



# **CERN's journey with OpenShift Origin and OKD**

Jack Henschel

OpenShift Commons Gathering @ KubeCon EU 2023

# Web services at CERN

# Web services at CERN



CERN main  
campus




# Web services at CERN





# Web services at CERN



# The Origins in 2014

Goal: enable **wider adoption of CI by facilitating deployment of Jenkins instances**

Historical context: Jenkins setup on dedicated VMs, provisioned with OpenStack and Puppet, requested via ticket

→ Looking for ways to **consolidate resources** and **simplify management** (especially for small software projects)



# Hosting Jenkins on OpenShift

- OpenShift's **Jenkins template** proved to be an **ideal place for getting started**
- Jenkins admins have/need a lot of control  
→ **offer not only a service, but a platform**



# Hosting Jenkins on OpenShift

- OpenShift's **Jenkins template** proved to be an **ideal place for getting started**
- Jenkins admins have/need a lot of control  
→ **offer not only a service, but a platform**

**OpenShift is capable of much more** – why not also use it as a **PaaS?**

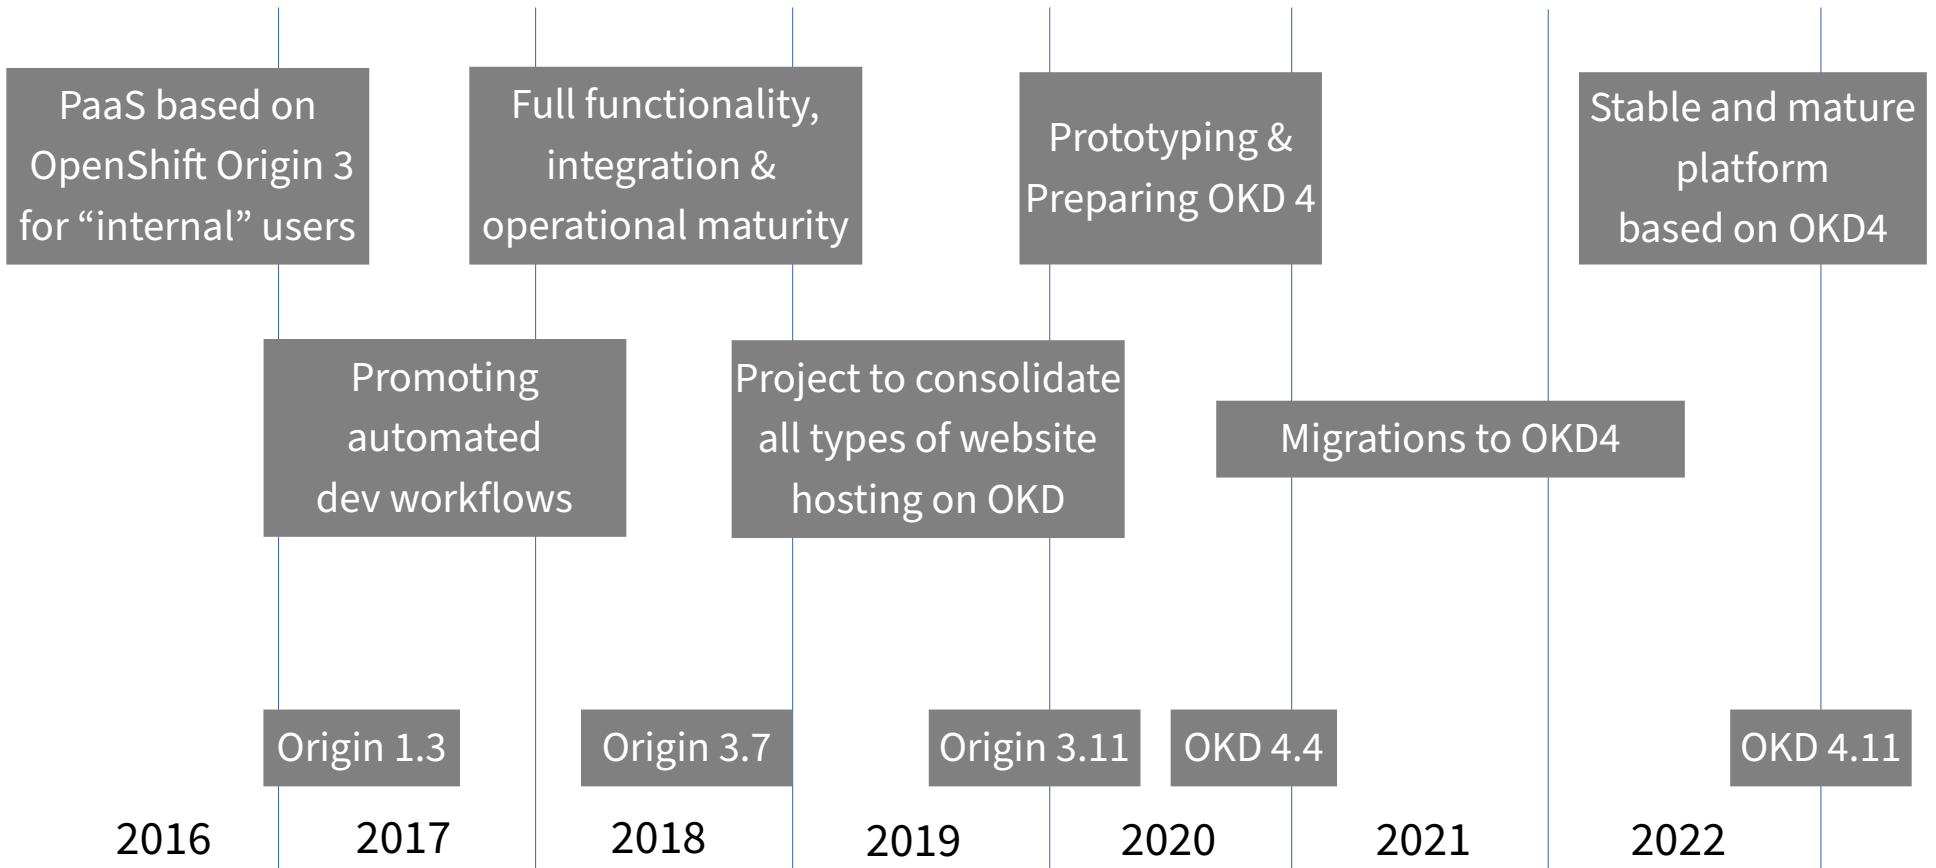
- Lots of **small web apps** (Python, PHP, Tomcat) that occupy **dedicated VMs**
- **Manual setup:** OpenStack VM, Puppet, Database, DNS, SSO, version control ...
- **Requirements:** low setup overhead, resource efficiency, minimal ongoing maintenance cost (OS and security management)



# Moving towards a general-purpose PaaS

- **Continuous Deployment** with **S2I workflow** (*BuildConfig + DeploymentConfig*)  
→ **automated deployments became easy!**
- **Knowledge sharing / teaching** was necessary:  
Containers, Images, Docker, Kubernetes ...
- **Integrations** needed to be developed:  
CERN's website management, DNS, Firewall, SSO, various storage systems

# Timeline



# Fast-forward to 2023

**OKD4** is the **foundation of Webservices Infrastructure** at **CERN**

Provides a **multi-tenant, highly-available** and **secure base**

**Enhanced** by us with additional features and integrations for:

- Hostname registration, DNS setup, certificates, backups
- Authentication and resource management
- Storage: CephFS, EOS, CVMFS
- Ingress router sharding

→ Using **operators, controllers** and **webhooks**



# OKD4 @ CERN

“Our” OKD provides **shared base** for different **cluster flavors**:

# OKD4 @ CERN


“Our” OKD provides **shared base** for different **cluster flavors**:



1400 projects, 96 nodes,  
1500 cores, 2.7 TiB memory

# OKD4 @ CERN

“Our” OKD provides **shared base** for different **cluster flavors**:



1400 projects, 96 nodes,  
1500 cores, 2.7 TiB memory




300 projects, 50 nodes,  
400 cores, 770 GiB memory



# OKD4 @ CERN


“Our” OKD provides **shared base** for different **cluster flavors**:



1400 projects, 96 nodes,  
1500 cores, 2.7 TiB memory




300 projects, 50 nodes,  
400 cores, 770 GiB memory



4100 projects, 20 nodes,  
270 cores, 600 GiB memory

# OKD4 @ CERN


“Our” OKD provides **shared base** for different **cluster flavors**:




1400 projects, 96 nodes,  
1500 cores, 2.7 TiB memory



300 projects, 50 nodes,  
400 cores, 770 GiB memory



4100 projects, 20 nodes,  
270 cores, 600 GiB memory



800 projects, 60 nodes,  
900 cores, 1.7 TiB memory

# Web Services Portal




**Stateless** web UI **federates** OKD clusters and offers entrypoint for **non-technical users**






# Web Services Portal

**Stateless** web UI **federates** OKD clusters and offers entrypoint for **non-technical users**



 <b>Content Management</b> Manage fully-fledged projects and organization sites with content management and WYSIWYG editor.	 <b>Documentation</b> Create structured documentation for a service, build your project's knowledge base, work on publications.	 <b>Communication</b> Allow your community to discuss topics and get their answers, surveys, send newsletters.
--	--	---

 <b>Application &amp; Site Hosting</b> Deploy your self-made website and share it with others.	 <b>Software development</b> Get support for the whole software lifecycle: issue tracking, version control, continuous integration and deployment, repository management, and others.	 <b>Monitoring Solution</b> Add analytics to your web application performance, operational problems with...
--	---	---


# Web Services Portal

**Stateless** web UI **federates** OKD clusters and offers entrypoint for **non-technical users**



 <b>Content Management</b> Fully-fledged projects and organization sites with content management and WYSIWYG editor.	 <b>Documentation</b> Create structured documentation for a service, build your project's knowledge base, work on publications.	 <b>Communication</b> Allow your community to discuss topics and get their answers, surveys, send newsletters.
--	---	--

 <b>Application &amp; Site Hosting</b> Deploy your self-made website and share it with others.	 <b>Software development</b> Get support for the whole software lifecycle: issue tracking, version control, continuous integration and deployment, repository management, and others.	 <b>Monitoring Solutions</b> Add analytics to your web application performance, operational problems with...
--	---	--



## Creating a WebEOS site

You're minutes away from getting your site up and running.

1. WebEOS sites have their content stored on EOS. Share the EOS location of your choice with user a : `wwwEOS` via the cernbox web interface.  
See how: [for personal sites](#) , [for project sites](#)
2. Create `index.html` file in the EOS location. [See example](#)
3. Fill in the form below and let us do the magic

Our recommendations for choosing a site name and category [↗](#)

Personal - deleted when owner leaves

my-site .docs.cern.ch

Documentation for my-site

/eos/user/j/jack/www

I have read and agreed to the [CERN Computing Rules](#) and taken into account the [design guidelines](#) for websites and the [website lifecycle policy](#).

Create

# Web Services Portal

Stateless web UI federates OKD clusters and offers entrypoint for **non-technical users**



### Content Management

Manage your fully-fledged projects and organizational sites with content management and WYSIWYG editor.

### Documentation

Create structured documentation for a service, build your project's knowledge base, work on publications.

### Communication

Allow your community to discuss topics and get their answers, surveys, send newsletters.

### Application & Site Hosting


Host your self-made website and share it 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 web application performance, operational problems with...



## Creating a WebEOS site

You're minutes away from getting your site up and running.

- WebEOS sites have their content stored on EOS. Share the EOS location of your choice with user a : `wwwEOS` via the cernbox web interface.  
See how: [for personal sites](#) , [for project sites](#)
- Create `index.html` file in the EOS location. [See example](#)
- Fill in the form below and let us do the magic

Our recommendations for choosing a site name and category [↗](#)

Personal - deleted when owner leaves

my-site .docs.cern.ch

Documentation for my-site

/eos/user/j/jack/www

I have read and agreed to the [CERN Computing Rules](#) and taken into account the [design guidelines](#) for websites and the [website lifecycle policy](#).

Create


Info	Aliases	Backup	Web Analytics	Advanced	
View site	drupal-demo.web.cern.ch ★				
URL*	drupal-der .web.cern.ch				
CERN Drupal version*	v9.4-2				
				Save	Delete

# Web Services Portal

**Stateless** web UI **federates** OKD clusters and offers entrypoint for **non-technical users**



 <b>Content Management</b> Fully-fledged projects and organizations with content management and WYSIWYG editor.	 <b>Documentation</b> Create structured documentation for a service, build your project's knowledge base, work on publications.	 <b>Communication</b> Allow your community to discuss topics and get their answers, surveys, send newsletters.
 <b>Application &amp; Site Hosting</b> Deploy your self-made website and share it with others.	 <b>Software development</b> Get support for the whole software lifecycle: issue tracking, version control, continuous integration and deployment, repository management, and others.	 <b>Monitoring Solutions</b> Add analytics to your web application performance, operational problems with...



## Creating a WebEOS site

You're minutes away from getting your site up and running.

1. WebEOS sites have their content stored on EOS. Share the EOS location of your choice with user a : `wwwEOS` via the cernbox web interface.  
See how: [for personal sites](#) , [for project sites](#)
2. Create `index.html` file in the EOS location. [See example](#)
3. Fill in the form below and let us do the magic

Our recommendations for choosing a site name and category [↗](#)

Personal - deleted when owner leaves	▼
my-site	.docs.cern.ch ▼
Documentation for my-site	
/eos/user/j/jack/www	

I have read and agreed to the [CERN Computing Rules](#) and taken into account the [design guidelines](#) for websites and the [website lifecycle policy](#).

Create

Info Aliases **Backup** Web Analytics Advanced

**★ Backup policy for the primary environment**  
This environment is the primary one and will be backed up automatically every 48 hours.

Create new backup	my-backup	Create
Restore from backup	Backup to restore ▼	Restore

# Behind the scenes

**apiVersion:** drupal.webservices.cern.ch/v1alpha1

**kind:** DrupalSite

**metadata:**

**name:** drupal-tools

**spec:**

**configuration:**

databaseClass: standard

diskSize: 1G

qosClass: standard

scheduledBackups: enabled

**siteUrl:**

- drupal-tools.web.cern.ch

**version:**

**name:** v9.4-2

releaseSpec: RELEASE-2023.02.13T13-47-51Z

**status:**

availableBackups: [...]

dbUpdatesLastCheckTimestamp: 'Feb 14, 2023 at 7:38am (UTC)'

expectedDeploymentReplicas: 1

# OKD cluster management

- Clusters are **pets**: production clusters are **stateful** since they run and store **user workload**
- Each cluster is completely **self-sufficient** and **isolated**
- Developed internal ***okdctl* tool** to facilitate common operations (creating/deleting clusters, replacing nodes)
- OKD4 **in-place cluster upgrades** are completely **automated** and **seamless**
- All **infra workloads** are **managed by ArgoCD**





# GitOps with ArgoCD



- **Natural extension** of Kubernetes' **continuous reconciliation** model
- Ensures all resources converge to the desired state (incl. orchestration)
  - despite manual actions in the cluster (troubleshooting, debugging etc.)
  - automatic alerts if this is not the case
- Fits the **operator-driven cluster management** of OKD
- **CLI & Web UI** are useful **for understanding** which resources are deployed and **what their state is**

# Automated provisioning & Integration tests

**Fully automated provisioning**  
(with custom tool) of **isolated**  
**clusters**

```
$ okctl/okctl provision-cluster --yes --cluster-name "${CLUSTER_NAME}"  
# Checking provided SSH key  
# Obtaining openshift-installer binary [11:18:31]  
# Generating ignition configs [11:18:41]  
# Uploading ignition data for bootstrap VM [11:18:48]  
# Creating the bootstrap VM [11:18:49]  
+ openstack server create --wait ci-28717890-boot-5wujq --format json --
```

# Automated provisioning & Integration tests

**Fully automated provisioning**  
(with custom tool) of **isolated clusters**

Paired with **integration tests** that **verify almost every feature** works as expected from a **user & admin perspective**

```
$ okdctl/okdctl provision-cluster --yes --cluster-name "${CLUSTER_NAME}"  
# Checking provided SSH key  
# Obtaining openshift-installer binary [11:18:31]  
# Generating ignition configs [11:18:41]  
# Uploading ignition data for bootstrap VM [11:18:48]  
# Creating the bootstrap VM [11:18:49]  
+ openstack server create --wait ci-28717890-boot-5wujq --format json --
```

```
$ bats -rpT tests/1-common/  
1-authz-operator.bats  
[ 1/22] SSO registration with application lifecycle and bootstrap application role ✓  
[ 2/22] Namespace is not deleted if ProjectLifecyclePolicy says not to delete it ✓  
[ 3/22] Namespace is blocked if ProjectLifecyclePolicy says 'BlockAndDeleteAfterGrad  
[ 4/22] Namespace is not deleted if AppReg is deleted ✓ [14s]  
2-cephfs-persistent-volumes.bats  
[ 5/22] Test provisioning, backup and deletion of cephfs PVs ✓ [332s]  
3-opa-cephfs.bats  
[ 6/22] Test OPA pv.opa.openshift.cern.ch/set-default-annotations ✓ [13s]  
[ 7/22] Test OPA pv.opa.openshift.cern.ch/set-default-labels ✓ [13s]  
4-opa-custom-ingress-default-hostname.bats  
[ 8/22] Generate default route hostname using domain in namespace annotation ✓ [8s]  
[ 9/22] Do not modify default route hostname generated by OKD if no domain specified  
[10/22] Do not modify user specified route hostname ✓ [5s]
```

# Automated provisioning & Integration tests

Fully automated provisioning  
(with custom tool) of **isolated clusters**

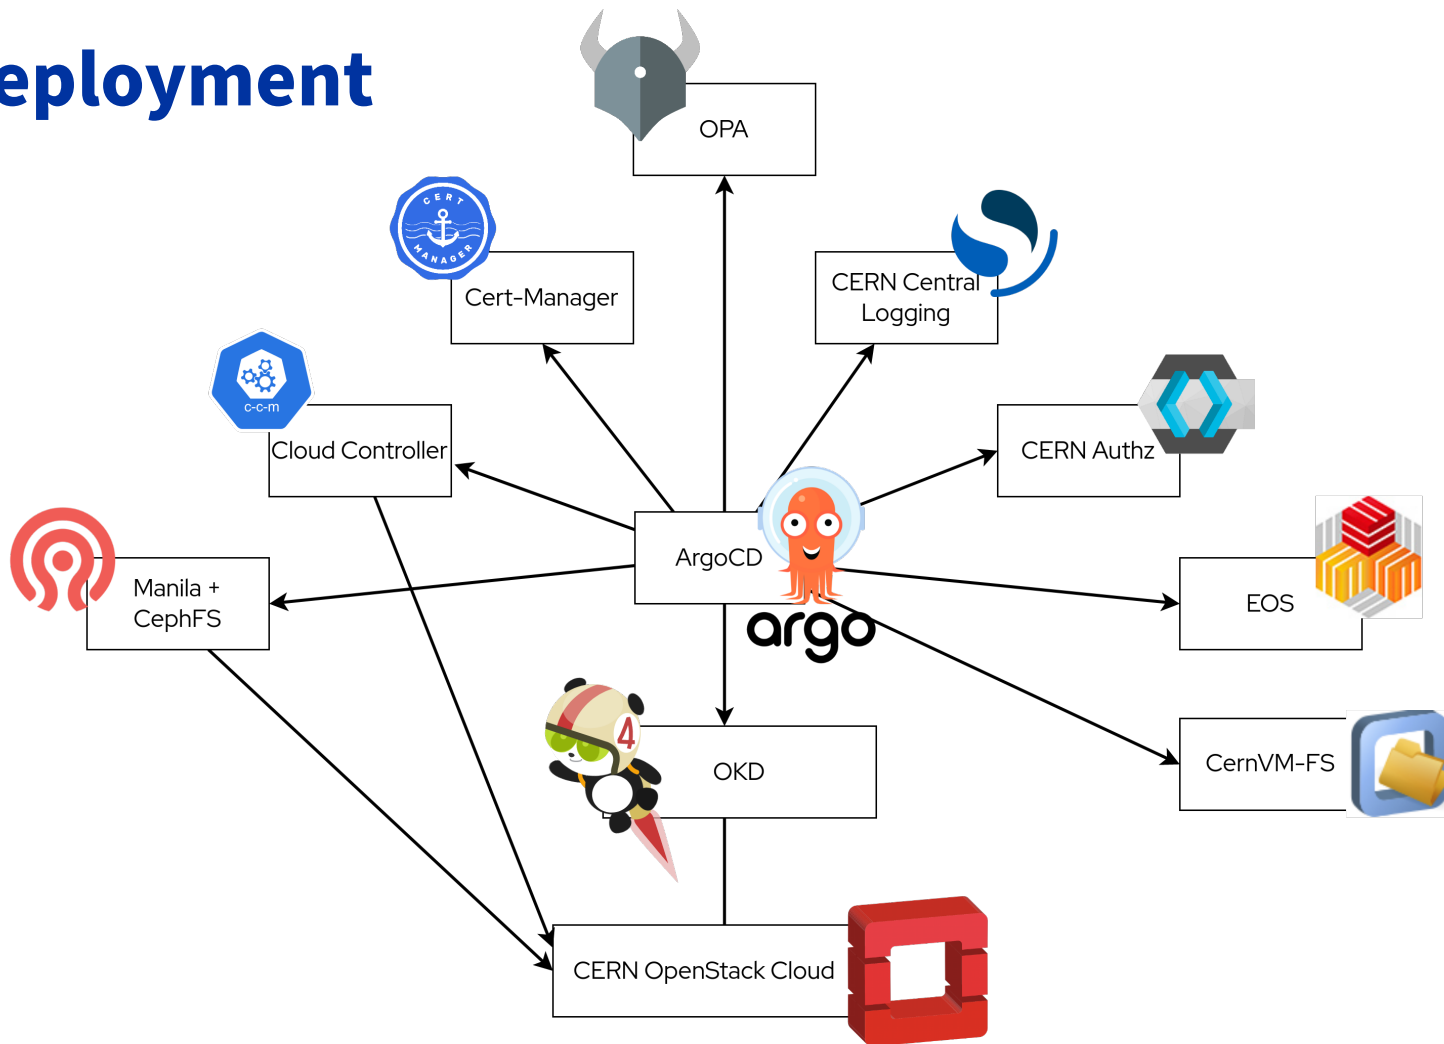
Paired with **integration tests** that **verify almost every feature** works as expected from a **user & admin perspective**

Allows **deploying changes frequently** and **predictably**

```
$ okdctl/okdctl provision-cluster --yes --cluster-name "${CLUSTER_NAME}"
# Checking provided SSH key
# Obtaining openshift-installer binary [11:18:31]
# Generating ignition configs [11:18:41]
# Uploading ignition data for bootstrap VM [11:18:48]
# Creating the bootstrap VM [11:18:49]
+ openstack server create --wait ci-28717890-boot-5wujq --format json --
```

```
$ bats -rpT tests/1-common/
1-authz-operator.bats
[ 1/22] SSO registration with application lifecycle and bootstrap application role ✓
[ 2/22] Namespace is not deleted if ProjectLifecyclePolicy says not to delete it ✓
[ 3/22] Namespace is blocked if ProjectLifecyclePolicy says 'BlockAndDeleteAfterGracePeriod' ✓
[ 4/22] Namespace is not deleted if AppReg is deleted ✓ [14s]
2-cephfs-persistent-volumes.bats
[ 5/22] Test provisioning, backup and deletion of cephfs PVs ✓ [332s]
3-opa-cephfs.bats
[ 6/22] Test OPA pv.opa.openshift.cern.ch/set-default-annotations ✓ [13s]
[ 7/22] Test OPA pv.opa.openshift.cern.ch/set-default-labels ✓ [13s]
4-opa-custom-ingress-default-hostname.bats
[ 8/22] Generate default route hostname using domain in namespace annotation ✓ [8s]
[ 9/22] Do not modify default route hostname generated by OKD if no domain specified ✓
[10/22] Do not modify user specified route hostname ✓ [5s]
```

# OKD deployment



# Spotlight: OpenPolicyAgent



**OPA** is used for a wide range of use cases (to **help admins & users**):

- **Unique hostnames** across all clusters
- **Ingress** sharding
- **Volume** labels & annotations (used for backups and mount permissions)
- **Network** visibility (Internet/Intranet/Private Networks)
- **Automation** of EOS mounts (initContainer + sidecar injection for authentication)



# What we like about OpenShift Origin & OKD

- **Strong multi-tenancy, security** and **high availability** out-of-the-box
- It's **stable**
- Simple yet **powerful web UI**
- It's fully **open source**
  - We can troubleshoot and fix issues ourselves!
  - We can contribute back!

openshift / router Public

<> Code Issues 9 Pull requests 14 Security Insights

## Respect route targetPort in dynamic config #7

Merged openshift-merg... merged 1 commit into openshift:master from alexcern:fix/dynamic\_config\_targetport

Conversation 16 Commits 1 Checks 0 Files changed 1

alexcern commented on Jan 16, 2019

Fix #6

openshift / oc Public

<> Code Issues 26 Pull requests 32 Projects Security Insights

## Allow triggers on batch/v1 CronJobs #1077

Merged openshift-merg... merged 1 commit into openshift:master from mic4ael:set-trig

Conversation 42 Commits 1 Checks 0 Files changed 1

mic4ael commented on Feb 21, 2022 • edited

openshift / router Public

<> Code Issues 9 Pull requests 14 Security Insights

## Bug 2093454: HAProxy: enable PROXY protocol for all list

Merged openshift-merg... merged 1 commit into openshift:master from jacksgt:fix-ipv4v6-proxy on Jun 10, 2022

Conversation 51 Commits 1 Checks 0 Files changed 1

jacksgt commented on Jan 19, 2022

Previously, the `accept-proxy` directive (which enables HAProxy to listen for PROXY protocol instead of regular HTTP traffic) was only appended to the last entry in the section. This worked fine for the IPv4 or IPv6 single-stack use-case, but not

openshift / openshift-ansible Public

<> Code Issues 4 Pull requests 4 Projects Security Insights

## Set nameservers on DHCPv6 event #2853

Merged sdodson merged 1 commit into openshift:master from alexcern:dhcpv6

Conversation 11 Commits 1 Checks 0 Files changed 1

alexcern commented on Nov 23, 2016

openshift / cluster-logging-operator Public

<> Code Issues 2 Pull requests 9 Projects Security Insights

## Fix namespace metadata in gen-olm-artifacts

Merged openshift-merg... merged 1 commit into openshift:master from jacksgt:patch-1

openshift / router Public

<> Code Issues 9 Pull requests 14 Security Insights

## Use ServiceAliasConfigKeys for dynamic config #17 / 20

Merged openshift-merg... merged 1 commit into openshift:master from alexcern:fix/dynamic\_route

Conversation 10 Commits 1 Checks 0 Files changed 1

# Takeaways

Users are very **happy about internal documentation**

**Operators** are a great way to provide automation **for users and admins**

- but they are also **very sharp tools** → **use soft deletion where possible!**

**Worthwhile effort** to fully automate **cluster provisioning** and **integration tests**

**For admins: not every manual operation needs to be automated**

# Thank you!



**Slides: <https://cern.ch/openshift-commons-2023>**



[home.cern](https://home.cern)