

Dynamic On-Demand Analysis Service (DODAS)

Daniele Spiga - INFN - (spiga@infn.it)

On behalf of DODAS Team

Jennifer2 Computing WorkShop - Task 5.1

CERN, Thursday 12 December 2019

- ❑ Introduction to DODAS
 - ❑ Main concepts

- ❑ DODAS and Kubernetes
 - ❑ Stateless sites
 - ❑ Managing Compute and Data

- ❑ A note on AuthN/Z

- ❑ Summary and Future work

What is DODAS



Dynamic On Demand Analysis Service: DODAS

A INFN solution designed with the goal to enable users **to create and provision infrastructure deployments, automatically and repeatedly**, on “any cloud provider” **with almost zero effort**.

- Implement the **infrastructure as code paradigm**: driven by a templating engine to specify high-level requirements. Declarative approach **allows to describe “What” instead of “How”**
 - Let the underlying system to abstract providers and automatically instantiate and setup the computing system(s)
- Allows to instantiate **on-demand container-based clusters** (Mesos/**Kubernetes**) to execute software applications:
 - E.g. HTCondor batch system, Spark cluster, Data Caches...
 - **But also composition of services e.g. to manage stateless (WLCG-compliant) sites**

... and where it comes from



DODAS was initially prototyped within the **INDIGO-DataCloud project (2017)**

- Having in mind a primary use case: to develop an **effective solution for dynamic resource provisioning@CMS** (targeting Opportunistic computing)

Since then it has been **evolved**:

- In terms of supported **use cases** (from HTCondor to BigData platforms)
- In terms of adopted **technologies** (Mesos/Marathon, Kubernetes)
- In terms of supported **communities** (see later)

Currently the project is also supported by **EOSC-hub H2020 EU project as a Thematic Service.**

Thematic Services of EOSC-hub

Thematic services of EOSC-hub are service which come from external partners willing to collaborate with the project.

The project has the objective of involving early adopters, stakeholder etc... aiming at reaching out to new user groups and service providers.

DODAS is one of the TS

marketplace@eosc-hub.eu

Log in Contact



CART 1 Product

COMPUTE DATA & STORAGE PLATFORMS THEMATIC SERVICES IDENTITY & SECURITY PROFESSIONAL SERVICES

Thematic services Dynamic On Demand Analysis Service

DYNAMIC ON DEMAND ANALYSIS SERVICE

Component Metadata Infrastructure

DARIAH

Dynamic On Demand Analysis Service

ENES Climate Analytics Service

LifeWatch ERIC Plants Identification App

OPENCoastS

WeNMR Suite for Structural Biology

VIEWED PRODUCTS



WeNMR Portal

WeNMR is a Virtual Research Community...

Dynamic On Demand Analysis Service

Simplify the access and management of computing resources

DODAS acts as cloud enabler designed for scientists seeking to easily exploit distributed and heterogeneous clouds to process data. Aiming to reduce the learning curve as well as the operational cost of managing community specific services running on distributed cloud, DODAS completely automates the process of provisioning, creating, managing and accessing a pool of heterogeneous computing and storage resources.

DODAS provides:

- A comprehensive approach to opportunistic computing, with the possibility of orchestrate multiple centers (e.g. campus facilities, public or private clouds, to gather all available computing and storage resources).
- A simple solution for elastic computing site extensions, e.g. extension of allocated resources in order to absorb peaks of usage.
- An easy and controlled procedure to dynamically instantiate a spot 'Data Analysis Facility', for example a mission specific site. This is meant as the generation of an ephemeral WLCG-Tier as a Service to share computing and data resources with collaborators.
- The support to create HTCondor batch systems and BigData Platform on demand over multi-backend IaaS cloud resources.

How to access the service

To access the DODAS service, users are first required to [register online](#).

See the online manual:

<https://dodas-ts.github.io/dodas-doc/>

Designed to:

- Support **user tailored** computing **environments**
- **Automate** configuration and deployment of custom services and/or dependencies
- Support declarative approach to define input parameters, customize workflows, **treating a collection of resources together as a single unit**

Highly flexible and modular solution enabling multiple usage patterns:

- Leverage clusters, possibly customising the stack, building highly reliable, highly scalable, applications
 - without worrying about creating and configuring the underlying infrastructure
 - **TOSCA + Ansible + Application setup (e.g Helm)**
- Generate Clusters on demand (**K8s on demand**), possibly customizing the underlying infrastructure
 - **TOSCA + Ansible**
 - And leave users **to focus just on applications**
 - **(e.g. Helm)**
- (Abstract the IaaS provisioning VMs/DB etc... not in scope with today's pre-GDB discussion)

DODAS internals

TOSCA

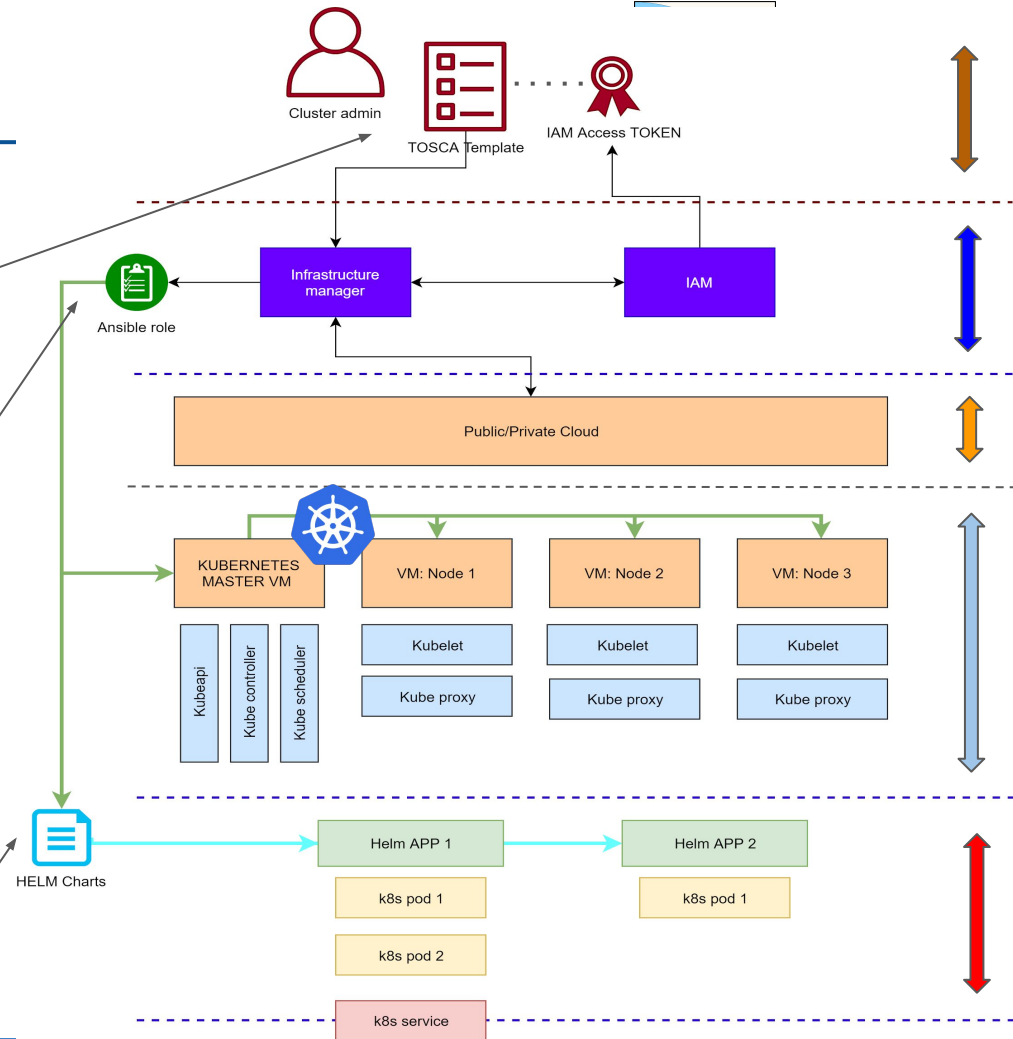
- Define the infrastructure (the HW)
- Define services (k8s) & Applications to setup (through Ansible)
- Declare (“any”) input parameters

Ansible based installation using:

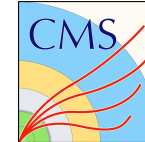
- Kubeadm (initialization)
- Flannel (default but others available)
- nginx ingress (optional)
- k8s dashboard (optional)

Helm (Applications layer)

- Configure at runtime. Dynamically load and compile values (from TOSCA through ansible)
- Install applications



Implementing **stateless sites** with DODAS

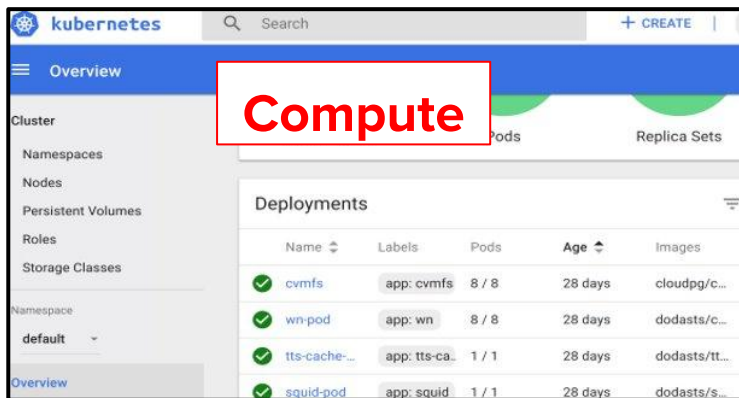


Provides remotely deployed CPU clusters and data Caches. Managed through K8s

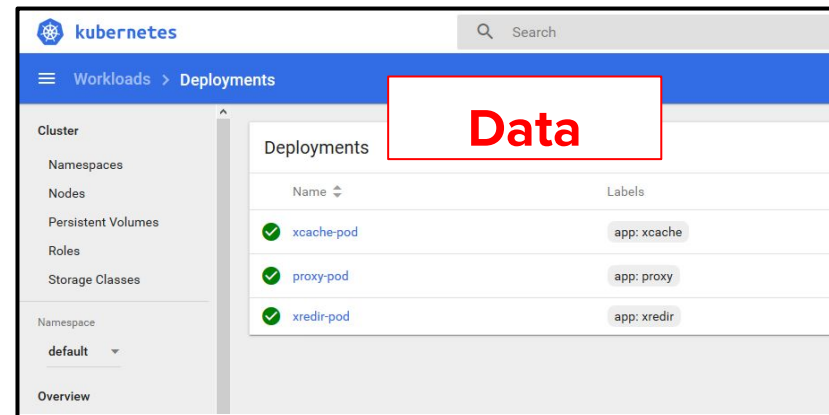
Applications
Layer



- **No Dedicated experiment support**
- **No site specific setup required**



Integrates Experiment Workload Management



Integrates Experiment Storages/Federation

Let's recap so far: key features



- DODAS is a solution to **Abstract the IaaS level** (“any” cloud)
- And allows to **automatically and repeatedly** instantiate computational stack
 - Toward the zeroOps model
- Computational **stack are described by templates**
 - Templates allow to customize any input parameter and which kind of service user want need
- **Highly customizable** and **experiment agnostic** by design

Adopting DODAS...



Where and how DODAS can take part in a experiment specific Workload Management (WM) need/want to integrate a cloud provider?

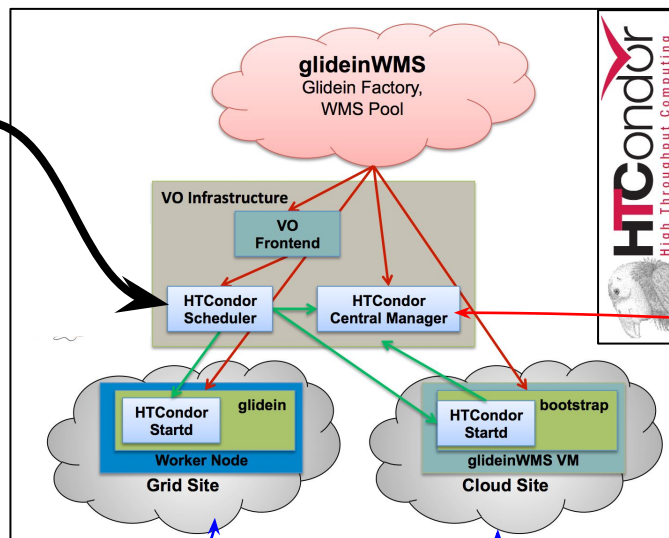
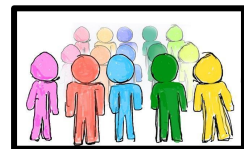
DODAS brings the cloud to the WM in a transparent and seamlessly way. WM doesn't need to do anything to include a new cloud. Two patterns are foreseen by DODAS:

- Either auto-deploy worker nodes which connect to the WM (to pull work, payload)
 - Assumption: community provides WN-daemon configuration/contextualization
 - e.g. here to go for getting work, etc
- Or providing the WM with a gateway as a conventional site would do
 - Assumption: community defines which is the preferred gateway
 - (e.g. Computing Element and batch). HTCondor is there

A concrete example

The CMS submission infrastructure

CMS Physicists



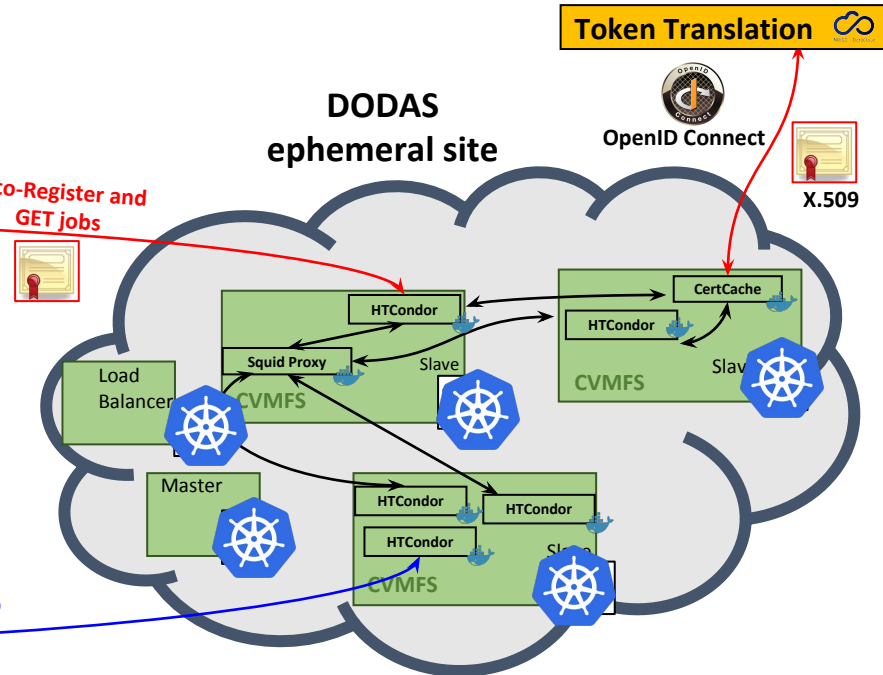
CMS Distributed Storages



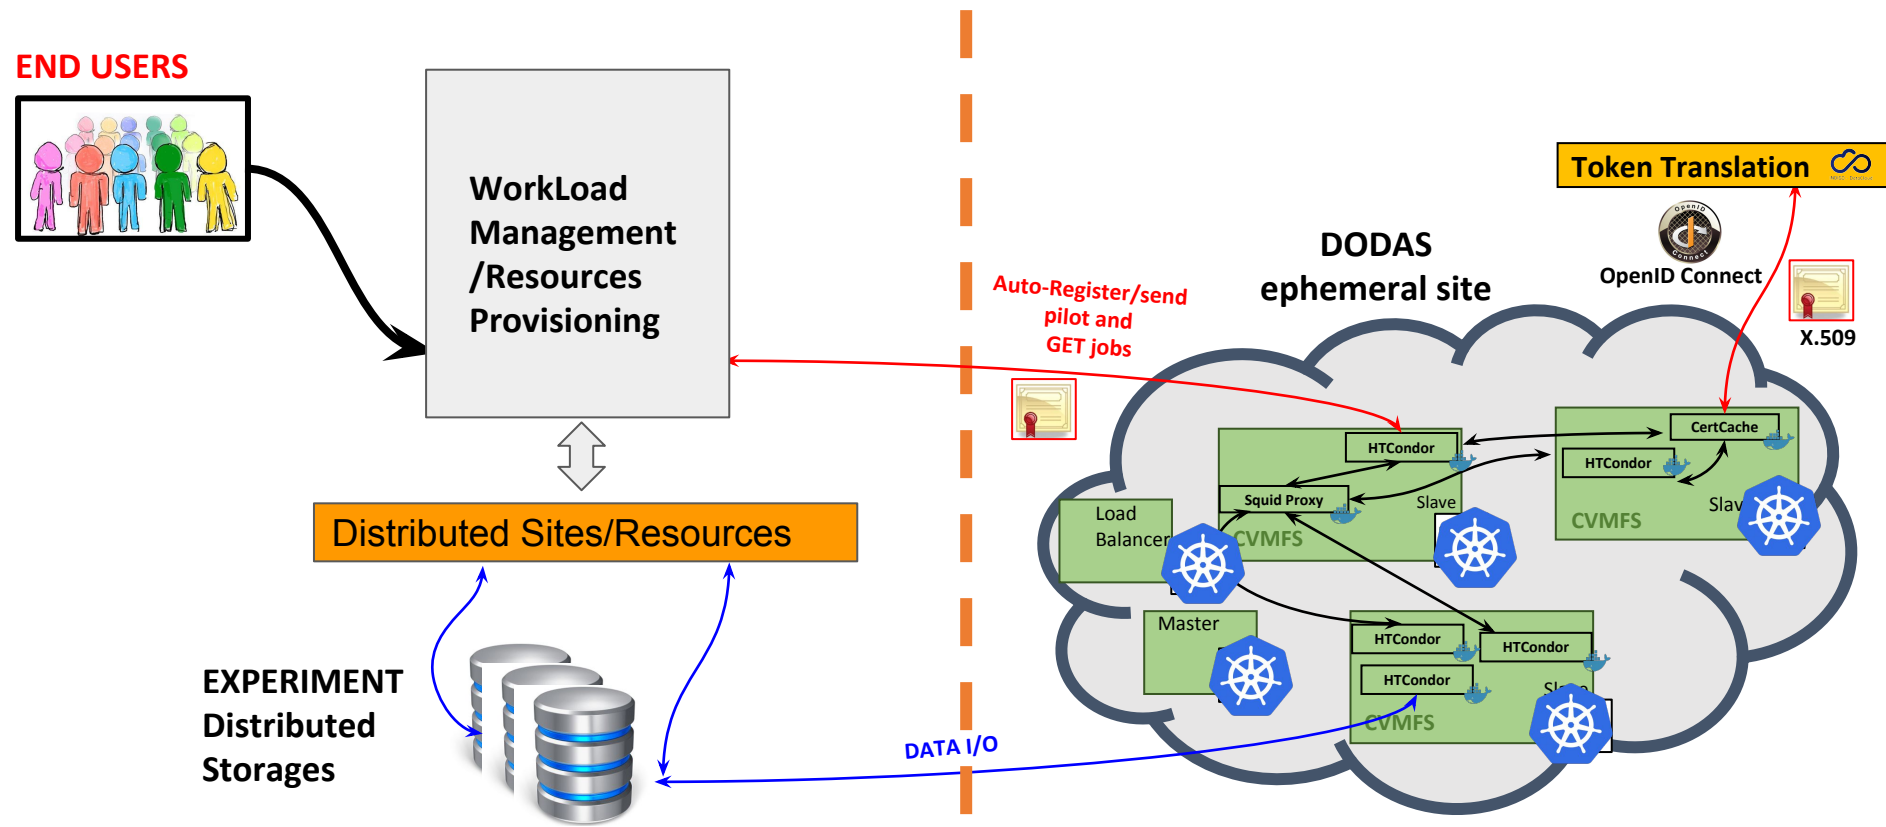
DATA I/O

- ✓ Completely transparent to CMS physicists
- ✓ Seamlessly integrating the global infrastructure

DODAS ephemeral site



And beyond CMS?



DODAS and Data Caches

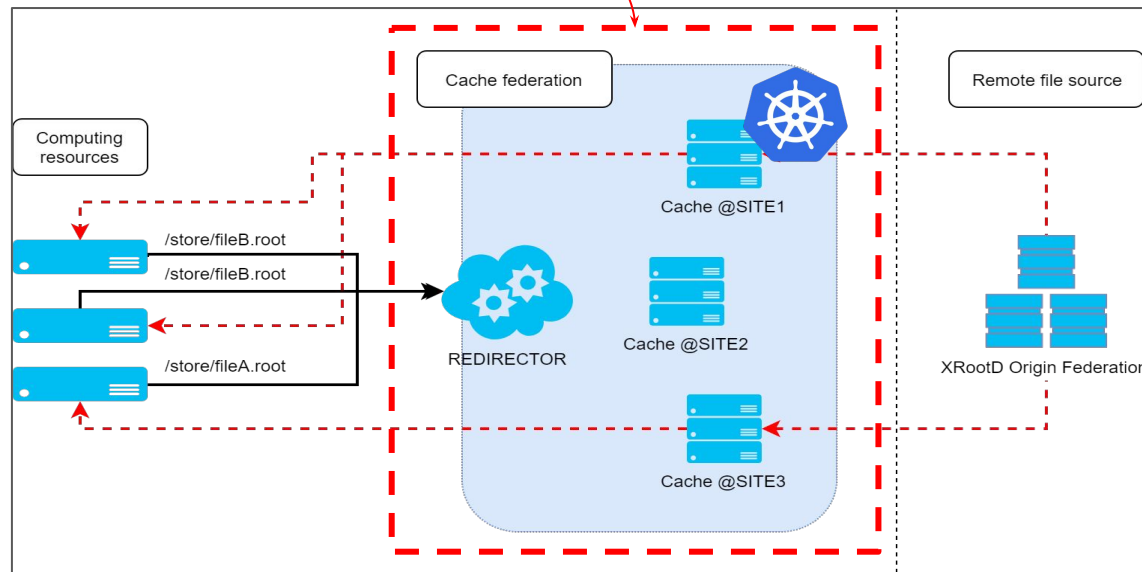
We rely on **XRootD** technology and we support configuration of a variety of services.

- **Data Server:**
- **Redirector:**
- **XCache:**



Cache development
in collaboration with
SoCaL US project

Generated by DODAS

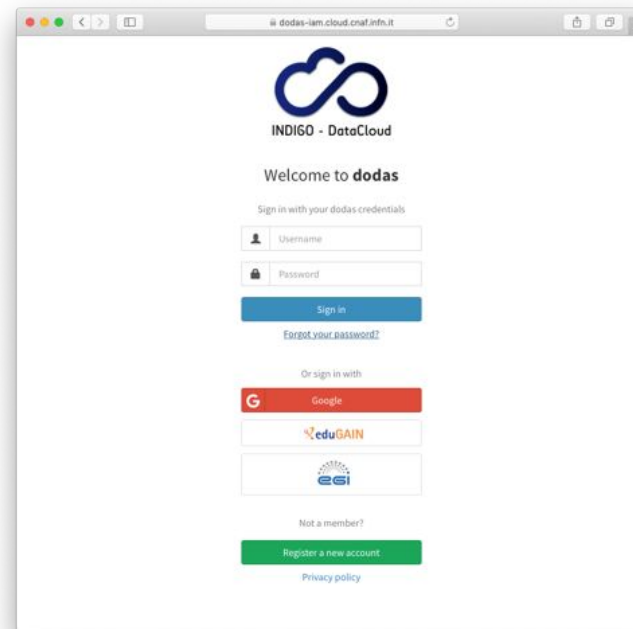


```
$> helm repo add cache https://cloud-pg.github.io/CachingOnDemand/
$> helm repo update
$> helm install cache/cachingondemand
```

DODAS relies on the INDIGO-Identity Access Management to manage Authentication and Authorization

A VO-scoped authentication and authorization service that

- supports multiple authentication mechanisms
- provides users with a persistent, VO-scoped identifier
- exposes identity information, attributes and capabilities to services via JWT tokens and
- standard OAuth & OpenID Connect protocols
- can integrate existing VOMS-aware services
- supports Web and non-Web access, delegation and token renewal



Finally: Not Only CMS



DODAS is also under evaluation (at different level of testing and integration) by **communities other than CMS**

- **AMS Experiment** is already testing/evaluating DODAS to run analysis over opportunistic resources
- **Fermi** analysts are already using (for daily activities) DODAS
- **Virgo** is integrating a pipeline for testing the whole flow



DODAS is a high modular deployer manager built on the concept of Infrastructure as a code to create and provision infrastructure deployments, **automatically and repeatedly**. Today

- We discussed where and how DODAS enter the game in the Cloud Integration
- We saw **stateless sites example on K8s**
 - CMS example and possible generalization
- We shown how DODAS can **provision K8s on demand**
 - Allowing the customisation of the underlying environment (e.g. dependencies) and the definition of the computational stack (DBs, FS, storages/caches)

DODAS also supports an **on-demand analysis facilities** offering (**on top of K8s**):

- HTCondor batch systems on demand, supporting HTCondor federations
 - Floking, routing and HTC/HPC mixing (CHEP2019)
- Spark cluster, TFaaS, etc..

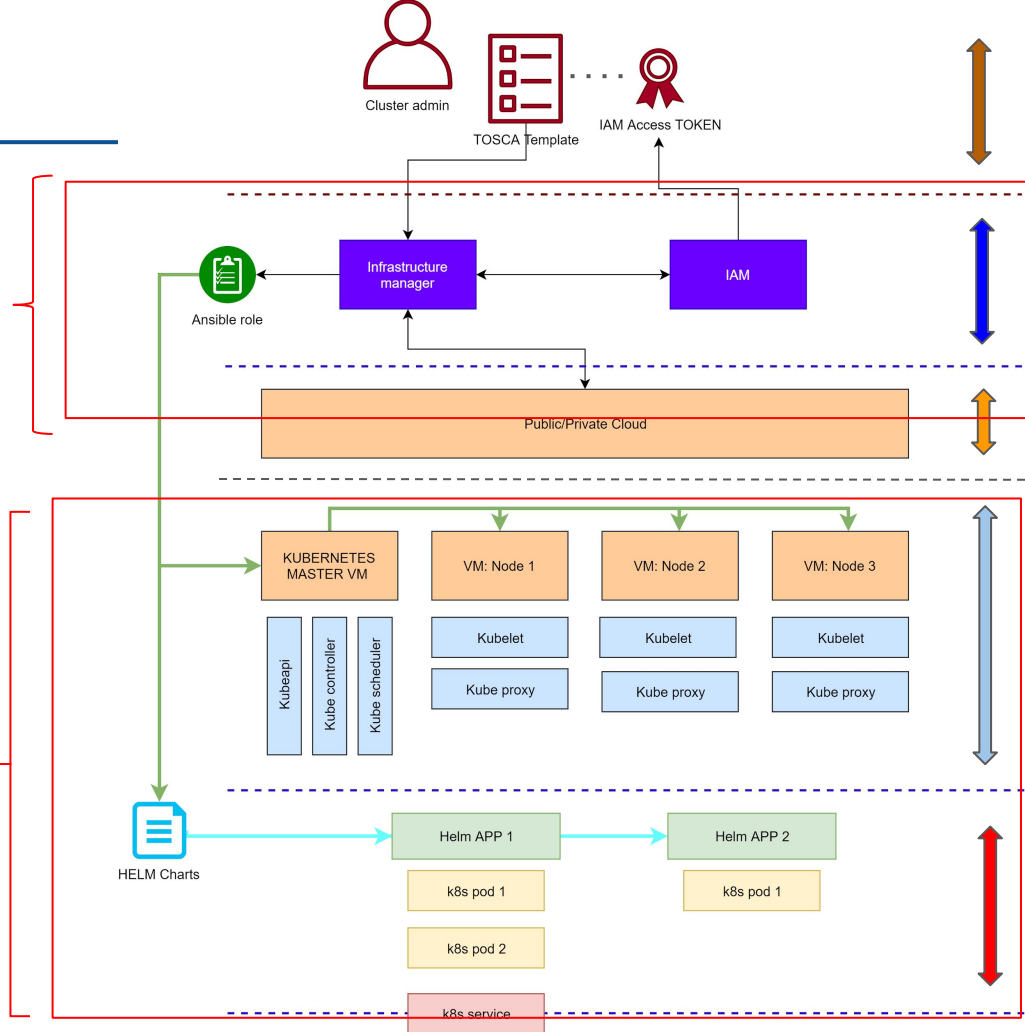
Future work and R&D

DODAS specific

- Improve and evolve the support for bare metal (instead of Cloud API)
- Improve/evolve User interface (GUI/CLI)

K8s oriented:

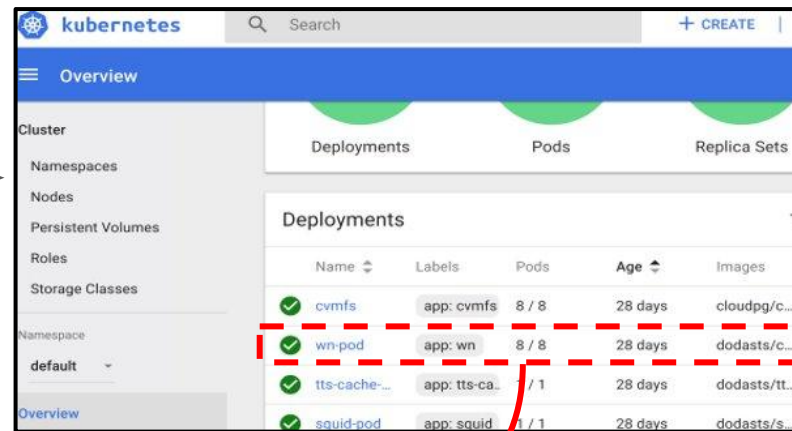
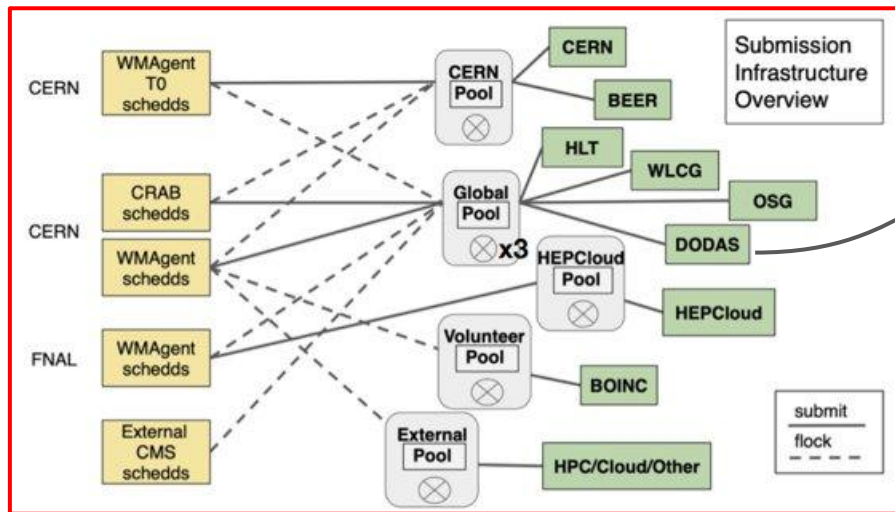
- Autoscalers with custom metrics
- Federated k8s
- Integrated AuthN/Z layer



Backup



Example: The CMS Integration

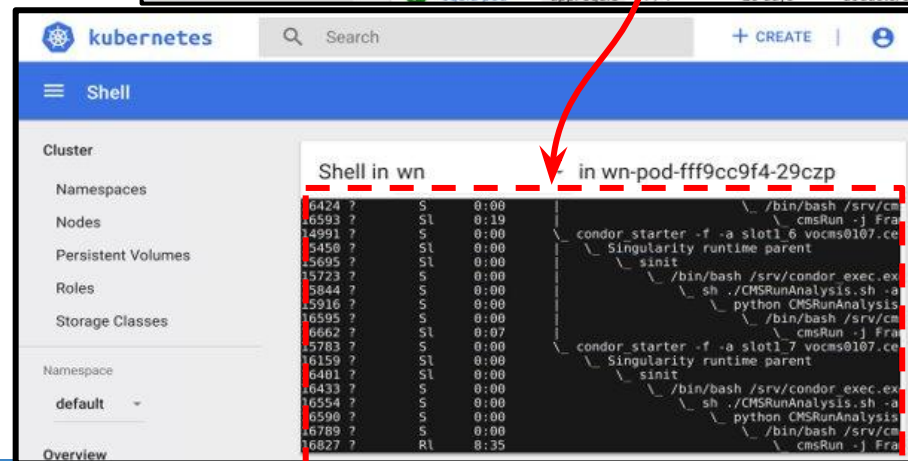


A Key component is the AuthN/Z:

DODAS is based on JWT (INDIGO-IAM). To integrate the CMS GlobalPool:

- start with JWT Token as incoming auth credential
- Implements security via IAM token exchange
- On-demand X.509 certificates generated from IAM tokens (and cached) to enable access to CMS

--> Global Pool authorization is based on DN mapping

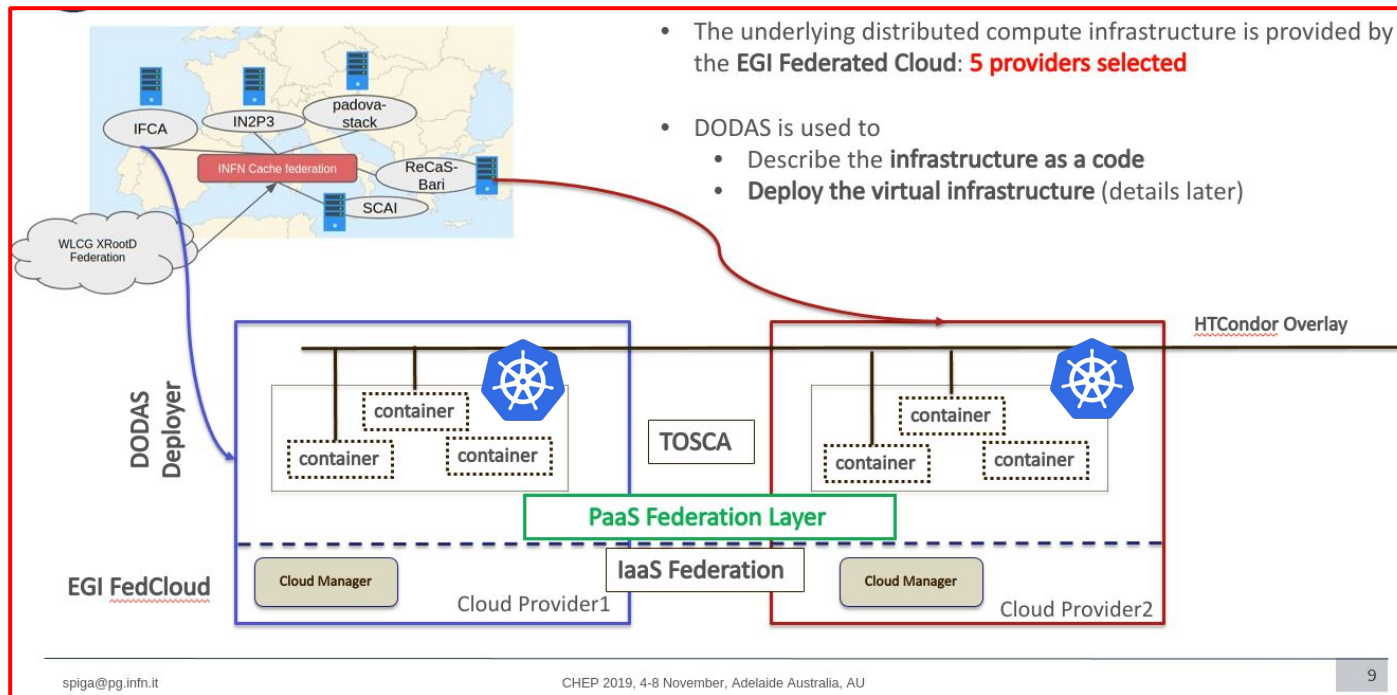


A recent setup

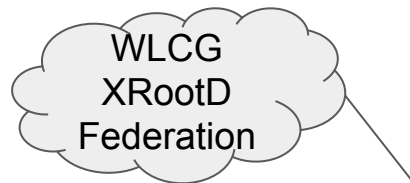


Many tests and deployments in the past two years (see results @ CHEP2018).
Recently we used DODAS to manage **5 stateless sites**... CMS just sees it as a single virtual site, thanks to the **HTCondor overlay**

From CHEP 2019



Putting everything together



Cloud resource provider

Opportunistic Storage Service

Ceph/HDFS/IOVolumes/?

Opportunistic Cache Service

Xcach Xcach Xcach
e e e

Redirector

WN WN WN

Opportunistic CMS startd Service

Generated by DODAS

Cache Network IN

2.8Gbps network inbound

Cache Network OUT

2.8Gbps network outbound

GitHub, Inc. [US] | <https://github.com/Cloud-PG/ansible-role-helm/blob/master/tasks/kube.yml>

23 lines (18 sloc) | 697 Bytes

Raw

Blame

History

```

1 ---
2 - name: Helm install cloudpg repo
3   command: helm repo add {{ item.name }} {{ item.url }}
4   with_items: "{{ repos }}"
5
6 # - name: Helm install cloudpg repo
7 #   command: helm repo add cloudpg https://cloud-pg.github.io/charts/
8
9 # - name: Helm install cache repos
10 #   command: helm repo add cache https://cloud-pg.github.io/CachingOnDemand/
11
12 - name: write values
13   get_url:
14     url: "{{ values_file }}"
15     dest: /tmp/values_{{ name }}-template.yml
16
17 - name: compile values
18   template:
19     src: /tmp/values_{{ name }}-template.yml
20     dest: /tmp/values_{{ name }}.yml
21
22 - name: Helm install chart {{ chart }}
23   command: "helm install --name {{ name }} -f /tmp/values_{{ name }}.yml {{ chart }}"

```

© 2019 GitHub, Inc. [Terms](#) [Privacy](#) [Security](#) [Status](#) [Help](#)

[Contact GitHub](#) [Pricing](#) [API](#) [Training](#) [Blog](#) [About](#)

Checkpoint.pdf hbr.pdf Formazione.pdf confirmation.pdf delunCinara.txt PIC Dashboard icon