

Beyond CMS, HTCondor, PaaS... a look to the future challenges

Daniele Spiga
HPC4L - Training. Beirut, Lebanon
21th22th October 2020

Outline

Introduction

DODAS and Big Data:

- Spark use case

Future challenges

- Serverless and Function as a Service

Goal of this last session

1. What we saw these 2 days.
 - a. Cloud, Paas, Automation and services orchestration
 - b. We saw how DODAS implement the infrastructure as a code paradigm
 - c. We saw two main use cases..
 - i. HTCodndor batch system
 - ii. CMS Site

2. What DODAS can do beyond these use cases
 - a. we 'll do the spark example

3. And finally what about evolution in cloud ? ...
 - a. From container / microservices ---> to functions

The Strategy based on a Lego Approach

There is a huge set of tools and solutions available, but there is NOT a one-size-fit-all solution

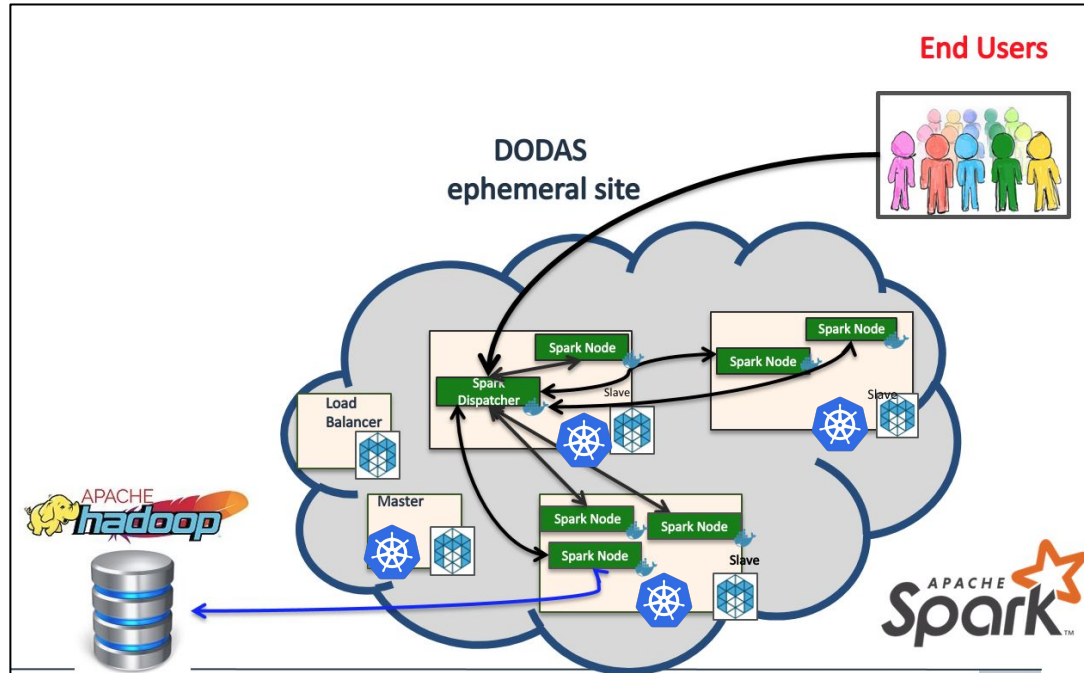
Open, Standard-based, flexible and extensible building blocks

Each use case can compose and customize
and customize



Let's spoil this talk then

Using DODAS to
automatically deploy
Spark on a cloud
environment

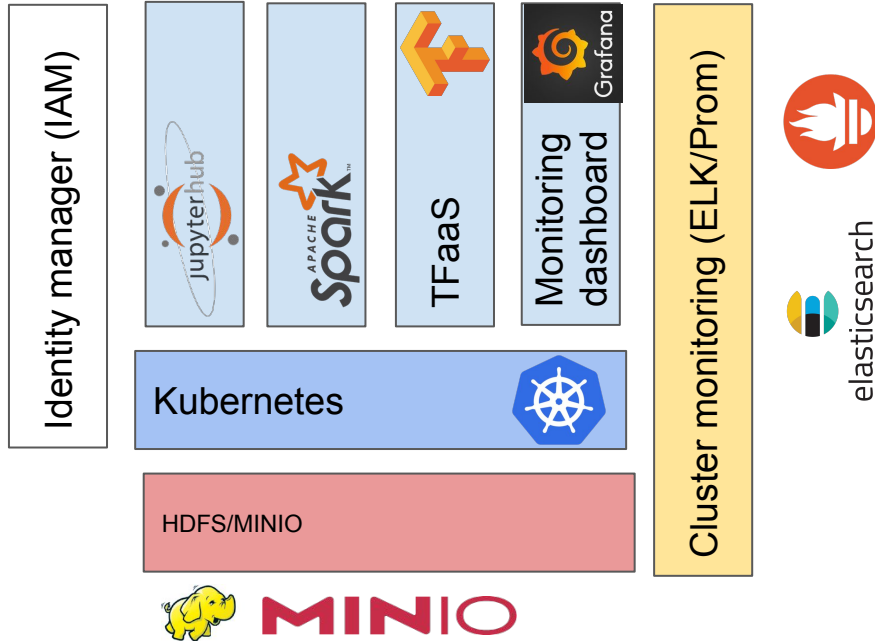


Big data infrastructure

“big data infrastructure entails the tools and agents that collect data, the software systems and physical storage media that store it, the network that transfers it, the application environments that host the analytics tools that analyze it and the backup or archive infrastructure that backs it up after analysis is complete.”



As example: something like this?



Most of what has been discussed this week in term of services and software, components
- plus something I will show in the next...

And how DODAS fits into this ?



DODAS main concepts (cont)



Provides a highly flexible and modular solution to enable several scenarios:

- Orchestrate and build computing stacks, following a “**all in one**” approach
 - From resources provisioning to application setup and management
 - **TOSCA + Ansible + Helm**


- 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)

Identity manager
(IAM)

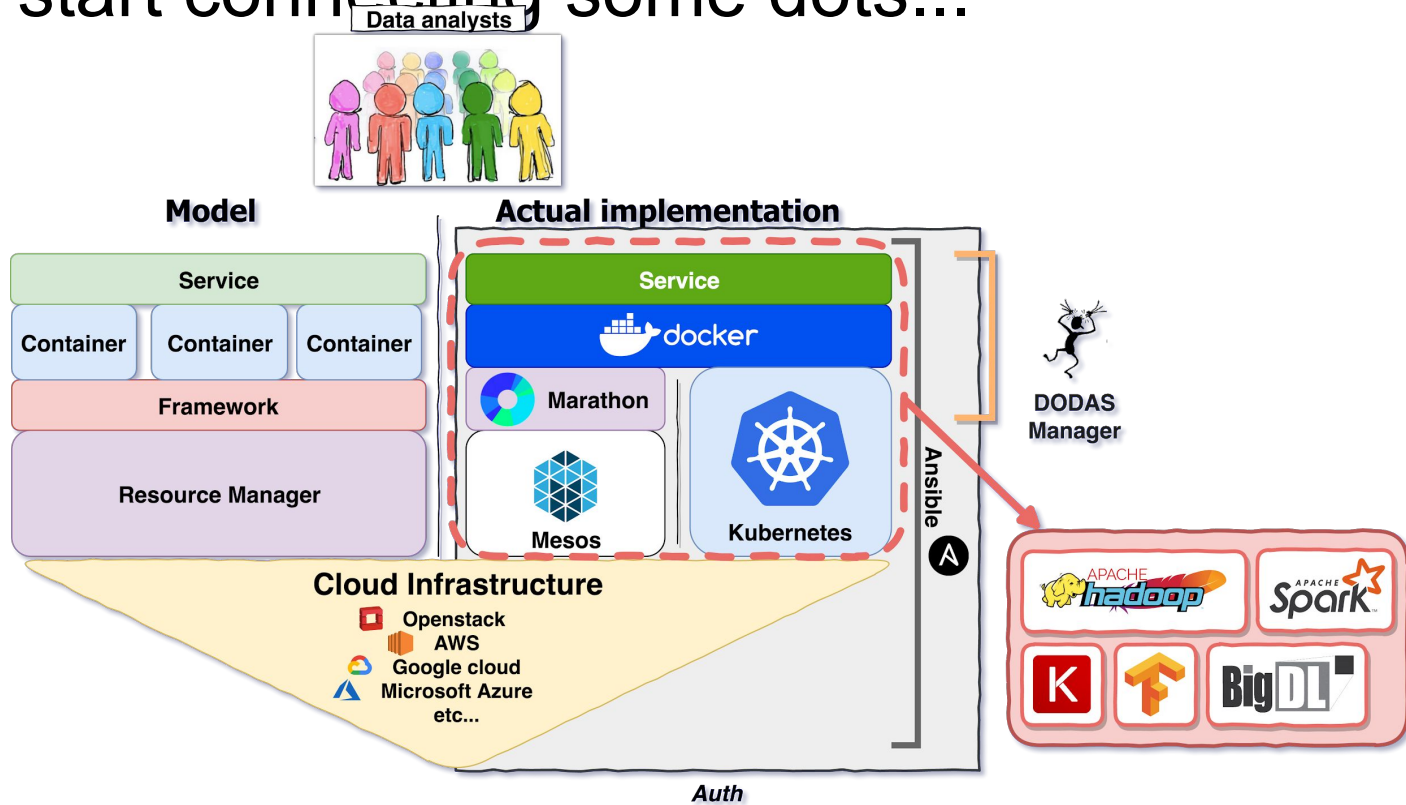


Cluster monitoring
(ELK/Prom)

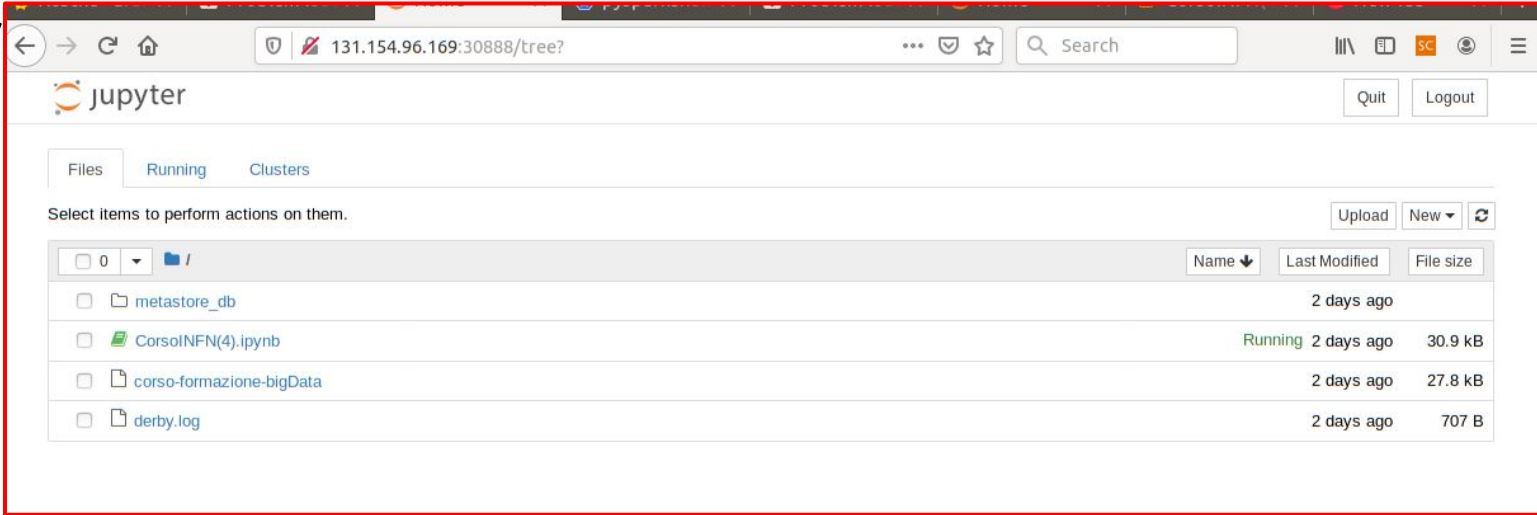


MINIO

Let's start connecting some dots...



Fr



Browser address: 131.154.96.169:30888/tree?

Page title: jupyter

Buttons: Quit, Logout

Tabs: Files, Running, Clusters

Select items to perform actions on them. [Upload] [New] [Refresh]

	Name	Last Modified	File size
<input type="checkbox"/>	metastore_db	2 days ago	
<input type="checkbox"/>	CorsoINFN(4).ipynb	Running 2 days ago	30.9 kB
<input type="checkbox"/>	corso-formazione-bigData	2 days ago	27.8 kB
<input type="checkbox"/>	derby.log	2 days ago	707 B

Accessed by



Runs on



```
10 inputs:
11
12   number_of_masters:
13     type: integer
14     default: 1
15
16   num_cpus_master:
17     type: integer
18     default: 2
19
20   mem_size_master:
21     type: string
22     default: "4 GB"
23
24   number_of_slaves:
25     type: integer
26     default: 1
27
28   num_cpus_slave:
29     type: integer
30     default: 2
31
32   mem_size_slave:
33     type: string
34     default: "4 GB"
35
36   server_image_slave:
```

```
65   k8s_master:
66     type: toasca.nodes.indigo.LRMS.FrontEnd.Kubernetes
67     properties:
68       admin_token: testme
69       kube_version: 1.14.0
70       kube_front_end_ip: { get_attribute: [ k8s_master_server, private_address, 0 ] }
71     requirements:
72       - host: k8s_master_server
73
```

```
73
74   k8s_wn:
75     type: toasca.nodes.indigo.LRMS.WorkerNode.Kubernetes
76     properties:
77       front_end_ip: { get_attribute: [ k8s_master_server, private_address, 0 ] }
78       kube_version: 1.14.0
79       nfs_master_ip: { get_attribute: [ k8s_master_server, private_address, 0 ] }
80     requirements:
81       - host: k8s_slave_server
82
```

TOSCA (cont)

```
k8s_master_server:  
  type: toska.nodes.indigo.Compute  
  capabilities:  
    endpoint:  
      properties:  
        network_name: PUBLIC  
      ports:  
        kube_port:  
          protocol: tcp  
          source: 6443  
        dashboard_port:  
          protocol: tcp  
          source: 30443  
        web_ui:  
          protocol: tcp  
          source: 30808  
        jupyter:  
          protocol: tcp  
          source: 30888  
    scalable:  
      properties:  
        count: { get_input: number_of_masters }  
  host:  
    properties:  
      num_cpus: { get_input: num_cpus_master }  
      mem_size: { get_input: mem_size_master }  
  os:  
    properties:  
      image: { get_input: server_image }
```

```
113 k8s_slave_server:  
114   type: toska.nodes.indigo.Compute  
115   capabilities:  
116     endpoint:  
117       properties:  
118         network_name: PRIVATE  
119     scalable:  
120       properties:  
121         count: { get_input: number_of_slaves }  
122   host:  
123     properties:  
124       num_cpus: { get_input: num_cpus_slave }  
125       mem_size: { get_input: mem_size_slave }  
126   os:  
127     properties:  
128       image: { get_input: server_image_slave }  
129
```

TOSCA but finally Spark

```
53     type: toska.nodes.indigo.HelmInstall
54     properties:
55         externalIP: { get_attribute: [ k8s_master_server, public_address,
56             name: "spark"
57         chart: "cloudpg/spark"
58         repos:
59             - { name: cloudpg, url: "https://cloud-pg.github.io/charts/"
60         values_file: { get_input: helm_values }
61     requirements:
62         - host: k8s_master
63         - dependency: k8s_wn
64
```

Compiling values at runtime and install

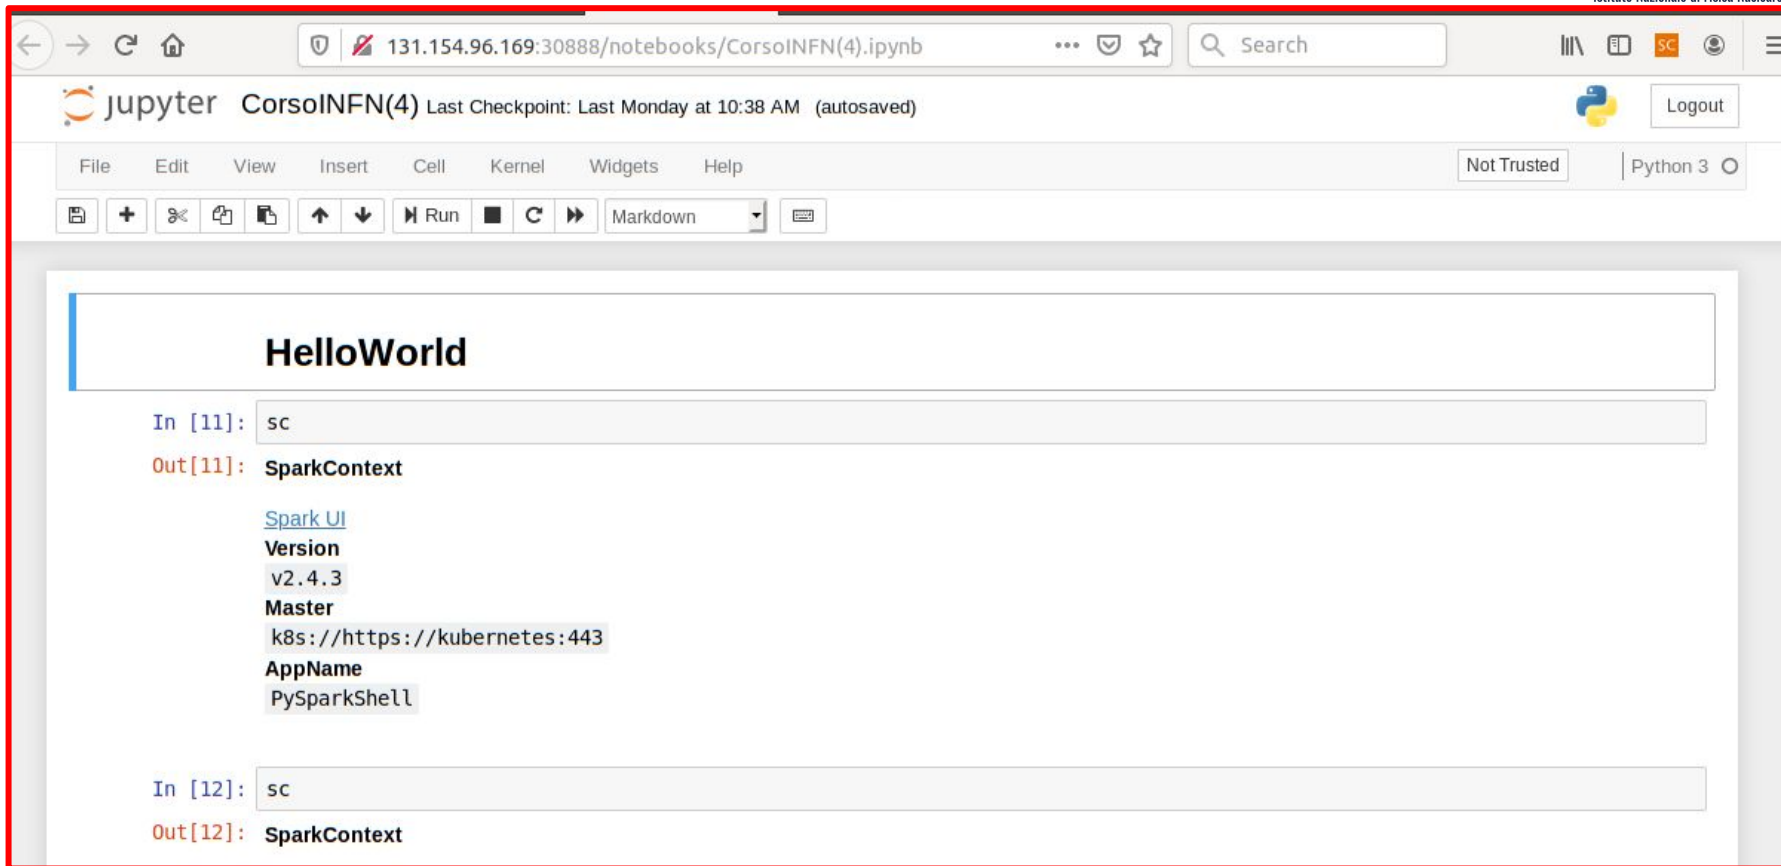
Who is doing this?

That's the last step..

Installing Spark on top of
k8s

```
23 lines (18 sloc) | 697 Bytes
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 }}"
```

And the result



The screenshot shows a Jupyter Notebook interface in a web browser. The browser address bar displays the URL `131.154.96.169:30888/notebooks/CorsoINFN(4).ipynb`. The notebook title is `CorsoINFN(4)` with a last checkpoint of `Last Monday at 10:38 AM (autosaved)`. The interface includes a menu bar (File, Edit, View, Insert, Cell, Kernel, Widgets, Help) and a toolbar with icons for file operations, running, and cell execution. The notebook content shows two code cells, each containing the command `sc`. The output for both cells is a `SparkContext` object with the following properties:

```
Spark UI
Version
v2.4.3
Master
k8s://https://kubernetes:443
AppName
PySparkShell
```

Cloud: quick reminder

<https://csrc.nist.gov/publications/detail/sp/800-145/final>


- **Infrastructure (IaaS → Infrastructure as a Service)**
- **IaaS**, the basic building blocks of a data center:
 - **Storage** I want to store data, lots of data, at low cost
 - **Compute** Give me a machine where I can host my services or run my applications
 - **Network** Create a “Software-Defined Network” infrastructure for me

→ No need to know details, no need to contact administrators to install something

Cloud definition (cont)

- **Platform (PaaS → Platform as a Service)**
 - **PaaS**, a computing platform providing you with several building blocks or components that you can request programmatically or statically. For example:
 - A cluster of systems with operating system and an entire execution environment installed and configured.
 - A web server (or a clusters) with database(s), virtual storage, load balancers...
- **Software (SaaS → Software as a Service)**
 - With **SaaS**, you are directly given access to some application software. You don't have to worry about the installation, setup and running of the application. You typically access SaaS apps via a web browser.
 - For example: gmail, social media such as
 - Facebook, Twitter, etc.

What matters at the end... are the applications.



Do we need something else?

While cloud environments made it convenient to build large-scale applications, there is still the downside of manual administration and operational interventions, such as:

- *Are the latest security fixes installed?"*
- *"When should we scale down/up?"*
- *"How many more servers do we need?"*

Ideally we would avoid all those administrative tasks, and we would like to simply focus on applications and related business value.

And thus yes: There is a digital transformation driven by the need for greater agility and scalability

- You saw containers as building blocks for Microservices as evolution of monolithic.
- We'll see now what come later

Containers

“Wouldn’t it be nice if one could **pack the application, with all its dependencies**, into a dedicated box and **run it anywhere**? No matter what software dependencies the host system has installed, or where and what the host system actually is?”


- **Yes it is and that is the idea of containerization allows all of this.**

Also, we know that containers are key pillars of microservices

Microservices architecture emerged as a **key method of providing development teams with flexibility** and other benefits, such as the ability **to deliver applications at warp speed** using infrastructure as a service (IaaS) and platform as a service (PaaS) environments.

Where we are so far...

Cloud Computing Infrastructure



Do we need something else?

So, is it all done?

- probably no, we need something else...

Wouldn't be nice if we could

- **divide our work into smaller pieces**
- **let the platform worry about manageability and autoscaling**

Great Ideally, but it's hard for the platform **to scale and manage the services**

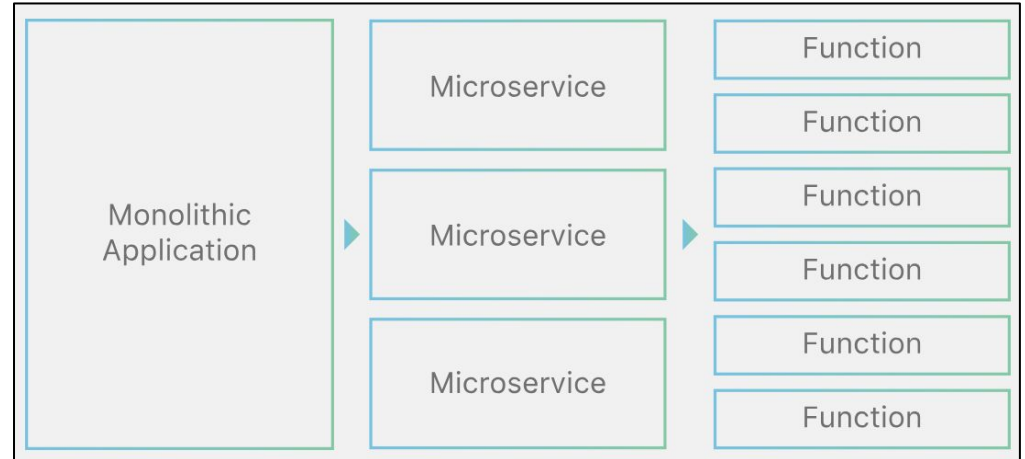
- So the suggestion might be to make them **stateless and smaller**

What is then something smaller than a “piece of application” running in a container?

Functions

With respect containers, **the basic idea of functions is to take the step further by making an application more granular to the level of functions (and events).**

Developers: difference here is to focus on a single function or module rather than a service with a large surface area like in the application runtime.



We've gone from monoliths to microservices to functions

Function as a Service (FaaS)

Extending the as a service model already presented we can define FaaS (Function as a Service) as

- **the ability to take a function and run it somewhere in the cloud**

Or maybe as

- the “**compute**” part of the **serverless** stack where you bring your own code.

The function contains a bespoke logic block. It is then called via some kind of registry like an API gateway, or it is scheduled or triggered by a cloud-related event (i.e., data written to Data Storage).

In other words: **FaaS is a form of serverless computing**, where you execute certain functions of your application in a **abstracted computing environment**.

Defining Server-less

Does server-less **means no servers**?

No, it is about deployment & operations model and **means worry-less about server operations and management**,



No servers to manage



Just code to develop and execute

Runs code **only** on-demand on a per-request basis over transparent & dynamically provisioned resources

Serverless vs FaaS

Let's consider serverless as an **amalgamation of two distinct points** as follows:

❑ **MBaaS, aka Mobile Backend-as-a-Service:**

The use of 3rd party services/applications (in the cloud) to handle the server-side logic and state.



❑ **FaaS, Functions-as-a-Service:** the use of 3rd party stateless compute containers to handle the server-side logic. These containers are event-triggered and may last for only one invocation i.e. ephemeral.

At its core, serverless computing provides runtimes to execute code, which is also known as function as a service (FaaS) platforms.

Ok, but in the end

What is serverless? Or better how we intend it in this lecture?

A cloud-native platform

for

short-running, **stateless** computation

and

event-driven applications

which

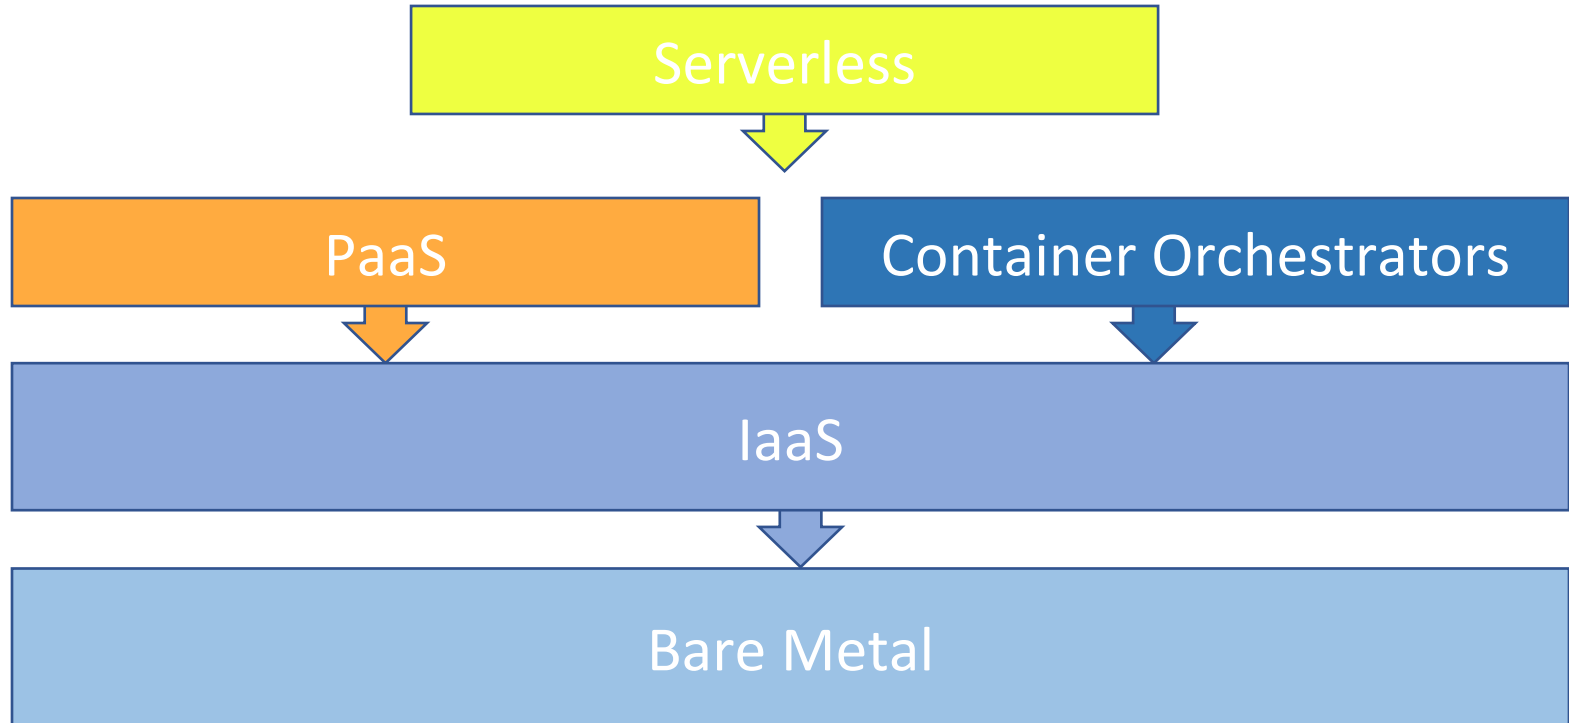
scales up and down instantly and automatically

and


(charges for actual usage at a millisecond granularity)

Where this positions ?

Cloud Computing Infrastructure



Cloud computing: server-less vs server-aware



If your PaaS can efficiently start instances in 20ms that run for half a second, then call it serverless.”

- Adrian Cockcroft (2016)

Why serverless ?

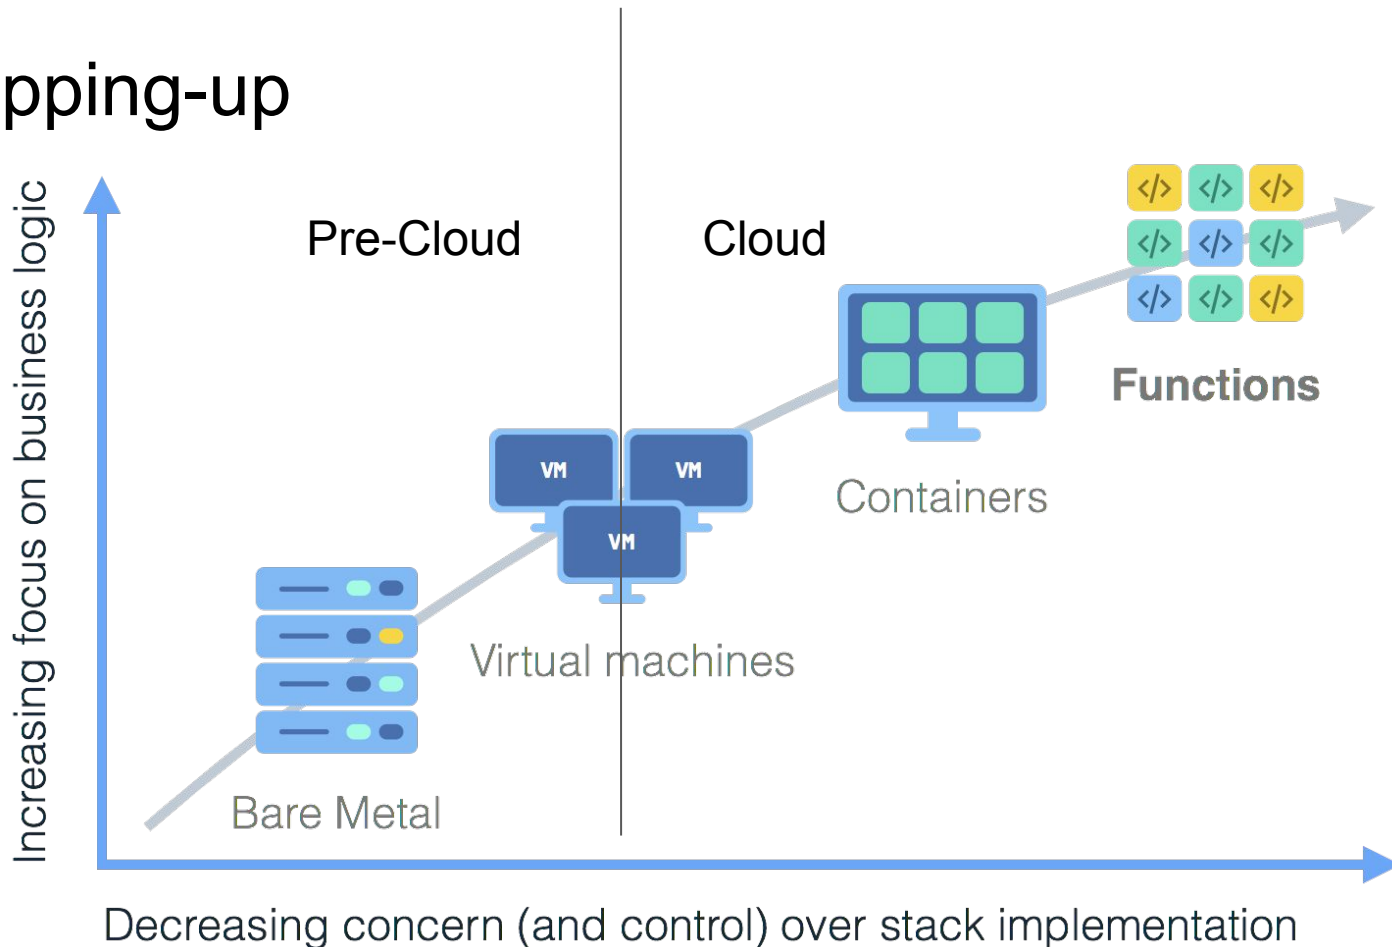
With serverless technologies, we perform another step toward automating and facilitating the use of Cloud resources.

Remember the inverted triangle we saw: what eventually matters are the applications, not the infrastructure.

Recall what happens with traditional Cloud applications, of which we have already seen several examples:


- We need to provision and manage the resources (e.g. VM1, VM2, the disks, the S3 buckets, etc.) for our applications.
- We are charged if we keep the resources up, even if they are doing nothing
- We are responsible to apply all the updates and security patches to our servers

Wrapping-up




Gaining attention...

Google Trends



FaaS



FaaS

VS

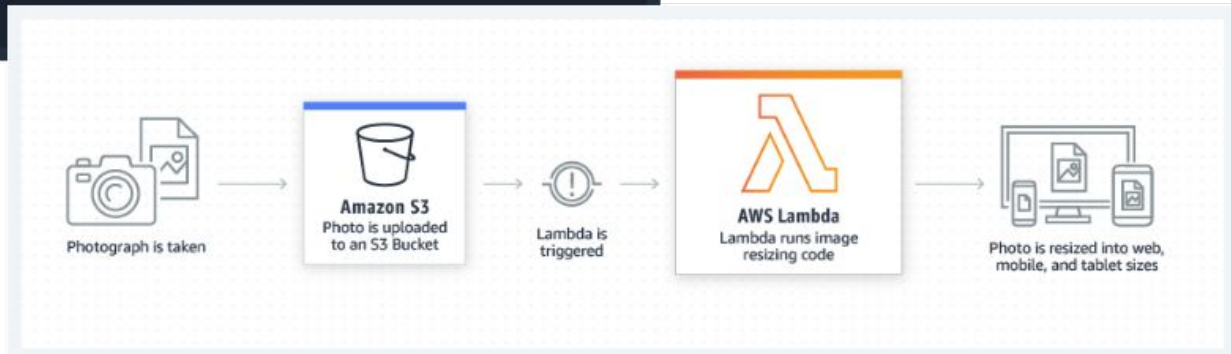
PaaS

Example

In the Amazon world, serverless computing is called AWS Lambda. This is how it works (picture from Amazon):



A simple AWS Lambda example:



Hands-on will provide more details

Some technologies (platforms for serverless)



Azure Functions



AWS
Lambda



OpenLambda



Kubernetes



IBM Cloud Functions



Red-Hat



Google Functions



What all of this is good for

good for

short-running
stateless
event-driven



Microservices



Mobile Backends



Bots, ML Inferencing



IoT



Modest Stream Processing



Service integration

not good for

long-running
stateful
number crunching



Databases



Deep Learning Training



Heavy-Duty Stream Analytics



Numerical Simulation



Video Streaming

One size fits all solutions?

No, generally speaking there is no a generic solutions which covers all the use cases...

This is true in general and for FaaS frameworks in particular

- Open problems
- Research challenges
- Questions
- ...

Challenges ahead of us: in a nutshell

- Can different cloud computing service models be mixed?
- Monitoring and debugging
 - Debugging is much different if instead of having one artifact (a micro-service or traditional monolithic app) developers need to deal with a myriad of smaller pieces of code ...
- Can legacy code be made running on serverless?
 - Hybrid model?
 - To what degree existing legacy code can be automatically or semi-automatically decomposed into smaller-granularity pieces to take advantage of these new economics?
- Is serverless fundamentally stateless?
 - Can there be serverless services that have stateful support built-in