

Providing support in **SuperQuota**

AGENDA

CASE 01 **Internal management quota requests**

CASE 02 **Org level advanced warnings**

CASE 03 **Solving short term rate quotas**

CASE 04 **Safe override**

Foundational **Research**



SuperQuota's mission is to provide a **central infrastructure** for **viewing, managing, and understanding quota usage** across all **Google APIs** and **services** for the **Google Cloud Platform**.

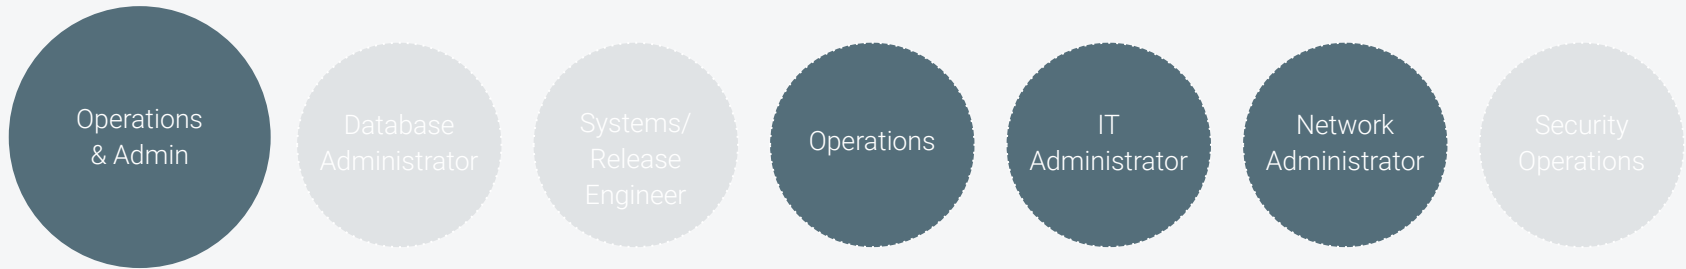


A quota is defined as having a **metric**, **unit**, **region** and **zone** combination.

USER ROLES

Tracks **infrastructure** and **resources**.

Responds to problems such as **quotas** reaching their **limits**.



CASE 01. Internal management quota requests

How can we **provide** a way for organizations (who are using services through adoption of org-level quotas) to handle **internal requests for quota increases** for **projects**?

Why is this a problem?

Today, **project owners** can only see a message telling them to “**contact your organization’s quota administrator**”.

But in a **large org**, that person **might not know** who to contact.

This is a **blocker** for adoption of org-level quotas and a **pain** for **Cloud Support** (they don’t know contact info for quota admins and don’t want to file support tickets).

There needs to be a **mechanism** to notify those in the **organization** with **adequate permissions** that a **project** needs **more quotas**.

CONSTRAINTS

* **AVOID AVOID AVOID** building a ticketing system.

SOLUTION INTEGRATING CLOUD FEED

What is Cloud Feed?

Comprehensive repository of all GCP events and activities

Easy fan-in for GCP teams to share activities with users

Human-understandable schema, including call to action

Human-driven configuration with notification and opt-in/out control

Why Cloud Feed?

1. Cloud Feed does not require **sender** (project owner) to know who the **recipients** (i.e. quota admins) are

2. Project owner can request several quota increases

The screenshot shows the Google Cloud Platform interface for managing quotas for project "foo-bar". The main table lists various quota types with their current and peak usage. Two quotas, "Networks" and "Routes", are selected for modification. A right-hand panel shows the configuration for these two selected quotas, including a warning about admin approval for increases above the limit and input fields for the new limits.

Quota [+ EDIT QUOTAS](#) [VIEW REQUESTS](#)

Quotas for project "foo-bar"
View, manage and understand all your quota usages across all APIs and services. Filter by quota type, service, metric and region and their zones.

Quota type: All quotas | Service: All services | Metric: All metrics | Region: All regions

<input type="checkbox"/> Service	Location	Current Usage & 7 Day Peak Usage	Current Usage	7 Day Peak
<input type="checkbox"/> Firewall rules Compute Engine API	us-west-1a		48	95
<input checked="" type="checkbox"/> Networks Compute Engine API	us-west-1		4	4
<input checked="" type="checkbox"/> Routes Compute Engine API	Global		100	120
<input type="checkbox"/> Number of indexes per day App Engine	us-west-1b		20	120
<input type="checkbox"/> Operation read requests per 100 seconds per use Compute Engine API	us-east-1		0	1,600

2 quotas selected

Google Compute Engine

Quota group for Routes *

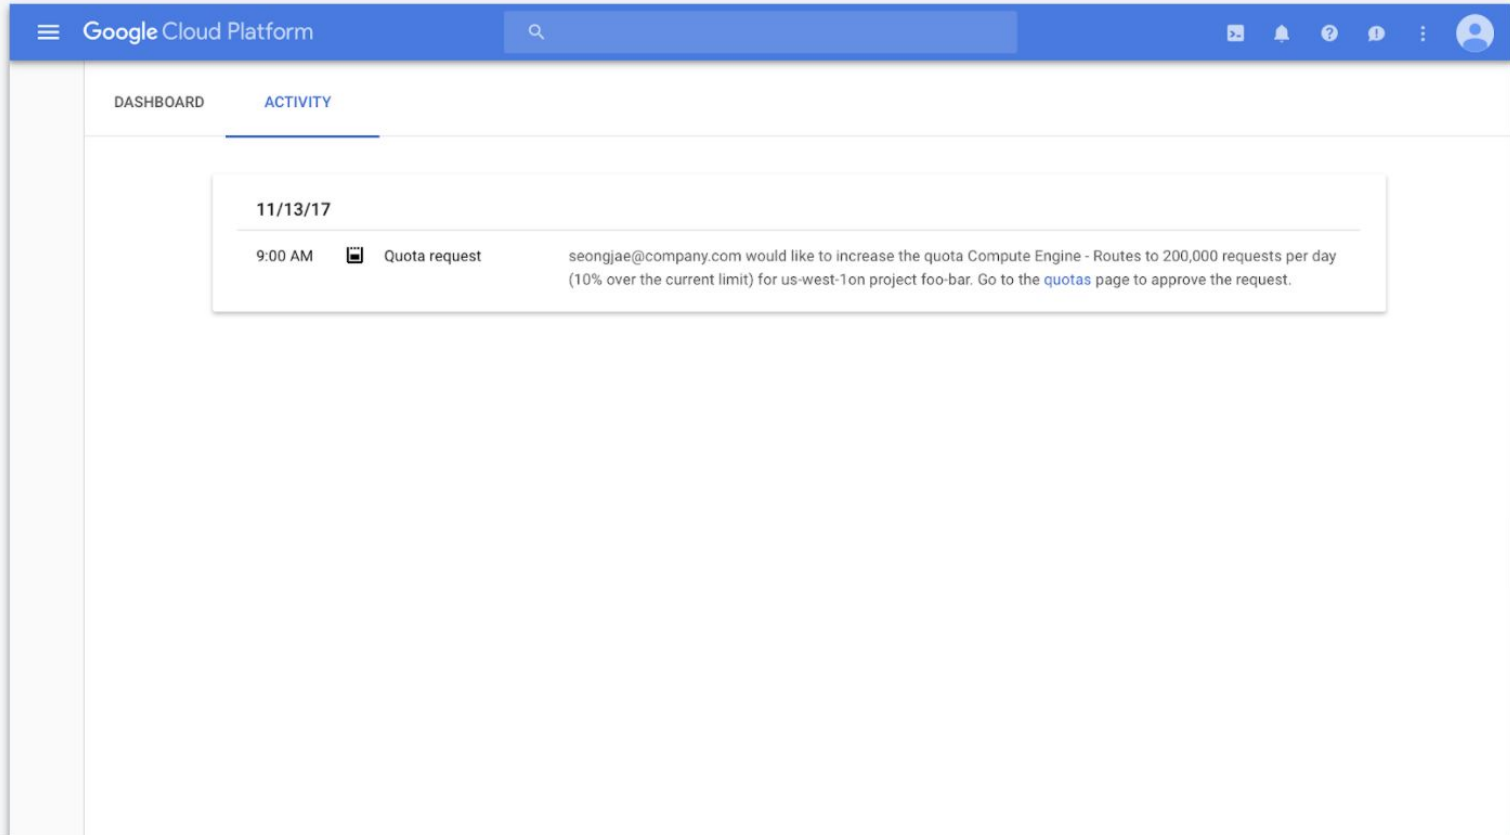
Change to
Enter a number between 0 and 100,000
Current limit: 100,000

Justification
Required

Quota group for Networks

Change to
Enter a number between 0 and 100,000
Current limit: 100,000

3. Provides a **time-stamp** + **message** to attach to it (view from admin side)




The screenshot shows the Google Cloud Platform interface. At the top, there is a blue header with the text "Google Cloud Platform" and a search bar. Below the header, there are two tabs: "DASHBOARD" and "ACTIVITY", with "ACTIVITY" being the active tab. The main content area displays a single activity entry for the date "11/13/17". The entry includes a timestamp "9:00 AM", a document icon, and the text "Quota request". The message details a request from "seongjae@company.com" to increase the quota for "Compute Engine - Routes" to 200,000 requests per day, which is 10% over the current limit for the project "foo-bar". The message concludes with a link to the "quotas" page to approve the request.

Google Cloud Platform

DASHBOARD ACTIVITY

11/13/17


9:00 AM  Quota request

seongjae@company.com would like to increase the quota Compute Engine - Routes to 200,000 requests per day (10% over the current limit) for us-west-1 on project foo-bar. Go to the [quotas](#) page to approve the request.

4. Supports **email notifications** at project-level

The screenshot shows an email notification from Google Cloud Platform. At the top left is the Google Cloud Platform logo, and at the top right is a link labeled "MANAGE QUOTAS". The main content of the email is titled "Quota request notification" and "Request for quota increase". It states that "seongjae@company.com has requested a a quota increase for project foo-bar". Under "Request details", it says "seongjae@company.com would like to increase the quota Compute Engine" and lists a detail: "- Routes to 200,000 requests per day (10% over the current limit) for us-west-1 on project foo-bar." At the bottom of the notification box is a blue button labeled "MANAGE QUOTAS". Below the notification box is the Google logo, followed by copyright information: "© 2016 Google Inc. 1600 Amphitheatre Parkway, Mountain View, CA 94043". At the very bottom, there is a footer note: "You have received this mandatory service announcement to update you about important changes to Google Cloud Platform or your account."

Google Cloud Platform MANAGE QUOTAS

 **Quota request notification**

Request for quota increase

seongjae@company.com has requested a a quota increase for project foo-bar

Request details

seongjae@company.com would like to increase the quota Compute Engine

- Routes to 200,000 requests per day (10% over the current limit) for us-west-1 on project foo-bar.

[MANAGE QUOTAS](#)

Google

© 2016 Google Inc.
1600 Amphitheatre Parkway, Mountain View, CA 94043

You have received this mandatory service announcement to update you about important changes to Google Cloud Platform or your account.

CASE 02. Org level advanced warnings

How can we **inform** users when their **quota usage** is being (or about to be) **limited** due to its **ancestor's quota limit**?

Why is this a problem?

Example

- 1 The folder's limit caps off at 48 VMs
- 2 The project limit caps off at 100 VMs
- 3 The folder only has this project (no other sibling projects)

Then further quota requests will be limited for this project.

...however, in this project scope view, **the user won't know** if it is being limited.

Provide **warning icon + banner**, with follow-up action to admin

The screenshot shows the Google Cloud Platform interface for managing quotas. At the top, there's a blue header with the Google Cloud Platform logo and a search bar. Below the header, the page title is "Quota" with two buttons: "+ EDIT QUOTAS" and "VIEW REQUESTS".

The main content area is titled "Quotas for project 'foo-bar'" and includes a sub-header: "View, manage and understand all your quota usages across all APIs and services. Filter by quota type, service, metric and region and their zones." Below this, a warning banner states: "Two quotas are being limited by other quotas in this project's folder. Click on 'view hierarchy' to learn more or notify an admin to let them know your concern." with a "CONTACT ADMIN" button.

Below the banner are four filter dropdowns: "Quota type" (All quotas), "Service" (All services), "Metric" (All metrics), and "Region" (All regions), along with a "Clear" button.

The main table displays quota usage for various services. The table has columns for "Service", "Location", "Current Usage & 7 Day Peak Usage", "Current Usage", "7 Day Peak Usage", and "Limit".

<input type="checkbox"/> Service	Location	Current Usage & 7 Day Peak Usage	Current Usage	7 Day Peak Usage	Limit	
<input type="checkbox"/> Firewall rules Compute Engine API	us-west-1a		48	95	100	View hierarchy
<input type="checkbox"/> Networks Compute Engine API	us-west-1	Quota is limited due to other quotas.	4	4	5	View hierarchy
<input type="checkbox"/> Routes Compute Engine API	Global		100	120	200	View hierarchy
<input type="checkbox"/> Number of indexes per day App Engine	us-west-1b		20	120	200	View hierarchy
<input type="checkbox"/> Operation read requests per	us-east-1		0	1,600	2,000	View hierarchy

Contact admin panel appears

Quota [+ EDIT QUOTAS](#) [VIEW REQUESTS](#)

Quotas for project "foo-bar"

View, manage and understand all your quota usages across all APIs and services. Filter by quota type, service, metric and region and their zones.

Two quotas are being limited by other quotas in this project's folder. Click on 'view hierarchy' to learn more or notify an admin to let them know your concern. [CONTACT ADMIN](#)

Quota type: All quotas | Service: All services | Metric: All metrics | Region: All regions

<input type="checkbox"/> Service	Location	Current Usage & 7 Day Peak Usage	Current Usage	7 Day Peak
<input type="checkbox"/> Firewall rules Compute Engine API	us-west-1a		48	95
<input type="checkbox"/> Networks Compute Engine API	us-west-1	 <small>Quota is limited due to other quotas in your folder.</small>		4
<input type="checkbox"/> Routes Compute Engine API	Global		100	120
<input type="checkbox"/> Number of indexes per day App Engine	us-west-1b		20	120

Notify Admin

Two resources are being limited due to other quotas in this project's folder. [Learn more.](#)

Select any or all quotas below to notify your organization's quota administrator of your concern.

[SELECT ALL](#)

- Firewall rules
Compute Engine API
- Routes
Compute Engine API

[NOTIFY ADMIN](#) [BACK](#)

Cloud Feed is deep linked from View Requests -- with quotas already filtered

The screenshot shows the Google Cloud Platform interface. At the top, the navigation bar includes the Google Cloud Platform logo, a search bar, and user profile icons. Below the navigation bar, there are two tabs: "Quota" and "Notify Admin". The "Quota" tab is active, and it contains two buttons: "+ EDIT QUOTAS" and "VIEW REQUESTS". The "VIEW REQUESTS" button is highlighted with a pink border. Below the tabs, the main content area displays "Quotas for project 'foo-bar'". A secondary window is overlaid on top of the main content, showing the "ACTIVITY" tab. The activity feed shows two entries for "11/13/17":

- 9:00 AM Quota usage impact seongjae@company.com is concerned about quota limit impact on Compute Engine API for Networks for project foo-bar
- 9:00 AM Quota usage impact seongjae@company.com is concerned about quota limit impact on App Engine Number of indexes per day for project foo-bar

CASE 03. Solving short term rate quotas

How can we show **short-term rate quotas** (per minute, per 100 seconds) in the quotas page as it is the **most common** in **APIs/services**?

Why is this a problem?

Suppose a customer uses APIs/services with **only short-term quotas**, all Quotas page **won't** give the **customer** any **usage data**.

RIGHT NOW

We provide a **question mark** with an **explanation text**.

i.e. "This quota is per minute, so your usage is not shown."



It's still **not** clear to our customers what this means.

CONSTRAINTS

It's **impossible** to show short-term usage because it is **constantly changing** -- the user will have to **refresh the page** every second // minute //etc to continue to see its **changing state**.

It's also not as meaningful, since the window shown is so short.


Users care more about **trends** rather than **snapshots**.

INTRODUCE 7 DAY PEAK AS AN ADDITIONAL INDICATOR

Pill usage component today:



Usage + 7 day peak usage:



Current usage Good | **7 Day Peak** Danger

Current usage Warning | **7 Day Peak** Warning

Current usage Good | **7 Day Peak** Good

W: 5px
H: 24px
Radius: 2

USER FEEDBACK



EXPLORATIONS

Bar & Line makes a little more sense. As I'm looking at them side by side I'm liking Bar & Line more.

P1

Google is pretty good at keeping things simple and I think it's simple and intuitive and easy to read and something that is actually useful and is not something that I have to learn how to use before looking at it.

P2

I think that one is the easiest to quickly look at and interpret your stats.

P10

All quotas page

Sortable by current usage or 7 day peak % [mocks](#)

Google Cloud Platform

Quota [+ EDIT QUOTAS](#) [VIEW REQUESTS](#)

Quotas for project "foo-bar"

View, manage and understand all your quota usages across all APIs and services. Filter by quota type, service, metric and region and their zones.

Quota type: All quotas | Service: All services | Metric: All metrics | Region: All regions | Clear

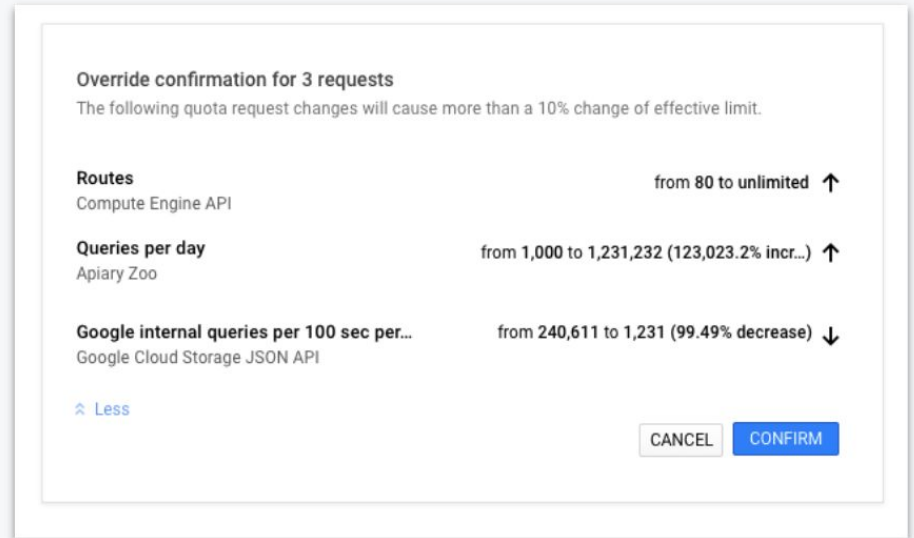
<input type="checkbox"/> Service	Location	Current Usage & 7 Day Peak Usage	Current Usage	7 Day Peak Usage	Limit	
<input type="checkbox"/> Firewall rules Compute Engine API	us-west-1a		48	95	100	View hierarchy
<input type="checkbox"/> Networks Compute Engine API	us-west-1		4	4	5	View hierarchy
<input type="checkbox"/> Routes Compute Engine API	Global		100	120	200	View hierarchy
<input type="checkbox"/> Number of indexes per day App Engine			20	120	200	View hierarchy
<input type="checkbox"/> Operation read requests per 100 seconds per use	us-east-1		0	1,600	2,000	View hierarchy

CASE 04. Safe override

How can we ensure that **high risk quota** changes be **safely** applied (> 10% increase)?

Modal dialog that acts as an extra **safeguard** for the user

Background research shows that users want to see **both %** and **#** when it comes to viewing and understanding their quotas.



THANK YOU

