

EMBL-EBI Cloud Portal for Everyone

Portal to clouds

David Yuan, Ph.D.

Cloud Bioinformatics Application Architect

Technology and Science Integration

European Bioinformatics Institute, EMBL

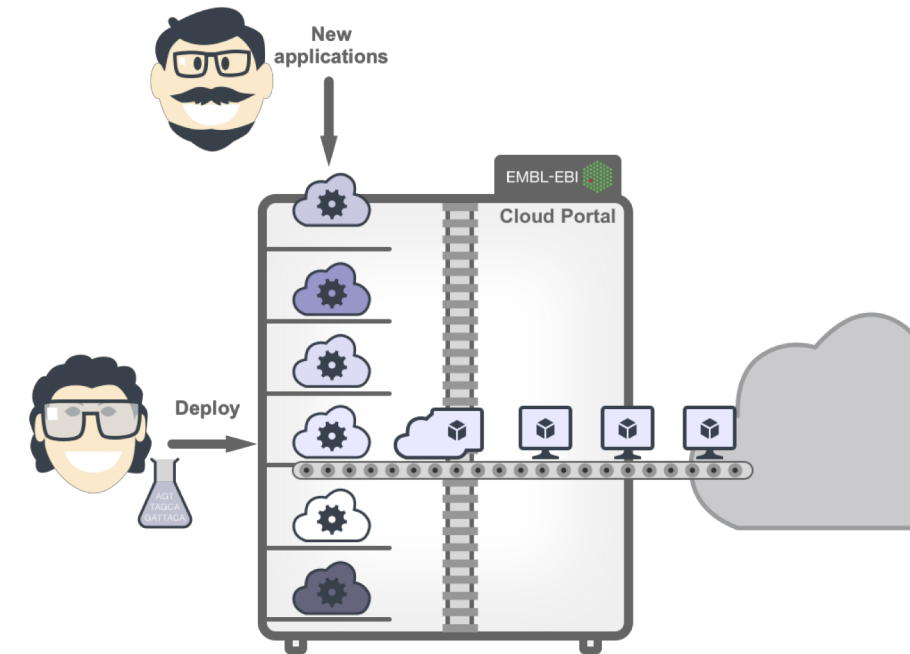
Technology and Science Integration

- European Bioinformatics Institute
 - Outstation of European Molecular Biology Laboratory
- Technology and Science Integration (TSI)
 - Led by Steven Newhouse
 - Cloud Consultancy Team – Tony Wildish
 - Software Development and Operations – Sarah Butcher
 - Project manager – Susheel Varma

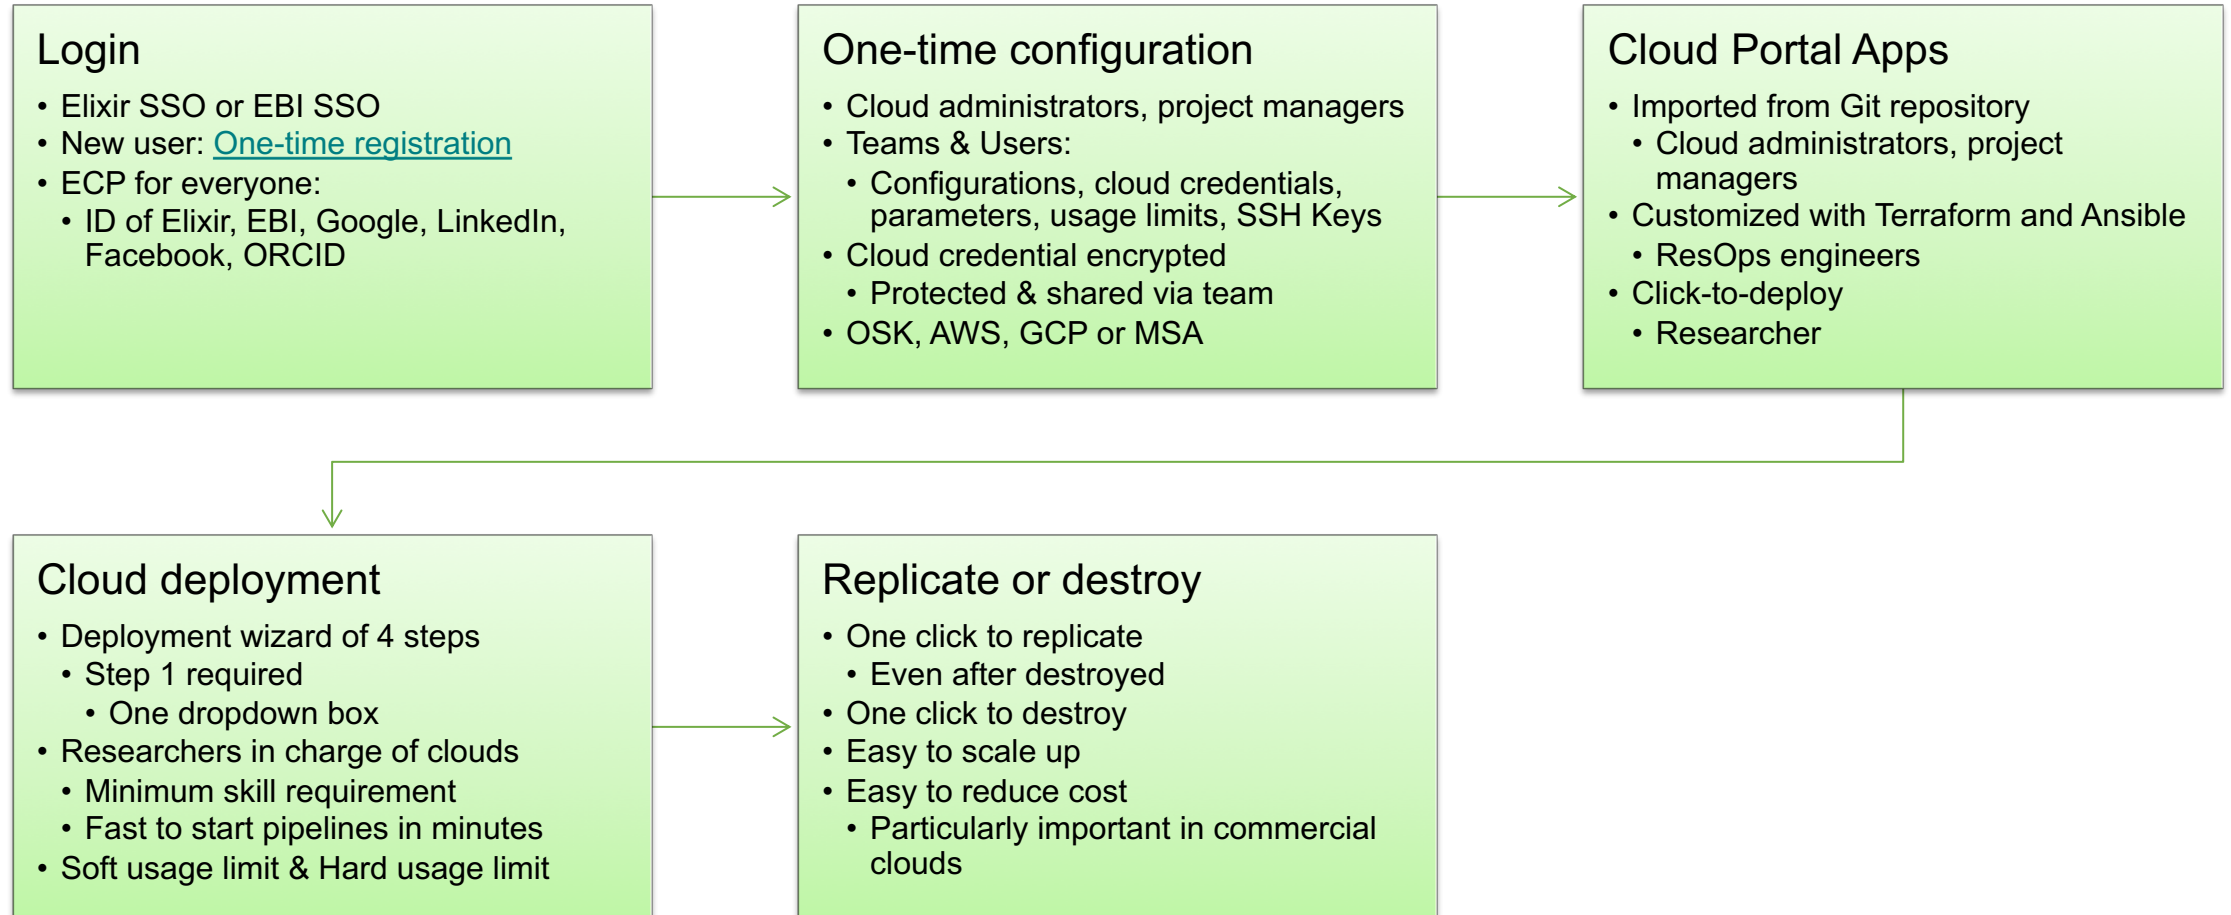


EMBL-EBI Cloud Portal (ECP)

- General consensus to move scientific pipelines to cloud
 - Which cloud?
 - OpenStack (OSK), Google (GCP), Amazon (AWS), Microsoft (MSA)
 - How to get there?
 - Very steep learning curve
 - Beyond one cloud?
 - Vendor lock-in – very difficult and expensive
 - Big data in Bioinformatics getting bigger!
 - Bring compute to data
- Cloud portal with common GUI, CLI & API and good usability
 - ECP for user-focused hybrid cloud strategy



EMBL-EBI Cloud Portal workflow



ECP: a portal to clouds for everyone

Researcher

- Use GUI on laptop or smart phone
 - [No programming](#)
 - [Same wizard](#) for multi-clouds
- [Click to deploy](#)
 - Infrastructure ready in minutes
- Run pipeline in favorite cloud
- [Click to duplicate](#)
 - Clone instances to scale up
- [Click to destroy](#)
 - Pay for actual use

Cloud administrator

- Use GUI and CLI
 - [Convenience and flexibility](#)
- Code in [Bash and Python](#)
 - Well known
 - Same script for multi-clouds
- [Click to import](#)
 - CPAs with Git ULRs
- [Click to define](#)
 - Configurations, deployment parameters & cloud credentials

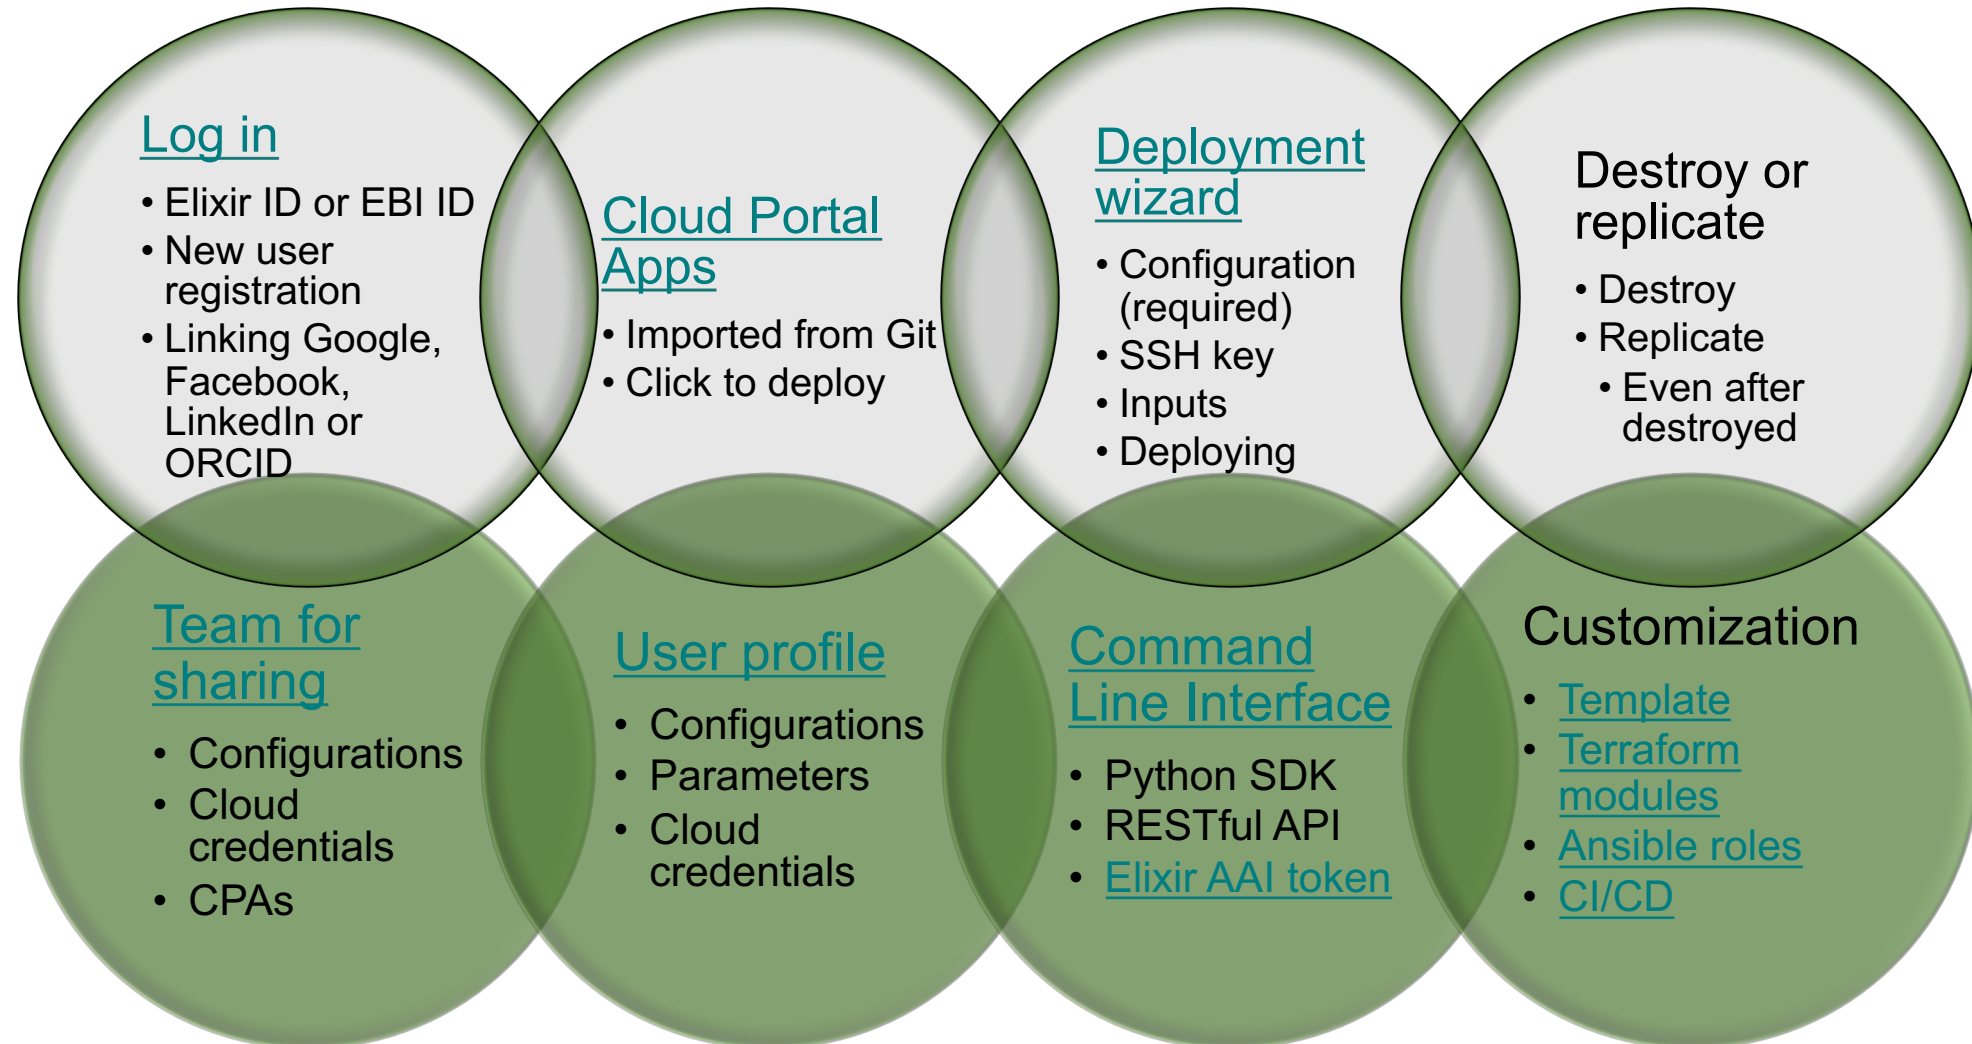
Project manager

- Create [project teams](#)
- Share for [consistency and usability](#)
 - CPAs, configurations
- Share for better [security](#)
 - Cloud credentials
- Reduce [cost](#)
 - Resource quota

ResOps engineer

- Customize CPAs
 - [Rich set available](#)
 - [From template or existing CPAs](#)
- Code in [Terraform and Ansible](#)
 - Well known
 - [Many modules and roles](#)
- Write same program for multi-clouds
 - [Terraform import](#)

ECP Live demo



Summary

- EMBL-EBI Cloud Portal is designed for everyone
 - Researchers, cloud administrators, project managers & ResOps engineers
 - Elixir AAI and EBI AAP for entire scientific community
 - Lower friction for accessing clouds for all groups of users
- ECP can help adopt a user-focussed hybrid cloud strategy
 - Same interfaces and scripts for multiple clouds to avoid vendor lock-in
 - Team sharing for consistency, usability and better security
 - Encrypted cloud credentials
 - Reusable by design: CPA catalogue, Terraform modules & Ansible roles
 - Cost reduction: replicating or destroying as needed, together with usage quota

Acknowledgement

- [EMBL-EBI](#) (European Molecular Biology Laboratory - European Bioinformatics Institute).
- [Cancer Research UK](#) funding to develop infrastructure to support the PanCancer collaboration.
- EC funded [BioExcel](#) project.
- EC funded [PhenomMeNal](#) project.
- [UKRI Cloud Working Group](#) for CI/CD toolchain with Kubernetes on GKE.

EMBL-EBI Cloud Portal for Everyone

Backup

David Yuan, Ph.D.

Cloud Bioinformatics Application Architect

Technology and Science Integration

European Bioinformatics Institute, EMBL

EMBL-EBI Cloud Portal (<https://cloud-portal.ebi.ac.uk/>)

EMBL-EBI Cloud Portal Sign On

A platform for better science


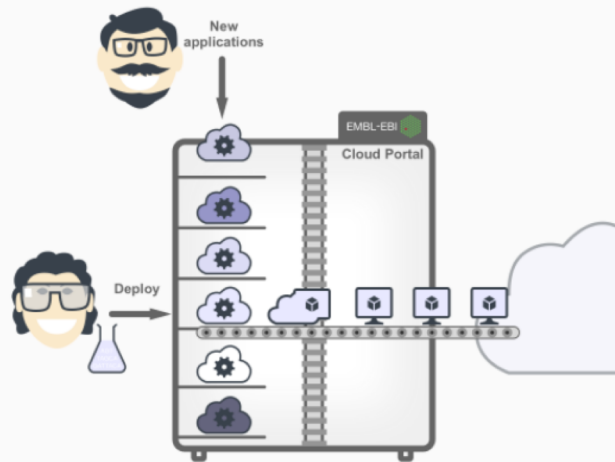
The EMBL-EBI Cloud Portal is platform for *cloud-based data analysis* that simplifies the experience for researchers when running their workloads across different cloud providers.

The portal enables cloud providers and research teams to package and share their key applications and data environments with their respective users. That way, these key applications can be deployed on demand onto the cloud providers the user is able to access.

At a glance:

- Deploy *domain-specific* applications and infrastructure, from your own repos or from the *EBI Cloud Portal community*.
- Use your own or *shared* cloud resources.
- Scale when you need to, securely.

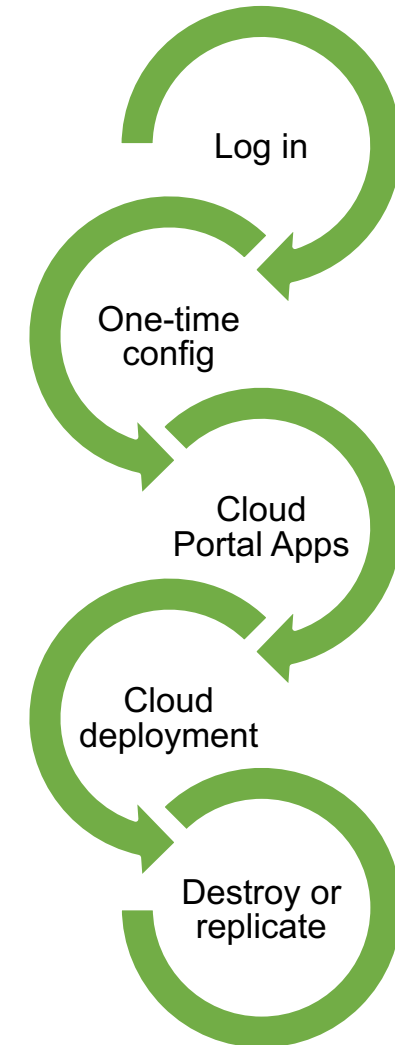
The EBI Cloud Portal has been developed by the [Technology and Science Integration](#) team at EMBL-EBI.



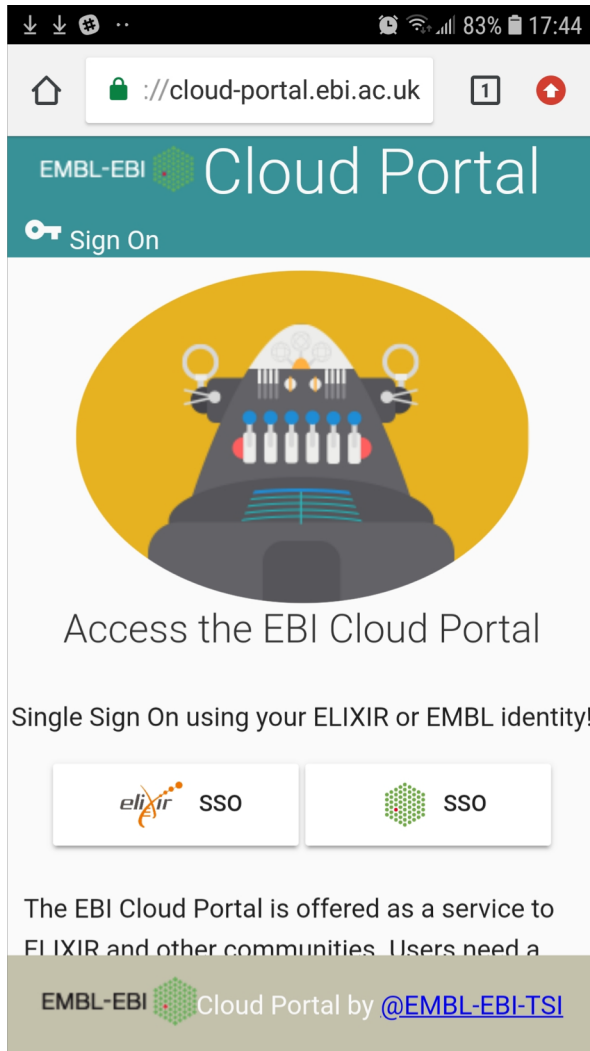
How to access to the EMBL-EBI Cloud Portal?

The EMBL-EBI Cloud Portal is offered as a service to ELIXIR and other communities. Users will need a freely available [ELIXIR identity](#) to access the portal.

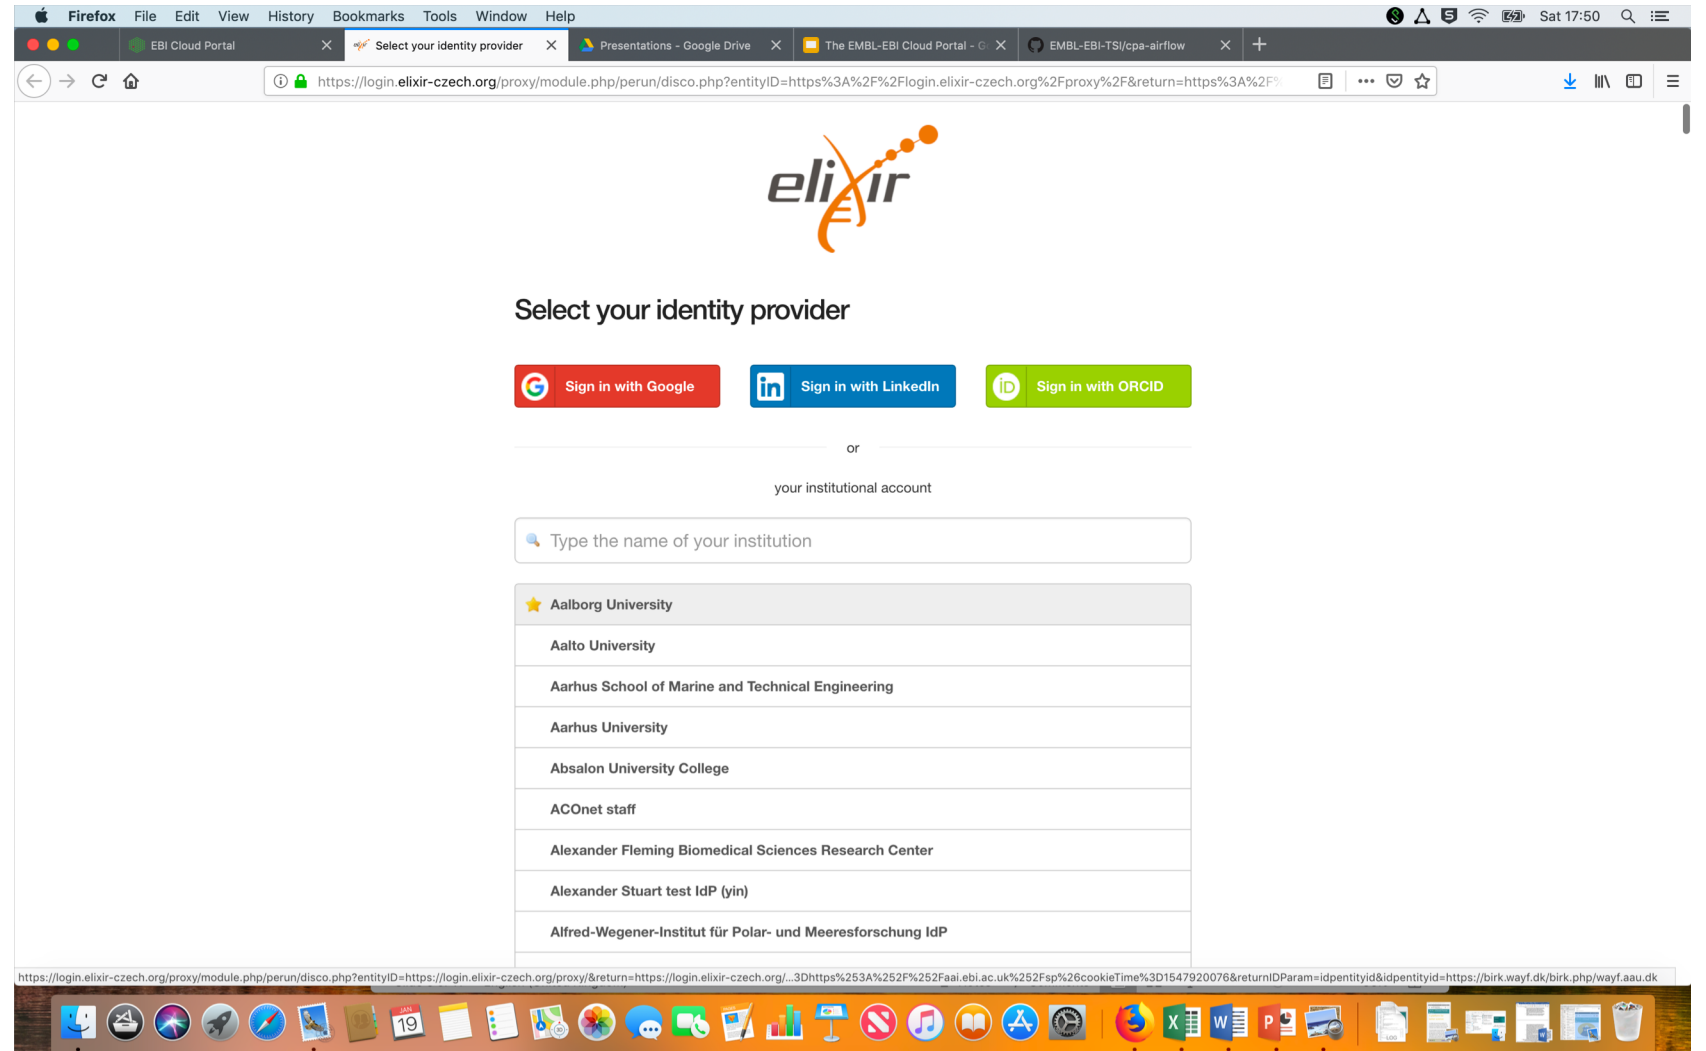
Obtain your ELIXIR identity [here](#). Then you can [login](#).



ECP login with Elixir or EBI ID

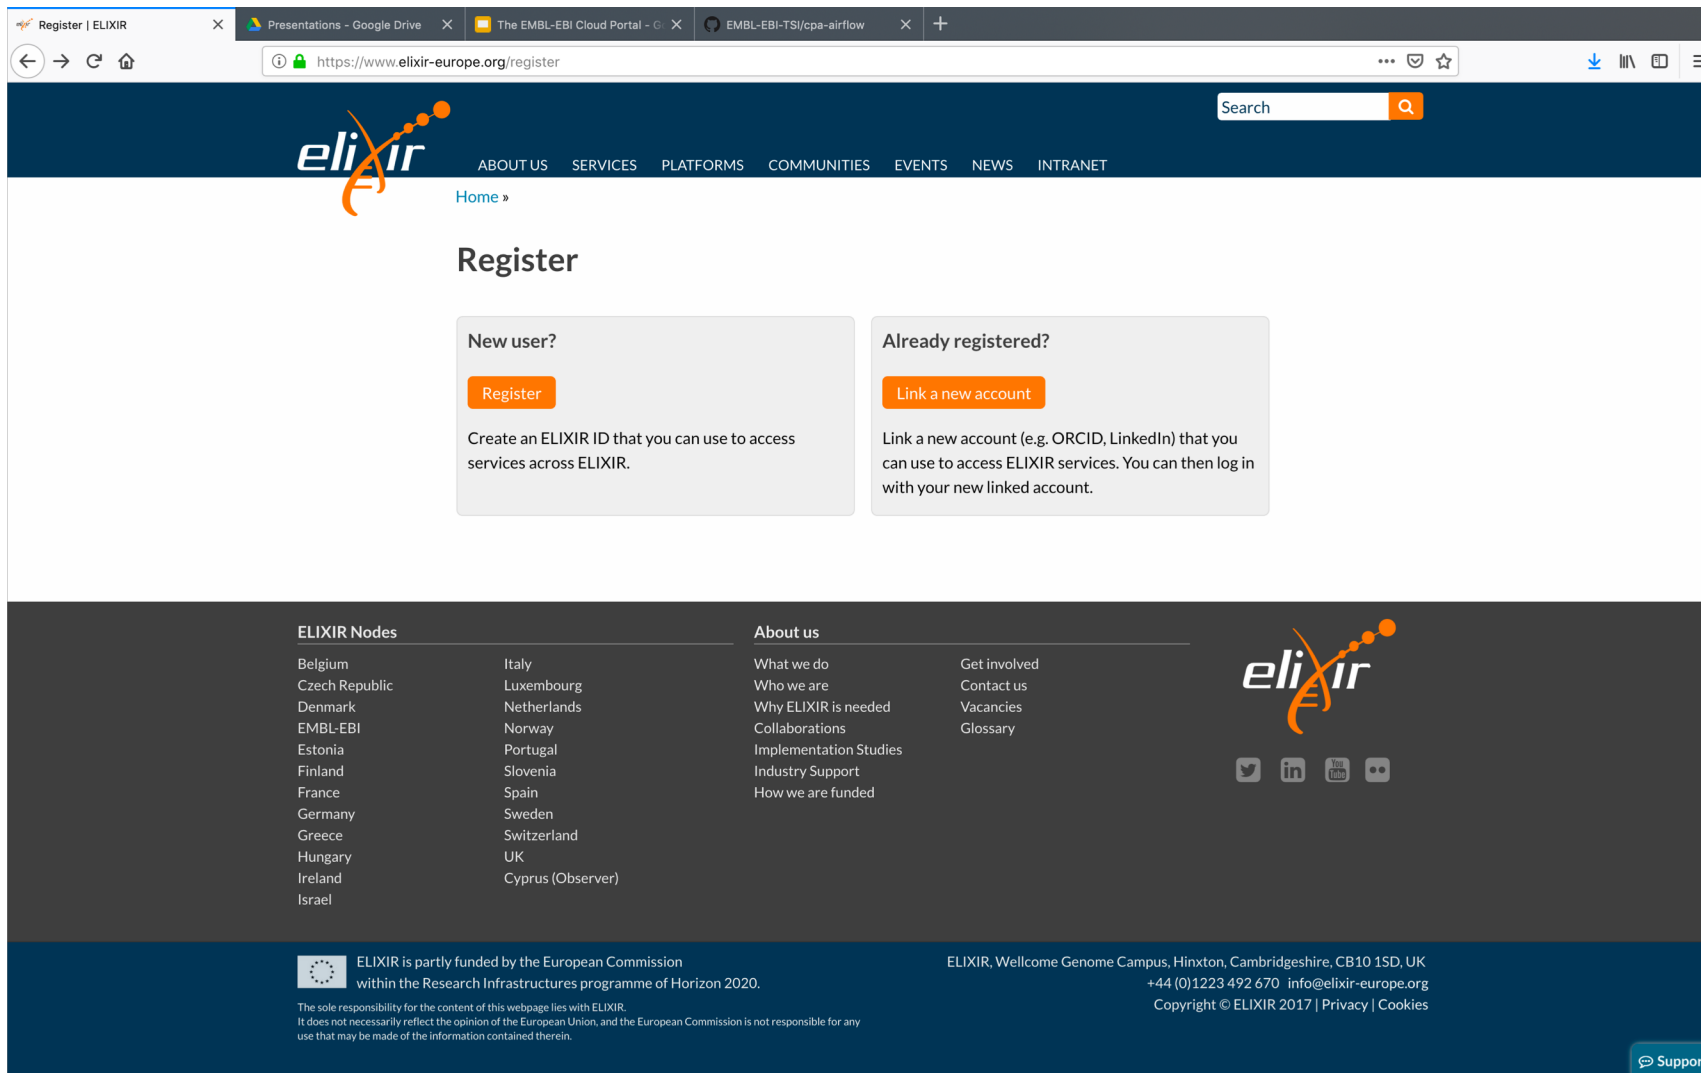


The screenshot shows the EMBL-EBI Cloud Portal homepage. At the top, there is a navigation bar with the EMBL-EBI logo and the text "Cloud Portal". Below this, there is a "Sign On" button. The main content area features a large circular graphic of a microscope. Below the graphic, the text reads "Access the EBI Cloud Portal" and "Single Sign On using your ELIXIR or EMBL identity!". At the bottom, there are two buttons: "elixir SSO" and "EMBL SSO". The footer contains the text "EMBL-EBI Cloud Portal by @EMBL-EBI-TSI".



The screenshot shows the Elixir login page. At the top, there is the Elixir logo. Below it, the text reads "Select your identity provider". There are three buttons: "Sign in with Google", "Sign in with LinkedIn", and "Sign in with ORCID". Below these buttons, there is a section for "your institutional account" with a search box labeled "Type the name of your institution". A list of institutions is displayed below the search box, including Aalborg University, Aalto University, Aarhus School of Marine and Technical Engineering, Aarhus University, Absalon University College, ACOnet staff, Alexander Fleming Biomedical Sciences Research Center, Alexander Stuart test IdP (yin), and Alfred-Wegener-Institut für Polar- und Meeresforschung IdP.

Registration for everyone



The screenshot shows the ELIXIR registration page in a browser. The URL is https://www.elixir-europe.org/register. The page features the ELIXIR logo and a navigation menu with links for ABOUT US, SERVICES, PLATFORMS, COMMUNITIES, EVENTS, NEWS, and INTRANET. A search bar is located in the top right. The main content area is titled 'Register' and contains two columns of options:

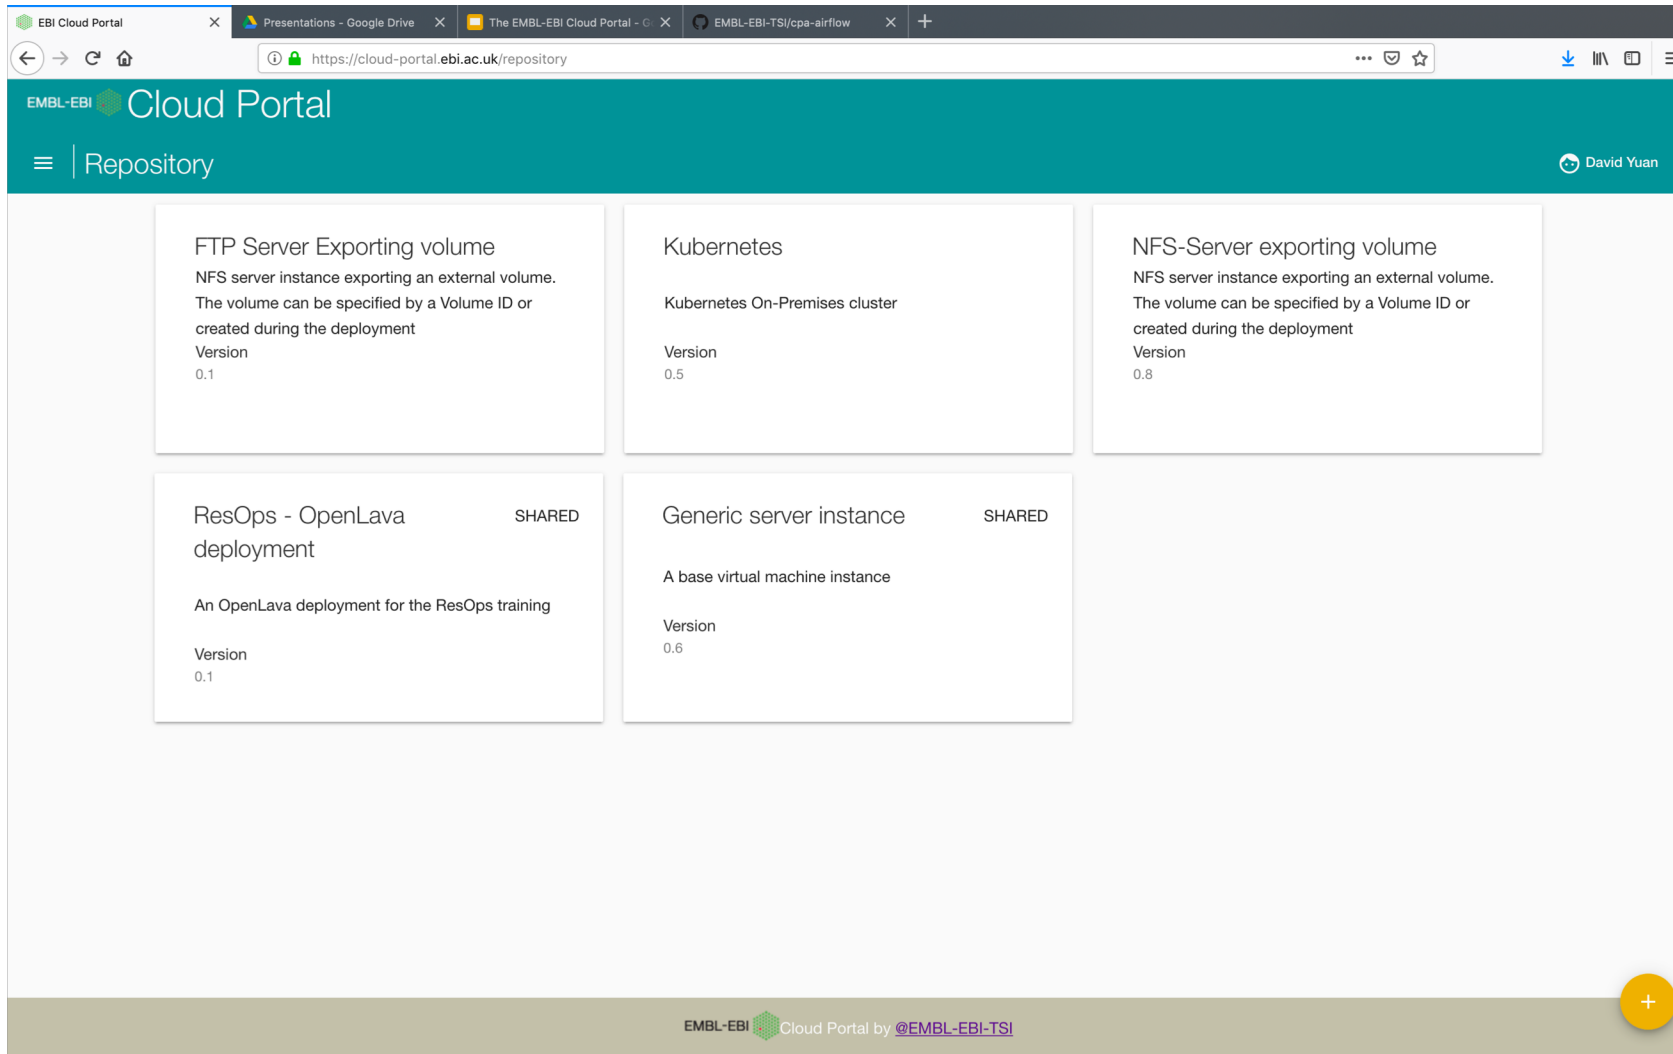
- New user?**
 - [Register](#)
 - Create an ELIXIR ID that you can use to access services across ELIXIR.
- Already registered?**
 - [Link a new account](#)
 - Link a new account (e.g. ORCID, LinkedIn) that you can use to access ELIXIR services. You can then log in with your new linked account.

The footer contains a list of ELIXIR Nodes (Belgium, Czech Republic, Denmark, EMBL-EBI, Estonia, Finland, France, Germany, Greece, Hungary, Ireland, Israel, Italy, Luxembourg, Netherlands, Norway, Portugal, Slovenia, Spain, Sweden, Switzerland, UK, Cyprus (Observer)), an 'About us' section (What we do, Who we are, Why ELIXIR is needed, Collaborations, Implementation Studies, Industry Support, How we are funded), and contact information (ELIXIR, Wellcome Genome Campus, Hinxton, Cambridgeshire, CB10 1SD, UK; +44 (0)1223 492 670; info@elixir-europe.org; Copyright © ELIXIR 2017 | Privacy | Cookies). There is also a 'Support' button in the bottom right corner.

Elixir ID

- 540 institution IDs
- Personal ID
 - Anyone in scientific community
- Linking ORCID, Google ID, Facebook ID or LinkedIn ID

Repository of Cloud Portal Apps



The screenshot shows a web browser window displaying the EMBL-EBI Cloud Portal repository. The page has a teal header with the EMBL-EBI logo and the text "Cloud Portal". Below the header, there is a navigation menu with "Repository" selected. The main content area displays five application cards arranged in two rows. Each card contains the application name, a brief description, and the version number. A yellow plus button is visible in the bottom right corner of the page.

Application Name	Description	Version	Shared
FTP Server Exporting volume	NFS server instance exporting an external volume. The volume can be specified by a Volume ID or created during the deployment	0.1	
Kubernetes	Kubernetes On-Premises cluster	0.5	
NFS-Server exporting volume	NFS server instance exporting an external volume. The volume can be specified by a Volume ID or created during the deployment	0.8	
ResOps - OpenLava deployment	An OpenLava deployment for the ResOps training	0.1	SHARED
Generic server instance	A base virtual machine instance	0.6	SHARED

CPA repository

- Researcher:
 - Click to deploy
- Cloud administrator:
 - Import a CPA
 - “+” button on bottom right
 - URL to a Git repository
- Project manager:
 - Share CPAs in a team

Deployment wizard (only step 1 required)

EMBL-EBI Cloud Portal

Repository | Metagenomics Pipeline

Dario Vianello

Metagenomics Pipeline

OSTACK GCP

1 Select configuration

2 SSH Key Optional

3 Inputs

4 Deploy

Your application is now ready to deploy.

Back DEPLOY

CPA deployed into the cloud

The image shows two overlapping browser windows. The left window displays the EBI Cloud Portal with a 'Deployments' section for Kubernetes. The right window shows the OpenStack Dashboard 'Instances' page with a table of three active instances.

EMBL-EBI Cloud Portal

Deployments

Hide destroyed

Kubernetes

Reference: TSI1545212780801 State: RUNNING

RED HAT OPENSTACK PLATFORM | Project | Identity

Compute | Network | Orchestration

Overview | **Instances** | Volumes | Images | Key Pairs | API Access

Project / Compute / Instances

Instances

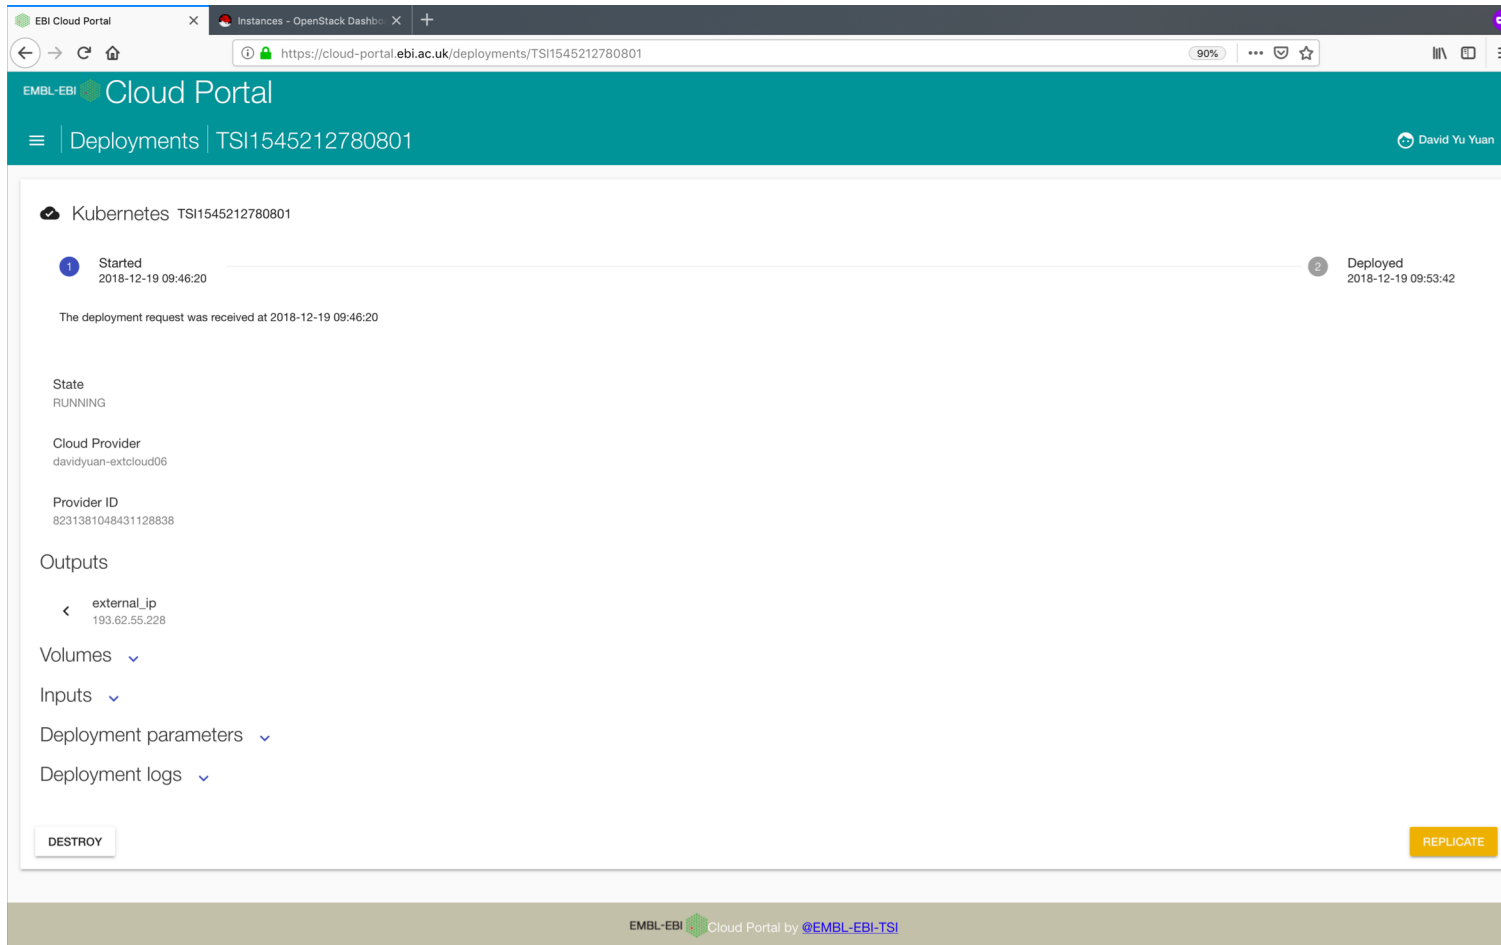
Instance Name: tsi1545212780801-k8s Filter Launch Instance Delete Instances More Actions

Displaying 3 items

<input type="checkbox"/>	Instance Name	Image Name	IP Address	Flavour	Key Pair	Status	Availability Zone	Task	Power State	Time since created	Actions
<input type="checkbox"/>	tsi1545212780801-k8s-s-node-2	centos7	192.168.0.6	s1.small	tsi1545212780801-keypair	Active	AZ_1	None	Running	1 month	Create Snapshot
<input type="checkbox"/>	tsi1545212780801-k8s-s-node-1	centos7	192.168.0.18	s1.small	tsi1545212780801-keypair	Active	AZ_1	None	Running	1 month	Create Snapshot
<input type="checkbox"/>	tsi1545212780801-k8s-s-master	centos7	192.168.0.16 Floating IPs: 193.62.55.228	s1.small	tsi1545212780801-keypair	Active	AZ_1	None	Running	1 month	Create Snapshot

Displaying 3 items

Replicating or destroying cloud instances



The screenshot shows the EMBL-EBI Cloud Portal interface. The browser address bar displays <https://cloud-portal.embl.ac.uk/deployments/TSI1545212780801>. The page title is "EMBL-EBI Cloud Portal" and the breadcrumb is "Deployments | TSI1545212780801". The user profile "David Yu Yuan" is visible in the top right.

The main content area shows a Kubernetes deployment for "Kubernetes TSI1545212780801". A progress bar indicates the deployment status: "Started" at 2018-12-19 09:46:20 and "Deployed" at 2018-12-19 09:53:42. Below the progress bar, it states "The deployment request was received at 2018-12-19 09:46:20".

Deployment details include:

- State: RUNNING
- Cloud Provider: davidyuan-extcloud06
- Provider ID: 8231381048431128838

Outputs section shows:

- external_ip: 193.62.55.228

Other sections include Volumes, Inputs, Deployment parameters, and Deployment logs, each with a dropdown arrow.


At the bottom of the deployment card, there are two buttons: "DESTROY" (grey) and "REPLICATE" (orange).

The footer of the page reads "EMBL-EBI Cloud Portal by @EMBL-EBI-TSI".

Pay-as-you-go




- Can be replicated as-is
 - Identical pipelines cloned to scale up
 - Limited by quota of cloud resources only
- Can be destroyed to recover cloud resources
 - Replicate even after destroyed
- One click

Team for sharing

EMBL-EBI  Cloud Portal

☰ | Teams | UKRI David Yu Yuan

UKRI ⓘ + 👤 🗑️

Applications	Configurations	Cloud Providers
 Kubernetes	 davidyuan-extcloud06	 davidyuan-extcloud06

Members

cdtiwari@ebi.ac.uk ✕
wildish@ebi.ac.uk ✕

User profile

The screenshot shows the EMBL-EBI Cloud Portal user profile page for David Yu Yuan. The page is titled "Profile" and shows the user's name, email (davidyuan@ebi.ac.uk), and join date (November 26, 2018). Below the profile information, there are three tabs: "Configurations", "Cloud Credentials", and "Deployment Parameters". The "Configurations" tab is active, showing two configurations: "davidyuan-extcloud06" and "davidyuan-extcloud06-data". Each configuration has three sections: "Cloud provider parameters", "SSH Public Key", and "Parameters". The "SSH Public Key" section for "davidyuan-extcloud06" shows a long string of characters: "ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQDbVJdgtiojWuTmlpa6GzIveaapJKwMb3zG7Y8IWgkhFRzILK+SfRoeg4VigFOMhpDRuMFbK5Ui61XJ/mDGNokFX8Zr4CZ8f+e5nzZWY/w58p5s2g2cbJcpJV249qKmlnNKQJi+qONAolczQN4...". The "Parameters" section for "davidyuan-extcloud06" shows "davidyuan-extcloud06". The "Parameters" section for "davidyuan-extcloud06-data" shows "Parameters". The page footer includes the EMBL-EBI logo and the text "Cloud Portal by @EMBL-EBI-TSI".

User profile

- Multiple configurations, cloud credentials & deployment parameters
 - Private or shared
 - Share buttons on top right
- Configuration = cloud credential + deployment parameter + public key + usage limit
- Cloud credentials are fully encrypted
- Cloud administrators define configurations, cloud credentials & deployment parameters for sharing
- Soft usage limit & hard usage limit to reduce cost

Configurations, cloud credential, deployment parameters

EMBL-EBI Cloud Portal

Profile | davidyuan-extcloud06 | David Yu Yuan

davidyuan-extcloud06

Included in teams
UKRI

Cloud provider
davidyuan-extcloud06

Parameters
davidyuan-extcloud06

SSH Public Key
ssh-rsa AAAAB3NzaC1yc2EAAAADAQABAAQDbVJdgtiojWuT...

Usage limits
Total usage: 1,701.97
Soft limit:
Hard limit:

Associated deployments

Reference	Status	Owner	Running time
TS11545212780801	RUNNING	David Yu Yuan	31 days 12h 25mi...

EMBL-EBI Cloud Portal by @EMBL-EBI-TSI

EMBL-EBI Cloud Portal

Profile | davidyuan-extcloud06 | David Yu Yuan

davidyuan-extcloud06

Provider
OSTACK

Included in teams
UKRI

Parameters

- OS_REGION_NAME
ENCRYPTED VALUE
- OS_PASSWORD
ENCRYPTED VALUE
- OS_AUTH_URL
ENCRYPTED VALUE
- OS_USERNAME
ENCRYPTED VALUE
- OS_USER_DOMAIN_NAME
ENCRYPTED VALUE

EMBL-EBI Cloud Portal by @EMBL-EBI-TSI

EMBL-EBI Cloud Portal

Profile | David Yu Yuan

David Yu Yuan
davidyuan@ebi.ac.uk
Joined November 26, 2018

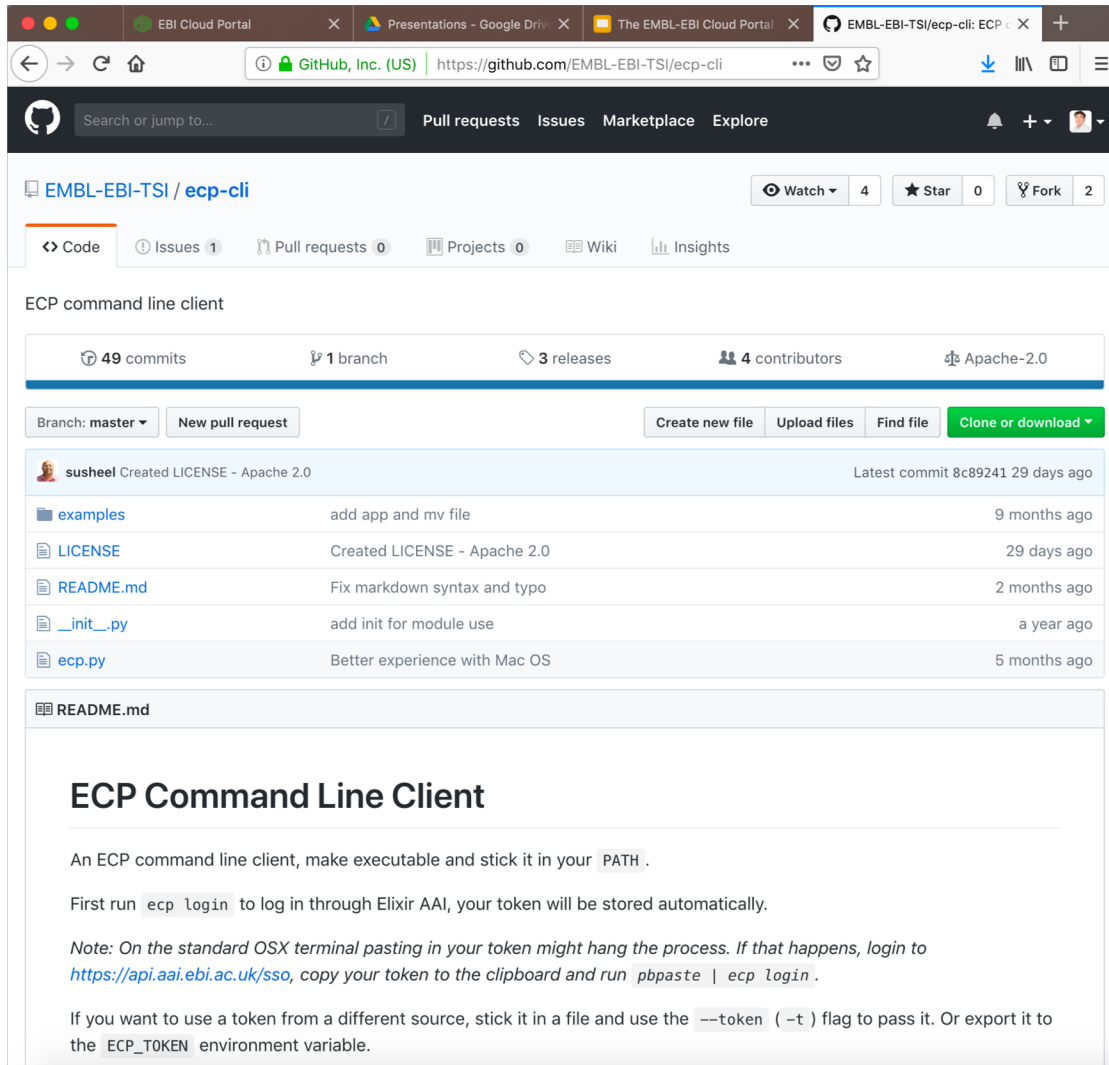
Configurations | Cloud Credentials | **Deployment Parameters**

davidyuan-extcloud06

- disk_image_name
centos7
- floating_ip_pool
ext-net-37
- machine_type
s1.small
- network_name
EBI-TSI-davidyuan_private
- create_volume
true

EMBL-EBI Cloud Portal by @EMBL-EBI-TSI

Command Line Interface for cloud administrators



EMBL-EBI-TSI / ecp-cli

49 commits 1 branch 3 releases 4 contributors Apache-2.0

Branch: master New pull request Create new file Upload files Find file Clone or download

susheel Created LICENSE - Apache 2.0 Latest commit 8c89241 29 days ago

File	Description	Time
examples	add app and mv file	9 months ago
LICENSE	Created LICENSE - Apache 2.0	29 days ago
README.md	Fix markdown syntax and typo	2 months ago
__init__.py	add init for module use	a year ago
ecp.py	Better experience with Mac OS	5 months ago

ECP Command Line Client

An ECP command line client, make executable and stick it in your `PATH`.

First run `ecp login` to log in through Elixir AAI, your token will be stored automatically.

Note: On the standard OSX terminal pasting in your token might hang the process. If that happens, login to <https://api.aai.ebi.ac.uk/sso>, copy your token to the clipboard and run `pbpaste | ecp login`.

If you want to use a token from a different source, stick it in a file and use the `--token (-t)` flag to pass it. Or export it to the `ECP_TOKEN` environment variable.

Same scripts for multiple clouds

- Git repository
 - <https://github.com/EMBL-EBI-TSI/ecp-cli>
- Python SDK
 - Wrapper of ECP RESTful API
- Bash
 - alias `ecp="python ~/IdeaProjects/ecp-cli/ecp.py"`
- CLI login
 - [Elixir AAI](#) token
 - `pbpaste | ecp login`
- CLI commands
 - Verbs: get, create, delete, stop, login
 - Nouns: cred, param, config, app, deployment, logs, destroylogs, status

CPA customization

The screenshot shows a GitHub search interface for repositories containing 'cpa-'. The search results are filtered to show 31 repositories. The top results are:

- cpa-instance**: TSI Cloud Portal - generic server instance. Updated on 2 May 2018.
- cpa-kubernetes**: Kubernetes On-Premises Cluster for the EBI Cloud Portal. Updated on 1 Nov 2018.
- cpa-airflow**: Private repository. Updated on 29 Oct 2018.
- cpa-malware**: A malware application for testing ECP API server vulnerability. Updated 3 days ago.

The interface also includes filters for top languages (HCL, Shell, Java, Python, TypeScript) and most used topics (ga4gh, ebi, tesk, aap, dsds).

Cloud Portal Apps

- Repositories (31 CPAs)
 - <https://github.com/EMBL-EBI-TSI/>
 - Implemented with Terraform & Ansible
- Jump start:
 - [CPA template](#)
 - [Terraform modules](#)
 - [Ansible roles](#)
 - Terraform import
- CI/CD toolchain for cloud deployment
 - Library for CI jobs
 - Library for Python functions
- Manual testing
 - GUI: <https://cloud-portal.ebi.ac.uk/>
 - CLI: <https://github.com/EMBL-EBI-TSI/ecp-cli>