To create an alerting policy, you must describe what is to be monitored, when the alerting policy is triggered, and how you want to be notified. This page contains settings that you can use to create alerting policies. Each section in this page has the following elements:

- **Title**: Lists the relevant product name and a brief description of the alerting policy.
- **Summary**: A brief description of the alerting policy. For full information, see the product documentation.
- **Steps to create an alerting policy**: Outline of the steps required to create an alerting policy. For detailed information on these steps, see Creating an alerting policy (/monitoring/alerts/using-alerting-ui#create-policy).
- **Target pane fields**: These fields specify what is being monitored and how the data is aggregated. If you are using the Google Cloud Console, these are the fields in the **Target** pane of the **Add Condition** dialog.
- **Configuration fields**: These fields specify when the alerting policy triggers. If you are using the Google Cloud Console, these are the fields in the **Configuration** pane of the **Add Condition** dialog.

You can use the settings in the target pane table when creating a chart or when using Metrics Explorer (/monitoring/charts/metrics-explorer).

To create an alerting policy that triggers when the 50th percentile of the execution time of a BigQuery (/bigquery/docs) query exceeds a user-defined limit, use the following settings:

Target pane Field	Value
Resource type	Global
Metric	Query execution time
Filter	
Aggregator	none
Advanced Aggregation	Aligner: 50th percentile Alignment Period: 1 m

Conditions pane Field	Value
Condition triggers if	Any time series violates
Condition	is above

Conditions pane Field	Value
`Threshold`	You determine the acceptable value. For this metric, consider a threshold that is about double the average value for the 50th percentile.
`For`	`1 minute`

To create an alerting policy that triggers when the ingested BigQuery (/bigquery/docs) metrics exceed a user-defined level, do the following:

Target pane Field	Value
`Resource type`	Enter `BigQuery Dataset`
`Metric`	Metrics specific to usage include `Stored bytes`, `Uploaded bytes`, and `Uploaded bytes billed`. However, this in only a partial list. For a full list of available metrics, see BigQuery metrics (/monitoring/api/metrics_gcp#gcp-bigquery).
`Filter`	**project_id**: Your Google Cloud project ID. **dataset_id**: Your dataset ID.
`Group By`	**dataset_id**: Your dataset ID.
`Aggregator`	`sum`
`Advanced Aggregation`	`Aligner`: `mean` `Alignment Period`: `1 m`

Conditions pane Field	Value
`Condition triggers if`	`Any time series violates`
`Condition`	`is above`
`Threshold`	You determine the acceptable value.
`For`	`1 minute`

Early Boot Validation shows the pass/fail status of the early boot portion of the last boot sequence. Early boot is the boot sequence from the start of the UEFI firmware until it passes control to the bootloader.

To create an alerting policy that triggers when the early boot sequence fails for any of your Compute Engine (/compute/docs) VM instances, use the following settings:

Target pane

Field	Value
Resource type	GCE VM Instance
Metric	Early boot validation
Filter	status = failed

Target pane Field	Value
Group By	status
Aggregator	sum
Advanced Aggregation	Use defaults.

Conditions pane Field	Value
Condition triggers if	Any time series violates
Condition	is above
Threshold	0
For	1 minute

Late Boot Validation shows the pass/fail status of the late boot portion of the last boot sequence. Late boot is the boot sequence from the bootloader until completion. This includes the loading of the operating system kernel.

To create an alerting policy that triggers when the late boot sequence fails for any of your Compute Engine (/compute/docs) VM instances, use the following settings:

Target pane Field	Value
Resource type	GCE VM Instance
Metric	late boot validation
Filter	status = failed
Group By	status
Aggregator	sum
Advanced Aggregation	Use defaults.

Conditions pane Field	Value
Condition triggers if	Any time series violates
Condition	is above
Threshold	0

Conditions pane Field	Value
For	1 minute

To create an alerting policy that triggers when your monthly log bytes ingested exceeds your user-defined limit for Stackdriver Logging (/logging/docs), use the following settings:

Target pane Field	Value
Resource type	Global
Metric	Monthly log bytes ingested
Filter	
Aggregator	sum
Advanced Aggregation	Aligner: max Alignment Period: 60m

Conditions pane Field	Value
Condition triggers if	Any time series violates
Condition	is above
Threshold	You determine the acceptable value.
For	Minimum acceptable value is 30 minutes.

To set up a Recommendations AI (/recommendations-ai/docs) prediction alert, use the following settings in the alerting policy:

Target pane Field	Value
Resource type	Consumed API
Metric	Request count
Filter	service = recommendationengine.googleapis.com method = google.cloud.recommendationengine.v1beta1. PredictionService.Predict response_code != 200
Aggregator	sum
Advanced Aggregation	Aligner: sum Alignment Period: 1m

Conditions pane Field	Value
Condition triggers if	Any time series violates
Condition	is above

Conditions pane Field	Value
Threshold	0
For	5 minutes

To set up a Recommendations AI (/recommendations-ai/docs) event recording reduction alert, use the following settings in the alerting policy:

Target pane Field	Value
Resource type	Consumed API
Metric	Request count
Filter	service = recommendationengine.googleapis.com method = google.cloud.recommendationengine.v1beta1.PredictionService.CollectUserEvent response_code != 200
Aggregator	sum
Advanced Aggregation	Aligner: sum Alignment Period: 1m

Conditions pane Field	Value
Condition triggers if	Any time series violates
Condition	is absent
Threshold	
For	10 minutes

To create an alerting policy that triggers when your high priority cpu utilization for Cloud Spanner (/spanner/docs) is above a recommended threshold, use the following settings:

Target pane Field	Value
Resource type*	Cloud Spanner Instance
Metric	CPU Utilization by priority
Filter	instance_id = *YOUR_INSTANCE_ID* priority = high
Aggregator	max
Advanced Aggregation	Aligner: mean Alignment Period: 10 m

* Selecting this resource and metric type is equivalent to entering the following value in the **Find resource type and metric** pane: `spanner.googleapis.com/instance/cpu/utilization_by_priority`.

Conditions pane Field	Value
`Condition triggers if`	`Any time series violates`
`Condition`	`is above`
`Threshold`	**`45%`** for multi-region instances; **`65%`** for regional instances.
`For`	`10 minutes`

To create an alerting policy that triggers when the 24 hour rolling average of your cpu utilization for Cloud Spanner (/spanner/docs) is above a recommended threshold, use the following settings:

Target pane Field	Value
Resource type*	Cloud Spanner Instance
Metric	Smoothed CPU utilization
Filter	instance_id = *YOUR_INSTANCE_ID*
Aggregator	sum
Advanced Aggregation	Aligner: mean Alignment Period: 10 m

* Selecting this resource and metric type is equivalent to entering the following value in the **Find resource type and metric** pane: `spanner.googleapis.com/instance/cpu/smoothed_utilization`.

Conditions pane Field	Value
Condition triggers if	Any time series violates
Condition	is above
Threshold	90%
For	10 minutes

To create an alerting policy that triggers when your storage for your Cloud Spanner (/spanner/docs) instance is above a recommended threshold, use the following settings:

Target pane Field	Value
Resource type*	Cloud Spanner Instance
Metric	Storage used
Filter	instance_id = *YOUR_INSTANCE_ID*
Aggregator	sum
Advanced Aggregation	Aligner: max Alignment Period: 10 m

* Selecting this resource and metric type is equivalent to entering the following value in the **Find resource type and metric** pane: `spanner.googleapis.com/instance/storage/used_bytes`.

Conditions pane Field	Value
`Condition triggers if`	`Any time series violates`
`Condition`	`is above`
`Threshold`	Set the threshold to 75% of the maximum storage per node, multiplied by the number of nodes. For the current node limits, see Cloud Spanner Quotas and limits (/spanner/docs/limits). For a 2 TB limit per node, the recommended threshold is: **1649267441664** multiplied by the number of nodes in your instance.
`For`	`10 minutes`

To create an alerting policy that triggers when your monthly Stackdriver Trace (/trace/docs) spans ingested exceeds your quota, use the following settings:

Target pane Field	Value
Resource type	Consumed API
Metric	Request count
Filter	service = cloudtrace.googleapis.com response_code = 429
Aggregator	sum
Advanced Aggregation	Aligner: sum Alignment Period: 1m

Conditions pane Field	Value
Condition triggers if	Any time series violates
Condition	is above
Threshold	0
For	1 minute

To create an alerting policy that triggers when your monthly Stackdriver Trace (/trace/docs) spans ingested exceeds a user-defined limit, use the following settings:

Target pane

Field	Value
Resource type	global
Metric	Monthly trace spans ingested
Filter	

Target pane Field	Value
Aggregator	sum
Advanced Aggregation	Aligner: max Alignment Period: 60m

Conditions pane Field	Value
Condition triggers if	Any time series violates
Condition	is above
Threshold	You determine the acceptable value.
For	Minimum acceptable value is 30 minutes.

To create an alerting policy for an uptime check, or to create a chart that displays the success or latency status of an uptime check, see Alerting on uptime checks (/monitoring/uptime-checks/uptime-alerting-policies).