


Pristup relacijskoj bazi putem datotečnog sustava

Dobrica Pavlinušić

dpavlin@rot13.org



<http://www.rot13.org/~dpavlin/>

Linux i perl čine teške stvari jednostavnima,
a nemoguće stvari izvedivima!


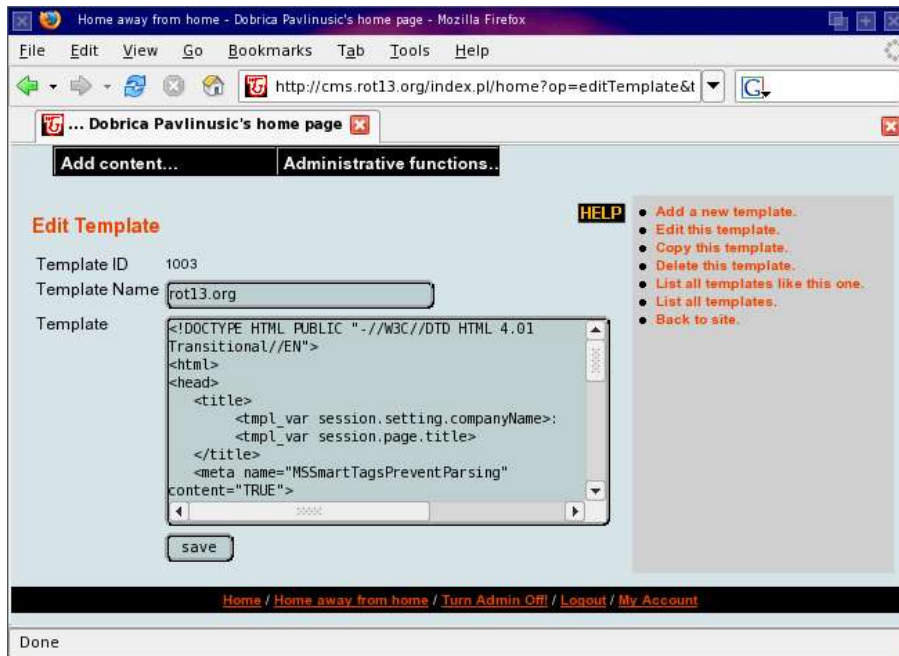


Zašto?

- Oracle je poznat po principu:
Stavimo sve u bazu!
- Ponekad podaci u bazi imaju strukturu koja se može predstaviti datotečnim sistemom
 - predlošci u CMS-ovima
 - datoteke spremljene u BLOB polja baze
 - razni strukturirani podaci u RDBMS-u
 - sjećate se hijerarhijskih baza podataka?




Ali stvarno, zašto?



Vidite li mali **textarea** kvadratić lijevo?


Da li Vam je brže editirati datoteke u **viju**?

Treba li Vas još uvjeravati?





Hmm, to bi moglo biti korisno!

- Kako prići bazi podataka iz editora?
 - dbext.vim script za vi (i 10 različitih baza)
 - http://www.vim.org/scripts/script.php?script_id=356
 - SQL filesystem za Linux
 - SQLFS: kao user-land NFS server
 - <http://sourceforge.net/projects/sqlfs/>
 - RelFS: korištenjem FUSE modula za kernel
 - <http://relfs.sourceforge.net/>
 - to je zapravo urbana legenda (WinFS?)
 - 38,400 hitova na groups.google.com
 - 138,000 hitova na www.google.com



Težak problem...

- Mapiranje relacijske SQL baze u datotečni sustav je težak problem
 - Pregledavanje podataka po tablicama je manje-više jednostavno
 - Komplicirani upiti preko više tablica su teški ili nemogući
 - Kako napraviti update ili insert podataka?
 - Što sa referencijalnim integritetom?
 - Datotečni sustav **nije** baza podataka!




Kako ono izgleda naš problem?

- Osnovni princip korištenja Linuxa je:

Don't panic


- Imamo jednostavnu (ili poznatu) strukturu baze
- Znamo napisati SQL upit koji vraća imena datoteka
- Znamo napisati SQL upit koji vraća sadržaj datoteke i zapisuje promijenjeni sadržaj




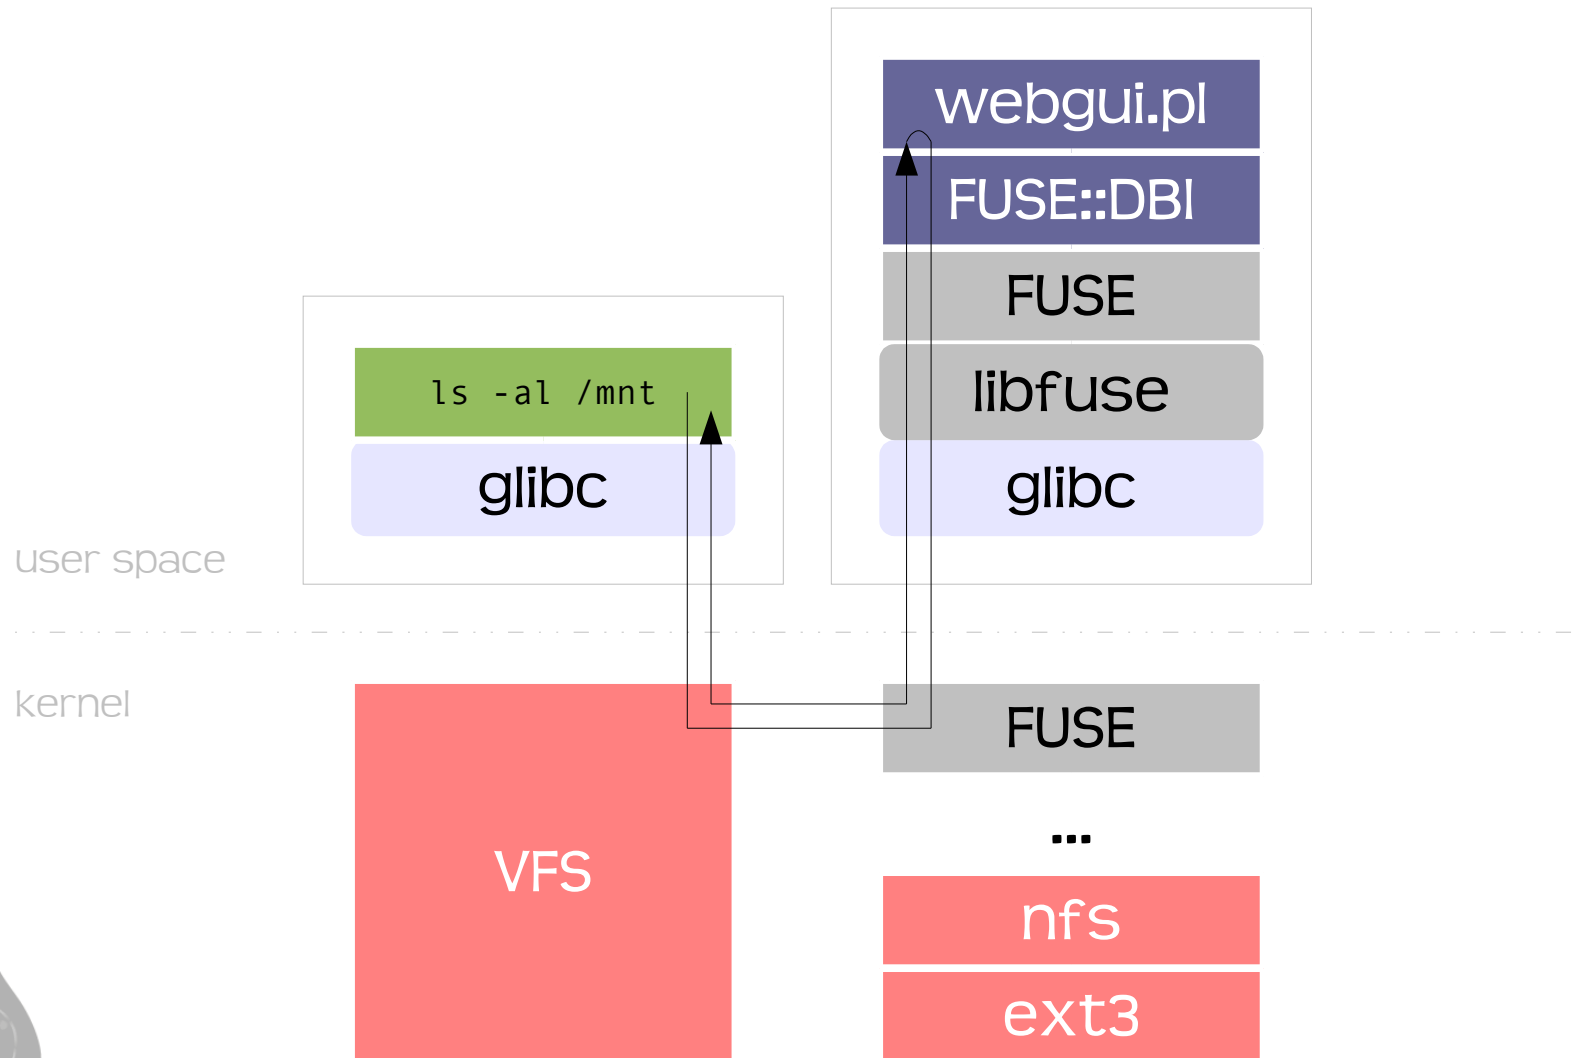
O.K., znamo sve o bazi...

- Ali, niti jedan postojeći filesystem ne podržava takav koncept!
 - KISS*: **napisati vlastiti!**
- Kako napisati Linux filesystem?
 - korištenjem FUSE (**f**ilesystem in **u**erspace) kernel modula sa FUSE perl bindingsima
 - <http://fuse.sourceforge.net/>
- **FUSE::DBI** perl modul

*) Keep It Simple and Stupid



Kako radi FUSE?



26 / 27.10.2004



Jednostavno... **Linux** 2004



FER / Fakultet Elektrotehnike i Racunarstva
Zagreb

FUSE::DBI primjer: WebGUI



<http://www.plainblack.com/webgui>

WebGUI is a content management platform built to allow average business users to build and maintain complex web sites. It is modular, pluggable, and platform independent. It was designed to allow the people who create the content, to manage it online, rather than content management taking up the time of the busy IT Staff.



```
mysql> describe template ;
```

Field	Type	Null	Key	Default	Extra
<code>templateId</code>	int(11)		PRI	0	
<code>name</code>	varchar(255)	YES		NULL	
<code>template</code>	mediumtext	YES		NULL	
<code>namespace</code>	varchar(35)		PRI	Page	
<code>isEditable</code>	int(11)			1	
<code>showInForms</code>	int(11)			1	

```
6 rows in set (0.01 sec)
```

Struktura tablice

```
mysql> describe template ;
```

Field	Type	Null	Key	Default	Extra
templateId	int(11)		PRI	0	
name	varchar(255)	YES		NULL	
template	mediumtext	YES		NULL	
namespace	varchar(35)		PRI	Page	
isEditable	int(11)			1	
showInForms	int(11)			1	

```
6 rows in set (0.01 sec)
```

```
mysql> select
    concat(templateid,name) as id,
    concat(namespace,'/',name,'.html') as filename,
    length(template) as size,
    iseditable as writable
from template;
```

id	filename	size	writable
1Default Site Map	SiteMap/Default Site Map.html	309	1
3Left Align Image	Article/Left Align Image.html	1311	1
2List with Thumbnails	FileManager/List with Thumbnails.html	2961	1
3Calendar Month (Small)	EventsCalendar/Calendar Month (Small).html	2093	1
1Default Event	EventsCalendar/Event/Default Event.html	756	1

Imena datoteka

```
mysql> describe template ;
```

Field	Type	Null	Key	Default	Extra
templateId	int(11)		PRI	0	
name	varchar(255)	YES		NULL	
template	mediumtext	YES		NULL	
namespace	varchar(35)		PRI	Page	
isEditable	int(11)			1	
showInForms	int(11)			1	

```
6 rows in set (0.01 sec)
```

```
mysql> select
  concat(templateid,name) as id,
  concat(namespace,'/',name,'.html') as filename,
  length(template) as size,
  iseditable as writable
from template;
```

id	filename	size	writable
1Default Site Map	SiteMap/Default Site Map.html	109	1
3Left Align Image	Article/Left Align Image.html	1511	1
2List with Thumbnails	FileManager/List with Thumbnails.html	2961	1
3Calendar Month (Small)	EventsCalendar/Calendar Month (Small).html	2693	1
1Default Event	EventsCalendar/Event/Default Event.html	755	1

**Čitanje i
zapisivanje
sadržaja**

```
mysql> select template from template where concat(templateid,name) = ?;
```

```
mysql> update template set template = ? where concat(templateid,name) = ?;
```



Odjećate se zbunjeno?

- **FUSE::DBI** implementira

```
my $mnt = Fuse::DBI->mount({  
    filenames => 'select id,name,size,writable from files as filenames',  
    read => 'sql read',  
    update => 'sql update',  
    dsn => 'DBI:Pg:dbname=webgui',  
    user => 'database_user',  
    password => 'database_password'  
});
```

- **webgui.pl** čita konfiguracijsku datoteku od WebGUI-ja, kreira SQL naredbe i poziva FUSE::DBI

– FUSE::DBI poziva FUSE modul koji poziva libfuse koja komunicira sa kernel modulom...



FUSE::DBI

- zahtijeva FUSE sa perl bindingsima
 - zahtijeva kompiliranje kernel modula i root prava za njegovo učitavanje
 - perl bindingsi dolaze u trenutnoj CVS verziji FUSEa
- dolazi sa primjerom
 - WebGUI CMS
 - očekujem još skripti koje demonstriraju mogućnosti
- na putu prema CPAN-u




Jednostavno...Linux (i perl, FUSE, FUSE::DBI)

```
dpavlin@mjesec:/mnt2$ find . -ls | head -20
  1 549366746 drwxr-xr-x  1 root  root           1 Oct 25 21:44 .
  2 549366746 drwxr-xr-x  1 root  root           1 Oct 25 21:44 ./AttachmentBox
34 549366746 -rw-r--r--  1 root  root          477 Oct 25 21:44 ./AttachmentBox/Attachment\ Box.html
  3 549366746 drwxr-xr-x  1 root  root           3 Oct 25 21:44 ./EventsCalendar
66 549366746 -rw-r--r--  1 root  root        2093 Oct 25 21:44 ./EventsCalendar/Calendar\ Month\ (Small).html
67 549366746 -rw-r--r--  1 root  root       1378 Oct 25 21:44 ./EventsCalendar/Events\ List.html
68 549366746 drwxr-xr-x  1 root  root           1 Oct 25 21:44 ./EventsCalendar/Event
70 549366746 -rw-r--r--  1 root  root         756 Oct 25 21:44 ./EventsCalendar/Event/Default\ Event.html
69 549366746 -rw-r--r--  1 root  root        2036 Oct 25 21:44 ./EventsCalendar/Calendar\ Month\ (Big).html
  4 549366746 drwxr-xr-x  1 root  root           3 Oct 25 21:44 ./Survey
168 549366746 -rw-r--r--  1 root  root       3011 Oct 25 21:44 ./Survey/Default\ Survey.html
169 549366746 drwxr-xr-x  1 root  root           1 Oct 25 21:44 ./Survey/Gradebook
172 549366746 -rw-r--r--  1 root  root       1411 Oct 25 21:44 ./Survey/Gradebook/Default\ Gradebook\ Report.html
170 549366746 drwxr-xr-x  1 root  root           1 Oct 25 21:44 ./Survey/Response
174 549366746 -rw-r--r--  1 root  root       1613 Oct 25 21:44 ./Survey/Response/Default\ Response.html
171 549366746 drwxr-xr-x  1 root  root           1 Oct 25 21:44 ./Survey/Overview
173 549366746 -rw-r--r--  1 root  root       3098 Oct 25 21:44 ./Survey/Overview/Default\ Overview\ Report.html
  5 549366746 drwxr-xr-x  1 root  root           1 Oct 25 21:44 ./prompt
148 549366746 drwxr-xr-x  1 root  root           1 Oct 25 21:44 ./prompt/Default\ WebGUI\ Yes
149 549366746 -rw-r--r--  1 root  root         232 Oct 25 21:44 ./prompt/Default\ WebGUI\ Yes/No\ Prompt.html
dpavlin@mjesec:/mnt2$ vi /mnt2/Survey/Response/Default\ Response.html
```



grep, find, cp, ftp, scp ...



Izbacite uljeza

- Operacije koje nisu podržane sa FUSE::DBI
 - kreiranje novih datoteka ili direktorija
 - previše bi zakompliciralo jednostavnu strukturu
 - jednostavno i intuitivno kreiranje novih templatea preko Web sučelja WebGUI-ja – KISS
 - brisanje datoteka ili direktorija
 - referencijalni integritet
 - u trenutnoj implementaciji invalidira in-memory cache
 - preimenovanje datoteka ili direktorija





In-memory cache? KISS?

- Neke operacije su jednostavno previše zahtjevne za bazu
 - select za čitanje datoteka kod readdir (ls)
- Operacije na filesystemu su u blokovima
 - Operacije na bazi nisu, pročitati podatak samo jednom iz baze (u transakciji)
- Kod otvaranja datoteke njezin sadržaj se pročita u in-memory cache (grep)
- Sadržaj se briše iz cache-a sa unlink (rm)




Goodies!

- transakcije
 - svaki put kod zapisivanja datoteke promjene se commitaju u bazu (rješava istovremeni pristup)
- perl funkcija za invalidiranje
 - omogućava brisanje još jedne cache-irane verzije podataka (npr. templatea na disku)
- fork opcija
 - omogućava pokretanje FUSE::DBI iz skripte i korištenje mountanog filesystema (npr. testovi)





Adrese na mreži

- Distribucija FUSE::DBI modula
 - http://www.rot13.org/~dpavlin/fuse_dbi.html
- Diskusija o FUSE::DBI modulu na PerIMonks
 - http://www.perlmonks.org/?node_id=397814
 - rasprava o drugim pristupima i problemima
- Weblog autora
 - <http://blog.rot13.org/>



Pitanja i diskusija





Jednostavno...Linux 2004

