
EXHIBIT 56 G



Facebook Photos Infrastructure

Share +

by Doug Beaver 1:39pm Monday, May 21

Facebook users love our photos application, it's one of the most popular areas on the site. The cool front-end features—like upload, tagging, and editing—are great, but I'd like to take a moment and talk about our photo infrastructure. It's the part of our photos feature that users never see, but it quietly works 24 hours a day to quickly serve all those photos to millions of users.

Allow me to geek out and share some photo infrastructure statistics with you. We have:

- 1.7 billion user photos
- 2.2 billion friends tagged in user photos
- 160 terabytes of photo storage used with an extra 60 terabytes available
- 60+ million photos added each week which take up 5 terabytes of disk space
- 3+ billion photo images served to users every day
- 100,000+ images served per second during our peak traffic windows

(For you sharp-eyed geeks out there, we actually store four image sizes for each photo, so if you want to know how many files we have, multiply the photo numbers by four. That's right, we're adding a quarter billion files every week.)

And that's just the storage side of things. We also have a fleet of servers who receive your photo uploads, happily scrubbing out EXIF data and tweaking the quality settings so your photos look their best. Another fleet of servers and caches who feed images to our CDN (content delivery network) partners. And hey, if we feel like it, we can serve some of the photo traffic directly ourselves too.

We have also developed our own specialized web servers that are tuned to serve files with as few disk reads as possible. Even with thousands of hard drive spindles, I/O (input/output) is still a concern because our traffic is so high. Our squid caches help reduce the load, but squid isn't nearly fast or efficient enough for our purposes, so we're writing our own web accelerator too.

That's just a taste of the current photo infrastructure, and while the system we've got is pretty sweet, there are even bigger things to come. Browser-side sampling of image response times so we can immediately know when ISPs are having issues, automated balancing of load between CDNs and the Facebook network (why pay a CDN to serve images in the middle of the night when we've got a huge photo serving tier that's quiet at that time?), better image performance for our international users, the list goes on and on...

Doug has a bet with the storage team that Facebook will have a petabyte of photos before April 2008, so please don't let him down.

To post a comment, you must be logged in. To register for Facebook, click here.

Archived Posts by Date:

- 2007**
- July (4)
- June (2)
- May (5)
- April (9)
- March (8)
- February (7)
- January (4)
- 2006**
- December (2)