Monitoring system for geographically distributed datacenters based on Openstack

Gioacchino Vino

INFN Bari

Tutor: Dott. Domenico Elia
Tutor: Dott. Giacinto Donvito
Borsa di studio GARR “Orio Carlini” 2016-2017

INFN BARI
Istituto Nazionale di Fisica Nucleare
Sezione di Bari

Index

- Motivations
- Project Overview
- Next steps
- Parallel derived Activity:
  - Implementing the monitoring system for the ALICE O2 Facility @ CERN

Motivations for a new monitoring system

- The **increasing** demand of computation resources for scientific purposes is leading to:
  - Datacenters increasing in **complexity** and **size.**
  - Taking advantages of **new technologies** like virtualization and cloud computing.
  - Datacenter **cooperation** needed in order to accomplish common goals.

- Geographically distributed datacenters
  - Goal: Increase the **computation capability** of the overall system.
  - Side effect: Increasing **complexity** for monitoring and control systems.

Project: Developing a monitoring system for geographically distributed datacenters.

Project Overview

- Goal:
    - Reduce the malfunction time
- Advances features are requested:
    - Anomaly detection
    - Root Cause Analysis
    - Graph data modeling
- Fully informative monitoring data are collected:
    - Service monitoring (HTTP server, DBs, ...)
    - Openstack and middleware monitoring
    - Hardware monitoring (servers, disks, disk controllers, network devices, PDU, ...)

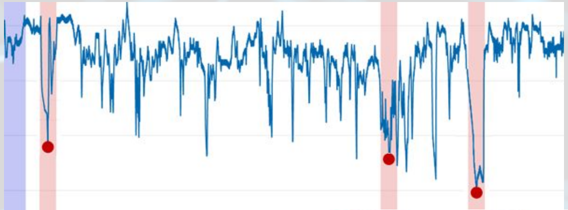

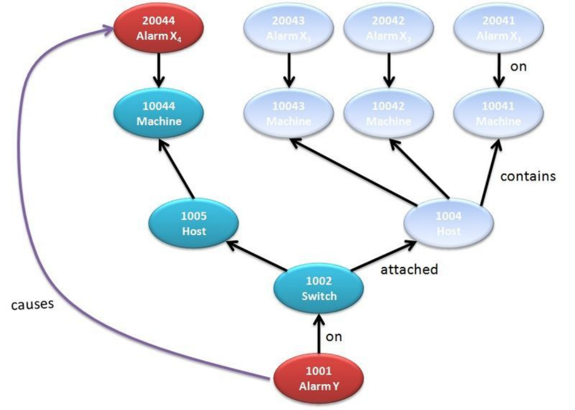

Project Testbed: ReCaS Bari

ReCaS Bari Datacenter

- More than 13.000 cores
- 7.1 PB Disk Storage
- 2.5 PB Tape storage
- HPC Cluster composed of 20 servers
- Dedicated network link: 10Gbps x2 to GARR, 20Gbps to Naples and 20Gbps to Bologna
- Cloud platform: OpenStack
- Batch system: HTCondor
  - 184 Worker Nodes
  - 350+ network connections
- Local Monitoring System: Zabbix
- Including ALICE and CMS Tier2s

Project Architecture

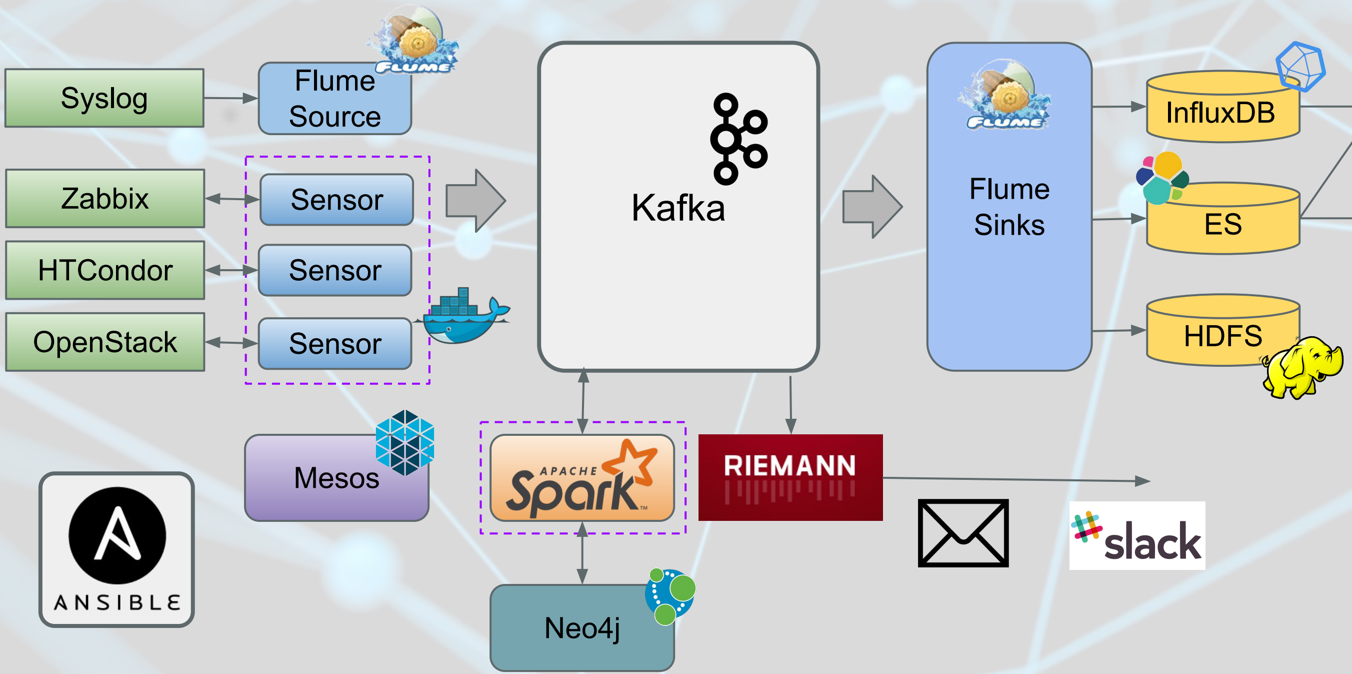

Key words:

- BigData
- Open Source
- Graph data modeling
- Machine Learning

Project Architecture

Syslog

Zabbix

HTCondor

OpenStack

Data Sources:

- **Syslog**: System processes and service information.
  - 2-3 millions of lines collected every day (500GB per year)
- **Zabbix**: Computation resource usage, service and Openstack monitoring.
  - +40k values sampled every 10 minutes (8GB from 19.07.2016)
- **HTCondor**: Scheduler, completed and running job state
  - +750k values sampled every 5 minutes (35GB from 19.07.2016)
- **OpenStack**: Information on server, images, flavors, volumes, network devices, ...
  - +50k values monitored every hour

Project Architecture

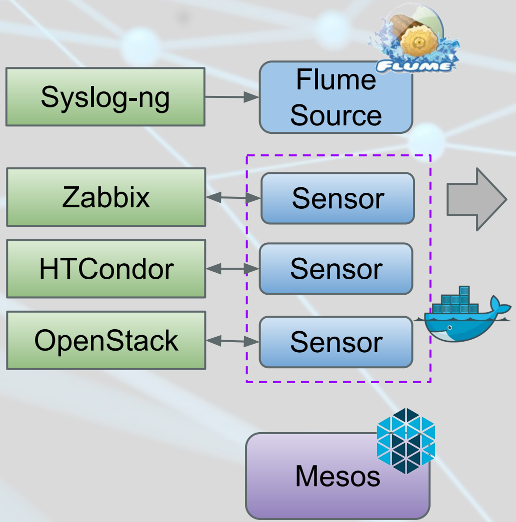

Metric collectors:

- **Apache Flume** Syslog Source.
- Python code inserted in **Docker**-container and executed periodically using **Apache Mesos.**

Apache Flume: a distributed and highly-reliable service for collecting, aggregating and moving large amounts of data in a very efficient way.
Apache Mesos: an open-source project to manage computer clusters.
Docker: a computer program that performs operating-system-level virtualization also known as containerization.

INFN BARI Istituto Nazionale di Fisica Nucleare Sezione di Bari

Project Architecture

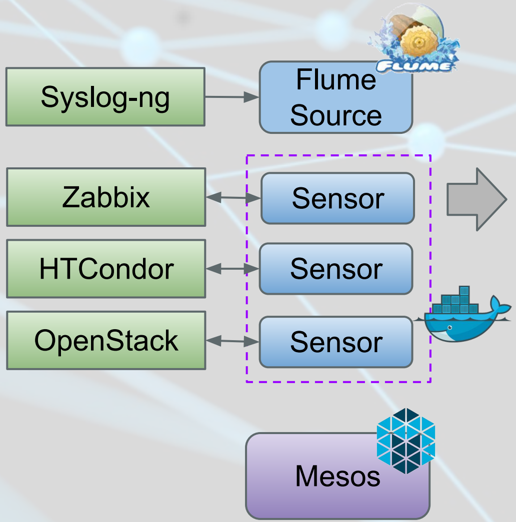

Transport Layer:

- Decouple all components.
- Increase the High Availability of system.

Apache Kafka: an open-source stream-processing software platform, provides a unified, high-throughput, low-latency platform for handling real-time data feeds.

Project Architecture

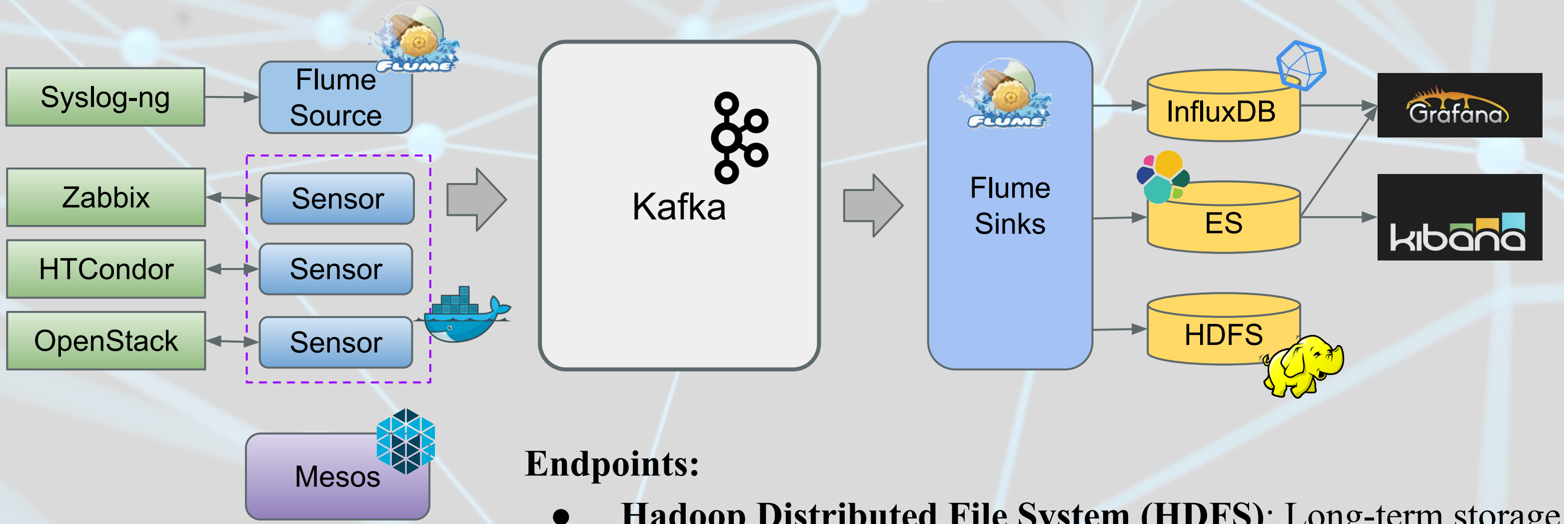

Endpoints:

- **Hadoop Distributed File System (HDFS)**: Long-term storage.
- **InfluxDB-Grafana**: Timeseries Dashboards.
- **ElasticSearch-Kibana**: Log Dashboards.

Project Architecture

Timeseries Dashboards

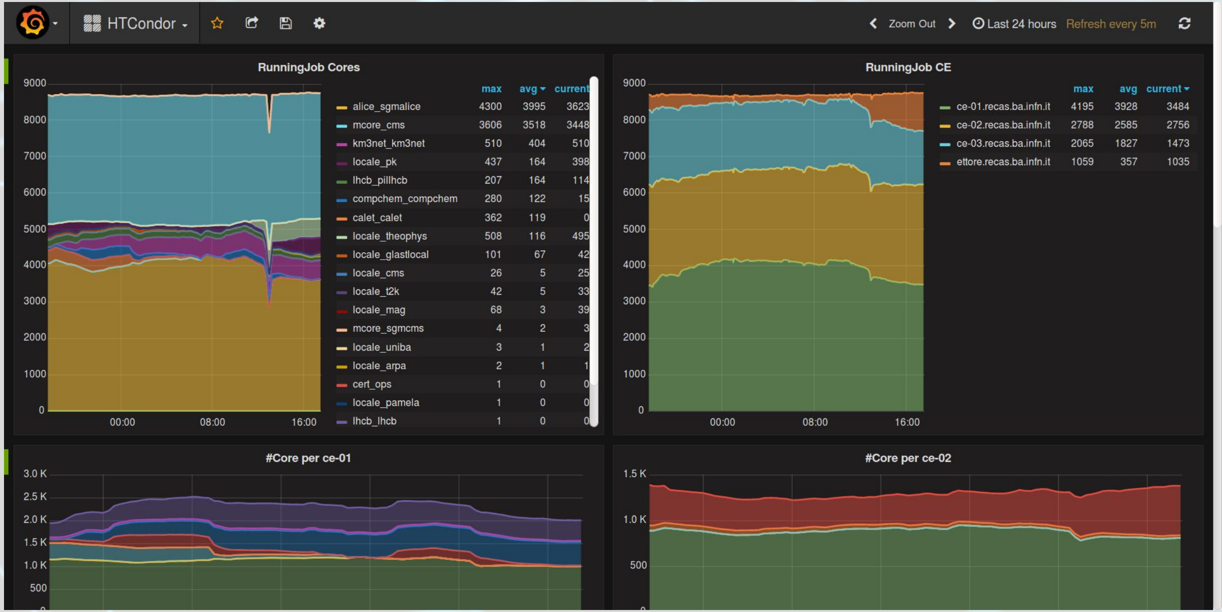

InfluxDB: a custom high-performance data store written specifically for time series data.
Grafana: Dashboards' builder for time-series data.

Project Architecture

Log Dashboards

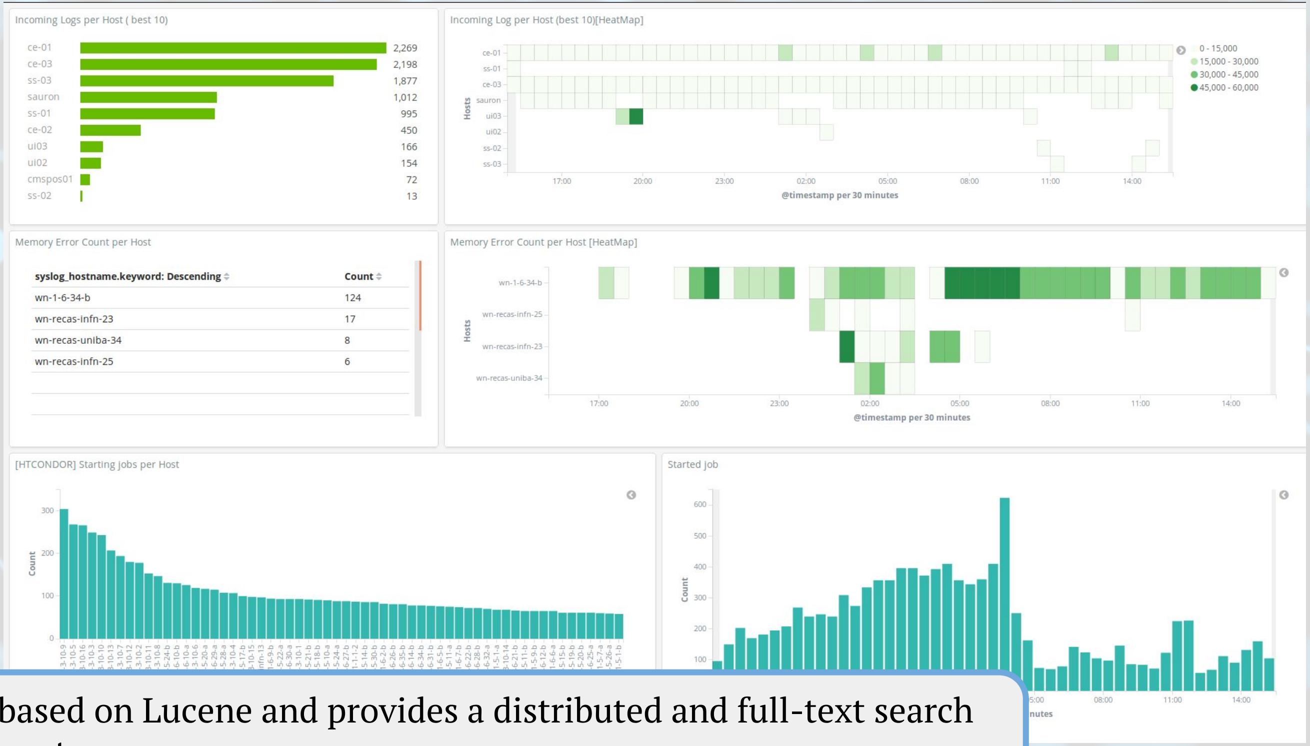

ElasticSearch: a search engine based on Lucene and provides a distributed and full-text search engine schema-free JSON documents.
Kibana: an open source data visualization plugin for Elasticsearch

INFN BARI
Istituto Nazionale di Fisica Nucleare
Sezione di Bari

Project Architecture

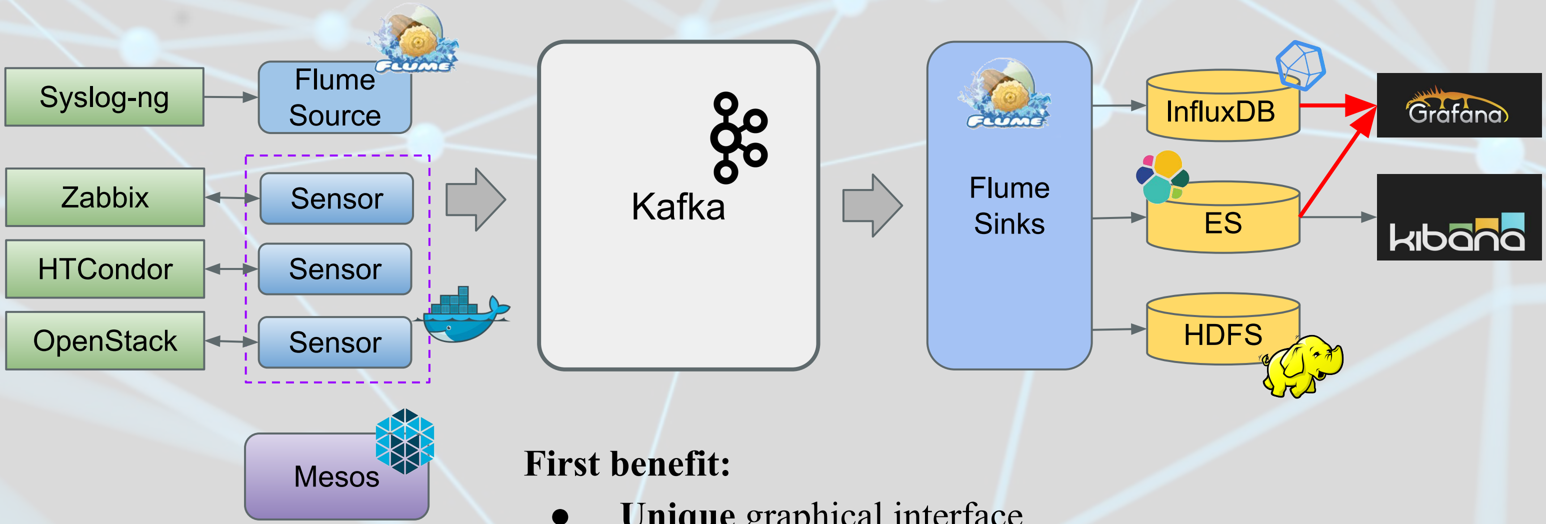

First benefit:

- **Unique** graphical interface

INFN BARI
Istituto Nazionale di Fisica Nucleare
Sezione di Bari

Project Architecture

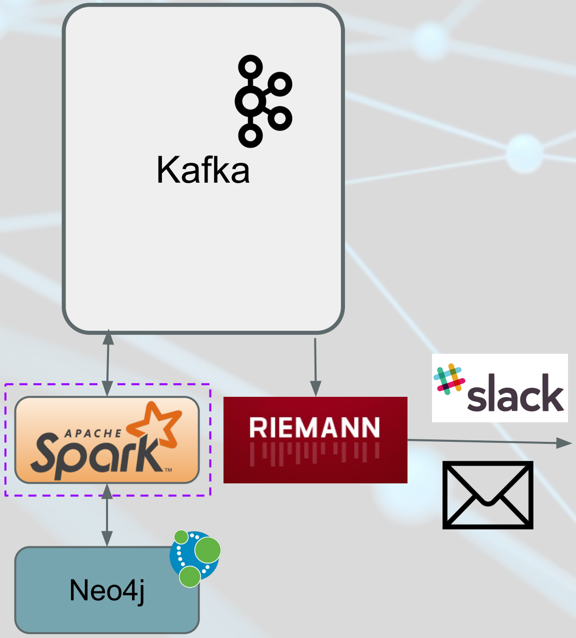

Riemann: aggregates events from your servers and applications with a powerful stream processing language.

Alarm dispatcher:

- Plugins: Email, Slack.
- Processes and filters events.

Information Structure:

- Classical monitoring is not enough.
- Relation information (Services, network, virtual-physical server, …)
    - Openstack data.
    - Open connections.
    - Other monitoring data.
- Advantages:
    - Eliminating the need for joins

Neo4j: High Performance native Graph Storage & Processing.

Project Architecture

Information Structure:

Subgraph example:

- Blues nodes: virtual machines.
- Yellow nodes: network interfaces.
- Red nodes: networks.

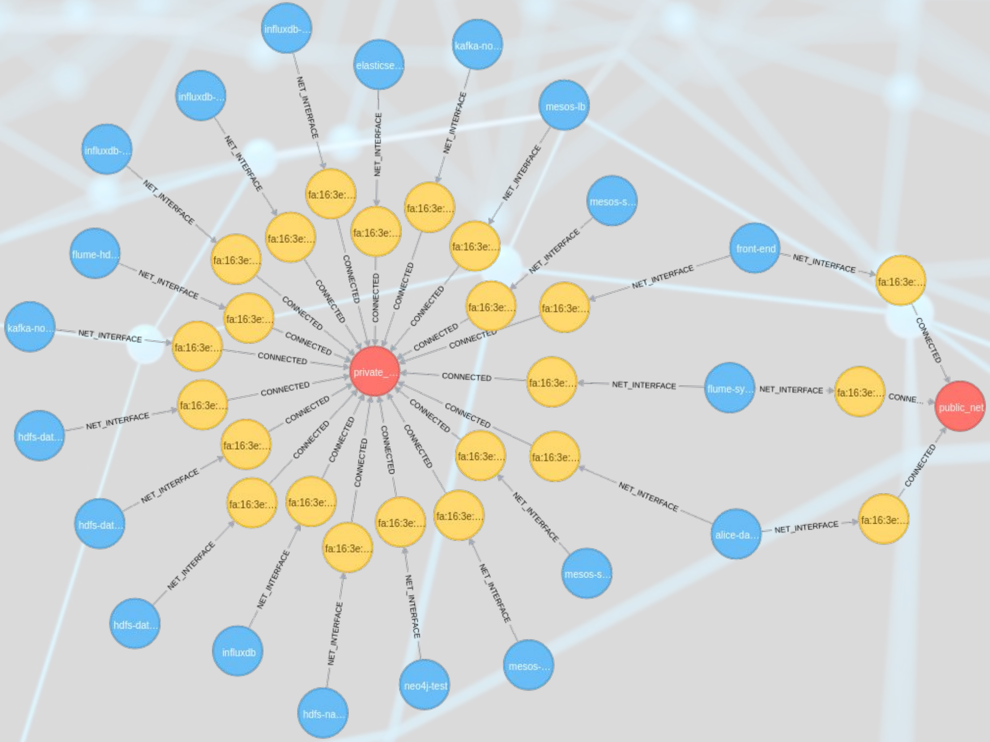

MATCH (s:server)-[:NET_INTERFACE]->(n:nic)-[CONNECTED]->(net:network)
WHERE s.name = “frontend” AND net.name = “public”
RETURN n.mac

Project Architecture

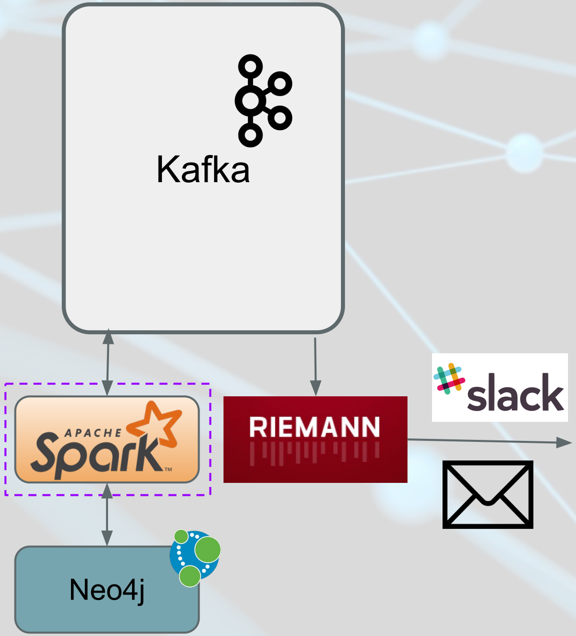

Processing Unit:

- Log Analyzer (Streaming and Batch).
- Anomaly Detection.
- Data Correlation.
- Root Cause Analysis.
- Pattern recognition.
- Graph Data Modeling.

Apache Spark: a fast and general engine for large-scale data processing.

Project Machine Learning Algorithms

Features:

- Adaptable to the data size.
- Combination of unsupervised and supervised algorithms.
- Incremental learning.
- Knowledge sharing.

Next steps

- Migrate all components in Mesos
- Improve the Machine Learning algorithms effectiveness
- Root Cause Analysis algorithm
- Integration with project management systems (OpenProject, Trello, ...)
- Actions

Monitoring system for the ALICE O2 Facility

- **ALICE** is a heavy-ion detector designed to study the physics of strongly interacting matter (the Quark–Gluon Plasma) at the CERN Large Hadron Collider (LHC).
- During the Long Shutdown 2 in the end of 2018, ALICE will start its **upgrade** to fully exploit the **increase** in luminosity.
- The current computer system (Data Acquisition, High-Level Trigger and Offline) will be replaced by a single, common **O2 (Online-Offline)** system.
- Some detectors will be **read out continuously**, without physics triggers.
- O2 Facility will **compress** the 3.4 TB/s of raw data to 100 GB/s of reconstructed data:
  - 268 First Level Processors
  - 500 Event Processing Nodes

- Development of a Monitoring System for ALICE O2 Facility:
  Modular Stack solution, with components and tools already used and tested in the MonGARR project
  (approved by the ALICE O2 TB last February)

Monitoring system for the ALICE O2 Facility

Requirements:

- Capable of handling O2 monitoring traffic – 600 kHz
- Scalable >> 600 kHz
- Low latency
- Compatible with CentOS 7
- Open Source, well documented, actively maintained and supported by developers
- Impose low storage size per measurement

Goals:

- Real-time Dashboards
- Historical Dashboards
- Alarm dispatcher

Monitoring system for the ALICE O2 Facility

Architecture:

- Sensors:
  - **Monitoring Library**
  - **CollectD**
- Transport Layer:
  - **Apache Flume**
- Time-series Database:
  - **InfluxDB**
- Visualization interface:
  - **Grafana**
- Alarming component:
  - **Riemann**
- Processing component:
  - **Apache Spark**

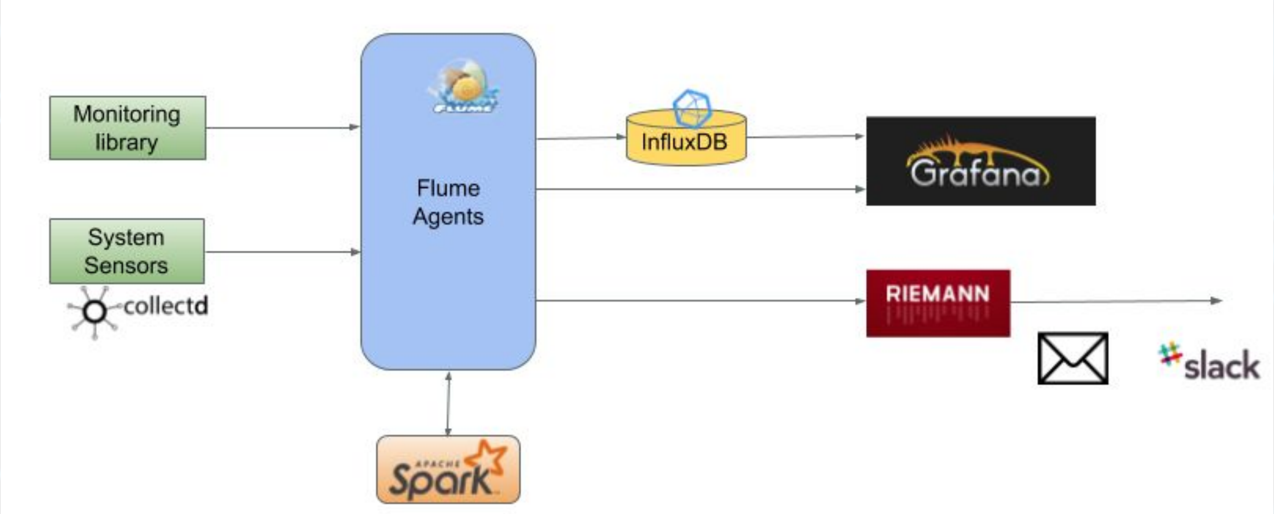

Monitoring system for the ALICE O2 Facility

Next Steps:

- System Validation using the TPC monitoring data, July 2018
- New functionalities will be added (new streaming analysis, alarming, log analysis)
- System Validation using ITS monitoring data, Dec 2018

THANKS FOR YOUR ATTENTION

INFN BARI
Istituto Nazionale di Fisica Nucleare
Sezione di Bari

Backup

Resource Usage for the monitoring system:

- 80 CPUs
- 150GB RAM
- 3 TB Disk
    - 1.5TB for HDFS in replica 3
    - 600 GB for Kafka nodes
    - No-volatile virtual machine volumes

Backup

Resource Usage for the most important hosts:

- ElasticSearch
  - 4 VCPUs
  - 8GB RAM
- InfluxDB
  - 2(8) VCPUs
  - 4(16)GB RAM

Backup

Apache Mesos:

Cluster:

- 3x Master (2 CPUs, 4GB RAM, 20GB Disk)
- 2x Slaves (4 CPUs, 8GB RAM, 20 GB Disk)
- 1x Load Balancer (2 CPUs, 4GB RAM, 20GB Disk)

Frameworks:

- Chronos
- Marathon
- Spark

Backup

Modular Stack Flume inner architecture