



JAVA, Docker, Kubernetes & OpenShift @ ISRIC



Jorge Mendes de Jesus
Paul van Genuchten

Geonetwork day 23 May



- **GeoCat B.V.** - Located in Bennekom, NL
- Jeroen Ticheler (CEO) creator of Geonetwork metadata catalogue
- **Geocat B.V.**, Major business on Geonetwork support, SDI on cloud and bridge between ArcGIS (ESRI) and OpenSource

http://nationaalgeoregister.nl

NGR Nationaal Georegister

Contact Help

Home Zoeken Kaart Actueel Over NGR Voor ontwikkelaars Inloggen

dé vindplaats van geo informatie van heel Nederland

Zoek in 7137 datasets, services en kaarten

Welk onderwerp? Op welke locatie?

Zoeken

Zoeken

Bekijk één van de categorieën

(civiele) structuren (414)	binnenwater (594)	economie (289)
geo wetenschappelijke data (362)	gezondheid (167)	grenzen (627)
hoogte (263)	klimatologie, meteorologie atmosfeer (147)	landbouw en veeteelt (264)
locatie (1022)	maatschappij (586)	militair (41)
natuur en milieu (2368)	nutsbedrijven communicatie (132)	oceanen (120)
planning kadaster (660)	referentie materiaal aardbedekking (219)	transport (867)

Privacy Cookies Copyright





ISRIC
World Soil Information

- **ISRIC** - World Soil Reference is located in Wageningen
- It has a mission to serve the international community with information about the world's soil.
- Spatial Data Infrastructures orient towards soil data serving
- SDI providing metadata catalogues, web services, REST

https://data.isric.org

The screenshot shows the top navigation bar of the ISRIC Soil Metadata Catalogue. It includes the site logo, 'Soil Metadata Catalogue', search and map icons, a 'Sign in' button, and a language dropdown set to 'English'. Below the navigation is a search bar with the placeholder text 'Search ...' and a search button. A message below the search bar states 'Search over 172 data sets, services and maps, ...'. The main content area is divided into 'Browse by topics' and 'Browse resources'. Under 'Browse by topics', there is a button for 'Geoscientific information' with a '172' badge. Under 'Browse resources', there is a 'Dataset' button with a '172' badge. Below these are 'Latest news' and 'Most popular' buttons. At the bottom, there are four preview cards for datasets: 'WISE derived soil properties on a 30 by 30 arc-seconds glo...', 'WISE - Global Soil Profile Data, version 3.1', 'Global Assessment of Human-induced Soil Degradation (G...', and 'Soil and Terrain Database (SOTER) for China'. Each card shows a map or data visualization.




- **ISRIC** metadata catalogue is used for **dataset distribution and metadata record search**
- Using **Geonetwork 3.6.0** supported by GeoCat B.V.

- Tomcat8 servers running Geonetwork and Geoserver
- Several **horror stories of migration** between physical server providers (different EU countries).



- **Hard to update and upgrade physical servers** and docs and audit only using written documents



- Started to use docker containers to **deploy server instances on physical servers**
- Instead of doc files we have now **Dockerfiles and README.md on gitlab/github**



OPENSIFT



WAGENINGEN
UNIVERSITY & RESEARCH


- Still needed a platform for containers to run - **Cloud**
- OpenShift is a container platform based on kubernetes



<https://kubernetes.io>


Kubernetes (K8s) is an open-source system for automating deployment, scaling, and management of containerized applications.

It groups containers that make up an application into logical units for easy management and discovery. Kubernetes builds upon [15 years of experience of running production workloads at Google](#), combined with best-of-breed ideas and practices from the community.



Planet Scale


Designed on the same principles that allows Google to run billions of containers a week, Kubernetes can scale without increasing your ops team.



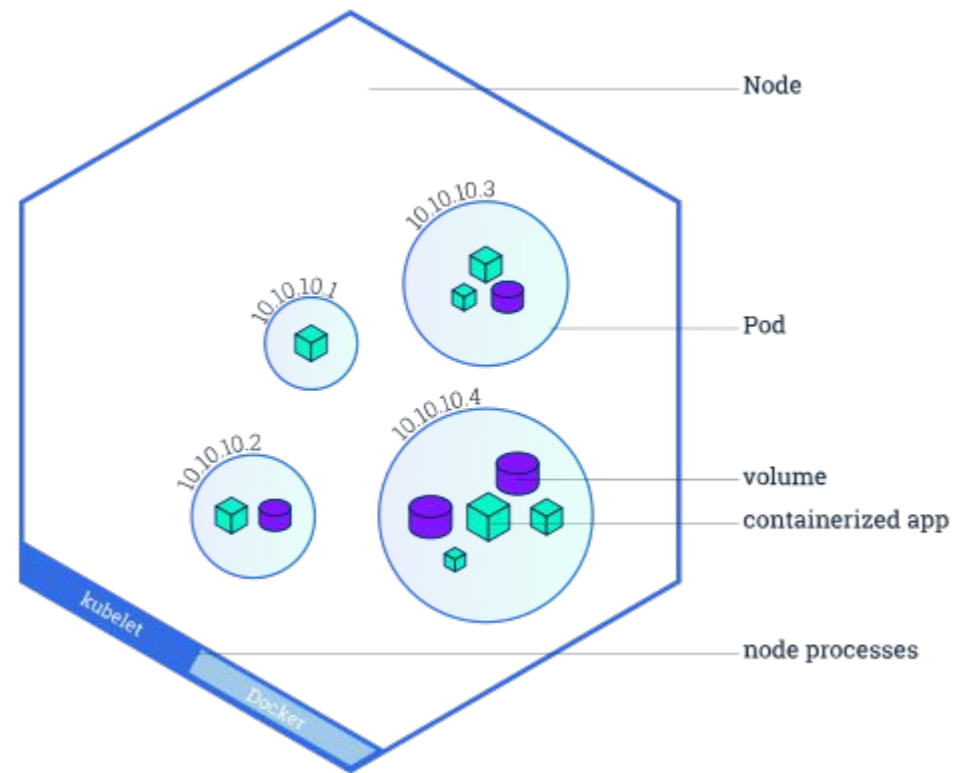
Versus




- K8 "rival" is Docker Swarm
- K8 can deploy other containers aside from Docker
- K8 supported by Google




- Pods are the **smallest unit** on kubernetes system.
- Pods contain container(s), volumes and internal network.
- Containers in **same pod share resources** (have access to)




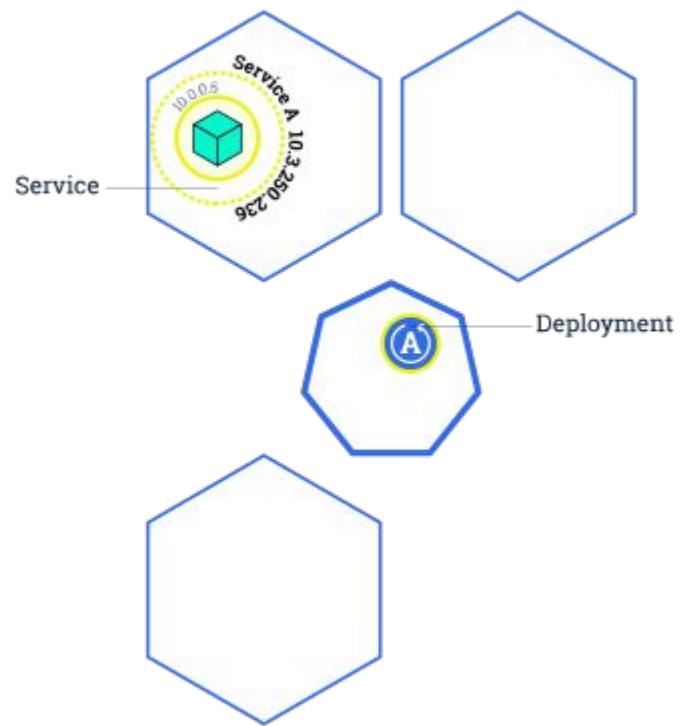
- Node is a VM or physical machine running the pods



- Nodes are grouped and monitored by resource schedulers
- Nodes and supporting management make a K8 cluster (Kube)



- Vertical versus Horizontal scaling
- Recommended is for horizontal scaling, also more resilient.



- K8 are totally oriented for horizontal scaling
- **Very easy to scale - extra pods and/or nodes**



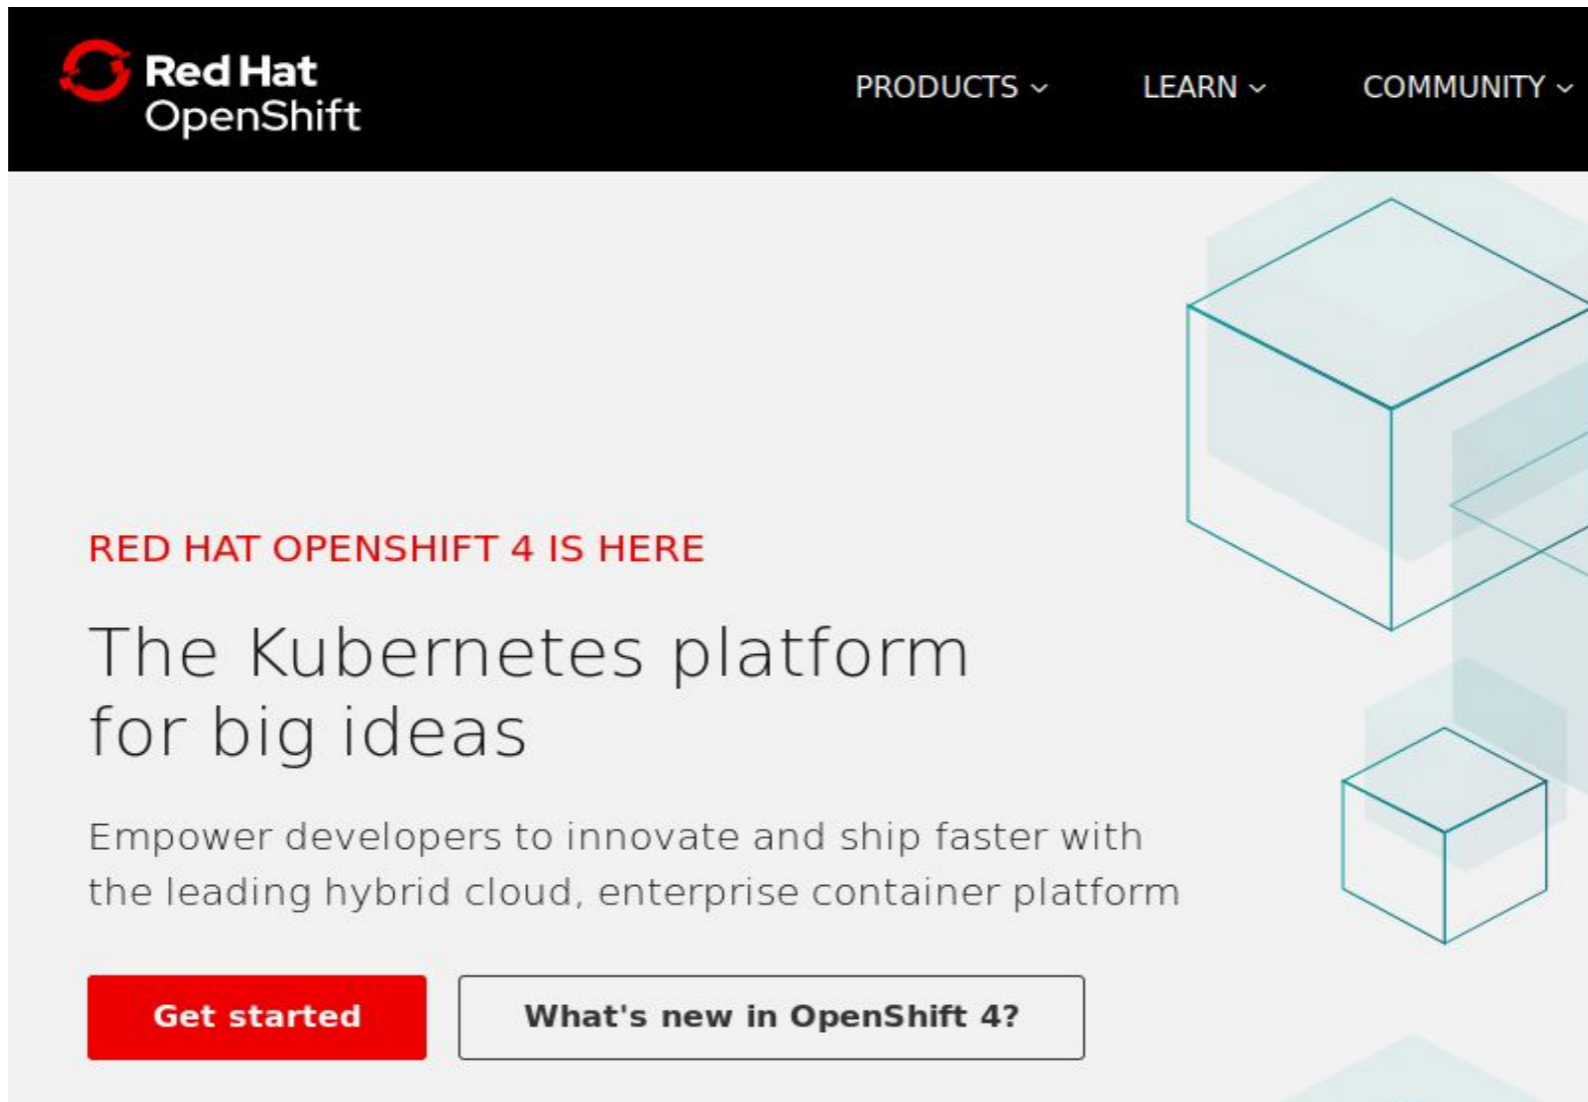
- K8 have a steep learning curve
- **Different systems and big number of yaml configuration files**

<https://bit.do/k8-deploy-example>

```
68 lines (67 sloc) | 1.34 KB
Raw Blame History
1  apiVersion: v1
2  kind: Service
3  metadata:
4    name: wordpress
5    labels:
6      app: wordpress
7  spec:
8    ports:
9      - port: 80
10   selector:
11     app: wordpress
12     tier: frontend
13   type: LoadBalancer
14   ---
15  apiVersion: v1
16  kind: PersistentVolumeClaim
17  metadata:
18    name: wp-pv-claim
19    labels:
20      app: wordpress
21  spec:
22    accessModes:
23      - ReadWriteOnce
24    resources:
25      requests:
26        storage: 20Gi
27  ---
```



<https://www.openshift.com/>



The screenshot shows the Red Hat OpenShift website homepage. At the top, there is a black navigation bar with the Red Hat OpenShift logo on the left and three menu items: 'PRODUCTS', 'LEARN', and 'COMMUNITY', each with a downward arrow. Below the navigation bar, the main content area features a light blue background with a large, stylized 3D cube graphic on the right side. The text on the page reads: 'RED HAT OPENSIFT 4 IS HERE' in red, followed by 'The Kubernetes platform for big ideas' in a large, dark font. Below this, a smaller line of text states: 'Empower developers to innovate and ship faster with the leading hybrid cloud, enterprise container platform'. At the bottom of the main content area, there are two buttons: a red button labeled 'Get started' and a white button with a black border labeled 'What's new in OpenShift 4?'.



- Container application platform
- Nice GUI to manage applications running on k8

OPENSIFT CONTAINER PLATFORM

Geonetwork

Search Catalog Add to Project

Name Filter by name List by Application

APPLICATION geonetwork <https://data2.isric.org/>

DEPLOYMENT CONFIG geonetwork, #42

3.0 Mib Memory -- Cores CPU -- Kib/s Network 1 pod

APPLICATION postgresql

DEPLOYMENT CONFIG postgresql, #13

60 Mib Memory < 0.01 Cores CPU 0.2 Kib/s Network 1 pod

- We have projects that manage, services, routes, resources, builds
- Project contains a geonetwork + postgresql instances


geonetwork-42-c8fxq created 3 minutes ago

Actions ▾

app geonetwork deployment geonetwork-42 deploymentconfig geonetwork

Details Environment Metrics Logs Terminal Events

Status

Status:  Running
Deployment: [geonetwork, #42](#)
IP: 10.130.2.29
Node: scomp1372.wurnet.nl (10.90.24.32)
Restart Policy: Always







Container geonetwork

State: Running since May 12, 2019 4:19:56 PM
Ready: true
Restart Count: 0

Template

Containers

geonetwork

 **Image:** [isric/geonetwork](#) c115558 562.1 MiB
 **Ports:** 8080/TCP
 **Mount:** volume-l17e1 → /usr/local/tomcat/webapps/geonetwork/WEB-INF/data read-write
 **Mount:** default-token-4lnm7 → /var/run/secrets/kubernetes.io/serviceaccount read-only
 **CPU:** 1300 millicores to 1400 millicores
 **Memory:** 3 GiB to 3 GiB

Volumes

volume-l17e1

Type: persistent volume claim (reference to a persistent volume claim)
Claim name: [geonetwork-data](#)
Mode: read-write

- Pod running a geonetwork image

Containers

geonetwork

-  **Image:** [isric/geonetwork](#) c115558 562.1 MiB
-  **Ports:** 8080/TCP
-  **Mount:** volume-l17e1 → /usr/local/tomcat/webapps/geonetwork/WEB-INF/data read-write
-  **CPU:** 1300 millicores to 1400 millicores
-  **Memory:** 3 GiB to 3 GiB


- We have full control CPU cycles and memory that container can use

Container: geonetwork Time Range: Last hour

About Compu



Memory

901 Available of 3072 MiB



CPU

1.39 Available of 1.4000000000000001 cores



Network




- Can monitor CPU/memory/Network

postgresql created 5 months ago

app postgresql

[History](#) Configuration Environment Events

 Deployment #13 is active. [View Log](#)
created 2 months ago

Filter by label	
Deployment	Status
#13 (latest)	 Active, 1 replica
#12	 Failed
#11	 Complete

- When we create a new container or change configurations we have a new deployment
- **It is possible to roll back to previous deployments**

postgresql created 5 months ago

app postgresql

[History](#) Configuration Environment Events

Deployment #13 is active. [View Log](#)
created 2 months ago

Filter by label

Deployment	Status
#13 (latest)	Active, 1 replica
#12	Failed
#11	Complete

- When we create a new container or change configurations we have a new deployment
- **It is possible to roll back to previous deployments (COOL FOR CATS)**

The screenshot displays the configuration for an external route in a Kubernetes dashboard. The left sidebar contains navigation options: Overview, Applications, Builds, Resources, Storage, Monitoring, and Catalog. The main content area shows the route details for 'geonetwork-external', which was created 3 months ago. The route is associated with the 'app' namespace and the 'geonetwork' service. The URL is 'https://data2.isric.org', and it is exposed on the 'router' service. The details section shows the path as '/', the service as 'geonetwork', and the target port as '8080-tcp'. A note indicates that this target port will route to Service Port 8080 → Container Port 8080 (TCP). The TLS settings section shows the termination type as 'Edge', insecure traffic as 'Redirect', and a certificate. The certificate content is displayed in a code block.

Routes > geonetwork-external

geonetwork-external created 3 months ago

app geonetwork

<https://data2.isric.org>

✓ Exposed on router 'router' 3 months ago

Details

Path: /

Service: [geonetwork](#)

Target Port: 8080-tcp

This target port will route to Service Port 8080 → Container Port 8080 (TCP).

TLS Settings

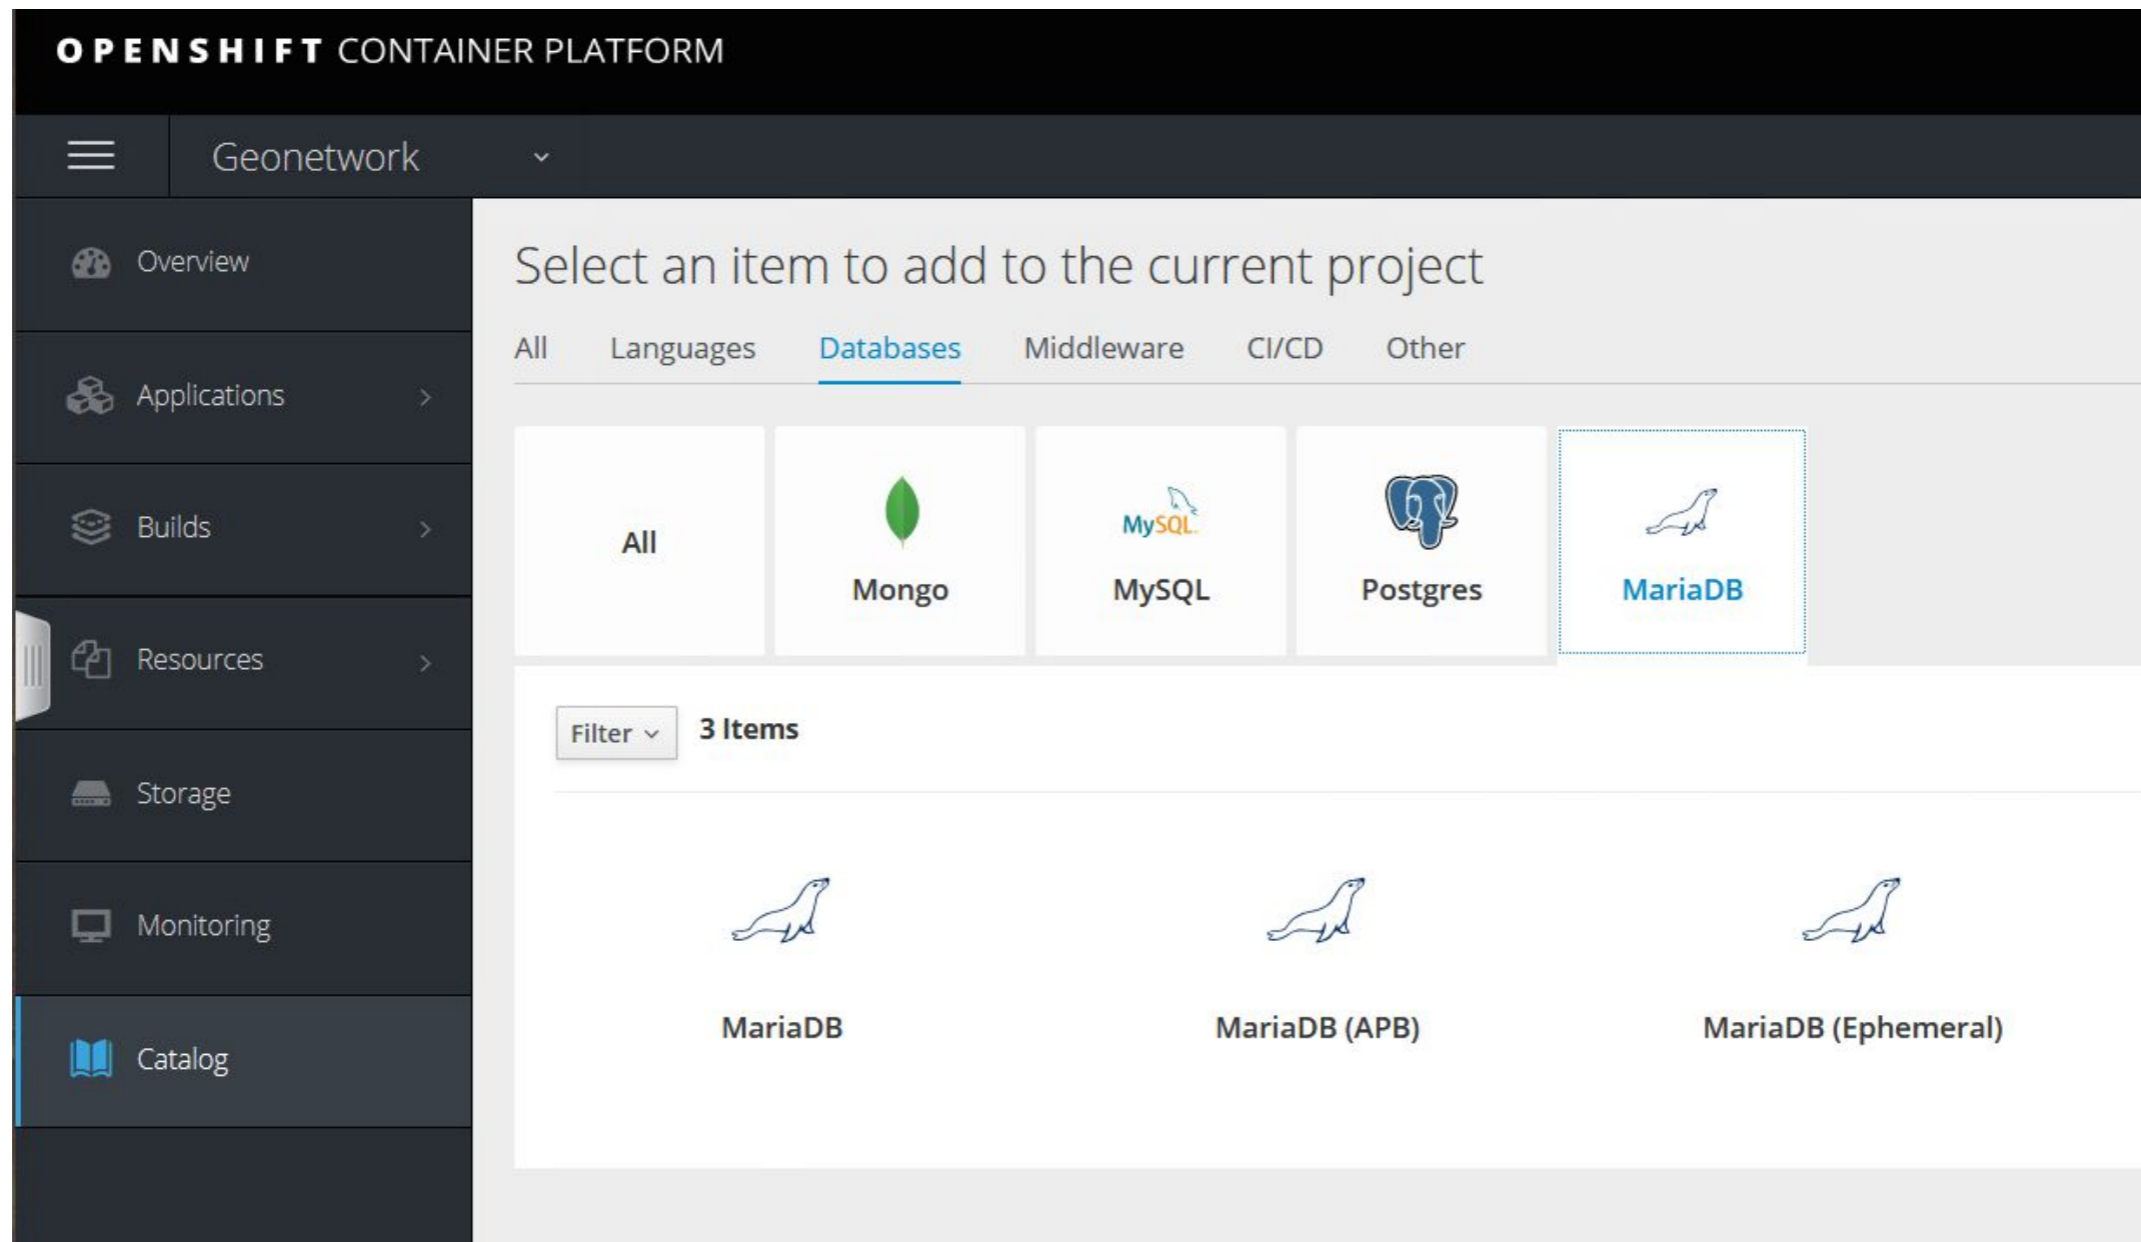
Termination Type: Edge

Insecure Traffic: Redirect

Certificate:

```
-----BEGIN CERTIFICATE-----  
MIIIXzCCB0egAwIBAgIRANaE+v1qGSQUod8upHe0LTgwDQYJKoZIhvcNAQELBQAw
```

- Router (HAproxy) can encrypt/control HTTPS traffic



- Possible to deploy containers from catalogue
- Containers optimized for OpenShift



- WUR Gitlab and Openshift are associated.
- Build a dockerfile extending Geonetwork image
- Current version: Geonetwork 3.6, Tomcat8.5, JRE8



geonetwork ☆

Docker Official Images

GeoNetwork is a FOSS catalog for spatially referenced resources.

↓ 1M+

Container Linux 386 ARM 64 ARM x86-64 IBM Z PowerPC 64 LE Application Services

Official Image

Linux - IBM Z (latest)

Copy and paste to pull this image

docker pull geonetwork

[View Available Tags](#)

- Geonetwork images on dockerhub made by GeoCat B.V
- GeoCat provides a service for automatic deployment of Geonetwork and Geoserver called GeoCat live

SDI as a Service

- Catalog with [CSW](#)
- Map Services with [WMS](#), [WFS](#), [WCS](#), [WPS](#)
- Spatial Database
- Hosting with [ISO/IEC 27001](#) certification
- Backups on a separated hosting with [ISO/IEC 27001](#) certification
- 24/7 monitoring* with automatic recovery
- GeoCat Bridge
- SSL connection
- High availability
- Quick response time
- Training: Online wiki/webinars, On-site in Europe
- Customised domain name
- High Data Allowance
- Enough storage for all your data
- INSPIRE Ready
- GDPR Compliant
- Customer Segmentation
- Security audited

What's included

Uptime

We provide provides a 99.95% uptime on our SLA for all our services. Furthermore, lost time is refunded back to your account.

Monitoring & Statistics

Review usage, uptime and incidents in the last months by downloading reports from <https://my.geocat.net>.

Also, we have an automatic recovery in place linked to monitoring to make sure the service is always fast and online.

Software updates

The software offered is carefully reviewed and major software releases will be installed after having been accepted by our testers. In addition, the software has nice features on top of the community releases of the software.

GeoCat Bridge

GeoCat Bridge is an Arcgis extension to publish data and metadata on open source servers. Along with GeoCat Live you get free licenses to publish data to your Live servers.

Custom domain name

Use your own DNS to access to our SDI as a Service.

INSPIRE Ready

Our services are fully INSPIRE compliant to make your life easier.



Container Registry

With the Docker Container Registry integrated into GitLab, every project can have its own space to store its Docker images.

Learn more about [Container Registry](#).

> [isric/geonetwork](#) 



> [isric/geonetwork/postgres](#) 



How to use the Container Registry

First log in to GitLab's Container Registry using your GitLab username and password. If you have [2FA enabled](#) you need to use a [personal access token](#):

```
docker login docker-registry.wur.nl
```

You can also use a [deploy token](#) for read-only access to the registry images.

Once you log in, you're free to create and upload a container image using the common `build` and `push` commands

```
docker build -t docker-registry.wur.nl/isric/geonetwork .  
docker push docker-registry.wur.nl/isric/geonetwork
```

- When we push an image to repository, the trigger make a new openshift deployment



- What have we learned :)
- Geonetwork consumes a lot of memory increasing the price of containers
- If there is not enough resources, pods get restarted **BUT users don't notice much**



- Good experience with redeployment in case of problems
- Take some time to set all configurations but when **it works**
- Increased speed by having more database pods



Let Op!!!

Pythonist commenting on JVM Options

I've had the time of my life and I've never felt this way before

Nobody puts Java in a container

January 2, 2018 Jörg Schad

#containers #java



[« In a microservices landscape; When do you update?](#)

[Part 1: OpenJ9 versus HotSpot »](#)



Written by Roy van Rijn (royvanrijn.com) on May 16, 2018 08:14:11 : [14 COMMENTS](#)

Java and Docker, the limitations

Java on Docker will no longer suck: improvements coming in Java 10

Wed 14 March 2018 - Richard Warburton & Sadiq Jaffer

- HOW DARE YOU!!! PUTTING JAVA ON A CONTAINER

- Major issue is that containers have memory/cpu resource limitation.

Memory 32 MiB min to 4 GiB max

Request



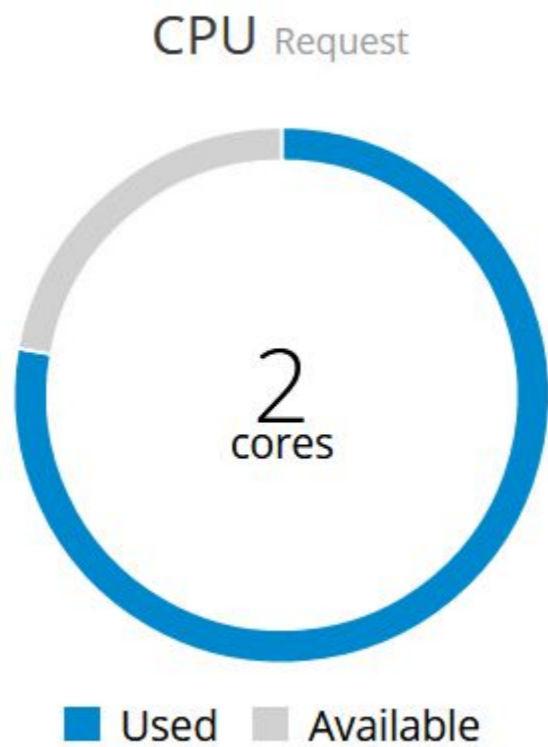
The minimum amount of memory the container is guaranteed.

Limit


The maximum amount of memory the container is allowed to use when running.

[What are MiB?](#)

- Major issue is that containers have memory/cpu resource limitation.



- JVM "sees" 4 Gigas of available memory and not the 3 Gigas allocated exclusively to Geonetwork





- JVM "sees" 4 Gigas of available memory and not the 3 Gigas allocated exclusively to Geonetwork
- Pods being restarted with no log information aside from the **OOMKilled (Out of Mememory Killed)**
- **OOMKilled** indicates a resource quota problem

- From Java 9 onwards and from 8u131+ onwards
- `-XX:+UnlockExperimentalVMOptions`
`-XX:+UseCGroupMemoryLimitForHeap`



- For JAVA 10 we no now have `-XX:UseContainerSupport`




- STILL OOMKill and strange memory leaks
- Some research until it was discovered:

-XX:+UnlockExperimentalVMOptions



- **-XX:+UnlockExperimentalVMOptions** not an option for Geonetwork but OK for Geoserver
- Solution to limit memory: `-XX:MaxRAM=3000m`
- Dockerfile `JAVA_OPTS` have to be in sync with resource allocation on OpenShift



APPLICATION

webdav

<https://files.isric.org>

>

DEPLOYMENT CONFIG

webdav, #47

760
Mib Memory

< 0.01
Cores CPU

59
Kib/s Network

1 pod


⋮

- ISRIC has a webDAV service for service files
- WebDAV running on Alpine OS + Apache2
- What is the size of the image?

webdav


 Image: [isric/webdav](#) f4c2bef 4.8 MiB

 Ports: 443/TCP

 Mount: webdav → /var/webdav read-only

 Mount: webdav, subpath general/public → /var/webdav_public read-write

 CPU: 756 millicores limit

 Memory: 756 MiB limit

4.8 Megas !!!!!

UBUNTU IS HEAVY !!!

Thank you!