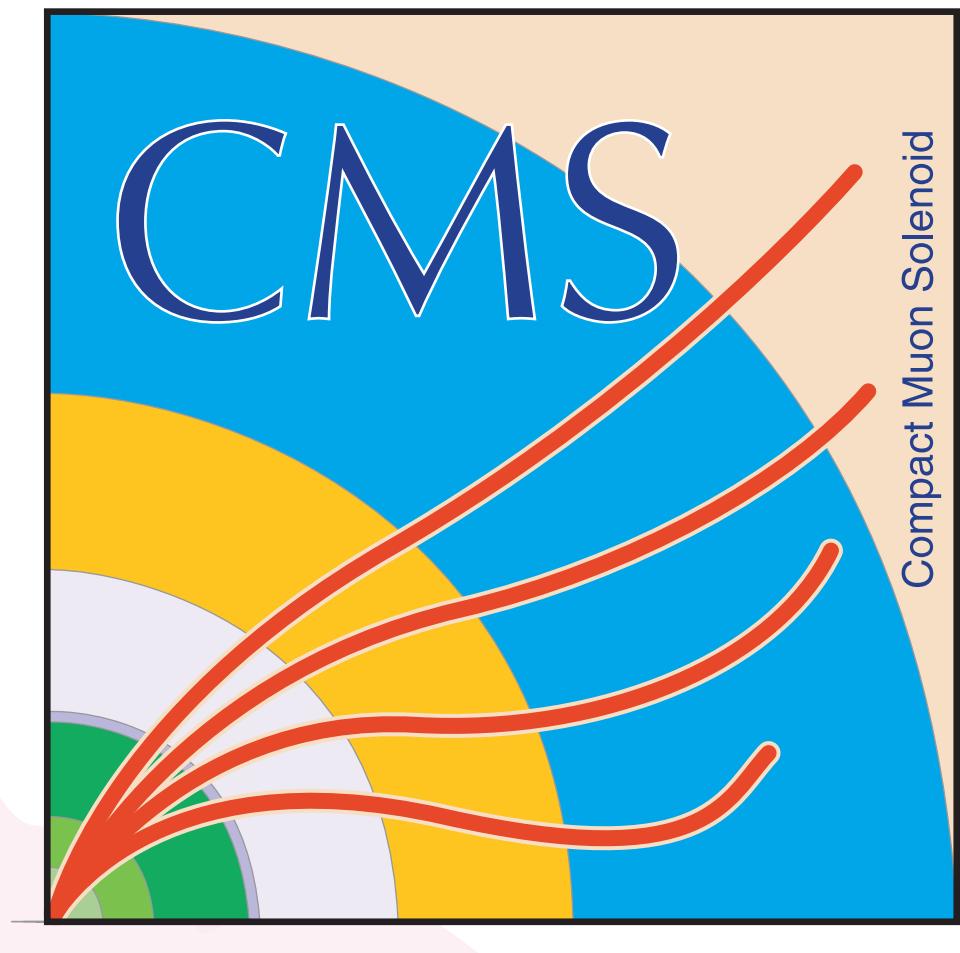


# CMS data quality monitoring web service

Lassi Tuura, Northeastern University  
 Giulio Eulisse, Northeastern University  
 Andreas Meyer, DESY and CERN  
 On behalf of the CMS collaboration



CHEP'09 – Prague, 21-27 March 2009

## Overview

### Web GUI for end-to-end DQM system

CMS developed the DQM GUI, a web-based user interface for visualising data quality monitoring data for two reasons. For one, it became evident we would much prefer a web application over a local one (Figure 1) [1, 2, 3]. Secondly, we wanted a single customisable application capable of delivering visualisation for all the DQM needs in all of CMS, for all subsystems, for live data taking as much as archives and offline workflows [4].

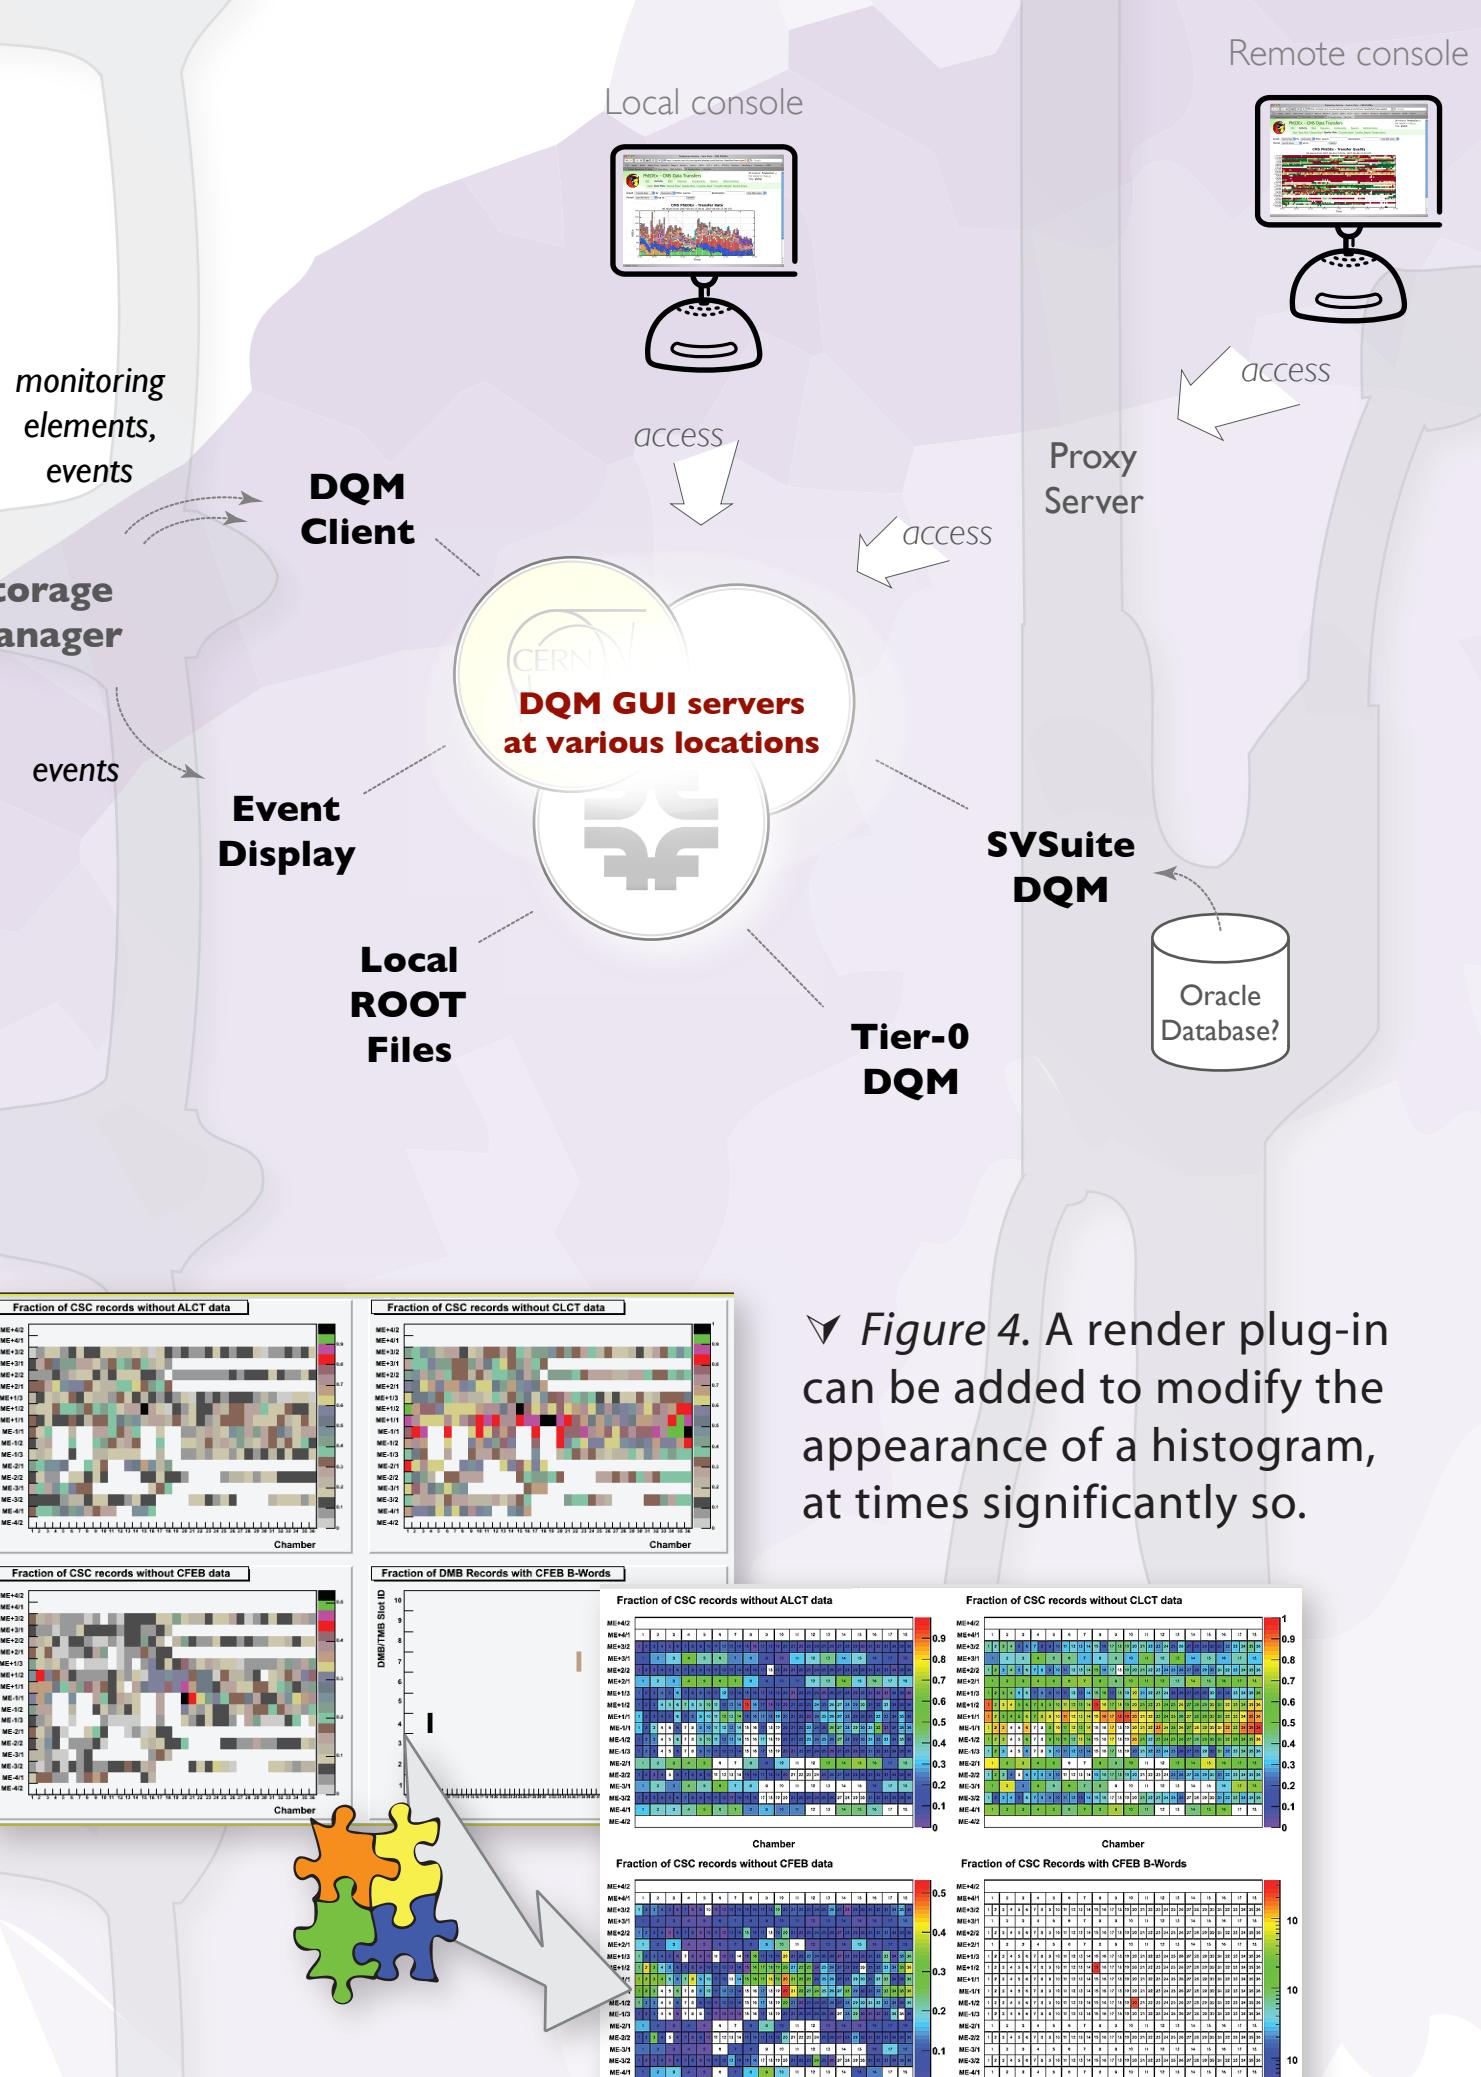
### Workspaces

Content is exposed as *workspaces* (Figure 2) from high-level summaries to shift views to expert areas, including even a basic histogram style editor. Event display snapshots are also accessible. The server configuration specifies the available workspaces.

### Layouts

Within a workspace histograms can be organised into layouts to bundle related information together. A layout defines not only the composition, but can also provide documentation (Figure 2, left top and bottom), change visualisation settings, and for example turn the reference histogram display on (Figure 2, top right). Shift views are usually defined as collections of layouts.

Figure 1. General DQM GUI architecture.



## Implementation

### Python, CherryPy and C++

Our server is built on CherryPy, a Python language web server framework [5]. The server configuration and the HTTP API are implemented in Python. The core functionality is in a C++ accelerator extension.

### JavaScript, CSS, AJAX and JSON

The client is a GUI-in-a-browser, written entirely in JavaScript. It fetches content from the server with asynchronous calls, a technique known as AJAX [2, 6]. The server responds in JSON [2, 6]. The browser code forms the GUI by mapping the JSON structure to suitable HTML+CSS content (Figures 2 and 3).

### State management

The user session state and application logic are held entirely on the web server; the browser application is “dumb”. User’s actions such as clicking on buttons are mapped directly to HTTP API calls such as “setRun?v=123”. The server responds to API calls by sending back a new full state in JSON, but only the minimum required to display the page at hand. The browser compares the new and the old states and updates the page incrementally. This arrangement trivially allows one to reload the web page, to copy

and paste URLs, or to resume the session later.

### Low-latency, parallel data processing

The server responds to most requests directly from memory, yielding excellent response time. All tasks which can be delayed are handled in background threads, such as receiving data from the network or flushing session state to disk. The server data structures support parallel traversal, permitting several HTTP requests to be processed concurrently.

### Distributed shared memory

The histograms are rendered in separate fortified processes to isolate the web server from ROOT’s instability. The server communicates with the renderer via low-latency distributed shared memory.

Live DQM data also resides in distributed shared memory (Figure 5). Each producer hosts its own histograms and notifies the server about updates. The renderer retrieves histograms asynchronously from the producers on demand; although single-threaded for ROOT, it can have dozens of operations in progress concurrently. Recently accessed views are re-rendered automatically on histogram update to reduce the image delivery latency.

```
((kind: 'AutoUpdate', interval: 300, stamp: 1237219783, serverTime: 96.78,
  kind: 'DQMHeaderRow', run: "77025", lumi: "47", event: "6'028'980",
  service: 'Online', workspace: 'Summary', page: 1, pages: 1, services: [...],
  workspaces: [{title: 'Summary', label: 'summary', category: 'Summaries', rank: 0, ...},
  runs: ["Live", "77057", ...], runid: 77025),
  kind: 'DQMQuality', items: [
    {label: 'CSC', version: 12362782330000000000,
      name: "CSC/EventInfo/reportSummaryMap",
      location: "archive/77025/Global/Online/ALL/Global/run",
      reportSummary: "0.998", eventTimeStamp: "1236244352", ...}])
```

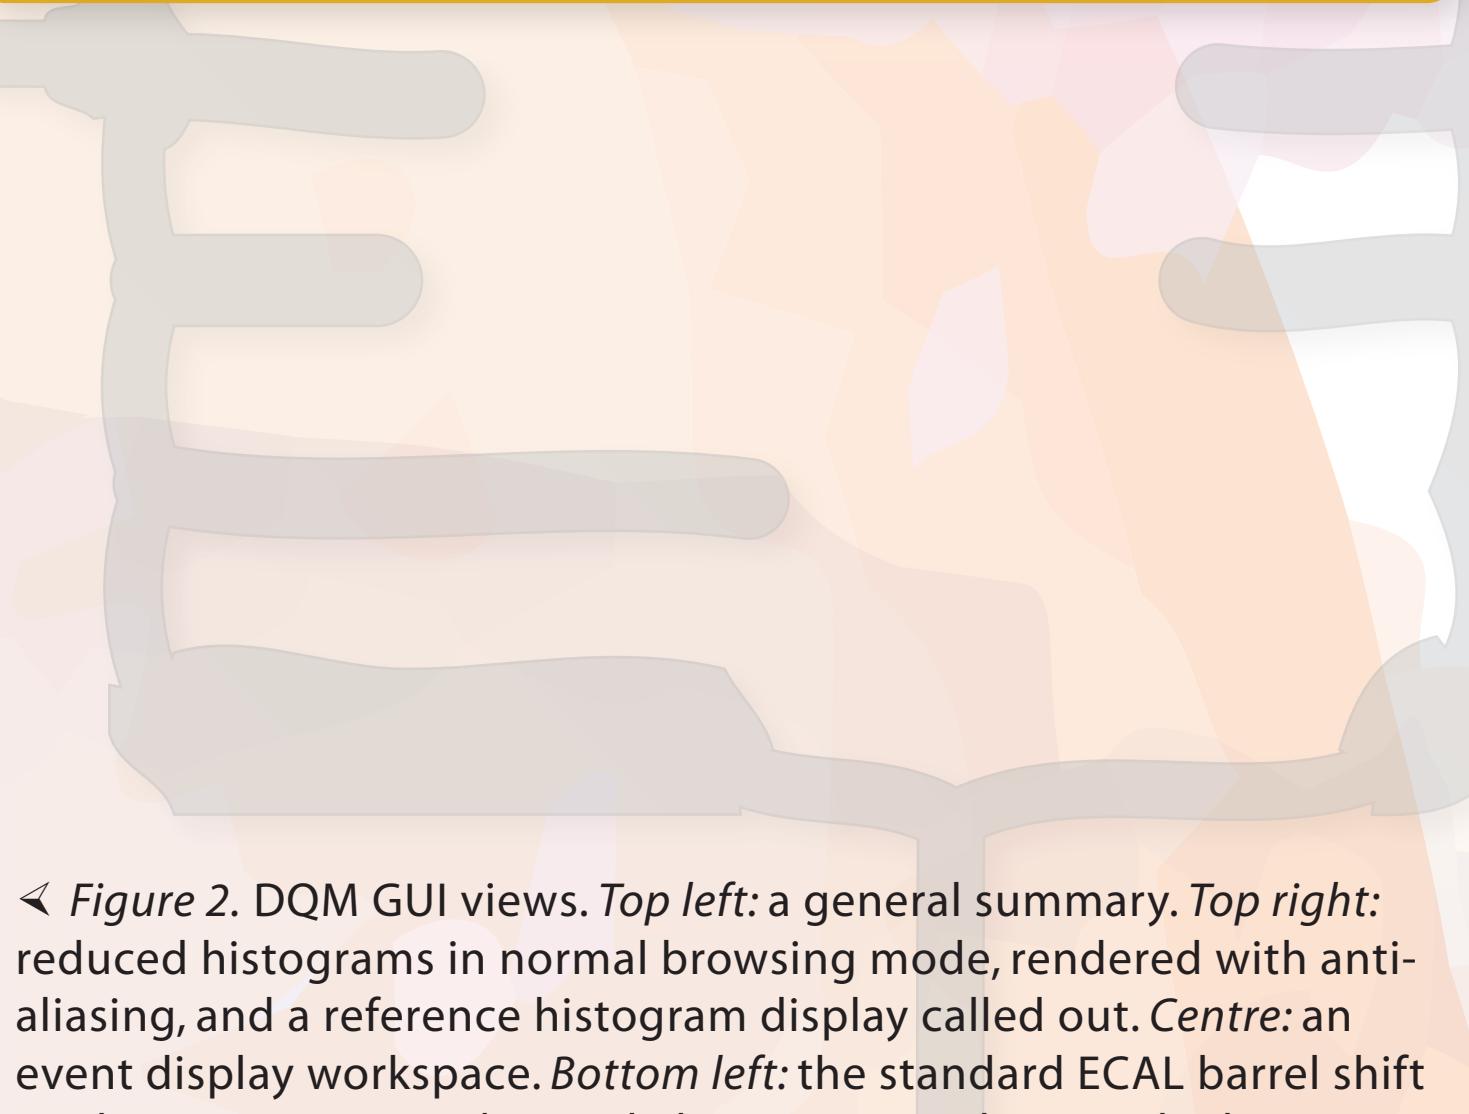


Figure 2. DQM GUI views. Top left: a general summary. Top right: reduced histograms in normal browsing mode, rendered with anti-aliasing, and a reference histogram display called out. Centre: an event display workspace. Bottom left: the standard ECAL barrel shift workspace. Bottom right: single histogram explorer and editor.

### Indexed ROOT files

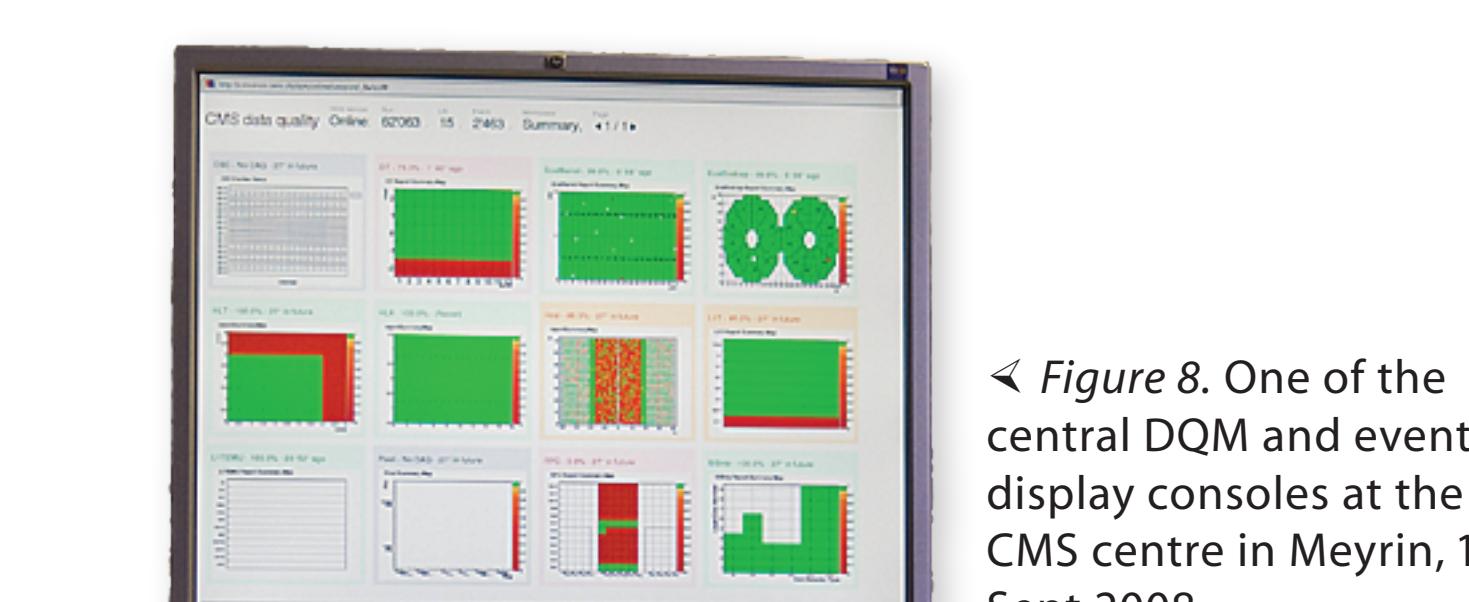
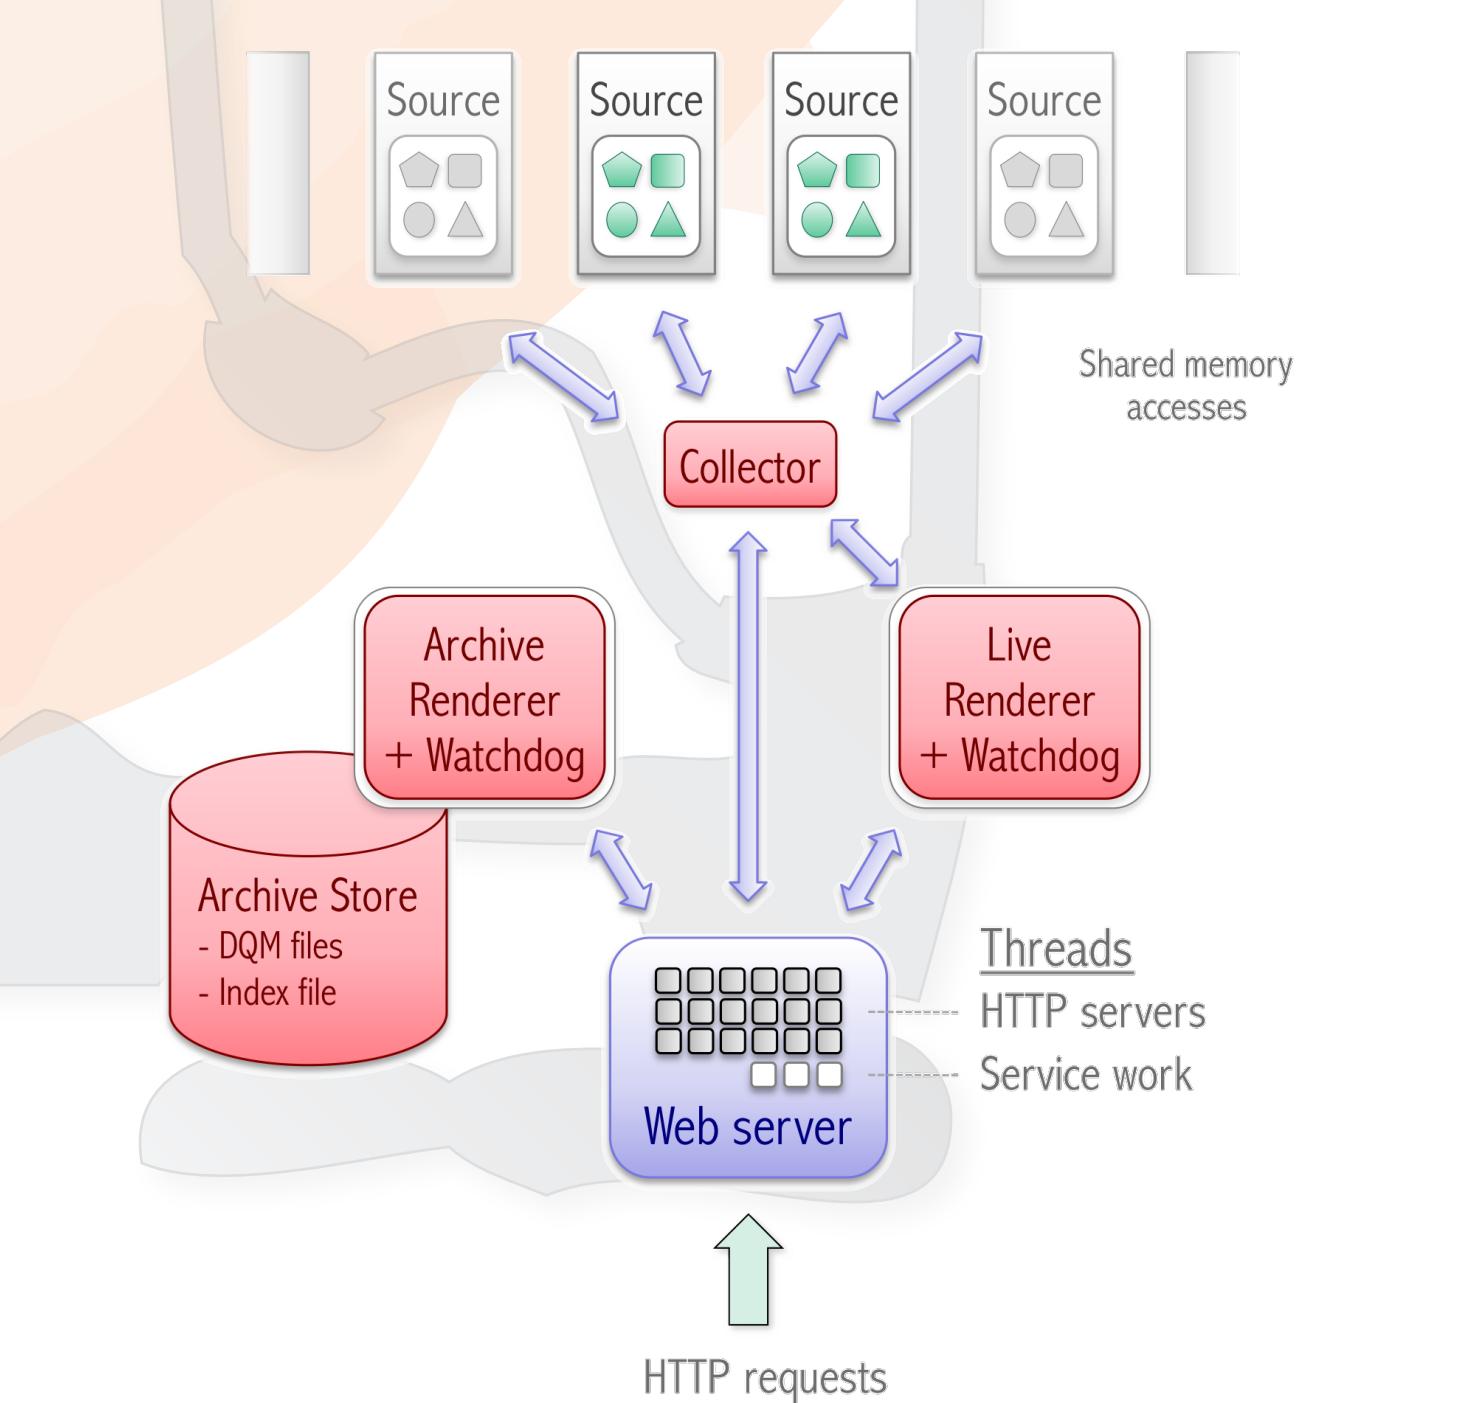
Our DQM data is archived in ROOT files [7]. As reading ROOT files in the web server itself would be too slow, use too much memory, prone to crash the server, and would seriously limit concurrency, we index the ROOT data files on upload. The GUI server computes its response using only the index. The ROOT files are accessed only to get the histograms for rendering in a separate process (Figure 5). The index is currently a simple SQLite database [8].

### High-quality extensible rendering

The server supports setting basic render options, such as linear vs. log axis, axis bounds and ROOT draw options (Figure 2, bottom right). These settings can be set interactively or as defaults in the subsystem layout definitions. The subsystems can further customise the default look and feel of the histograms by registering C++ render plug-ins, which are loaded on server start-up (Figure 4).

We improve image quality significantly by generating images in ROOT in much larger size than requested, then reducing the image to the final smaller size using a high-quality anti-alias filter.

Figure 5. The distributed shared memory system.



### Reliability

Early on it became abundantly evident ROOT was neither robust nor suitable for long-running servers. Some three quarters of all the effort on the entire DQM GUI has gone into debugging ROOT instabilities and producing countermeasures. We are very pleased with the robustness of the rest of the DQM GUI system.

### Capacity and performance

CMS typically creates circa 50'000 histograms per run. The average GUI HTTP response time is around 270 ms (Figure 6), which we find satisfactory. The production server has scaled to about 750'000 HTTP requests per day with little apparent impact on the server response time (Figure 7).

Interestingly the vast majority of the accesses are to the online production server from outside the online environment. This indicates the web-based monitoring and visualisation solution applies well to the practical needs of the experiment.

We have exercised the GUI with up to 300'000

## References

- [1] Data Quality Monitoring for the CMS Silicon Strip Tracker, D. Giordano et al, CHEP'07.
- [2] Interactive Web-based Analysis Clients using AJAX: with examples for CMS, ROOT and GEANT4, G. Eulisse et al, CHEP'06.
- [3] CMS Offline Web Tools, S. Metson et al, CHEP'07.
- [4] CMS data quality monitoring: systems and experiences, L. Tuura et al, CHEP'09.
- [5] <http://cherrypy.org>
- [6] <http://json.org>
- [7] <http://root.cern.ch>
- [8] <http://sqlite.org>

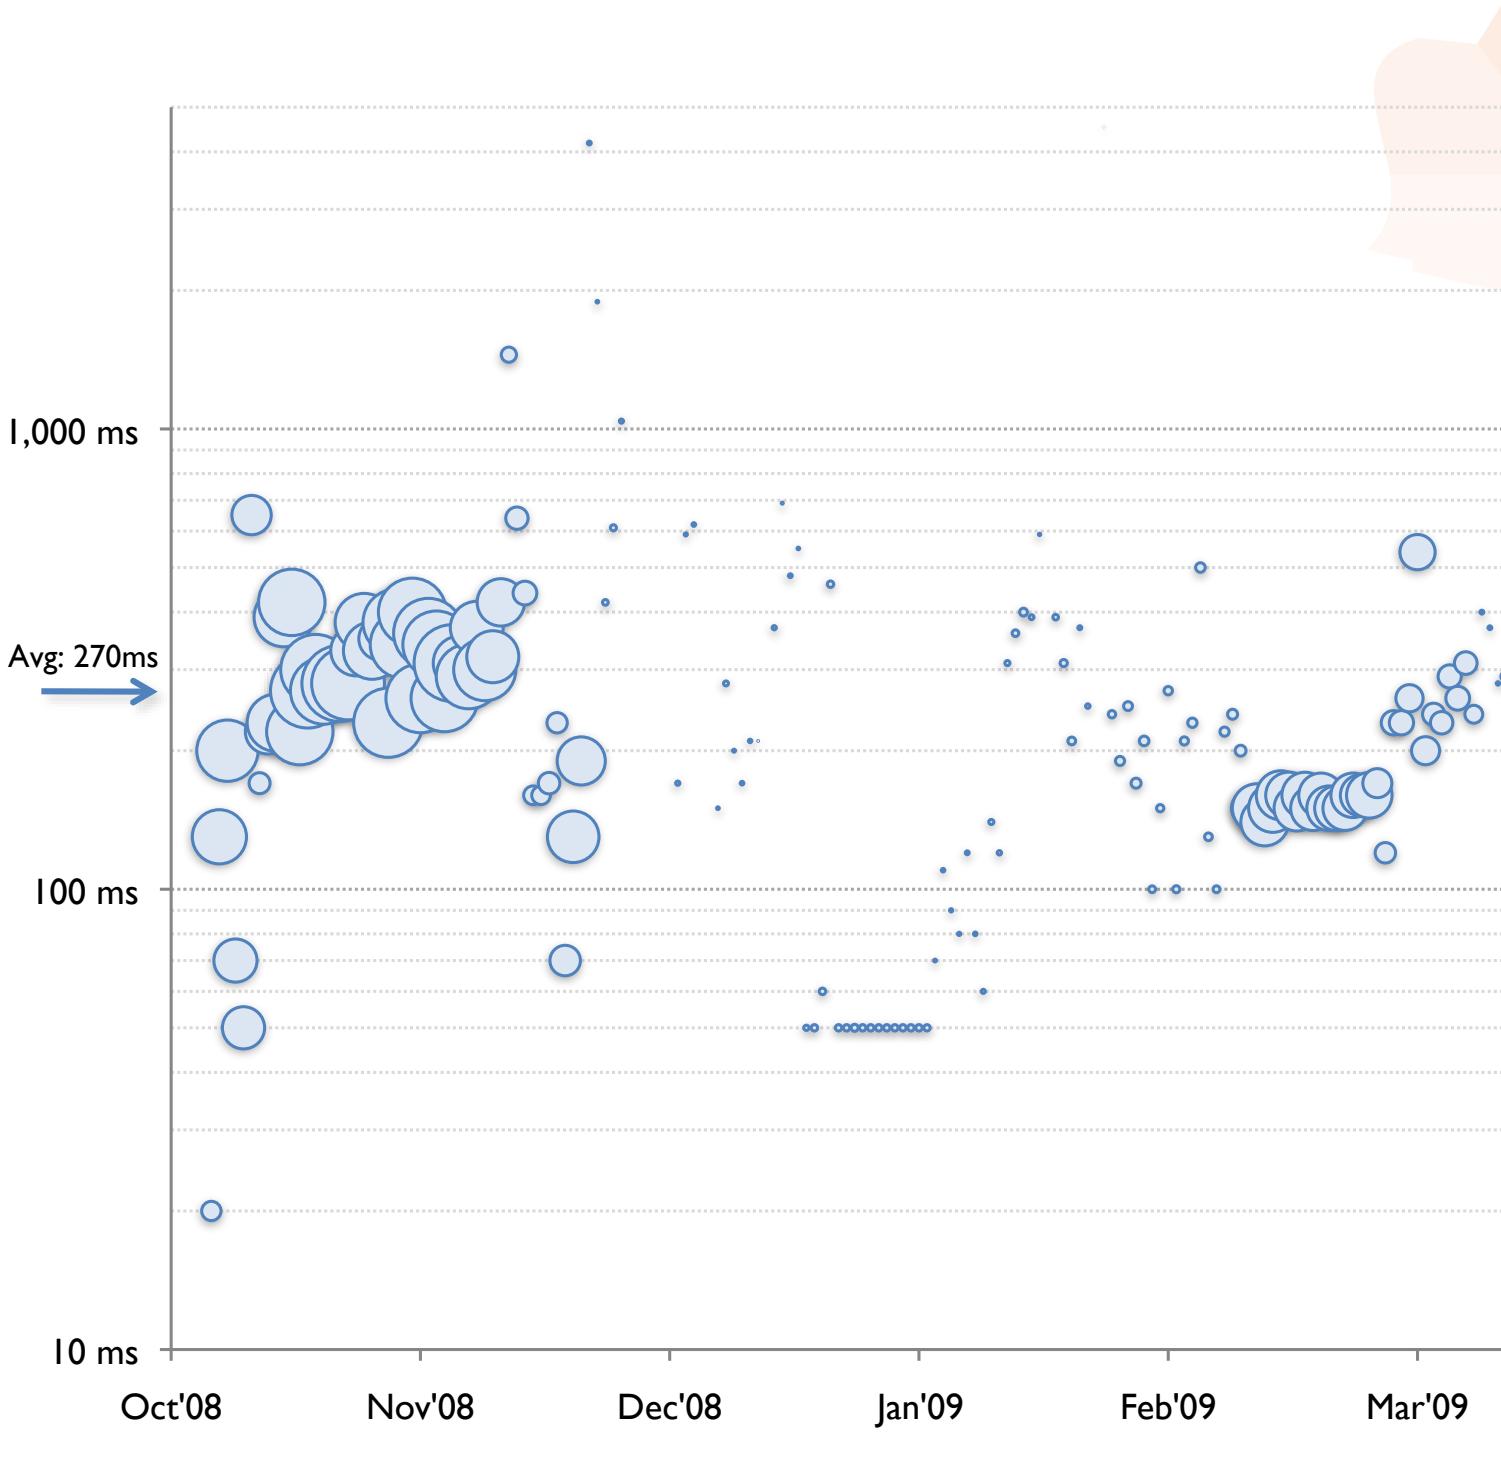


Figure 6. The daily average DQM GUI response time versus date and number of requests per day.

histograms per run. The GUI remains usable although there is a perceivable interaction delay. We plan to optimise the server further such that it has ample capacity to gracefully handle growing histogram archives and special studies with large numbers of monitored entities.

## Operation and experience

### Deployments

CMS centrally operates four DQM GUI instances for online and offline each, an instance per purpose for the existence of data: Tier-0, CAF, release validation, and so on. In addition at least four instances are operated by detector subsystems in online for exercises private to the detector. Most DQM developers also run a private GUI instance while testing. A picture of a live station is shown in Figure 8.

CMS typically creates circa 50'000 histograms per run. The average GUI HTTP response time is around 270 ms (Figure 6), which we find satisfactory. The production server has scaled to about 750'000 HTTP requests per day with little apparent impact on the server response time (Figure 7).

Interestingly the vast majority of the accesses are to the online production server from outside the online environment. This indicates the web-based monitoring and visualisation solution applies well to the practical needs of the experiment.

We have exercised the GUI with up to 300'000