dockercon EU
BARCELONA • 16-17 NOVEMBER 2015

Containing IoT Sensor Telemetry

http://lect.europalab.com/dociotsen/

Samuel Cozannet, Canonical
Michael Schloh, Europalab

In this hour...

Dive into IoT

Sensor telemetry
Actuator telecommand
Computing device nodes
Gateways and routing

Added Value

Isolation and reliability
Transports and protocols
Maintenance reduction
Data management

SDP Workflows

Embedded challenges
Computing node images
Research and develop
Deployment strategies

Show and Tell

ARMHF image tour
Application layers
Sensory data flows
IoT service deployments

Dive Into IoT

Defining the Internet of Things

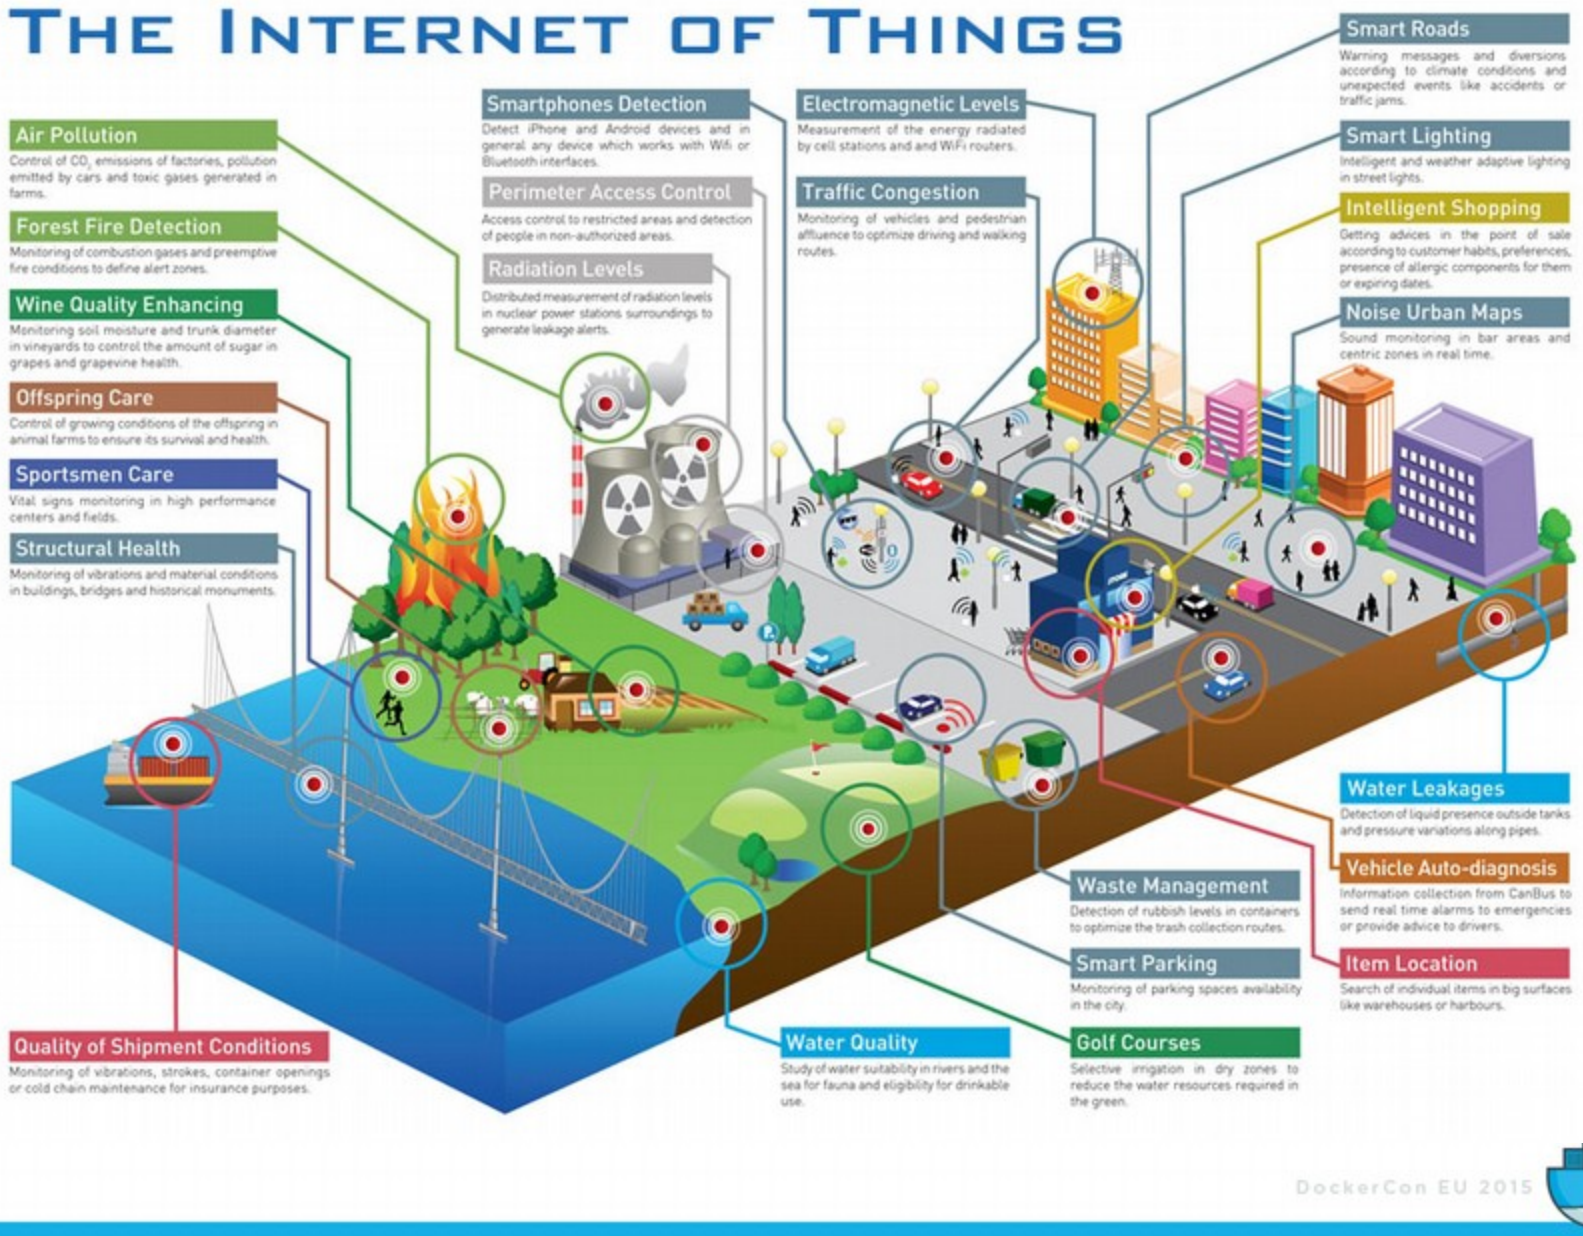
THE INTERNET OF THINGS
Air Pollution
Control of CO_2 emissions of factories, pollution emitted by cars and toxic gases generated in farms.
Forest Fire Detection
Monitoring of combustion gases and preemptive fire conditions to define alert zones.
Wine Quality Enhancing
Monitoring soil moisture and trunk diameter in vineyards to control the amount of sugar in grapes and grapevine health.
Offspring Care
Control of growing conditions of the offspring in animal farms to ensure its survival and health.
Sportsmen Care
Vital signs monitoring in high performance centers and fields.
Structural Health
Monitoring of vibrations and material conditions in buildings, bridges and historical monuments.
Quality of Shipment Conditions
Monitoring of vibrations, strokes, container openings or cold chain maintenance for insurance purposes.
Smartphones Detection
Detect iPhone and Android devices and in general any device which works with Wifi or Bluetooth interfaces.
Perimeter Access Control
Access control to restricted areas and detection of people in non-authorized areas.
Radiation Levels
Distributed measurement of radiation levels in nuclear power stations surroundings to generate leakage alerts.
Electromagnetic Levels
Measurement of the energy radiated by cell stations and and WiFi routers.
Traffic Congestion
Monitoring of vehicles and pedestrian affluence to optimize driving and walking routes.
Water Quality
Study of water suitability in rivers and the sea for fauna and eligibility for drinkable use.
Waste Management
Detection of rubbish levels in containers to optimize the trash collection routes.
Smart Parking
Monitoring of parking spaces availability in the city.
Golf Courses
Selective irrigation in dry zones to reduce the water resources required in the green.
Smart Roads
Warning messages and diversions according to climate conditions and unexpected events like accidents or traffic jams.
Smart Lighting
Intelligent and weather adaptive lighting in street lights.
Intelligent Shopping
Getting advices in the point of sale according to customer habits, preferences, presence of allergic components for them or expiring dates.
Noise Urban Maps
Sound monitoring in bar areas and centric zones in real time.
Water Leakages
Detection of liquid presence outside tanks and pressure variations along pipes.
Vehicle Auto-diagnosis
Information collection from CanBus to send real time alarms to emergencies or provide advice to drivers.
Item Location
Search of individual items in big surfaces like warehouses or harbours.
DockerCon EU 2015

44


nest
69
kWh
MULT BY
1

- Sensor telemetry
- Actuator telecommand
- Computing device nodes
- Gateways and routing

"Telemetry concerns remotely determining what a device senses, while telecommand directs a remote device to act" —Michael Schloh von Bennewitz

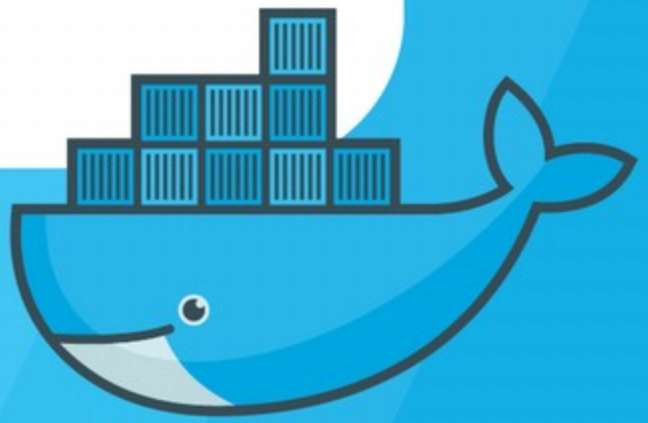

dockercon EU

Use Cases

Docker container roles in Internet of Things research and development

- Release engineering
- SDK and IDE deployment
- Firmware infrastructure
- Runtime management

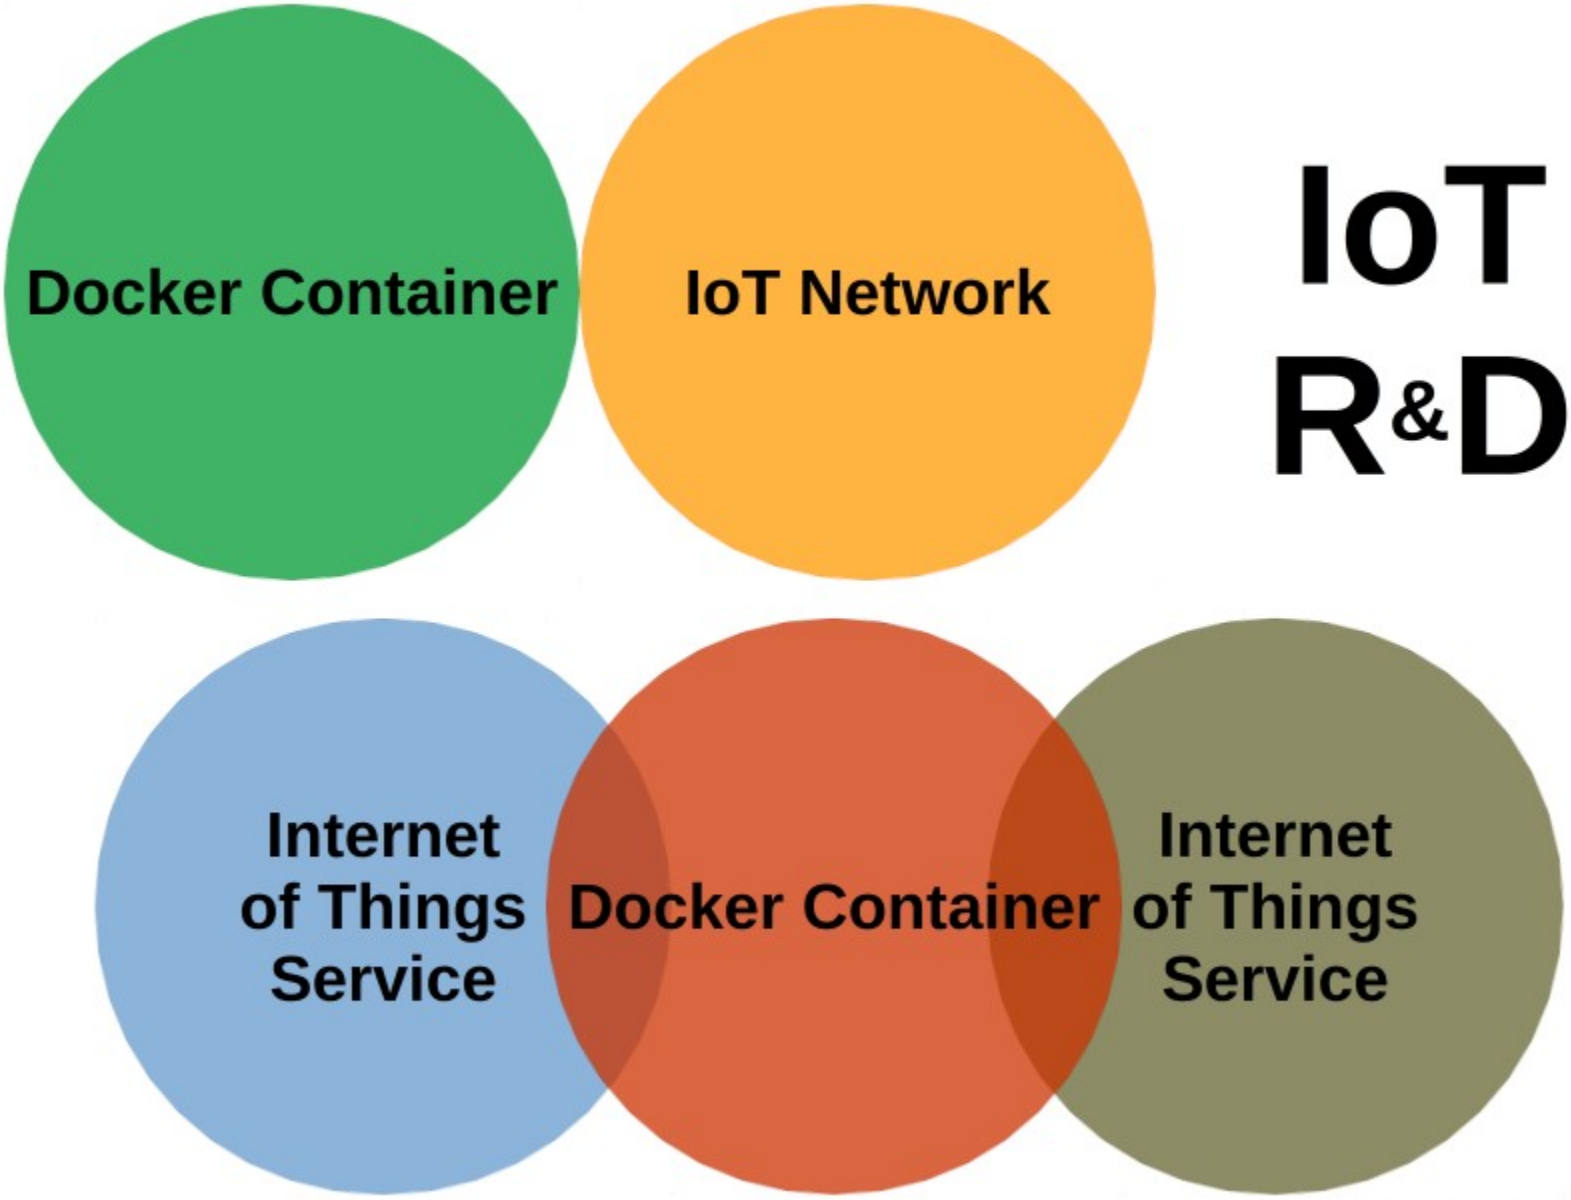
Docker Container
IoT Network
IoT
R&D
Internet
of Things
Service
Docker Container
Internet
of Things
Service

dockercon EU

New Workflows

Software Development Process departure from classic devops

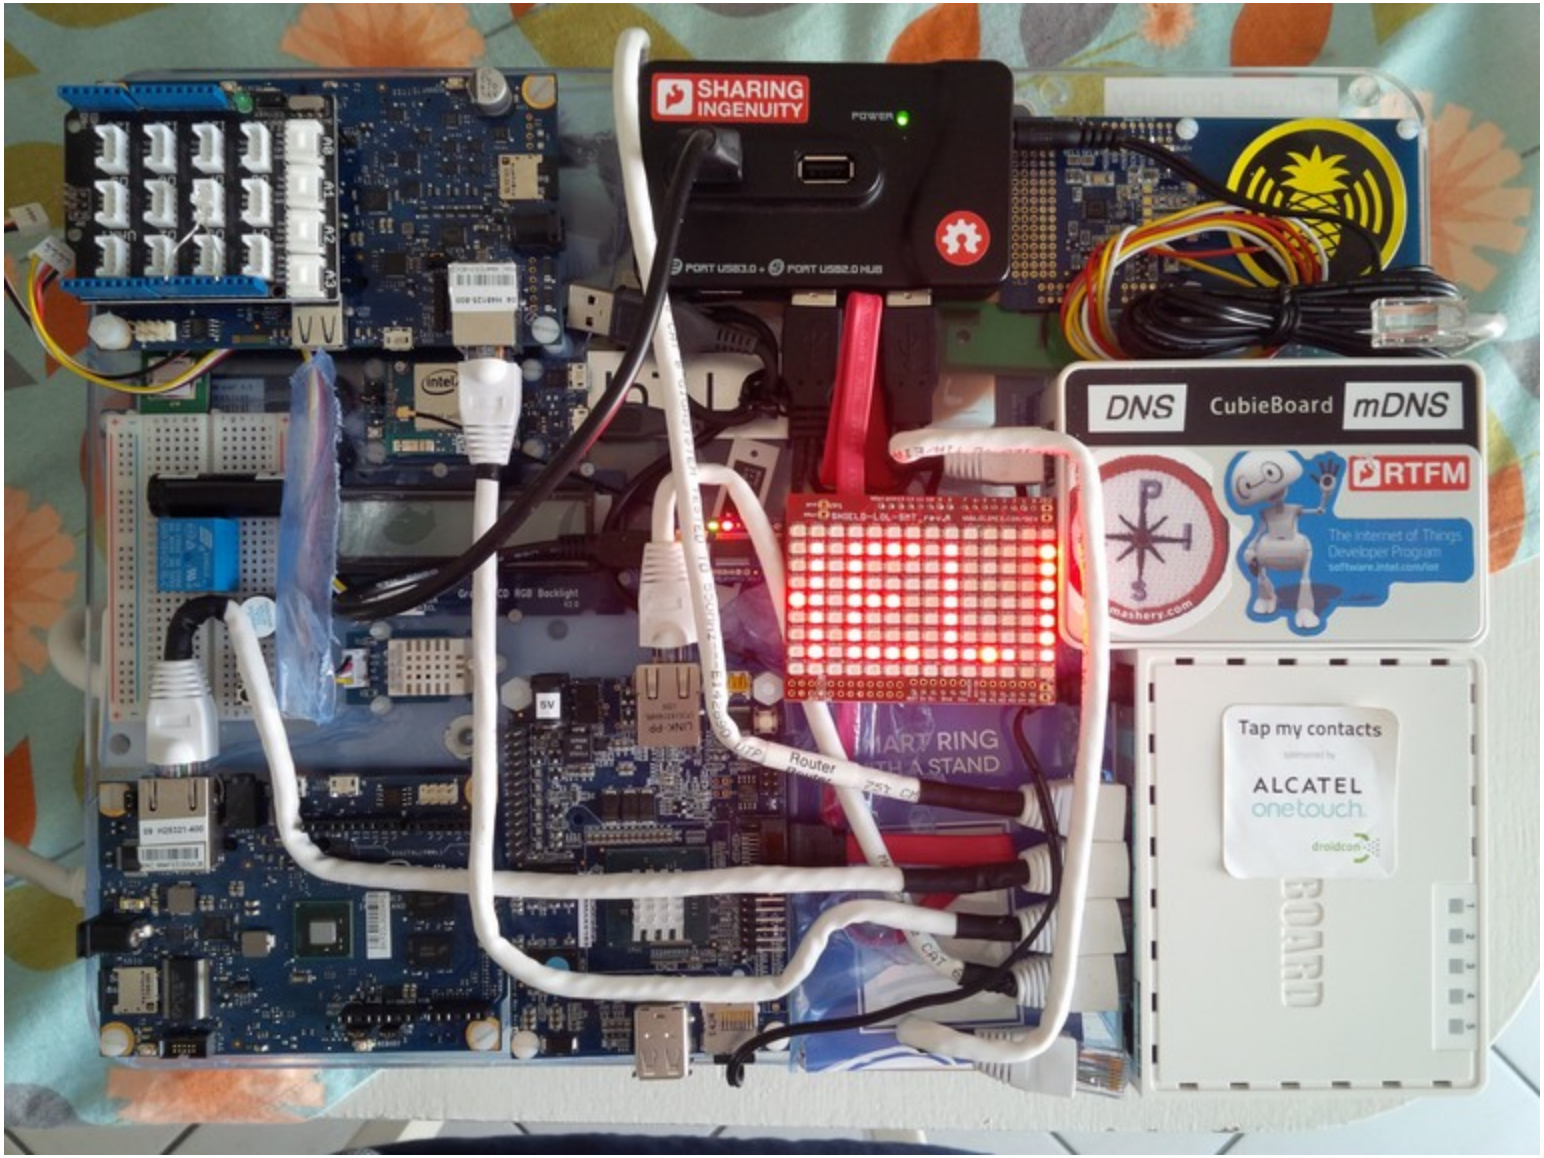
SHARING INGENUITY
POWER
DNS
CubieBoard
mDNS
RTFM
The Internet of Things Developer Program
Tap my contacts
ALCATEL
onetouch
droidcon
BOARD

- Embedded challenges
- Computing node images
- Research and develop
- Deployment strategies

Adding Value

Smart engineering of ubiquitous systems in contrast to containers

- Isolation and reliability
- Transports and protocols
- Maintenance reduction
- Data management

```
$ docker pull msvb/armhf-iotempire
$ docker pull msvb/ia32-iotempire
$ docker run -t -i msvb/armhf-iotempire
```

- -i: interactive stdin
- -t: pseudo-tty stdin
- Images are alpha tested

Sample image contents:

- **MQTT**
- **AMQP**
- **LWM2M**
- **ZeroMQ**

Show and Tell

Whirlwind tour of Docker assisted networks of cloudy sensor things

- ARMHF image tour
- Application layers
- Sensory data flows
- IoT service deployments

dockercon EU

Front End

Human to Machine (H2M) interfaces

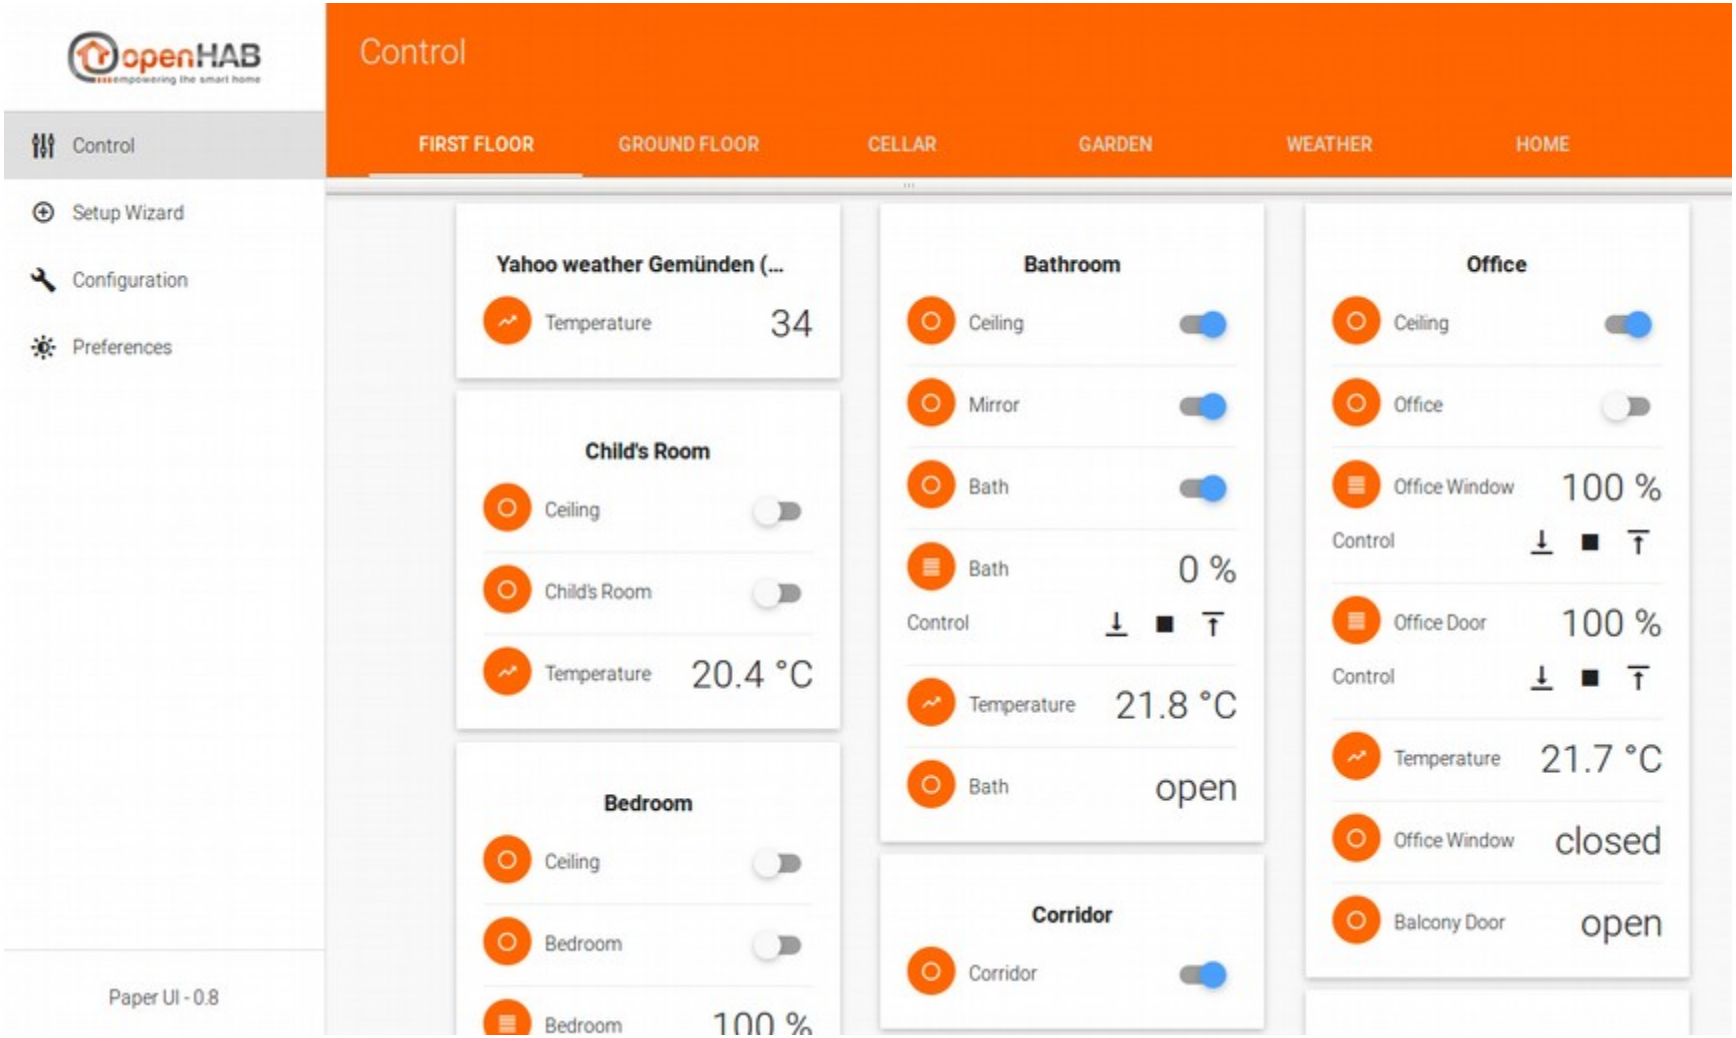
openHAB
empowering the smart home
Control
Setup Wizard
Configuration
Preferences
Paper UI - 0.8
Control
FIRST FLOOR
GROUND FLOOR
CELLAR
GARDEN
WEATHER
HOME
Yahoo weather Gemünden (...
Temperature
34
Child's Room
Ceiling
Child's Room
Temperature
20.4 °C
Bedroom
Ceiling
Bedroom
Bedroom
100 %
Bathroom
Ceiling
Mirror
Bath
Bath
0 %
Control
Temperature
21.8 °C
Bath
open
Corridor
Corridor
Office
Ceiling
Office
Office Window
100 %
Control
Office Door
100 %
Control
Temperature
21.7 °C
Office Window
closed
Balcony Door
open

Back End

How to store & process all that data?

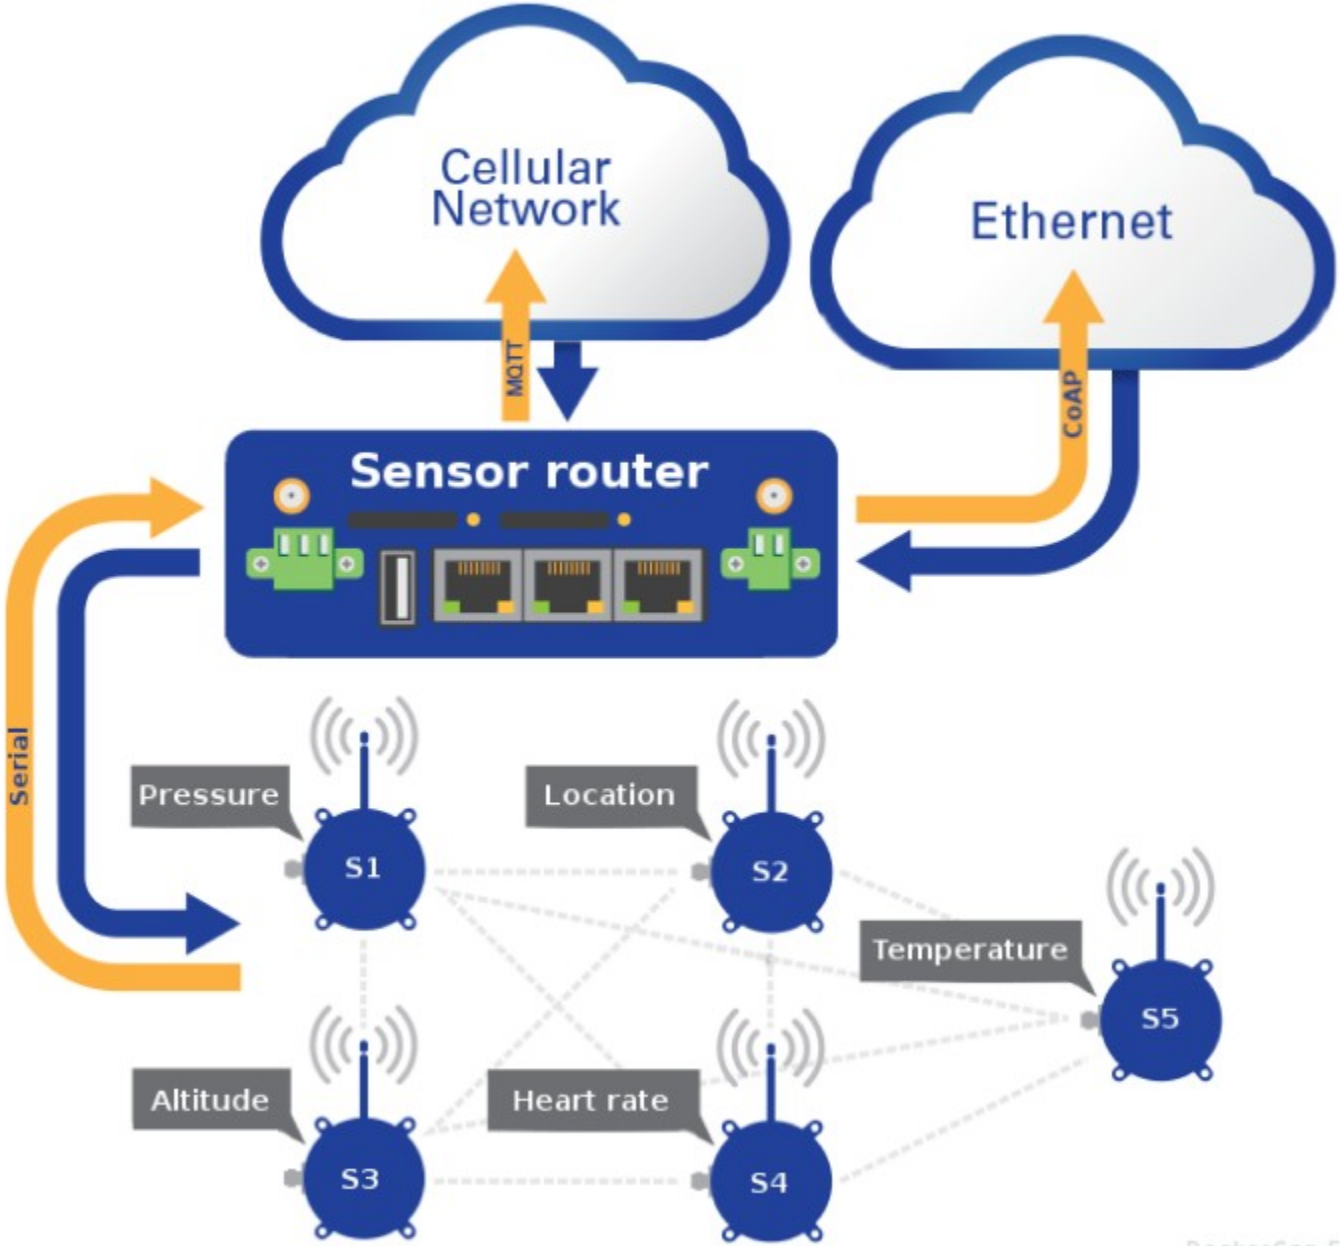
Cellular Network
Ethernet
MQTT
CoAP
Sensor router
Serial
Pressure
S1
Location
S2
Temperature
S5
Altitude
S3
Heart rate
S4

- 1 Million devices
- 1 tweet every 5 seconds

2.4PB / day

Lambda Architecture

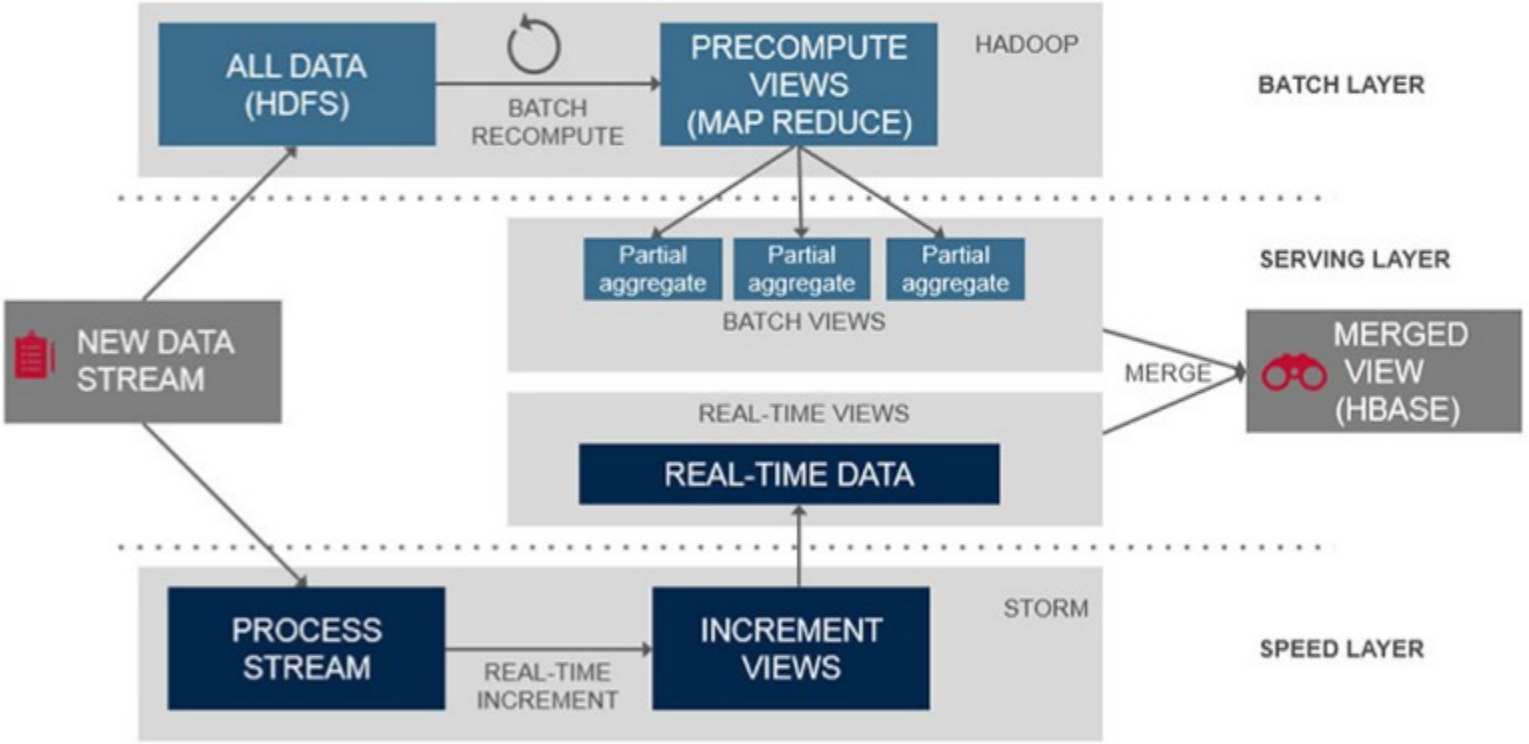

Source https://www.mapr.com

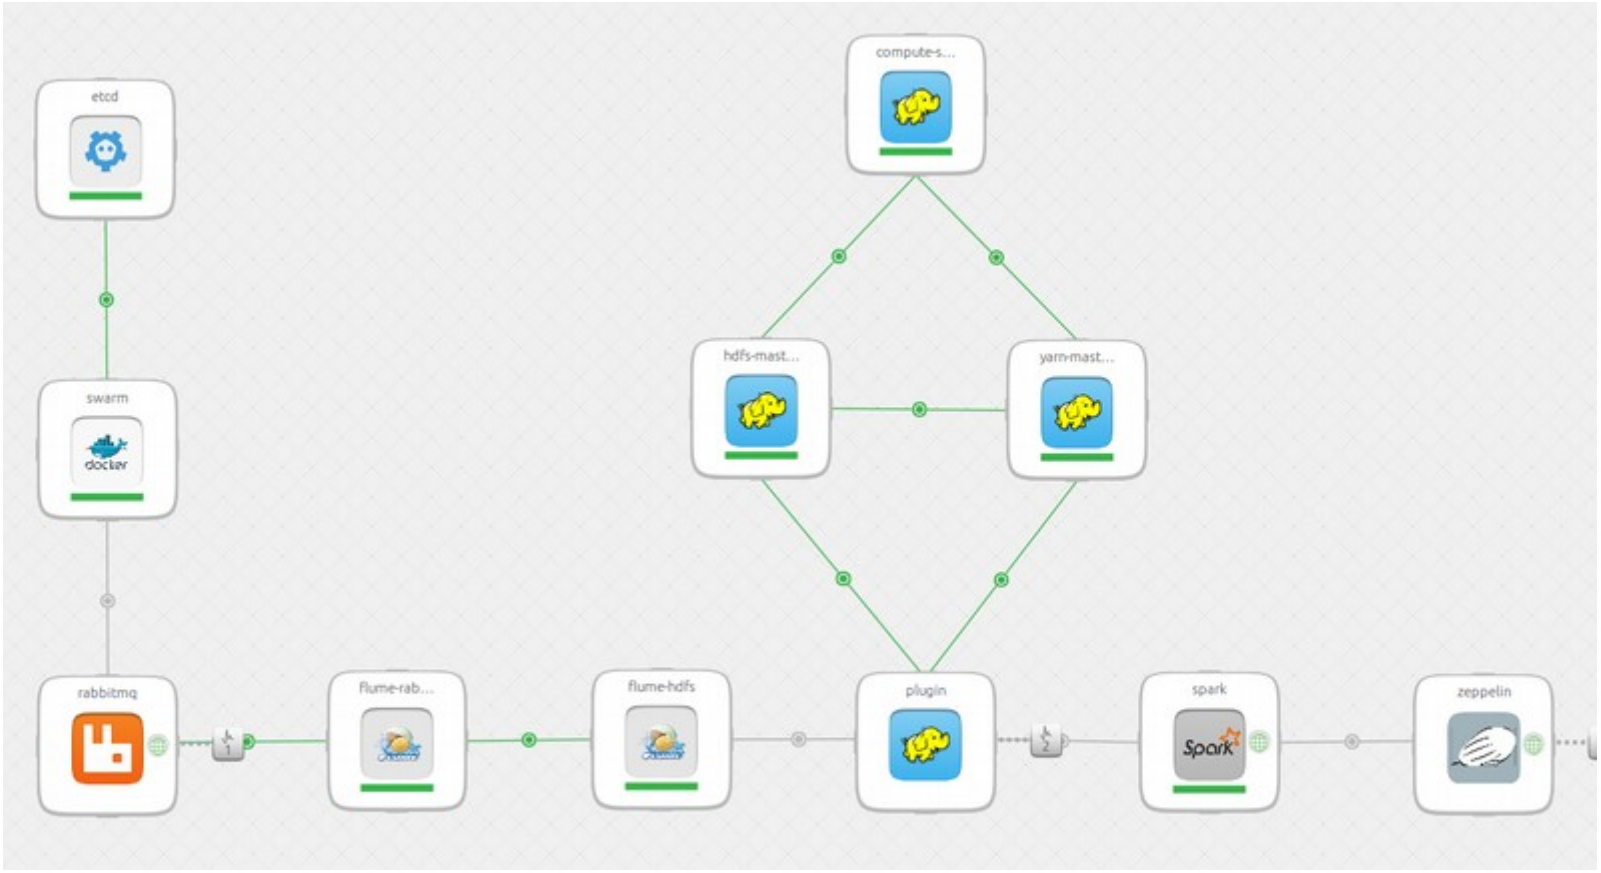
etcd
compute-s...
hdfs-mast...
yarn-mast...
swarm
rabbitmq
flume-rab...
flume-hdfs
plugin
spark
zeppelin

RabbitMQ

Overview | Connections | Channels | Exchanges | Queues | Admin

Queue rabbitmq

Overview

Queued messages (chart: last minute) (?)

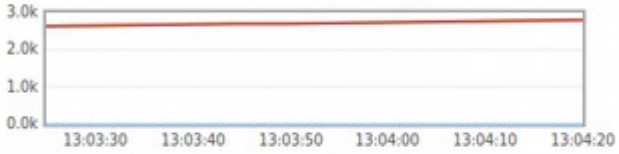

Ready	2,801
Unacked	0
Total	2,801

Message rates (chart: last minute) (?)

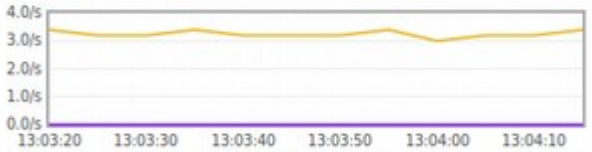

Publish	3.4/s	Acknowledge	0.00/s
Deliver	0.00/s	Get	0.00/s
Redelivered	0.00/s		

Details

Features	State	running
Policy	Consumers	0
	Consumer utilisation (?)	N/A

	Total	Ready	Unacked	In memory	Persistent
Messages (?)	2,801	2,801	0	2,801	0
Message body bytes (?)	489kB	489kB	0B	489kB	0B
Process memory (?)	3.9MB				

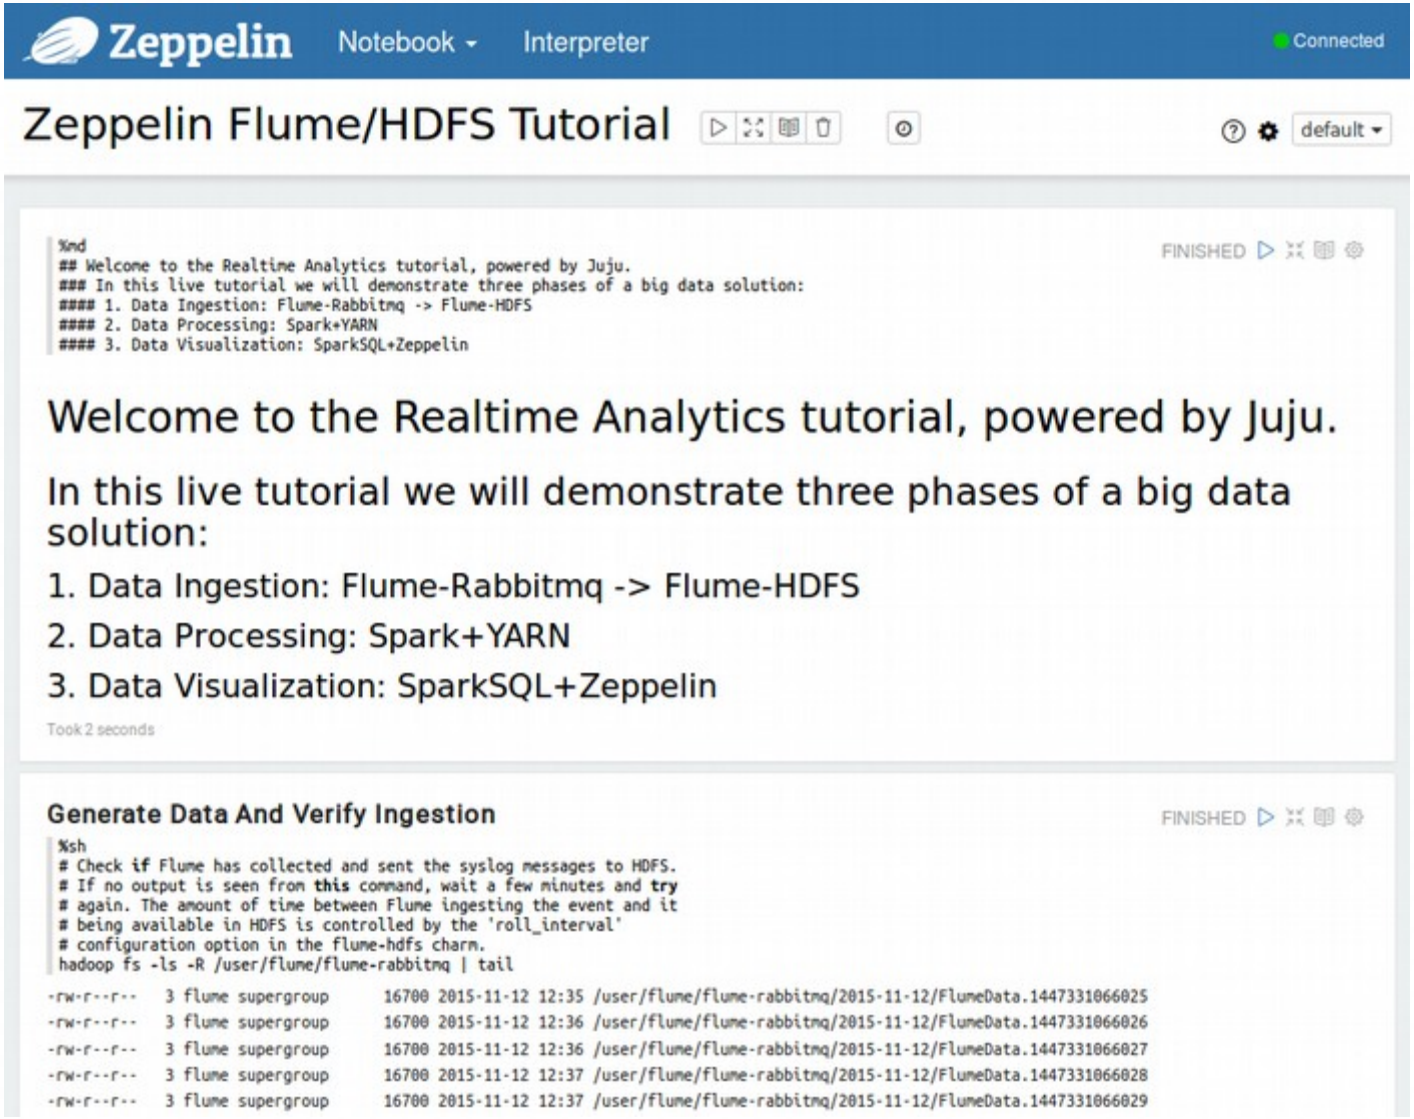
Zeppelin
Notebook
Interpreter
Connected
Zeppelin Flume/HDFS Tutorial
default
%md
Welcome to the Realtime Analytics tutorial, powered by Juju.
In this live tutorial we will demonstrate three phases of a big data solution:
1. Data Ingestion: Flume-Rabbitmq -> Flume-HDFS
2. Data Processing: Spark+YARN
3. Data Visualization: SparkSQL+Zeppelin
FINISHED
Welcome to the Realtime Analytics tutorial, powered by Juju.
In this live tutorial we will demonstrate three phases of a big data solution:
1. Data Ingestion: Flume-Rabbitmq -> Flume-HDFS
2. Data Processing: Spark+YARN
3. Data Visualization: SparkSQL+Zeppelin
Took 2 seconds
Generate Data And Verify Ingestion
FINISHED
%sh
Check if Flume has collected and sent the syslog messages to HDFS.
If no output is seen from this command, wait a few minutes and try
again. The amount of time between Flume ingesting the event and it
being available in HDFS is controlled by the 'roll_interval'
configuration option in the flume-hdfs charm.
hadoop fs -ls -R /user/flume/flume-rabbitmq | tail
-rw-r--r-- 3 flume supergroup 16700 2015-11-12 12:35 /user/flume/flume-rabbitmq/2015-11-12/FlumeData.1447331066025
-rw-r--r-- 3 flume supergroup 16700 2015-11-12 12:36 /user/flume/flume-rabbitmq/2015-11-12/FlumeData.1447331066026
-rw-r--r-- 3 flume supergroup 16700 2015-11-12 12:36 /user/flume/flume-rabbitmq/2015-11-12/FlumeData.1447331066027
-rw-r--r-- 3 flume supergroup 16700 2015-11-12 12:37 /user/flume/flume-rabbitmq/2015-11-12/FlumeData.1447331066028
-rw-r--r-- 3 flume supergroup 16700 2015-11-12 12:37 /user/flume/flume-rabbitmq/2015-11-12/FlumeData.1447331066029

dockercon EU
BARCELONA • 16-17 NOVEMBER 2015

Thank you!

Michael Schloh Von Bennewitz

michael@schloh.com
Europalab Networks

Samuel Cozannet

samuel.cozannet@canonical.com
Canonical, Ltd.

http://lect.europalab.com/dociotsen/