Implementing a complete Workshops-on-Demand infrastructure in less than an hour

HPE Developer Community Meetup

Frederic Passeron (HPE DEV Community)
Bruno Cornec (Former HPE Linux Distinguished Technologist)

17th December 2025

Agenda

Implementing a complete Workshop-on-Demand infrastructure in less than an hour

01	A bit of Background on WoD
02	Live Deployment
03	Architecture and Communications
04	QA

Workshops-on-Demand

A bit of background

WOD Background

In a nutshell, we went from:

- Remote Terminal Server sessions with pdf to Jupyter Notebooks.
- An HPE DEV only solution to an open-source project (validated by HPE OSRB Process).
- A single playbook on a JupyterHub server to a multi tier application.

Project Repository: https://github.com/Workshops-on-Demand

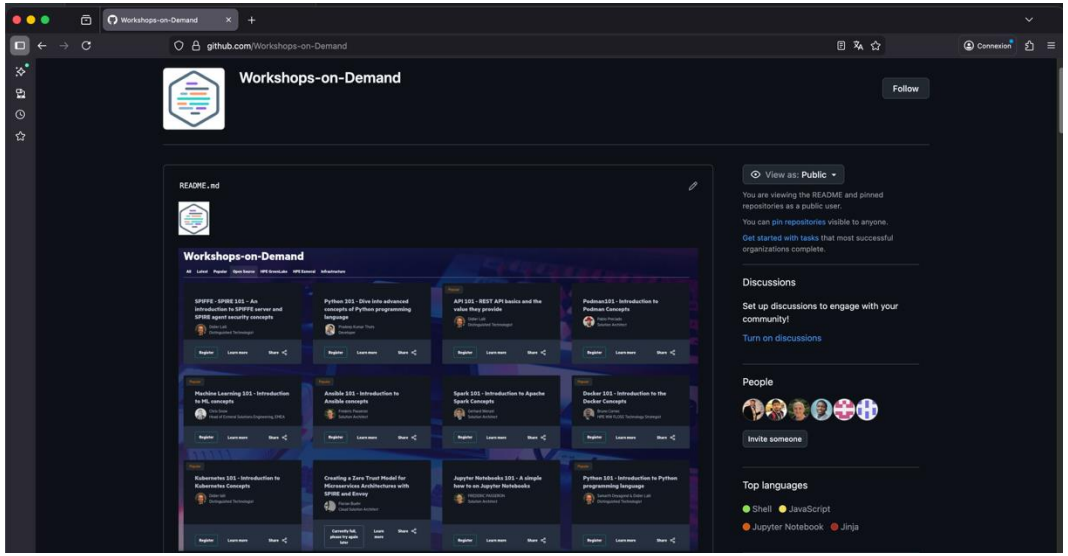

Workshops-on-Demand

Deployment Mechanism

In order to install a full infrastructure, a set of preliminary steps are required:

- Install a VM/Physical machine with Ubuntu 24.04 LTS minimal and default setup. (20.04 or 22.04 should still work, while less tested these days).
- On each VM the user ubuntu is being created and can be used for the initial setup. For access to the machine and account refer to the Ubuntu documentation or how a VM template was set up. Ensure that this user as root access (either via sudo or with password access)
- Then you have to ensure minimal dependencies are met to run the installation. We only need git. All other dependencies are installed by the installer. Issue on each machine: sudo apt install git
- Once installed you can use it to clone the installer wod-install repository) again on each machine: git clone https://github.com/Workshops-on-Demand/wod-install
- Then you use the installer to install your WoD infrastructure: cd wod-install/install ; ./install.sh -h

Reading the example for the full infrastructure at the end of the help message should give you the required guidance to set it up.

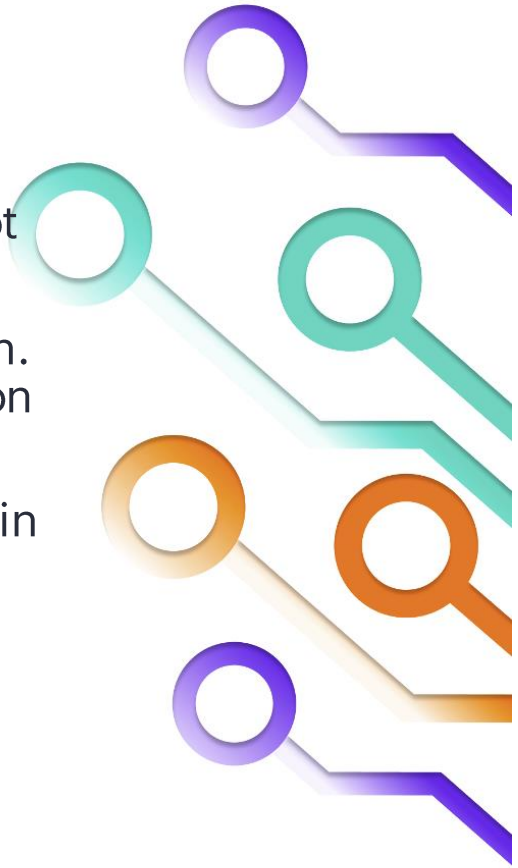

Deployment Mechanism

Full installation example of a stack with:

- **2 backend servers be1 and be2 using port 8010**
- **1 api-db server apidb on port 10000**
- **1 frontend server front on port 8000**
- all declared on the .local network
- internal **postfix server** running on **port 9000**
- **e-mail sender** being **wodmailer@local**
- **ansible groupname** being **test**
- **management user** being **wodmgr**

On the be1 machine:

```
sudo ./install.sh -a apidb.local:10000::https -f front.local:8000:https \
-g test -u wodmgr -p 9000 -s wodmailer@local\
-b be1.local:8010 -n 1 -t backend \
```

On the be2 machine:

```
sudo ./install.sh -a apidb.local:10000:https -f front.local:8000:https \
-g test -u wodmgr -p 9000 -s wodmailer@local\
-b be2.local:8010 -n 2 -t backend \
```

On the apidb machine:

```
sudo ./install.sh -t apidb -b .local:10000 -f front.local:8000 \
-g test -u wodmgr -p 9000 -s wodmailer@local\
-b be1.local:8010,be2.local:8010 -t api-db \
```

On the frontend machine:

```
sudo ./install.sh -a apidb.local:10000:https -f front.local:8000:https \
-g test -u wodmgr -p 9000 -s wodmailer@local\
-t frontend \
```

Workshops-on-Demand

Architecture and Communications

Architecture and Communications

The WoD infrastructure comprises 3 different systems to work, that are usually spread across 3 machines:

•a wod-backend machine, hosting the Jupyter Hub and the WoD templates to generate the real workshop that a given student will run, wit their metadata. This machine may also interact with appliances for WoD needed one, such as Docker e.g. You may have multiple wod-backends in case of a large setup. Corresponding software on the repos wod-backend, wod-notebooks templates, wod-private optional setup

•a wod-api-db machine hosting the WoD API service and a PostgreSQL database to store live information about the running platform. Corresponding software on the repo wod-api-db

•a wod-fronted machine, hosting the Web interface to see the list of available workshops and book one. Corresponding software on the repo wod-frontend

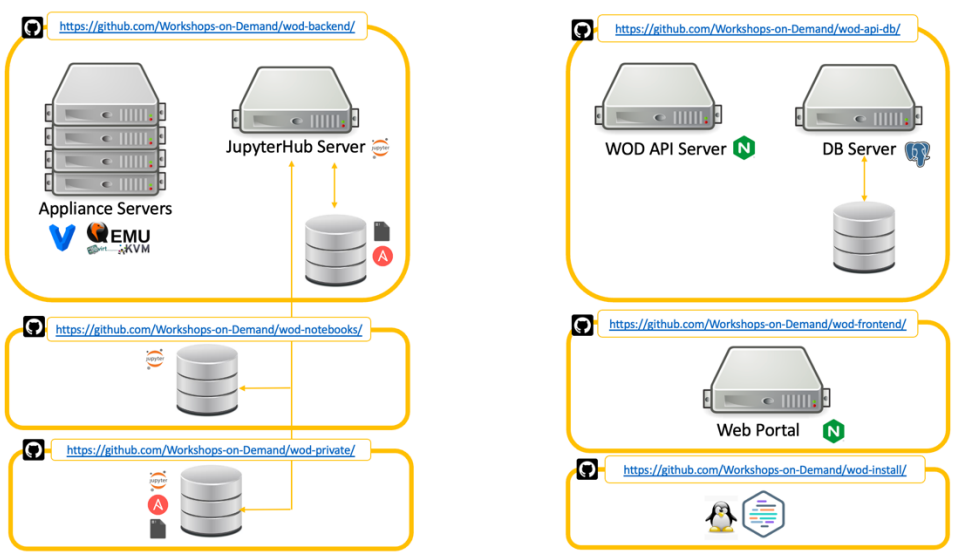

Architecture and Communications

Before a notebook is made available to the user requesting it, there are a certain number of communications happening between machines

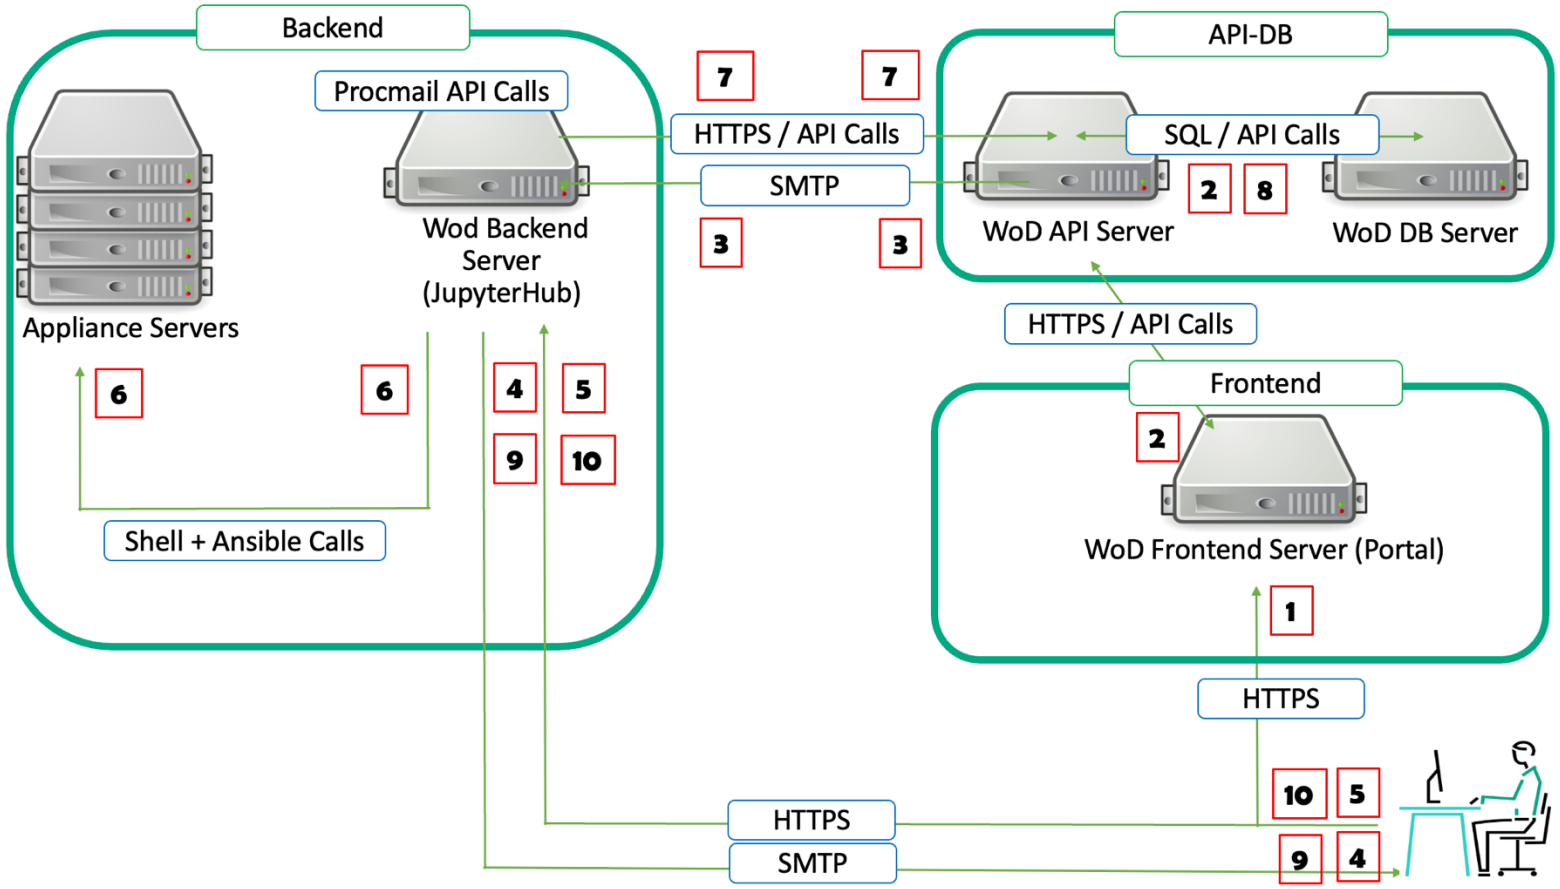

Workshops-on-Demand

QA.

Additional resources

- **Workshops-on-Demand project:**

https://github.com/Workshops-on-Demand

- **Workshops-on-Demand documentation**

Here are a series of blogs published on the HPE Developer Community Portal describing the genesis and the architecture of the project. It also covers the deloyment phase of the different components as well as the creation of content.

0/ From Jupyter Notebooks-as-a-Service to HPE DEV Workshops-on-Demand

1/ Open sourcing Workshops-on-Demand part 1: Why and How

2/ Open Sourcing Workshops-on-Demand part 2: How to Deploy the infrastructure

3/ Open Sourcing Workshops-on-Demand part 3: Understand the Backend

4/ Open Sourcing Workshops-on-Demand part 4: Manage the Backend

5/ Open Sourcing Workshops-on-Demand part 5: Create new workshops

You can also refer to the various presentations made on this project:

LCA 2021

OSXP 2021

JDLL 2024

AlpOSS 2024

Thank You