

Who uses PostgreSQL at UCLA?

The purpose of this page is to create a directory of people who are using PostgreSQL and what advanced features they are using. Please try to mention how you have found certain features useful so that you can give other people ideas on how to better utilize the database.

Jason Fong

Database's purpose: collecting large sets of data, and querying data (usually big long queries) to discover trends or other interesting results

Total database size: approximately 500 GB

Largest table: approximately 500 million rows

Advanced features in use:

- Procedural languages:
 - PL/pgSQL
 - PL/perl: This is very useful— it lets one run bits of Perl code inside of a query. This can also be used to create a Perl function index on a table so that the Perl code doesn't even need to be run for future queries.
- Triggers: I use these for a number of things:
 - Automated auditing: a table can have a trigger that inserts a row into a log table whenever any insert/update/delete occurs
 - Duplicate row detection: A table can have a trigger on inserts that checks for a duplicate primary key and then compares the other column values to the matching row. If all the columns are the same, the insert can be dropped. If any column is different, then the rows can be merged (or the new row could be inserted into a separate table to keep track of these duplicate primary keys with different column values).

Revision #4

Created Tue, May 30, 2006 6:19 PM by Fong, Jason

Updated Wed, May 31, 2006 12:10 AM by Fong, Jason