



# SKA TANGO Operator

M. Di Carlo, P. Harding et al.  
INAF – Osservatorio Astronomico d'Abruzzo



# SKA Telescopes




## Mid Telescope

Location: South Africa

350 Mhz to 15.3 GHz

197 dishes - max baseline: 150km<sup>(1)</sup>

<sup>(1)</sup> Data for SKA1 implementation



## Low Telescope

Location: Australia

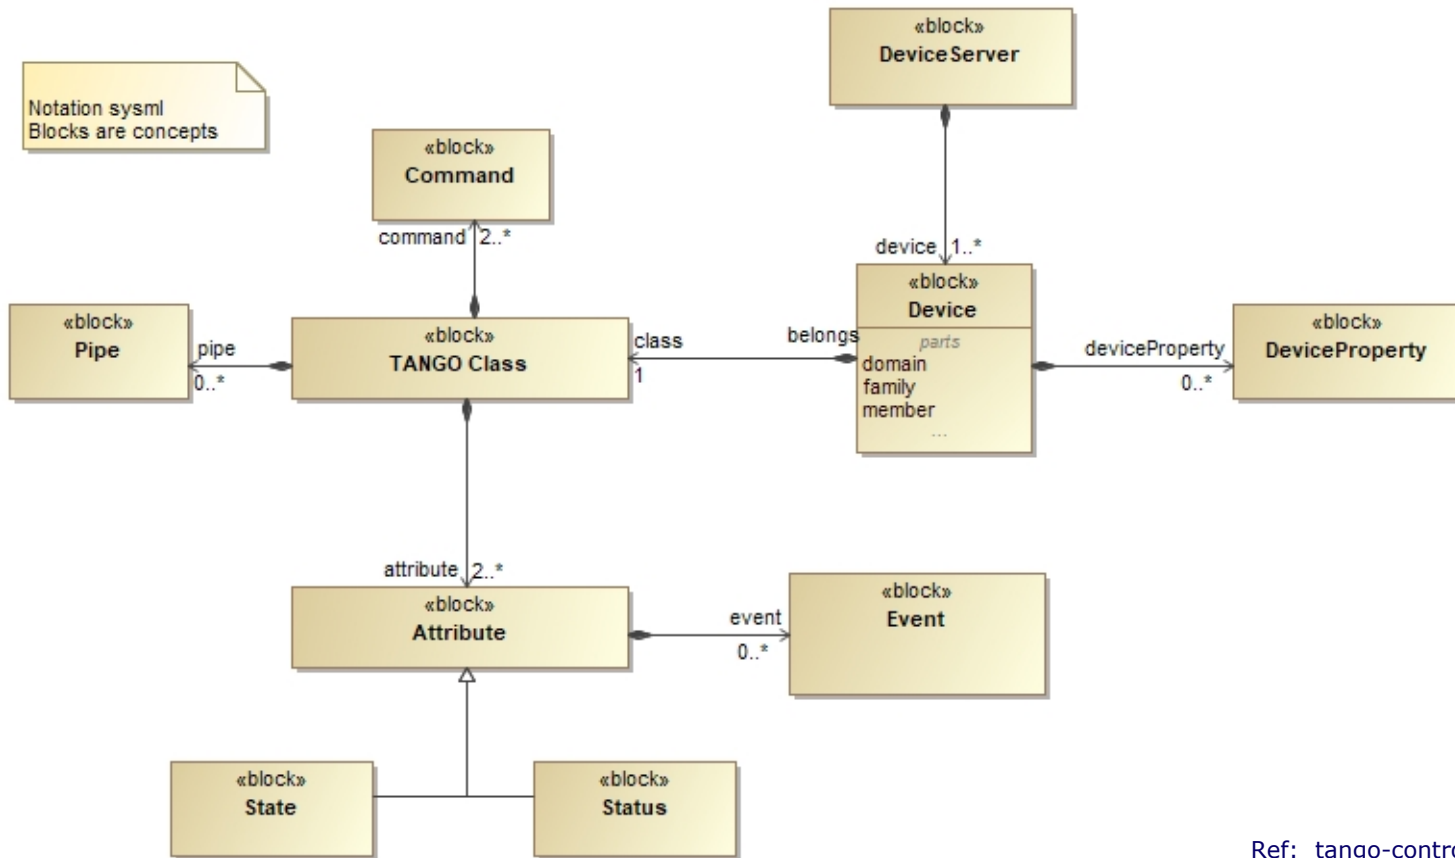
50 MHz to 350 GHz

Ref: [skao.int](http://skao.int)

# SKA Development and Deployment Practices

- Kubernetes (**k8s**) for container orchestration ([kubernetes.io](https://kubernetes.io))
- **Helm** for packaging and deploying SKA Software ([helm.sh](https://helm.sh))
  - Tool for managing charts, where a **chart** is a recipe to deploy several **k8s** resources (i.e., containers, storage, networking components, etc) required for an application to run
  - Works on templates, allows to adapt generic configurations to different environments (i.e., the different SKA datacentres)

# TANGO-controls framework



Ref: [tango-controls.readthedocs.io](https://tango-controls.readthedocs.io)

# Kubernetes

- Open-source system for **automating** the deployment, **scaling**, and **management** of containerized applications
  - Very extensible with multiple plugin points
  - Collection of primitives: an abstraction of compute, network, and storage
  - Supports common scheduling models: stateful, replicas (stateless), daemons and jobs
- Declarative “model” of operation, with various controller components managing the lifecycle of the declared resources

# Imperative vs Declarative

- "Imperative" is a command
  - Add this device to the DB
  - Execute this python command
  - ...
- "Declarative" is a statement of the desired state
  - Deploy these devices, with this configurations, in this specific order, with those dependencies...

# Setting up a TANGO Device server (in general)

- For every **device server**:
  - Deploy the code (i.e. to a VM)
    - Configuration (i.e., environment variables)
    - Resources (i.e., persistent storage, networking)
  - Include all of the devices in the TANGO **database**:
    - Properties
    - Polling definition
    - ...
  - Start devices in **order** (if there are dependencies, wait for each of them)

# Device server in k8s

- A device server will be composed by the following k8s resources:
  - **Job** for the initialization of the TANGO database
  - **Service** (exposes the application to the network)
  - **Statefulset** (in k8s this is an object used to manage stateful applications)
    - Pod (in k8s groups one or more containers into a unit)
    - InitContainers (specialized containers that run before app containers in a Pod) to resolve the DS dependencies
      - Waiting for the configuration job to complete
      - Waiting on devices to be up in TANGO



# Declaring a DS in k8s

Name is the k8s name of the resources

The list of dependencies: devices or simple host:port

Command or entry points of the device servers (if more than one entry point is specified, we are referring to a multi-device DS)

The server definition, containing the list of instances, devices and classes to be ran

The container image to use

How/when to check if the device server is ready or in failure

```
1 name: "name-used-in-k8s"
2 function: description-text
3 domain: description-text
4 legacy_compatibility: false
5 instances: ["01"]
6 depends_on:
7   - device: sys/database/2
8   - device: sys/motor/1
9 entrypoints:
10  - path: "<optional-path-to-python-file.py>"
11    name: "module.ClassName"
12  command: "<optional: the python command to use>"
13 server:
14   name: "server-name"
15   instances:
16     - name: "01"
17     classes:
18       - name: "ClassName"
19     devices:
20       - name: "test/mydevice/3"
21 image:
22   registry: artefact.skao.int
23   image: ska-tango-examples
24   tag: 0.4.24
25   pullPolicy: IfNotPresent
26 livenessProbe:
27   initialDelaySeconds: 0
28   periodSeconds: 10
29   timeoutSeconds: 3
30   successThreshold: 1
31   failureThreshold: 3
32 readinessProbe:
33   initialDelaySeconds: 0
34   periodSeconds: 10
35   timeoutSeconds: 3
36   successThreshold: 1
37   failureThreshold: 3
```

# Declaring a DS in k8s

Name is the k8s name of the resources

The list of dependencies: devices or simple host:port

Command or entry points of the device servers (if more than one entry point is specified, we are referring to a multi-device DS)

The server definition, containing the list of instances, devices and classes to be ran

The container image to use

How/when to check if the device server is ready or in failure

```
1 name: "name-used-in-k8s"
2 function: description-text
3 domain: description-text
4 legacy_compatibility: false
5 instances: ["01"]
6 depends_on:
7   - device: sys/database/2
8   - device: sys/motor/1
9 entrypoints:
10  - path: "<optional-path-to-python-file.py>"
11    name: "module.ClassName"
12    command: "<optional: the python command to use>"
13 server:
14   name: "server-name"
15   instances:
16     -
17       pullPolicy: IfNotPresent
18       livenessProbe:
19         initialDelaySeconds: 0
20         periodSeconds: 10
21         timeoutSeconds: 3
22         successThreshold: 1
23         failureThreshold: 3
24       readinessProbe:
25         initialDelaySeconds: 0
26         periodSeconds: 10
27         timeoutSeconds: 3
28         successThreshold: 1
29         failureThreshold: 3
```

Partial set of parameters to set!



# Challenges with ska-tango-util helm library chart

- Creates **extra** resources in the Kubernetes cluster (i.e., extra init-containers to solve dependencies)
- Leaves behind **spent** resources (i.e., job pods that have completed)
- The Crash Loop Backoff behaviour that exists in the Kubernetes cluster can slow the Device Server startup in case of multiple dependencies
- An environment might end up with **different versions** of the k8s resource declarations, at the choosing of the developers (i.e. developers can select the version they like)

# Extending k8s - The Operator pattern

- *"The operator pattern aims to capture the key aim of a human operator who is managing a service or set of services. Human operators who look after specific applications and services have deep knowledge of how the system ought to behave, how to deploy it, and how to react if there are problems"*
- Extends the **Control Plane** to give custom **behaviours**
- Uses **Custom Resource Definitions** (extending the API) to create new manageable resources
- Uses the **control loop** pattern (a non-terminating loop that regulates the state of a system) to reconcile the resources to their desired state

Ref: [kubernetes.io/docs/concepts/extend-kubernetes/operator](https://kubernetes.io/docs/concepts/extend-kubernetes/operator)

# The SKA TANGO Operator

- Kubernetes Operator capable of managing TANGO resources (DeviceServer and DatabaseDS), controlling their lifecycle within the Kubernetes' native control/event loop
- Allows for a **lighter** deployment, due to avoiding init-containers for dependency resolution
- Optimised startup time for Device Servers, as the operator can directly tap into the TANGO environment and retrieve information on dependent devices and the TANGO **host** itself

# Why introduce an Operator ?

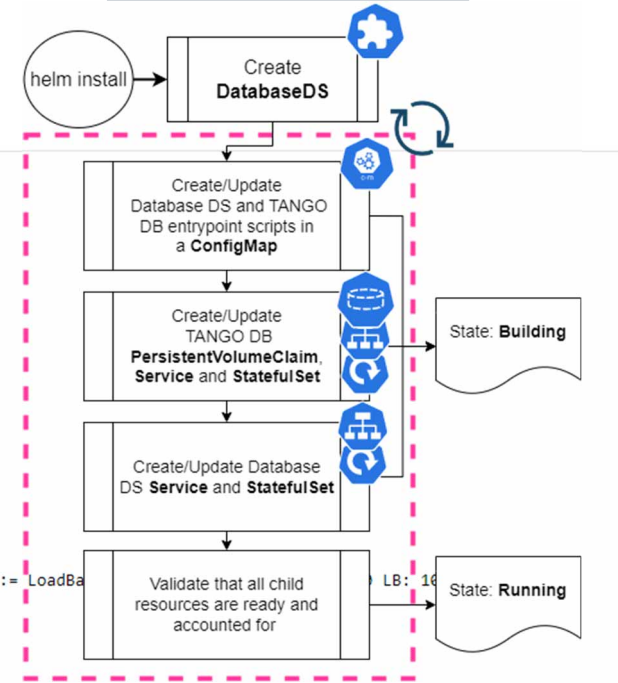

- Developers know Device Servers, not StatefulSets and Pods (or any k8s resource for that matter), as those are components relevant to the **platform** and not the **application**
- Making TANGO components **first-class** citizens of Kubernetes
- Automating much of the tasks a human would do to operate a TANGO resource, quicker and more reliably
- Allows to collect custom metrics on the CRDs, giving information on the system in a TANGO domain instead of a k8s domain (i.e., Device servers instead of StatefulSets)

# Custom Resource Definition (CRD)

## *databaseds.tango.tango-controls.org*

- TANGO DB StatefulSet, Service and PersistentVolume
- Database DS StatefulSet and Service

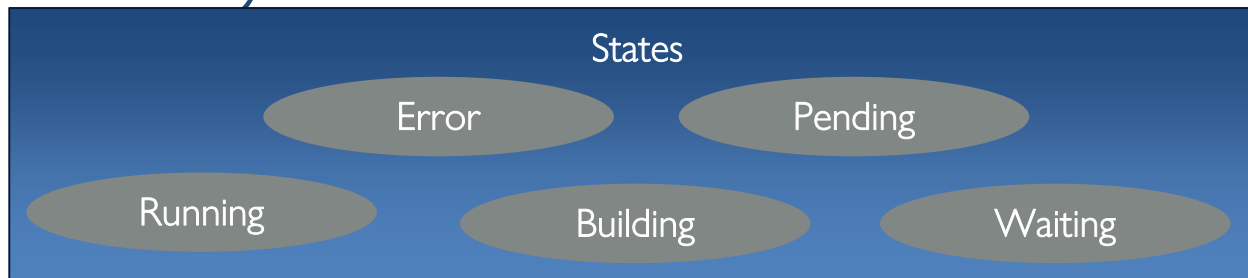
```
1 > kubectl -n ci-ska-skampi-3649827723-mid describe databaseds databaseds-tango-base
2 -----
3
4 API Version: tango.tango-controls.org/v1
5 Kind: DatabaseDS
6 Metadata:
7   Creation Timestamp: 2023-01-24T09:26:16Z
8 Spec:
9   Cluster Domain: cluster.local
10  Enable Load Balancer: true
11  Image Pull Policy: IfNotPresent
12  Orbport: 10000
13  Tango DB Storage Class: nfss1
14 Status:
15   Ds: 1 of 1
16   Replicas: 2
17   Resources: databaseds-tangodb-databaseds-tango-base := ClusterIP/10.98.139.14, mysql/3306 - databaseds-tango-base := LoadBalancingService
18   State: Running
19   Succeeded: 2
20   Tangodb: 1 of 1
```





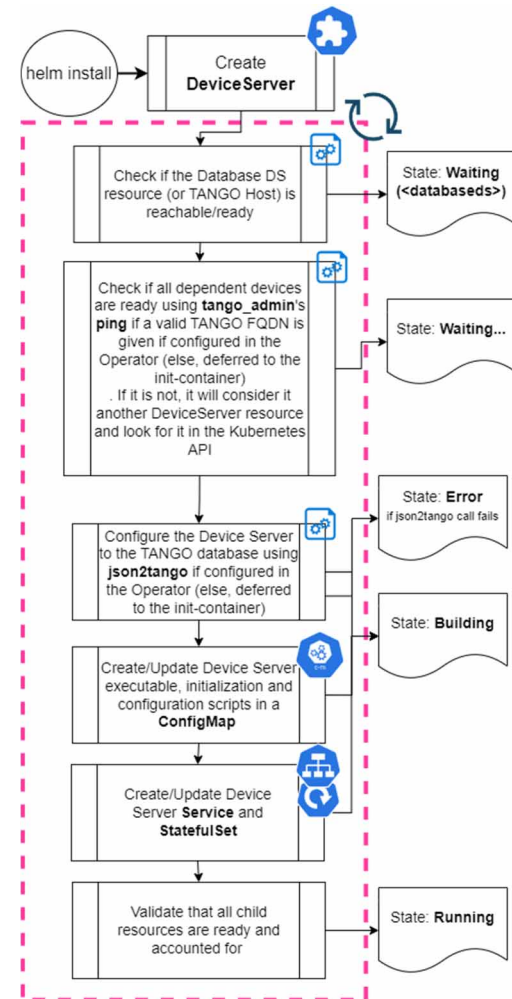
# Custom Resource Definition (CRD): *deviceservers.tango.tango-controls.org*

- Device Server StatefulSet and Service and Configurations (DS script used to run the device)

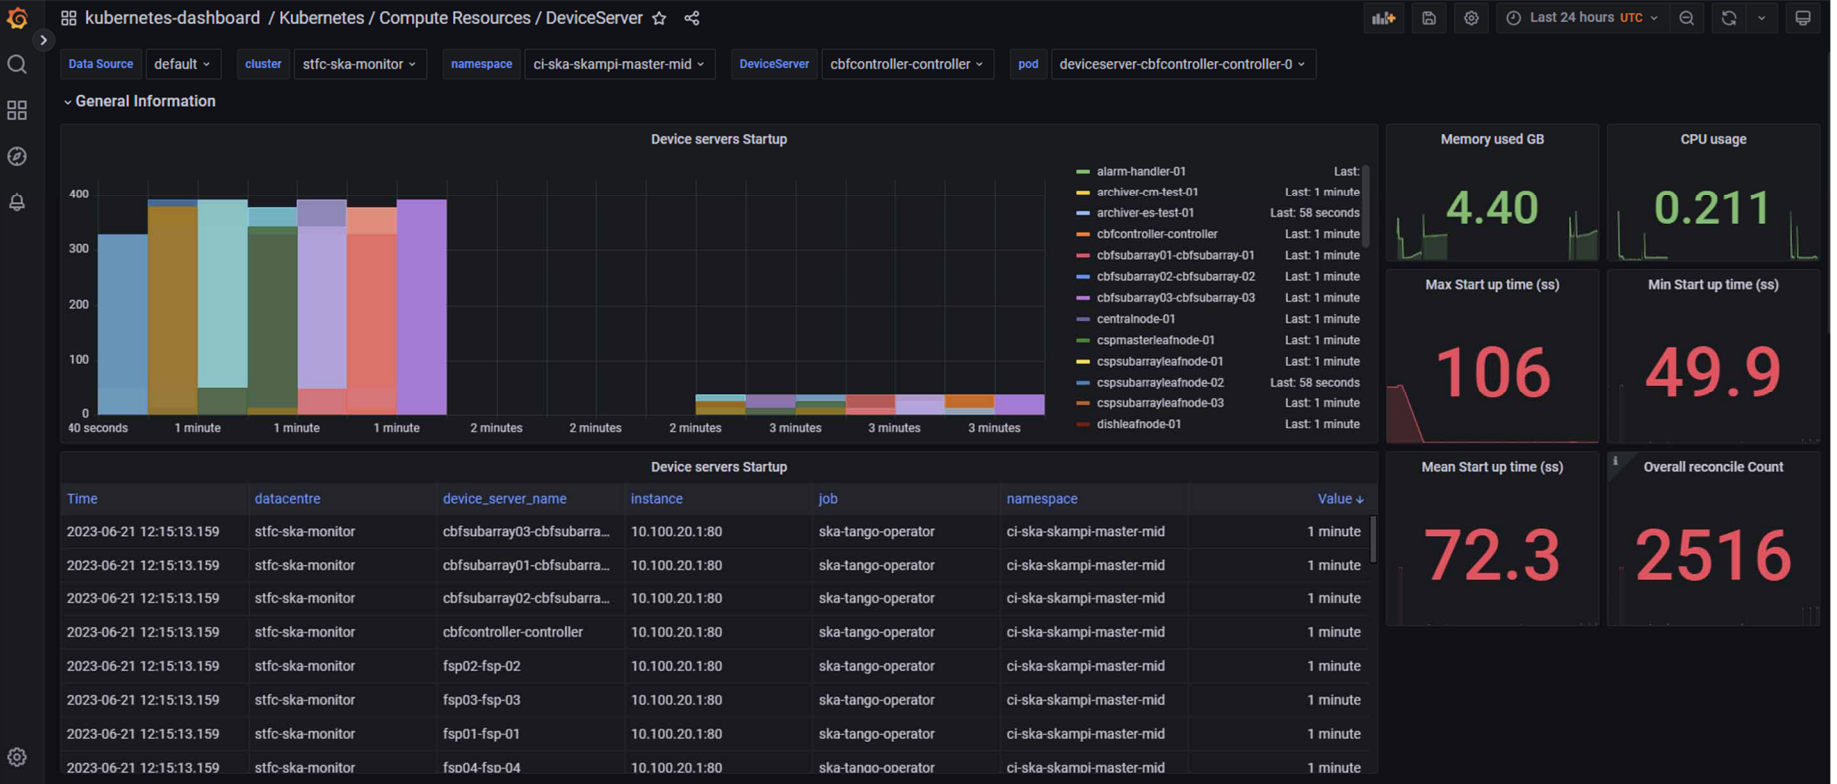


```

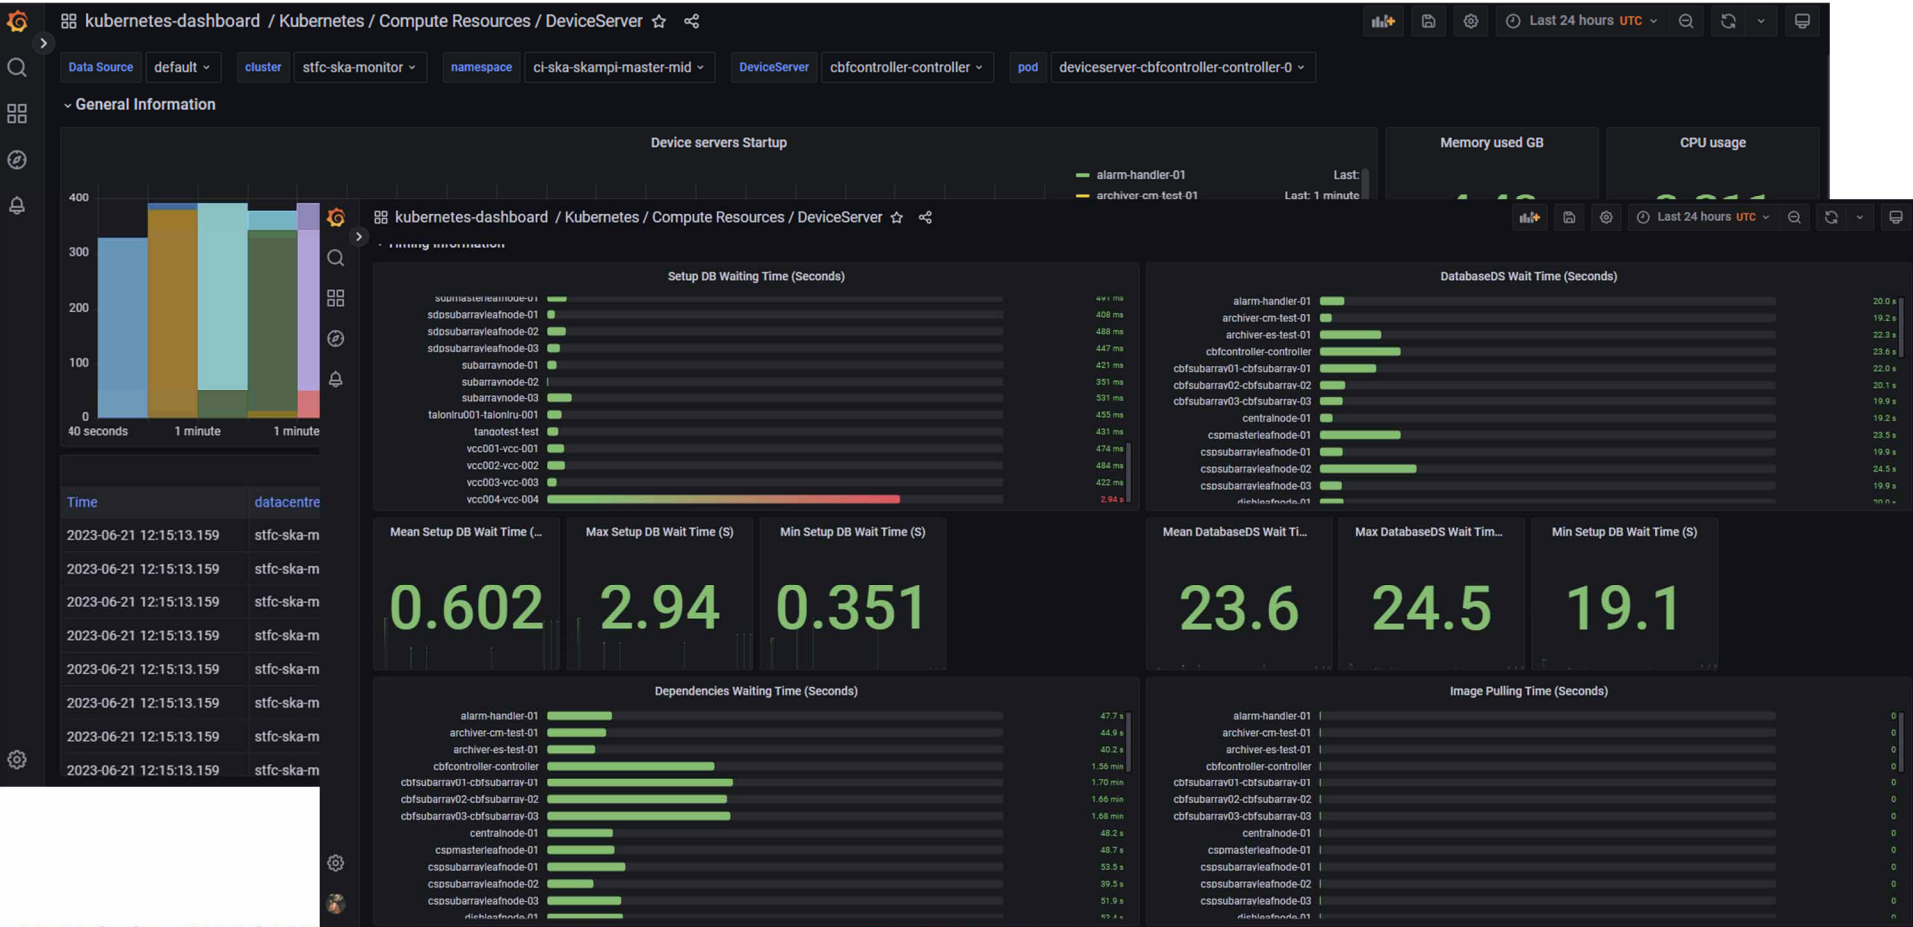
1 > kubectl get deviceserver --all-namespaces
2
3 -----
4 NAMESPACE                NAME                                COMPONENTS  SUCCEEDED  AGE    STATE
5 ci-ska-skampi-3650792176-mid  archiver-cm-archiver-test-01      2           2           2m30s  Running
6 ci-ska-skampi-3650792176-mid  archiver-es-archiver-test-01      2           2           2m30s  Running
7 ci-ska-skampi-3650792176-mid  cbfcontroller-controller          2           2           2m30s  Running
8 ci-ska-skampi-3650792176-mid  cbfsubarray01-cbfsubarray-01      2           2           2m30s  Running
9 ci-ska-skampi-3650792176-mid  cbfsubarray02-cbfsubarray-02      2           2           2m30s  Running
10 ...
  
```



# Grafana support for the Operator metrics



# Grafana support for the Operator metrics



# Summary

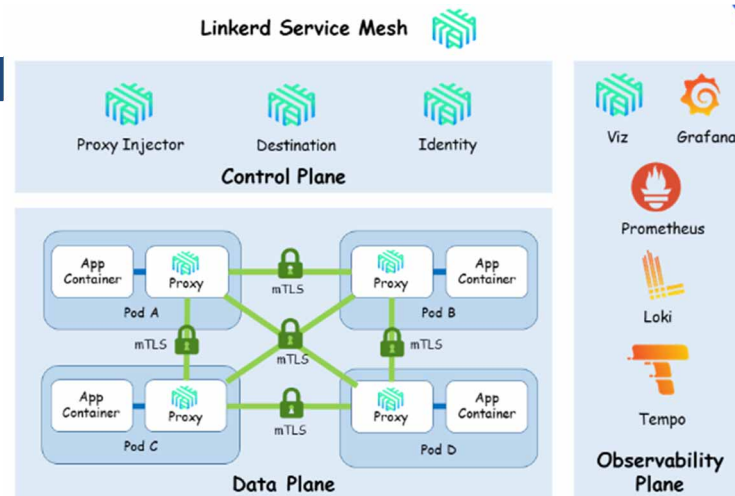
- Kubernetes managed device servers!
- Optimized initialization times and startup reliability
- Configuration delegated to a specialized service
- Less resources usage (memory, cpu, storage and even network)
- Better metrics (and dashboard) to evaluate our device servers
- Abstraction of TANGO domain-specific components from the platform ones
- And... Security (almost) for free!

# Securing Tango - With Linkerd



linkerd

- Linkerd is a service mesh:
  - It provides a lightweight proxy that is attached as a SideCar to Pods
  - It captures in and outbound traffic
  - The traffic is encrypted/decrypted
  - Policy controlled
  - Traffic is monitored and logged



# Securing Tango - Linkerd, a Universal Mesh

- Linkerd can be used in Cluster and between Clusters
- This approach enables the end to end security of the Tango Controls hierarchy - all without modification to the underlying application
- It can also be used for non-Tango traffic (any TCP), wherever it is challenging to retrofit mTLS



Thanks for you attention!



[www.oa-abruzzo.inaf.it](http://www.oa-abruzzo.inaf.it)

[matteo.dicarlo@inaf.it](mailto:matteo.dicarlo@inaf.it)

