



# Elasticsearch Service on Elastic Cloud

The official service from the creators of Elasticsearch

Nov 23, 2020



# SaaS and Managed Services

## SPEED

---



### Fast time to value

Instant global infrastructure & operations encourage experiments and get you to market.



### Agility

Pivot quickly to respond to competitive pressure, opportunity, or changing customer expectations.

## SIMPLICITY

---



### Operational Simplicity

Focus on driving business value versus managing infrastructure, operations or integrating products.



### Development Simplicity

A single programming model and integrated product portfolio that can address search, observability, security.

## SECURITY & RISK REDUCTION

---



### Secure from the start

All deployments automatically have the best security by default, no configuration. Elastic compliance included.



### Reduce Risk

Source global infrastructure, ops and support experts from the originators of the Elasticsearch technology on any cloud.

# Elastic Technology

3 solutions



Elastic Enterprise Search



Elastic Observability



Elastic Security

Powered by the Elastic Stack



Deployed anywhere



SaaS



Elastic Cloud Enterprise



Elastic Cloud on Kubernetes

Orchestration

# Only with Elastic Cloud

---



Lens &  
Canvas



Machine  
learning



Index life cycle  
management



Integrated stack  
security



Native sql  
engine



Spaces



Rollups



Cross cluster  
search



Alerting



Maps

---

Exclusive  
Features



ELASTIC ENTERPRISE SEARCH



ELASTIC OBSERVABILITY



ELASTIC SECURITY

---

Exclusive  
Solutions

# Elasticsearch Service on Elastic Cloud

*Fast, Simple, Secure Cloud for your mission critical apps*



## Managed Elasticsearch and Kibana

Deploy, operate, and scale our solutions, Elasticsearch, and Kibana in an instant on Azure.

The screenshot displays the Elastic Cloud console interface. On the left, a sidebar shows navigation options: Deployments, Edit, Elasticsearch, Logs, Snapshots, API Console, Kibana, APM, Activity, Security, Performance, and Custom plugins. The main area is titled 'First cluster Edit' and shows configuration for a 'Data' cluster. The cluster name is 'gcp.data.highio.1' and it is described as 'A Kibana instance.'. The 'Fault tolerance' section has radio buttons for '1 zone' (selected), '2 zones', and '3 zones'. Below this is a 'RAM per Node' slider ranging from 1 GB to 64 GB, with a current selection at 8 GB. A 'Summary' panel on the right provides cluster details:

Summary	
Name	First Cluster
Version	v7.0.1
ES data memory	24 GB
ES data storage	1.25 TB
Total memory	25.5 GB
Total storage	1.25 TB
Hourly rate	\$0.8281
Monthly rate	\$604.51

Below the summary is an 'Architecture' diagram showing two zones. Zone 1 contains three nodes: 'gcp.data.hi...' (8 GB RAM), 'gcp.data.hi...' (4 GB RAM), and 'gcp.kibana.1' (1 GB RAM). Zone 2 contains two nodes: 'gcp.data.hi...' (8 GB RAM) and 'gcp.data.hi...' (4 GB RAM). In the foreground, a 'Kibana 1 configuration' dialog is open, showing settings for 'gcp.kibana.1'. It includes 'Fault tolerance' (1 zone selected), 'RAM per Node' (slider at 1 GB), and a 'Summary' section showing '1 GB RAM x 1 instance x 1 zone = 1 GB RAM'. A 'User settings overrides' link is visible at the bottom of the dialog.

# Elasticsearch Service on Elastic Cloud: Benefits

For Elastic Stack and Solutions

## SPEED



Fast time to value



Agility

## SIMPLICITY



Operational Simplicity



Development Simplicity

## SECURITY & RISK REDUCTION



Secure from the start



Reduce Risk

# Elastic Best Practices for Your Use Case

With a click

## How long does it take to deploy a new cluster, hot-warm, the architecture you want?

- Predefined deployment templates use Elastic tested and recommended instance types and configurations
- Hot-warm deployment with index lifecycle management for large scale time-series use case
- Independently scale Elasticsearch node roles for improved ingest and search performance

### Select a hardware profile

**I/O Optimized** Recommended

New to Elasticsearch? This template is suitable for all-purpose workloads that don't require more specialized resources. [See details](#)

**Compute optimized**

Run CPU-intensive workloads or run smaller workloads cost-effectively when you need less memory and storage. [See details](#)

**Cross Cluster Search**

Search data across one or more associated deployments. [See details](#)

**Hot-Warm Architecture**

Useful for time-series analytics that benefit from automatic index curation. [See details](#)

**Memory Architecture**

Perform memory-intensive operations efficiently, including workloads with frequent aggregations. [See details](#)

# Simple Cluster Management

Ease of use

Do you have long upgrade times, big deployments, multiple use cases?

- Configure, manage and scale multiple deployments with a single console
- In-place configuration changes for faster and more reliable cluster changes
- Automatic snapshots with configurable frequency and retention
- Deploy / turn on built-in monitoring cluster with a click in Kibana

The screenshot displays the Elastic Cloud console interface for creating a new deployment. The left sidebar contains navigation links for Deployments, Custom plugins, Account, and Help. The main content area is titled 'Create deployment' and includes a description of the template, a 'Data' configuration card, and a 'Machine Learning' configuration card. A 'Summary' table on the right provides a breakdown of resources and costs.

**Deployments** / Create

## Create deployment

Take the template that pre-configures the Elastic Stack and make it yours. Adjust capacity and performance, change the level of fault tolerance, add more features, and much more. [Learn more](#)

**Data** 1 configuration

Store, search, and analyze big volumes of data quickly. [Learn more](#)

**azure.data.highio.i32sv2** **Data** **Coordinating** **Master**

An I/O optimized Elasticsearch instance running on an Azure I32sv2.

**Fault tolerance**

1 zone  2 zones  3 zones

**RAM per Node** **Nodes**

1 GB 2 GB 4 GB **8 GB** 15 GB 29 GB 58 GB

**RAM per Zone**

8 GB

**Summary**

8 GB RAM 240 GB storage × 1 node × 2 zones =

16 GB RAM 480 GB storage

> User setting overrides

**Machine Learning** 1 configuration

Automatically model the behavior of your Elasticsearch data — trends, periodicity, and more. [Learn more](#)

**azure.ml.d64sv3** **Machine Learning**

An Elasticsearch machine learning instance running on an Azure D64sv3.

**Fault tolerance**

1 zone  2 zones  3 zones

**RAM per Node** **Nodes**

Summary	
Version	v7.6.0
ELASTICSEARCH	
Memory	16 GB
Storage	480 GB
Master memory	1 GB
Hourly rate	\$0.3658
KIBANA	
Memory	1 GB
Hourly rate	FREE
ML	
Memory	1 GB
Hourly rate	FREE
APM	
Memory	512 MB
Hourly rate	FREE
TOTAL	
Total memory	19.5 GB
Total storage	480 GB
Hourly rate	\$0.3658

## Architecture

**Zone 1**

**Zone 2**

azure.data.highio.i32...

# Simple Cluster Management

REST API, High Availability, Upgrades

**Do you experience downtime with upgrades, use APIs, hit cloud quotas?  
How often do you upgrade?**

- Zero downtime upgrades with a few clicks
- CI/CD pipeline integration via Public REST API and ecctl command line
- High availability across multiple AZs
- Smooth scaling with no quota delays from your cloud provider

The screenshot displays the Elastic Cloud console interface for a deployment named "GCP\_Demo".

**Deployment Details:**

- Deployment name:** GCP\_Demo (with an Edit button)
- Deployment status:** Healthy
- Deployment version:** v7.2.0 (with an Upgrade button)
- Applications:** Elasticsearch, Kibana, and APM are listed with their respective icons and "Launch" buttons. Each application has a "Copy Endpoint URL" button.
- Cloud ID:** GCP\_Demo : dXmtd2VzdEuZ2NwLmNsb3VkLmVzLm1vJDEyODgzMmJhMjY4MjQ5OWNhNDUwZjYyNzFhMDd1ZGI4JDhmOGYwODc2NWE2MjQ3ZDF1Zjc4MzQ2ZTZ1MDkyMTdi

**Instances:**

There are three instance groups: gcp.data.highio.1 (1 instance), gcp.kibana.1 (1 instance), and gcp.apm.1 (1 instance).

**Instance Details:**

- us-west1-b:**
  - GCP.DATA.HIGHIO.1:** Instance #0, v7.2.0, 8 GB RAM. Roles: master, data, ingest. JVM memory pressure: 3%. Disk usage: 0% of 240 GB. Stop routing button.
  - GCP.KIBANA.1:** Instance #2, v7.2.0, 2 GB RAM. Stop routing button.
- us-west1-c:**
  - GCP.APM.1:** Instance #0, v7.2.0, 1 GB RAM. Stop routing button.

# Continuous Compliance

In the cloud

**What customer type do you need to serve?**

- HIPAA
- CSA Star Level 2
- SOC 2 Type 1, Type 2 and SOC 3
- ISO 27001, ISO 27107, ISO 27018
- FedRAMP authorized at Moderate impact level, deployable to AWS GovCloud (US)
- Elastic Cloud operates in compliance with GDPR principles

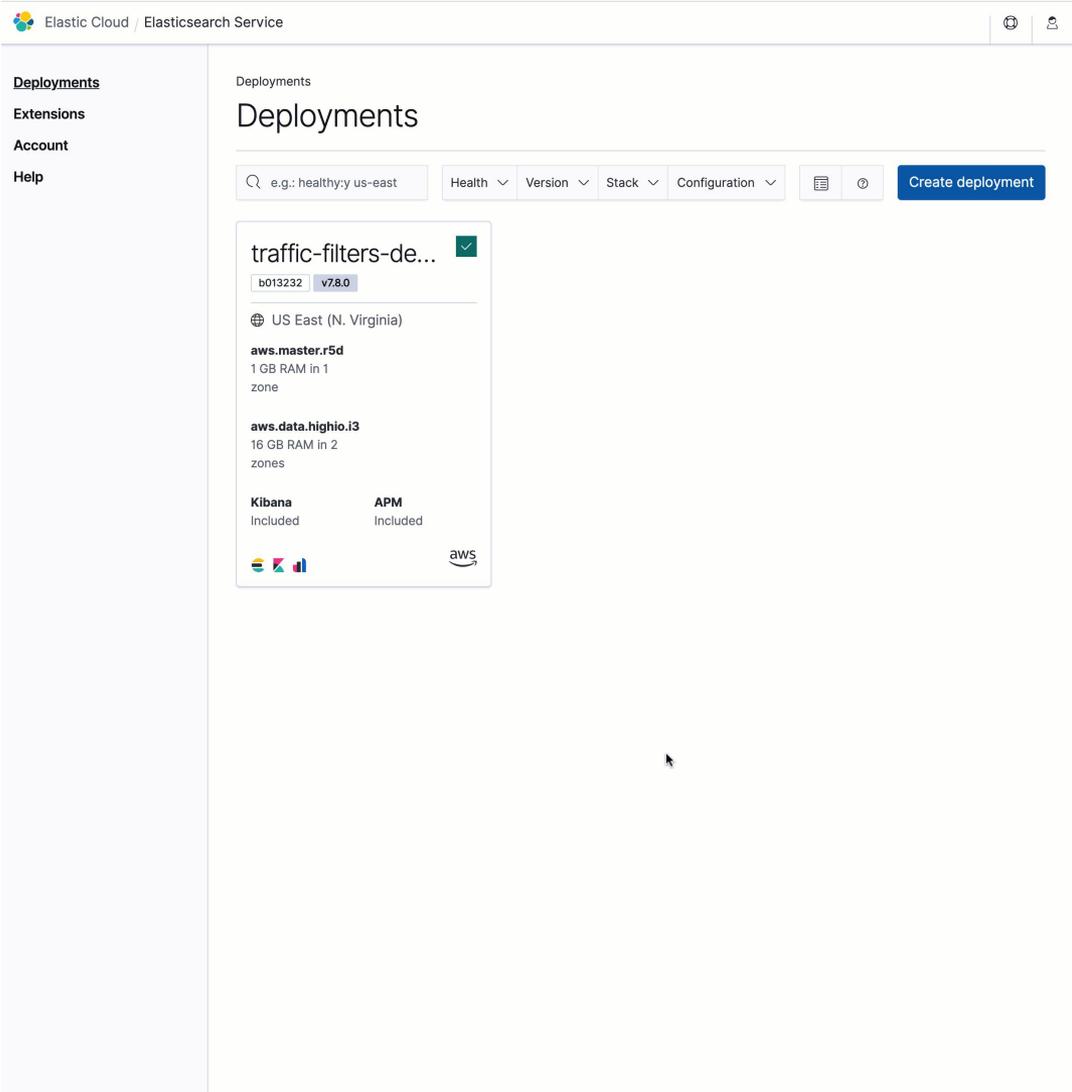


# Secure by Default

Native security, fully integrated

**How long do you spend configuring security, how do you ensure you are protected?**

- Native Authentication: SAML & OIDC, Kerberos Auth and more
- Native Authorization: Powerful role and attribute based access control, field and document level security
- Network security: IP filtering and AWS PrivateLink integration
- Elastic Cloud is always up to date with any new Stack and solution security updates



The screenshot displays the Elastic Cloud interface for managing deployments. The top navigation bar includes 'Elastic Cloud' and 'Elasticsearch Service'. A left sidebar contains navigation links for 'Deployments', 'Extensions', 'Account', and 'Help'. The main content area is titled 'Deployments' and features a search bar with the placeholder 'e.g.: healthy:us-east', several dropdown menus for 'Health', 'Version', 'Stack', and 'Configuration', and a 'Create deployment' button. A deployment card is shown for 'traffic-filters-de...' with a green checkmark. The card includes the ID 'b013232', version 'v7.8.0', and location 'US East (N. Virginia)'. It lists two AWS instance types: 'aws.master.r5d' (1 GB RAM in 1 zone) and 'aws.data.highio.i3' (16 GB RAM in 2 zones). Below the instance list, it indicates that 'Kibana' and 'APM' are included. The card also features a small bar chart icon and the AWS logo.

# Secure by Default

Encryption, patching, hardening

## How long does it take to configure encryption, perform security patches, how much downtime?

- Data and snapshot encryption at rest (EAR)
- TLS encryption of data to, from, and within your deployments
- Multi-Factor Authentication for Elastic Cloud console
- [CIS](#) Level 1 & 2 Server Profile OS hardening
- OS kernel patches in ~48 hours from CVE publication in our virtual images

### Multi-factor authentication

Add an extra layer of security by setting up Google authenticator or text messaging on a mobile device.

[Learn more](#)

Add a device to enable multi-factor authentication.



### Authenticator app

#### 1 Scan QR code

Use the Google authenticator app to scan the QR code below



#### 2 Enter passcode

Enter the 6 digit auth.elastic.co passcode

E.g. 123456

Enable device

Cancel



### Text message

Use your mobile phone to receive security codes

Add a phone number

# Consolidate your Cloud Bills

One bill from Azure + Elastic

---

## Billing integration for major cloud providers

- Subscribe and pay using your Azure account
- Draw down from Azure spend commitment
- Support for monthly standard and premium tiers



# Elasticsearch Service

The official managed Elasticsearch Service



## 10 Azure Regions and growing

### Americas

- westus2 (Washington)
- centralus (Iowa)
- eastus2 (Virginia)

### Europe

- francecentral (Paris)
- northeurope (Ireland)
- euksouth (London)
- westeurope (Netherlands)

### Asia Pacific

- southeastasia (Singapore)
- japaneast (Tokyo)
- australiaeast (New South Wales)

# Elasticsearch Service

*Fast, Simple, Secure* Cloud from Elastic

## 8+ years

Experience operating  
Elasticsearch as a service

## 22,000+ Clusters, 39 Regions

Created & maintained for every  
possible use case around the world

## Always the latest version

New Stack/Solution releases are available on Elastic  
Cloud and as downloadable software the same day

The screenshot shows the Elastic Cloud interface for Elasticsearch Service. At the top, there's a navigation bar with the Elastic Cloud logo and user profile icons. The main content area features a large card for 'Elasticsearch Service' with a 'Create deployment' button. Below this are sections for 'Documentation', 'Webinars', 'News', and 'Training'. The 'News' section lists recent releases of Elastic Stack 7.6.0 and Elasticsearch Service on Google Cloud Platform (GCP) in Mumbai and Montréal. The 'Training' section promotes getting certified in the Elastic Stack.

**Elasticsearch Service**

Create your first deployment

Create your first deployment to manage an Elasticsearch cluster on the cloud platform of your choice. Add additional Elastic products to your deployment like Kibana, machine learning, or APM.

[Create deployment](#)

**Documentation**

Help me find...

- [Elasticsearch Service on Elastic Cloud documentation](#)
- [Elasticsearch documentation](#)
- [Elasticsearch REST API](#)

**Webinars**

- [Elastic Stack 7.0 o...](#)  
Get an in-depth look at the latest in...
- [Logging and...](#)  
Learn how to use the purpose-built...
- [Index lifecycle...](#)  
We'll cover how to use the index...
- [Kibana for...](#)  
In this webinar, we'll share how we use...

**News**

- [Elastic Stack 7.6.0 released](#)  
FEBRUARY 11, 2020 **New!**
- [Elasticsearch Service is now available on Google Cloud Platform \(GCP\) in Mumbai](#)  
JANUARY 16, 2020 **New!**
- [Elasticsearch Service is now available on Google Cloud Platform \(GCP\) in Montréal](#)  
DECEMBER 23, 2019 **New!**

**Training**

**Get certified!**

The Elastic Stack is versatile enough to tackle any use case. We'll teach you how to harness the power of that versatility and become an Elastic expert.

[Elasticsearch Engineer I](#)

[Kibana Data Analyst](#)



# Thank You

Elastic is a Search Company.

[www.elastic.co](http://www.elastic.co)