

Monitoring system for large and federated datacenters

Gioacchino Vino



OUTLINE

- **Initial development:** Dashboard for ALICE computing in Italy
- **Evolution:** Monitoring for large and distributed centers
- **Application for O2:** Contribution to WP8 (modular stack)
- Outlook

DASHBOARD FOR THE ALICE COMPUTING IN ITALY

Motivation:

- Concentrate in a single graphical interface all the information concerning the ALICE activity in each site (MonALISA, local Batch system, local Monitoring system metrics)
- Concentrate in a custom graphical interface all the needed information concerning the ALICE activity in Italy
- Provide a better debug tool using real-time value coming from **multiple sources**


DASHBOARD FOR THE ALICE COMPUTING IN ITALY

- The Bari site was used as testbed and the Dashboard is active and running from Oct 2014
- Currently it is running in all ALICE T2 and WLCG sites in Italy from Nov 2016
- Presented to CHEP' 16

DASHBOARD FOR THE ALICE COMPUTING IN ITALY

The Dashboard system consists of :

- **InfluxDB**, an open source time-series database
- **Grafana**, dashboard builder with powerful visualization features for time series data
- **Sensors**, python scripts able to gather data from datasources and send them to the database



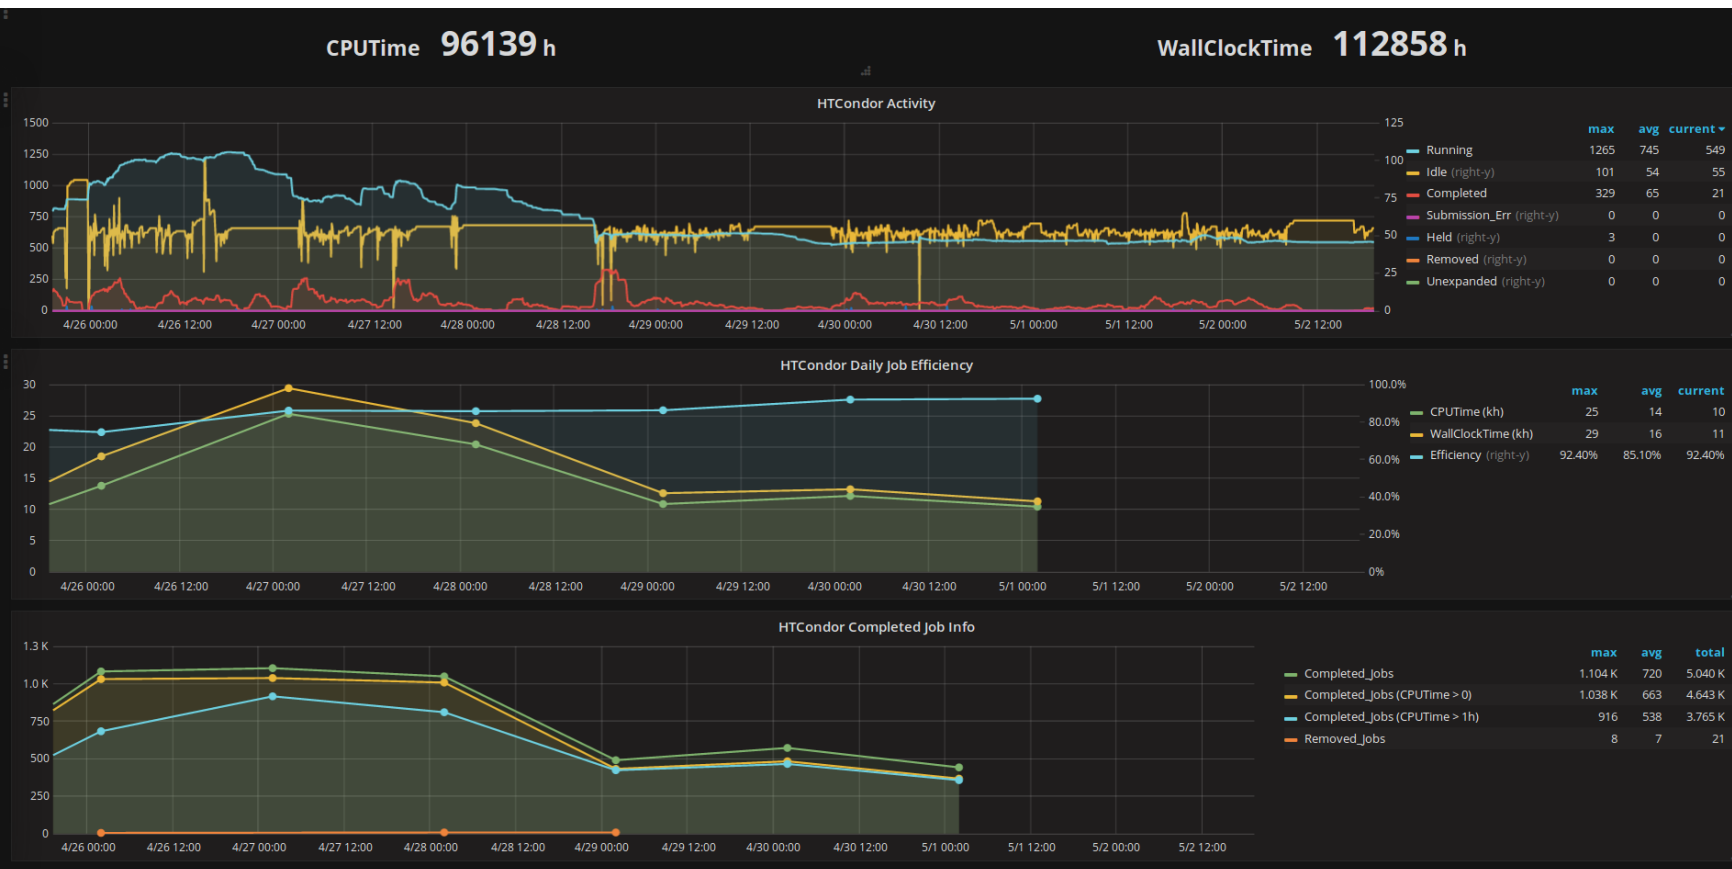
DASHBOARD FOR THE ALICE COMPUTING IN ITALY

Bari Storage activity



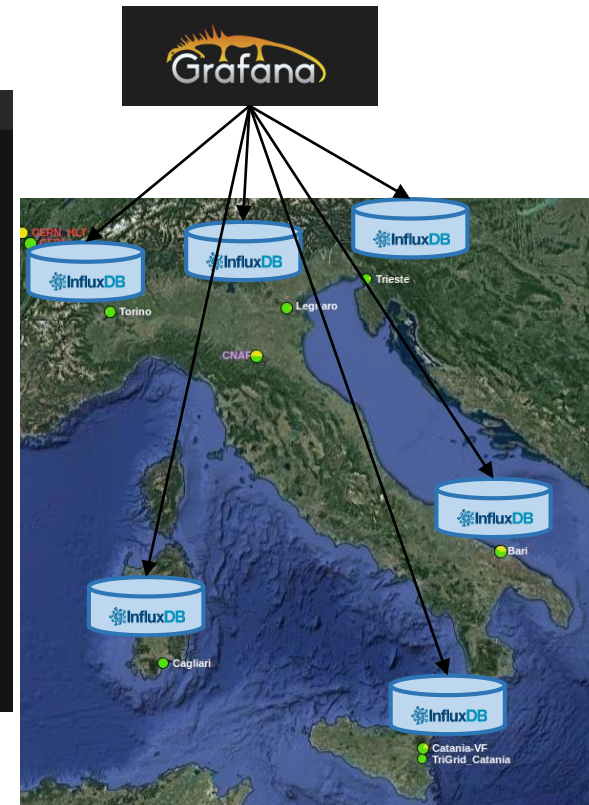
DASHBOARD FOR THE ALICE COMPUTING IN ITALY

Bari Batch system activity



DASHBOARD FOR THE ALICE COMPUTING IN ITALY

Italian computing activity



MONITORING FOR LARGE AND DISTRIBUTED CENTERS

Designing of a monitoring system able to support the management of large and distributed datacenters

Key features:

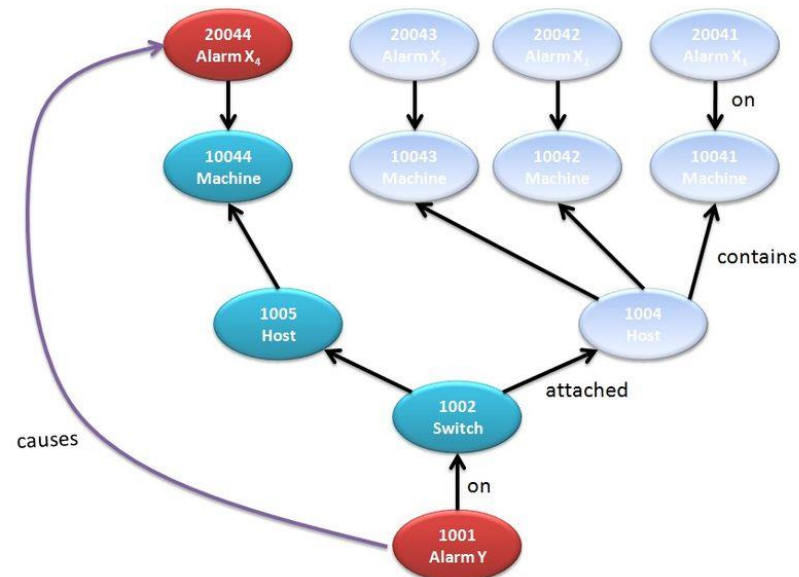
- Collecting heterogenous data from different data sources:
 - Services
 - Cloud platform (OpenStack)
 - Hardware Devices
- Analysis on the gathered data:
 - Anomaly Detector
 - Root Cause Analysis

MONITORING FOR LARGE AND DISTRIBUTED CENTERS

- Anomaly Detector



- Root Cause Analysis



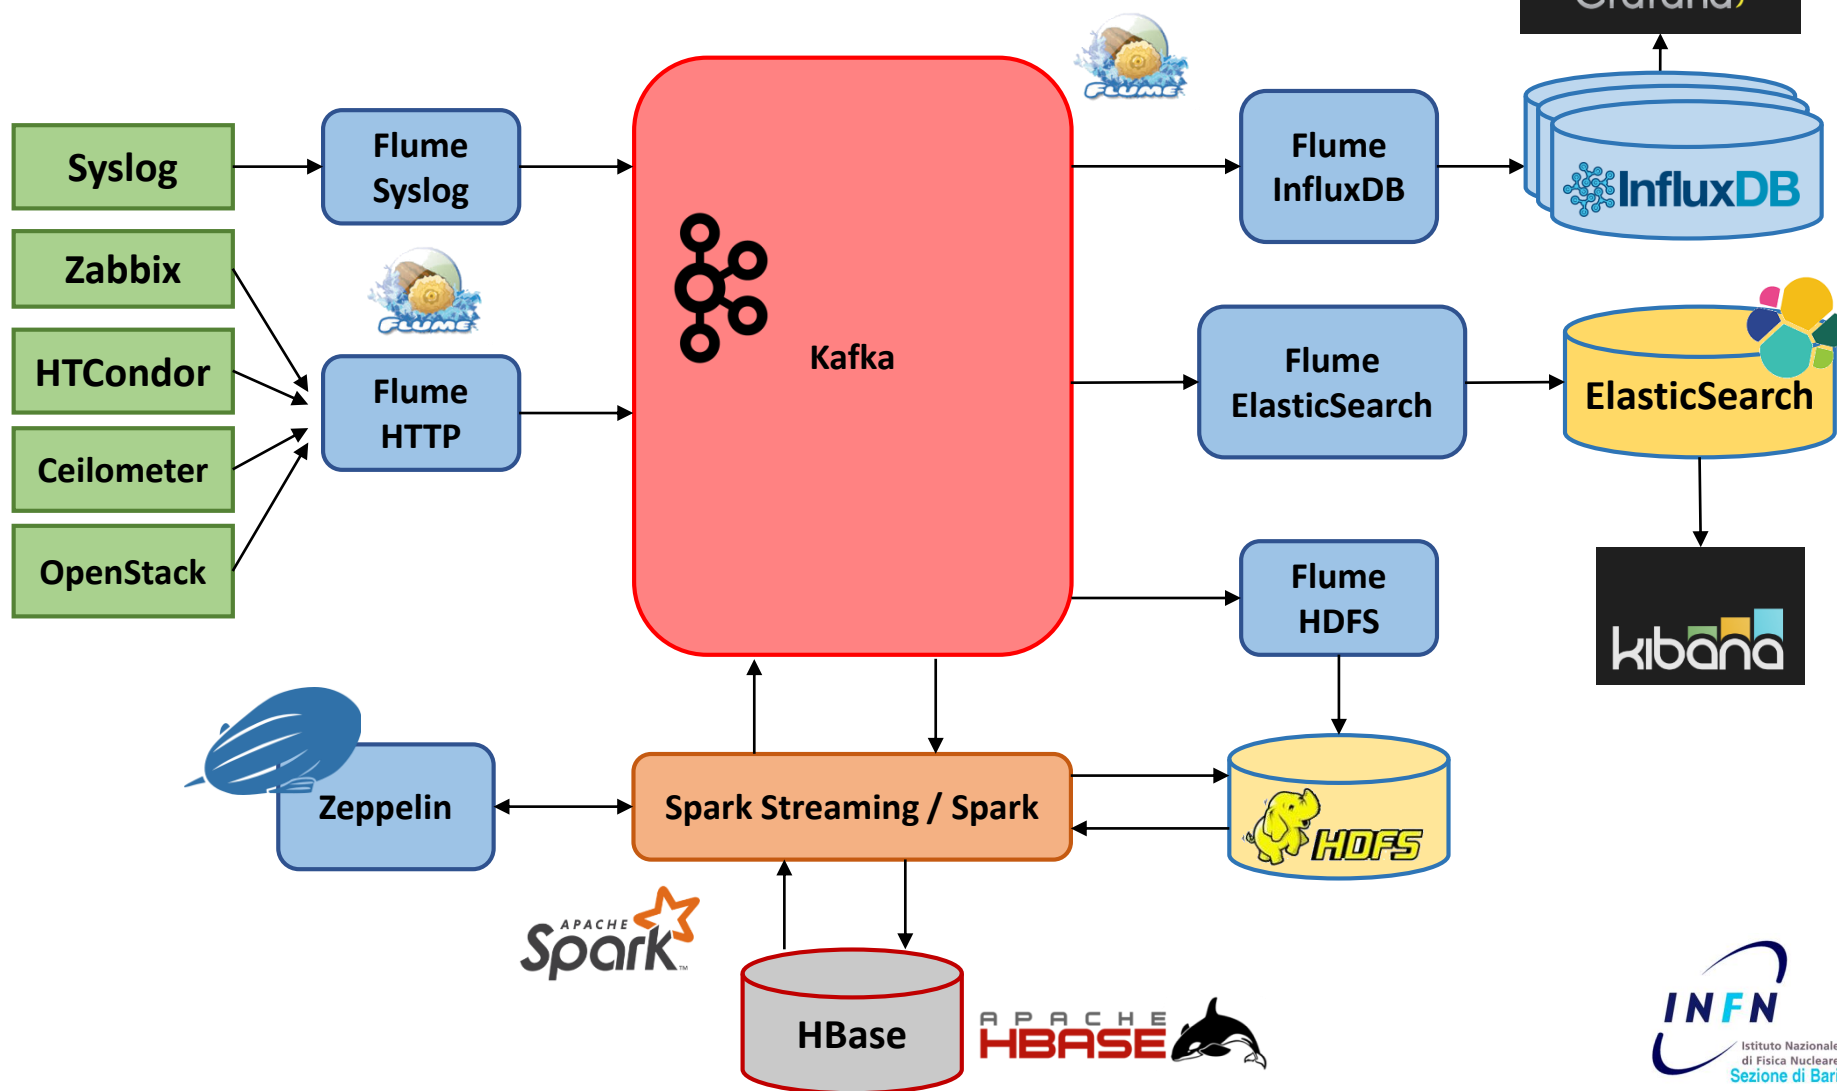
MONITORING FOR LARGE AND DISTRIBUTED CENTERS

Testbed: Datacenter ReCaS in Bari

- 128 server with 8192 cores
- Disk space: 3.5 PB
- Tape: 2.5 PB
- Cloud platform: OpenStack
- Cluster HPC composed of 20 servers with 800 cores



MONITORING FOR LARGE AND DISTRIBUTED



MONITORING FOR LARGE AND DISTRIBUTED CENTERS

Data sources:

- Syslog:
 - Information on system processes
 - 5 - 6 million of logs per day
 - Stored more than 70 GB starting from 18 November 2016

MONITORING FOR LARGE AND DISTRIBUTED CENTERS

Data sources:

- Syslog
- Zabbix:
 - Resource usage of nodes, information on OpenStack components and services
 - Sensor written in Python
 - Sampled 42000 values every 10 minutes
 - Collected 3 GB starting from 19 July 2016

MONITORING FOR LARGE AND DISTRIBUTED CENTERS

Data sources:

- Syslog
- Zabbix
- HTCondor:
 - Scheduler states, completed and running job information
 - Sensor written in Python
 - Sampled 750000 values every 5 minutes
 - Collected 11 GB starting from 18 July 2016

MONITORING FOR LARGE AND DISTRIBUTED CENTERS

Data sources:

- Syslog
- Zabbix
- HTCondor
- Openstack + Ceilometer:
 - Resource usage and services information
 - Sensor being written in Python

MONITORING FOR ...

Transport layer:

- Apache Flume

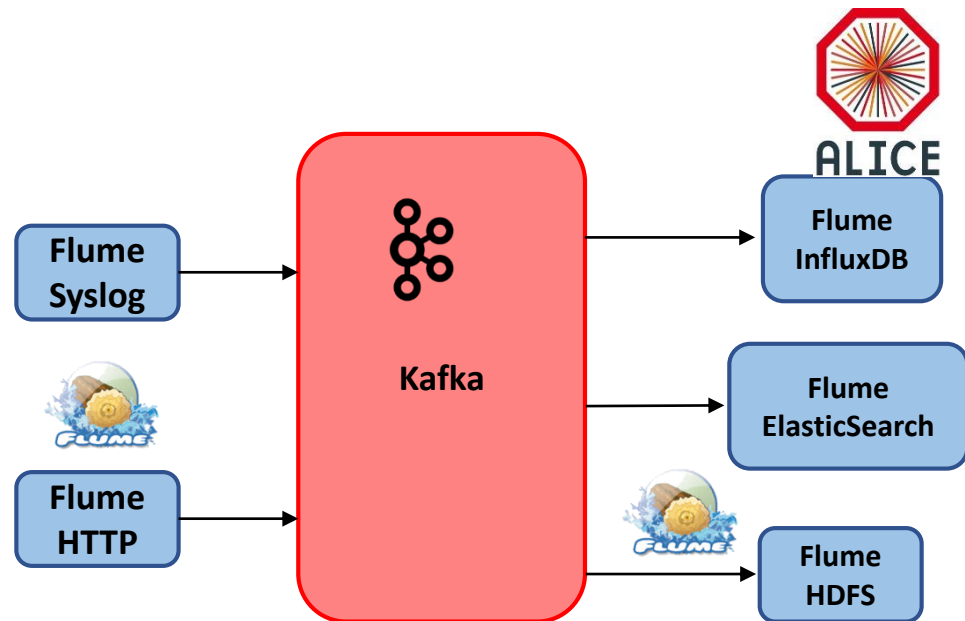
- Distributed, reliable, and available service for efficiently collecting, aggregating and moving large amounts of log data.
- Robust, fault tolerant and provides ready-to-use interfaces

- Apache Kafka

- Distributed streaming platform, reliable and allows data replication on multiple nodes

- Apache Flume + Kafka (aka Flafka)

- Take advantage of both



MONITORING FOR LARGE AND DISTRIBUTED CENTERS

Storage:

- HDFS (Hadoop Distributed File System)
 - Used as long term storage of batch jobs
- HBase
 - Very fast key-value database on top of HDFS
 - Serve real-time requests
- InfluxDB
 - With Grafana, used to visualize time-series data
- ElasticSearch
 - With Kibana, used to plot information about log data



MONITORING FOR LARGE AND DISTRIBUTED CENTERS

Processing Components:

- Apache Spark:
 - Execute batch jobs on data stored in HDFS
- Apache Spark Streaming:
 - Execute real-time analysis on acquired data

Support Components:

- Spark SQL, Spark GraphX, Spark MLlib, Apache Zeppelin

O2 WP8 CONTRIBUTION - MISSION

- **Data Collection** of system monitoring, infrastructure monitoring and application monitoring (~600 kHz)
- **Processing** like Data suppression, Data enrichment, Data aggregation and Data correlation.
- **Storage**
- **Graphical display**

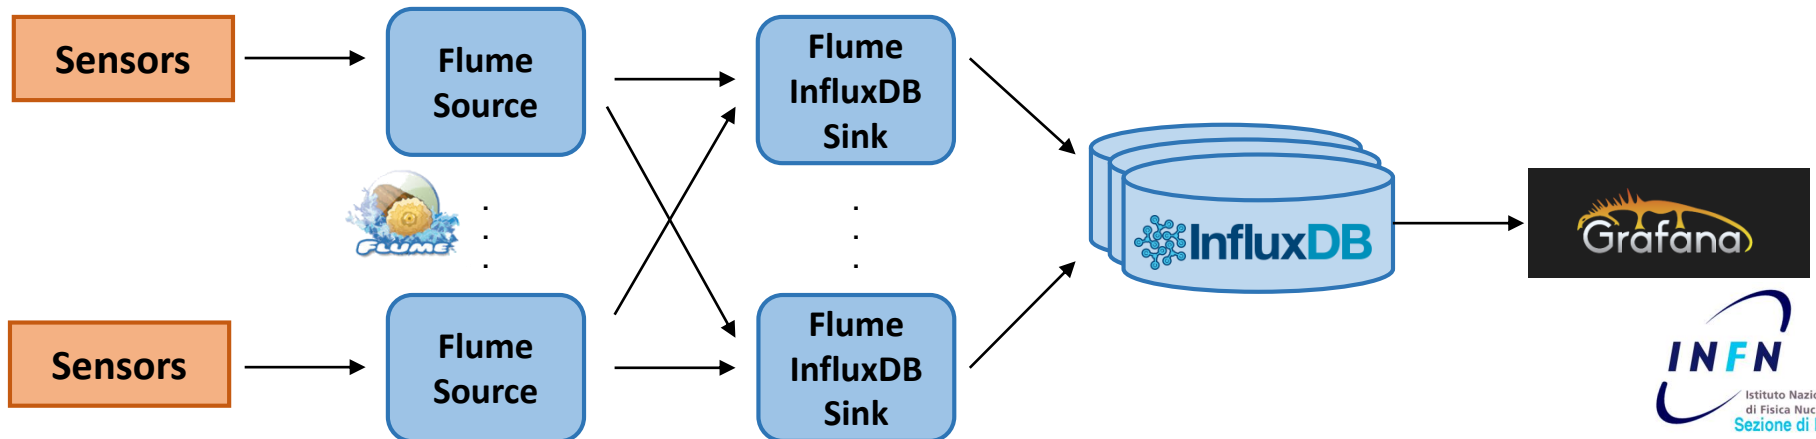
Three main alternative options currently under evaluation:

- MonALISA, Modular Stack, Zabbix

O2 WP8 CONTRIBUTION – MODULAR STACK

Different tools used to accomplish the goal:

- CollectD, used to collect host information
- Apache Flume, used as transport layer
- InfluxDB, used as TimeSeries Database
- Grafana, used as Dashboard for Timeseries data



OUTLOOK

- Implement algorithms for **Anomaly Detector** and **Root Cause Analysis**
 - Use **Apache Mesos** or **DC/OS** as resource manager
 - Design and implement **bottleneck analysis**
 - Test the project on multiple datacenters
-
- Finalize system choice for **O2 monitoring**
 - Upgrade of the **Dashboard of ALICE activity in Italy** using the knowledge acquired on Apache components

**THANKS
FOR YOUR
ATTENTION**