

↳ Convenciones:

En todos los nodos como 'sudo su'.
[root@srv1 ~]# Solo en servidor 'srv1' → como 'sudo su'.
[root@srv2 ~]# Solo en servidor 'srv2' → como 'sudo su'.

→ Escenario Principal -Anterior- :**=> /etc/hosts:****192.168.10.150 pgsq1-vip****192.168.10.160 pgsq1-alt-vip****192.168.10.151 srv1****192.168.10.152 srv2****192.168.10.161 srv1-alt****192.168.10.162 srv2-alt****=> Pacemaker => Resources => Active/Passive - DRBD:****Resource Group: vipGROUP****pgsq1-vip (ocf::heartbeat:IPaddr2): Started srv1****pgsq1-alt-vip (ocf::heartbeat:IPaddr2): Started srv1****Master/Slave Set: drbdctrlDataClone [drbdctrlData]****Masters: [srv1]****Slaves: [srv2]****Master/Slave Set: pgsq1DataClone [pgsq1Data]****Masters: [srv1]****Slaves: [srv2]****Master/Slave Set: odooDataClone [odooData]****Masters: [srv1]****Slaves: [srv2]****Resource Group: pgsq1GROUP****pgsq1FS (ocf::heartbeat:Filesystem): Started srv1****odooFS (ocf::heartbeat:Filesystem): Started srv1****pgsq1DB (ocf::heartbeat:pgsql): Started srv1****Clone Set: pgadmin4HTTP-clone [pgadmin4HTTP]****Started: [srv1]**

Stopped: [srv2]

→ [Paquetería Java y Configuración previa:](#)

```
# yum install java-1.8.0-openjdk java-1.8.0-opendk-devel -y
# cat > /etc/profile.d/java8.sh <<FIN
export JAVA_HOME=$(dirname $(dirname $(readlink $(readlink $(which javac))))))
export PATH=$PATH:$JAVA_HOME/bin
export CLASSPATH=.:$JAVA_HOME/jre/lib:$JAVA_HOME/lib:$JAVA_HOME/lib/tools.jar
FIN
# cat /etc/profile.d/java8.sh
export JAVA_HOME=/usr/lib/jvm/java-1.8.0-openjdk-1.8.0.191.b12-1.el7_6.x86_64
export PATH=$PATH:$JAVA_HOME/bin
export CLASSPATH=.:$JAVA_HOME/jre/lib:$JAVA_HOME/lib:$JAVA_HOME/lib/tools.jar
[root@srv1 ~]# java -version
openjdk version "1.8.0_191"
OpenJDK Runtime Environment (build 1.8.0_191-b12)
OpenJDK 64-Bit Server VM (build 25.191-b12, mixed mode)
```

→ [Instalación de -elasticsearch- y configuración básica:](#)

```
# cat > /etc/yum.repos.d/elasticsearch.repo <<FIN
[elasticsearch-6.x]
name=Elasticsearch repository for 6.x packages
baseurl=https://artifacts.elastic.co/packages/6.x/yum
gpgcheck=1
gpgkey=https://artifacts.elastic.co/GPG-KEY-elasticsearch
enabled=1
autorefresh=1
type=rpm-md
FIN
# firewall-cmd --permanent --add-port=9200/tcp
# firewall-cmd --reload
# yum install elasticsearch -y
```

```
# systemctl enable --now elasticsearch.service
# curl http://127.0.0.1:9200
{
  "name" : "3iCb6ca",
  "cluster_name" : "elasticsearch",
  "cluster_uuid" : "vK6F7z21SoCFdvPOz9J3jw",
  "version" : {
    "number" : "6.5.4",
    "build_flavor" : "default",
    "build_type" : "rpm",
    "build_hash" : "d2ef93d",
    "build_date" : "2018-12-17T21:17:40.758843Z",
    "build_snapshot" : false,
    "lucene_version" : "7.5.0",
    "minimum_wire_compatibility_version" : "5.6.0",
    "minimum_index_compatibility_version" : "5.0.0"
  },
  "tagline" : "You Know, for Search"
}
```

```
[root@srv1 ~]# vim /etc/elasticsearch/elasticsearch.yml
```

```
...
bootstrap.memory_lock: true
network.host: 192.168.10.151
http.port: 9200
```

```
[root@srv2 ~]# vim /etc/elasticsearch/elasticsearch.yml
```

```
...
bootstrap.memory_lock: true
network.host: 192.168.10.152
http.port: 9200
```

```
...
# vim /usr/lib/systemd/system/elasticsearch.service
```

```
...
```

```
# Añadimos.
LimitMEMLOCK=infinity
...
# vim /etc/sysconfig/elasticsearch
...
MAX_LOCKED_MEMORY=unlimited
...
# systemctl daemon-reload
# systemctl restart elasticsearch.service
[root@srv1 ~]# netstat -plntpu
...
tcp6    0    0 192.168.10.151:9200 :::*          LISTEN       12145/java
...
[root@srv2 ~]# netstat -plntpu
...
tcp6    0    0 192.168.10.152:9200 :::*          LISTEN       12145/java
...
# curl -XGET '192.168.10.151:9200/_nodes?filter_path=**.mlockall&pretty'
{
  "nodes" : {
    "3iCb6cabR6-pooF-IxVV8w" : {
      "process" : {
        "mlockall" : true
      }
    }
  }
}
# curl -XGET '192.168.10.151:9200/?pretty'
{
  "name" : "3iCb6ca",
  "cluster_name" : "elasticsearch",
  "cluster_uuid" : "vK6F7z21SoCFdvPOz9J3jw",
  "version" : {
```

```
"number" : "6.5.4",
"build_flavor" : "default",
"build_type" : "rpm",
"build_hash" : "d2ef93d",
"build_date" : "2018-12-17T21:17:40.758843Z",
"build_snapshot" : false,
"lucene_version" : "7.5.0",
"minimum_wire_compatibility_version" : "5.6.0",
"minimum_index_compatibility_version" : "5.0.0"
},
"tagline" : "You Know, for Search"
}
```

```
[root@srv1 ~]# vim /etc/elasticsearch/elasticsearch.yml
```

```
...
```

```
cluster.name: cluster-elk
```

```
node.name: srv1
```

```
...
```

```
[root@srv2 ~]# vim /etc/elasticsearch/elasticsearch.yml
```

```
...
```

```
cluster.name: cluster-elk
```

```
node.name: srv2
```

```
...
```

```
# systemctl restart elasticsearch.service
```

```
# curl -XGET '192.168.10.151:9200/?pretty'
```

```
{
```

```
  "name" : "srv1",
```

```
  "cluster_name" : "cluster-elk",
```

```
  "cluster_uuid" : "vK6F7z21SoCFdvPOz9J3jw",
```

```
  "version" : {
```

```
    "number" : "6.5.4",
```

```
    "build_flavor" : "default",
```

```
    "build_type" : "rpm",
```

```
    "build_hash" : "d2ef93d",
```

```
"build_date" : "2018-12-17T21:17:40.758843Z",
"build_snapshot" : false,
"lucene_version" : "7.5.0",
"minimum_wire_compatibility_version" : "5.6.0",
"minimum_index_compatibility_version" : "5.0.0"
},
"tagline" : "You Know, for Search"
}
# curl -XGET '192.168.10.152:9200/?pretty'
{
  "name" : "srv2",
  "cluster_name" : "cluster-elk",
  "cluster_uuid" : "w0pkMHHiT92mMWSmPHTD6w",
  "version" : {
    "number" : "6.5.4",
    "build_flavor" : "default",
    "build_type" : "rpm",
    "build_hash" : "d2ef93d",
    "build_date" : "2018-12-17T21:17:40.758843Z",
    "build_snapshot" : false,
    "lucene_version" : "7.5.0",
    "minimum_wire_compatibility_version" : "5.6.0",
    "minimum_index_compatibility_version" : "5.0.0"
  },
  "tagline" : "You Know, for Search"
}
```

[→ Instalación y configuración de Kibana.](#)

```
# firewall-cmd --permanent --add-port=5601/tcp
```

```
# firewall-cmd --reload
```

```
# yum install kibana -y
```

```
[root@srv1 ~]# cat /etc/kibana/kibana.yml |grep -v "#"
```

```
server.port: 5601
```

```
server.host: "192.168.10.151"
```

```
server.name: "kibana-srv1"
```

```
elasticsearch.url: "http://192.168.10.151:9200"
```

```
[root@srv2 ~]# cat /etc/kibana/kibana.yml |grep -v "#"
```

```
server.port: 5601
```

```
server.host: "192.168.10.152"
```

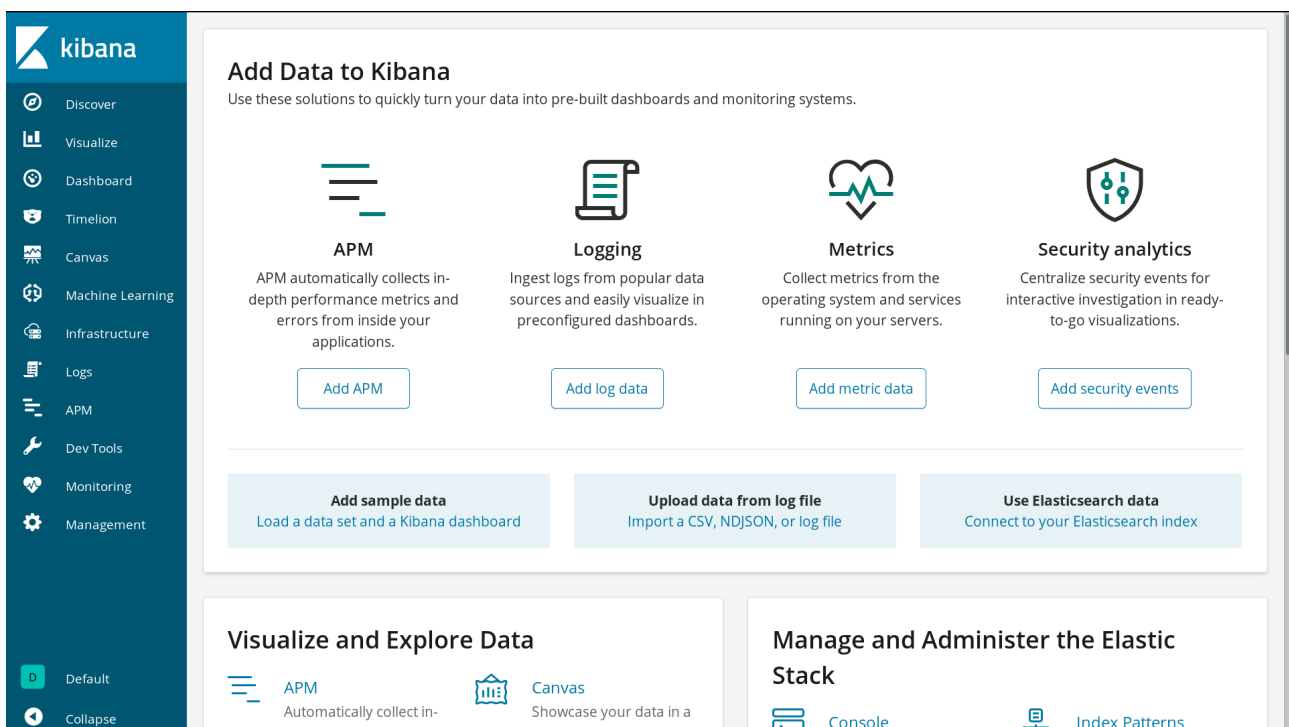
```
server.name: "kibana-srv2"
```

```
elasticsearch.url: "http://192.168.10.152:9200"
```

```
# systemctl enable --now kibana.service
```

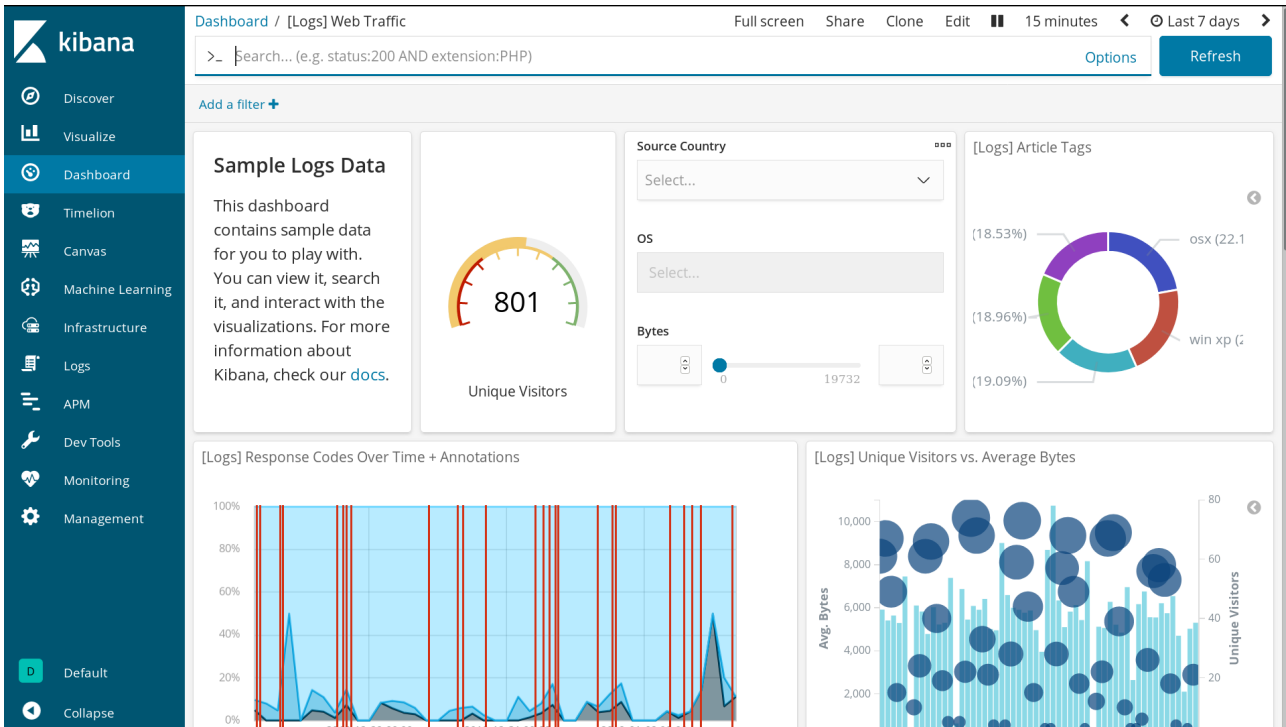
<http://192.168.10.151:5601>

<http://192.168.10.152:5601>



The screenshot shows the Kibana dashboard interface. On the left is a dark blue sidebar with the 'kibana' logo and a navigation menu with items: Discover, Visualize, Dashboard, Timelion, Canvas, Machine Learning, Infrastructure, Logs, APM, Dev Tools, Monitoring, and Management. Below the menu are 'Default' and 'Collapse' options. The main content area is titled 'Add Data to Kibana' and includes a subtitle: 'Use these solutions to quickly turn your data into pre-built dashboards and monitoring systems.' There are four primary data sources: APM (Automatically collect in-depth performance metrics and errors from inside your applications), Logging (Ingest logs from popular data sources and easily visualize in preconfigured dashboards), Metrics (Collect metrics from the operating system and services running on your servers), and Security analytics (Centralize security events for interactive investigation in ready-to-go visualizations). Each source has an 'Add' button. Below these are three secondary options: 'Add sample data' (Load a data set and a Kibana dashboard), 'Upload data from log file' (Import a CSV, NDJSON, or log file), and 'Use Elasticsearch data' (Connect to your Elasticsearch index). The bottom section is split into two: 'Visualize and Explore Data' (with APM and Canvas options) and 'Manage and Administer the Elastic Stack' (with Console and Index Patterns options).

↳ Se utiliza la Base de Datos de Ejemplo de Kibana, un mundo ...



→ Despliegue de cabeceras => elasticsearch-head. => NO OK.

```
# firewall-cmd --permanent --add-port=9100/tcp
```

```
# firewall-cmd --reload
```

↳ Continuar en esta parte creando índices, buckets y alias previos para -elasticsearch-...

```
# git clone git://github.com/mobz/elasticsearch-head.git
```

```
# cd elasticsearch-head/
```

```
# npm install
```

```
# npm run start
```

```
@open@ "http://localhost:9100/":http://localhost:9100/
```

```
# open http://192.168.10.152:9100
```

This will start a local webserver running on port 9100 serving elasticsearch-head

→ Despliegue LogStash ==> NO INICIADO => NO OK.

REFERENCIAS:

<https://tecadmin.net/install-java-8-on-centos-rhel-and-fedora/>

<https://www.ochobitshacenunbyte.com/2018/09/12/instalar-elasticsearch-en-centos-7/>

<https://gryzli.info/2018/03/30/elasticsearch-installation-how-to/>

<https://tecadmin.net/install-elasticsearch-on-linux/>

<https://computingforgeeks.com/install-elasticsearch-6-centos-7-kibana/>

<https://www.howtoforge.com/tutorial/how-to-install-elastic-stack-on-centos-7/>

https://www.server-world.info/en/note?os=CentOS_7&p=elasticstack6&f=1

<https://www.elastic.co/guide/index.html>

<https://computingforgeeks.com/easy-way-to-delete-elasticsearch-index-data/>

<https://www.elastic.co/es/downloads>

wget <https://artifacts.elastic.co/downloads/elasticsearch/elasticsearch-6.5.4.rpm>

<https://mobz.github.io/elasticsearch-head/>

Creative Commons

Reconocimiento-NoComercial-CompartirIgual 3.1 ESPAÑA

© 2019 by carlos briso. Usted es libre de copiar, distribuir y comunicar públicamente la obra y hacer obras derivadas bajo las condiciones siguientes:

a) Debe reconocer y citar al autor original.

b) No puede utilizar esta obra para fines comerciales (incluyendo su publicación, a través de cualquier medio, por entidades con fines de lucro.

c) Si altera o transforma esta obra o genera una obra derivada, sólo puede distribuir la obra generada bajo una licencia idéntica a ésta. Al reutilizar o distribuir la obra, tiene que dejar bien claro los términos de la licencia de esta obra.

Alguna de estas condiciones puede no aplicarse si se obtiene el permiso del titular de los derechos de autor. Los derechos derivados de usos legítimos u otras limitaciones no se ven afectados por lo anterior. Licencia completa en castellano.

→ La información contenida en este documento y los derivados de éste se proporcionan tal cual son y los autores no asumirán responsabilidad alguna si el usuario o lector hace mal uso de éstos.