



Installazione presa da un nodo di test 3D (IRONMAN)

Dopo aver aggiornato il sistema operativo ed installato i software (JDK ed Elastic) procediamo come segue

Limits.conf

```
...
elasticsearch - nofile 65536
elasticsearch - nproc 2048
elasticsearch soft memlock unlimited
elasticsearch hard memlock unlimited
#elasticsearch - memlock unlimited
```

Elastic

jvm.options

Esempio

```
## JVM configuration

#####
## IMPORTANT: JVM heap size
#####
##
## You should always set the min and max JVM heap
## size to the same value. For example, to set
## the heap to 4 GB, set:
##
## -Xms4g
## -Xmx4g
##
## See https://www.elastic.co/guide/en/elasticsearch/reference/current/heap-size.html
## for more information
##
#####

# Xms represents the initial size of total heap space
# Xmx represents the maximum size of total heap space

-Xms1g
-Xmx1g

#####
## Expert settings
#####
##
## All settings below this section are considered
## expert settings. Don't tamper with them unless
## you understand what you are doing
##
#####

## GC configuration
-XX:+UseConcMarkSweepGC
-XX:CMSInitiatingOccupancyFraction=75
-XX:+UseCMSInitiatingOccupancyOnly

## optimizations

# disable calls to System#gc
-XX:+DisableExplicitGC

# pre-touch memory pages used by the JVM during initialization
-XX:+AlwaysPreTouch

## basic
```



```
# force the server VM (remove on 32-bit client JVMs)
-server

# explicitly set the stack size (reduce to 320k on 32-bit client JVMs)
-Xss1m

# set to headless, just in case
-Djava.awt.headless=true

# ensure UTF-8 encoding by default (e.g. filenames)
-Dfile.encoding=UTF-8

# use our provided JNA always versus the system one
-Djna.nosys=true

# use old-style file permissions on JDK9
-Djdk.io.permissionsUseCanonicalPath=true

# flags to configure Netty
-Dio.netty.noUnsafe=true
-Dio.netty.noKeySetOptimization=true
-Dio.netty.recycler.maxCapacityPerThread=0

# log4j 2
-Dlog4j.shutdownHookEnabled=false
-Dlog4j2.disable.jmx=true
-Dlog4j.skipJansi=true

## heap dumps
# PROXY Options
#-Dhttp.proxyHost=proxy.equitaliaservizi.org
#-Dhttp.proxyPort=80
#-Dhttps.proxyHost=proxy.equitaliaservizi.org
#-Dhttps.proxyPort=80

# generate a heap dump when an allocation from the Java heap fails
# heap dumps are created in the working directory of the JVM
-XX:+HeapDumpOnOutOfMemoryError

# specify an alternative path for heap dumps
# ensure the directory exists and has sufficient space
#-XX:HeapDumpPath=${heap.dump.path}

## GC logging

#-XX:+PrintGCDetails
#-XX:+PrintGCTimeStamps
#-XX:+PrintGCDateStamps
#-XX:+PrintClassHistogram
#-XX:+PrintTenuringDistribution
#-XX:+PrintGCApplicationStoppedTime

# log GC status to a file with time stamps
# ensure the directory exists
#-Xloggc:${loggc}

# By default, the GC log file will not rotate.
# By uncommenting the lines below, the GC log file
# will be rotated every 128MB at most 32 times.
#-XX:+UseGCLogFileRotation
#-XX:NumberOfGCLogFiles=32
#-XX:GCLogFileSize=128M

# Elasticsearch 5.0.0 will throw an exception on unquoted field names in JSON.
```



```
# If documents were already indexed with unquoted fields in a previous version
# of Elasticsearch, some operations may throw errors.
#
# WARNING: This option will be removed in Elasticsearch 6.0.0 and is provided
# only for migration purposes.
#-Elasticsearch.json.allow_unquoted_field_names=true
```



La memoria assegnata ad Elastic non deve superare il 50% della memoria totale disponibile.

elasticsearch.yml

Esempio

```
# ===== Elasticsearch Configuration =====
#
# NOTE: Elasticsearch comes with reasonable defaults for most settings.
#       Before you set out to tweak and tune the configuration, make sure you
#       understand what are you trying to accomplish and the consequences.
#
# The primary way of configuring a node is via this file. This template lists
# the most important settings you may want to configure for a production cluster.
#
# Please consult the documentation for further information on configuration options:
# https://www.elastic.co/guide/en/elasticsearch/reference/index.html
#
# ----- Cluster -----
#
# Use a descriptive name for your cluster:
#
cluster.name: 3delk
#
# ----- Node -----
#
# Use a descriptive name for the node:
#
node.name: ironman
#
# Add custom attributes to the node:
#
#node.attr.rack: r1
#
# ----- Paths -----
#
# Path to directory where to store the data (separate multiple locations by comma):
#
#path.data: /path/to/data
#
# Path to log files:
#
#path.logs: /path/to/logs
#
# ----- Memory -----
#
# Lock the memory on startup:
#
bootstrap.memory_lock: true
#
# Make sure that the heap size is set to about half the memory available
# on the system and that the owner of the process is allowed to use this
# limit.
#
# Elasticsearch performs poorly when the system is swapping the memory.
#
# ----- Network -----
```



```
#
# Set the bind address to a specific IP (IPv4 or IPv6):
#
network.host: 0.0.0.0
#
# Set a custom port for HTTP:
#
http.port: 9200
transport.tcp.port: 9300
network.publish_host: 192.168.200.103
#
# For more information, consult the network module documentation.
#
# ----- Discovery -----
#
# Pass an initial list of hosts to perform discovery when new node is started:
# The default list of hosts is ["127.0.0.1", ":::1"]
#
discovery.zen.ping.unicast.hosts:
["192.168.200.103:9300", "192.168.1.104:9300", "192.168.1.105:9300"]
#
# Prevent the "split brain" by configuring the majority of nodes (total number of master-eligible
nodes / 2 + 1):
#
discovery.zen.minimum_master_nodes: 2
#
# For more information, consult the zen discovery module documentation.
#
# ----- Gateway -----
#
# Block initial recovery after a full cluster restart until N nodes are started:
#
#gateway.recover_after_nodes: 3
#
# For more information, consult the gateway module documentation.
#
# ----- Various -----
#
# Require explicit names when deleting indices:
#
#action.destructive_requires_name: true
node.master: true
node.data: true
```

Le cose importanti sono:

- Nome del cluster;
- nome del nodo;
- bootstrap memory lock;
- discovery unicast;
- node.master;
- node.data;
- max locked memory.

/usr/lib/systemd/system/elasticsearch.service

```
...
# Specifies the maximum number of bytes of memory that may be locked into RAM
# Set to "infinity" if you use the 'bootstrap.memory_lock: true' option
# in elasticsearch.yml and 'MAX_LOCKED_MEMORY=unlimited' in /etc/sysconfig/elasticsearch
LimitMEMLOCK=infinity
...
```

/etc/sysconfig/elasticsearch



```
...
# The maximum number of bytes of memory that may be locked into RAM
# Set to "unlimited" if you use the 'bootstrap.memory_lock: true' option
# in elasticsearch.yml.
# When using Systemd, the LimitMEMLOCK property must be set
# in /usr/lib/systemd/system/elasticsearch.service
MAX_LOCKED_MEMORY=unlimited
...
```