

# OpenRiskNet

RISK ASSESSMENT E-INFRASTRUCTURE

## Deploying Applications to an OpenRiskNet Virtual Environment

The OpenRiskNet Consortium

OpenRiskNet: Open e-Infrastructure to Support Data Sharing, Knowledge Integration and *in silico* Analysis and Modelling in Risk Assessment  
Project Number 731075


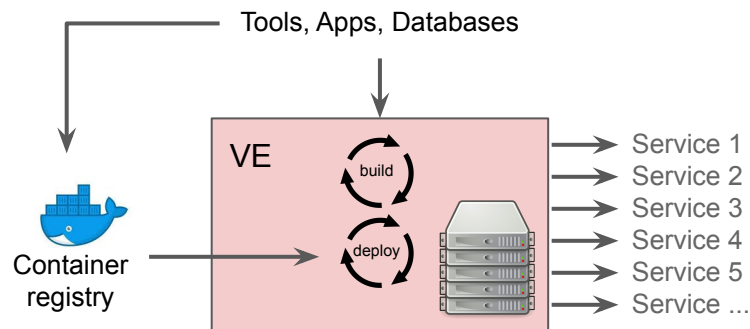




	Topic	Date & Time
<b>Past events</b>	Introduction sessions to the OpenRiskNet e-infrastructure	Webinar recordings: <ul style="list-style-type: none"> <li>• <a href="#">Session 1</a> (24 Sep 2018)</li> <li>• <a href="#">Session 2</a> (27 Sep 2018)</li> <li>• <a href="#">Session 3</a> (4 Oct 2018)</li> <li>• <a href="#">Session 4</a> (30 Oct 2018)</li> </ul>
	Learn how to deploy the OpenRiskNet virtual research environment	<a href="#">Webinar recordings</a> (25 Feb 2019)
	Demonstration on data curation and creation of pre-reasoned datasets in the OpenRiskNet framework	<a href="#">Webinar recordings</a> (18 Mar 2019)
	Identification and linking of data related to AOPWiki (an OpenRiskNet case study)	<a href="#">Webinar recordings</a> (26 March 2019)
	The Adverse Outcome Pathway Database (AOP-DB)	<a href="#">Webinar recordings</a> (8 April 2019)
	How to describe OpenRiskNet services and their functionality by semantic annotation	<a href="#">Webinar recordings</a> (13 May 2019)
	Use of Nextflow tool for toxicogenomics-based prediction and mechanism identification in OpenRiskNet e-infrastructure	<a href="#">Webinar recordings</a> (27 May 2019)
	Demonstration on OpenRiskNet approach on modelling for prediction or read across (ModelRX case study)	Tuesday, <b>11 June 2019</b> , 16:00 CEST Registration: <a href="https://openrisknet.org/events/67/">https://openrisknet.org/events/67/</a>
Combining neXtProt and WikiPathways strengths using SPARQL federated queries	Wednesday, <b>12 June 2019</b> , 20:00 CEST Registration: <a href="https://openrisknet.org/events/73/">https://openrisknet.org/events/73/</a>	
<b>Current event</b>	Deploying Applications to an OpenRiskNet Virtual Environment	Monday, <b>24 June 2019</b> , 16:00 CEST Registration: <a href="https://openrisknet.org/events/66/">https://openrisknet.org/events/66/</a>
<b>Future events</b>	AOPlink workflow	Monday, <b>15 July 2019</b> , 16:00 CEST Registration: <a href="https://openrisknet.org/events/70/">https://openrisknet.org/events/70/</a>

# OpenRiskNet Virtual Environment (VE)

- Computational infrastructure into which applications can be deployed
- Includes environment for building and testing those applications
- Includes compute, security, storage, monitoring ...
- Can be deployed to range of infrastructures

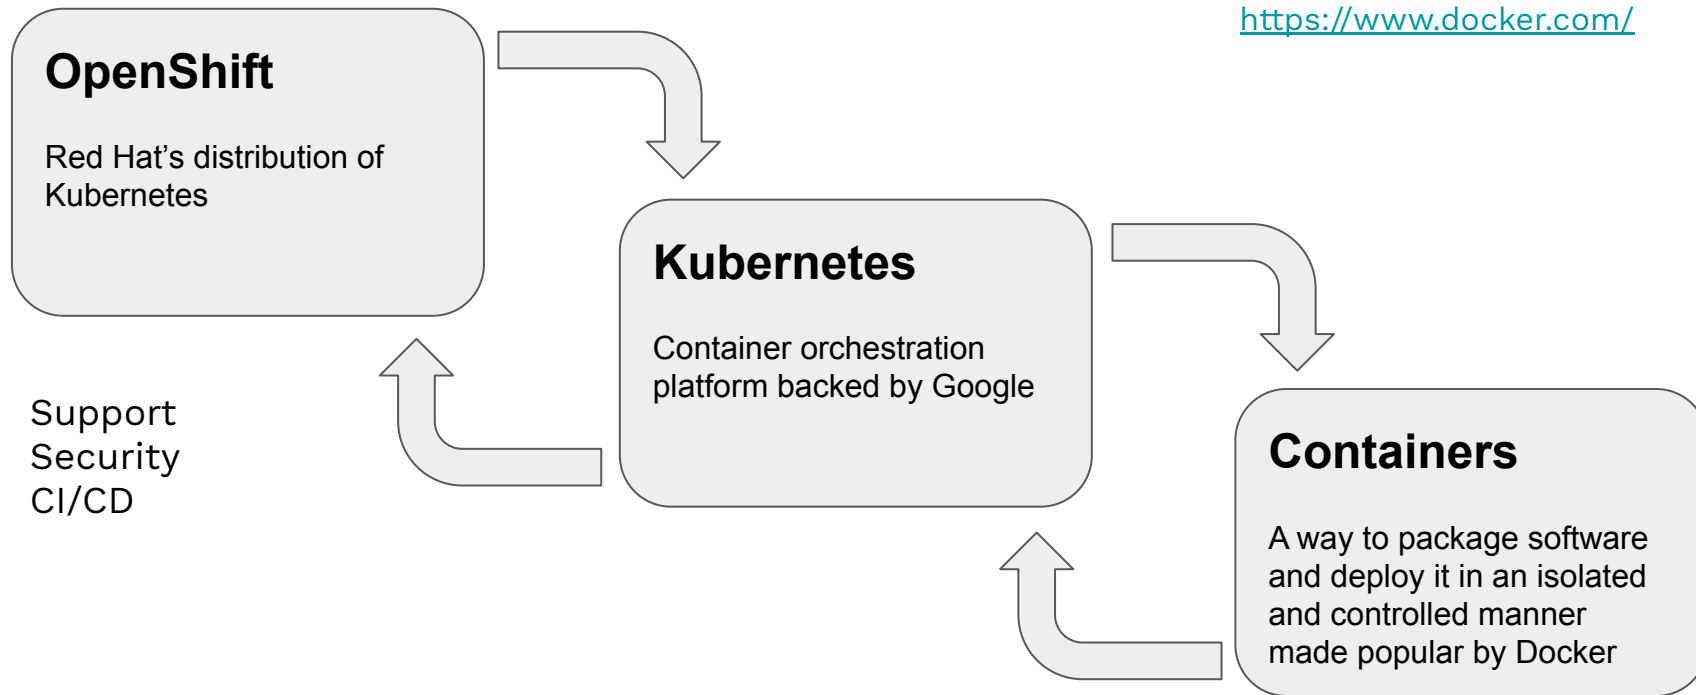


# What forms a VE

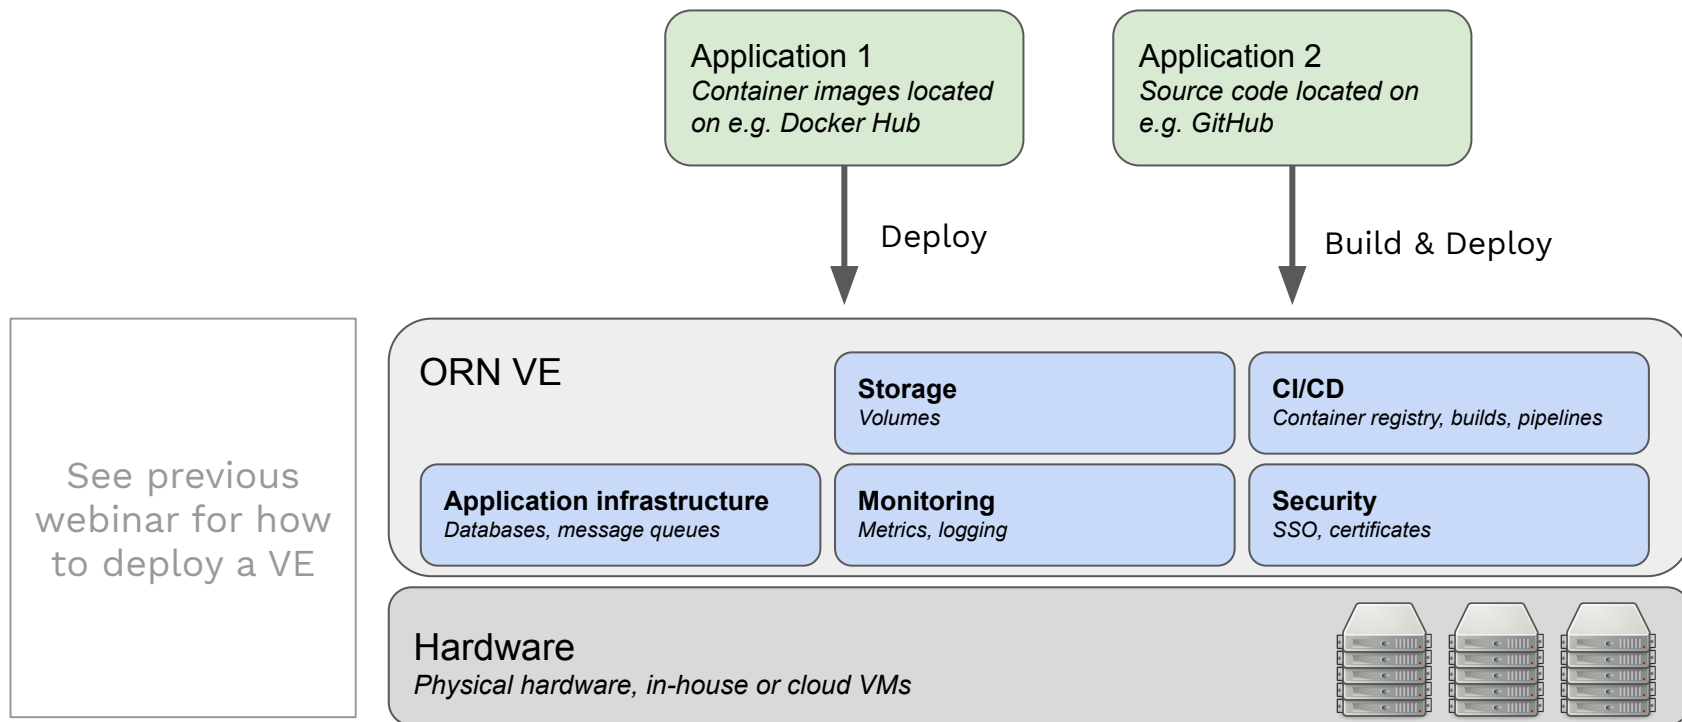
<https://www.openshift.com/>

<https://kubernetes.io/>

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



# App Deployment to a VE



# Introduction to containers

- A **container** is a set of Linux processes running in an isolated environment that is managed by features of the Linux kernel
- A bit like **virtual machines** but much more lightweight and efficient
- The software and data for those processes are packed into a container **image** that can be distributed



```
tmbo@xps:~$ ps -ef
root      15175   2  0 12:04 ?        00:00:00 [kworker/4:2-cgr]
tmbo     15249  6015  4 12:04 tty2    00:01:42 /opt/google/chrome/chrome --
tmbo     15362  6015  0 12:05 tty2    00:00:06 /opt/google/chrome/chrome --
root     15390   2  0 12:05 ?        00:00:00 [kworker/3:0-eve]
tmbo     15554  6015 15 12:08 tty2    00:05:41 /opt/google/chrome/chrome --
root     15885   2  0 12:19 ?        00:00:00 [kworker/7:0-eve]
root     16100   2  0 12:25 ?        00:00:00 [kworker/2:1-eve]
tmbo     16248  2964  0 12:30 tty2    00:00:07 /usr/bin/python2 /usr/bin/te
tmbo     16257 16248  0 12:30 pts/0   00:00:00 /bin/bash
root     16420   2  0 12:31 ?        00:00:00 [kworker/u16:2-e]
root     16435   2  0 12:31 ?        00:00:00 [kworker/4:1-eve]
root     16533   2  0 12:31 ?        00:00:00 [kworker/3:1-rcu]
root     16910   2  0 12:32 ?        00:00:00 [kworker/6:1-cgr]
root     16933   2  0 12:32 ?        00:00:00 [kworker/2:2-eve]
tmbo     17020  6015  0 12:33 tty2    00:00:05 /opt/google/chrome/chrome --
root     17885   2  0 12:34 ?        00:00:00 [kworker/7:2-rcu]
root     17179   2  0 12:35 ?        00:00:00 [kworker/0:0-eve]
root     18024   2  0 12:37 ?        00:00:00 [kworker/1:2-eve]
root     18025   2  0 12:37 ?        00:00:00 [kworker/1:4-eve]
root     18115   2  0 12:37 ?        00:00:00 [kworker/u16:3-e]
root     18373   2  0 12:38 ?        00:00:00 [kworker/4:0-eve]
root     18374   2  0 12:38 ?        00:00:00 [kworker/7:3-eve]
root     18446   2  0 12:38 ?        00:00:00 [kworker/3:2-eve]
root     19070   2  0 12:39 ?        00:00:00 [kworker/6:0-eve]
tmbo     19833  2798  1 12:40 ?        00:00:05 /usr/bin/nautilus --gapplica
root     19919   2  0 12:40 ?        00:00:00 [kworker/6:2]
root     19920   2  0 12:40 ?        00:00:00 [kworker/4:3]
root     19995   2  0 12:41 ?        00:00:00 [kworker/3:3-eve]
tmbo     20174  5993  0 12:42 tty2    00:00:00 /opt/google/chrome/chrome --
root     20195   2  0 12:43 ?        00:00:00 [kworker/0:2-eve]
tmbo     20211  6015  6 12:43 tty2    00:00:09 /opt/google/chrome/chrome --
tmbo     20313  6015  0 12:44 tty2    00:00:00 /opt/google/chrome/chrome --
root     20334   2  0 12:44 ?        00:00:00 [kworker/2:0-eve]
root     20391   2  0 12:45 ?        00:00:00 [kworker/u16:0]
tmbo     20406  2790 27 12:45 ?        00:00:01 /usr/libexec/tracker-extract
tmbo     20617 16257  0 12:45 pts/0   00:00:00 ps -ef
tmbo@xps ~$
tmbo@xps ~$
tmbo@xps ~$
tmbo@xps ~$
tmbo@xps ~$
tmbo@xps ~$ docker pull centos:
```

# Packaging container images

- An image can package up pretty well anything you want
- Package multiple components into one container or each component into separate containers and let them communicate with each other
- Typically defined using a 'Dockerfile'

<https://github.com/alanbchristie/PySimple>

```
FROM python:3.7.3-slim
WORKDIR /app
ADD . /app
EXPOSE 8080
RUN pip install -r requirements.txt
USER nobody
CMD ["python", "app.py"]
```

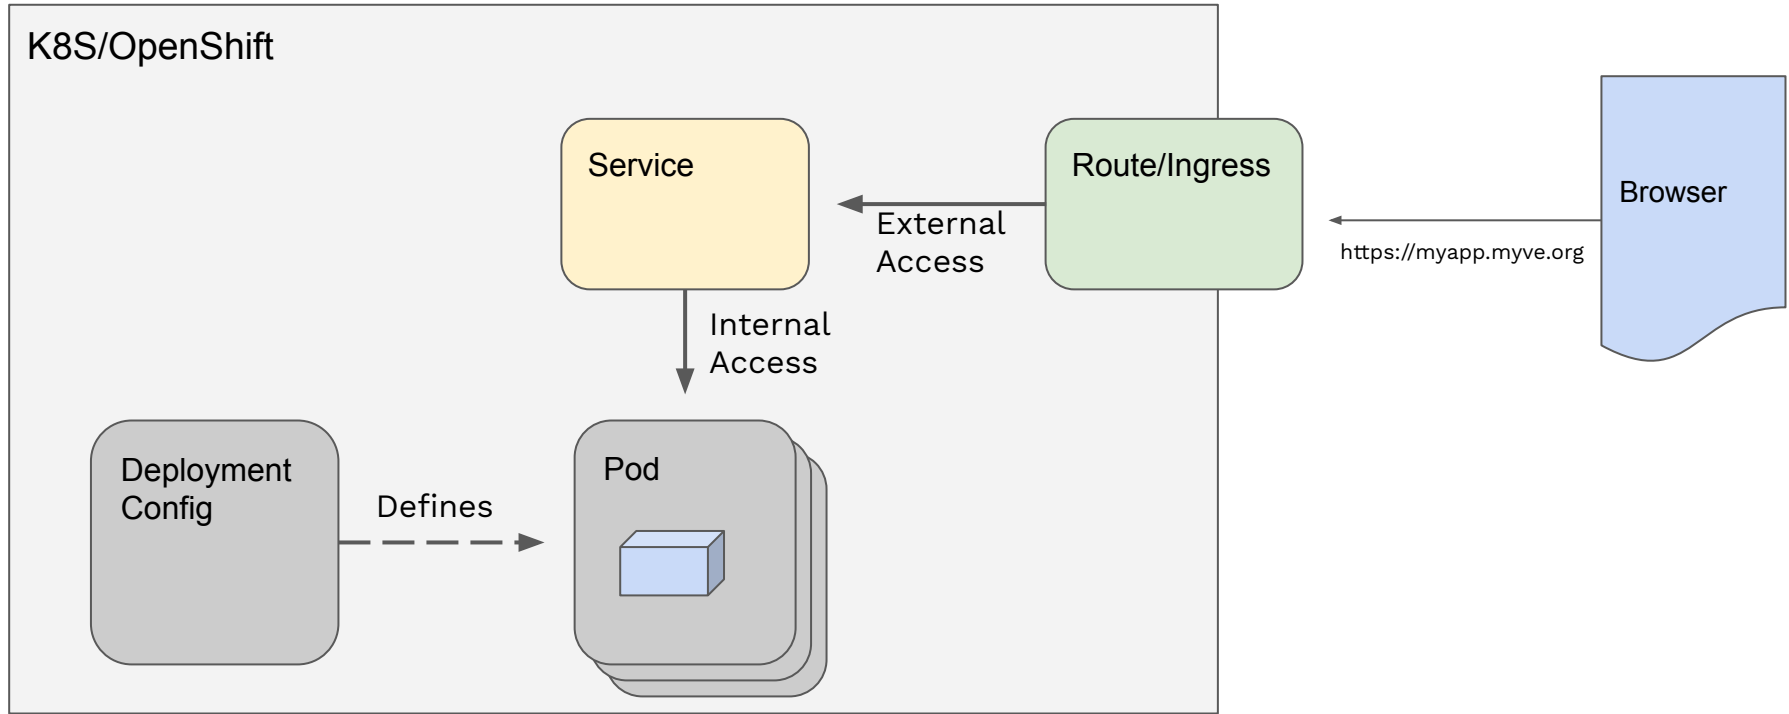
```
$ git clone git@github.com:alanbchristie/PySimple.git
$ cd PySimple
$ docker build -t tdudgeon/pysimple .
$ docker push tdudgeon/pysimple

...

$ docker pull tdudgeon/pysimple
$ docker run -d -p 8080:8080 tdudgeon/pysimple

$ curl http://localhost:8080/
```

# Deploying container images on an ORN VE





# Procedure for deploying applications

Step 1: Create your container images

Step 2: Deploy to OpenShift - multiple approaches possible

# Step 1: creating container images

Create your container images externally and push to registry such as DockerHub

or

Use OpenShift's CI/CD mechanisms to build the container images and push to OpenShift's own container registry running in the VE

# Step 2: Deploy to OpenShift

Multiple approaches possible

1. Web console vs. CLI vs. REST API
2. Manual/interactive procedure
3. Templates
4. Operators

We will show some examples.

# Deployment examples

1. Deploy app through web console
2. Deploy app using CLI
3. Deploy Lazar using CLI
  - a. The Lazar template
  - b. Deploying
4. Deploying Lazar from web console
5. Deploying Squonk using Ansible
  - a. Templates
  - b. Playbooks



# Anatomy of the Lazar template

<https://github.com/OpenRiskNet/home/tree/master/openshift/deployments/lazar>

## Template

### Metadata

e.g.  
name = lazar

### Labels

e.g.  
app = lazar

### Parameters

e.g.  
LAZAR\_SERVICE\_PORT = 8088

## Objects

**ImageStream**

**DeploymentConfig**

**Service**

**Route**

# Lazar template - Parameters

parameters:

```
- name: IMAGE_TAG
```

Parameter name

```
description:
```

```
    The lazar docker image tag
```

```
value: latest
```

Optional default value

```
- name: ROUTE_NAME
```

```
description:
```

```
    The name of the service endpoint.
```

```
    This is typically *lazar* but if you want different instances  
    you can pass another name like *lazar-2*.
```

```
value: lazar
```

```
- name: ROUTES_BASENAME
```

```
description:
```

```
    The base name of the service endpoint.
```

```
    This is typically the PROD or DEV URL basename.
```

```
value: prod.openrisknet.org
```

Parameters allow to  
configure the  
deployment

# Lazar template - DeploymentConfig

```
- kind: DeploymentConfig
  apiVersion: v1
  ...
  spec:
    ...
    template:
      ...
      spec:
```

Container specification

```
  containers:
  - name: lazar
    image: docker.io/gebele/lazar-rest: ${IMAGE_TAG}
    ports:
    - containerPort: 8088
      protocol: TCP
    ...
    imagePullPolicy: Always
```

Use of a parameter

```
  readinessProbe:
    httpGet:
      path: "/"
      port: 8088
      scheme: "HTTP"
    failureThreshold: 4
    initialDelaySeconds: 30
    periodSeconds: 30
    timeoutSeconds: 4
  livenessProbe:
    ...
```

Readiness and liveness probes

```
  resources:
    requests:
      cpu: ${CPU_REQUEST}
      memory: ${MEMORY_REQUEST}
    limits:
      cpu: ${CPU_LIMIT}
      memory: ${MEMORY_LIMIT}
```

Resource requests and limits

# Lazar template - Service and Route

```
- kind: Service
```

```
  apiVersion: v1
```

```
  metadata:
```

```
    name: lazar
```

```
    ...
```

```
  spec:
```

```
    ports:
```

```
      - name: lazar
```

```
        protocol: TCP
```

```
        port: ${LAZAR_SERVICE_PORT}
```

```
        targetPort: 8088
```

```
        nodePort: 0
```

```
    selector:
```

```
      name: lazar
```

```
    type: ClusterIP
```

```
    sessionAffinity: None
```

Port spec

Pods to service

```
- kind: Route
```

```
  apiVersion: v1
```

```
  metadata:
```

```
    name: ${ROUTE_NAME}
```

```
  annotations:
```

```
    kubernetes.io/tls-acme: ${TLS}
```

```
  spec:
```

```
    host: ${ROUTE_NAME}.${ROUTES_BASENAME}
```

```
    to:
```

```
      kind: Service
```

```
      name: lazar
```

```
    tls:
```

```
      termination: edge
```

```
      insecureEdgeTerminationPolicy: Redirect
```

Public hostname



# Lazar Template - Template Service Broker

metadata:

name: lazar

annotations:

```
openshift.io/display-name: lazar toxicity prediction service
openshift.io/provider-display-name: Johannes Gutenberg University Mainz - JGU, in silico toxicology gmbh - IST
openshift.io/documentation-url: https://github.com/OpenRiskNet/home.git
openshift.io/support-url: https://github.com/OpenRiskNet/home/issues
description: lazar (lazy structure-activity relationships) is a modular framework for predictive toxicology.
  Similar to the read across procedure in toxicological risk assessment, lazar creates local
  QSAR (quantitative structure-activity relationship) models for each compound to be predicted.
iconClass: ''
tags: lazar,prediction,rest
```

These annotations are used by the Template Service Broker to allow the template to be deployed easily through the web console - demo coming later

# Other ORN app templates

<https://github.com/OpenRiskNet/home/tree/master/openshift/deployments>

Contains templates for most of the ORN partner applications plus some additional 3rd party applications such as JupyterHub.

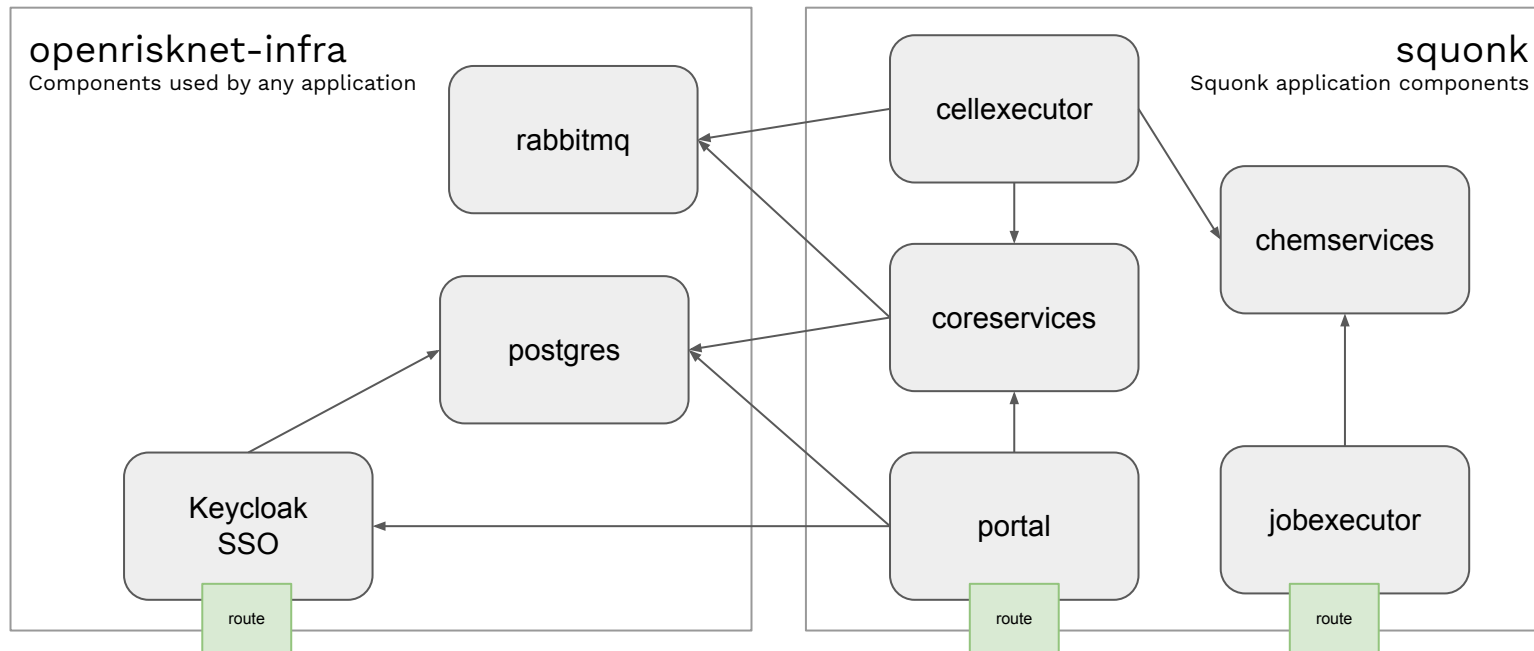
Simple examples: [bridgedb](#), [jguweka](#)

More complex examples: [jupyterhub](#), [squonk](#), [jaqpot](#)

We'll now look at Squonk as a more complex example.

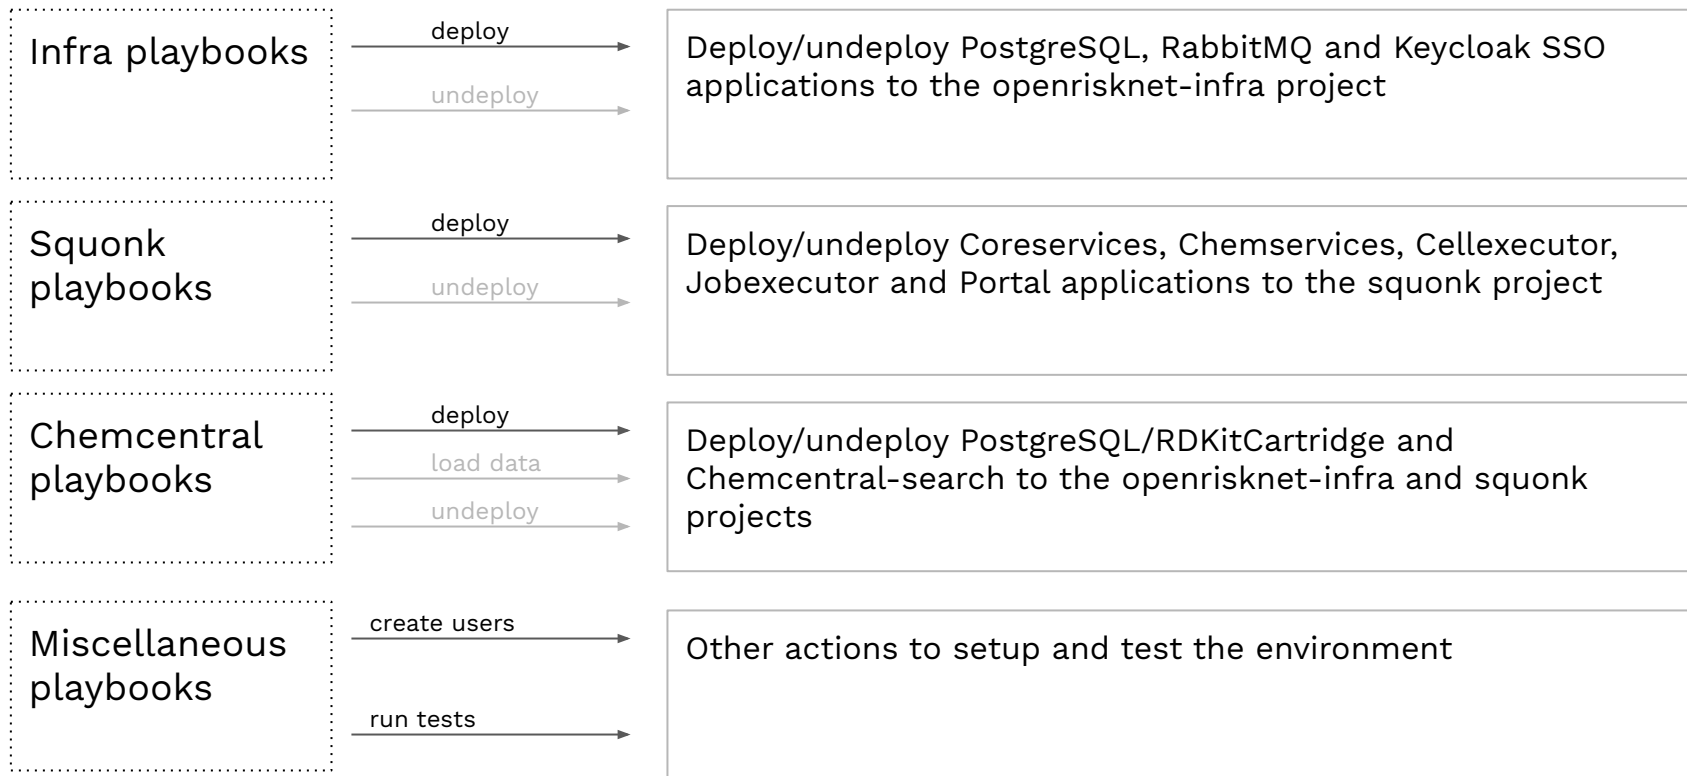
# Squonk architecture overview

some of the squonk components



# Squonk deployment

<https://github.com/InformaticsMatters/squonk/tree/master/openshift/ansible>



# Best practices

- Consider security aspects of your containers
- Try to create small containers
- Consider how much resource (CPU, memory) your containers need
- Use SSO for authentication

Guidelines are provided here:

<https://github.com/OpenRiskNet/home/wiki/Deployment-Guidelines>

For full details of the ORN partner and 3rd party applications that can be deployed to an ORN VE look here: <https://home.prod.openrisknet.org/>

### OpenRiskNet and Thrid-Party Workflow Managers and Scripting Tools

→ [Squonk Computational Notebook](#)

→ [Jupyter Notebooks](#)

Please note that the jupyter container is very large and needs some time to be deployed on a specific node of the reference instance. Please press the "refresh" button of your browser until the interface is appearing. [Example workflows can be accessed here.](#)

### Graphical User Interface Access to OpenRiskNet Applications

→ [Lazar Toxicity Predictions](#)

### OpenRiskNet Data Sources

→ [Nanomaterial database](#)

→ [Data Explorer serving ToxCast, ToxRefDB and TG-Gates data](#)

### Example Workflows based on OpenRiskNet Tools

→ [Jupyter Notebook: Access TG-Gates data for seleted compounds, select differentially expressed genes and identifier relevant pathways](#)

→ [Jupyter Notebook: Cleaning LTKB data prior to generating predictive models](#)

# Conclusion

- OpenShift/Kubernetes is a powerful application platform
- A wide range of options for deploying applications
- Support for simple and complex application topologies
- Can also include building applications from source
- Significant learning curve is involved
- But provides excellent approach for robust and automated deployment of applications

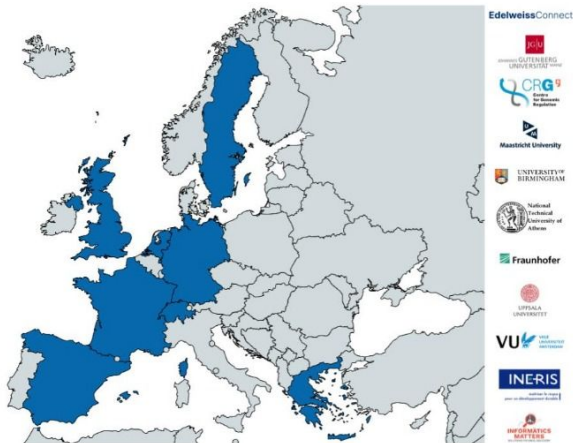


# Acknowledgements

**OpenRiskNet** (Grant Agreement 731075) is a project funded by the European Commission within Horizon 2020 Programme



## Project partners:



- P1 Edelweiss Connect GmbH, Switzerland (EwC)
- P2 Johannes Gutenberg-Universität Mainz, Germany (JGU)
- P3 Fundacio Centre De Regulacio Genomica, Spain (CRG)
- P4 Universiteit Maastricht, Netherlands (UM)
- P5 The University Of Birmingham, United Kingdom (UoB)
- P6 National Technical University Of Athens, Greece (NTUA)
- P7 Fraunhofer Gesellschaft Zur Foerderung Der Angewandten Forschung E.V., Germany (Fraunhofer)
- P8 Uppsala Universitet, Sweden (UU)
- P9 Medizinische Universität Innsbruck, Austria (MUI)
- P10 Informatics Matters Limited, United Kingdom (IM)
- P11 Institut National De L'environnement Et Des Risques INERIS, France (INERIS)
- P12 Vrije Universiteit Amsterdam, Netherlands (VU)

