



## BBC News Website uses MySQL to Monitor Reader Interest

Millions of people every day visit the BBC News website for their daily fix of world news and analysis. Stories from hundreds of dedicated journalists in all time zones are constantly being fed to the site, which currently attracts 35 million unique users and receives over 800 million page impressions each month.

The BBC News website was started in 1997, but a system for reporting in real time which stories were proving most popular had yet to be invented.

The BBC wanted a dynamic system, working in real time, which would give audiences a real sense of the news stories on the website being picked up most by other users.

The challenge was to balance a demand for high performance with value for money on a limited budget. The BBC's solution was to develop the 'BBC News Live Stats' system in-house, supported by an open source software stack consisting of Apache web servers, Linux as the operating system and the MySQL database for data management. This strategy reduced the cost outlay to the hardware alone.

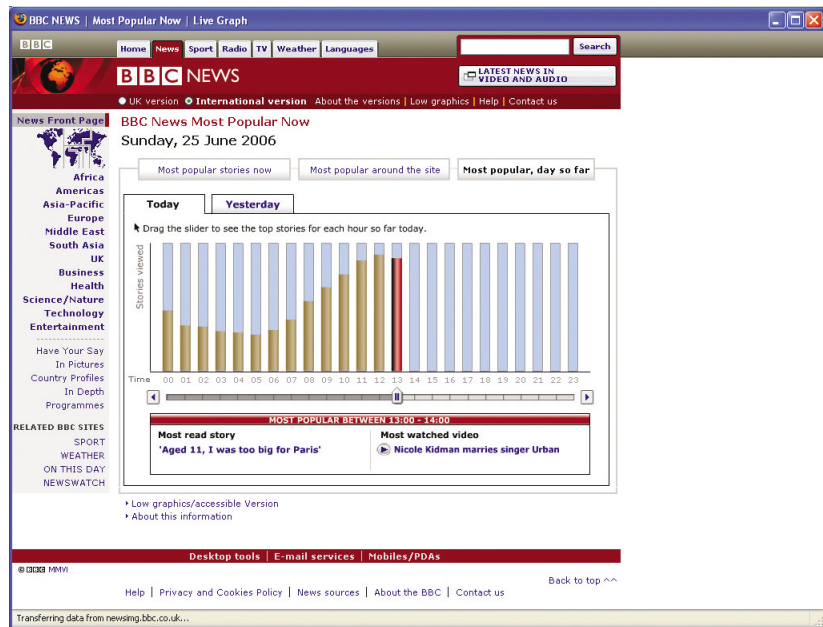
The project involved setting up an entirely separate content network that houses four Apache web servers behind load balancers that serve nothing but real time data collection logic. Data is passed to Java applications on Linux application boxes via a messaging queue and sampled from one in seven users by the Java applications.



## Open Source Stack Serves Statistics on News

The applications also carry out geographical IP-look-ups with two dedicated geo-IP servers before passing the processed data to a MySQL load data infile – the MySQL database is then queried by hundreds of processes all over BBC News via the system’s live stats web service. Typically there’s a ten second delay from a user hitting one of the BBC News pages to it contributing to the live charts online. Due to privacy reasons, no personally identifiable information is stored.

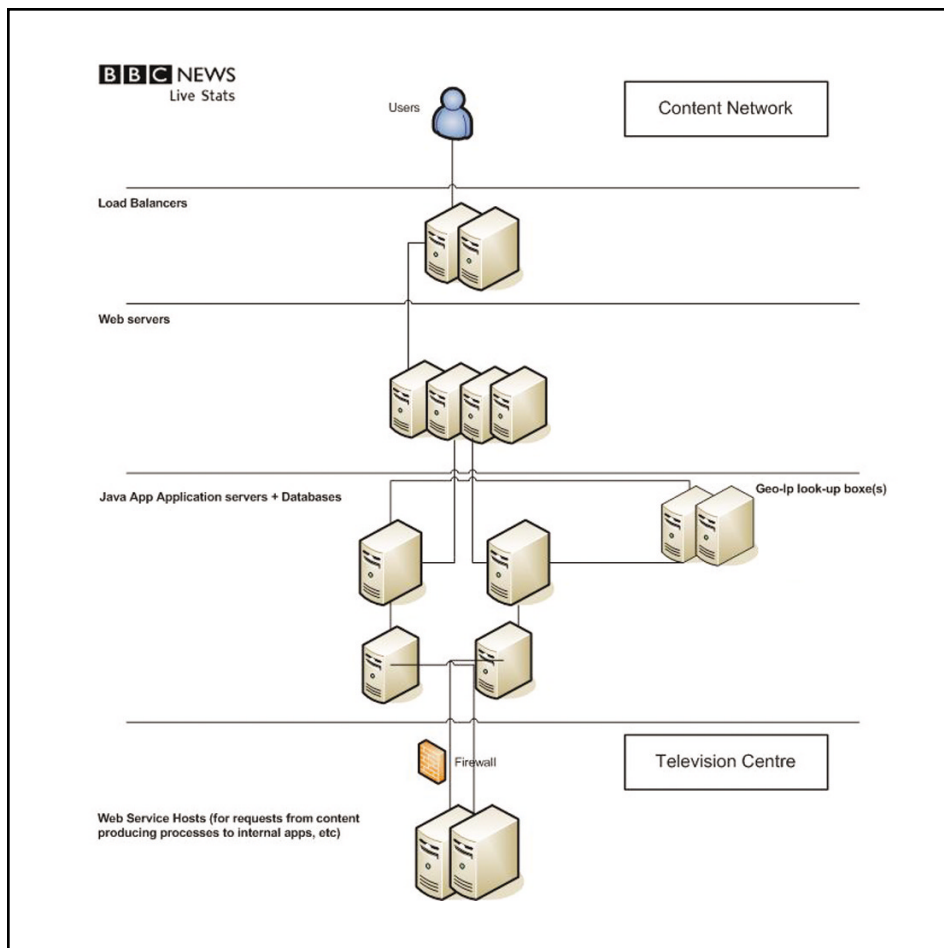
The BBC News stats application has been live since spring 2006. It can be reached from any BBC News web page by clicking the “Most popular stories now” links. In the default view, the current top ten stories worldwide are listed. By clicking a region in the map, the application visually reshuffles the list to reflect interest in that particular part of the world. In another view it is possible to monitor the number of stories read hour by hour during the day as well as which story and video have been the top during each hour of the day so far.



*The BBC’s solution was to develop the ‘BBC News Live Stats’ system in-house, supported by an open source software stack consisting of Apache web servers, Linux as the operating system and the MySQL database for data management. This strategy reduced the cost outlay to the hardware alone.*

## Technical Setup

Input from users browsing the Content network is directed via two load balancers to the four Apache web servers. Two database boxes with MySQL 5.08 on Linux OS pick up the data from the inhouse Java application boxes after sampling, and add geographical data retrieved from two Geo-IP Look-up application servers. Two application servers running in-house applications on Windows feed internal requests.



## Technical Environment

Number of tables in MySQL database:  
24, excluding merge tables

Number of records in Largest Table:  
up to 4 million

Size of database: around 8 million rows

Number of transactions:  
around 30,000 inserts/minute, 4000 selects/hour.

MySQL server version 5.08 is currently used. The MyISAM merge tables function proved to speed up performance significantly. Development work and performance have been good. MySQL Replication is also used to replicate data from live databases to query fielding database boxes.

## About MySQL

MySQL AB develops and supports a family of high performance, affordable database products – including MySQL Network, a comprehensive set of certified software and premium support services. The company's flagship product is the MySQL Server, the world's most popular open source database, with more than 10 million active installations. Many of the world's largest organizations, including Yahoo!, Alcatel, The Associated Press, Suzuki and NASA are realizing significant cost savings by using MySQL to power high-volume Web sites, business-critical enterprise applications and packaged software.

With headquarters in Sweden and the United States — and operations around the world — MySQL AB supports both open source values and corporate customers' needs in a profitable, sustainable business.

For more information about MySQL, please go to **[www.mysql.com](http://www.mysql.com)**.



The World's Most Popular Open Source Database