CERN Agile Infrastructure Monitoring

Pedro Andrade
CERN – IT/GT

HEPiX Spring 2012

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Introduction

- Motivation
  - Several independent monitoring activities in IT
    - Similar overall approach, different tool-chains, similar limitations
  - High level services are interdependent
    - Combination of data from different groups necessary, but difficult
  - Understanding performance became more important
    - Requires more combined data and complex analysis
  - Move to a virtualized dynamic infrastructure
    - Comes with complex new requirements on monitoring
- Challenge
  - Find a shared architecture and tool-chain components while preserving/improving our investment in monitoring

Monitoring @ CERN/IT

- >30 monitoring applications
  - Number of producers: ~40k
  - Input data volume: ~280 GB per day
- Covering a wide range of different resources
  - Hardware, OS, applications, files, jobs, etc.
- Application-specific monitoring solutions
  - Using different technologies (including commercial tools)
  - Sharing similar needs: aggregate metrics, get alarms, etc
- Limited sharing of monitoring data
  - Hard to implement complex monitoring queries

Architecture Constraints

- Aggregate monitoring data in a large data store
  - For storage and combined analysis tasks
- Make monitoring data easy to access by everyone
  - Not forgetting possible security constraints
- Select a simple and well supported data format
  - Monitoring payload to be schema free
- Rely on central metadata services
  - To discover the computer center resources information
  - E.g. which is the physical node running virtual machine A

Data

Architecture Constraints

- Follow a tool chain approach
  - Each tool can be easily replaced by a better one
- Provide well established solutions for each layer of the architecture: transport, storage/analysis, etc.
  - Adopt whenever possible existing tools
  - Avoid home grown solutions
- Allow a phased transition to the new architecture
  - Existing applications can be gradually integrated

Technology

Architecture Constraints

- Guarantee that all essential monitoring use cases are considered and satisfied
  - Fast and Furious (FF)
    - Get metrics values for hardware and selected services
    - Raise alarms according to appropriate thresholds
  - Digging Deep (DD)
    - Curation of hardware and network historical data
    - Analysis and statistics on batch job and network data
  - Correlate and Combine (CC)
    - Correlation between usage, hardware, and services
    - Correlation between job status and grid status

Use Cases

Architecture Overview

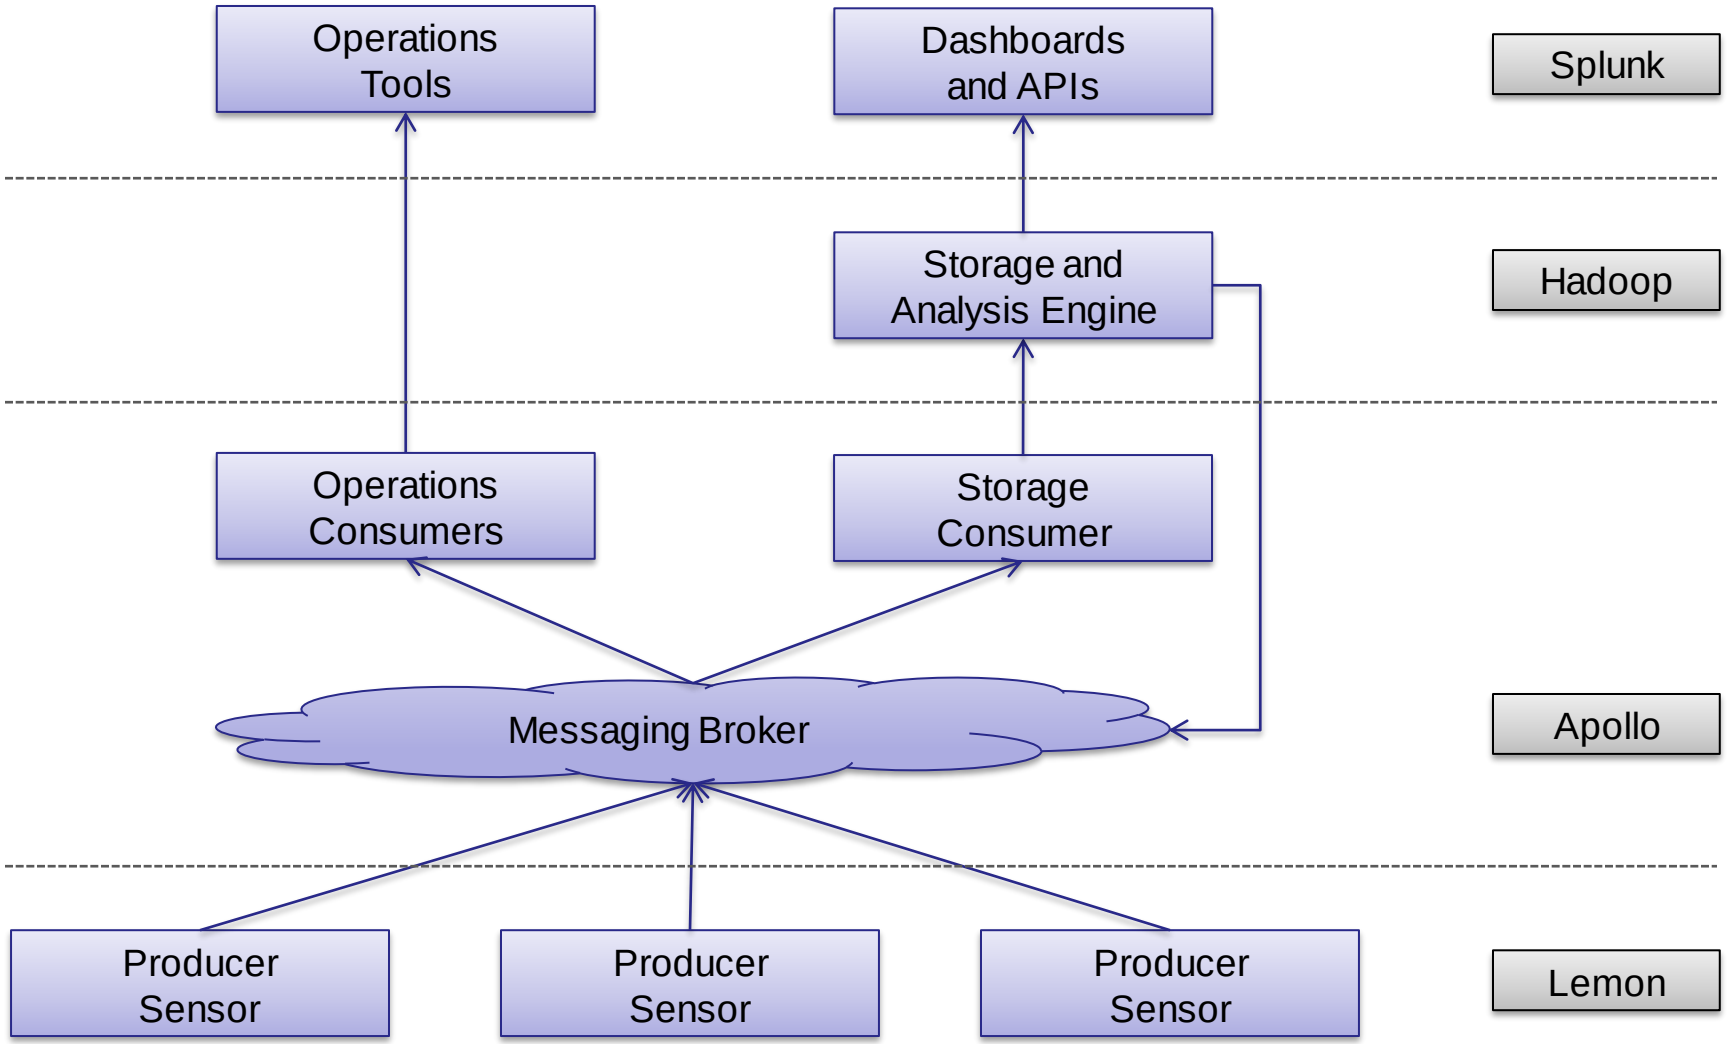

Architecture Overview

- All components can be changed easily
  - Including the messaging system (standard protocol)
- Messaging and storage as central components
  - Tools connect either to the messaging or storage
- Scalability can be addressed
  - Horizontally scaling
  - Adding additional layers (e.g. pre-aggregation)
- Messages based on a common format (JSON)
  - A simple common schema is being defined to guarantee cross-reference between data sets: timestamp, node, etc.

Producers and Sensors

- Monitoring data generated by all resources
  - Monitoring metadata available at the node
  - Published to messaging using common libraries
  - May also be produced as a result of pre-aggregation or post-processing tasks
- Support and integrate closed monitoring solutions
  - By injecting final results into the messaging layer or exporting relevant data at an intermediate stage

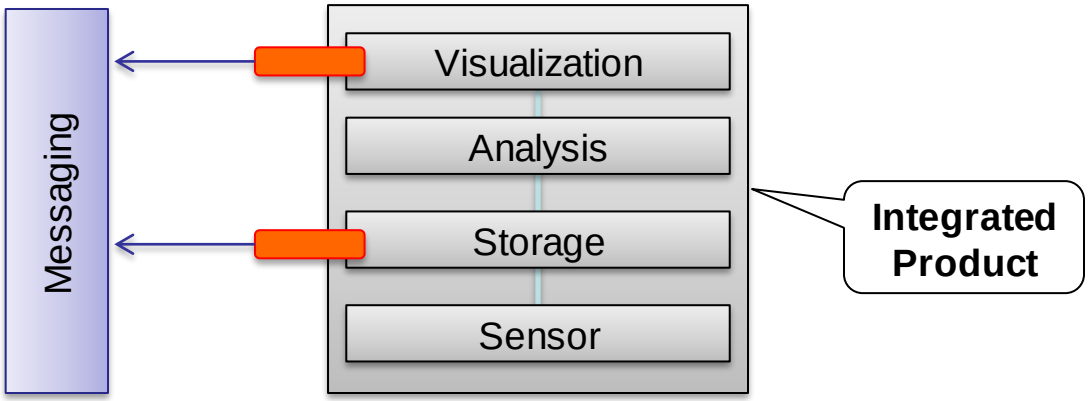

Messaging Broker

- Monitoring data transported via messaging
  - Provide a network of messaging brokers
  - Support for multiple configurations
  - The needs of each monitoring application must be clearly analyzed and defined
    - Total number of producers and consumers
    - Size of the monitoring message
    - Rate of the monitoring message
  - Realistic testing environments are required to produce reliable performance numbers
- First tests with Apollo (ActiveMQ)
  - Prior positive experience in IT and the experiments

Storage and Analysis

- Monitoring data stored in a common location
  - Easy the sharing of monitoring data and analysis tools
  - Allows feeding into the system data already processed
  - NoSQL technologies are the most suitable solutions
    - Focus on column/tabular solutions
- First tests with Hadoop (Cloudera distribution)
  - Prior positive experience in IT and the experiments
  - Map-reduce paradigm is a good match for the use cases
  - Has been used successfully at scale
  - Several related modules available (Hive, HBase)
- For particular use cases a parallel relational database solution (Oracle) can be considered

Notifications and Dashboards

CERN IT Department

- Provide an efficient delivery of notifications
  - Notifications directly sent to correct consumer targets
  - Possible targets: operators, service managers, etc.
- Provide powerful dashboards and APIs
  - Complex queries on cross-domain monitoring data
- First tests with Splunk

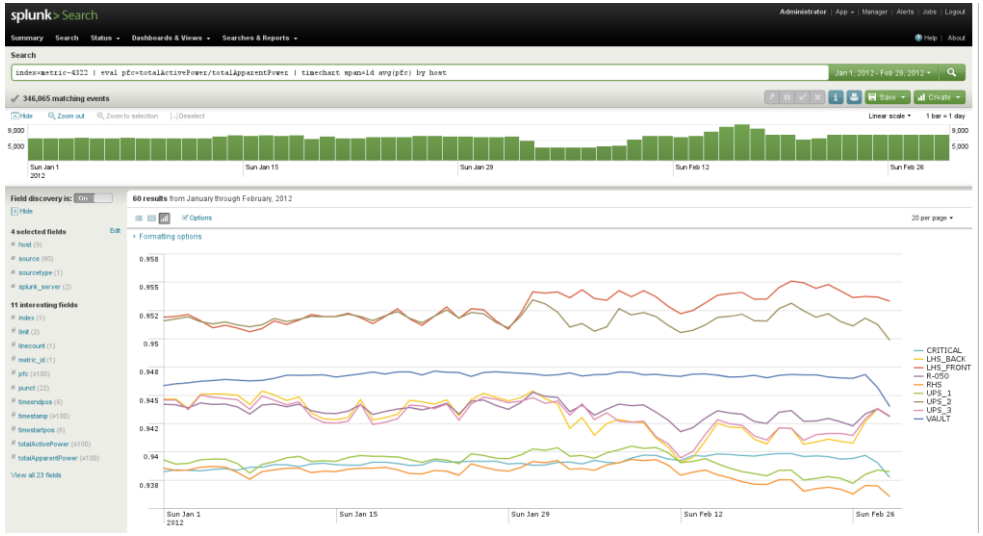

Status and Plans

Q1 2012
- Definition of monitoring model for notifications
- Definition and porting of Lemon as monitoring producer
- Deployment and initial testing of Apollo, Hadoop, Splunk

Q2 2012
- Enable operations notifications via the messaging infrastructure
- Test Splunk as dashboard for operations notifications
- Consume monitoring data to Hadoop and initial testing

Q4 2012
- Several applications publishing to messaging infrastructure
- More complex data store/analysis with different data sets on Hadoop

Q4 2013
- All monitoring data managed via the messaging infrastructure
- Large scale data store/analysis on Hadoop cluster

Summary

- A monitoring architecture has been defined
  - Promotes sharing of monitoring data between apps
  - Tool-chain approach based on few core components
  - Several existing external technologies identified

- A concrete implementation plan is ongoing
  - It assures a smooth transition for today's applications
  - It enables the AI to be continuously monitored
  - It allows moving towards a common system

Thank You !

Backup Slides

Links

- Monitoring WG Twiki (new location!)
  - https://twiki.cern.ch/twiki/bin/view/MonitoringWG/
- Monitoring WG Report (ongoing)
  - https://twiki.cern.ch/twiki/bin/view/MonitoringWG/MonitoringReport

- Agile Infrastructure TWiki
  - https://twiki.cern.ch/twiki/bin/view/AgileInfrastructure/
- Agile Infrastructure JIRA
  - https://agileinf.cern.ch/jira/browse/AI