

Project A: Extending Microblog

In this first project, you will spend an intensive two weeks understanding the implementation of a small web log (“blog”) application, Microblog, and extending it to add additional functionality.

The project must be done individually, although you are strongly encouraged to ask for help from the TAs, classmates, and friends. The goal here is for you to understand, top-to-bottom how a database-backed web application works. In project B, you will design and implement your own database-backed web application from scratch.

Before you start

Read the handout “Using Oracle in the TLAB”. This will explain how to log in to the tlab server, how to configure your environment, and how to access oracle using SQL*Plus. Make sure that your environment is working correctly.

Getting and installing Microblog

To install Microblog, log into your account on tlab-server and do the following:

```
cd ~/public_html
tar xvfz ~pdinda/HANDOUT/microblog.tgz
cd microblog
more README
```

The README file will give you detailed instructions on how to configure Microblog and verify that it is working. You should be able to visit your Microblog via <http://tlab-server/~you/microblog/blog.html>. At this point, you should be able to log in as the root user and do the following: add and delete users, add and revoke user permissions, write messages, see a summary of all messages, and see all the messages.

How does Microblog work?

The state of Microblog (users, their permissions, and the messages) lives entirely in the database. That state consists of a sequence and three tables (see blog.sql). The `blog_users` table contains the username, password, and email address of each Microblog user. The password is stored in cleartext, which is insecure. The `blog_actions` table contains the names of all the possible actions that a user may take in Microblog. The `blog_permissions` table maps from Microblog usernames to the actions that they have permission to take. An action should only be taken if the action exists in the `blog_actions` table, the user exists in the `blog_users` table, and the user has permission for the action in the `blog_permissions` table. By default, the following actions are made available: `manage-users`, `query-messages`, `delete-any-messages`, `delete-own-messages`, and `write-messages`. Two users are added by default: `root`, with password `rootroot`, who can do anything, and `none`, with password `nonenone`, who can do nothing.

The `blog_message_id` sequence provides a way of getting a unique number for tagging each message (the message id). The `blog_messages` table contains the actual messages. In addition to a message id, each message has an author (who must exist in the `blog_messages` `blog_users` table), a subject, a timestamp, and the text of the message (which can be HTML). A message also has a field called “`respid`”, which is the id of the message to which it is responding. This referential behavior means that the messages form a tree. By default, there is a message with id zero, written by user none, and referring to itself that is installed into the database. This message, which is never displayed, is the root of the message tree.

Microblog has a notion of user sessions. The state of a session is kept in the web browser, in a cookie named `MicroblogSession`. When users log into Microblog they are given a cookie that contains their login name and password in cleartext (this is not secure) and is set to expire in about one hour. Without a cookie, the user can only log in. When presented with a cookie, Microblog uses its contents to validate the user and check to see if he has the necessary permissions to do what he wants.

The following is a more detailed description of a Microblog session:

- ⇒ The user visits <http://tlab-server/~you/microblog/blog.html>
- ⇒ `blog.html` is loaded
 - It creates an HTML frameset, consisting of a left frame and a right frame
 - The left frame is filled with `actions.html`, whose different links are targeted at the right frame.
 - The right frame is filled by executing `blog.pl` with the “`act`” parameter set to “`query`”
 - `blog.pl` notices that there is no cookie and so forces the “`login`” action.
 - The login form is displayed.
- ⇒ The user fills out the login form and hits submit.
- ⇒ `blog.pl` is invoked with the “`act`” parameter set to “`login`”, and the “`loginrun`” parameter set to one.
 - `blog.pl` extracts the “`user`” and “`password`” parameters provided by the form, and does a SQL query to see if the combination exists in the database.
 - If the combination exists, it creates a cookie with the combination and passes it back to the user’s browser. Next, it displays the query form (provided the user is permitted `query-messages`).
 - If the combination does not exist, `blog.pl` complains, does not return a cookie, and displays the login form again.
- ⇒ At this point, the user’s browser has the cookie, which is good for one hour.
- ⇒ The default “`query`” action, if there is a cookie, and the user/password combination in the cookie is valid, and the user has `query-messages` permission, displays the query form and then does a SQL query to generate a summary of all the messages in the database which it then displays as an HTML table.
- ⇒ The user fills out the query form and hits submit.
- ⇒ Seeing that the “`queryrun`” parameter exists, `blog.pl`, if there is a cookie, and the user/password combination in the cookie is valid, and the user has `query-`

- messages permission, executes a SQL query that fetches all the matching messages. It then prints them formatted in HTML.
- ⇒ On every interaction, the cookie is refreshed, so the one hour time limit on the cookie is basically the maximum idle time before the user is logged out.
 - ⇒ The other actions are very similar to the query action.

Project Steps

In this project, you will extend Microblog to provide the following functionality. Each extension is marked with a percentage that is intended to reflect its perceived difficulty level. You may do the extensions in any order, although we believe the order presented will be the easiest. We believe the best approach to building each extension is the following:

1. Read and understand `blog.sql` and `blog.pl`
2. Design the SQL statements that are needed
3. Test the SQL statements using SQL*Plus.
4. Embed the SQL statements into Perl functions (see `UserAdd()`, for example)
5. Write the Perl logic to call the functions at the appropriate times.

You should avoid becoming bogged down in the Perl and, especially, the HTML, parts of the code. As this is a databases course, you will be graded on correctness, not the appearance of the final product.

Implement query (20%) : As handed out, message queries completely ignore all the arguments and simply return all messages. Extend `MessageQuery` so that it is possible to query messages by author, by date range, and by both. You will find that it is easiest to add query-by-author first. Query by date will require that you become familiar with `Time::ParseDate`.

Add reply and delete (30%) : In the message summary display and the message query display, it should be possible to reply to or delete a message in a single click. To do so, modify `MessageSummary` and `MessageQuery` so that when they print a message, they also print an anchor tag that leads to deletion or update. For example:

```
<a href="blog.pl?act=delete&deleterun=1&id=533">delete</a>  
<a href="blog.pl?act=reply&respid=533">reply</a>
```

These urls will then point to your delete and reply functionality. A user is allowed to delete any of his own messages if he has `delete-own-messages` permission, and any messages if he has `delete-any-messages` permission. In order to reply, he must have `write-messages` permission. A message can only be deleted if no other messages refer to it.

Add tree display (20%) : Modify message summary so that it displays the messages as a tree. You may find the Oracle-specific “CONNECT BY” feature to be useful here, although it is not necessary.

Add multi-page display (30%) : Currently, message summaries and message query results are displayed as a single web page. This will become unworkable as the number of messages increases. In the multi-screen version of MessageSummary and MessageQuery, you will display matching messages n at a time, where n will be a form parameter supplied by the user. At the end of each display, you will include “next” and “previous” links that will allow you to scroll through the matching messages. For example:

```
<a
href="blog.pl?act=query&queryrun=1&from=date1&to=date2&by=b
ywhom&startwith=25&endwith=50">prev</a>
<a
href="blog.pl?act=query&queryrun=1&from=date1&to=date2&by=b
ywhom&startwith=75&endwith=100">next</a>
```

You will find the “rownum” feature of Oracle very useful here:

```
select id from blog_messages where [query clause]
and rownum>=25 and rownum<50
```

Where to go for help

The goal of this project is to learn how a database-backed web application works. Don't fall into the trap of spending lots of time generating pretty HTML or particularly elegant Perl. Get the SQL right first and make the Perl return what you need. We don't want you to get stuck on the details of Perl, CGI, or HTML, so please use the following resources:

- ⇒ Blackboard discussion groups. If you have a question, please check here first, and post it if it hasn't already been answered. We strongly encourage the class to talk on the discussion group, share ideas and insights, and code snippets.
- ⇒ Office hours. Make sure to use the office hours made available by the instructor and the TAs.
- ⇒ Handouts: we have prepared handouts on using Oracle in the TLAB and a high-level introduction to the Perl programming language.
- ⇒ Introductory sessions. We will hold two evening sessions in the first week. One to bring everyone up to speed on Linux, the other to introduce the fundamental concepts behind CGI and Perl.
- ⇒ Web content. You will find many examples of Perl CGI and DBI programming on the web.

Hand-in

To hand in the project, you will do the following:

- ⇒ Get a copy of your extended blog up and running in your account on tlab-server.

- ⇒ Email the instructor and the TAs with a copy of all the files that you have modified or added. At minimum, this will be blog.pl. In the email, also supply a URL for your running blog and the root account name and password.

Extra Credit (30% Maximum)

You can gain extra credit by one or both of the following extensions:

Interface: The Microblog user interface is “primitive” (ugly!). Enhance the interface to make it suitable for ordinary users.

Security: Microblog is horrendously insecure because passwords are stored in cleartext in the database and cookie file, and are sent in cleartext over the network. Read about more secure approaches to session management and implement one of them.