



Effective Monitoring of Web Services with Heartbeat and Uptime

AJ Pahl

Manager, Solutions Architecture

January 2020



Housekeeping & Logistics

- **Slides and recording** will be available following the webinar
- Chat via IRC `#elastic-webinar`
 - `#elastic-webinar` @ Freenode
 - Click "Join the Chat" link, create an IRC account
- Please select high resolution in the YouTube video player



AJ Pahl
@ajpahl1008

But first...



apm

+

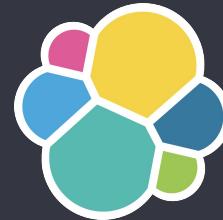


logging

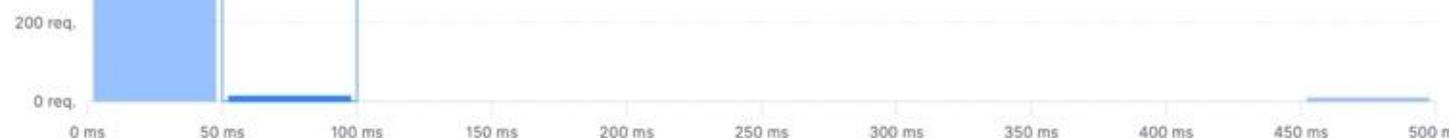
=



metrics



elastic
observability



Trace sample

Actions ▾

[View full trace](#)

16 minutes ago | 82 ms (89.9% of trace)

GET <http://opbeans-java:3000/api/products>

500 Internal Server Error

2 Errors

Other

[Timeline](#) [Metadata](#)

Services opbeans-java opbeans-ruby opbeans-python

0 ms 10 ms 20 ms 30 ms 40 ms 50 ms 60 ms 70 ms 82 ms

HTTP 5xx **DispatcherServlet#doGet** 82 ms

GET opbeans-ruby 79 ms

HTTP 5xx **Rack** 73 ms

GET opbeans-python 70 ms

HTTP 5xx **GET opbeans.views.products** 2 66 ms

GET opbeans-node:3000 6,828 µs

Logs			
	Search for log entries... (e.g. host.name:host-1)	Configuration	Customize
		Highlights	07/17/2019 10:24:55 AM
		Stream live	
Timestamp	Message		
Jul 17, 2019 @ 10:24:55.497	0", "stack": "Error: connect ECONNREFUSED 172.18.0.9:3000\n at TCPConnectWrap.afterConnect [as oncomplete] (net.js:1191:14)", "type": "Error", "v": 1} {"level":50,"time":1563373495496,"pid":786,"hostname":"f559ddf5cb6b","msg":"Application encountered an uncaught exception. Flushing Elastic APM queue and exiting...","v":1}	Wed 17	
Jul 17, 2019 @ 10:24:55.497	{"level":50,"time":1563373495496,"pid":786,"hostname":"f559ddf5cb6b","code": "ECONNRESET", "msg": "socket hanging up", "stack": "Error: socket hang up\n at createHangUpError (_http_client.js:342:15)\n at Socket.socketCloseListener (_http_client.js:377:23)\n at emitOne (events.js:121:20)\n at Socket.emit (events.js:211:7)\n at TCP._handle.close [as _onclose] (net.js:561:12)", "type": "Error", "v": 1} {"level":50,"time":1563373495496,"pid":786,"hostname":"f559ddf5cb6b","msg":"Application encountered an uncaught exception. Flushing Elastic APM queue and exiting...","v":1}	03 AM	
Jul 17, 2019 @ 10:24:55.497	{"level":50,"time":1563373495496,"pid":786,"hostname":"f559ddf5cb6b","msg": "Elastic APM queue flushed!", "v": 1}	06 AM	
Jul 17, 2019 @ 10:24:55.507	2019-07-17 14:24:55.507 UTC [10307] LOG: could not receive data from client: Connection reset by peer		
Jul 17, 2019 @ 10:24:55.507	2019-07-17T14:24:55: PM2 log: App [server:1] exited with code [1] via signal [SIGINT]		
Jul 17, 2019 @ 10:24:55.509	2019-07-17T14:24:55: PM2 log: App [server:1] starting in -fork mode-	12 PM	
Jul 17, 2019 @ 10:24:55.512	14:24:55 dotnet.1 SUCCESSES: 0 FAILURES: 12083 WORKERS: 1		
Jul 17, 2019 @ 10:24:55.513	14:24:55 dotnet.1 SUCCESSES: 0 FAILURES: 12083 WORKERS: 1		
Jul 17, 2019 @ 10:24:55.513	14:24:55 dotnet.1 SUCCESSES: 0 FAILURES: 12083 WORKERS: 1	03 PM	
Jul 17, 2019 @ 10:24:55.513	14:24:55 dotnet.1 ClientConnectorError(111, 'Connection refused')		
Jul 17, 2019 @ 10:24:55.513	14:24:55 dotnet.1 File "/usr/local/lib/python3.7/site-packages/molotov/worker.py", line 206, in step		
Jul 17, 2019 @ 10:24:55.513	14:24:55 dotnet.1 **scenario['kw'])		
Jul 17, 2019 @ 10:24:55.513	14:24:55 dotnet.1 File "molotov_scenarios.py", line 34, in scenario_products_top	06 PM	
Jul 17, 2019 @ 10:24:55.513	14:24:55 dotnet.1 async with session.get(join(SERVER_URL, 'api', 'products', 'top')) as resp:		
Jul 17, 2019 @ 10:24:55.513	14:24:55 dotnet.1 File "/usr/local/lib/python3.7/site-packages/aiohttp/client.py", line 843, in __aenter__		
Jul 17, 2019 @ 10:24:55.514	14:24:55 dotnet.1 self._resp = await self._coro	09 PM	



Visualize Create



Save Share Inspect Refresh

Last 30 minutes

Show dates

Refresh



Time Series Metric Top N Gauge Markdown Table



Auto apply

The changes will be automatically applied.



Data Panel options Annotations



Label



Metrics Options



Aggregation



Count



Group by

Everything



D

Metrics

[Inventory](#) [Metrics Explorer](#) [Settings](#)

View: Kubernetes

Search for infrastructure data... (e.g. host.name:host-1)

10/28/2019 5:20:18 PM

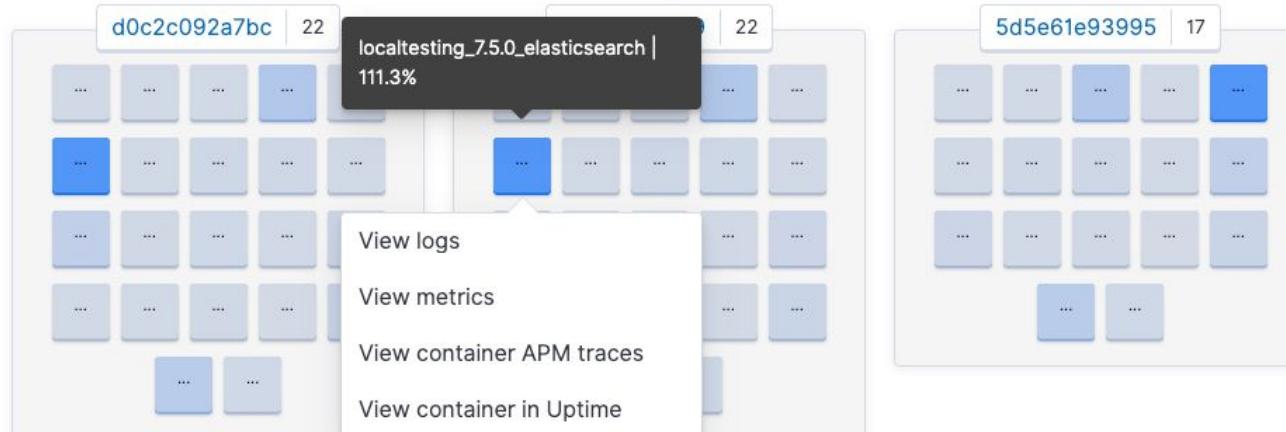
Auto-refresh

Metric: CPU usage

Group By: All

[Save](#) [Load](#)[Map view](#)[Table view](#)

Showing the last 1 minute of data at the selected time



0.3%

139.6%





Elastic Approach to Observability

Dev & Ops Teams



Log Data

Web Logs
App Logs
Database Logs
Container Logs

Metrics Data

Container Metrics
Host Metrics
Database Metrics
Network Metrics
Storage Metrics

APM Data

Real User Monitoring
Txn Perf Monitoring
Distributed Tracing

Uptime Data

Uptime
Response Time



Why is this important?

Weekends

KPI

Home Life

< Key Performance Indicators >

NPS

< Net Promoter Score >



Focus on
Projects



Stop Wasting
Time

SLA

< Service Level Agreements >

Get Day Job
Done

MTTR

<Mean Time to Resolution>

Need Sleep

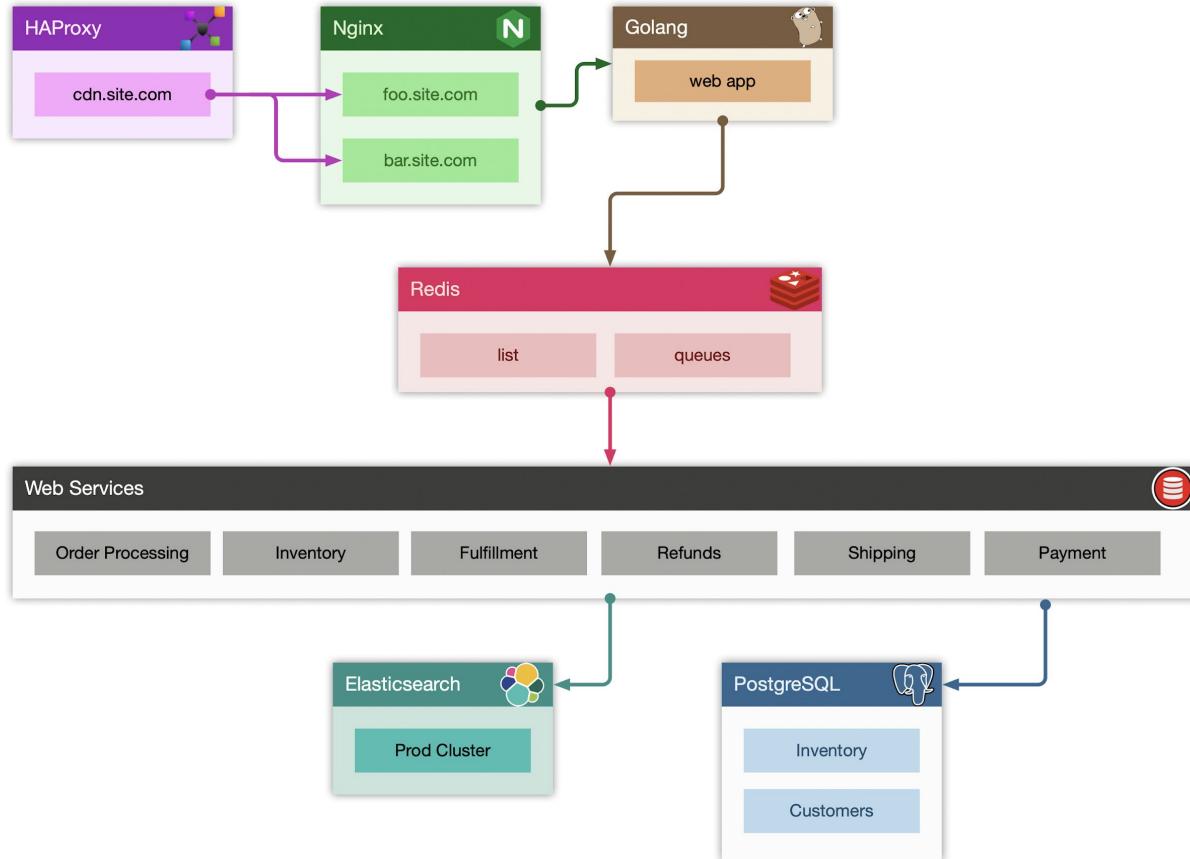
Lost Productivity



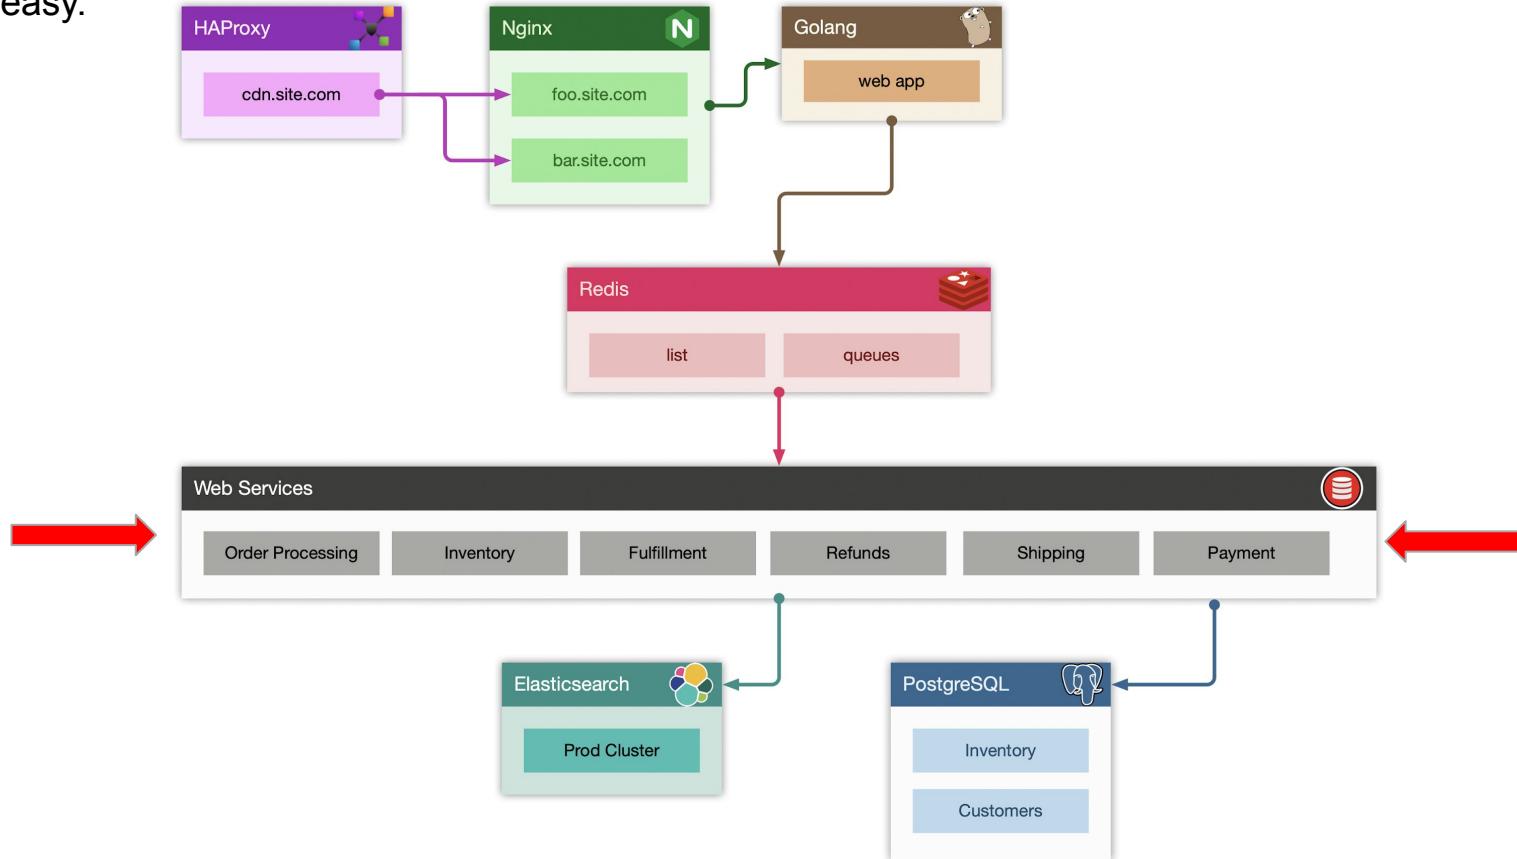
Easy.



Not so easy.



Not so easy.



Even still...

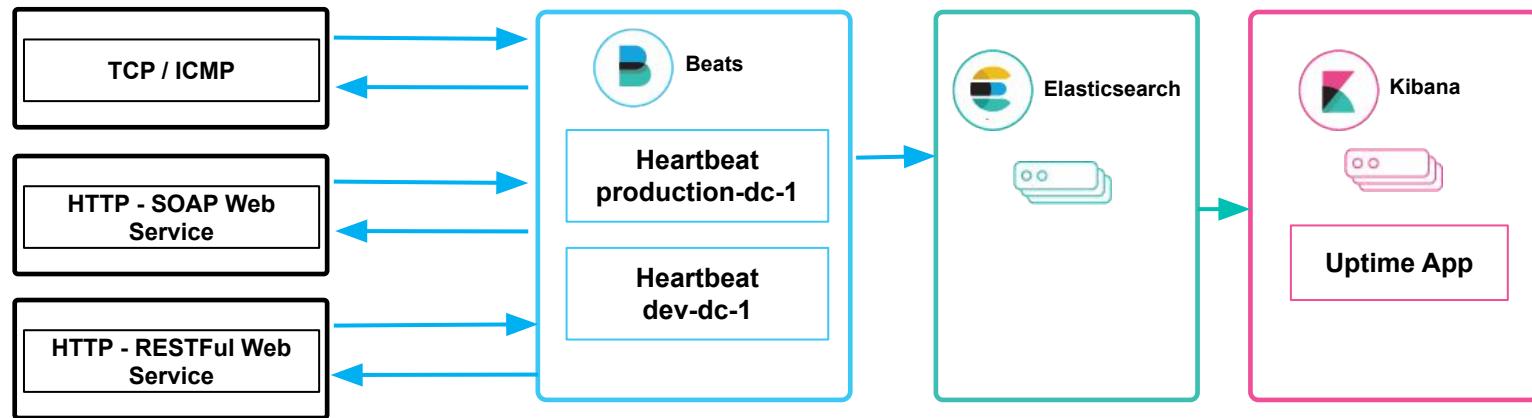


So how can this be made it easy?



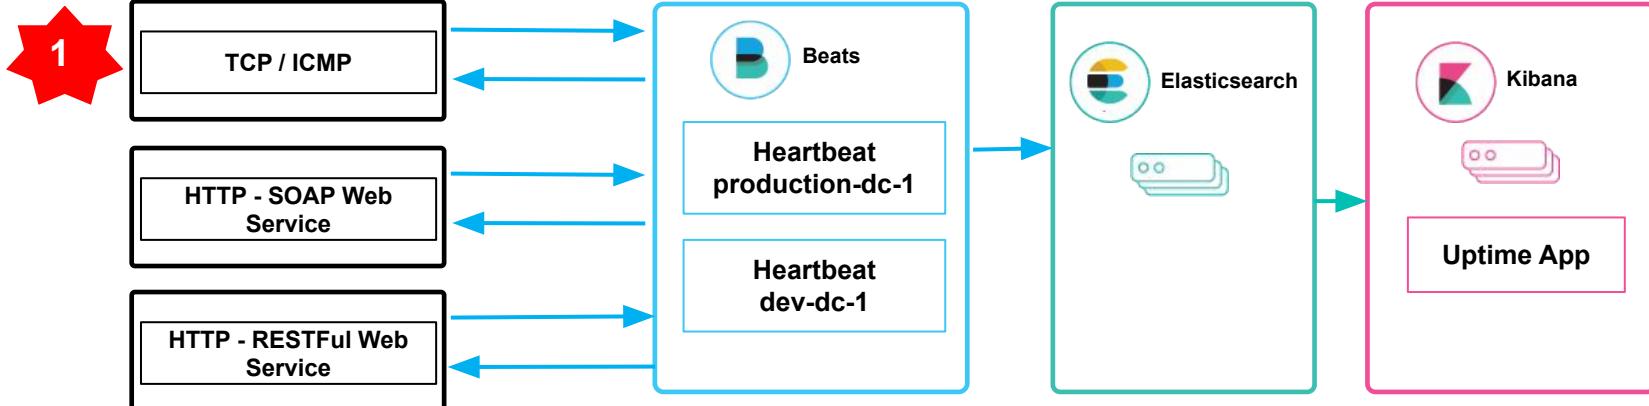
Demo Architecture

Elastic Stack



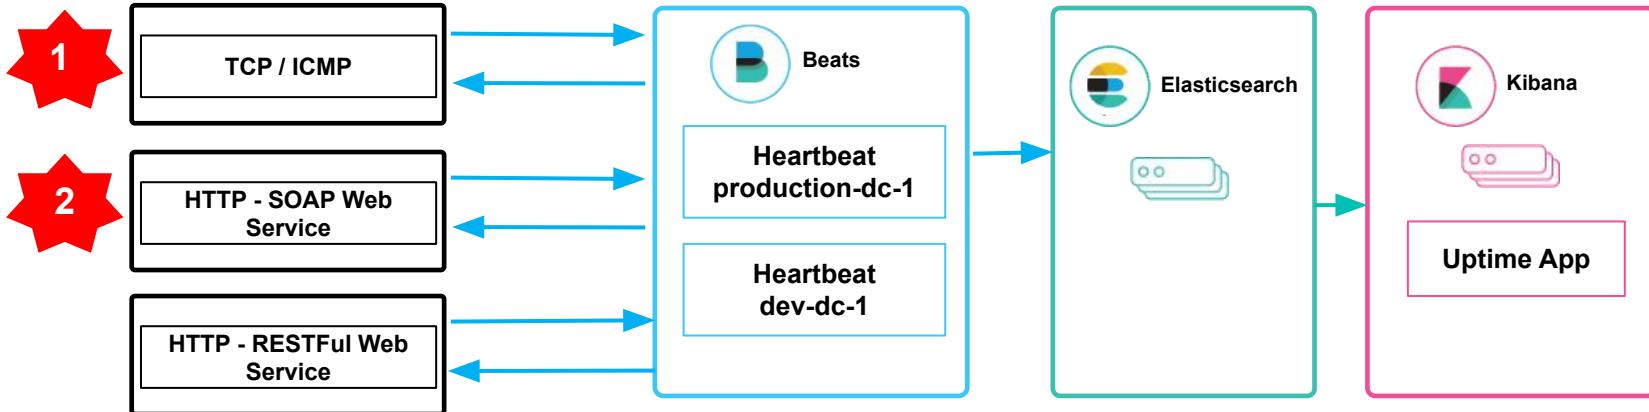
 Elasticsearch Service

Elastic Stack



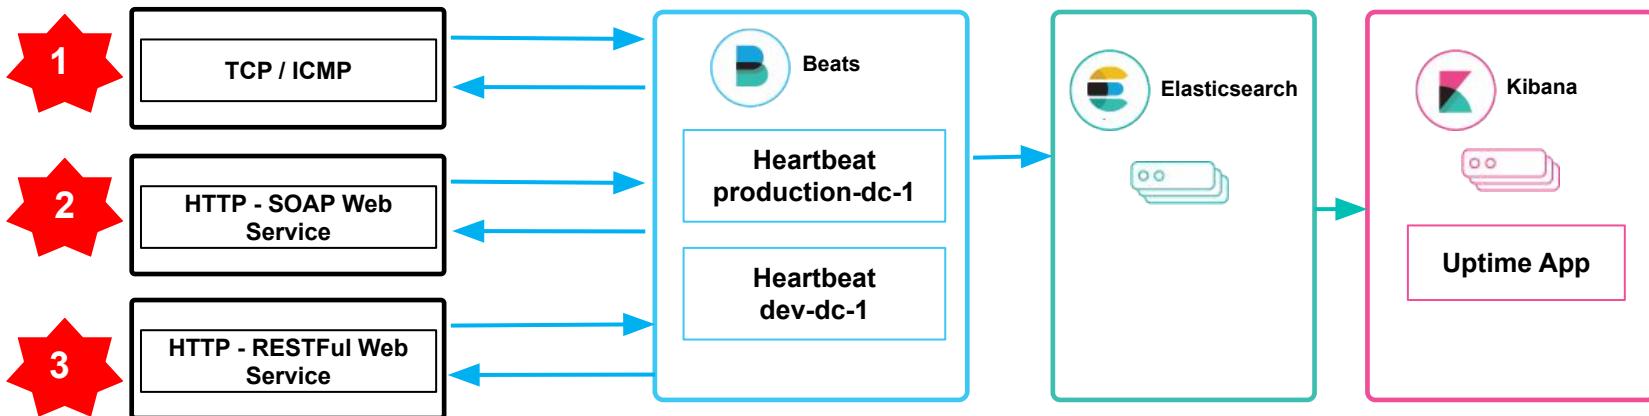
Elasticsearch Service

Elastic Stack



Elasticsearch Service

Elastic Stack



Elasticsearch Service

Elasticsearch Service

Hosted Elasticsearch, from the creators. No one hosts it better.

Deployments

web-prod

Edit

US West 1 (Oregon)

Data 1 configuration

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

gcp.data.highio.1 Data Ingest Master

An I/O optimized Elasticsearch instance.

You are at risk of data loss with fault tolerance set to a single zone. [Learn more](#)

Fault tolerance

1 zone 2 zones 3 zones

RAM per Node

Nodes: 1 RAM per Zone: 16 GB

Summary: 16 GB RAM x 1 node x 1 zone = 16 GB RAM, 480 GB storage

User setting overrides

gcp.data.highstorage.1 Data Ingest Master

A storage optimized Elasticsearch instance.

Enable

Allows the storage of indices that require less frequent querying on warm configurations, saving on professing costs. [Learn more](#)

Summary

Name	web-prod
Version	v7.4.0
ES data memory	16 GB
ES data storage	480 GB
Total memory	21 GB
Total storage	490 GB
Hourly rate	\$0.5776

Architecture

Zone 1

gcp.data.highio.1 (data)
16 GB RAM x 1, 480 GB storage x 1

gcp.kibana.1 (kibana)
1 GB RAM x 1

gcp.ml.1 (ml)
4 GB RAM x 1

Installation & Start

Downloading Heartbeat

```
wget https://artifacts.elastic.co/downloads/beats/heartbeat/heartbeat-7.5.1-darwin-x86\_64.tar.gz
(Other operating environments)
# wget https://artifacts.elastic.co/downloads/beats/heartbeat/heartbeat-7.5.1-linux-x86\_64.tar.gz
# wget https://artifacts.elastic.co/downloads/beats/heartbeat/heartbeat-7.5.1-windows-x86\_64.zip

tar zxvf heartbeat-7.5.1-darwin-x86_64.tar.gz
cd heartbeat-7.5.1-darwin-x86_64
```

Configure heartbeat monitors

```
heartbeat.monitors:
- type: http
  urls: ["https://httpbin.org"]
  schedule: '@every 30s'

# set name
- type: http
  urls: ["https://httpbin.org"]
  schedule: '@every 30s'
  name: "my-other-httpbin-check"

# add tags
- type: http
  urls: ["https://httpbin.org"]
  schedule: '@every 30s'
  tags: "my-httpbin-tag"

# always fails
- type: http
  urls: ["https://httpbin.org/status/404"]
  check.request.method: "POST"
  schedule: '@every 30s'
```

Configure heartbeat monitors

```
# expect 404 status
- type: http
  urls: ["https://httpbin.org/status/404"]
  schedule: '@every 30s'
  check.request.method: "POST"
  check.response.status: 404

# send a certain body
- type: http
  urls: ["https://httpbin.org/anything"]
  schedule: '@every 30s'
  check.request.method: "POST"
  check.send: 'Hello World'

# expect a certain body
- type: http
  urls: ["https://httpbin.org/anything>Hello+World"]
  schedule: '@every 30s'
  check.request.method: "POST"
  check.receive: 'Hello World'
```

Configure heartbeat monitors

```
# auth test
- type: http
  urls: ["https://httpbin.org/basic-auth/my_user/my_password"]
  schedule: '@every 30s'
  username: "my_user"
  password: "my_password"
  check.request.method: "POST"

# check response headers
- type: http
  urls: ["https://httpbin.org/response-headers?freeform=my_value"]
  schedule: '@every 30s'
  check.response.headers:
    'freeform': 'my_value'
```

Configure heartbeat monitors

```
# ICMP test against google DNS server
- type: icmp
  schedule: '@every 30s'
  hosts: [ '8.8.8.8' ]

# TCP test against google DNS server
- type: tcp
  schedule: '@every 30s'
  hosts: [ 'tcp://8.8.8.8:53' ]
```

Demo Time!



Thank You

- Web : www.elastic.co
- Demos: demo.elastic.co
- Products : <https://www.elastic.co/products>
- Forums : <https://discuss.elastic.co>
- Community : <https://www.elastic.co/community/meetups>
- Twitter : @elastic





Questions?





Thanks Again!!!



