

# Daily Elastic Byte - Elastic Cloud Autoscaling

---

David Pilato | @dadoonet | david@elastic.co



**Elastic  
B(y|i)te**

daily lunch appetizer



[Deployments](#)

Features

Support

# Deployments

🔍 e.g.: healthy:y us-east

Health ▾

Version ▾

Stack ▾

Configuration ▾



Create deployment

## Elastic Daily Bytes S04 ✓

ed32a7b v8.2.0

🌐 Belgium (europe-west1)

**gcp.es.master.n2.68x32x45**  
1 GB RAM in 1 zone

**gcp.es.datahot.n2.68x10x45**  
2 GB RAM in 2 zones

**gcp.kibana.n2.68x32x45**  
1 GB RAM in 1 zone

2 other configurations

...



## devoxxfr ✓

372f4f1 v8.1.0

🌐 Belgium (europe-west1)

**gcp.es.master.n2.68x32x45**  
1 GB RAM in 1 zone

**gcp.es.datahot.n2.68x10x45**  
2 GB RAM in 2 zones

**gcp.kibana.n2.68x32x45**  
1 GB RAM in 1 zone



Deployments

Elastic Daily Bytes S04

- Edit
- Elasticsearch
  - Snapshots
  - API console
- Kibana
- Integrations Server
- Enterprise Search
- Logs and metrics
- Activity
- Security
- Performance

Features

Support

# Elastic Daily Bytes S04

 Belgium (europa-west1)

**Autoscaling is available** ✕

You can now enable autoscaling on your deployment. When the deployment approaches the disk allocation threshold, storage capacity will scale up automatically. Additionally, for machine learning nodes RAM can scale both up and down. [Learn more](#)

[Enable autoscaling](#)

Deployment name

Elastic Daily Bytes S04 [Edit](#)

Deployment ID: ed32a7b

Deployment status

● Healthy

Manage ▾

[Open Kibana](#)

Custom endpoint alias [?](#)

elastic-daily-bytes-s04 [Edit](#)


Deployment version

v8.2.0

Applications [?](#)

<b>Elasticsearch</b>	<a href="#">Copy endpoint</a>	<a href="#">Copy cluster ID</a>
<b>Kibana</b>	<a href="#">Open</a>	<a href="#">Copy endpoint</a> <a href="#">Copy cluster ID</a>
<b>APM</b>	<a href="#">Open</a>	<a href="#">Copy endpoint</a> <a href="#">Copy cluster ID</a>
<b>Fleet</b>	<a href="#">Open</a>	<a href="#">Copy endpoint</a> <a href="#">Copy cluster ID</a>
<b>Enterprise Search</b>	<a href="#">Copy endpoint</a>	<a href="#">Copy cluster ID</a>

Cloud ID [?](#)

Elastic\_Daily\_Bytes\_S04:ZXVyb3BLLXdlc3QxLmdjcC5jbG91ZC5lcy5pbYQ00DdjZTRlMDk2ZDI0NTEzYWwMNTkwNGU2ZUW2MjdhYyRkMTRk0YTI1MTczYTc0YjQ050WUzZDE2NTkzNzQ2ZTAzMg== 

## Deployments

## Elastic Daily Bytes S04

Edit  
Elasticsearch  
Snapshots  
API console  
Kibana  
Integrations Server  
Enterprise Search  
Logs and metrics  
Activity  
Security  
Performance

## Features

## Support

## Elastic Daily Bytes S04

Belgium (europe-west1)

## Autoscaling is available

You can now enable autoscaling for machines in your deployment.

[Enable autoscaling](#)

### Enable autoscaling

Autoscaling increases deployment storage capacity automatically. For ML nodes both storage capacity and RAM increase and decrease. Future capacity needs are forecasted based on past usage.

You can set the maximum size per zone in order to avoid excess costs. For ML nodes you can set a maximum and minimum.

[Cancel](#)[Enable autoscaling](#)

## Deployment name

Elastic Daily Bytes S04

## Deployment ID

## Custom endpoint

elastic-daily-bytes-s04

## Applications

Elasticsearch	<a href="#">Copy endpoint</a>	<a href="#">Copy cluster ID</a>
Kibana	<a href="#">Open</a>	<a href="#">Copy endpoint</a> <a href="#">Copy cluster ID</a>
APM	<a href="#">Open</a>	<a href="#">Copy endpoint</a> <a href="#">Copy cluster ID</a>
Fleet	<a href="#">Open</a>	<a href="#">Copy endpoint</a> <a href="#">Copy cluster ID</a>
Enterprise Search	<a href="#">Copy endpoint</a>	<a href="#">Copy cluster ID</a>

## Cloud ID

Elastic\_Daily\_Bytes\_S04:ZXVyb3BLLXdlc3QxLmdjcc5jbG91ZC5lcys5bpyQ00ddjZTRlMdk2ZDI0NTEzYWwWNTkwNGU2ZWU2MjdhYyRkNTk0YTI1MTczYTc0YjQ5Q50WUzZDE2NTkzNzQ2ZTAzMg==

Deployments

Elastic Daily Bytes S04

- Edit
- Elasticsearch
  - Snapshots
  - API console
- Kibana
- Integrations Server
- Enterprise Search
- Logs and metrics
- Activity**
- Security
- Performance

Features

Support

# Activity

🇧🇪 Belgium (europa-west1)

🔍 e.g.: healthy\_configuration:y apm

Successful
Unsuccessful
Pending
Finished
System
Source
▼

All
20
🇧🇪 Elasticsearch
5
📊 Kibana
6
📄 Integrations Server
5
🔍 Enterprise Search
4

Change	Summary	Applied by	Actions
	🕒 Applied 3 minutes ago, took a few seconds		
🟢 🇧🇪 #elasticsearch-4	Enable autoscaling No significant Elasticsearch configuration changes	david@elastic.co	<span>Reapply</span> <span>▼</span> <span>Details</span> <span>^</span>

Step	Started	Ended	Duration
🟢 Plan successful	2022-05-04T13:56:54.651Z	2022-05-04T13:56:54.651Z	0 milliseconds
🟢 Performing clean-up	2022-05-04T13:56:54.641Z	2022-05-04T13:56:54.641Z	0 milliseconds
🟢 Post plan cleanup	2022-05-04T13:56:54.240Z	2022-05-04T13:56:54.633Z	393 milliseconds
🟢 Configuring autoscaling policy in Elasticsearch	2022-05-04T13:56:52.679Z	2022-05-04T13:56:54.230Z	around a second
🟢 Apply plan settings	2022-05-04T13:56:52.641Z	2022-05-04T13:56:52.670Z	29 milliseconds
🟢 Updating cluster settings	2022-05-04T13:56:51.865Z	2022-05-04T13:56:52.633Z	768 milliseconds
🟢 Setting maintenance mode	2022-05-04T13:56:51.803Z	2022-05-04T13:56:51.826Z	23 milliseconds
🟢 Applying X-Pack monitoring configuration	2022-05-04T13:56:51.783Z	2022-05-04T13:56:51.792Z	9 milliseconds
🟢 Applying X-Pack license to deployment	2022-05-04T13:56:51.701Z	2022-05-04T13:56:51.775Z	74 milliseconds

Restore Status - Snapshot and Restore | [ted32a7] Overview - Elastic Cloud | VM instances - Compute Engine

elastic-daily-bytes-s04.kb.europe-west1.gcp.cloud.es.io:9243/app/management/data/snapshot\_restore/res...

elastic Find apps, content, and more. Ex: Discover

Stack Management > Snapshot and Restore > Restore Status

## Management

- Ingest
  - Ingest Pipelines
  - Logstash Pipelines
- Data
  - Index Management
  - Index Lifecycle Policies
  - Snapshot and Restore**
  - Rollup Jobs
  - Transforms
  - Cross-Cluster Replication
  - Remote Clusters
- Alerts and Insights
  - Rules and Connectors
  - Cases
  - Reporting
  - Machine Learning
  - Watcher
- Security
  - Users
  - Roles
  - API keys
  - Role Mappings
- Kibana
  - Data Views
  - Saved Objects

## Snapshot and Restore

Use repositories to store and recover backups of your Elasticsearch indices and clusters.

[Snapshot and Restore docs](#)

Snapshots   Repositories   Policies   **Restore Status**

Refresh data every 30 seconds


Index	Status	Last activity	Shards completed	Shards in progress
demo-person-1	In progress	now	0	1


ID	Stage	Start time	End time	Duration	Repository	Snapshot	Version	Target host	Target node	Bytes	Files
0 P	INDEX	May 04, 2022 4:04 PM GMT+2		68 seconds	gcs	demo-backup	7.12.0	10.43.0.139	instance-0000000001	0.0%	0.0%

demo-person Complete May 04, 2022 3:39 PM GMT+2 1

Rows per page: 20

< 1 >







Index Management - Elastic x Edit deployment - Elastic Cloud x elastic-daily-bytes - IAM & Ad...

cloud.elastic.co/deployments/ed32a7ba14f11459fa2359402f293c600/edit

elastic

Cloud Deployments Elastic Daily Bytes S... Edit

Deployments

Elastic Daily Bytes S04

Edit

Elasticsearch

Snapshots

API console

Kibana

Integrations Server

Enterprise Search

Logs and metrics

Activity

Security

Performance

Features

Support

## Edit

**Autoscale this deployment**  
Scale the capacity of your data and machine learning nodes automatically. [Learn more](#)

### Elasticsearch

**Hot data and Content tier** [info](#)  
Nodes in this tier ingest and process frequently queried data.

Size per zone	45 GB storage   1 GB RAM	<a href="#">Edit settings</a>
Availability zones	<input type="radio"/> 1 zone <input checked="" type="radio"/> 2 zones <input type="radio"/> 3 zones	
Total (size x zone)	90 GB storage   2 GB RAM	Up to 5 vCPU

**Warm data tier**  
Store less-frequently queried data on a warm tier for increased cost savings.  
[Edit capacity and settings](#)

**Cold data tier**  
Store rarely queried data on a cold tier for increased cost savings.  
[Edit capacity and settings](#)

**Frozen data tier**  
Maximize savings by archiving data on a frozen tier.

## Autoscaling settings

**gcp.es.datahot.n2.68x10x45**

**How does autoscaling work?**  
By monitoring the instance storage usage, we predict future capacity needs and scale up the storage size automatically.

### Settings

**Maximum size per zone**  
5.63 TB storage | 128 GB RAM

We will never autoscale this tier beyond the maximum size per zone.

**Current size per zone**  
45 GB storage | 1 GB RAM [Force set size](#)

To maintain this size while autoscaling is enabled, set the maximum size per zone to match this value.

**Advanced settings**

**Forecast window**  
30 minutes

Make capacity changes based on the usage during this previous period of time.

[Close](#) [Update](#)

You must click the save button on the main page for these changes to take effect.