Kubernetes operators for web hosting at CERN

Hepix fall 2022

Alex Lossent

Outline

- Web hosting infrastructure:
  - Vision
  - Backbone
  - Skin
  - Heart
- Timeline & future plans

The Vision

Legacy webservices portal (since 2004!)

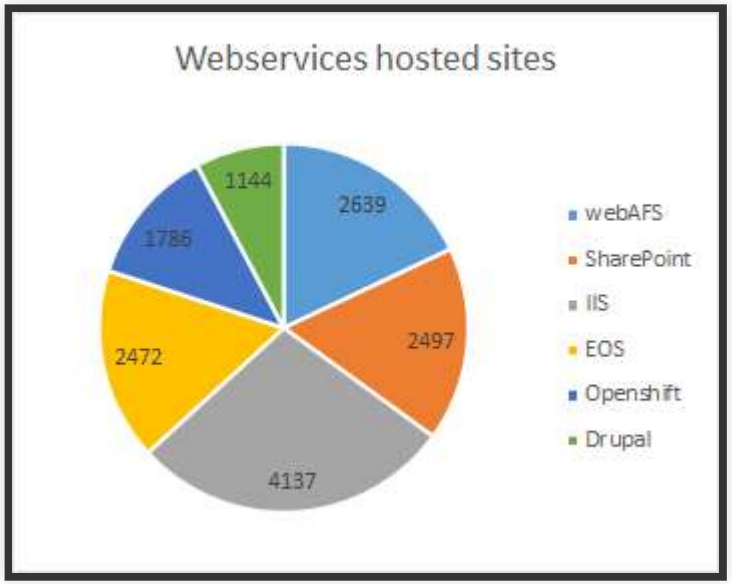

Strategy (2018)

“strategy for web hosting at CERN is to provide a range of solutions fulfilling the diverse needs of the community, from simple/static sites to content management systems (Drupal) to complex web applications (Platform-as-a-Service) using a shared infrastructure (Kubernetes/OpenShift)”

Why discontinue the old portal?

- MALT considerations:
  - webservices portal & site management logic (`ASP.NET`)
  - web site lifecycle (`FIM`)
  - Windows IIS web hosting (`DFS`)
  - replace SharePoint with Application Templates: Discourse, Wordpress…
- Move 10k+ sites/apps to new SSO
- 20-year old design choices and hardcoded behaviors (DNS domains…)

Plan new architecture (end 2019)

- Common container-based infrastructure for all types of site
- Kubernetes Operators to implement site management logic
- Integrate with new Application Portal for new SSO and lifecycle of user applications
- New web user portal front-end
    - design goal: provide more guidance to users

Scope

Use this new infra for:

- static/CGI web site hosting (webeos, gitlab pages)
- Drupal sites
- replacement of Openshift 3 PaaS
  - custom applications
  - application templates

The Backbone

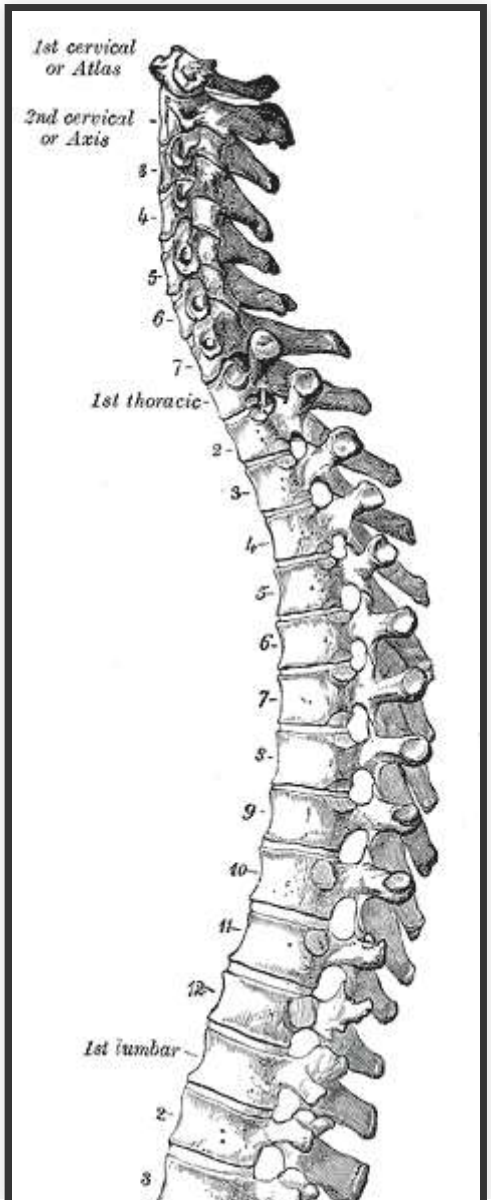

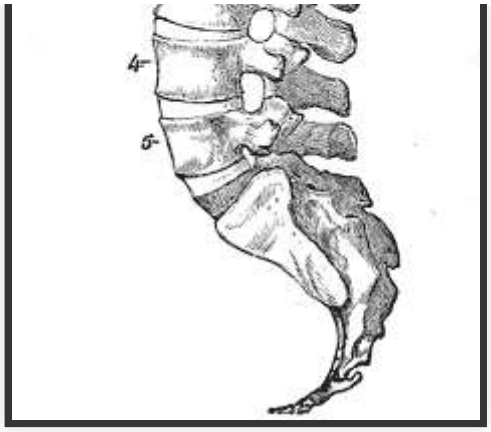

Public Domain

OKD4 infrastructure

- OKD4 Openshift Kubernetes Distribution
  - Adds multi-tenancy features to Kubernetes
  - PaaS: nice automation for custom apps, stronger security/isolation
- 4 production OKD4 clusters for the different use cases
  - Sharing a common design and infrastructure

Webeos (3700 projects)

Serve static/CGI web sites

- from user-provided EOS folders (e.g. `/eos/user/a/alossent/www`)
- documentation sites based on GitLab Pages

Drupal (850 projects)

Content Management System

- Official content: CERN homepage…
- Help site owners with Drupal version updates
- Support the large ecosystem of modules

PaaS (1250 projects)

Host users' custom web applications

- deploy from upstream or custom Docker images
- S2I: build applications from code (PHP, Python, NodeJS…)
- high level of integration with CERN computing environment: SSO, storage (EOS, CVMFS, CephFS), DNS, firewall, TN integration…
- very large applications out of scope: use a dedicated Kubernetes cluster (Openstack Magnum)

App-Catalogue (260 projects)

Provide a catalogue of self-service application templates

- Grafana, Sentry, Nexus…
- MALT: Wordpress, Discourse…

4.6

Common infrastructure

- One infra, multiple clusters
  - use cases: webeos, drupal, paas, app-catalogue
  - environments: prod, staging, dev, CI
  - decentralized: each cluster is 100% self-sufficient
- Clusters managed with gitops
  - ArgoCD, Helm charts
  - end-to-end tests for every change

4 . 7

CERN computing environment

Made OKD4 work at CERN using:

- OKD4 customization (Openstack support…)
- Extra upstream components: OPA, ExternalDNS, velero, restic…
- Shared components with Kubernetes team: EOS, CVMFS, CephFS…
- New development to integrate with LanDB, SSO…

4 . 8

The Skin

5 . 1

New webservices portal

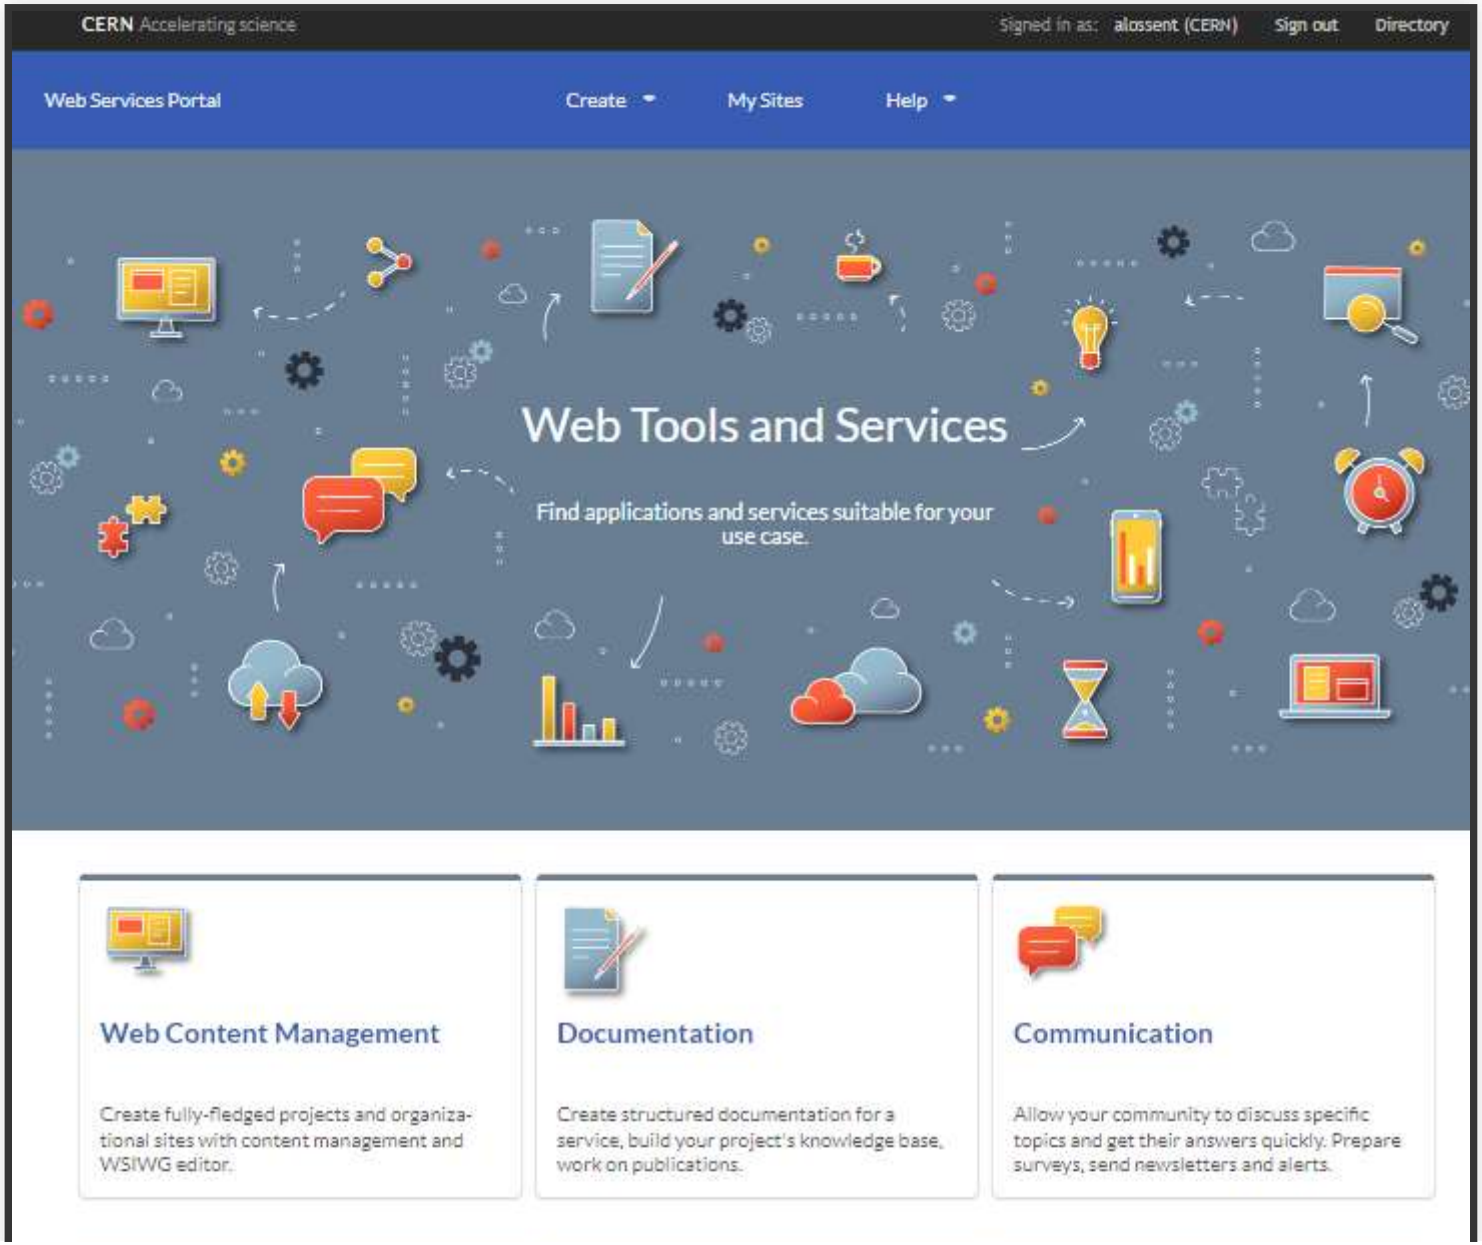


Web Application & Site Hosting
Expose your self-made website and share content with others.
Software development
Get support for the whole software lifecycle: issue tracking, version control, continuous integration and deployment, repository management, and others.
Monitoring Solutions
Add analytics to your website, get insight into your application performance and identify operational problems with no hassle.

5 . 2

My sites

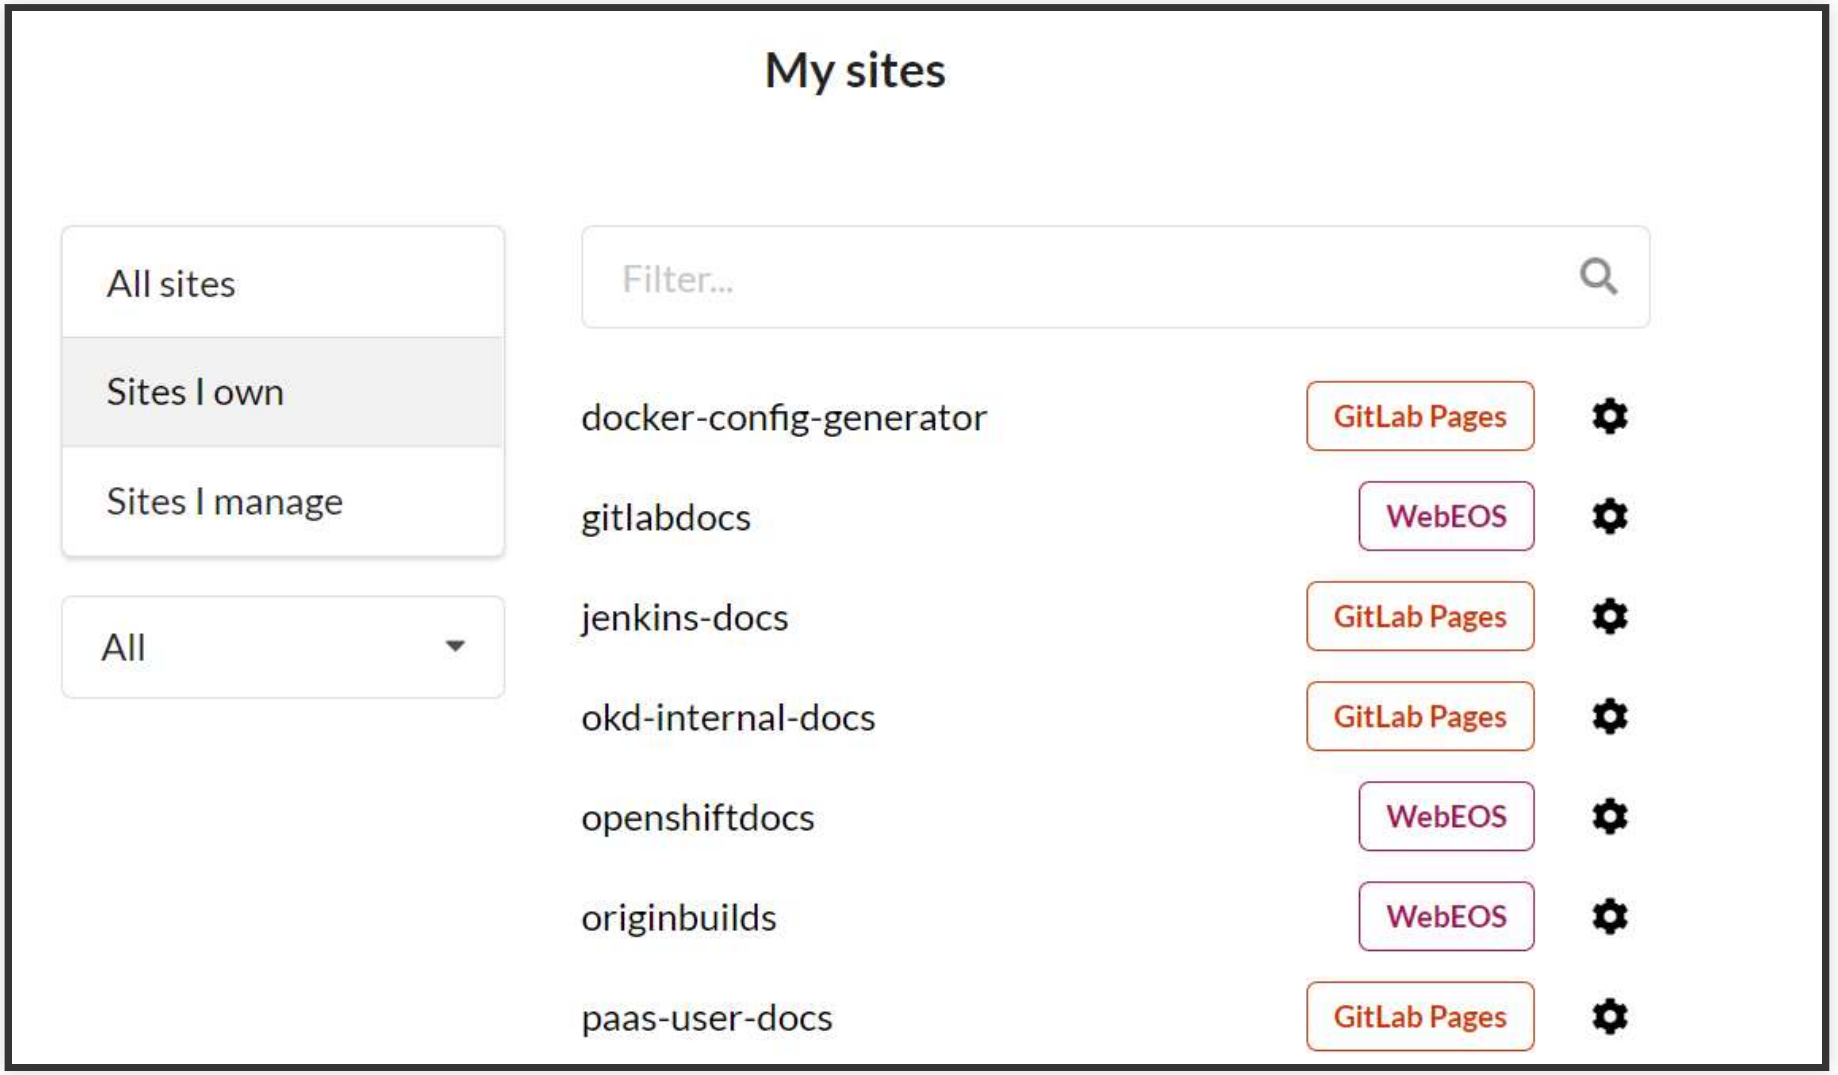
My sites
All sites
Sites I own
Sites I manage
All
Filter...
docker-config-generator
GitLab Pages
gitlabdocs
WebEOS
jenkins-docs
GitLab Pages
okd-internal-docs
GitLab Pages
openshiftdocs
WebEOS
originbuilds
WebEOS
paas-user-docs
GitLab Pages

5 . 3

Site management

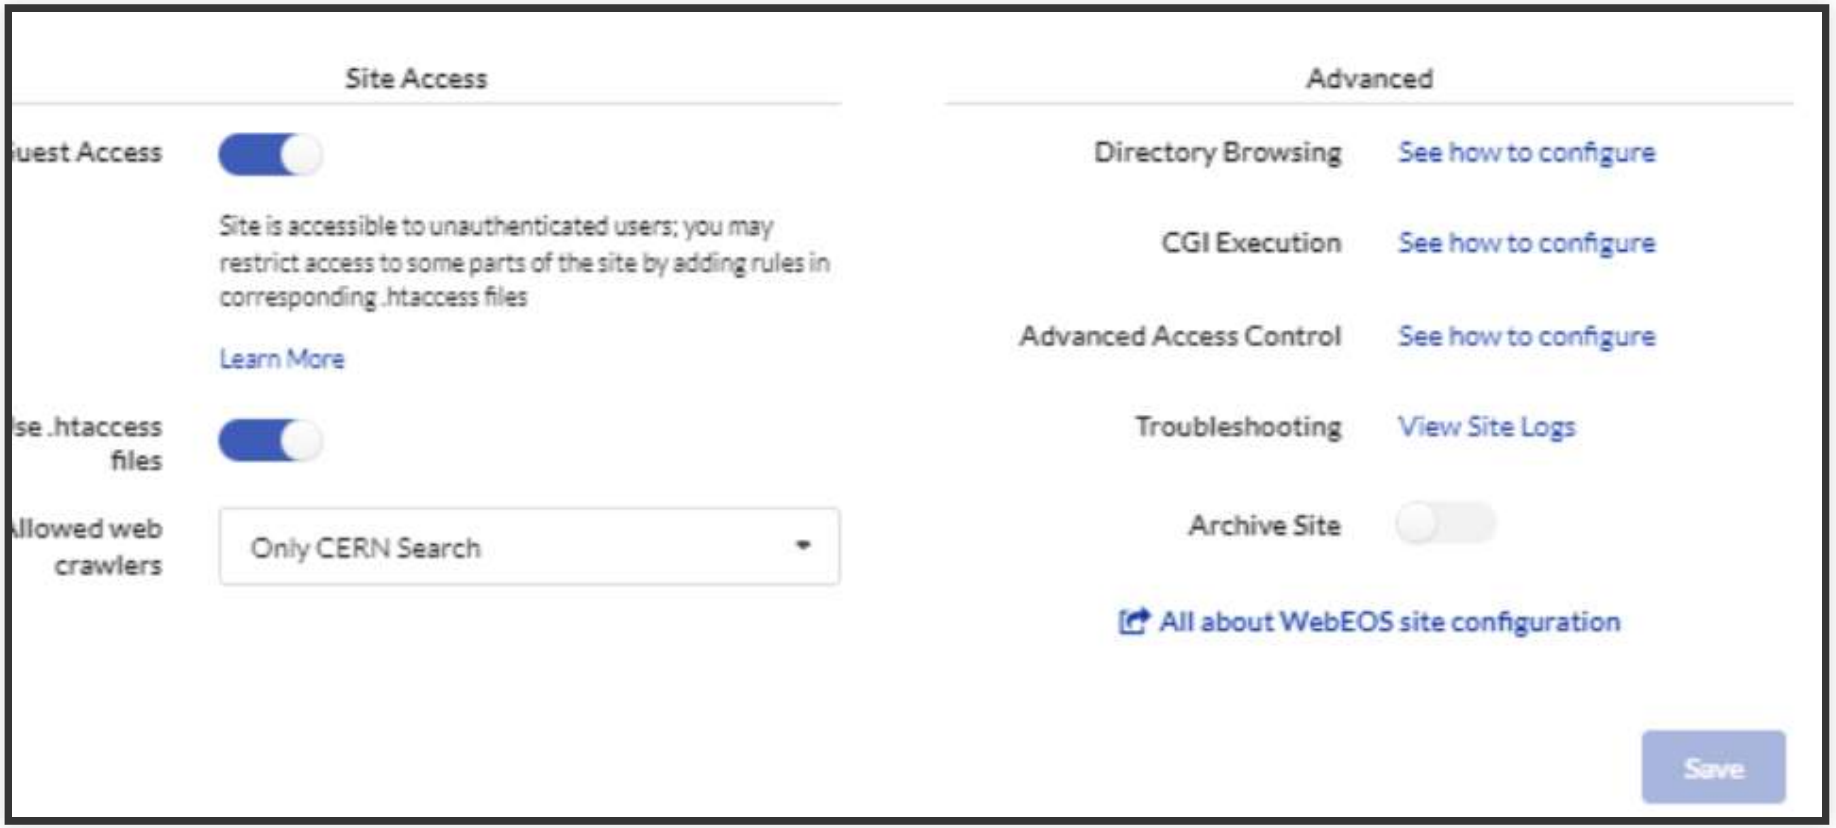

5 . 4

Behind the scenes

- Each web site/application is an OKD4 project (= Kubernetes namespace) in one OKD4 cluster
  - 4 production clusters: webeos, paas, drupal, app-catalogue
- The portal is essentially a stateless front-end for the Kubernetes API
  - Aggregated view of owned projects in all clusters
  - Kubernetes operators implement all the logic for site provisioning, configuration changes

5 . 5

Application Portal integration

Each OKD4 project registers itself into the Application Portal.

The Application Portal provides:

- SSO registration for each web site/application
- Management of application roles
- Lifecycle, e.g. what happens when the owner leaves CERN

The application in the Application Portal is the actual computing resource you own.

The Heart

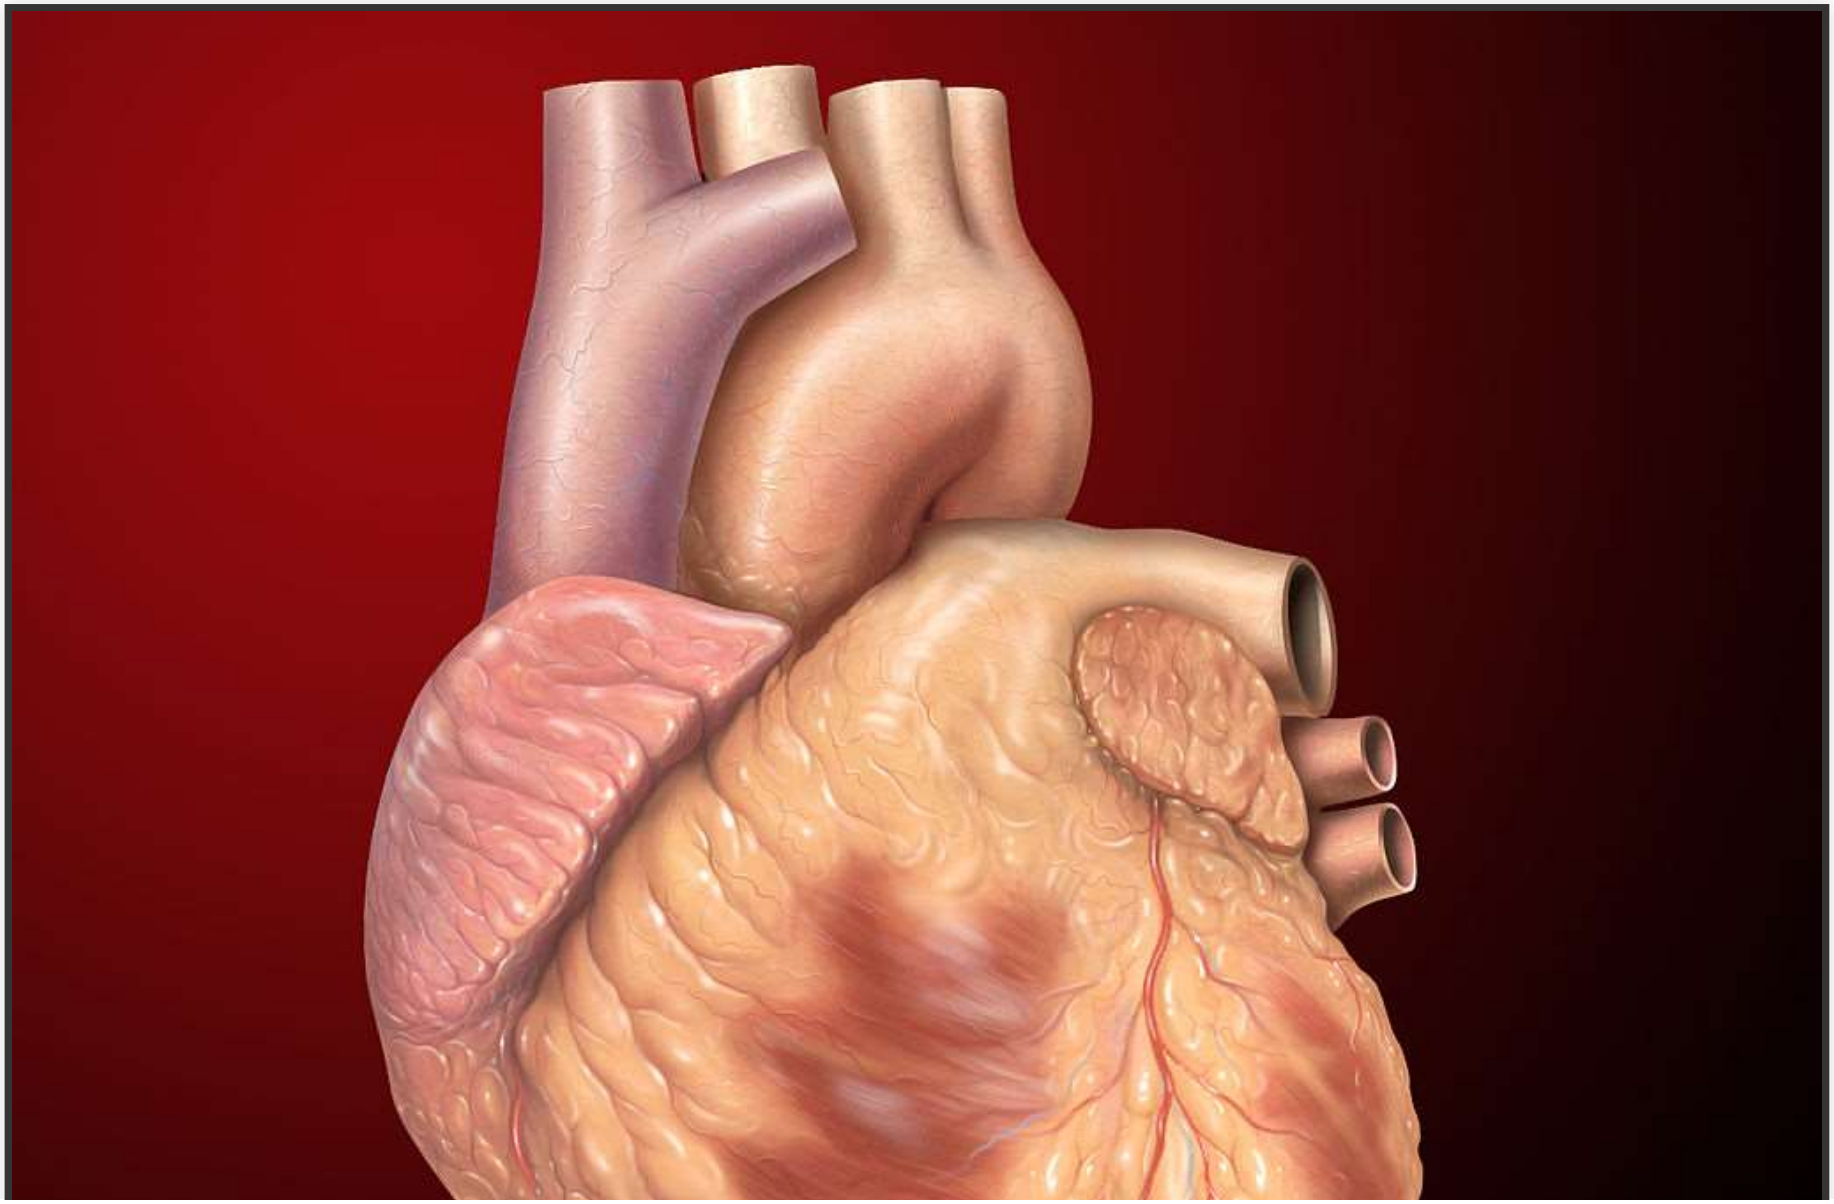

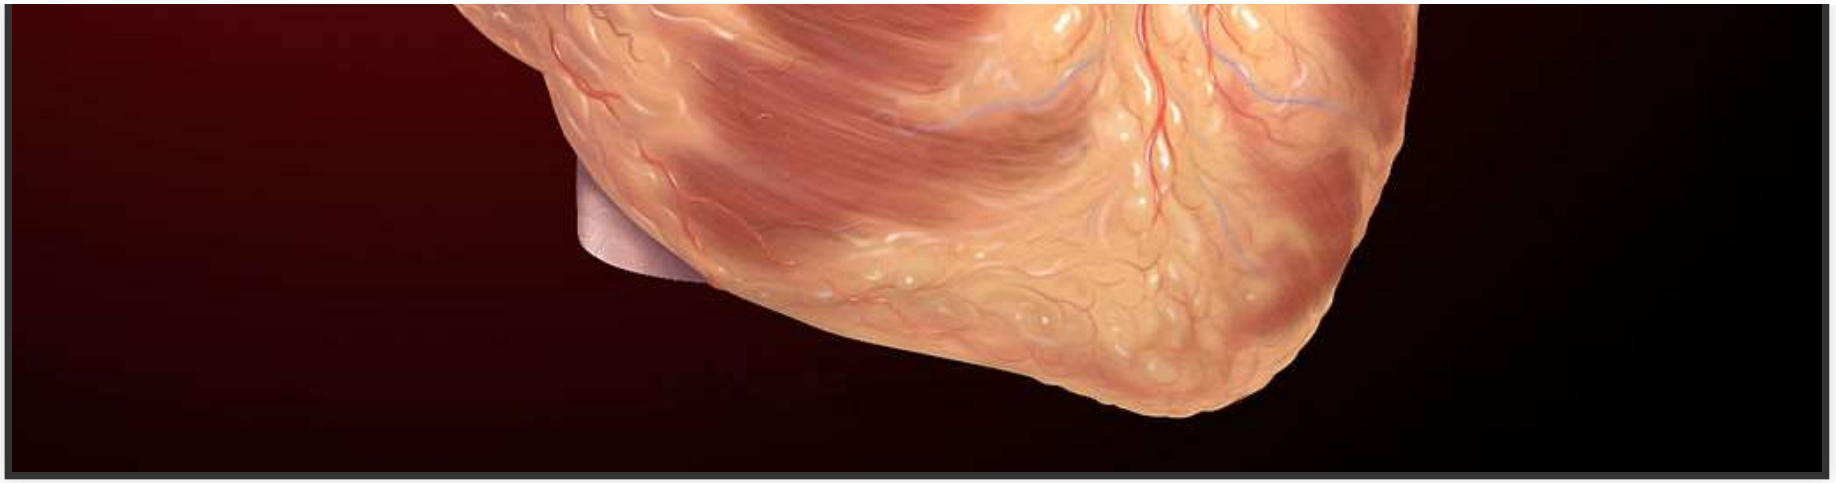

Kubernetes operators

Software extensions to Kubernetes, following the same design principle as Kubernetes itself

- Declarative API: operators extend the Kubernetes API with custom resource types (e.g. `"DrupalSite"`) that describe *what we want*
- A *control loop* implements reconciliation logic to enforce the desired state

DrupalSite example

```
apiVersion: drupal.webservices.cern.ch/v1alpha1
kind: DrupalSite
  name: home
spec:
  configuration:
    databaseClass: critical
    diskSize: 10Gi
    qosClass: critical
    scheduledBackups: enabled
  siteUrl:
    - home.web.cern.ch
    - home.cern
    - press.cern
    - ...
  version:
    name: v9.4-1
    releaseSpec: RELEASE-2022.09.29T12-31-15Z
status:
  availableBackups:
```

6.3

Drupal operator control loop

- Create container deployments, database, DNS records, SSO...
- Manage site configuration
- Automate tasks: clone sites, upgrades, backup/restore...
- Leverage Kubernetes ecosystem, e.g.
  - `Tekton` to perform async tasks
  - `Velero` for resource backup

6.4

Operators for site types

Every hosted site/application is described by a Kubernetes resource inside an OKD4 project, e.g.:

- `UserProvidedDirectory` describes a webeos site
- `GitlabPagesSite`, `DrupalSite`, `WordPress`, `Grafana`...

Each resource type has an associated operator.
The operator's control loop brings the site/application to the desired state.

6.5

Web portal

The new web portal provides the UI to edit these resources

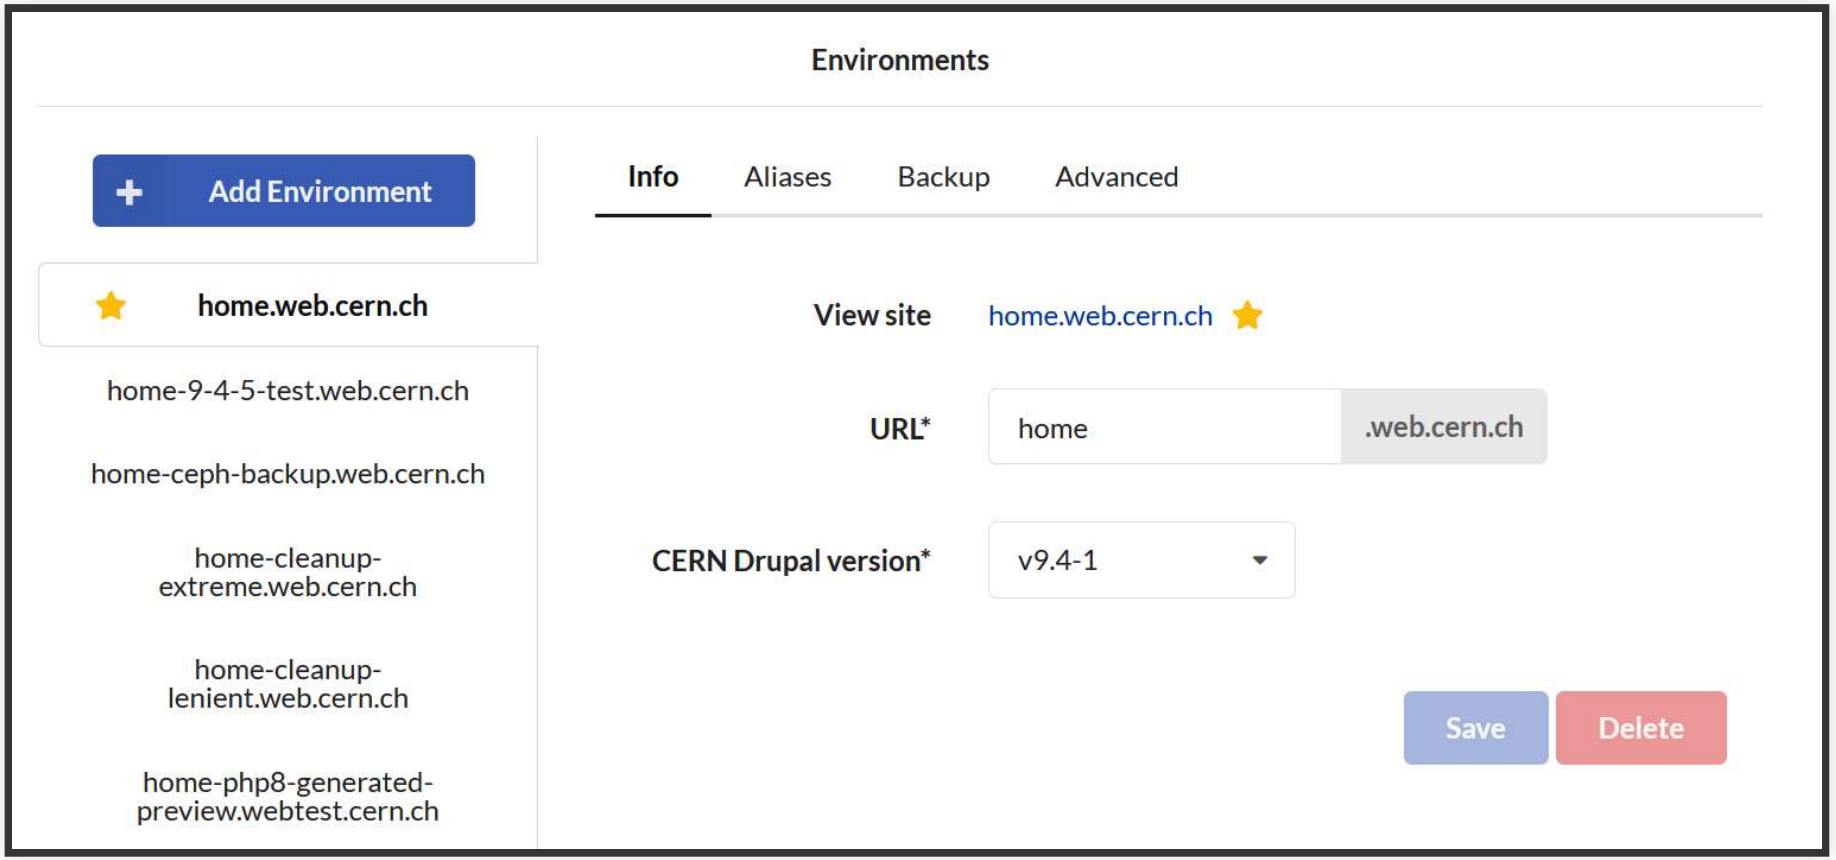

6.6

Infrastructure operators (examples)

- LanDB operator
  - Custom Resources: `LandbSet`, `DelegatedDomain`…
- App Portal operator
  - Manages SSO registration and lifecycle for user applications
  - Custom Resources: `ApplicationRegistration`, `ProjectLifecyclePolicy`…

Operator - languages

Operator SDK

- Application templates: “Helm operators”
- Go for the rest

Policies

Policy enforcement by CNCF Open Policy Agent, e.g.

- unique hostname across all clusters
- security team approvals (TN visibility…)
- automation of EOS mounts, publication of hostnames in DNS…

Timeline & future plans

Public domain

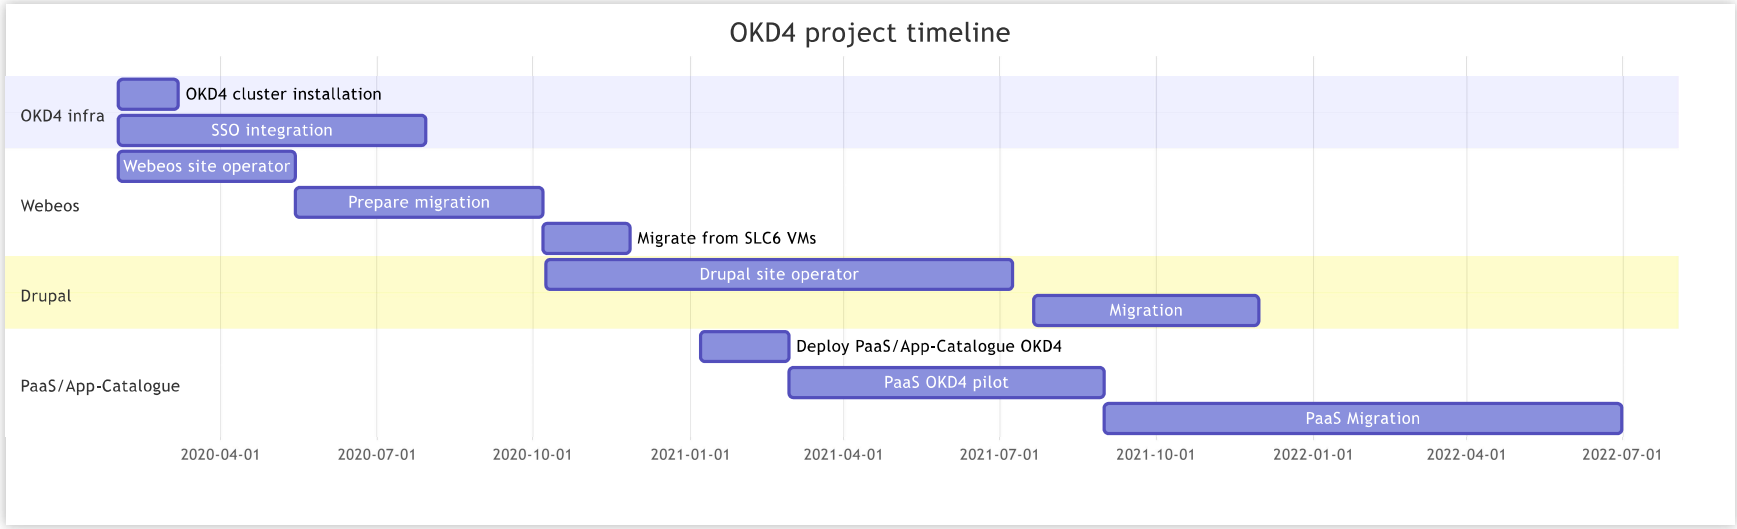
OKD4 project timeline
OKD4 infra
OKD4 cluster installation
SSO integration
Webeos site operator
Webeos
Prepare migration
Migrate from SLC6 VMs
Drupal
Drupal site operator
Migration
Deploy PaaS/App-Catalogue OKD4
PaaS/App-Catalogue
PaaS OKD4 pilot
PaaS Migration
2020-04-01
2020-07-01
2020-10-01
2021-01-01
2021-04-01
2021-07-01
2021-10-01
2022-01-01
2022-04-01
2022-07-01

7 . 2

PaaS - plans

Usability improvements

- gitops for user applications
- logging, monitoring services for user apps

WebEOS - plans

Managed EOS folders

- Provision & manage workspaces under `/eos/web`
- UI to manage folder properties (access control, CGI...)

Goal: facilitate migration of web sites hosted on DFS, AFS

Consolidate remaining site types

- Migrate web sites still managed by legacy web services:
  - WebAFS => WebEOS
  - WebDFS (Windows IIS) => WebEOS (static, PHP), PaaS (dotnet)
  - SharePoint => SharePoint Online

7 . 5

Conclusion

- New architecture for web hosting services
  - Built around OKD4 and Kubernetes Operators
  - Supports the range of hosting services: simple/static sites to content management to complex containerized web apps
- Kubernetes operator model: fit for use
  - Collection of small software development projects interacting trough K8s API
  - Composition, reusability