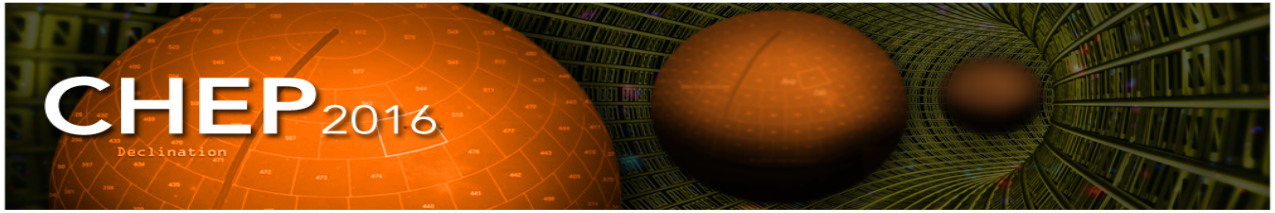

22nd International Conference on Computing in High Energy and Nuclear Physics, Hosted by SLAC and LBNL, Fall 2016

Monitoring performance of a highly distributed and complex computing infrastructure in LHCb

Christophe Haen, Federico Stagni and Zoltan Mathe

On behalf of the LHCb collaboration

DIRAC Monitoring system

- Why?
  - Need of a robust real time monitoring system
    - WMS History accounting
      - Not designed for real time monitoring (more for accounting)
      - Can not manage semi structured data
      - Not for real time analysis
      - Does not scale to hundred millions rows (more than 500 million). It requires ~1200 second to generate one month duration plot

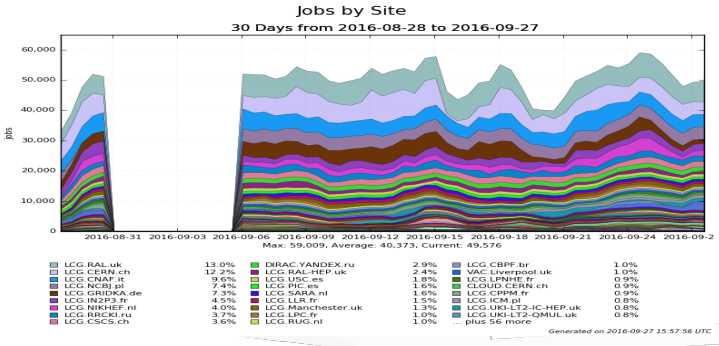

- Activity Monitor
  - Not easy to extract information
  - Not user friendly
  - Uses very old technology

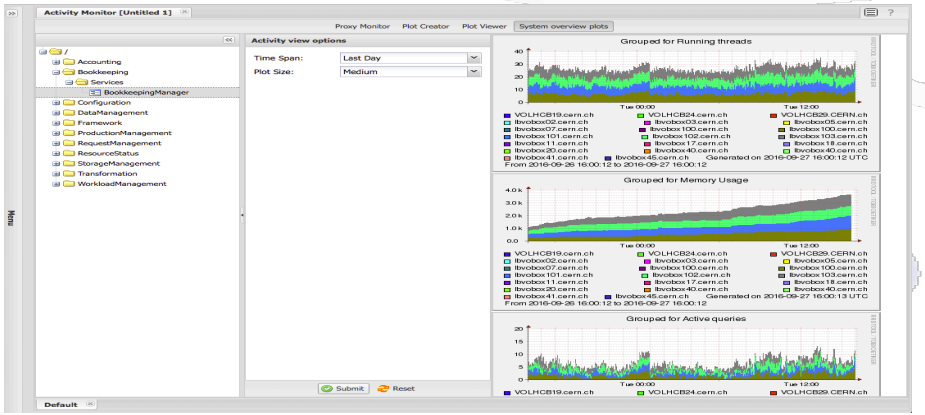

- Based on DIRAC framework and designed for:
  - Real time monitoring (WMS jobs, DIRAC components, etc.)
  - Managing semi-structured data (in our case JSON)
  - Efficient data storage, data analysis and retrieval
  - Provide good quality reports

- Use the following technologies:
  - Elasticsearch distributed search and analytic engine
  - DIRAC Graph library based on Matplotlib
  - DIRAC web framework
  - Messaging queue system as failover (stomp)

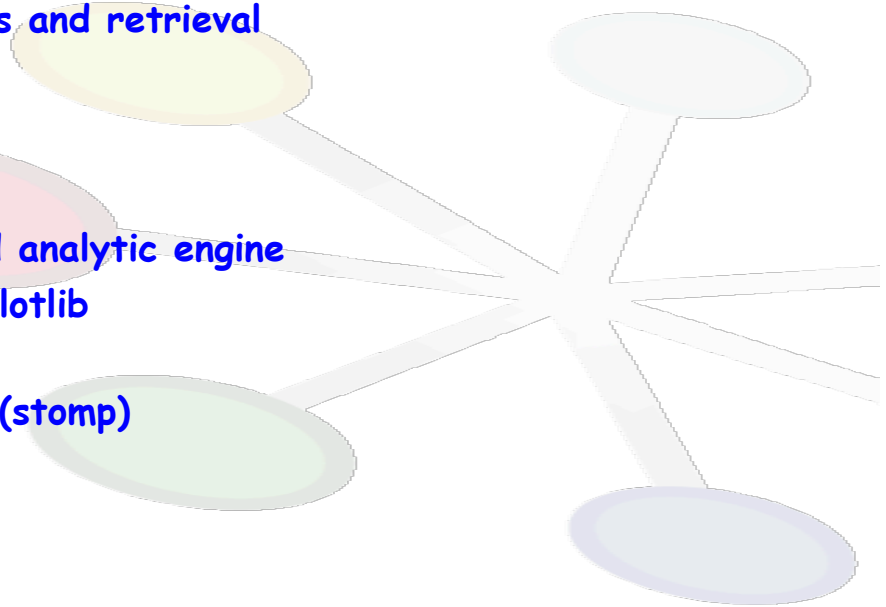

Overview of the System

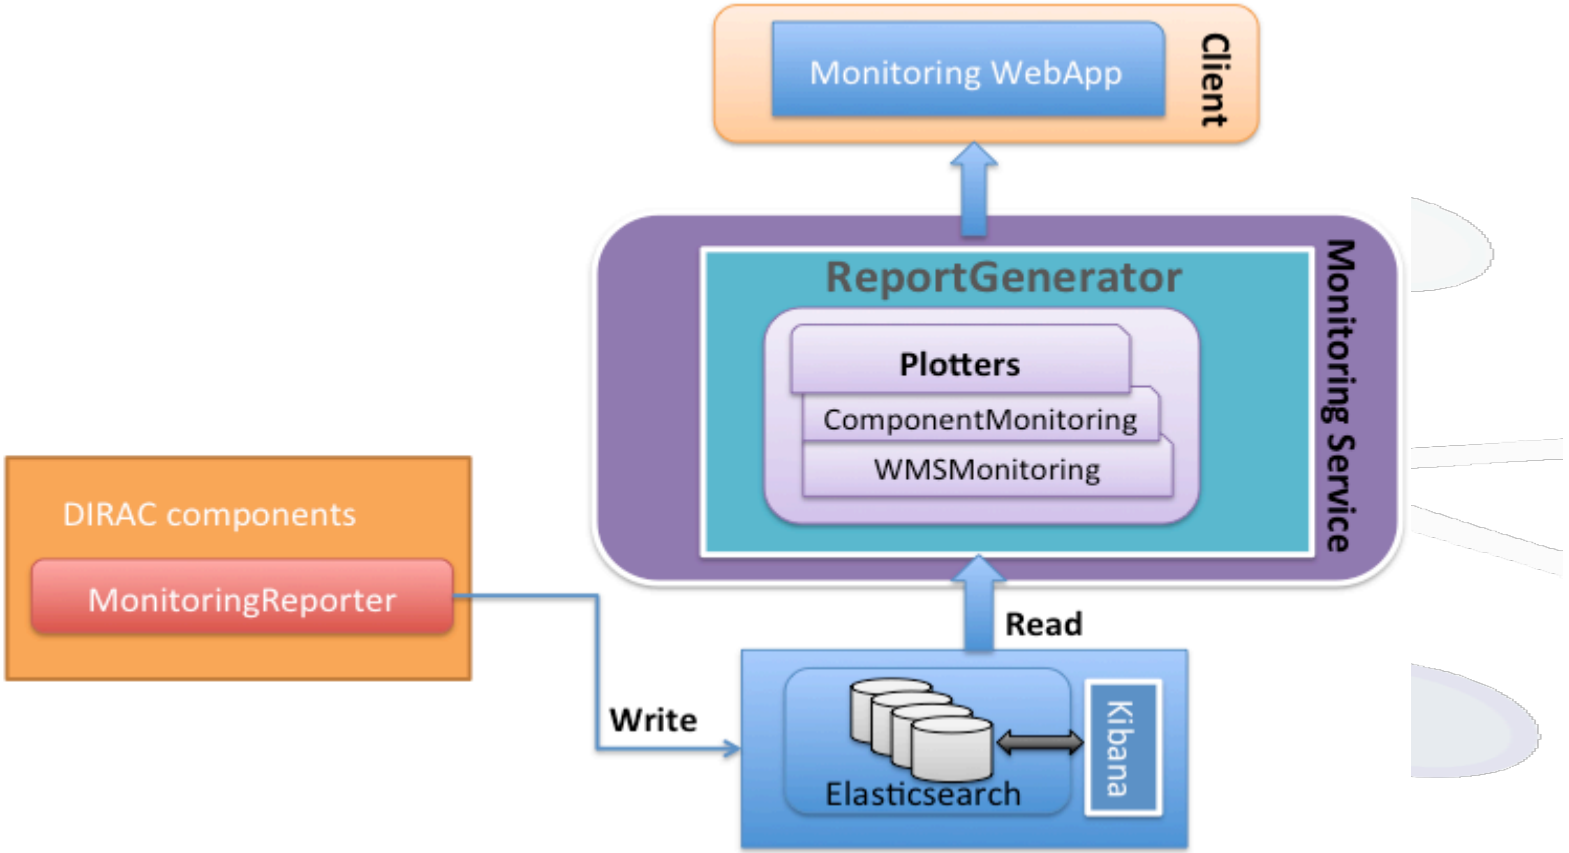

Data storage

- Data format is key/value pairs defined by the Monitoring type:
  - Monitoring type is a configurable which contains:
    - Data Retention
    - Key fields (selectable conditions)
    - Monitoring fields (values which will be plotted)
    - Timestamp
  - Example: {timestamp:t,key_1:$value_1$,key_2:$value_2$, ... ,key_n:$value_n$,$mkey_1$:$mvalue_1$,... $mkey_n$:$mvalue_n$}
- Two Monitoring types:
  - WMS history
  - Component Monitoring
- Records are stored in daily indexes per Monitoring type
- Infrastructure: 3VMs (4 processors, 8GB memory) provided by CERN OpenStack, plan to move to the CERN provided Elasticsearch cluster
- DB size: 400 Million records, 35 GB required, equivalent to 6 month monitoring information
- Performance: ~20 second to create 1 month plot. See CHEP2015 paper.

- Elasticsearch-dsl high level python library for creating and running queries
- Metric, bucket and pipeline aggregations
  - Pipeline aggregation for *dynamic bucketing*

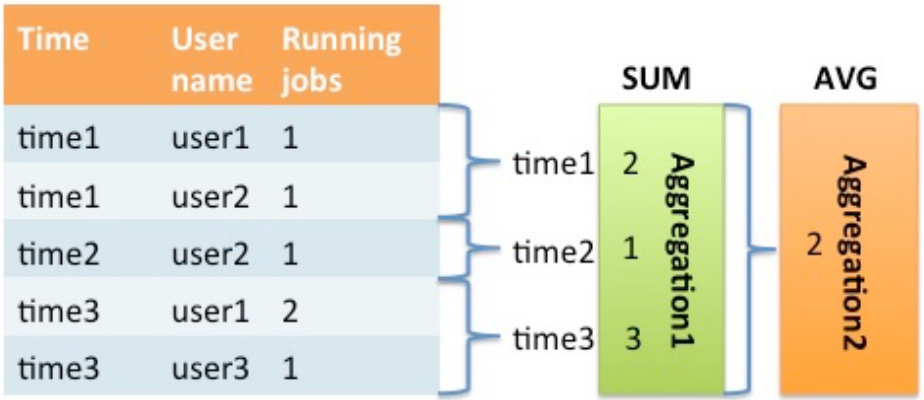

Plot creation within DIRAC framework

- Dedicated Plotter for each Monitoring type
- ReportGenerator based on DIRAC Graph library used to create the plots using the appropriate Plotter
- Plots are created on the service side using two level caching mechanisms:
  - DataCache: data used to create the plots kept in memory
  - FileSystem: plots stored in the file system
- The DIRAC web framework provided all functionalities used to visualize the plots (see CHEP2015 paper):
  - Does not require to learn external tools
  - Very simple selectors
  - Use of existing tools
  - Customization
  - Good quality plots
  - Authentication and authorization
  - Plot sharing/export mechanism

Data visualization within DIRAC

DIRAC Monitoring system

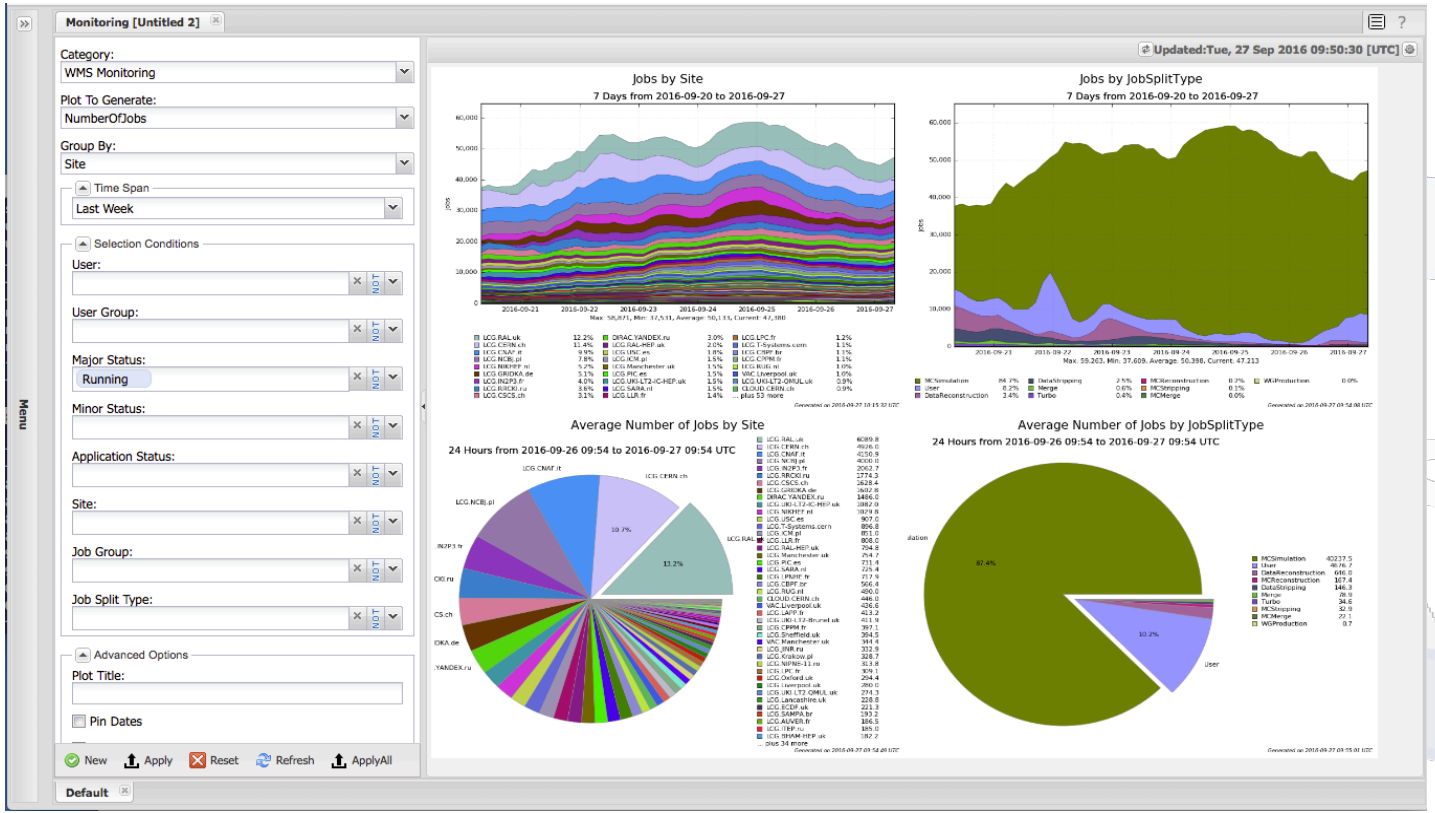

DIRAC Monitoring system

- Kibana also accessible for experts

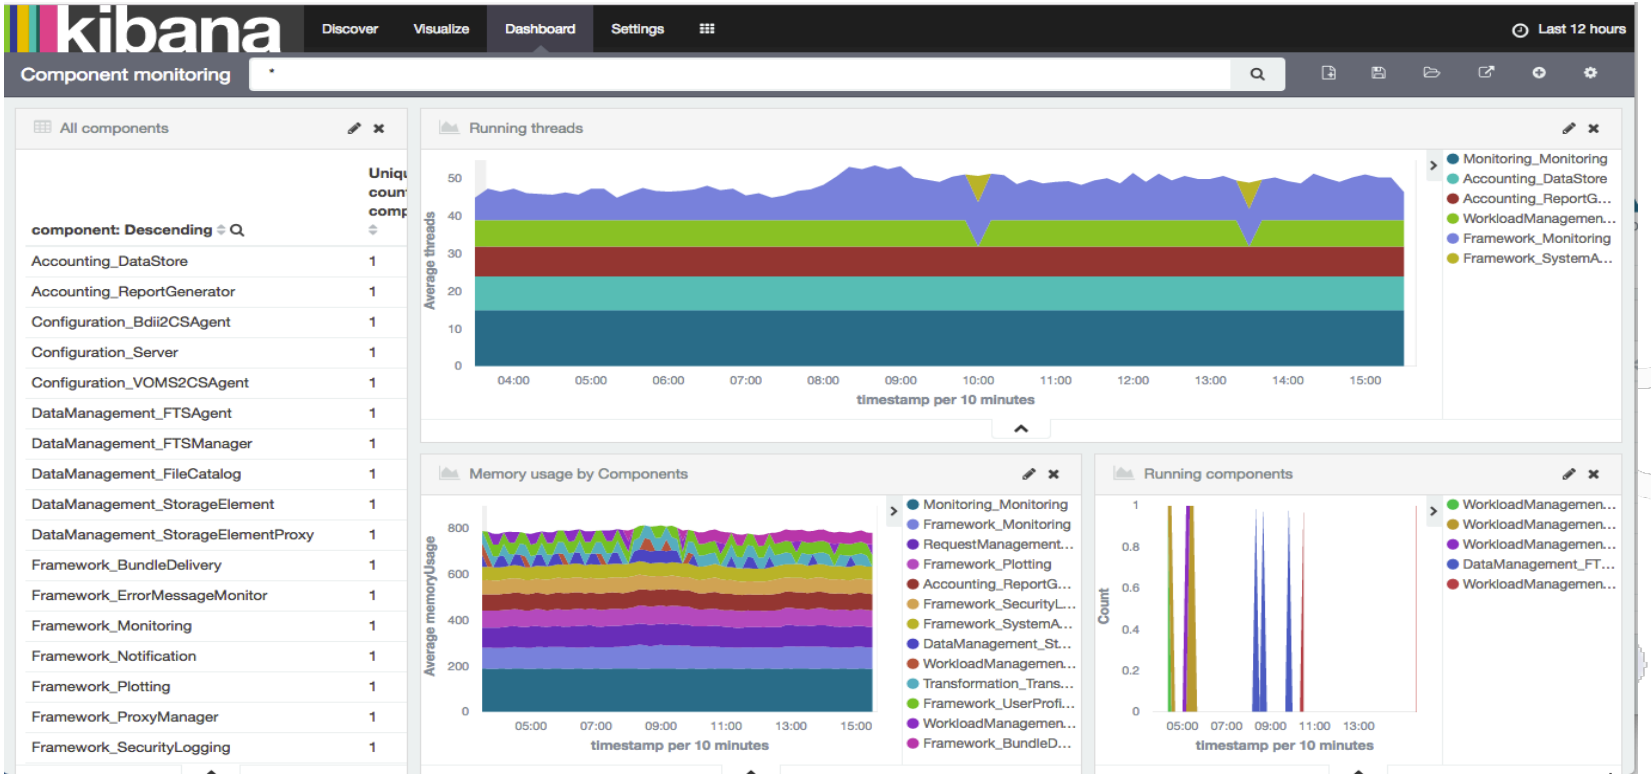

Conclusions

- **Real time:**
  - DIRAC Monitoring system can be used for real time monitoring
  - Performance issues solved
  - Easy to extend or add new Monitoring type to the system
  - Elasticsearch allows bucketing on the fly
- **Visualization:**
  - DIRAC Monitoring web application still the main visualization tool
  - Kibana available for the experts

Thank you!

Questions, comments

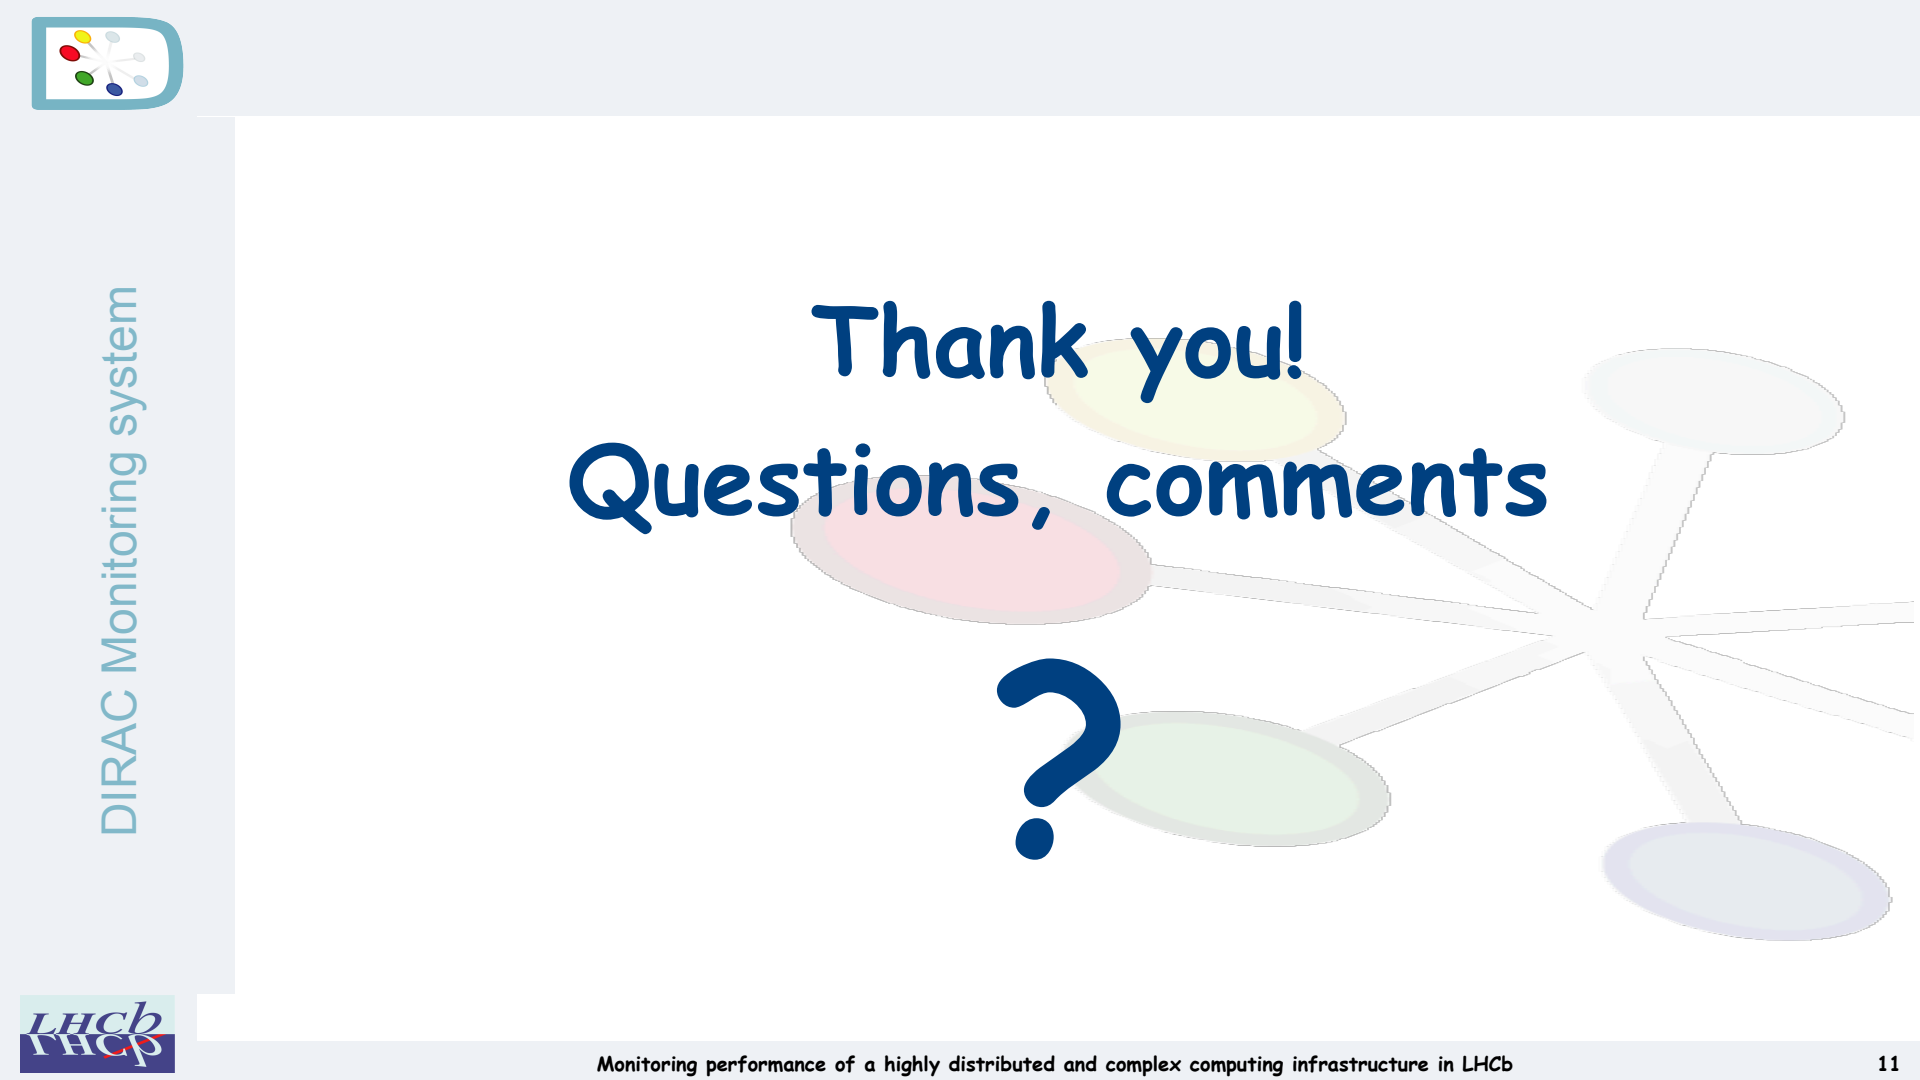

?

DIRAC Monitoring system