

: How does the Facebook database operate?

by Daniel J. Power

Editor, DSSResources.COM

In March 2013, Facebook, the social networking megaplatform, reported it had 1.11 billion active users. 665 million people were using the application each day. In late 2012, a post at Stackoverflow.com, estimated Facebook was storing approximately 400 Billion posts and photos on more than 30,000 servers. How can Facebook rapidly query, categorize and display relevant posts for each user when they signon the site?

Imagine the situation posed at Stackoverflow.com:

What if a simple select query like:

```
SELECT post_message FROM Post_table WHERE user_id = 1500;
```

takes .02 of a second and this query is executed for 400 million users who see at least 10 posts per day every day and that means 46,720 queries per second. How is that possible?

First, Facebook distributes the queries across more than 30,000 servers. So the load for 400 million users per day is about 13,333 user queries per server. Anton at Stackoverflow.com responded that Facebook is not necessarily loading the posts you see from the database – they might end up fetching data from memcached without touching the DB at all for 90% of users. Second, "please do not underestimate the power of splitting the data across different servers. If you have a table with 1B records on one server and have to select 10 rows from it – it is not nearly the same as if you have 100 servers hosting 10M records each." You might have to ask 3 servers instead of one, but they would reply 100 times faster and you can parallel the query if you like. The third thing is you should never rely on assumptions. Two tenths of a second is a good query execution time, but it is not the best possible. "In reality you can get data in as little as a tenth of a millisecond, but again it depends. If we sum everything up, it is proper software (with aggressive caching) with proper hardware (tons of servers)."

Facebook is a customized LAMP site that uses MySQL primarily as a key-value persistent storage moving joins and logic onto its web servers. It used Memcached, a distributed memory caching system. According to the Royal Pingdom blog, "Facebook runs thousands of Memcached servers

: *How does the Facebook database operate?*

with tens of terabytes of cached data at any one point in time. It is likely the world's largest Memcached installation." Also Facebook uses Cassandra (<http://cassandra.apache.org>), a distributed storage system with no single point of failure. Cassandra's data model provides "the convenience of column indexes with the performance of log-structured updates, strong support for denormalization and materialized views, and powerful built-in caching."

References

Associated Press, "Number of active users at Facebook over the years," at <http://news.yahoo.com/number-active-users-facebook-over-230449748.html>

<http://stackoverflow.com/questions/13313282/why-is-the-facebook-database-so-fast>

<https://www.facebook.com/MySQLatFacebook>

Royal Pingdom blog, "Exploring the software behind Facebook, the world's largest site," June 18, 2010 at URL <http://royal.pingdom.com/2010/06/18/the-software-behind-facebook/>

Author: Daniel Power

Last update: 2013-08-24 11:30