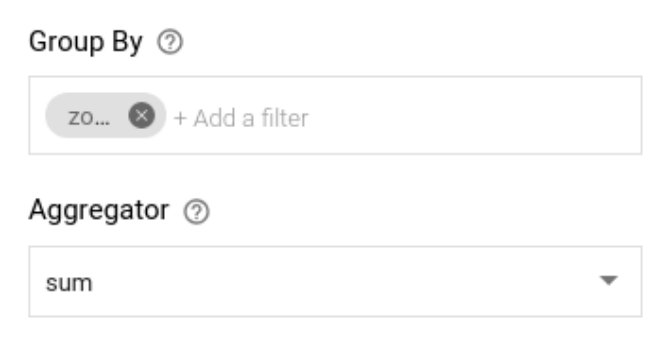

This creates a new time series for each zone, with values computed from the sum of all the values for individual time series in that zone.

You can group by multiple labels, as well. When you group by multiple labels, you get a time series for each combination of labels. The order in which you specify the labels does not matter.

To remove a group-by condition, you must:

1. Delete the group-by labels.
2. Set the aggregation method back to **none**.

If you only delete the group-by labels, your chart will display a single line representing all of the data after the aggrega
en applied. To undo grouping, you must remove the group-by labels *and* reset the aggregation to **none**.

The **Aggregator** option lets you combine time series by using common functions. This results in fewer lines on the chart displaying the metric, which can improve the performance of the chart.

Click in the **Aggregator** field to see a list of the available aggregation options. These are the functions, or *reducers*, that can be used to combine the time series.

The available reducing functions depend on the type of values the metric captures, but they commonly include choices like mean, max or min, standard deviation, assorted percentile values, and so forth. For more information about these dependencies, see Metrics, time series, and resources (/monitoring/api/v3/metrics).

When used without grouping, the reducing function is applied across all the selected time series, combining them to a single time series consisting of the mean, sum, or other measure as calculated across all the time series. When used with grouping, the function is applied to the time series within each group.

For more information on aggregation, see `Aggregation` (/monitoring/api/ref_v3/rest/v3/projects.alertPolicies#Aggregation) in the API reference.

The **Group By** option automatically applies aggregation to compute statistics within each group. The lines on a **Group By** chart already represent aggregations. **Group By** chooses a default aggregation function based on the type of data being grouped, but you can change this selection.

Time series must be aligned before they can be aggregated. Both the **Aggregator** and **Group by** options use a default ıent strategy, but you can change this. See Alignment (#alignment) for more information on this.

You can also apply aggregation to a set of time series that you have filtered. As with the unfiltered time series, aggregation of filtered time series will reduce all of the lines to one that reflects the chosen aggregation function, unless you also use grouping.

There are additional options for selecting a metric. These options have default values, but you can expose these options and override the defaults. To see the additional options, click **Show advanced options**.

A time series is a set of data points in temporal order. To align a time series is to break the data points into regular buckets of time, the *alignment period*. Multiple time series must be aligned before they can be combined.

Alignment is a prerequisite to aggregation across time series, and it is applied to each time series individually. Because alignment is a prerequisite for charting the data, Stackdriver does it automatically, by using default values. You can override these defaults by using the alignment options, **Aligner** and **Alignment Period**:

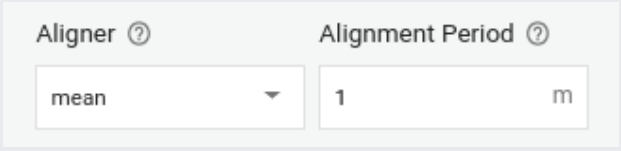

Alignment Period: The alignment period determines the length of time for subdividing the time series. For example, you can break a time series into one-minute chunks or one-hour chunks. The data in each period is summarized so that a single value represents that period. The default alignment period is one minute.

Although you can set the alignment interval for your data, time series might be realigned when you change the time interval displayed on a chart or change the zoom level.

Aligner: The aligner is a function that determines how to summarize the data in each alignment period. Aligners include the sum, the mean, and so forth. Valid aligner choices depend on the kind and type of metric data a time series stores. That is, aligner choice depends on the `MetricKind` (/monitoring/api/ref_v3/rest/v3/projects.metricDescriptors#MetricKind) and `ValueType` (/monitoring/api/ref_v3/rest/v3/projects.metricDescriptors#ValueType) of the time series.

Some aligners both align the data and convert it from one metric kind or type to another. For more information on the available aligners, see `Aligner` (/monitoring/api/ref_v3/rest/v3/projects.alertPolicies#Aligner) in the API reference.

When you have multiple time series that already represent aggregations, like the examples illustrating the **Group By** option, you can then aggregate across them by choosing a **Secondary Aggregator**:

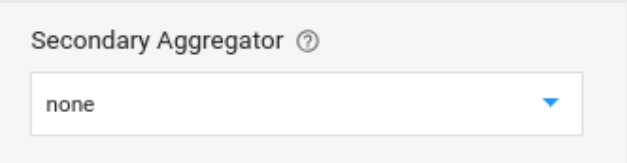

Secondary aggregation reduces all the time series on the chart to a single time series.

The **Legend Template** field lets you customize a description for the time series on your chart. These descriptions appear on the hover card for the chart and on the chart legend in the **Name** column.

By default, these descriptions are created for you from the values of different labels in your time series. Because the system selects the labels, the results might not be helpful to you. You can use

this field to build a template for the descriptions.

The **Legend Template** field accepts the following:

- Plain text.
- Variables that represent labels in your time series. The available labels depend on the time series you are charting, but the variable syntax generally look similar to this: `${resource.labels.project_id}`. The `${}` notation causes the variable to be evaluated in your description.

If you supply only text, then the descriptions of the time series on the chart will be identical. Using variables lets you pull values from labels in the time series into the description.

If you know the variable names, you can manually type them into the template field. You can also select variables for the available labels by using the **+ Add a filter** widget in the field. This approach ensures that the variable syntax is correct.

The following screenshot shows a template consisting of a string and a single variable:

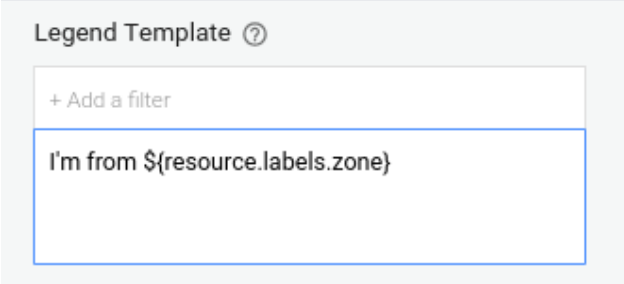

In the chart legend, the values generated from the template appear in a column with the header **Name**. The values also appear on the hover card.

The following screenshot shows the **Name** values and hover card that result from the template:

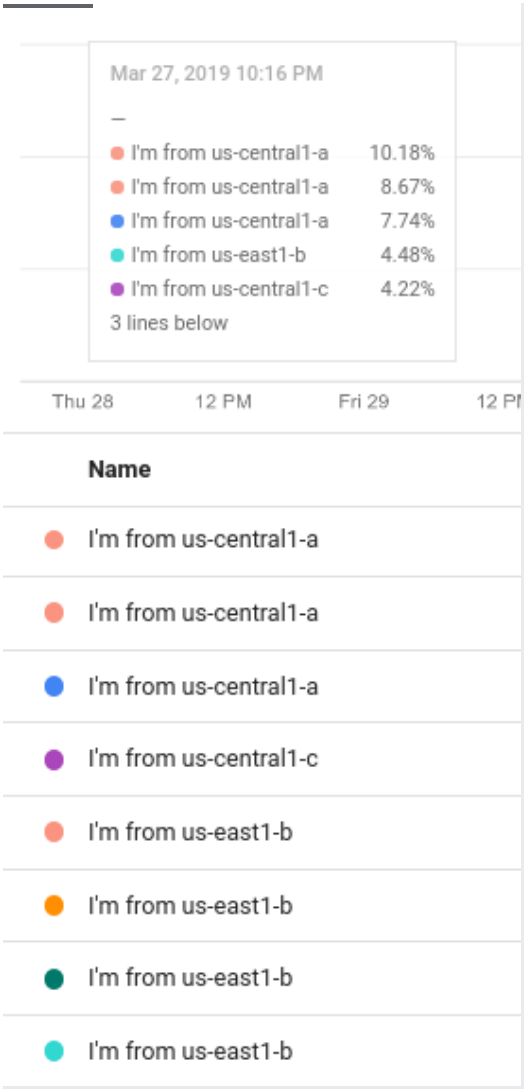

You can create templates that use multiple strings and variables, but the display space available on the hover card is limited.