Harnessing the power of Ubuntu and MicroK8s:

Case studies from
Center for Free Electron Lasers (CFEL). Germany
San Diego Supercomputer Center (SDSC). U.S.

Igor Khokhriakov
Ubuntu Summit 2023 | Riga, Latvia
Nov 3 – 5, 2023

LET ME INTRODUCE MYSELF...

- 15+ yrs in Software Development
- 10+ yrs in Scientific Software Development
- 3+ yrs of contribution into a kernel of an open source SCADA framework
- 1+ yrs in strategic committee for that framework
- 500+ citations
- Event-driven systems architect
- Reactive programming advocate

BIG RESEARCH FACILITIES

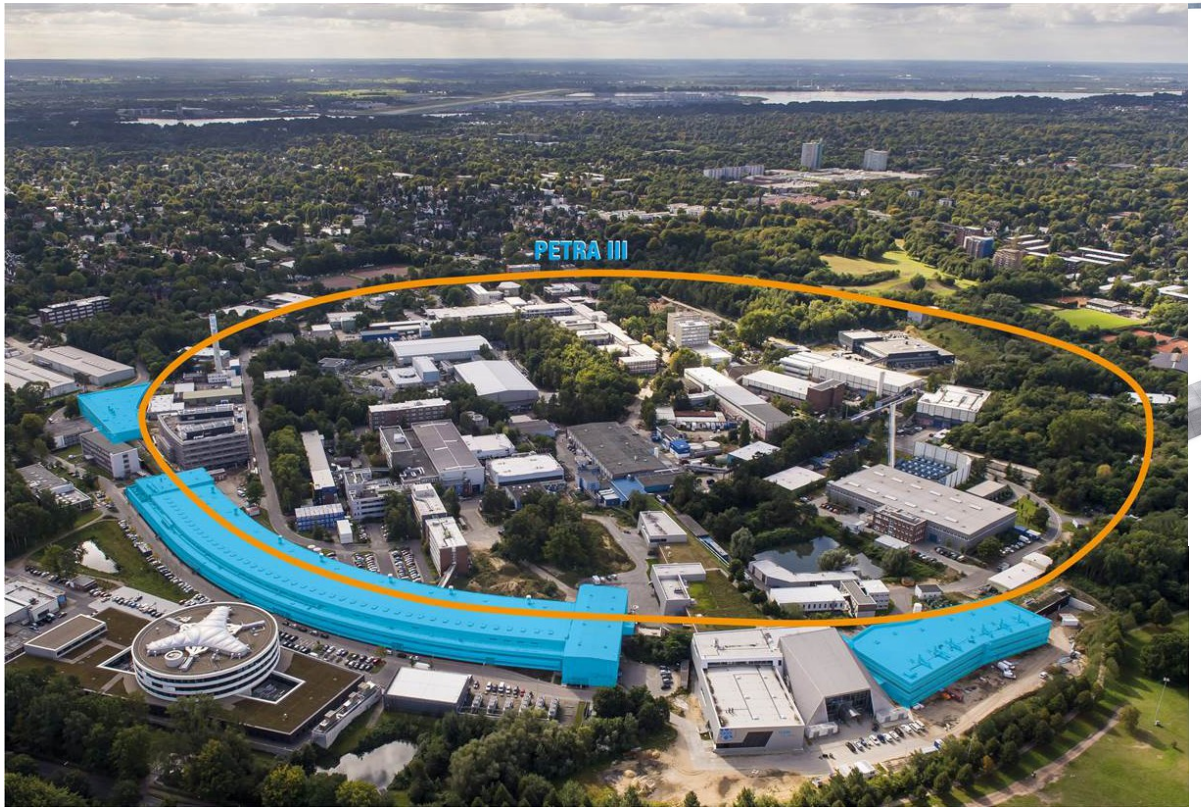

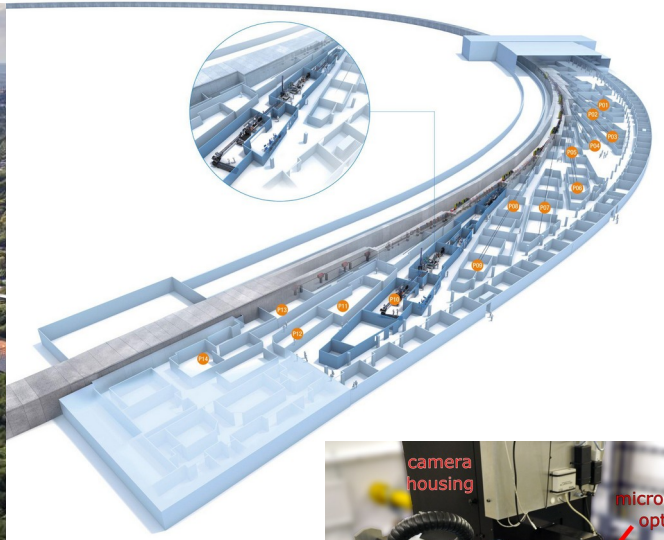

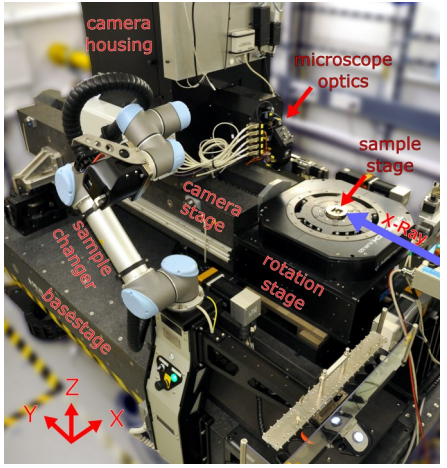

Do not try to read this :) Ain't important

WHAT THE HELL IS SCADA

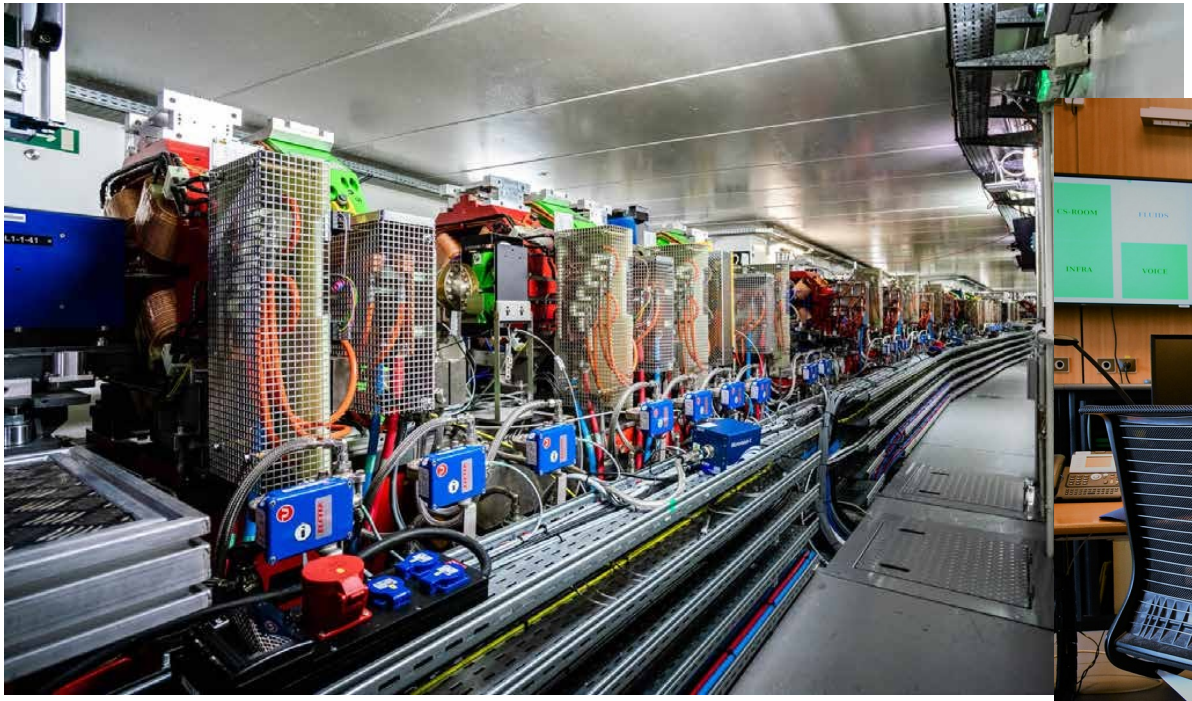

*ESRF control room

CASE I: SLOW CONTROLS SYSTEM FOR A SPECTROMETER AT CFEL, GERMANY

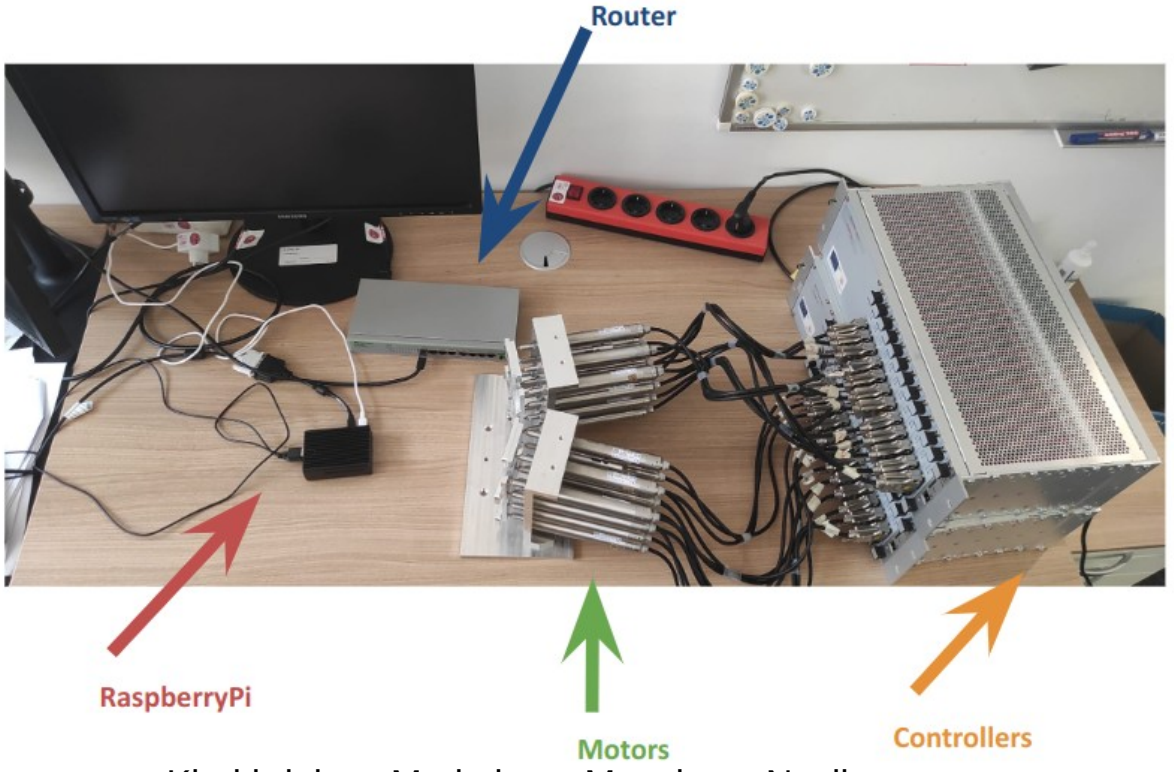

JOURNAL OF SYNCHROTRON RADIATION

Khokhriakov, Merkulova, Mazalova, Nozik.

“A novel solution for controlling hardware components of accelerators and beamlines”

Volume 29 | Part 3 | May 2022 | | 10.1107/S1600577522002685


DESY.
Sweden
Finland
Denmark
United Kingdom
Ireland
Belarus
Poland
Germany
Ukraine
PETRA III
PAL
POHANG ACCELERATOR LABORATORY
Kyrgyzstan
Pakistan
Mongolia
South Korea
Japan
Sea of Japan
East China Sea


MicroK8s

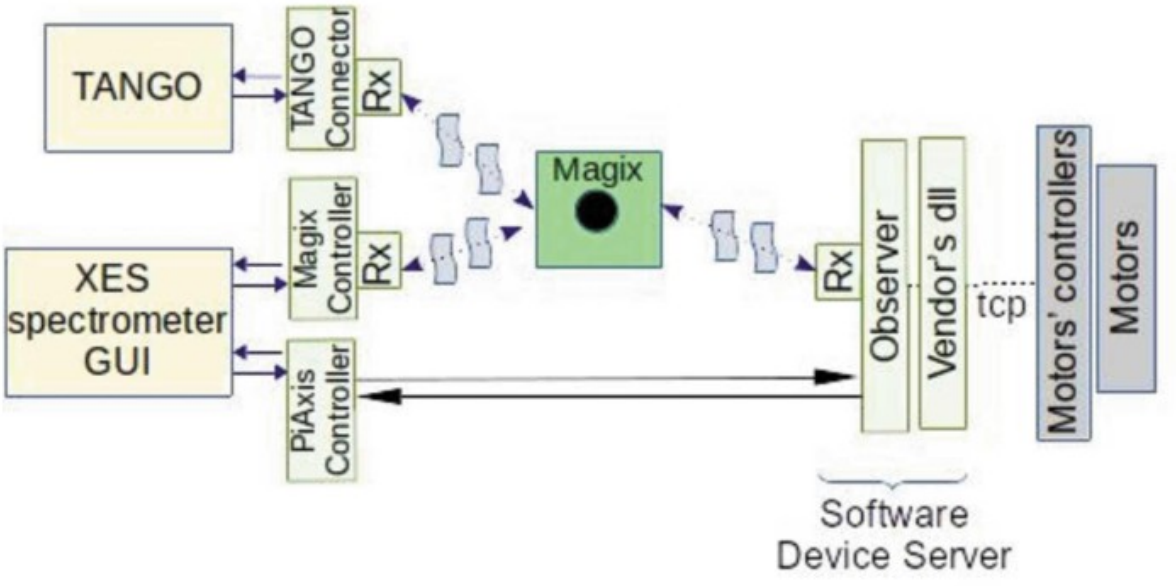
TANGO
TANGO Connector
Rx
Magix
Magix Controller
Rx
XES spectrometer GUI
PiAxis Controller
Rx
Observer
Vendor's dll
tcp
Motors' controllers
Motors
Software Device Server

DEPLOYED TO MICROK8S

```
ubuntu@luigi-etl-dev-igor-o3566465:~/axsis-kube$ microk8s kubectl get all --all-namespaces
NAMESPACE     NAME                                                 READY   STATUS    RESTARTS   AGE
kube-system   pod/coredns-64c6478b6c-9tcwx                         1/1     Running   0          5d2h
kube-system   pod/calico-node-29zzq                                1/1     Running   0          5d2h
kube-system   pod/calico-kube-controllers-76cc49f584-tdgbr         1/1     Running   0          5d2h
kube-system   pod/kubernetes-dashboard-585bdb5648-rm6lb            1/1     Running   0          5d2h
kube-system   pod/hostpath-provisioner-7764447d7c-4z2h5            1/1     Running   0          5d2h
kube-system   pod/dashboard-metrics-scraper-69d9497b54-8r55q       1/1     Running   0          5d2h
kube-system   pod/metrics-server-679c5f986d-r89nk                  1/1     Running   0          5d2h
ingress       pod/nginx-ingress-microk8s-controller-zf546          1/1     Running   0          4d4h
default       pod/magix-6f97865b8f-mn7zp                           1/1     Running   0          72m
default       pod/axsis-8547fc55d5-dmhcl                           1/1     Running   0          62m
default       pod/web-65bf5cf4c4-k2kb4                             1/1     Running   0          57m
default       pod/axsis.magix-6fb8647ff9-gm8xm                     1/1     Running   0          55m
default       pod/tango-host-7d95f46d96-8lvpf                      1/1     Running   0          37m
default       pod/tango-server-748fdb644f-xpvtp                    1/1     Running   0          36m

NAMESPACE     NAME                                  TYPE        CLUSTER-IP       EXTERNAL-IP   PORT(S)                  AGE
default       service/kubernetes                    ClusterIP   10.152.183.1     <none>        443/TCP                  5d2h
kube-system   service/kube-dns                      ClusterIP   10.152.183.10    <none>        53/UDP,53/TCP,9153/TCP   5d2h
kube-system   service/metrics-server                ClusterIP   10.152.183.241   <none>        443/TCP                  5d2h
kube-system   service/kubernetes-dashboard          ClusterIP   10.152.183.55    <none>        443/TCP                  5d2h
kube-system   service/dashboard-metrics-scraper     ClusterIP   10.152.183.71    <none>        8000/TCP                 5d2h
default       service/magix                         NodePort    10.152.183.132   <none>        8080:32407/TCP           72m
default       service/axsis                         NodePort    10.152.183.60    <none>        80:30493/TCP             62m
default       service/web                           NodePort    10.152.183.70    <none>        80:32600/TCP             57m
default       service/pi-controller-1               ClusterIP   None             <none>        50000/TCP                47m
default       service/pi-controller-2               ClusterIP   None             <none>        50000/TCP                47m

NAMESPACE     NAME                                        DESIRED   CURRENT   READY   UP-TO-DATE   AVAILABLE   NODE SELECTOR
kube-system   daemonset.apps/calico-node                  1         1         1       1            1           kubernetes.io/
ingress       daemonset.apps/nginx-ingress-microk8s-controller   1         1         1       1            1           <none>

NAMESPACE     NAME                                       READY   UP-TO-DATE   AVAILABLE   AGE
kube-system   deployment.apps/calico-kube-controllers    1/1     1            1           5d2h
kube-system   deployment.apps/coredns                    1/1     1            1           5d2h
kube-system   deployment.apps/kubernetes-dashboard       1/1     1            1           5d2h
kube-system   deployment.apps/hostpath-provisioner       1/1     1            1           5d2h
kube-system   deployment.apps/dashboard-metrics-scraper  1/1     1            1           5d2h
kube-system   deployment.apps/metrics-server             1/1     1            1           5d2h
default       deployment.apps/magix                      1/1     1            1           72m
default       deployment.apps/axsis                      1/1     1            1           62m
default       deployment.apps/web                        1/1     1            1           57m
default       deployment.apps/axsis.magix                1/1     1            1           55m
default       deployment.apps/tango-host                 1/1     1            1           37m
default       deployment.apps/tango-server               1/1     1            1           37m

NAMESPACE     NAME                                                 DESIRED   CURRENT   READY   AGE
kube-system   replicaset.apps/calico-kube-controllers-76cc49f584   1         1         1       5d2h
kube-system   replicaset.apps/coredns-64c6478b6c                   1         1         1       5d2h
kube-system   replicaset.apps/kubernetes-dashboard-585bdb5648      1         1         1       5d2h
kube-system   replicaset.apps/hostpath-provisioner-7764447d7c      1         1         1       5d2h
kube-system   replicaset.apps/dashboard-metrics-scraper-69d9497b54   1         1         1       5d2h
kube-system   replicaset.apps/metrics-server-679c5f986d            1         1         1       5d2h
default       replicaset.apps/magix-6f97865b8f                     1         1         1       72m
default       replicaset.apps/axsis-8547fc55d5                     1         1         1       62m
default       replicaset.apps/web-65bf5cf4c4                       1         1         1       57m
default       replicaset.apps/axsis.magix-6fb8647ff9               1         1         1       55m
default       replicaset.apps/tango-host-7d95f46d96                1         1         1       37m
default       replicaset.apps/tango-server-748fdb644f              1         1         1       37m
ubuntu@luigi-etl-dev-igor-o3566465:~/axsis-kube$
```

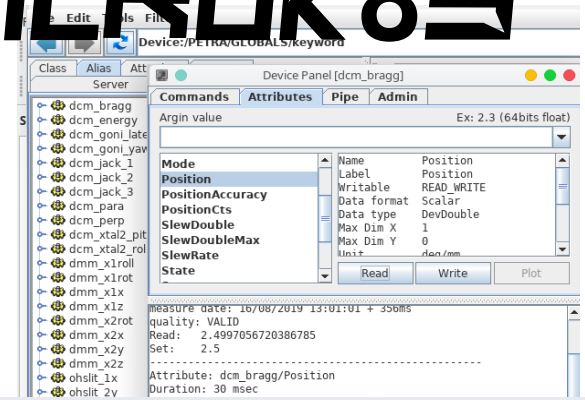

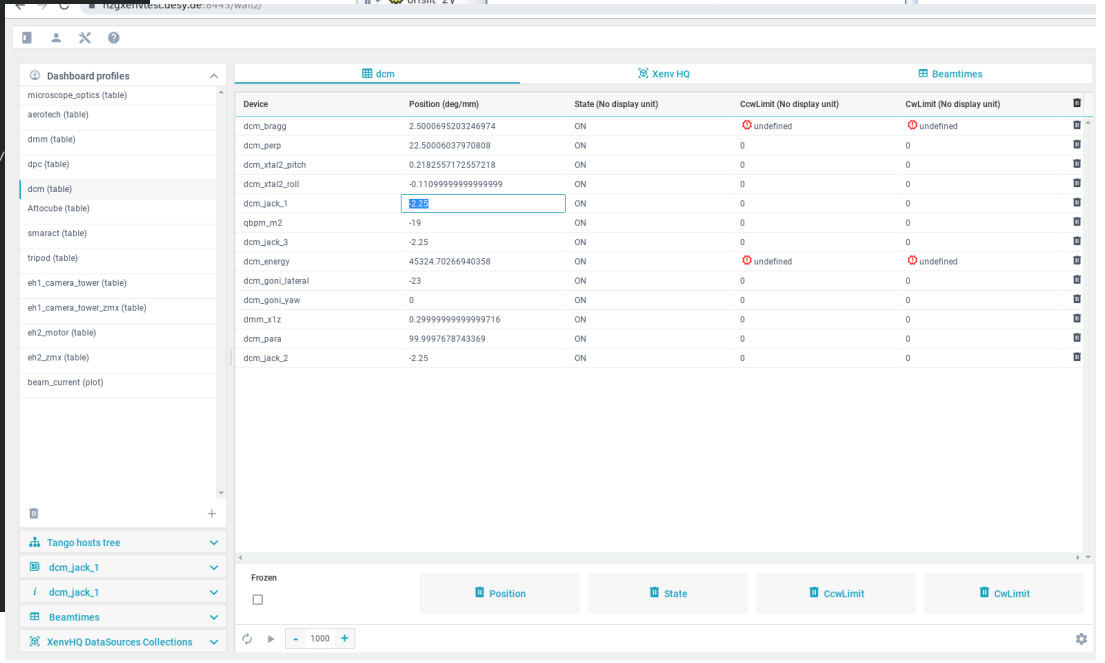

Also helps scientists to fastern in software development

- Easy to deploy and develop on a local machine
- Native CI/CD integrations
- Embrace serverless design
- Easy to share (*Helm*)
- Opens window to integrate 3rd party open source tools e.g. observability, service meshes etc

Case 2: RCSB PDB search engine at San Diego Supercomputer Center, US

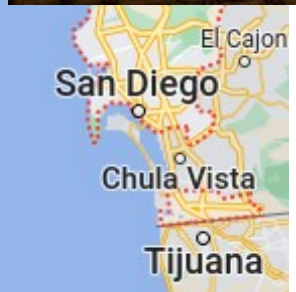

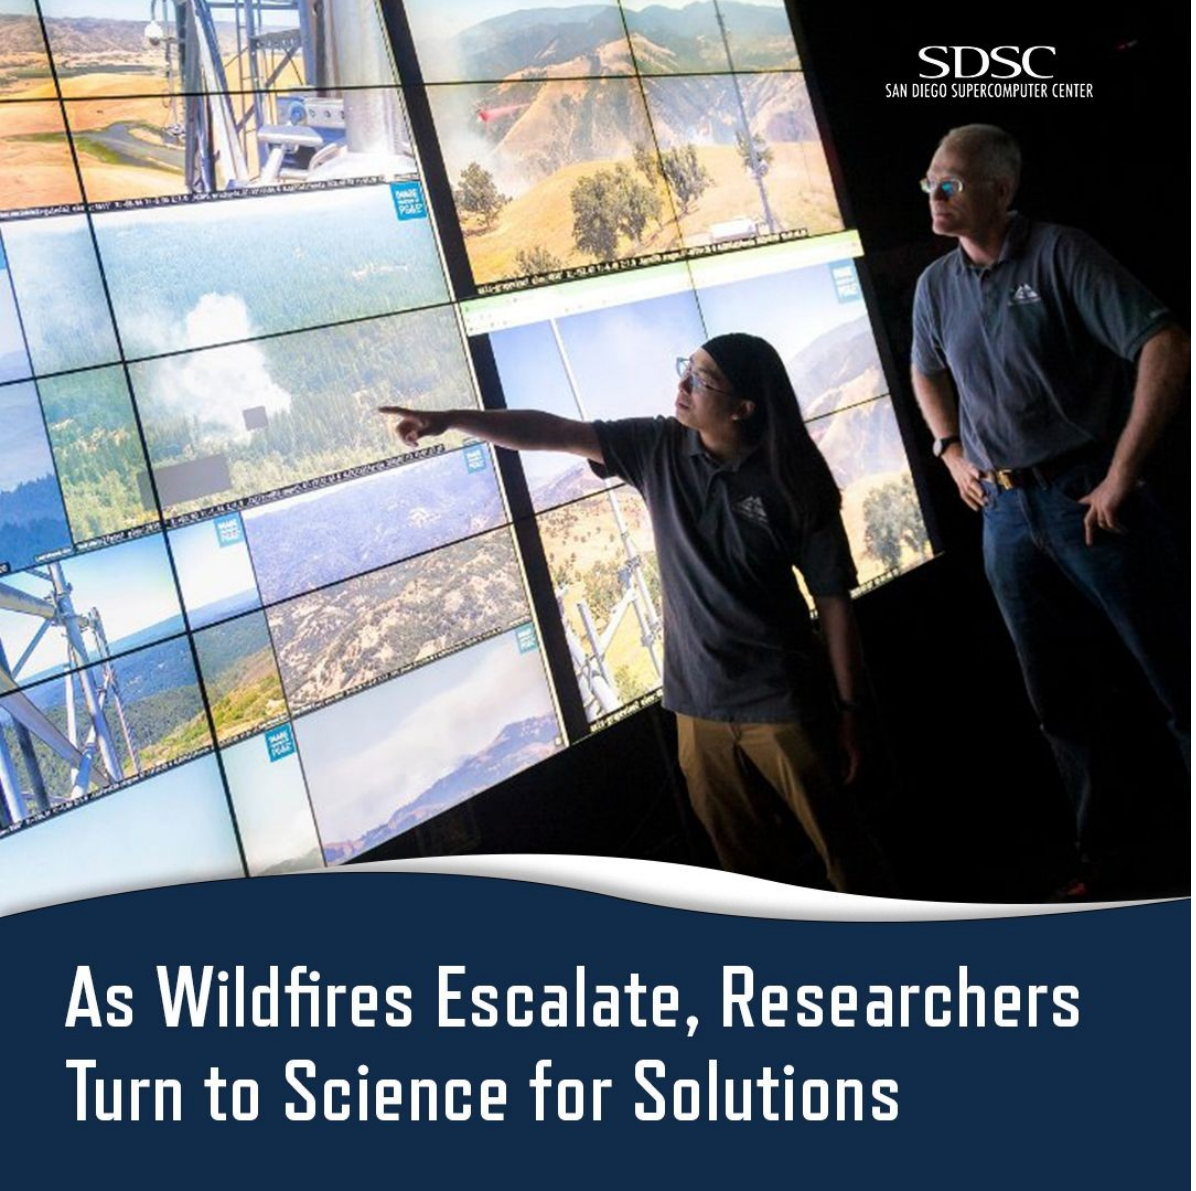

- Increased temperatures, weather extremes and hot and dry conditions have intensified the threat of wildfires in California and around the world.
- Wildfire is an extremely complex problem and research and mitigation have never been more important.

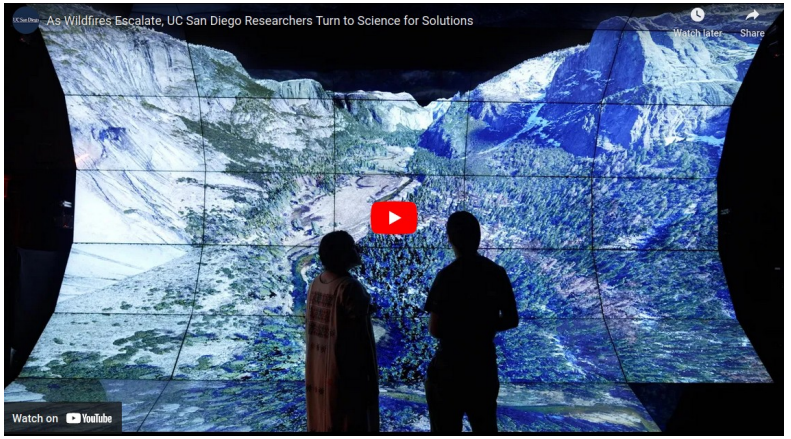

https://youtu.be/NE7_8zqnbR0

RCSB PDB PROTEIN DATA BANK

201,515 Structures from the PDB

1,068,577 Computed Structure Models (CSM)

3D Structures | Enter search term(s), Entry ID(s), or sequence | Include CSM

Advanced Search | Browse Annotations | Help

PDB-101 | wwPDB | EMDataResource | NUCLEIC ACID DATABASE | wwPDB Foundation

New: More Computed Structure Models (CSM) available Learn more

- Welcome
- Deposit
- Search
- Visualize
- Analyze
- Download
- Learn

RCSB Protein Data Bank (RCSB PDB) enables breakthroughs in science and education by providing access and tools for exploration, visualization, and analysis of:

- Experimentally-determined 3D structures from the **Protein Data Bank (PDB)** archive
- **Computed Structure Models (CSM)** from AlphaFold DB and ModelArchive

These data can be explored in context of external annotations providing a structural view of biology.

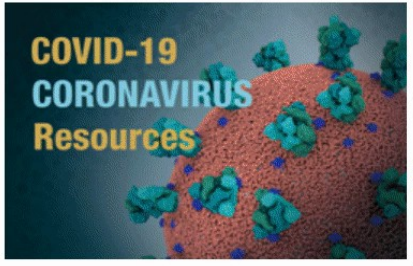

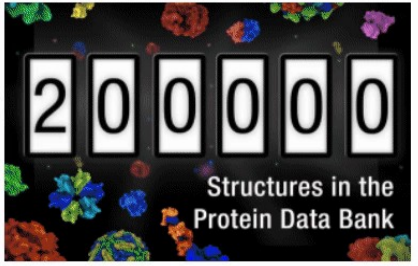

February **Molecule of the Month**

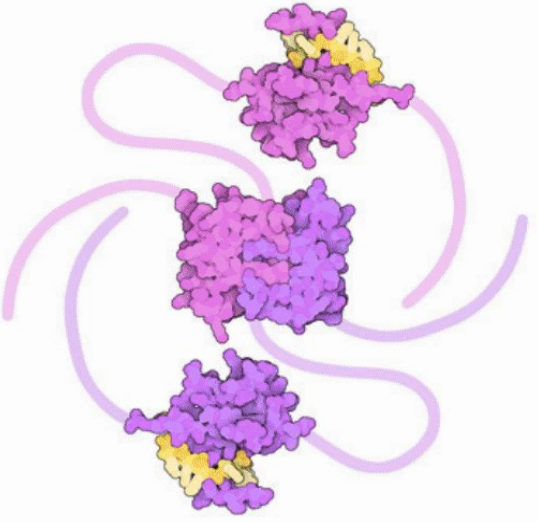

SARS-CoV-2 Nucleocapsid and Home Tests

RCSB PDB
PROTEIN DATA BANK

Data API - Counts and average duration of events over time [RCSB lb-web-dataapi]

Search API - Counts and average duration of events over time [RCSB lb-web-searchapi]

Data API - Counts of error events over time [RCSB lb-web-dataapi]

Search API - Counts of error events over time [RCSB lb-web-searchapi]

Data API - Traffic origin and stage [RCSB lb-web-dataapi]

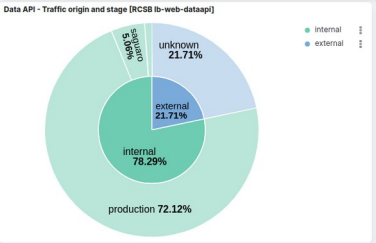

Search API - Internal VS external traffic [RCSB lb-web-searchapi]

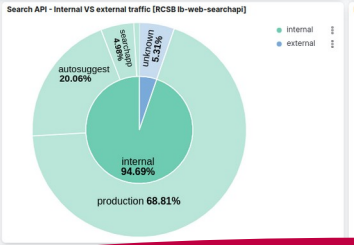

PDB vs CSM

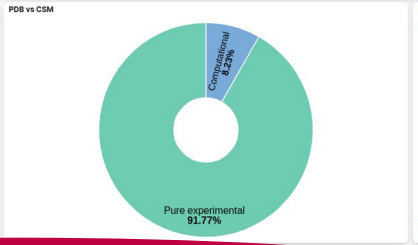

Search Services breakdown

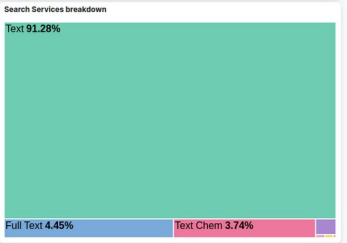

Data API - Counts of events metric [RCSB lb-web-dataapi]

Events
2,287,910
Error events: 2,419

Search API - Counts of events metric [RCSB lb-web-searchapi]

Events
1,130,766
Error events: 6,457

3,5M R/Day

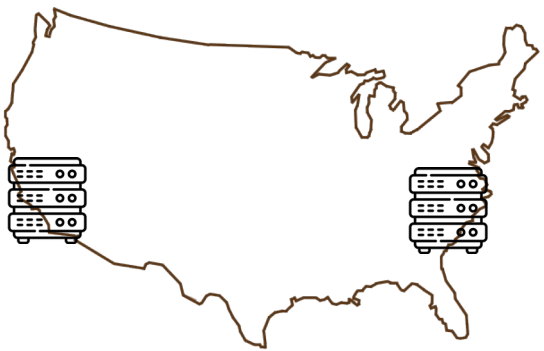

MY ROLE AT RCSB PDB

- **Pilot project of migration to K8s**
- **Designing and implementing a Kubernetes platform**
- **Leading the migration process from OpenStack to K8s**
- Designing and implementing a backend architecture based on Elastic Search to integrate AlphaFold computed structure models (over 200 million)
- Designing, integrating, and introducing various analytics metrics including APM, business KPIs, and domain-specific metrics

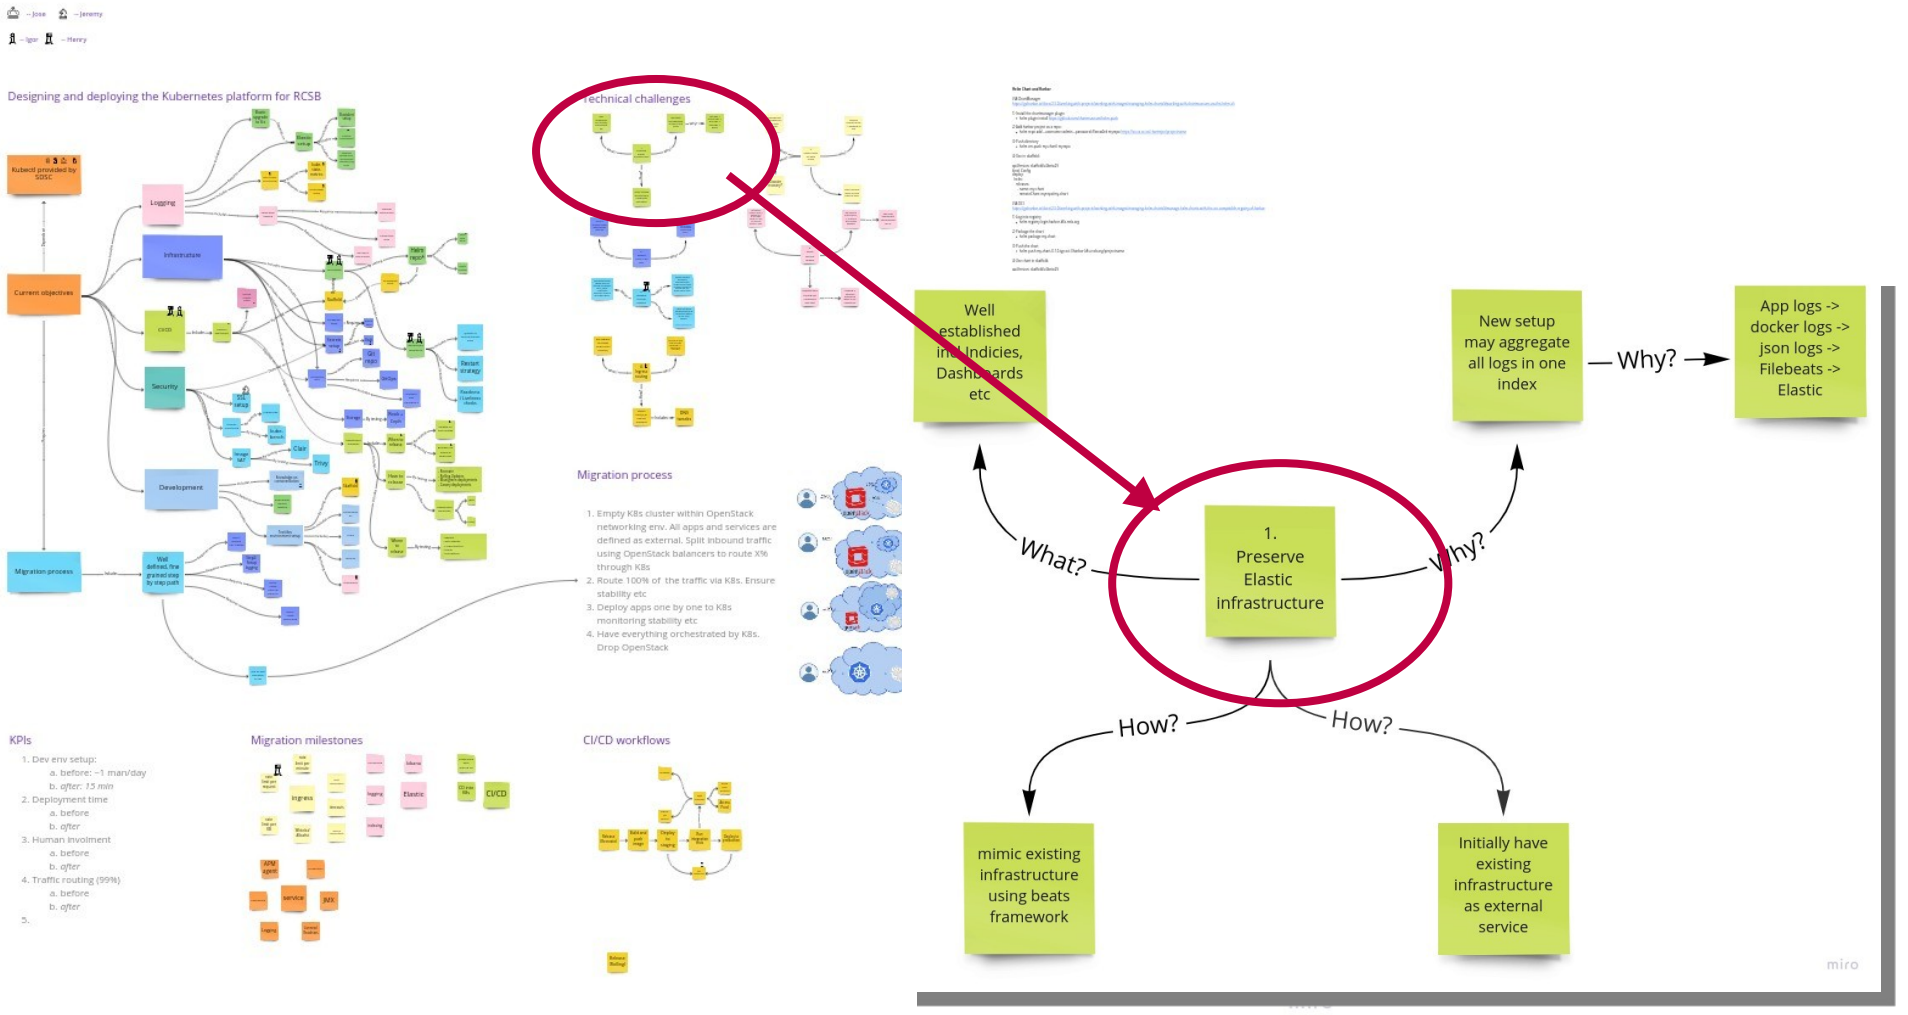

https://medium.com/@ingvord.mail/project-management-gamefication-c7cfdee3cefd

DOMAIN SPECIFIC ANALYTICS

RCSB PDB PROTEIN DATA BANK

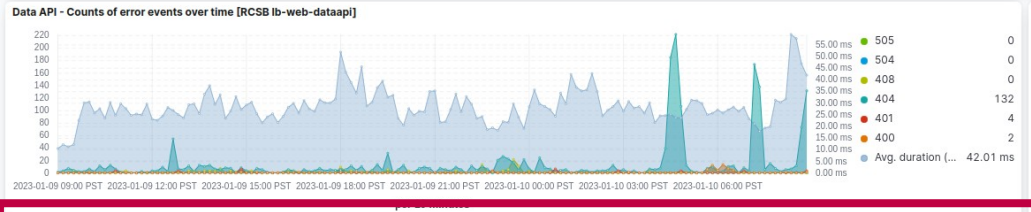

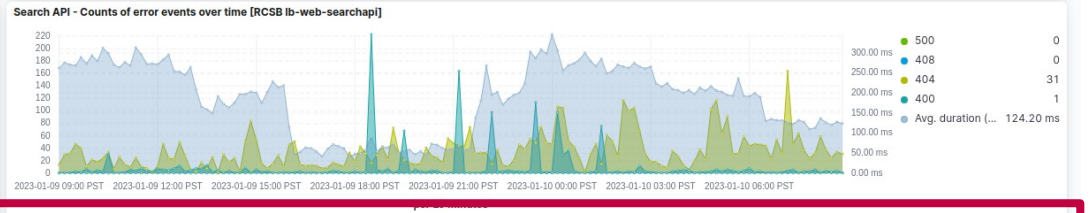

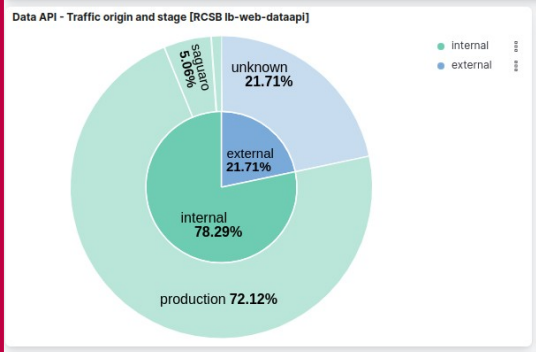

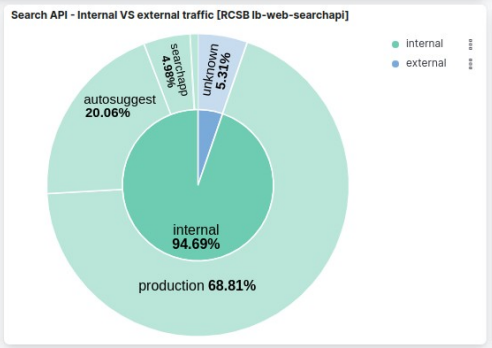

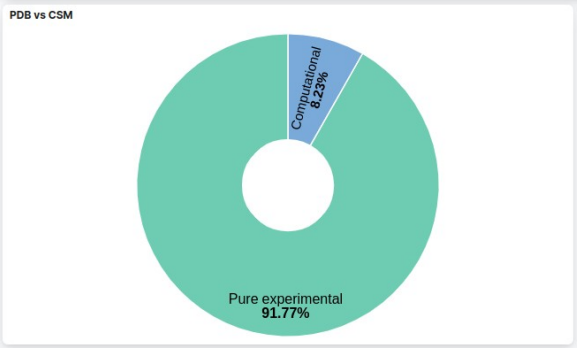

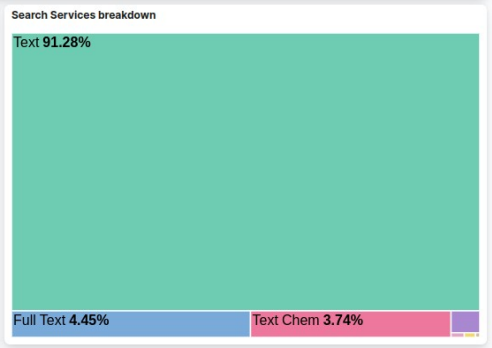

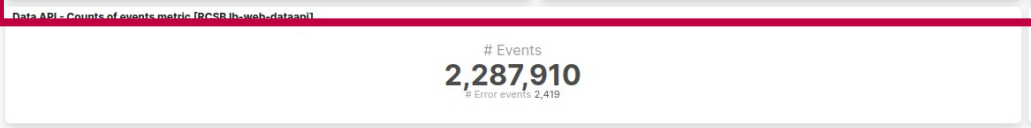

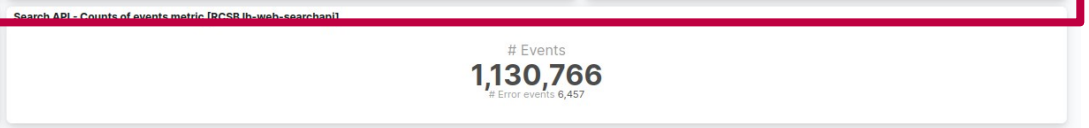

How MicroK8s saved RCSB $200K USD...

- Internal Analytics Engine based on HAProxy logs, regexp and filebeat Elastic ingestor
- Quite clear efforts estimation (1 yr) developer as per previous experience
- Essential for K8s migration
- Required for **fundraising** i.e. grant proposal

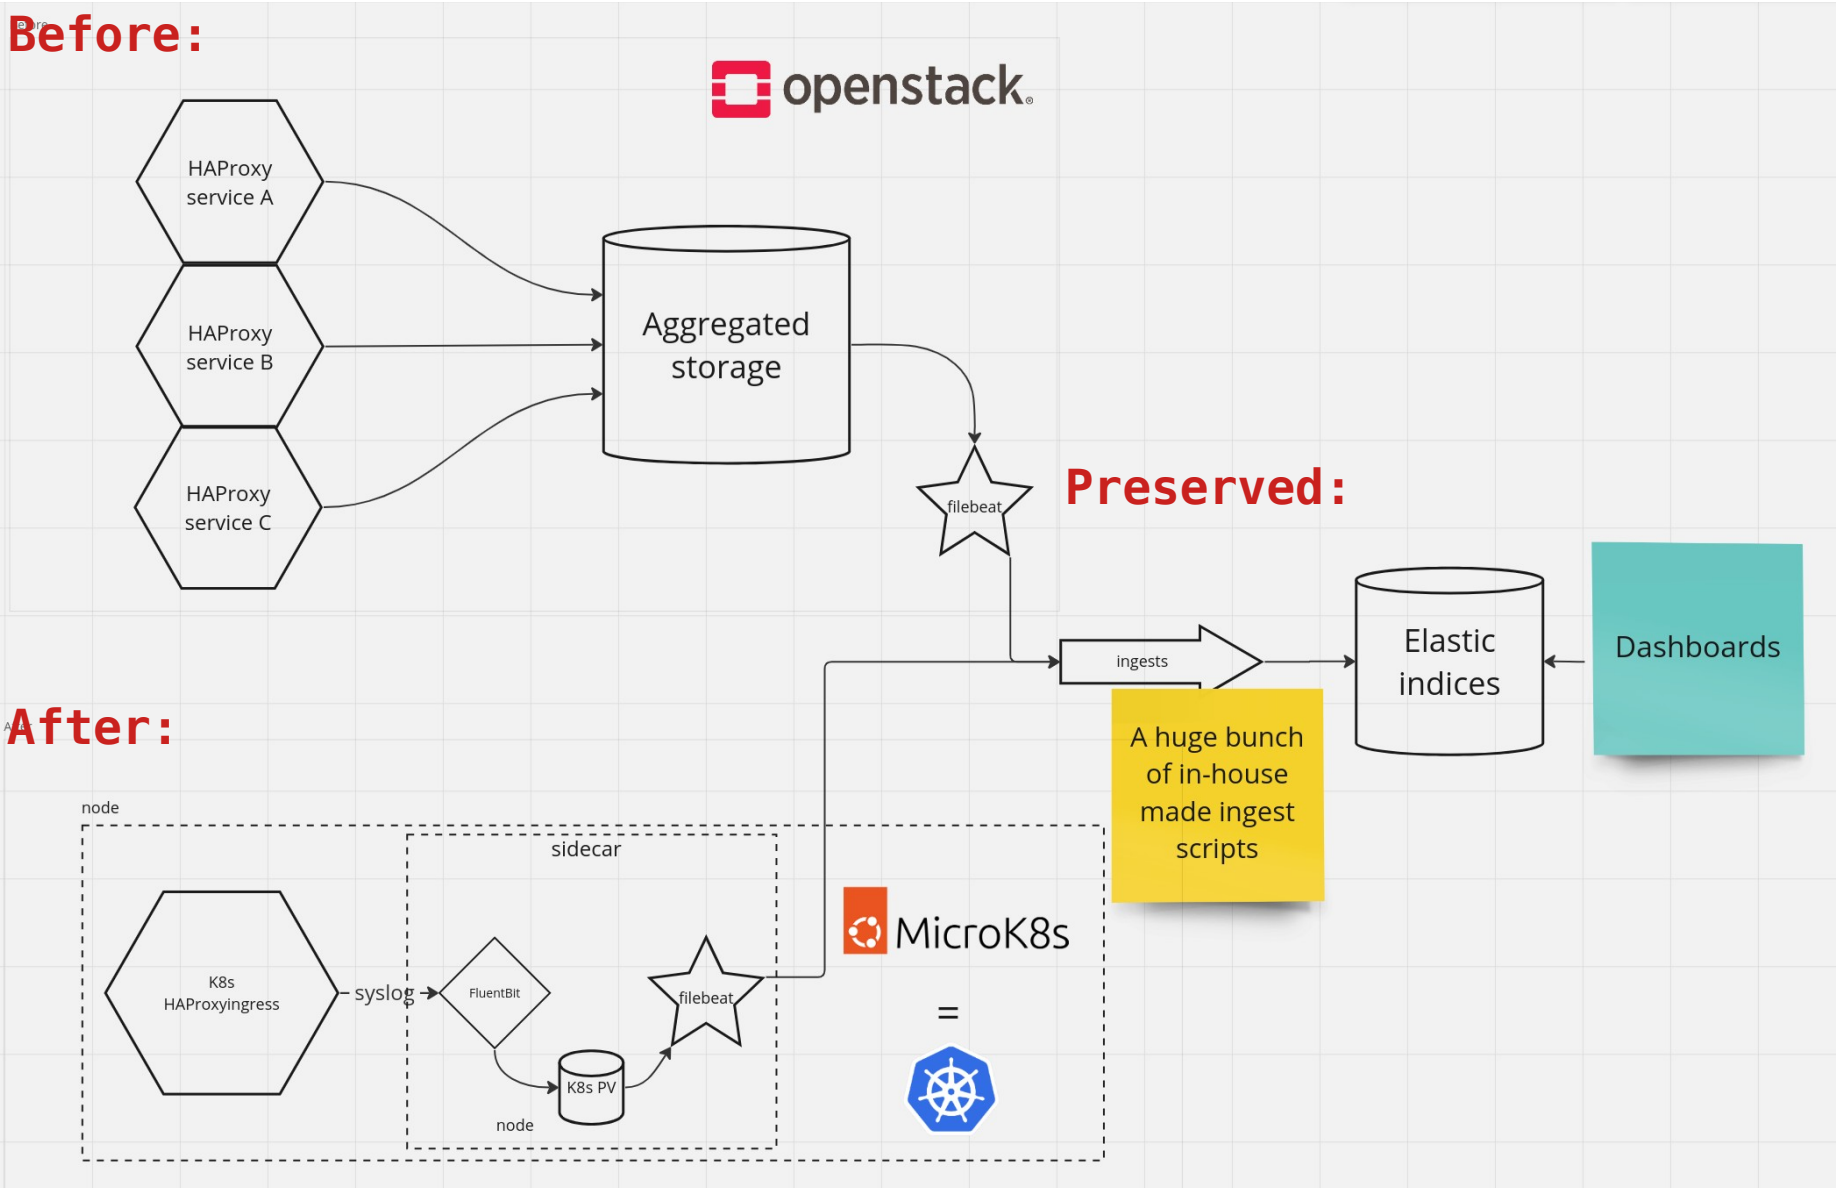
Before:
openstack
HAProxy service A
HAProxy service B
HAProxy service C
Aggregated storage
filebeat
Preserved:
ingests
Elastic indices
Dashboards
A huge bunch of in-house made ingest scripts
After:
node
sidecar
K8s HAProxyingress
syslog
FluentBit
filebeat
K8s PV
node
MicroK8s
=

BUT HOW... ?

- Sandbox in MicroK8s
  - Echo servers
  - Different ingress controllers
  - Ksniff for inter-pods traffic analysis
  - Non-production Elastic cluster (external at first)
- Allows very fine grained tuning
- Efficient problems isolation

MicroK8s

WHAT IS NEXT?

- K8s becomes de facto a standard solution for deploying applications
- MicroK8s provides great starting point and beyond
- I will invest my efforts into advocating MicroK8s
  - Workshops
  - Show cases
  - etc

Open
Source
.Science

LET'S CONTINUE THE CONVERSATION

LinkedIn

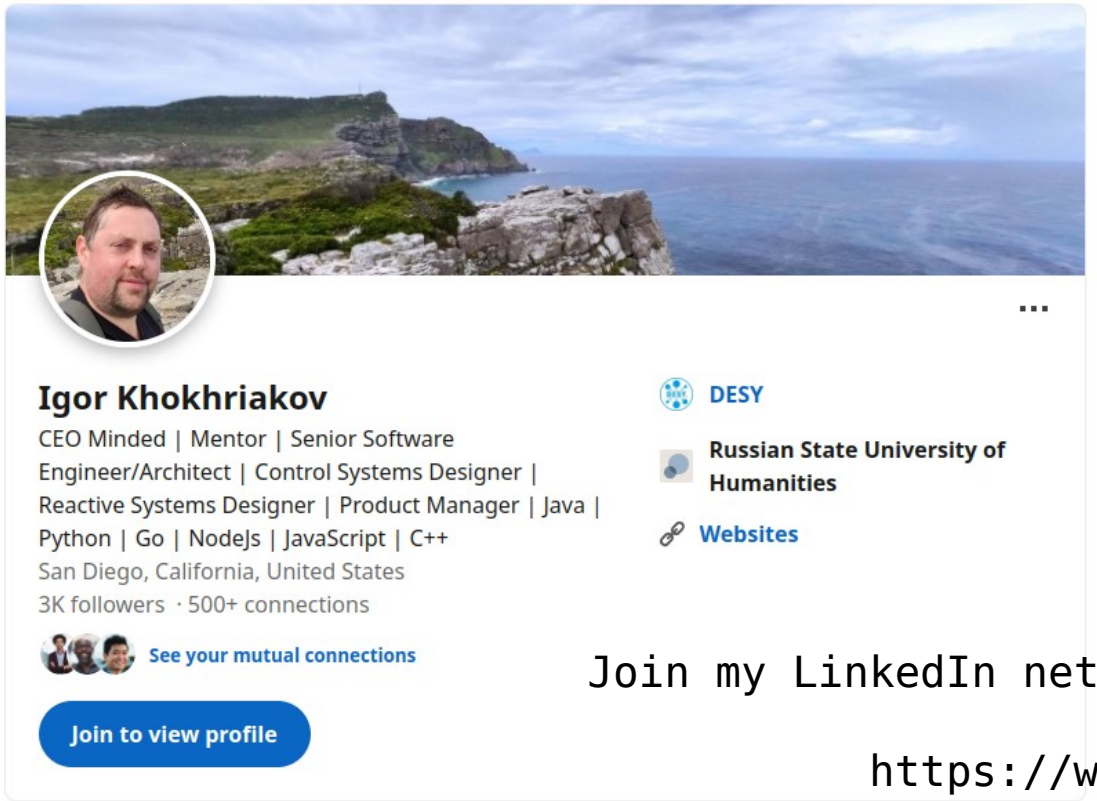

Join my LinkedIn network of 3,000+ professionals...

https://www.linkedin.com/in/ikhokhryakov

THANKS!

QUESTIONS?

Ingvord.mail@gmail.com

Igor.khokhriakov@desy.de