Introduction to Docker & OpenShift

Neale Ferguson

2017-06-24

http://download.sinenomine.net/clefos/epel7/Getting_Started_with_OpenShift_on_z.pdf

Preface

- Examples built and run using ClefOS 7.3
  - CentOS Clone with name change
  - Available for z Systems
- However, as we will see this is irrelevant

What is Docker

SINE NOMINE ASSOCIATES

- An open source project to pack, ship and run any application as a lightweight container
- Container: self-contained receptacle
  - Filesystem
  - Apps
  - Static data
  - Network

Docker

UNDERLYING TECHNOLOGIES

cgroups…

SINE NOMINE ASSOCIATES

- A kernel feature that limits, accounts for, and isolates the resource of a collection of processes
- Similar to processes:
  - They are hierarchical
  - Child cgroups inherit certain attributes from their parent cgroup
- Difference: multiple cgroup hierarchies

…cgroups…

SINE NOMINE ASSOCIATES

- Can span multiple “subsystems”
  - blkio —sets limits on input/output access to and from block
  - cpu —uses the scheduler to provide cgroup tasks access to the CPU
  - cpuacct —generates automatic reports on CPU resources used by tasks in a cgroup.

…cgroups…

SINE NOMINE ASSOCIATES

- Can span multiple “subsystems”
  - cpuset —assigns individual CPUs (on a multicore system) and memory
  - devices —allows or denies access to devices by tasks in a cgroup
  - freezer —suspends or resumes tasks in a cgroup

…cgroups…

SINE NOMINE ASSOCIATES

- Can span multiple “subsystems”
  - memory —sets limits on memory use by tasks in a cgroup, & generates automatic reports on memory
  - net_cls —tags network packets with a class identifier (classid) that allows the Linux traffic controller (tc) to identify packets originating from a particular cgroup task.

…cgroups…

SINE NOMINE ASSOCIATES

- Can span multiple “subsystems”
  - net_prio —provides a way to dynamically set the priority of network traffic per network interface
  - ns — the namespace subsystem

Namespaces…

SINE NOMINE ASSOCIATES

- CLONE_NEWIPC: IPC Namespaces: SystemV IPC and POSIX Message Queues can be isolated.
- CLONE_NEWPID: PID Namespaces: PIDs are isolated, meaning that a virtual PID inside of the namespace can conflict with a PID outside of the namespace. PIDs inside the namespace will be mapped to other PIDs outside of the namespace. The first PID inside the namespace will be '1' which outside of the namespace is assigned to init

…Namespaces…

SINE NOMINE ASSOCIATES

- CLONE_NEWNET: Network Namespaces: Networking (/proc/net, IPs, interfaces and routes) are isolated. Services can be run on the same ports within namespaces, and "duplicate" virtual interfaces can be created.
- CLONE_NEWNS: Mount Namespaces. We have the ability to isolate mount points as they appear to processes. Using mount namespaces, we can achieve similar functionality to chroot() however with improved security.

…Namespaces

SINE NOMINE ASSOCIATES

- CLONE_NEWUTS: UTS Namespaces. This namespaces primary purpose is to isolate the hostname and NIS name.
- CLONE_NEWUSER: User Namespaces. Here, user and group IDs are different inside and outside of namespaces and can be duplicated.

Copy-on-Write

SINE NOMINE ASSOCIATES

- Allows Docker to instantiate containers very quickly
- Instead of having to make full copies of the which files comprise a container, it can use “pointers” back to existing files
- Containers are easily “linked” (or “stacked” or “layered”) to other containers

Docker Registry (optional)

SINE NOMINE ASSOCIATES

- A stateless, highly scalable server-side application that stores and distributes Docker images
- Enables:
  - Tight control where images are stored
  - Full ownership of distribution pipeline
  - Integration of image storage & distribution into an in-house development workflow

Docker Daemon

SINE NOMINE ASSOCIATES

- Manages containers
  - Creates volumes
  - Starts/stops containers

```
docker daemon -H tcp://0.0.0.0:4243 -H
unix:///var/run/docker.sock
```

OR

```
systemctl enable docker
systemctl start docker
```

Docker

BUILDING CONTAINERS

Creating a Starter System

SINE NOMINE ASSOCIATES

- Base image: containers built from it or its descendants
- Create a chroot-like environment
  - File system including /dev
  - yum install packages
  - Trim unwanted stuff
  - Create tar ball
  - Import to Docker

The Dockerfile

SINE NOMINE ASSOCIATES

- A recipe for building a container
- Build from an existing container
- Install requirements
- Define network and volume requirements
- Specify command to run on startup


```
FROM    clefos-base-s390x:latest

MAINTAINER The ClefOS Project <neale@sinenomine.net>
LABEL   Vendor="ClefOS" License="GPLv2" Version="8.0-10.1"

COPY    ibm-java-sdk-8.0-1.10-s390x-archive.bin java.rsp dummy-
java-1.8-0.el7.noarch.rpm /

RUN     yum install -y tar zip && \
        mkdir -p /opt/ibm && \
        echo "Installing IBM JDK" && \
        /ibm-java-sdk-8.0-1.10-s390x-archive.bin -f /java.rsp -i silent && \
        yum install -y dummy-java-1.8-0.el7.noarch.rpm && \
        yum erase -y tar zip vim-minimal && \
        yum clean all && \
        rm -f /*.rpm /java.rsp /*.bin

ENV     JAVA_HOME=/opt/ibm/java PATH=$JAVA_HOME/bin:$PATH
```


```
FROM    sinenomine/clefos-nodejs-s390x
MAINTAINER      The ClefOS project <neale@sinenomine.net>
ADD     epel.repo /etc/yum.repos.d/epel.repo
RUN     yum install -y git tar gcc gcc-c++ make mongodb mongodb-server \
                mongo-tools krb5-devel perl-Digest-SHA && \
                npm install -g express && \
                npm install -g mongodb && \
                npm install -g tar mkdirp
WORKDIR /mean
EXPOSE  27017 28017
VOLUME  /mongodb/data
RUN     echo "mongod --fork --logpath /mongodb/data/log/mongod.log \
        --dbpath /mongodb/data --smallfiles --noprealloc --httpinterface --rest \
        > /start.sh && echo "node \$1" >> /start.sh && \
        yum erase -y git tar gcc gcc-c++ make perl-Digest-SHA && \
        rm -f /etc/yum.repos.d/epel.repo && \
        rm -rf /tmp/* /var/cache/yum/* /root/* /root/.[a-zA-Z0-9]* /src
ENV     NODE_PATH=/opt/ibm/nodejs/lib/node_modules:/mean/node_modules
ENTRYPOINT ["sh", "/start.sh"]
```

Building Images

SINE NOMINE ASSOCIATES

- Each step corresponds to a layer
- Stop build at one point
- Rebuild starts from last change

Managing Images

SINE NOMINE ASSOCIATES

```
[root@docker docker]# docker images
REPOSITORY                          TAG     IMAGE ID     CREATED       VIRTUAL SIZE
sinenomine/fluentd-s390x            latest  b3e3d646f313 4 days ago    515.2 MB
sinenomine/amhub-s390x              latest  76a2c4a387f0 7 days ago    795 MB
sinenomine/ade-s390x                latest  5dc6c7c6191c 5 weeks ago   645.8 MB
sinenomine/compose-ui-s390x         latest  ff5b9eda68ec 8 weeks ago   315.9 MB
sinenomine/nginx-1.8-s390x          latest  4f87e1292531 8 weeks ago   211 MB
sinenomine/clefos71-base-s390x      latest  60ef3a8ba174 3 months ago  110.5 MB
clefos-base-s390x                   latest  60ef3a8ba174 3 months ago  110.5 MB
sinenomine/clefos71-nodejs-s390x    latest  d76f12128dde 5 months ago  548.7 MB
sinenomine/mariadb-5.5-s390x        latest  91233ea5a5c1 5 months ago  311.3 MB
sinenomine/clefos71-java-s390x      latest  3cb8ef8fd562 5 months ago  480.2 MB
```

Making Images Available

SINE NOMINE ASSOCIATES

```
[root@docker ~]# docker push sinenomine/fluentd-s390x:latest
The push refers to a repository [docker.io/sinenomine/fluentd-s390x] (len: 1)
b3e3d646f313: Pushed
1b11901fbead: Pushed
5f6ab7c78e8b: Pushed
288d092713a6: Pushed
f86e5eb99f4b: Pushed
:
d69fc3fad8fa: Pushed
732e18ef67b6: Pushed
7196f6de1451: Pushed
7118afa06d84: Pushed
ec3ec425b681: Pushed
60ef3a8ba174: Pushed
latest: digest:
sha256:120519d3d8f0cf00a0caddb3fd8c0c6148b8145dbf6fed2897b36e965d35424d size: 29665
```

Dockerhub

Repositories

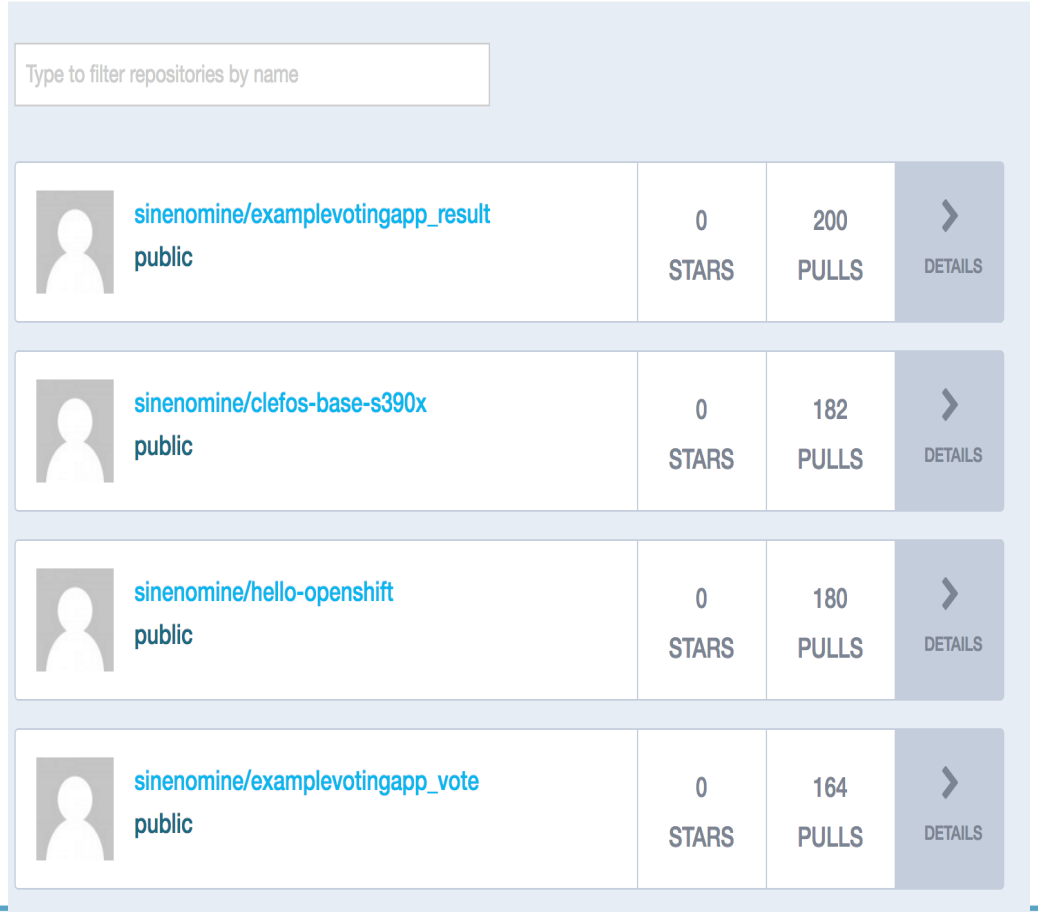

Running Containers

- Persistent data goes to [a] volume[s]
- Run a standalone container
  - All functionality within the container
- Run a “swarm” of containers
  - Typically database server
  - Web server
  - Application server

Running Containers

SINE NOMINE ASSOCIATES

- `docker run --name=mariadb -v /var/local/mariadb:/var/lib/mysql -d -p 3306:3306 -e MYSQL_ROOT_PASSWORD=passw0rd sinenomine/mariadb-5.5-s390x:latest mysqld_safe --connect-timeout=30`

- `docker run --rm -i -t --name=ade -p 8022:22 --link=mariadb -v /var/local/ade:/var/local/ade -e MARIADB_ROOT_PASSWORD=passw0rd -e MARIADB_ADE_PASSWORD=passw0rd sinenomine/ade-s390x`

Running Containers

SINE NOMINE ASSOCIATES

- Containers run as daemons or interactively
- Multiple containers wanting to use same port?
  - Docker can remap:
    -p <host port>:<container port>

Running Containers

SINE NOMINE ASSOCIATES

- What about environment variables?
  - -e option
  - Dockerfile
- What is my container doing:
  - docker top <*image id*>
  - top
- What is my container config?
- docker inspect <*image*>


```
 "Networks": {
        "bridge": {
            "EndpointID":
"0448a9a68ed5a9c5f89435b3d62d78bbc42d4f601a25e65b2e1
49ed8f694993c",
            "Gateway": "172.18.0.1",
            "IPAddress": "172.18.0.3",
            "IPPrefixLen": 16,
            "IPv6Gateway": "",
            "GlobalIPv6Address": "",
            "GlobalIPv6PrefixLen": 0,
            "MacAddress": "02:42:ac:12:00:03"
        }
```

Running Containers

SINE NOMINE ASSOCIATES

- Command line
  - docker run
  - kubernetes
- GUIs
  - Compose-UI
  - AMHub
  - **OpenShift**
- Images are automatically downloaded

Running Containers

SINE NOMINE ASSOCIATES

- Build on ClefOS / Run on Ubuntu
- Build on ClefOS / Build upon image on Ubuntu
- Builders meet all pre-requisites
- Self-contain requirements
  - No conflicts with other containers
  - Unlike multi-tenancy apps

Docker on z

DEMO TIME

Docker Compose UI

Projects

- hello-node
- nginx
- node-redis
- volumes-demo

filter by project name:

Reload projects

Create new project

Settings

nginx

view yml

Actions

build | pull | logs | kill | stop | start | up | down

Services

no containers found

SINE NOMINE
ASSOCIATES

nginx

close yml

```
nginx:
  image: brunswickheads/nginx-1.8-s390x
  container_name: nginx
  hostname: nginx
  ports:
  - 80:80
  - 443:443
```

nginx

close yml

```
nginx:
  image: brunswickheads/nginx-1.8-s390x
  container_name: nginx
  hostname: nginx
  ports:
  - 80:80
  - 443:443
```

Services

nginx scale

nginx logs details

Network

0.0.0.0:443 : 443/tcp

0.0.0.0:80 : 80/tcp

nginx / nginx

Key	Value
name	nginx
ip	172.18.0.4
number	1
short id	5a9fb7a01a
log config	json-file
command	nginx -g daemon off;
environment	• PATH = /usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin:/sbin:/bin
labels	• com.docker.compose.config-hash = fe4fe73c044ff3dd65431737d3340ed20a0ad8ab1c7288b30e69fbb416da19fe • com.docker.compose.container-number = 1 • com.docker.compose.oneoff = False • com.docker.compose.project = nginx • com.docker.compose.service = nginx • com.docker.compose.version = 1.6.0

Welcome to **nginx** on ClefOS!

This page is used to test the proper operation of the **nginx** HTTP server after it has been installed. If you can read this page, it means that the web server installed at this site is working properly.

Website Administrator

This is the default `index.html` page that is distributed with **nginx** on ClefOS. It is located in `/usr/share/nginx/html`.

You should now put your content in a location of your choice and edit the `root` configuration directive in the **nginx** configuration file `/etc/nginx/nginx.conf`.

nginx

close yml

```
nginx:
  image: brunswickheads/nginx-1.8-s390x
  container_name: nginx
  hostname: nginx
  ports:
  - 80:80
  - 443:443
```

SINE NOMINE ASSOCIATES

nginx

close yml

```
nginx:
  image: brunswickheads/nginx-1.8-s390x
  container_name: nginx
  hostname: nginx
  ports:
  - 80:80
  - 443:443
```

Actions

build | pull | logs | kill | stop | start | up | down

Services

no containers found

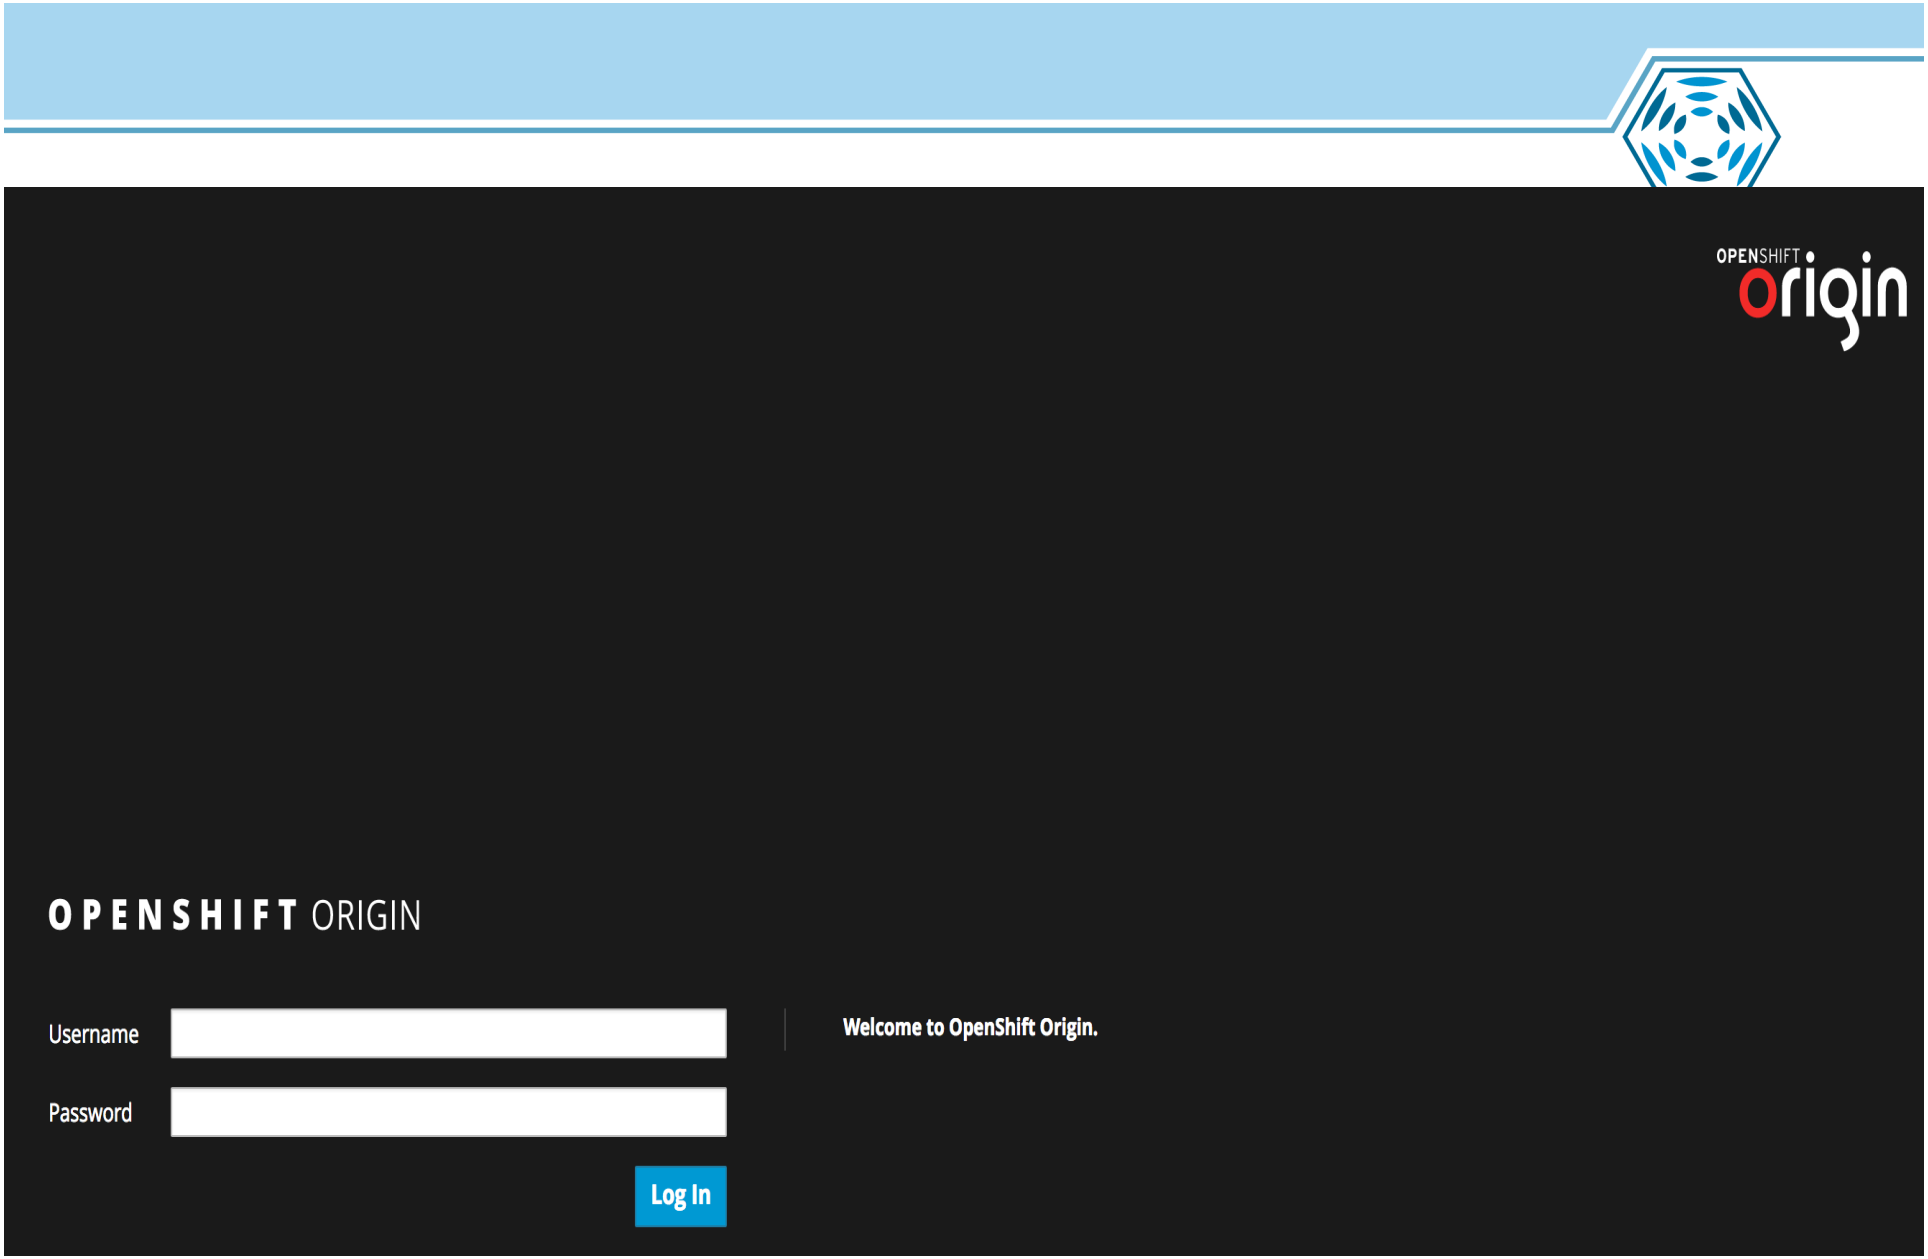
OPENSHIFT origin
OPENSHIFT ORIGIN
Username
Password
Log In
Welcome to OpenShift Origin.

SINE NOMINE ASSOCIATES

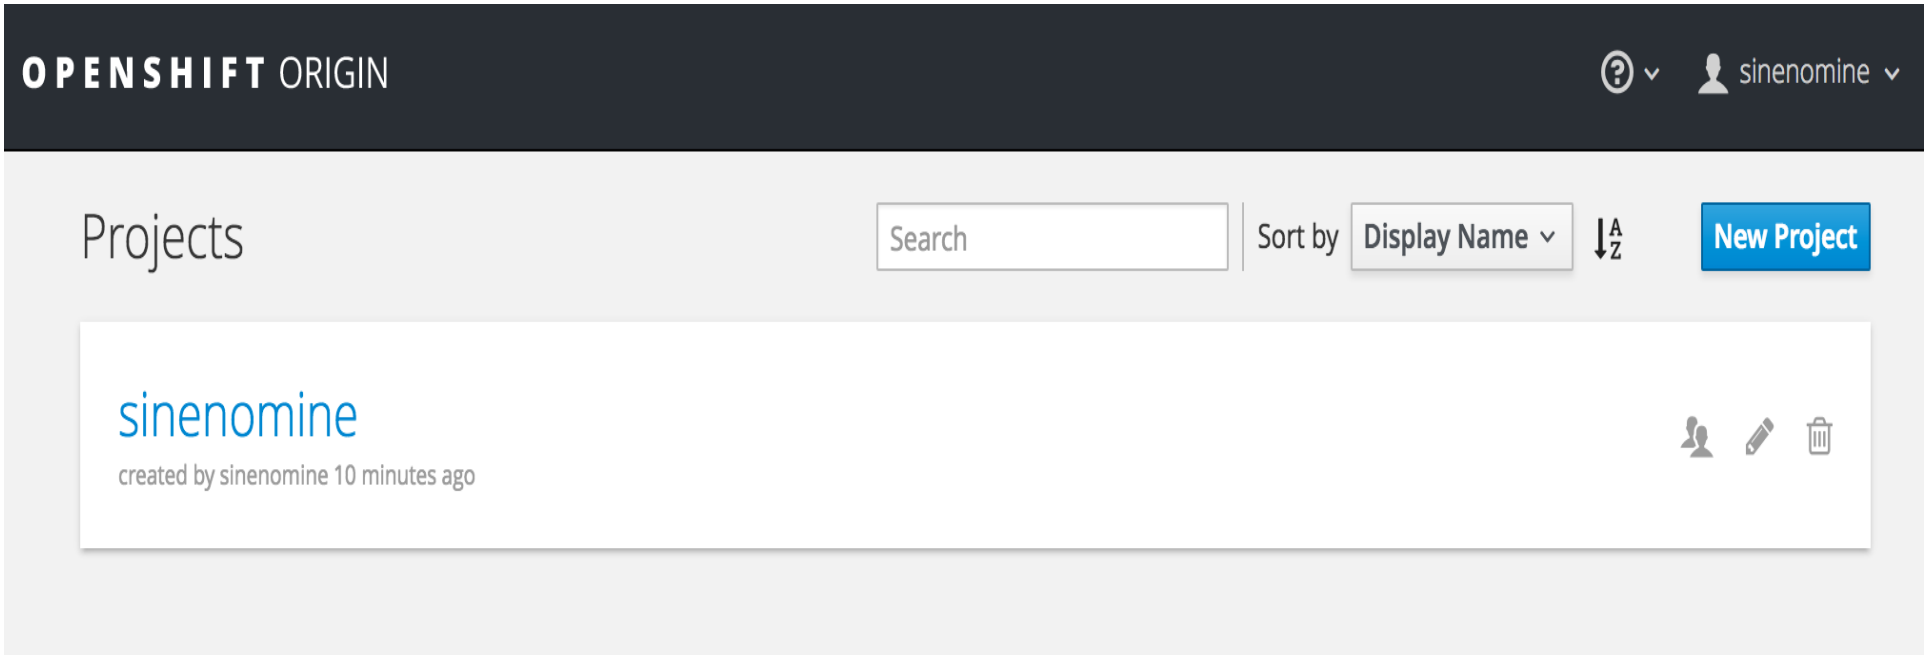

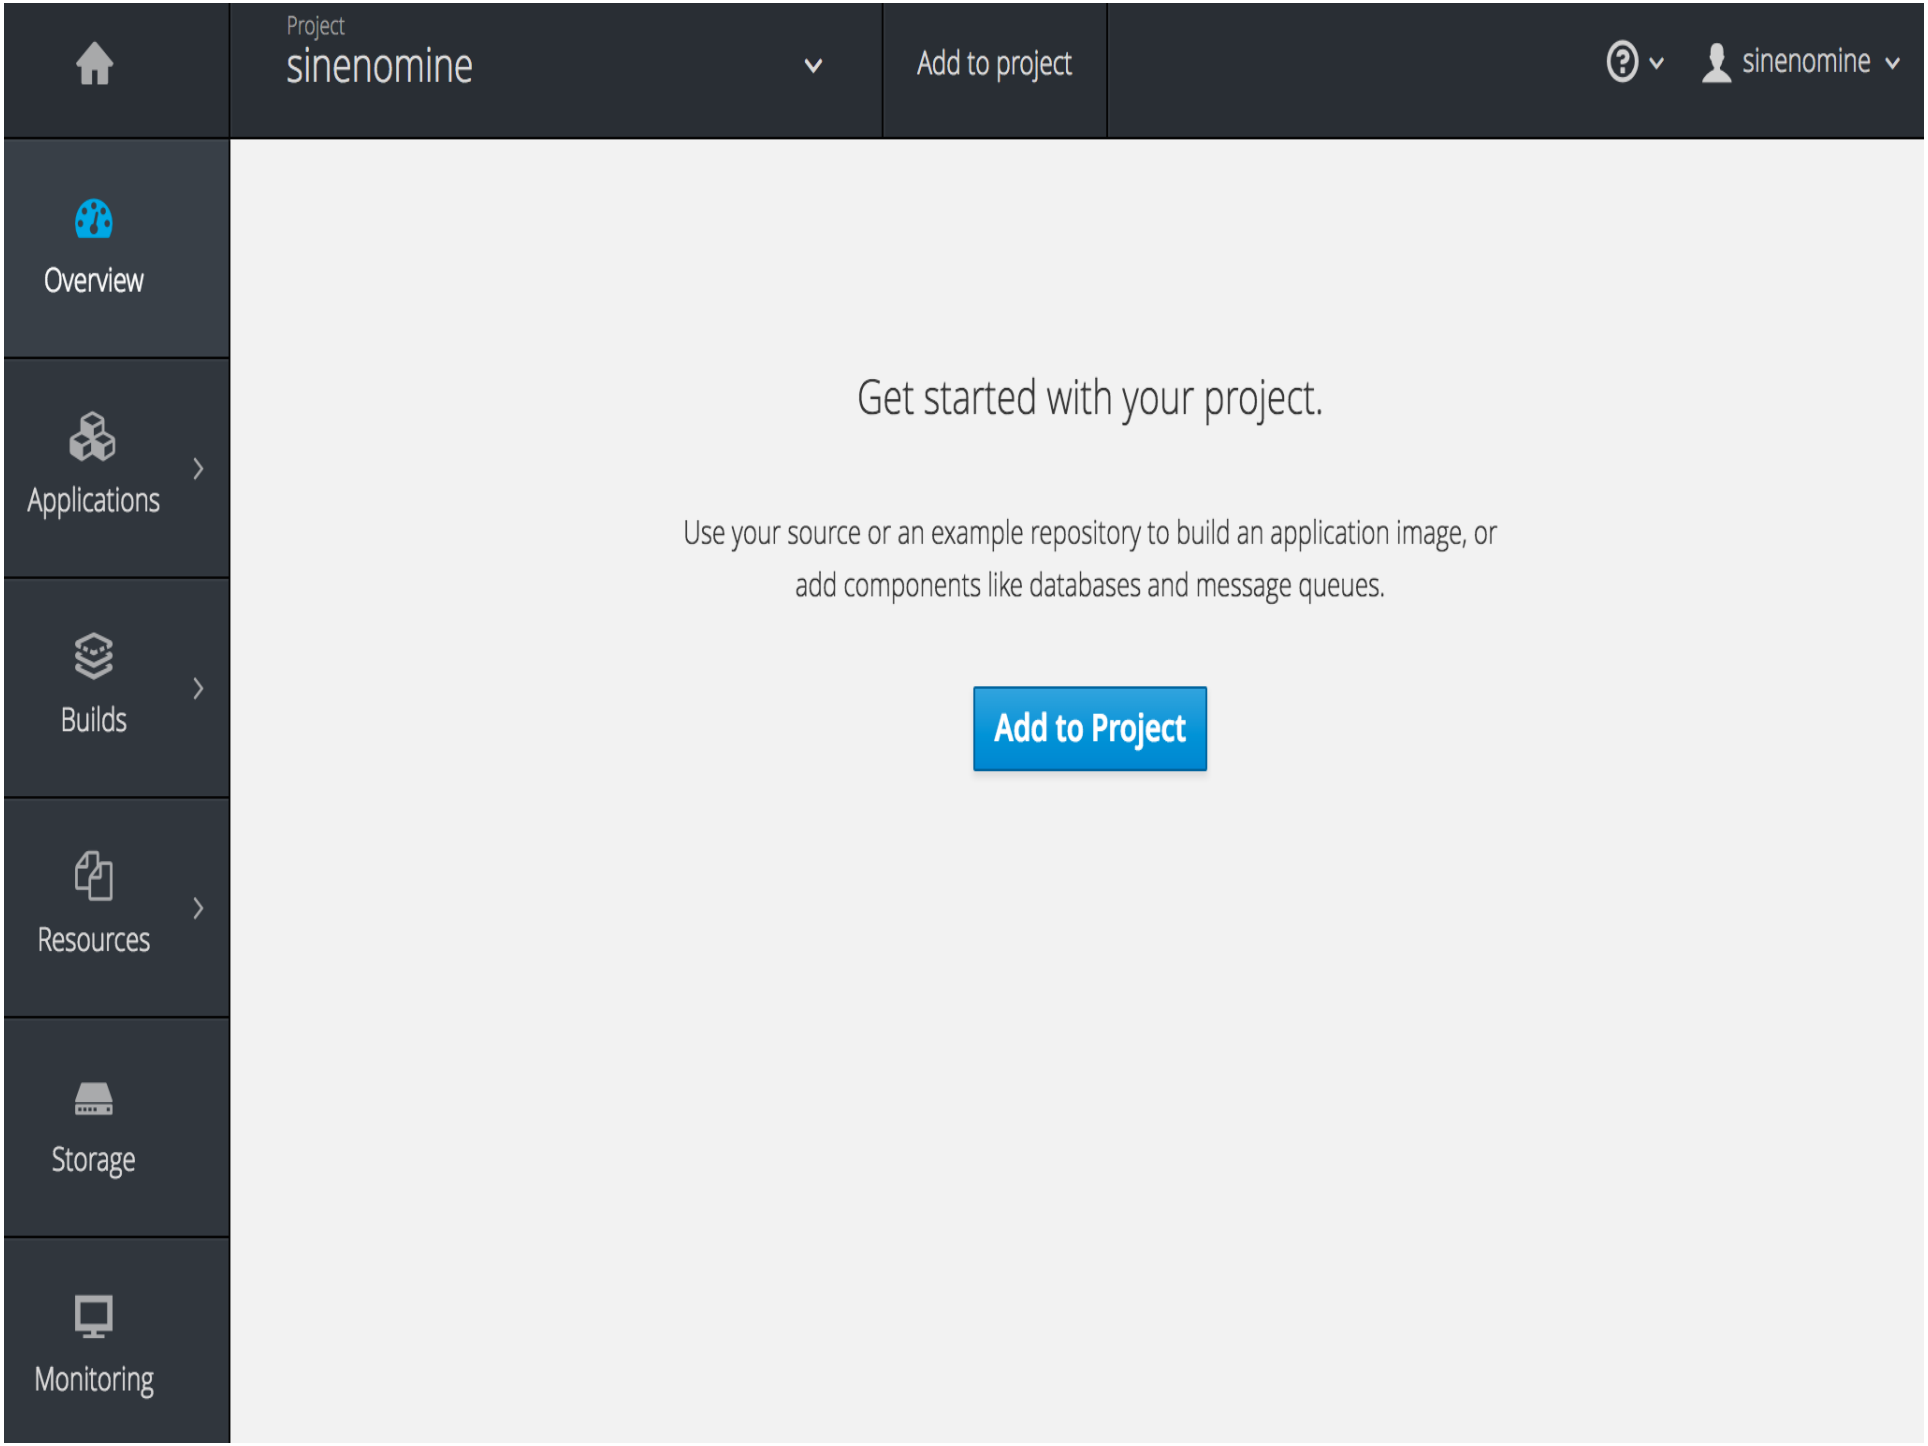
Project
sinenomine
Add to project
sinenomine
Overview
Applications
Builds
Resources
Storage
Monitoring
Get started with your project.
Use your source or an example repository to build an application image, or
add components like databases and message queues.
Add to Project

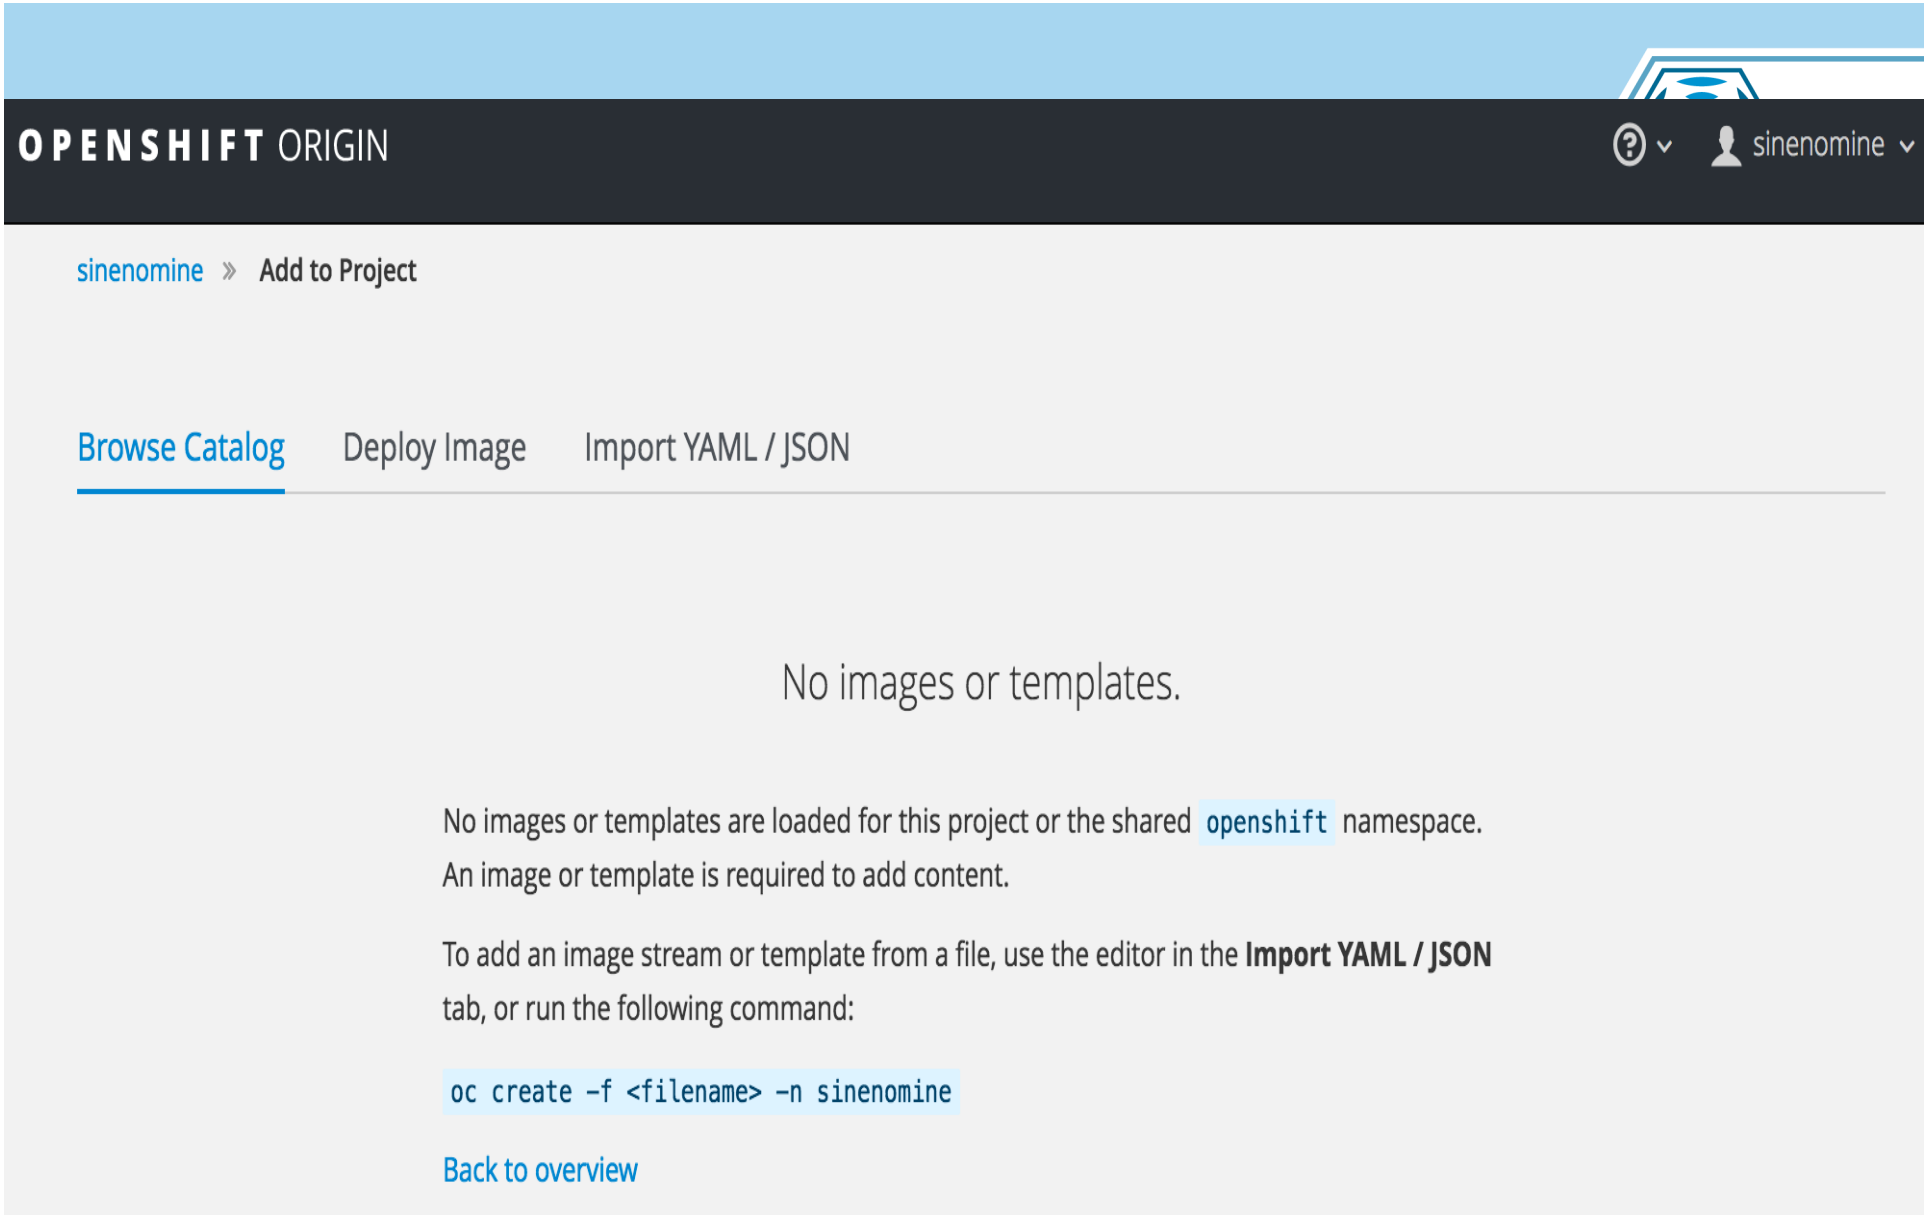
OPENSHIFT ORIGIN
sinenomine
sinenomine » Add to Project
Browse Catalog
Deploy Image
Import YAML / JSON
No images or templates.
No images or templates are loaded for this project or the shared openshift namespace.
An image or template is required to add content.
To add an image stream or template from a file, use the editor in the Import YAML / JSON tab, or run the following command:
oc create -f <filename> -n sinenomine
Back to overview

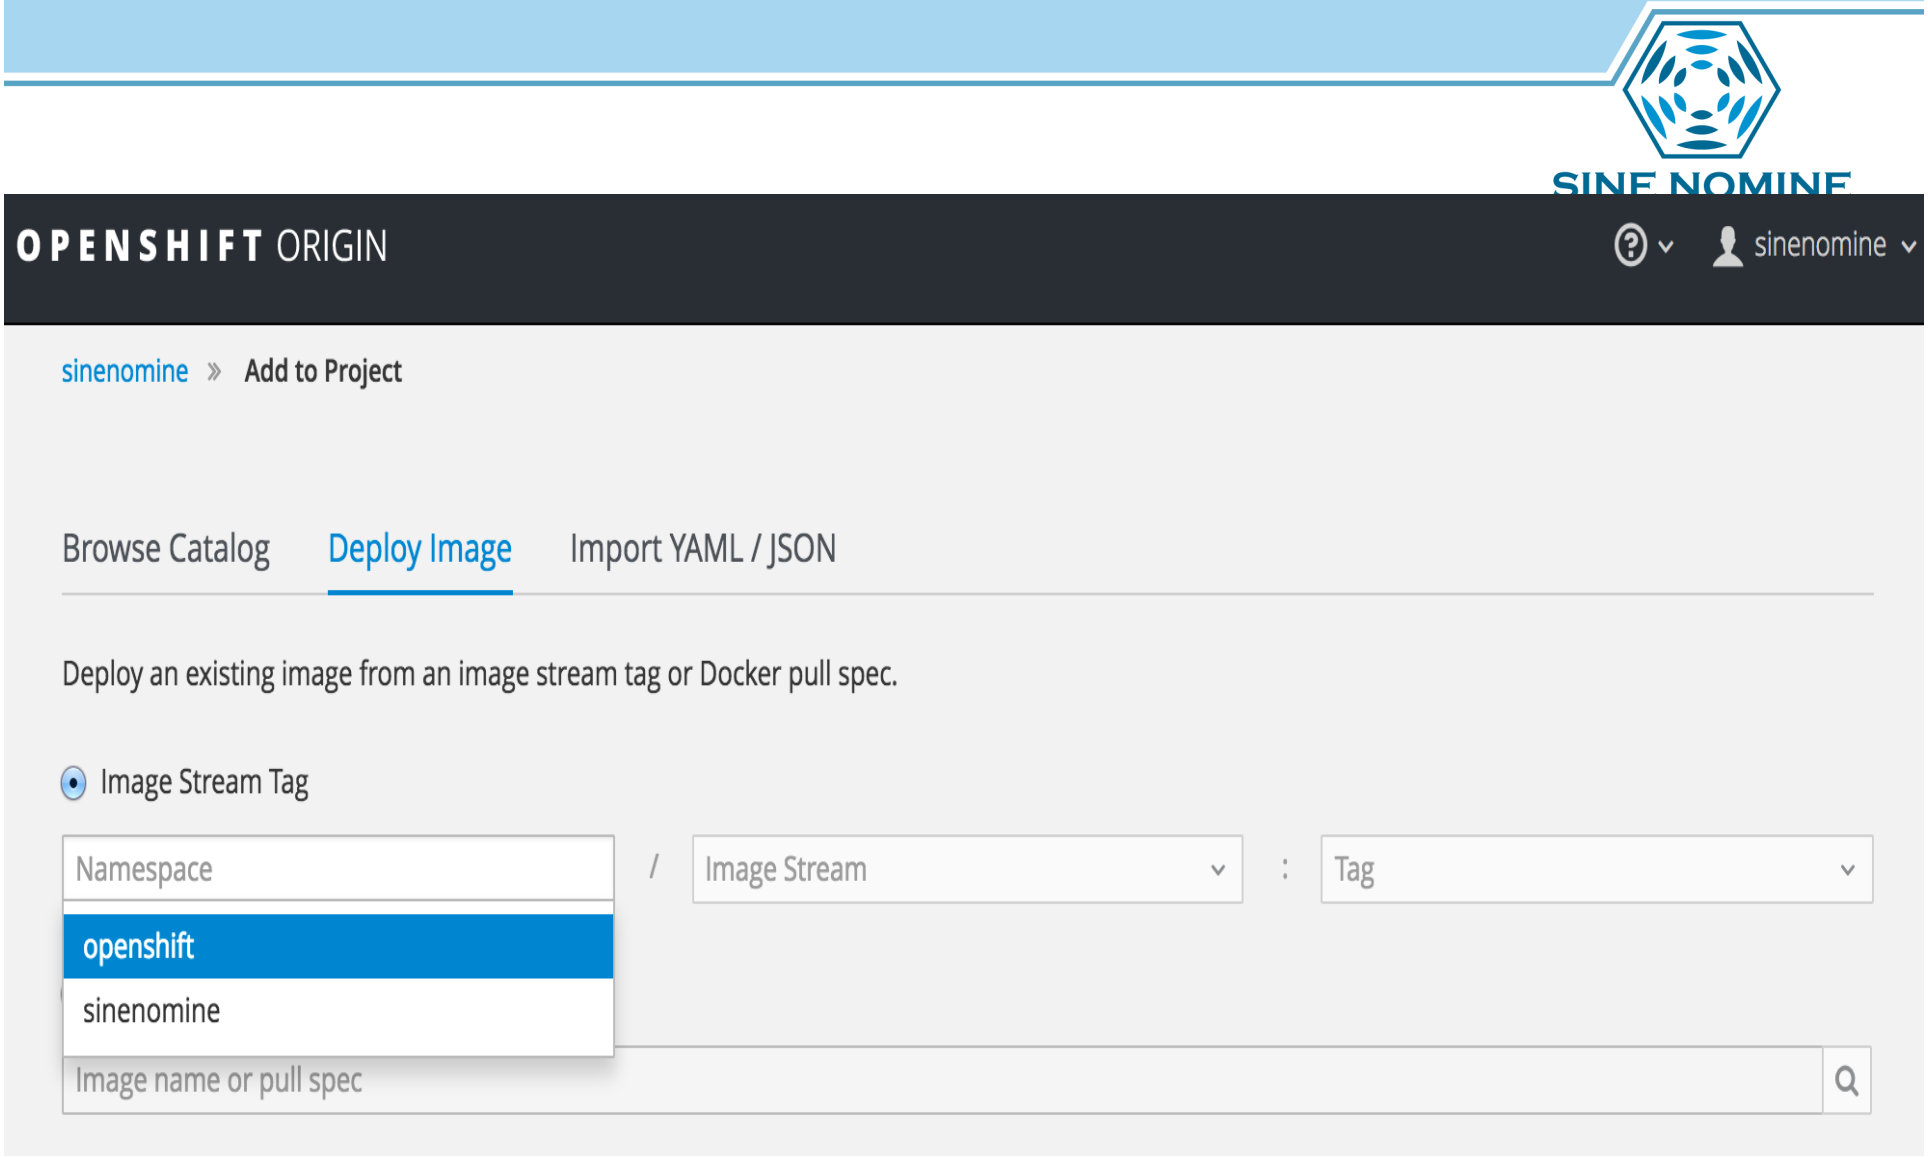
SINE NOMINE
OPENSHIFT ORIGIN
sinenomine
sinenomine » Add to Project
Browse Catalog
Deploy Image
Import YAML / JSON
Deploy an existing image from an image stream tag or Docker pull spec.
Image Stream Tag
Namespace
/
Image Stream
:
Tag
openshift
sinenomine
Image name or pull spec

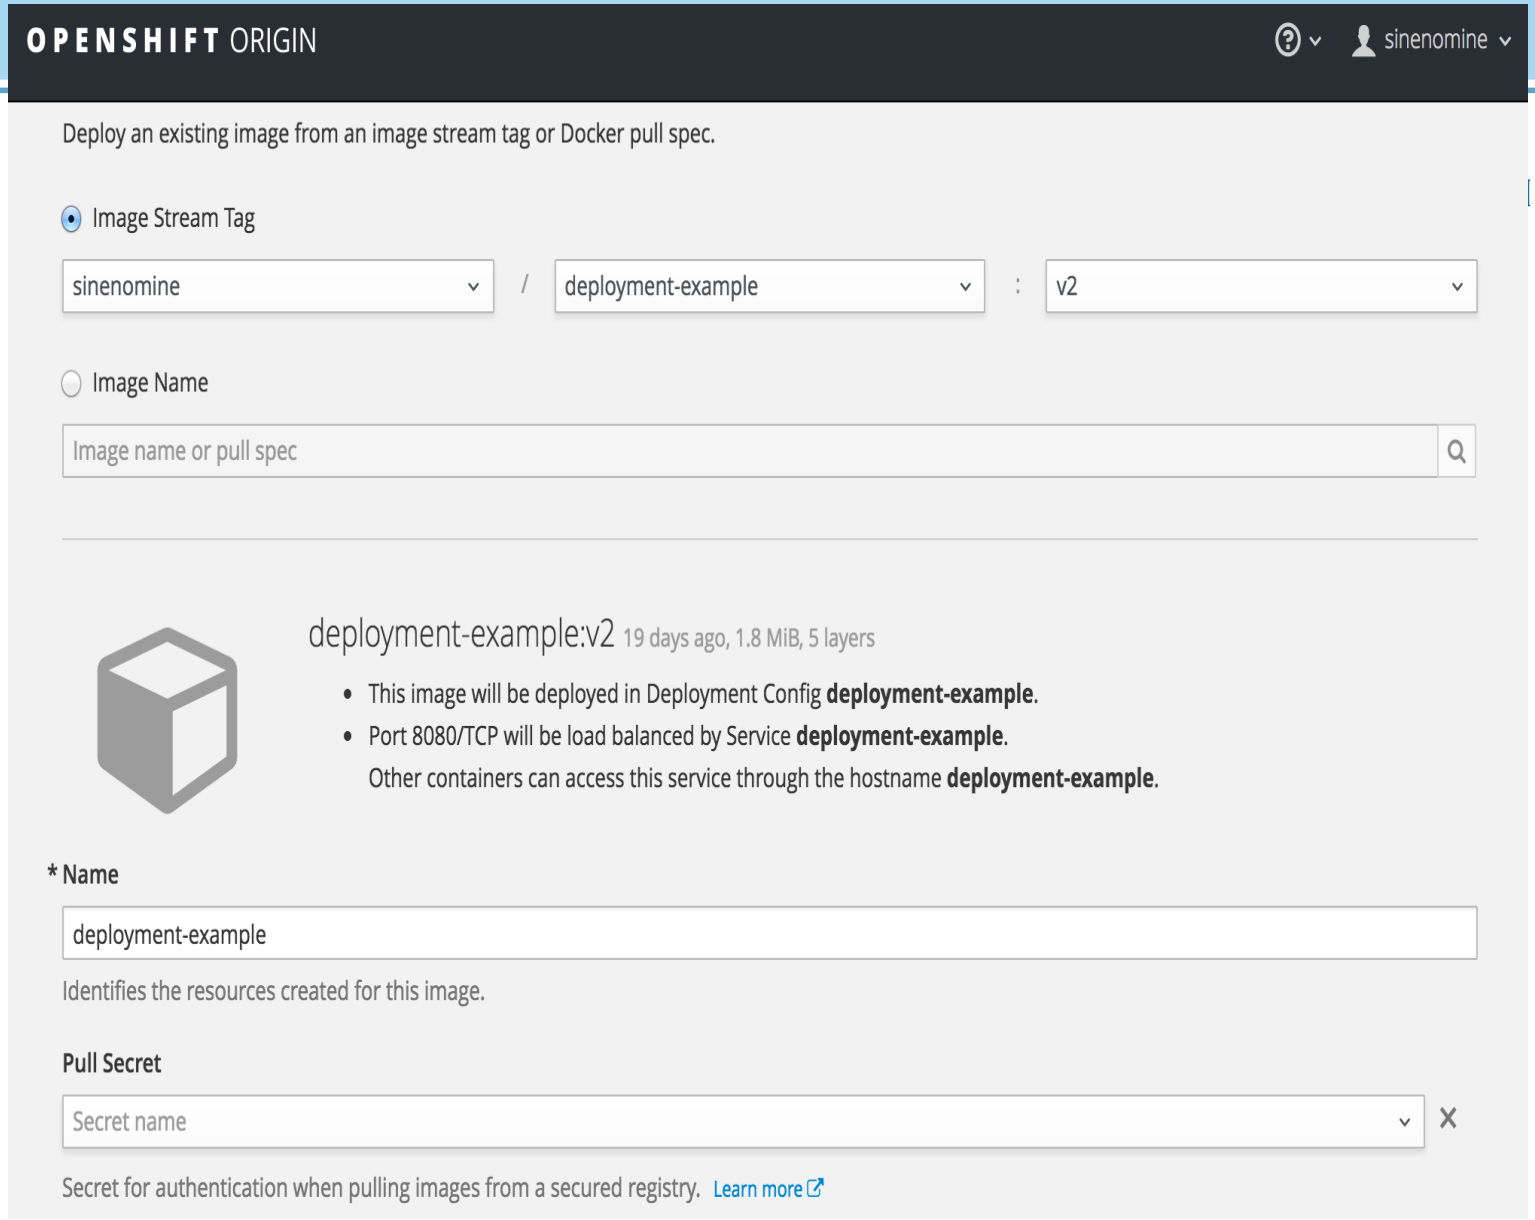

NE NOMINE ASSOCIATES

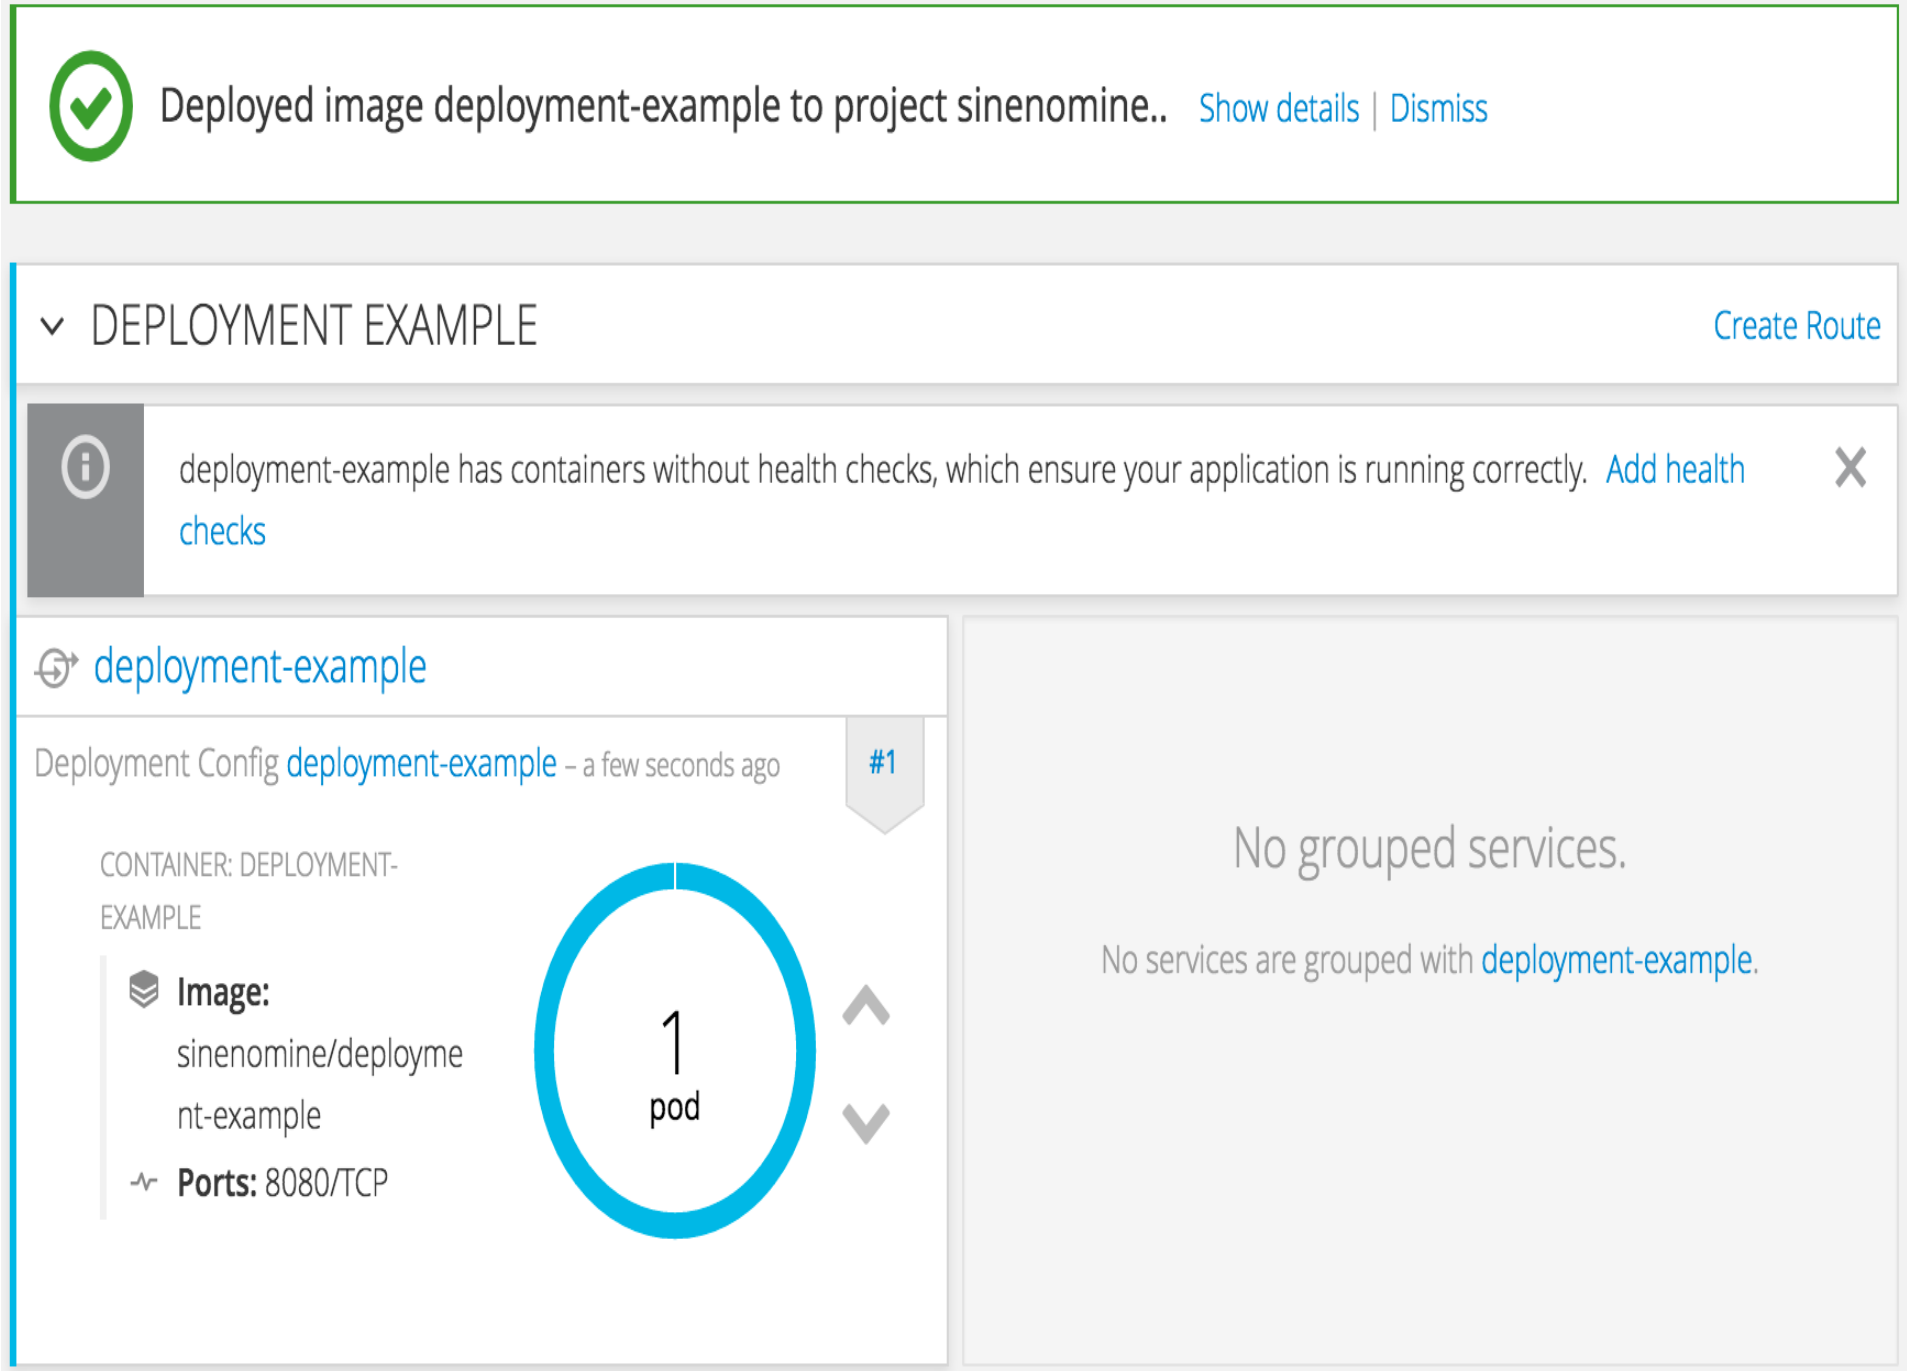
Deployed image deployment-example to project sinenomine.. Show details | Dismiss
DEPLOYMENT EXAMPLE
Create Route
deployment-example has containers without health checks, which ensure your application is running correctly. Add health checks
deployment-example
Deployment Config deployment-example – a few seconds ago
#1
CONTAINER: DEPLOYMENT-EXAMPLE
Image:
sinenomine/deployment-example
Ports: 8080/TCP
1
pod
No grouped services.
No services are grouped with deployment-example.

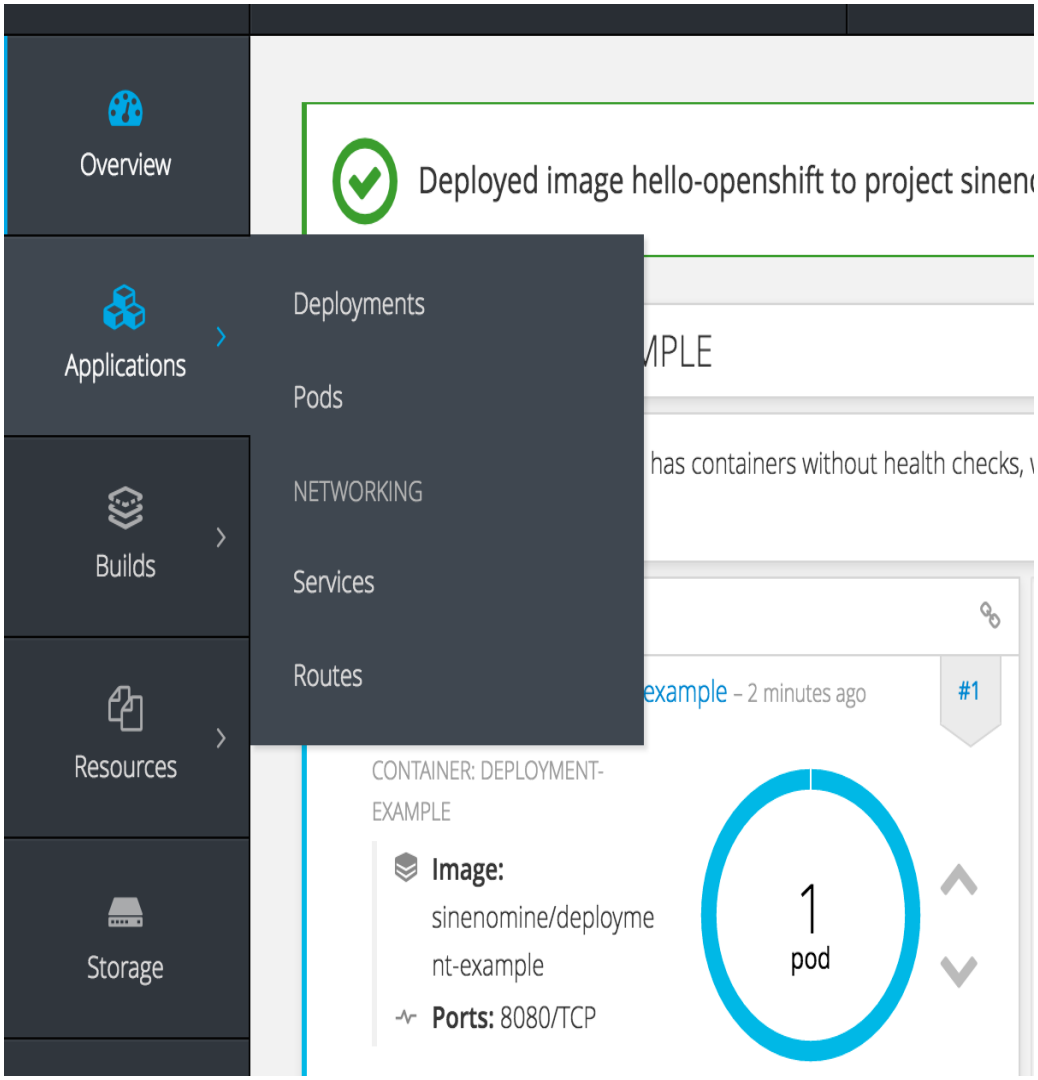
Overview
Applications
Builds
Resources
Storage
Deployments
Pods
NETWORKING
Services
Routes
Deployed image hello-openshift to project sinen
MPLE
has containers without health checks,
example – 2 minutes ago
#1
CONTAINER: DEPLOYMENT-EXAMPLE
Image:
sinenomine/deployment-example
Ports: 8080/TCP
1
pod

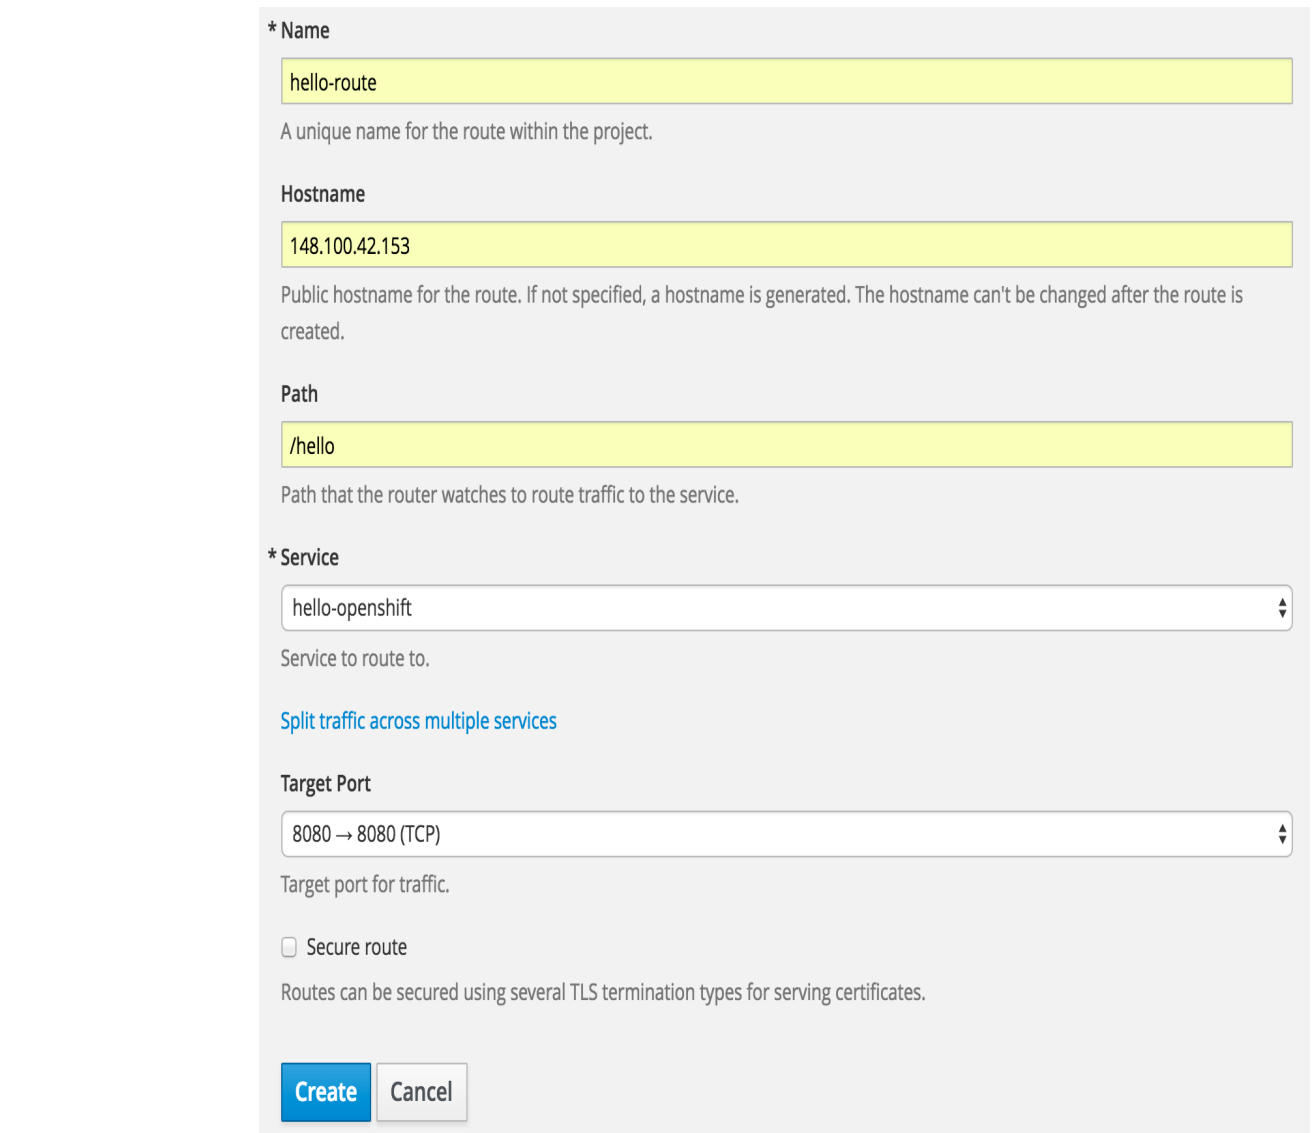

SINE NOMINE ASSOCIATES

Routes

Create Route

Filter by label | Add

Name	Hostname	Routes To	Target Port	TLS Termination
deployment	http://148.100.42.153/deployment	deployment-example	8080-tcp	
hello-route	http://148.100.42.153/hello	hello-openshift	8080-tcp	

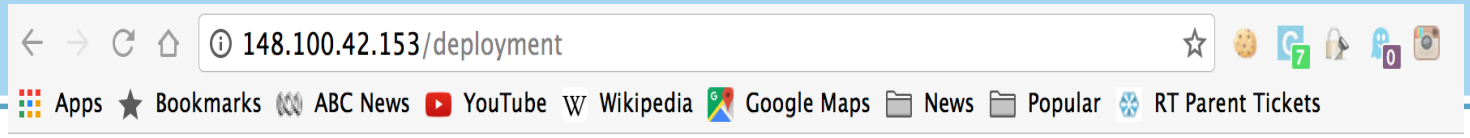

SINE NOMINE ASSOCIATES

v2

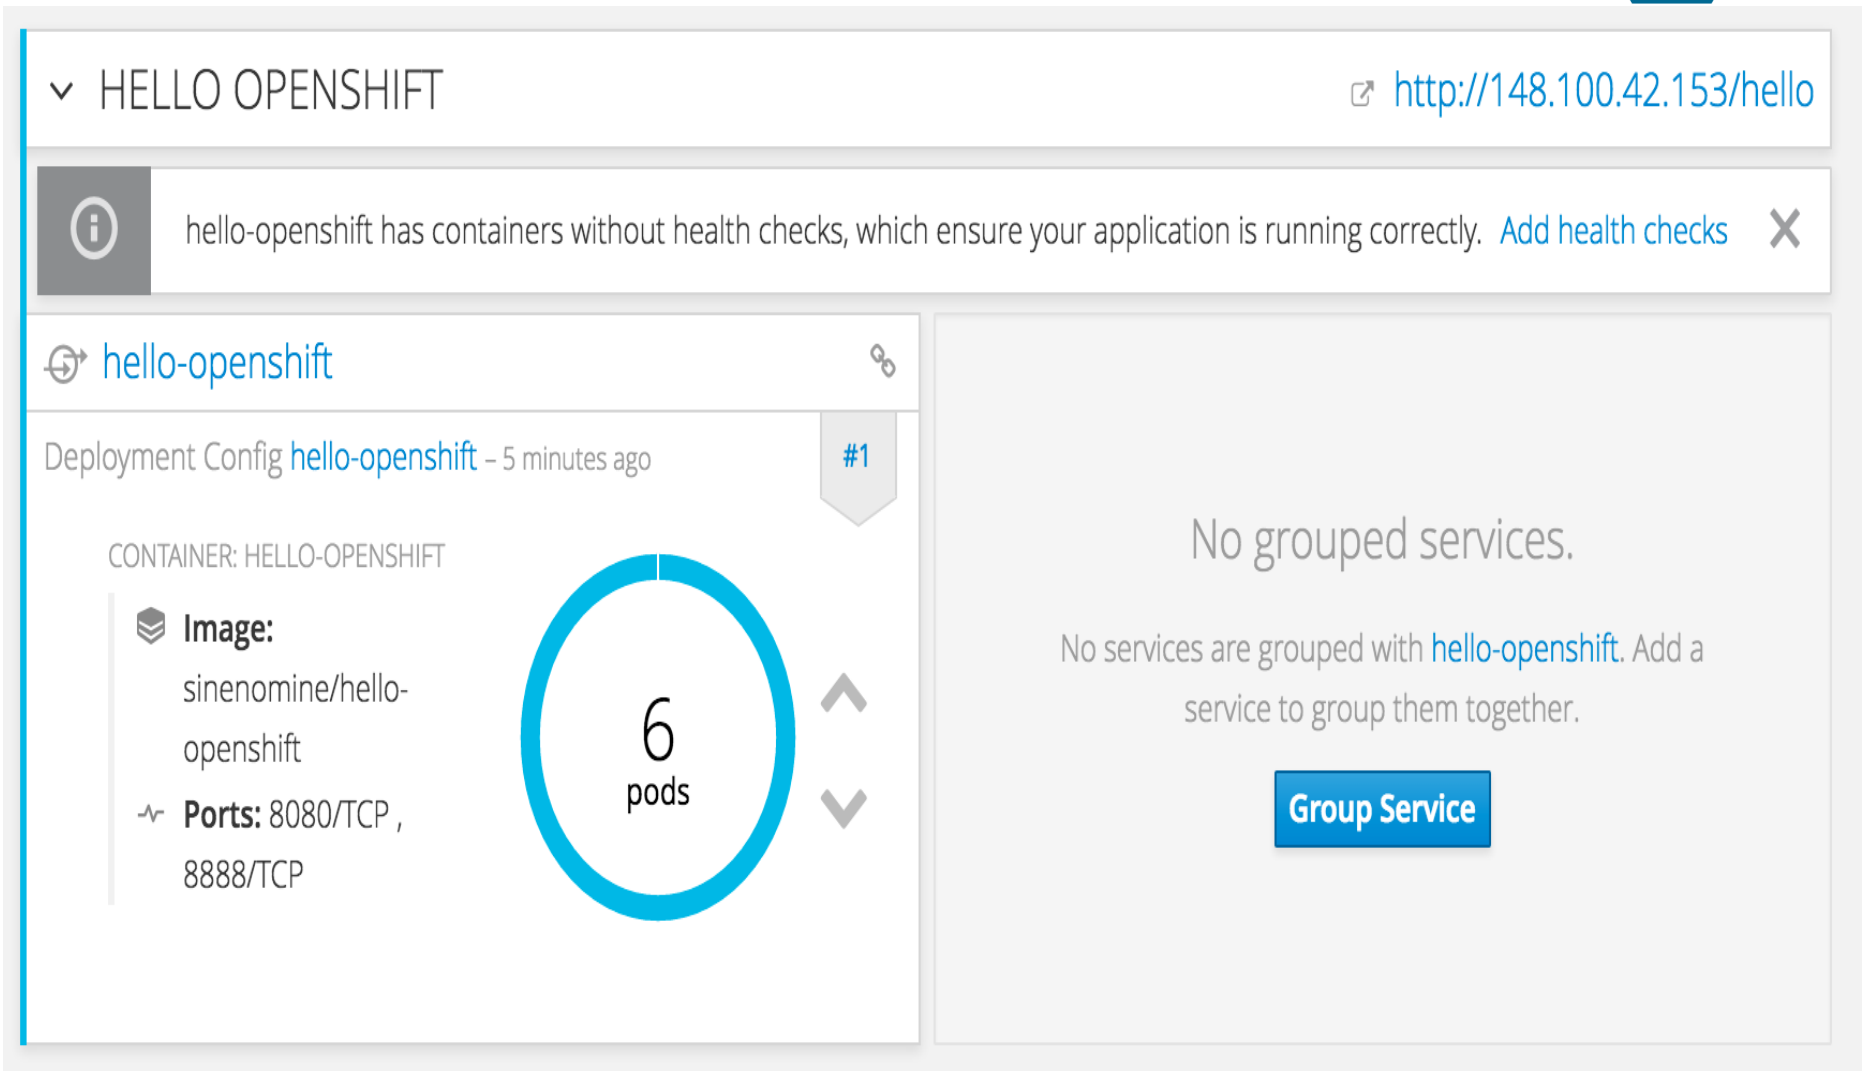
HELLO OPENSHIFT
http://148.100.42.153/hello
hello-openshift has containers without health checks, which ensure your application is running correctly. Add health checks
hello-openshift
Deployment Config hello-openshift – 5 minutes ago
#1
CONTAINER: HELLO-OPENSHIFT
Image:
sinenomine/hello-openshift
Ports: 8080/TCP , 8888/TCP
6
pods
No grouped services.
No services are grouped with hello-openshift. Add a service to group them together.
Group Service


```
# docker exec -it origin bash
bash-4.2# oc login
Authentication required for https://148.100.42.153:8443 (openshift)
Username: sinenomine
Password:
Login successful.

You have one project on this server: "sinenomine"

Using project "sinenomine".
bash-4.2# oc status
In project sinenomine on server https://148.100.42.153:8443

http://148.100.42.153 to pod port 8080-tcp (svc/deployment-example)
  dc/deployment-example deploys istag/deployment-example:v2
    deployment #1 deployed 8 minutes ago - 1 pod

http://148.100.42.153 to pod port 8080-tcp (svc/hello-openshift)
  dc/hello-openshift deploys istag/hello-openshift:latest
    deployment #1 deployed 6 minutes ago - 6 pods
```


```
# oc get pods
NAME                          READY     STATUS    RESTARTS   AGE
deployment-example-1-ypclg    1/1       Running   0          32m
hello-openshift-1-b61vl       1/1       Running   0          25m
hello-openshift-1-jqkvb       1/1       Running   0          25m
hello-openshift-1-tu25u       1/1       Running   0          25m
hello-openshift-1-v7ra3       1/1       Running   0          30m
hello-openshift-1-vxss9       1/1       Running   0          25m
hello-openshift-1-yw0z7       1/1       Running   0          25m
```


```
# oc new-app docker.io/sinenomine/lighttpd-s390x:latest
--> Found Docker image 3d4758d (3 weeks old) from docker.io for "docker.io/sinenomine/lighttpd-
s390x:latest"

    * An image stream will be created as "lighttpd-s390x:latest" that will track this image
    * This image will be deployed in deployment config "lighttpd-s390x"
    * Port 8091/tcp will be load balanced by service "lighttpd-s390x"
      * Other containers can access this service through the hostname "lighttpd-s390x"
    * WARNING: Image "docker.io/sinenomine/lighttpd-s390x:latest" runs as the 'root' user which
      may not be permitted by your cluster administrator

--> Creating resources ...
    imagestream "lighttpd-s390x" created
    deploymentconfig "lighttpd-s390x" created
    service "lighttpd-s390x" created
--> Success

svc/lighttpd-s390x - 172.30.232.241:8091
  dc/lighttpd-s390x deploys istag/lighttpd-s390x:latest
    deployment #1 deployed 50 seconds ago - 1 pod
```