



ICAT-k8s

**A Helm chart for deploying ICAT in
Kubernetes**

Rodrigo Cabezas Quirós

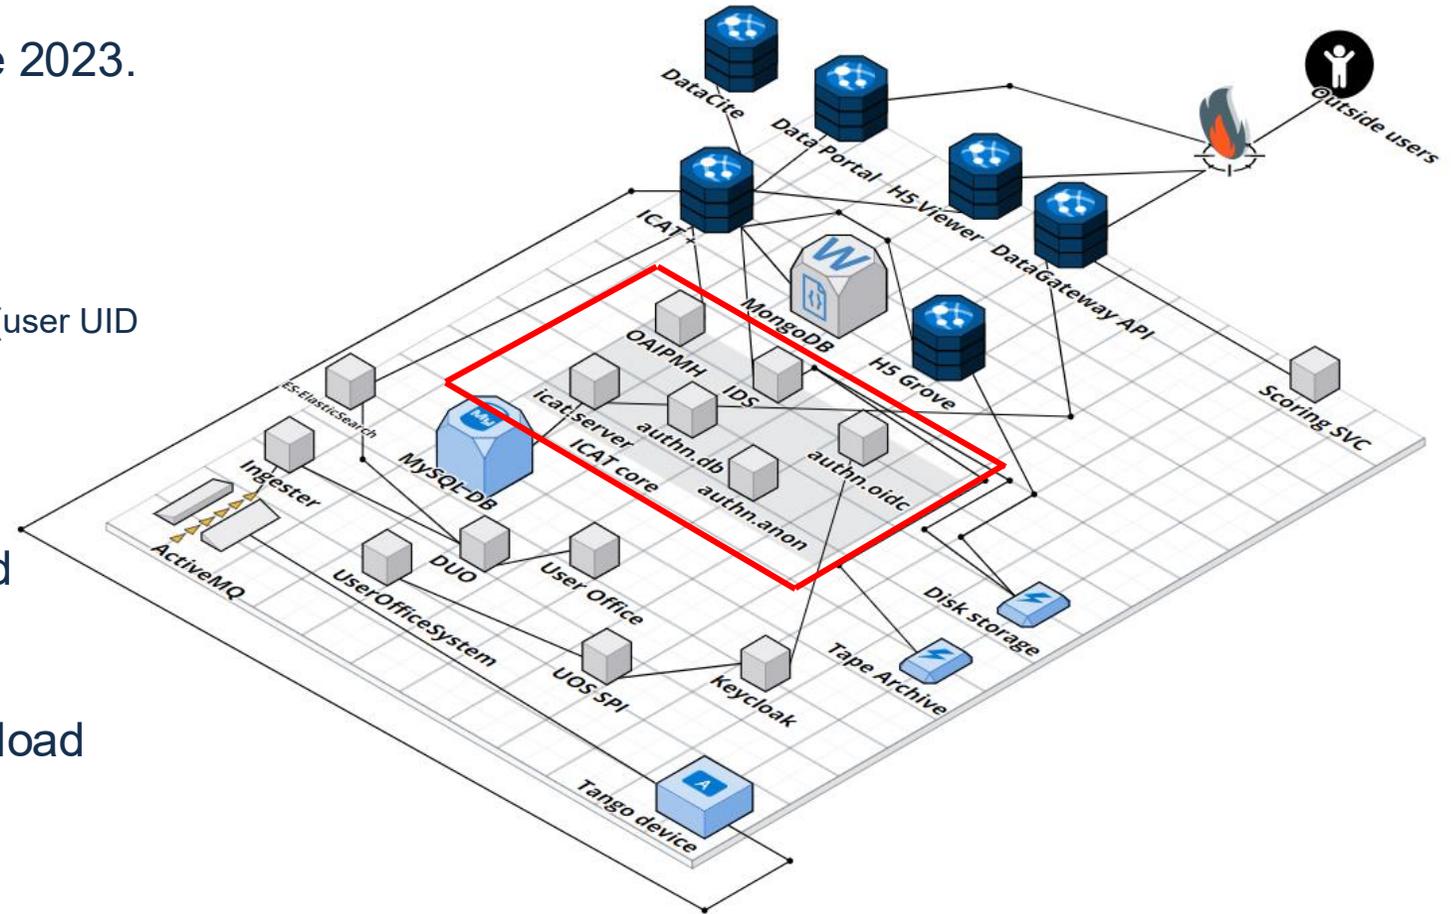


Table of contents

- Context and motivation.
- Helm chart overview.
- Implementation.
- Deployment.
- Next steps.

Context

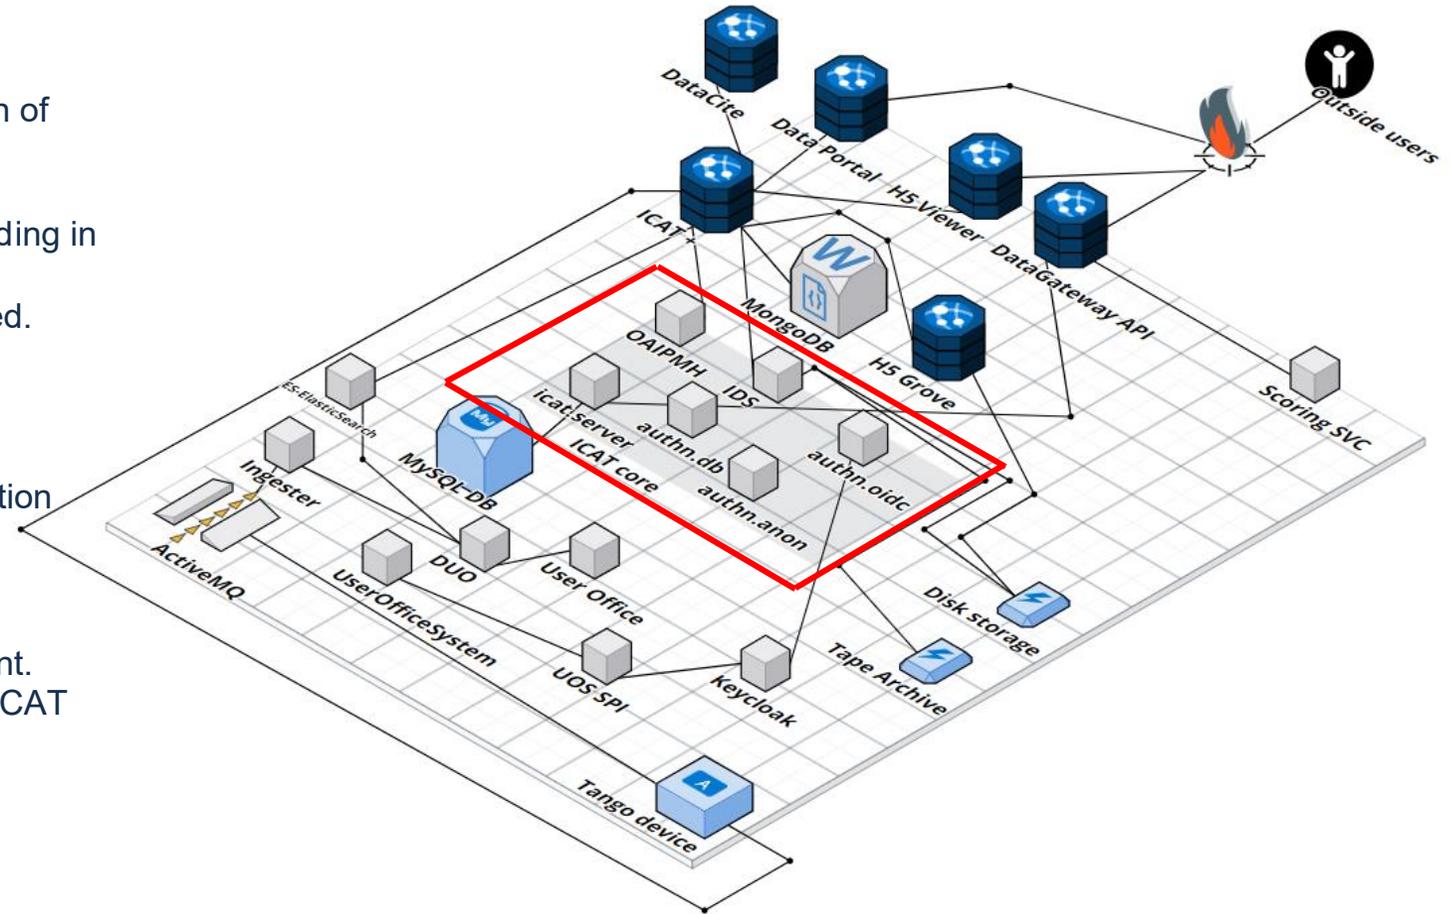
- ICAT deployed in Kubernetes since 2023.
- Using a slightly modified version of `docker.icat1`:
 - Not all core components are deployed.
 - Base image modified to support NFS mounts (user UID adjustment).
 - Elasticsearch APM monitoring added.
- The cluster was running an older Kubernetes version (~3–4 years old at the time).
- ICAT was the first large-scale workload deployed in the cluster.



1. <https://github.com/RKrahl/docker.icat>

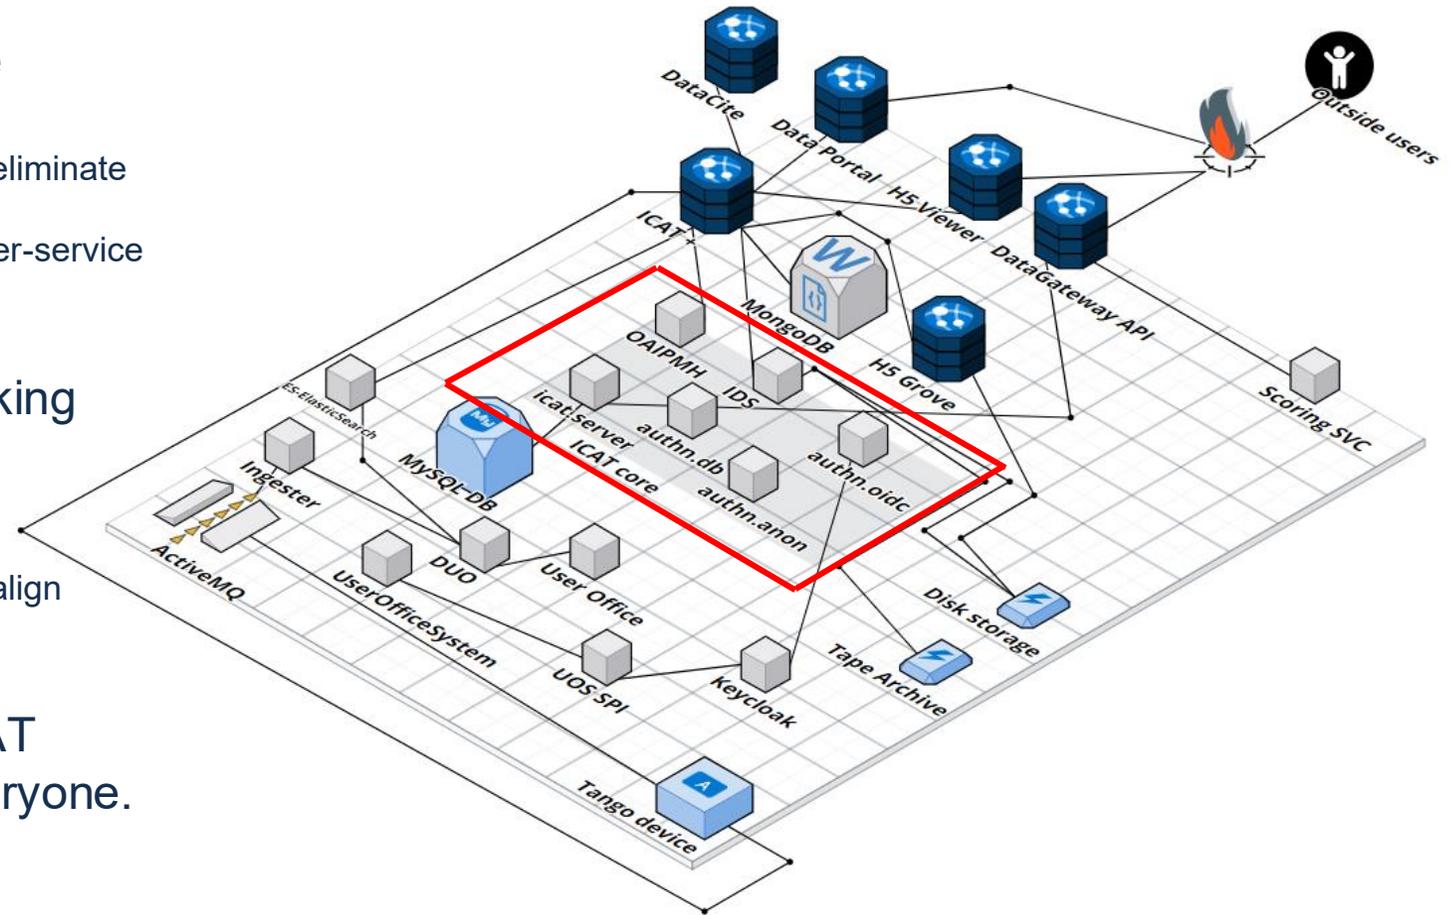
Context

- Cluster-related issues during 2024:
 - Failure to prune completed pods led to eviction of running services.
 - Network plugin experienced certificate errors.
 - CoreDNS issues led to excessive request flooding in some applications.
 - Partial high-availability clusters underperformed.
- Issues with ICAT's deployment:
 - ICAT pods were repeatedly evicted, causing significant downtime and impacting data ingestion / access.
 - Long ICAT startup times.
 - Resource usage per component could not be measured due to bundled container deployment.
 - Unable to scale individual components of the ICAT stack.



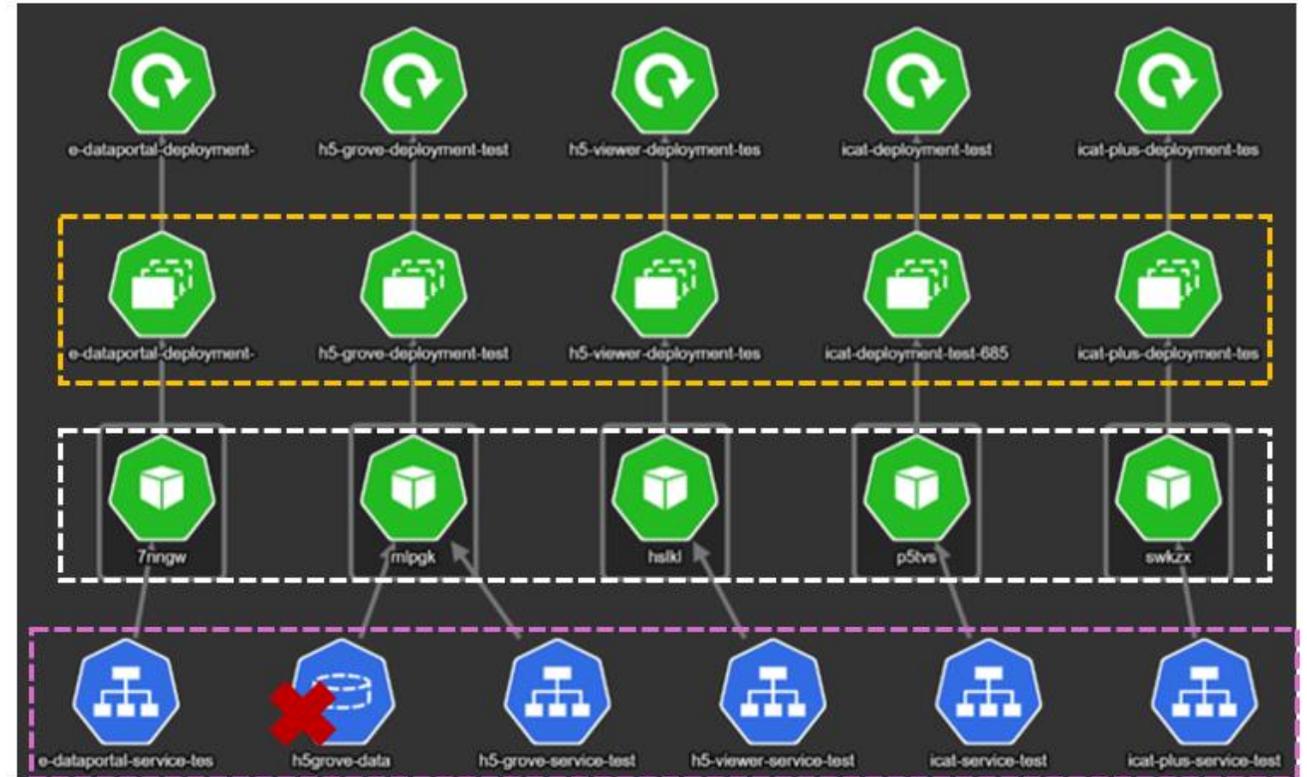
Motivation

- ICAT architecture and performance improvements:
 - Split ICAT core services into multiple pods to eliminate monolithic architecture.
 - Optimize ICAT startup time, accounting for inter-service dependencies.
- Upgrade ICAT and some of its backing components:
 - Migrate database from MySQL to MariaDB.
 - Upgrade ICAT to version 6.
 - Update user and proposal synchronization to align with the new User Office application.
- Enable easy and standardized ICAT deployment on Kubernetes for everyone.



Helm chart overview

- Typical Kubernetes deployment:
 - Deployment / StatefulSet: pod blueprint.
 - Service: internal load balancer for pods.
 - Ingress: external access to a service.
- ICAT deployment context:
 - Each Kubernetes component is defined in a YAML file.
 - ICAT orchestration requires 17 separate YAML files.
- The problem? WTMY or YAMLpocalypse.
 - High complexity and risk of misconfiguration
 - Difficult to maintain consistency and reuse components.
 - Mentally taxing and error-prone.
 - Troubleshooting nightmares.
 - Small changes, automation, and parameterization are difficult to implement safely.



Deployments, pods and **ingresses**.

Helm chart overview



What is a Helm Chart?

A package format for Kubernetes applications

- Chart.yaml
- values.yaml
- templates/
- charts/

Key Concepts

- Chart** A packaged application
- Release** A deployed instance
- Values** Configuration parameters

Advantages

- ✓ Consistency
- ✓ Reuse
- ✓ Versioning

Define → Package → Deploy → Manage



Deploy

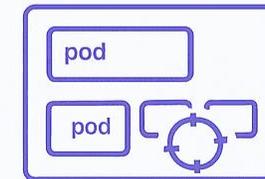
Manage



Advantages

- ✓ Consistency
- ✓ Reuse
- ✓ Versioning
- ✓ Templating

Example Use Case



simplyblock.io

Helm chart goals

- **Single-file configuration:**
 - Centralize all ICAT parameters in a single 'values.yaml' file.
 - Eliminate separate configuration files (e.g., logback.xml, *.properties, *.xslt).
- **Container strategy:**
 - Use standard Payara Docker images; avoid custom modifications*.
 - Prefer smaller Payara Micro images to improve startup times and reduce resource usage.
- **Deployment simplicity:**
 - Enable easy, one-line deployment.

```
helm install icat-test2 icat-k8s -n icat-core-test -f values.yaml
```

Implementation

- For lightweight components use Payara

Micro:

- authn.anon
- authn.db
- authn.oidc
- icat.oaipmh

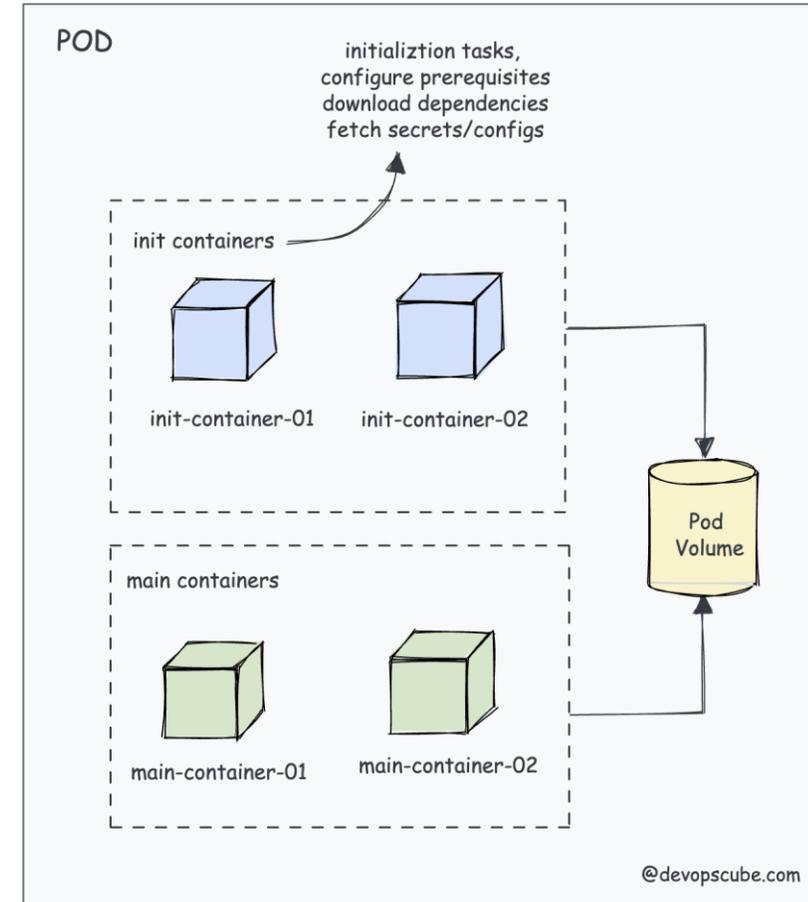
- For heavier components*, use Payara Server Full:

- icat.server
- ids.server
- *: or any other component that relies on a message broker

	Payara Micro	Payara Server Full
Purpose	Microservices, lightweight containers	Traditional Java apps
Startup time	Very fast	Slower
Image size	Small	Larger (full EE stack)
Resource usage	Lower CPU / RAM	Higher
Configuration	CLI options, environment vars	Admin console, config files
Deployment model	One service per container	Multiple apps per server
Built-in message broker	No	Yes

Implementation

- All components setup.py scripts bundled into a single file.
- Use of init-containers:
 - For setting up each component:
 1. Prepare each component by fetching artifacts and libraries from the repository.
 2. Execute setup.py to add or replace *.properties, logback.xml and *.xslt files within the WAR artifacts and generate 'post_boot_asadmin_commands' files.
 - Manage startup order by waiting for dependent services (e.g., icat.server waits for authn.anon, authn.db, and authn.oidc; ids.server waits for icat.server).



Implementation

- Example for IDS server:

State	Ready	Name	Image	Init Container	Restarts	Started
Waiting	✗	icat-core-ids-server-test	registry.cells.es/mis/icat/payara/server-full:6.2025.3-jdk21	—	0	-
PodInitializing						
Terminated	✓	init-ids-server-test-setup	registry.cells.es/docker/python:3.12.9-bullseye	✓	0	-
Completed						
Running	✗	init-test-wait-for-dependen- svc-availability	registry.cells.es/docker/curlimages/curl:8.13.0	✓	0	12 secs ago

1

```
badassfish@icat-core-ids-server-deployment-test-icat-test2-df97dc7b7-jl7ct:/opt/payara/deployments$ cat post_boot_asadmin_commands
add-library --type app /opt/payara/libs/elastic-apm-agent.jar
add-library --type app /opt/payara/libs/ids-storage_file_alba.jar
create-jms-resource --restype jakarta.jms.Topic jms/IDS/log
```

2

```
Thu, Feb 5 2026 11:35:13 am Checking http://icat-core-icat-server-service-test-icat-test2:8080/icat/version/ - Status: 000
Thu, Feb 5 2026 11:35:15 am Checking http://icat-core-icat-server-service-test-icat-test2:8080/icat/version/ - Status: 000
Thu, Feb 5 2026 11:35:17 am Checking http://icat-core-icat-server-service-test-icat-test2:8080/icat/version/ - Status: 000
Thu, Feb 5 2026 11:35:19 am Checking http://icat-core-icat-server-service-test-icat-test2:8080/icat/version/ - Status: 000
Thu, Feb 5 2026 11:35:21 am Checking http://icat-core-icat-server-service-test-icat-test2:8080/icat/version/ - Status: 200
```

Implementation

- values.yaml:

```
1 namespace: icat-core-test2
2 environment: test
3
4 ingress:
5   enabled: false
6   host: icatatest.cells.es
7   ingressClassName: nginx-main
8   tlsSecretName: icatatest.cells.es-tls
9
10 images:
11   container:
12     micro: '<<REDACTED>>.cells.es/mis/icat/payara/micro:6.2025.2-jdk21'
13     serverFull: '<<REDACTED>>.cells.es/mis/icat/payara/server-full:6.2025.3
14               -jdk21'
15   initContainers:
16     python: '<<REDACTED>>.cells.es/docker/python:3.12.9-bullseye'
17     curl: '<<REDACTED>>.cells.es/docker/curlimages/curl:8.13.0'
18 imagePullSecretName: harbor-registry-key
19 ...
```

Implementation

- values.yaml:

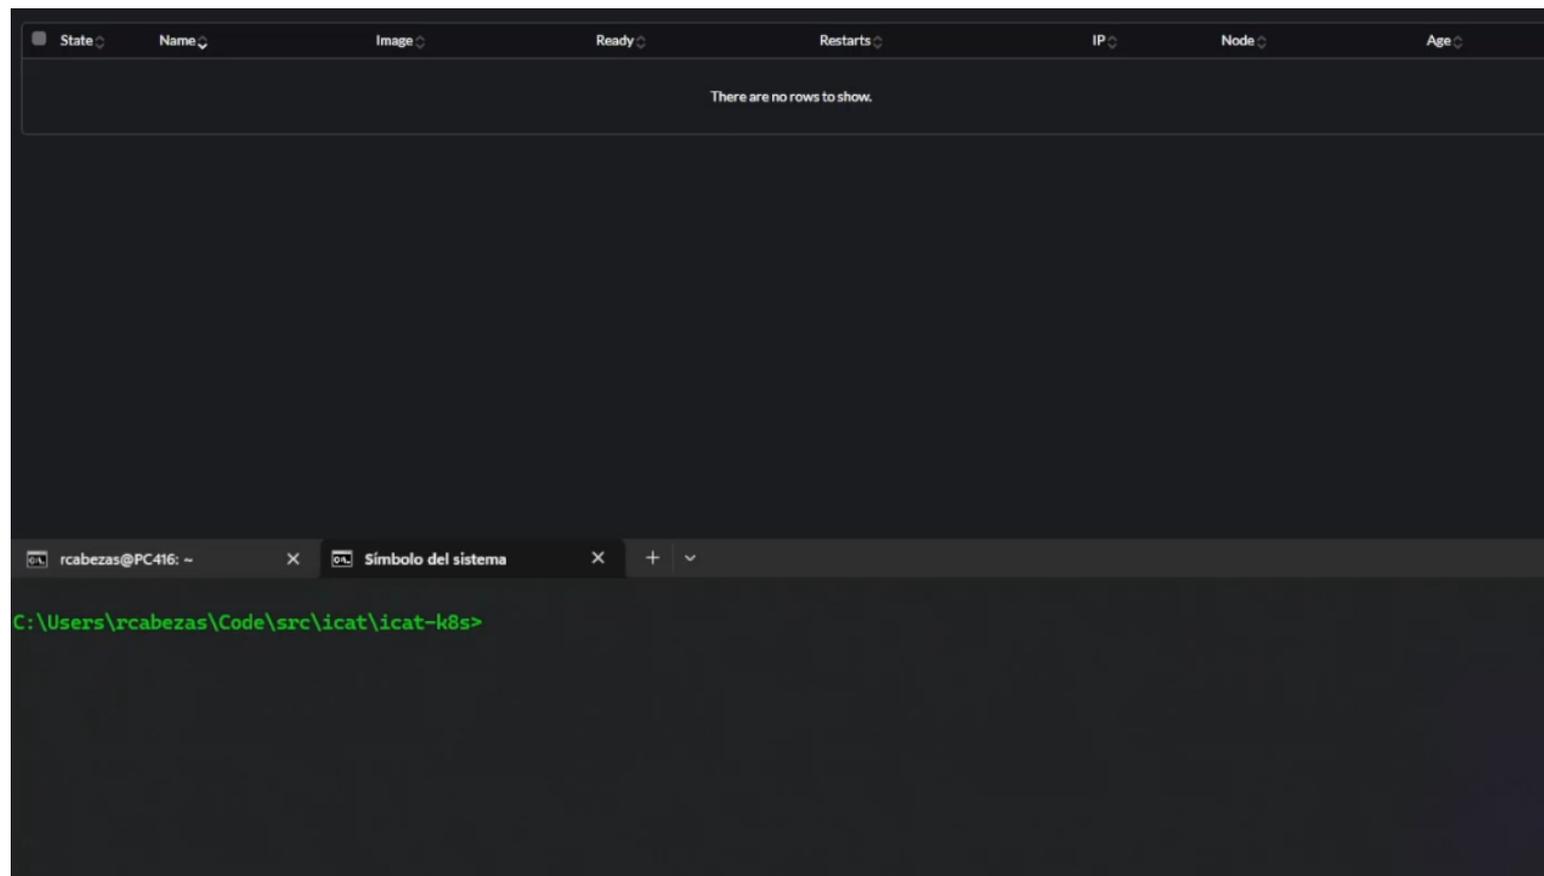
```
21 artifacts:
22   rootRepositoryURL: "https://<<REDACTED>>.cells.es/nexus/repository/icat-alba"
23   databaseConnector:
24     artifactPath: "/org/icatproject/mariadb-java-client/3.5.2/mariadb-java
25       -client-3.5.2.jar"
26     filename: "mariadb-java-client.jar"
27     type: "lib"
28     #repositoryURL: "https://youcanoverridethisifyouwant.nice"
29   elasticAPM:
30     artifactPath: "/elastic/elastic-apm-agent/1.50.0/elastic-apm-agent-1.50.0
31       .jar"
32     filename: "elastic-apm-agent.jar"
33     type: "lib"
34     #repositoryURL: "https://youcanoverridethisifyouwant.nice"
35   idsStorageFile:
36     artifactPath: "/org/icatproject/ids.storage_file_alba/1.4.4_ALBA/ids
37       .storage_file_alba-1.4.4_ALBA.jar"
38     filename: "ids-storage_file_alba.jar"
39     type: "lib"
40   icatProjectComponents:
41     baseArtifactPath: "/org/icatproject/"
42     #repositoryURL: "https://youcanoverridethisifyouwant.nice"
43   initScripts:
44     url: "https://<<REDACTED>>.cells.es/mis-devops/icat-public/icat-k8s/-/raw
45       /master/setup_scripts"
46
47 deploymentPriorityClassName: "icat-priority"
48
49 timezone: "Europe/Madrid"
```

Implementation

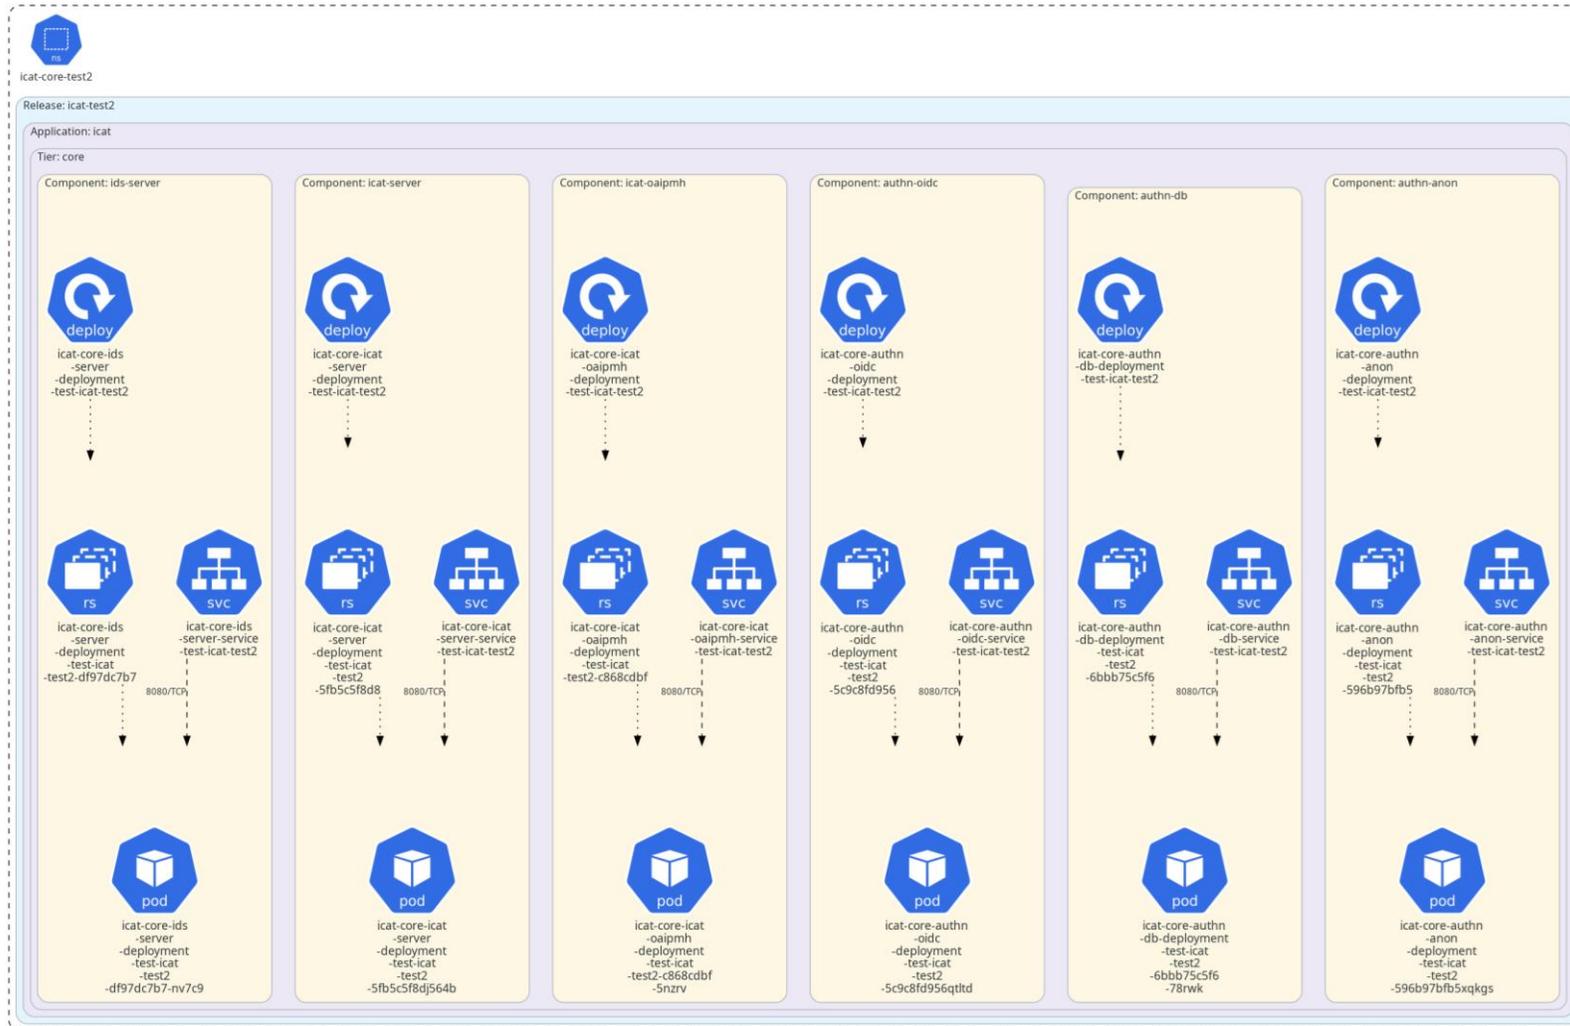
- values.yaml:

```
49 components:
50 # ids.server -- ICAT DATA SERVICE (IDS) SERVER COMPONENT
51 - name: "ids-server"
52   containerType: "serverFull"
53   repoKey: "ids.server"
54   version: "2.1.1"
55   replicas: 1
56   startupDependencies:
57     - name: "icat-server"
58       path: "/icat/version/"
59       port: 8080
60   resources:
61     limits:
62       memory: "1Gi"
63       cpu: "4.0"
64     requests:
65       memory: "1Gi"
66       cpu: "1.0"
67   readinessProbe:
68     httpGet:
69       path: "/ids/ping/"
70       port: 8080
71     failureThreshold: 1
72     periodSeconds: 2
73   livenessProbe:
74     httpGet:
75       path: "/ids/ping/"
76       port: 8080
77     periodSeconds: 20
78   mountInstrumentStorage: false
79   setupProperties: |
80     ...
81   runProperties: |
82     icat.url = http://icat-core-icat-server-service-test.icat-core-test:8080
83     ...
84   logbackXml: |
85     <configuration>
```

Deployment



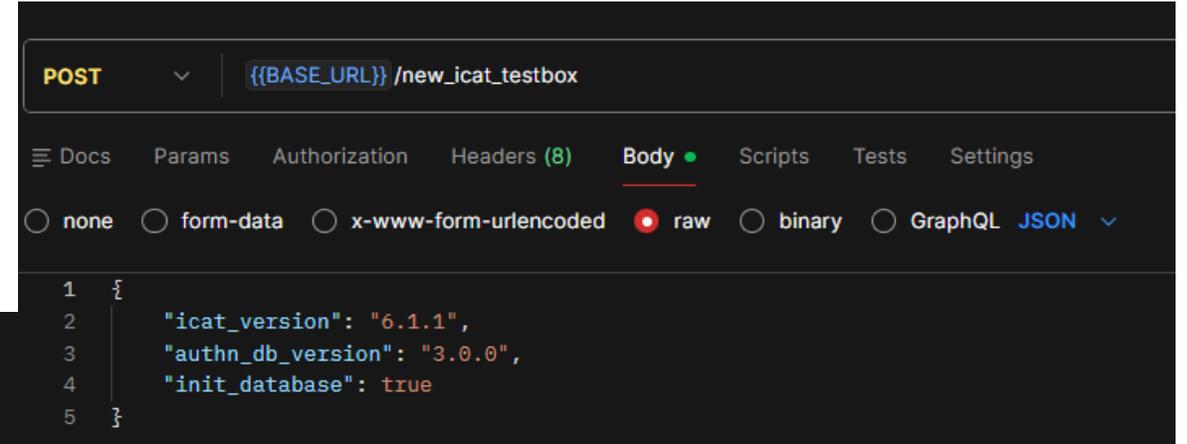
Deployment



Next steps

- Chart publication:
 - Publish the Helm chart to GitHub and Artifact Hub for public use.
 - Validate the chart with other ICAT components (e.g., Lucene).
- icat-testbox publication.
 - Publish icat-testbox, a lightweight project for creating ephemeral ICAT instances for unit and integration testing.
 - Initially developed to support testing of the new ingestion system (PACER); work in progress.

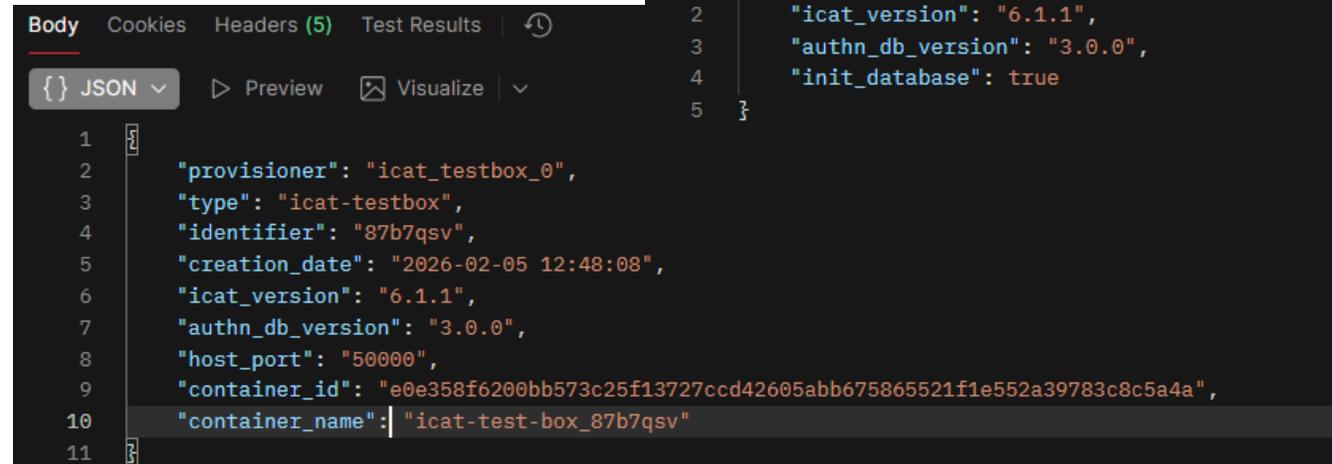
<https://github.com/ALBA-Synchrotron/icat-k8s>



```

POST {{BASE_URL}}/new_icat_testbox
Body
  none
  form-data
  x-www-form-urlencoded
  raw
  binary
  GraphQL
  JSON
  {
    "icat_version": "6.1.1",
    "authn_db_version": "3.0.0",
    "init_database": true
  }

```



```

Body
  Cookies
  Headers (5)
  Test Results
  {} JSON
  Preview
  Visualize
  {
    "provisioner": "icat_testbox_0",
    "type": "icat-testbox",
    "identifier": "87b7qsv",
    "creation_date": "2026-02-05 12:48:08",
    "icat_version": "6.1.1",
    "authn_db_version": "3.0.0",
    "host_port": "50000",
    "container_id": "e0e358f6200bb573c25f13727ccd42605abb675865521f1e552a39783c8c5a4a",
    "container_name": "icat-test-box_87b7qsv"
  }

```





cells.es

Cerdanyola del Vallès (Barcelona)

Spain

Tel: (+34) 93 592 43 00