

GRNET eScience platform for Big Data management *Codename: orka*

Monday, February 1, 2016



European Union
European Regional
Development Fund



digitalgreece
Everything is possible
Operational Programme
"Digital Convergence"



NSRF
2007-2013
Research & Innovation
Quality of life for everyone

The project is co-financed by Greece and the European Union

Project Vision

- Data-Intensive Science (store and process big data, at Petabyte scale)
- Scientific workflows
- Virtual Research Environment
- Data streaming

Big data

- The problem: data deluge
- Solution:
 - PaaS over
 - ~okeanos (VM, processing)
 - Pithos+ (storage)

Hadoop project

- Most popular implementation for the MapReduce programming paradigm
- Open source, commodity hardware
- Hadoop core (MapReduce, Hadoop distributed file system)
- Rich ecosystem (Pig, Hive, Hbase, many more)
- Researcher focuses on the **algorithm** and not the software install/maintain/scale etc.

Hadoop cluster with ~orka

- GUI, CLI, REST on top of ~okeanos to:
 - Create cluster (with configurable options) from a range of Hadoop distro's (aka images)
 - Transfer your data
 - Submit, execute, monitor jobs
 - Delete cluster
 - Start/stop/format cluster
 - Scale cluster, add/remove nodes
 - Save cluster creation metadata for reproducibility

Hadoop cluster with ~orka

~Okeanos Project Selection

Projects

esience.grnet.gr VM:14CPU:18RAM:24576MBDisk:200GB

Apply Last Cluster Configuration

Common Settings

Available Images

Debian Base

Cluster Name

cluster_name

Only with Workflow

public keys

Select public key

Cluster size

Select size

Storage:

Standard

~okeanos project maximum cluster size: 25

Master Settings

Predefined VM Flavors:

Small Medium Large

CPU: 0

RAM: 0

Disk Size: 0

1x 2x 4x 8x

512 1024 2048 4096 6144 8192

5 10 20 40 60

CPU: (18 available)

RAM size (24576 MB available)

Disk size (200 GB available)

Slaves Settings

Predefined VM Flavors:

Small Medium Large

CPU: each: 0

RAM: each: 0

Disk Size: each: 0

1x 2x 4x 8x

512 1024 2048 4096 6144 8192

5 10 20 40 60

CPU: (18 available)

RAM size (24576 MB available)

Disk size (200 GB available)

Hadoop Configuration

Replication Factor

2 (default)

HDFS blocksize in MB

128 (default)

Default HDFS Configuration

Cancel

Create Cluster



European Union
European Regional
Development Fund



digitalgreece
Everything is possible
Operational Programme
"Digital Convergence"



The project is co-financed by Greece and the European Union

Add-ons to basic Hadoop

- Other components & runtimes
 - Spark
- Apache Hadoop-based distro's
 - Cloudera
 - Hue (HDFS explorer, Oozie web editor)
- Storage backend
 - Pithos ⇔ HDFS connector (analogous to Amazon S3 Filesystem for Hadoop)

Scientific Workflows

- Orchestration of atomic jobs
- Apache Oozie
- Apache Pig
 - Built-in in orka images

Collaborative scientific research

- Virtual Research Environment
- Complete system for teams and projects
- Components:
 - Research/Project home page (portal, wiki)
 - Project Management
 - Teleconference
 - Digital repositories
- Implemented as Docker images

Virtual Research Environment



Category	Software stack
Portal / CMS	Drupal (v7.37)
Wiki, blog, forum	Mediawiki (v1.2.4)
Project management	Redmine (v3.04)
Web conferencing	BigBlueButton (v0.81)
Digital repositories	DSpace (v5.3)




Reproducible Research

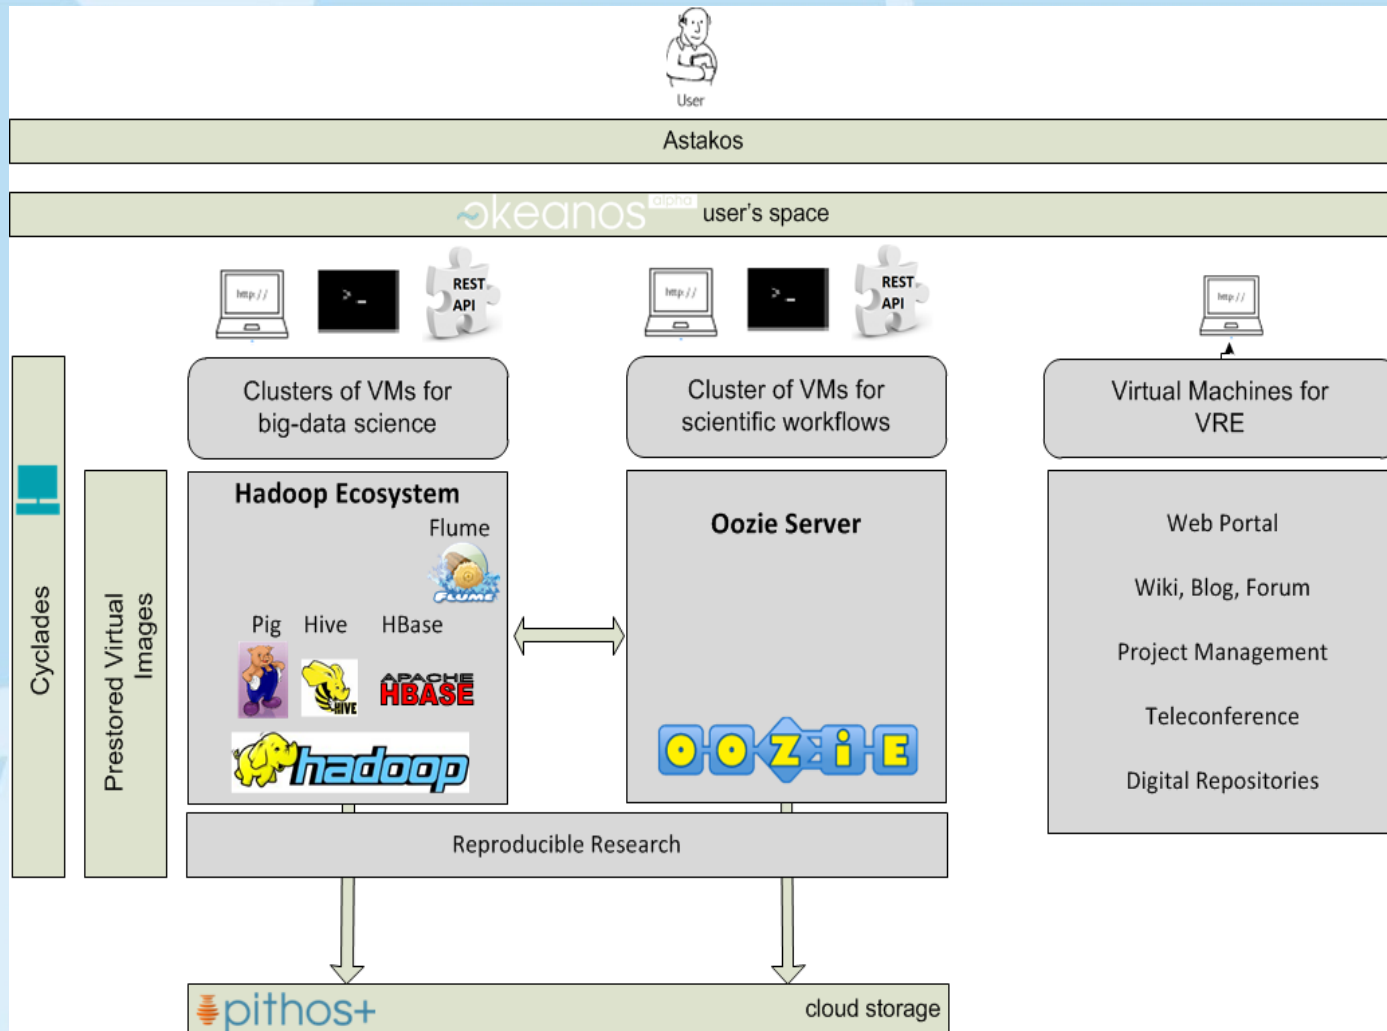
- Save your experiment's metadata as a bundle
- Domain Specific Language (DSL) that fully describes an experiment/job
- Text editor => simple YAML file
- Re-play, possibly with different parameters
- Save bundle to Pithos
- Share your bundle with other ~okeanos users

Data streams into HDFS

- Apache Flume
- Integrated into the Hadoop ecosystem
- Focus on streaming data

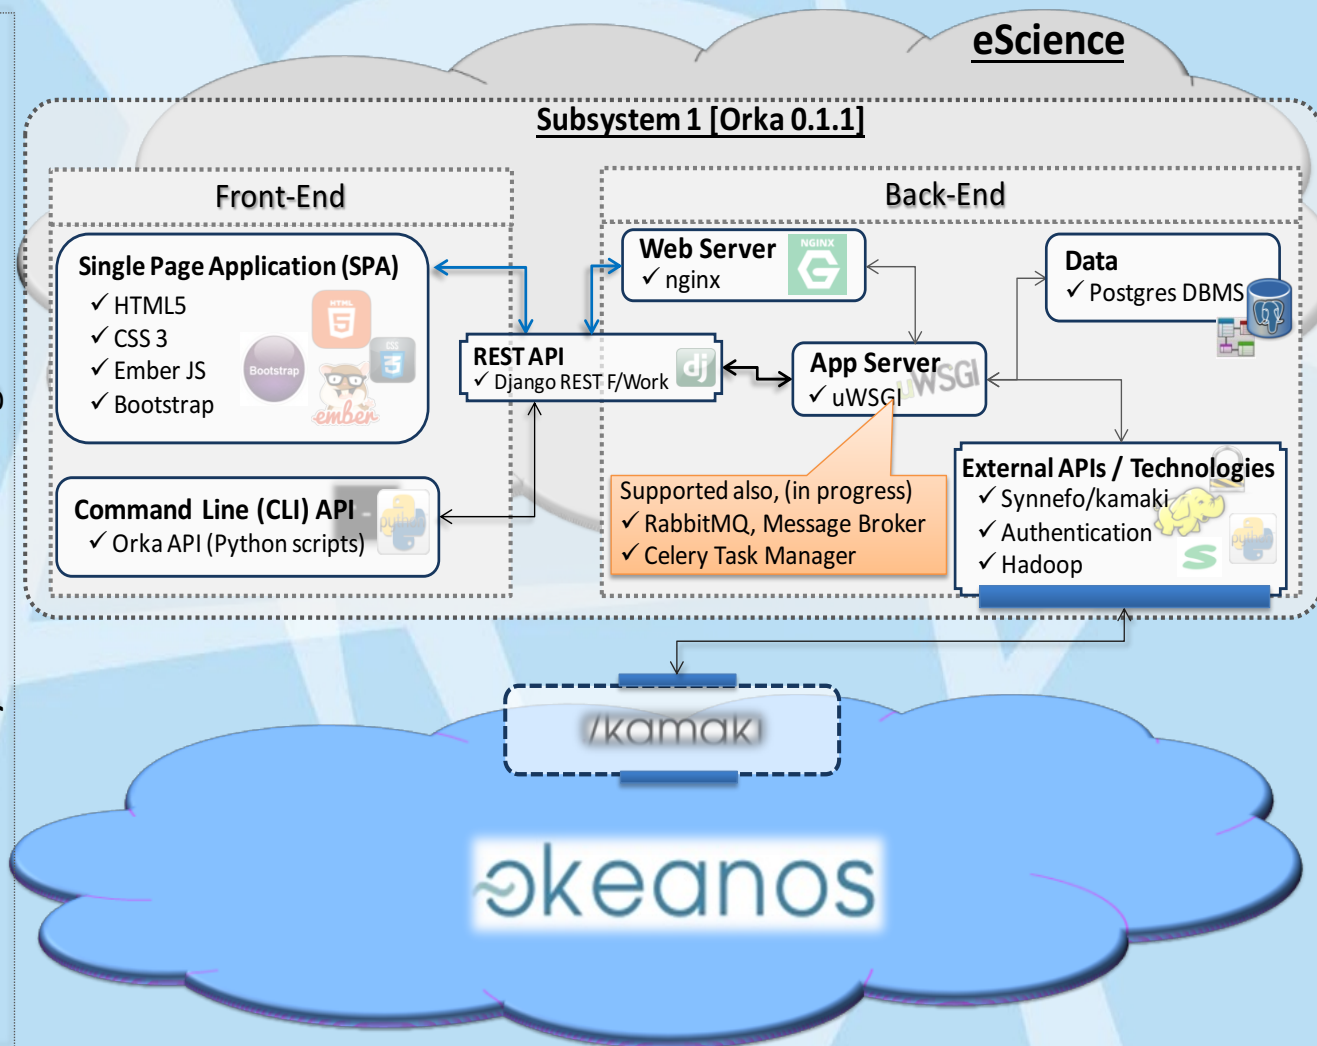


High-level Architecture



Technology Stack

Orka SubSystem: Technologies Overview



Current state

- github.com/grnet/e-science 
- escience.grnet.gr 



Simplifying Computing

The lambda architecture

a

a useful framework to think about designing big data applications

b

a robust framework for ingesting real-time streams of data while providing efficient stream and batch analytics.

c


fault-tolerant against both hardware failures and human errors

d

serves a wide range of use cases, and in which low-latency reads and updates are required

λ : lambda architecture

The Lambda Architecture solves the problem of computing arbitrary functions on arbitrary data in realtime by decomposing the problem into three layers: the batch layer, the serving layer, and the speed layer.



Batch Layer

The batch layer has two functions: (i) managing the master dataset (an immutable, append-only set of raw data), and (ii) pre-computing arbitrary query functions, called batch views.

Serving layer

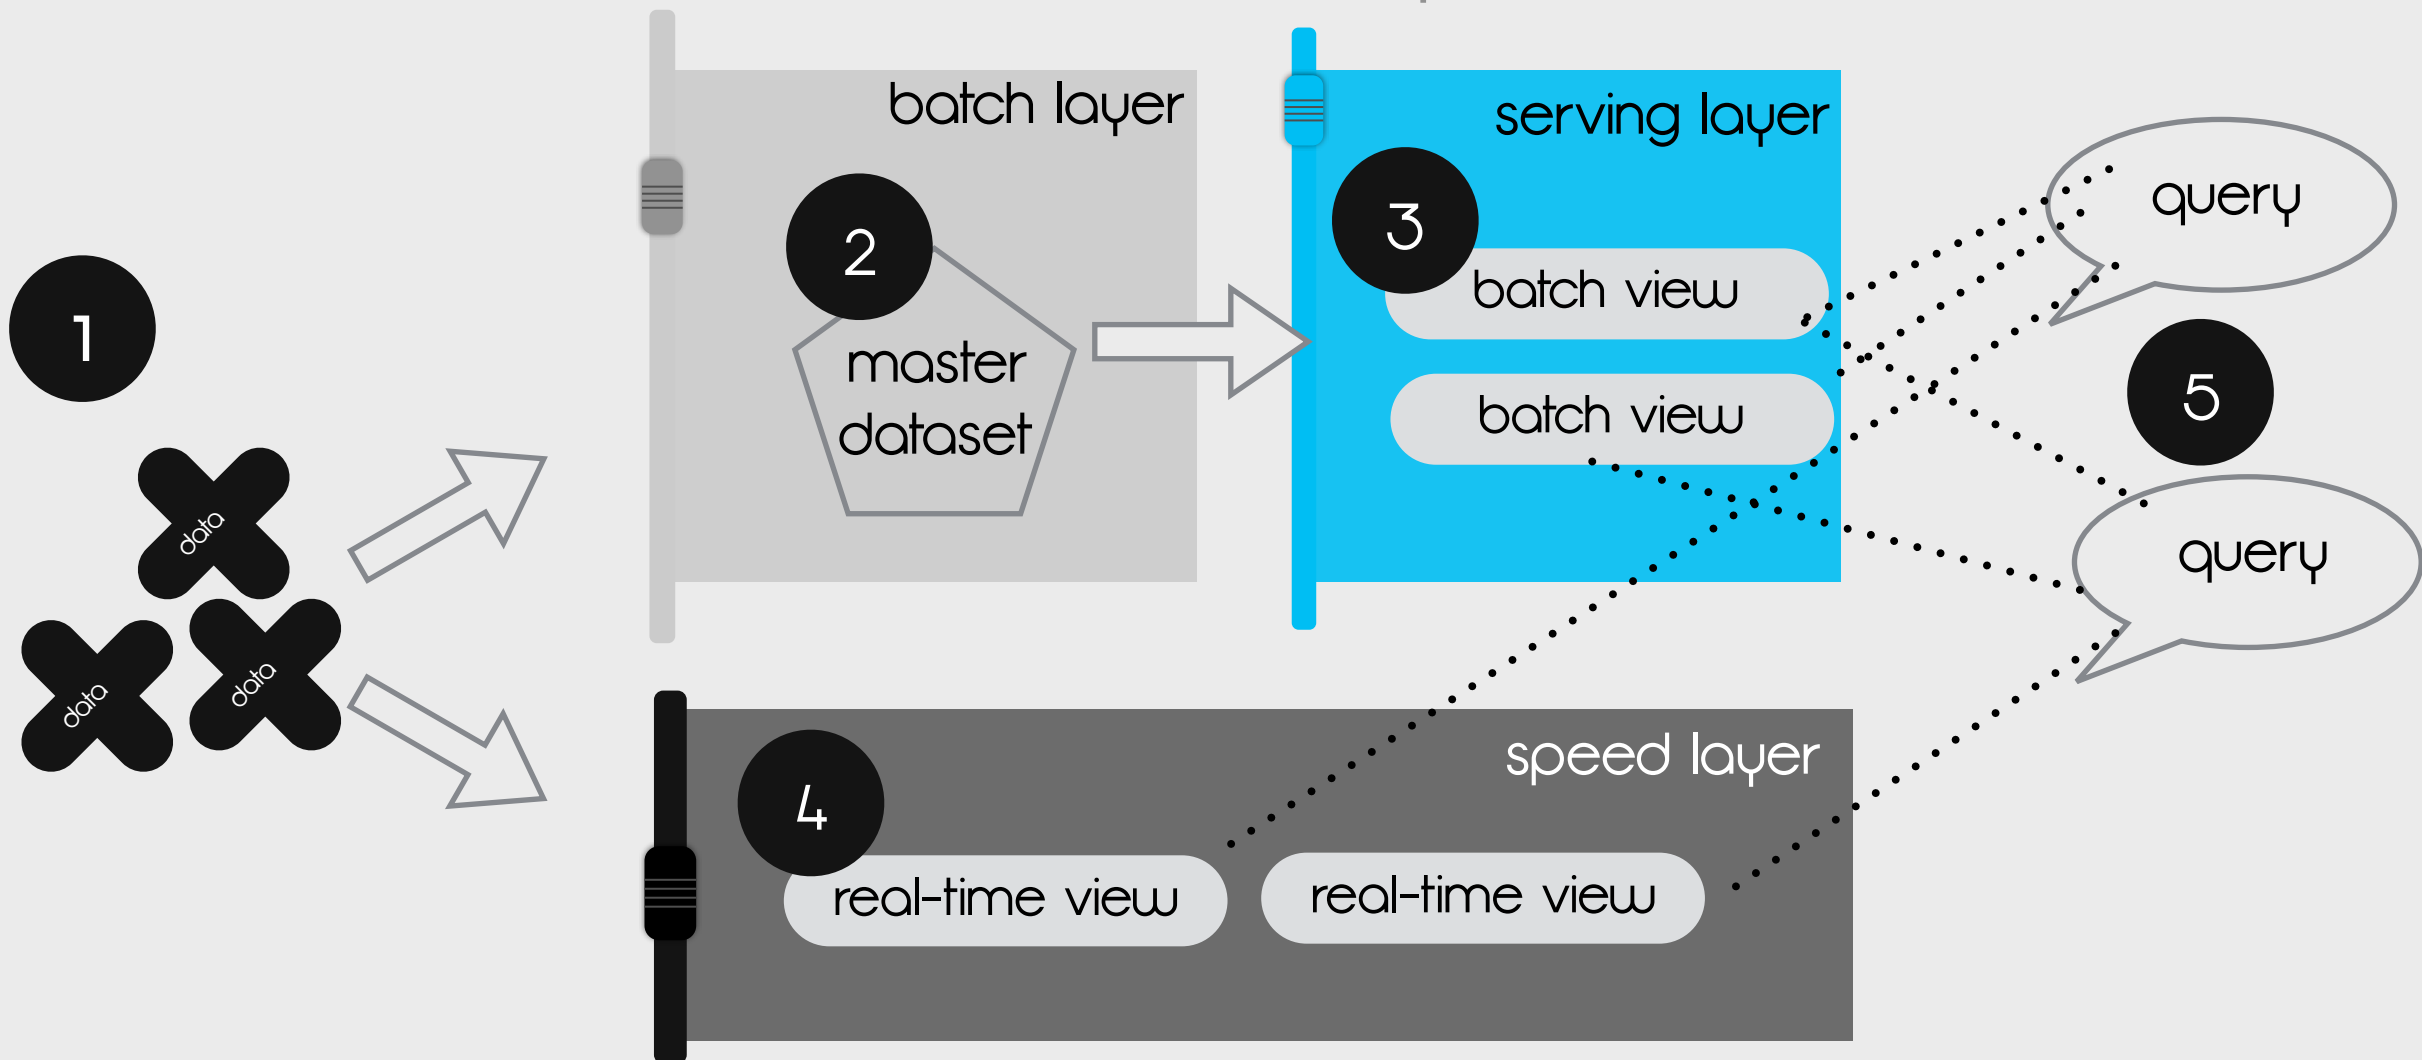
The serving layer indexes the batch views so that they can be queried in low-latency, ad-hoc way.

Speed layer

The speed layer compensates for the high latency of updates to the serving layer and deals with recent data only.

λ : lambda architecture

an example



1 data is dispatched to batch and speed layer for processing.

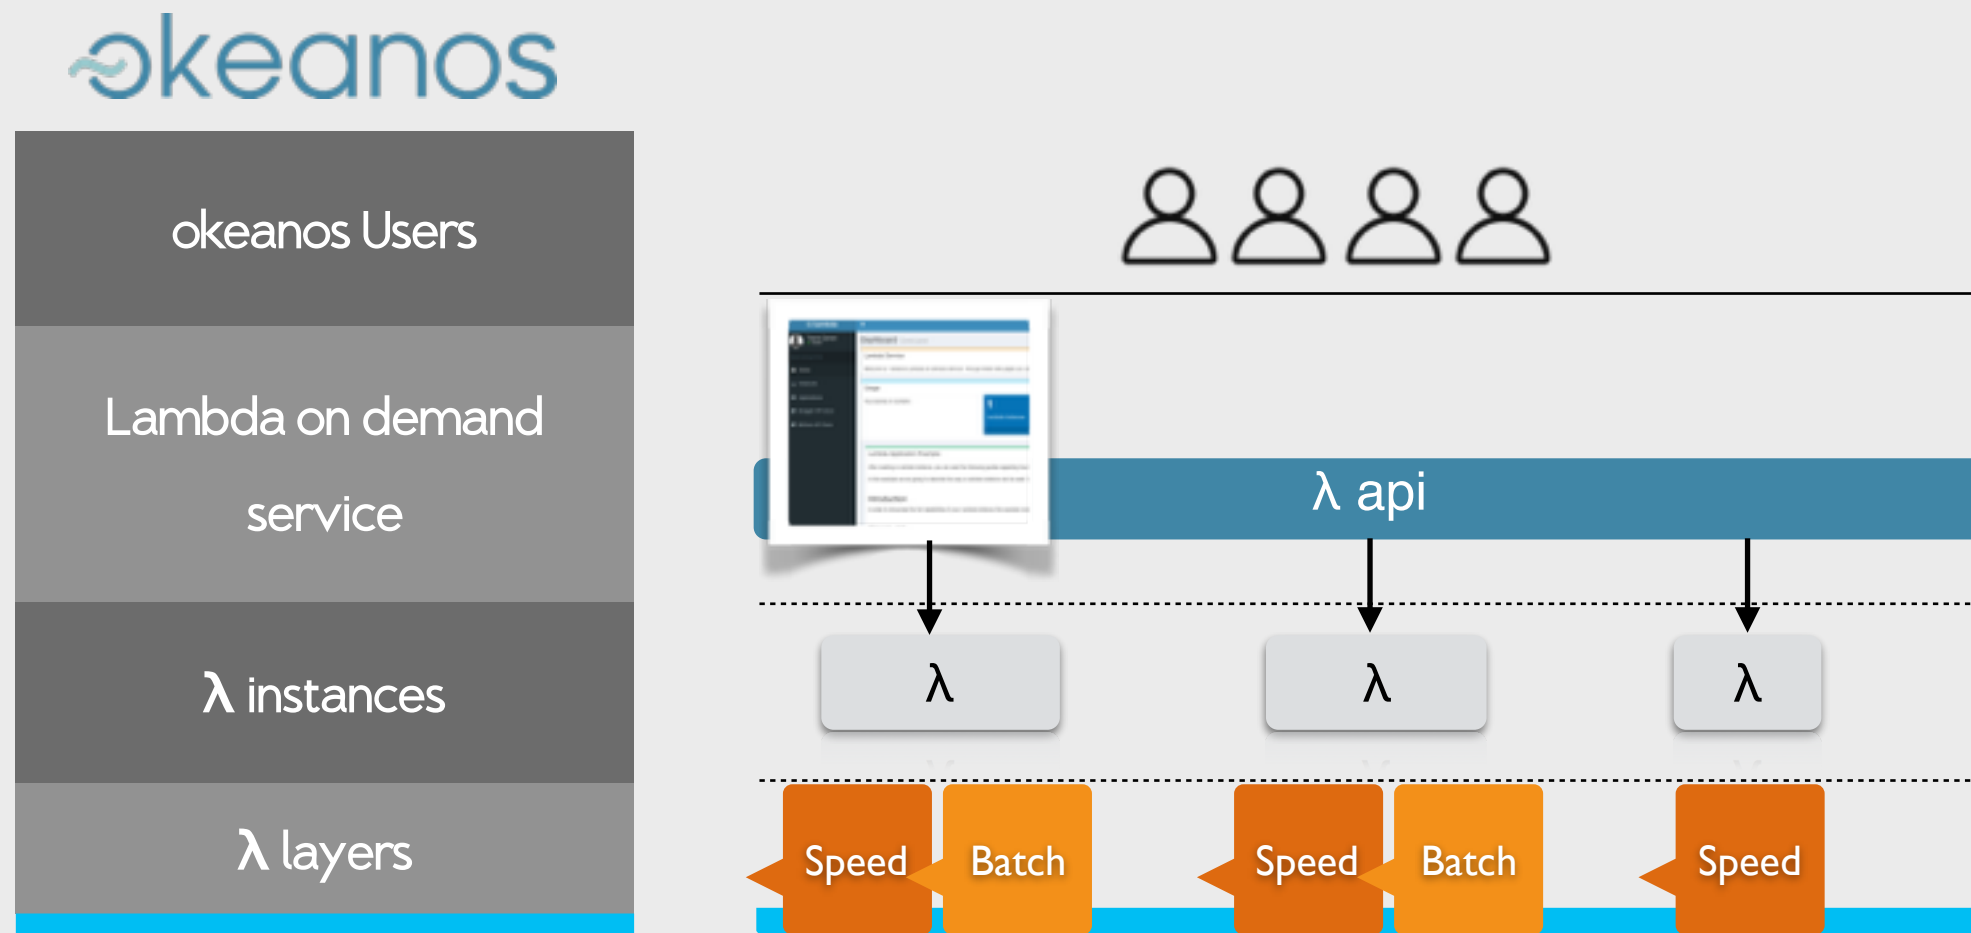
2 precomputes the batch views

3 indexes the batch views


4 deals with recent data only.

5 Any incoming query can be answered by merging results from batch views and real-time views.

Provisioning a λ instance

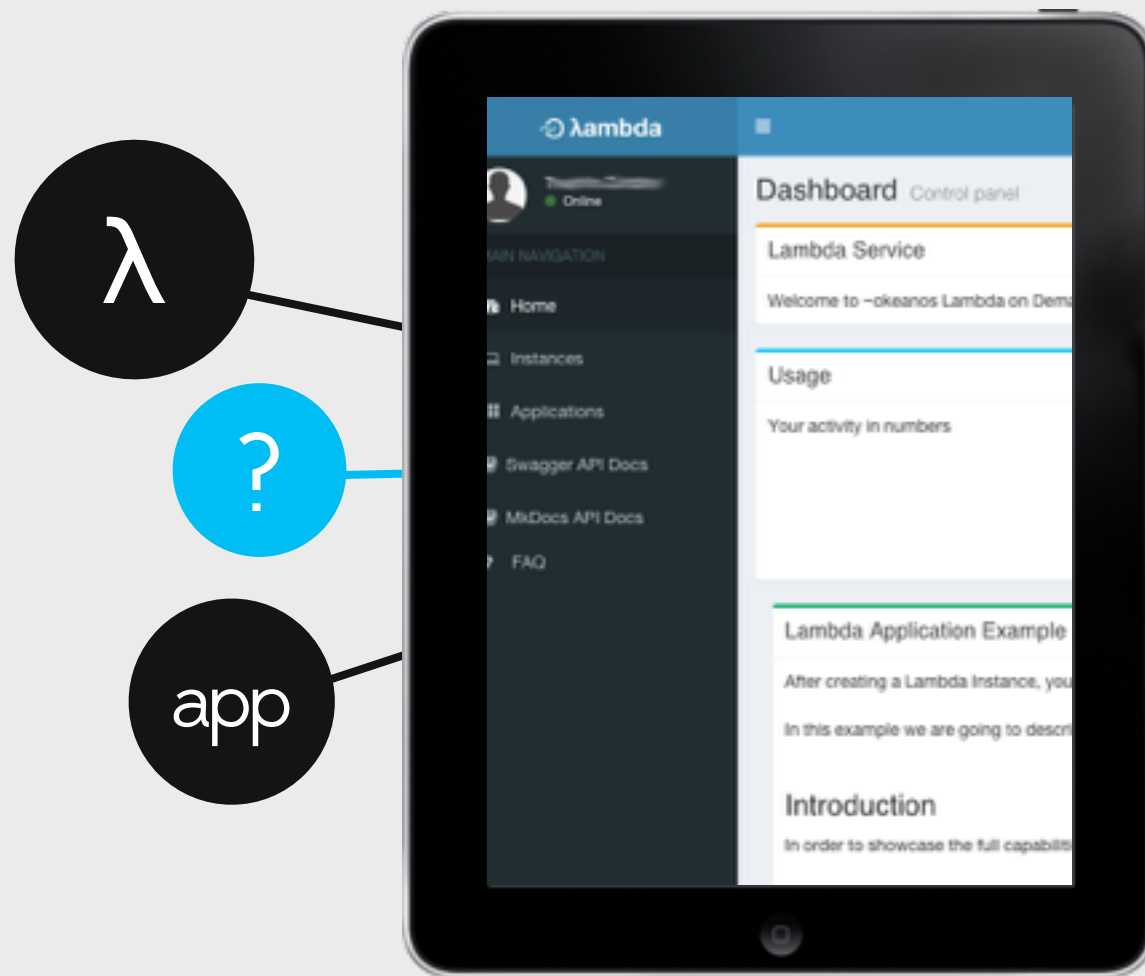


Based on



lambda UI

Dashboard, Instances, Applications and help



λ - Instances

manage lambda instances

Create your lambda instances based on your needs. Manage , deploy applications and start your lambda instance.

Applications

manage your applications

Upload your Java or Scala application for streaming and batch jobs. Your applications are stored on the Pithos+ storage service.

Help

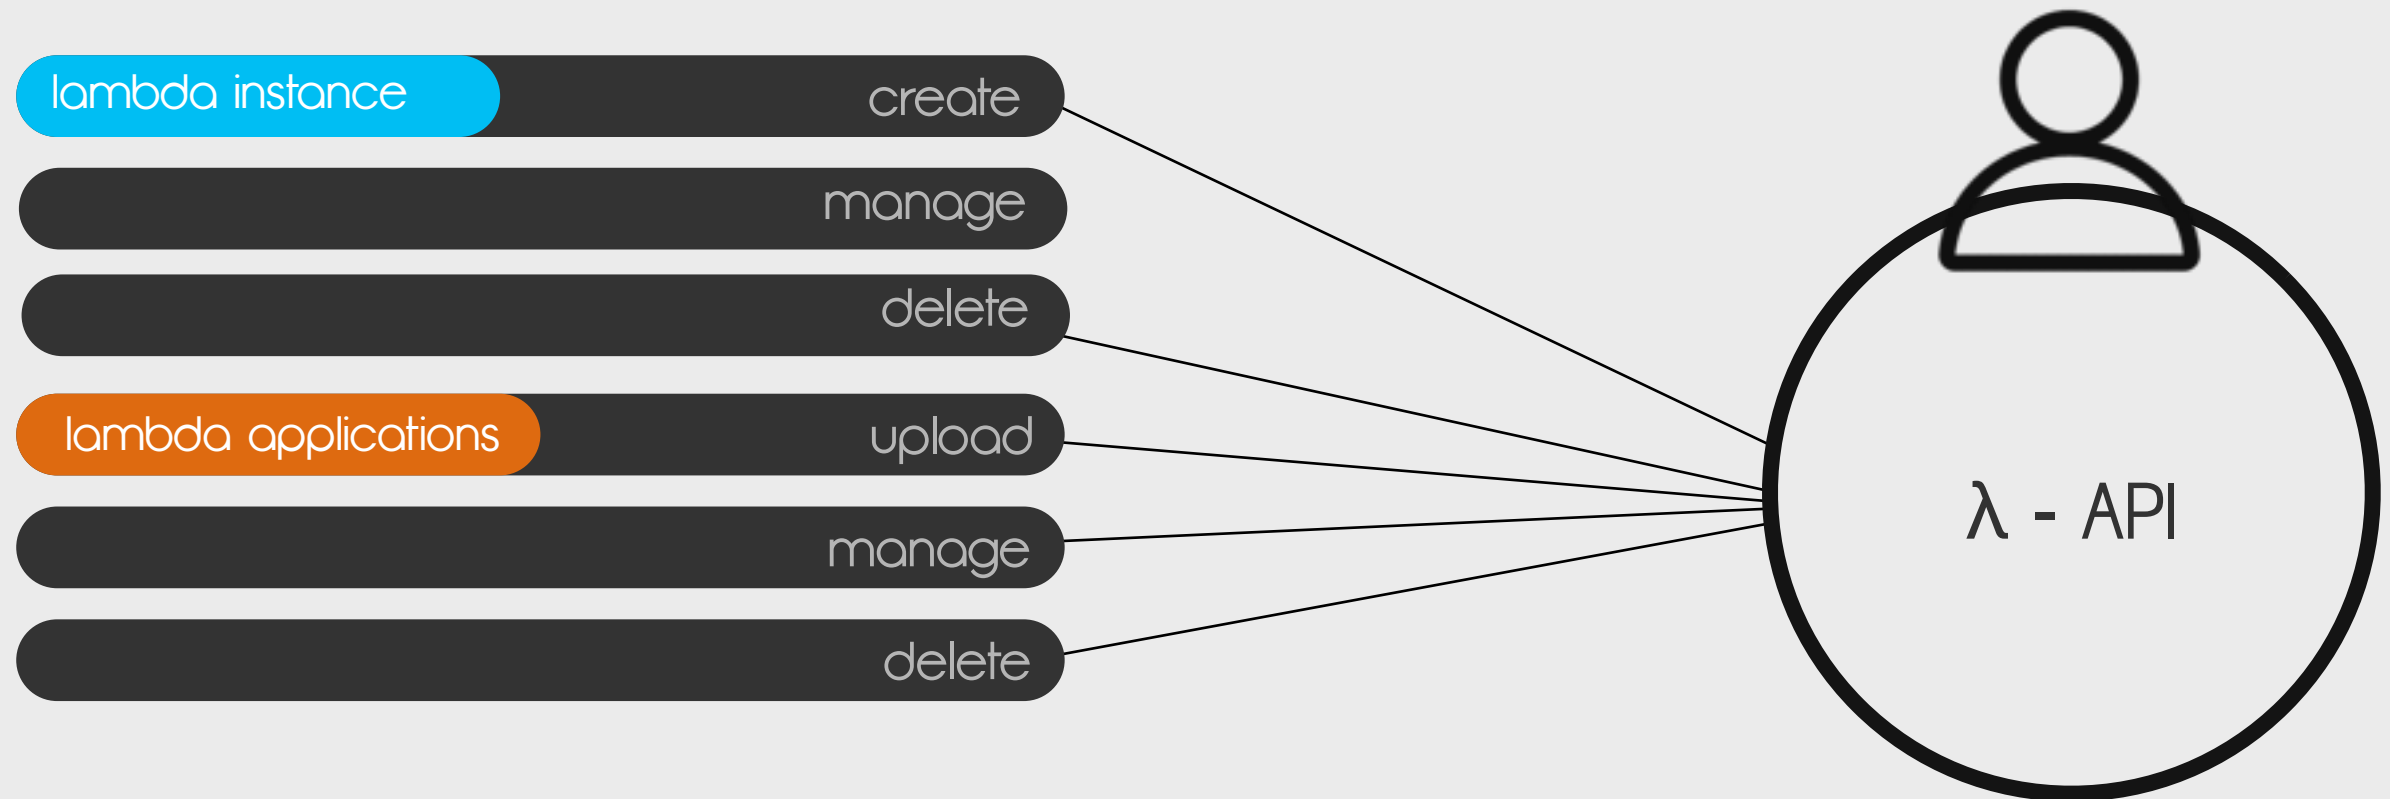
Informational guides

Short guides on how to

- I) deploy, run and manage your lambda instances.
- II) deploy, run and manage your applications
- III) export and view your results

Experienced User

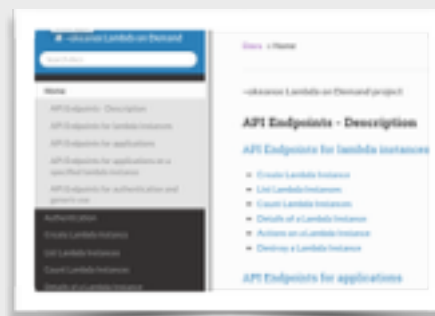
Use the λ ambda API



well documented
with



Swagger



mkdocs doc

e-science vs λ

Lambda λ : focuses on analysing streaming Data

e-Science: focuses on existing data + offers a pre-installed collaborative tools to handle data

Questions ?