

Implementing SQL on files

Source: <http://linux.derkeiler.com/Mailing-Lists/Kernel/2004-02/3411.html>

From: Anuradha Ratnaweera (*anuradha_at_linux.lk*)

Date: 02/14/04

To: LKML <linux-kernel@vger.kernel.org>

Date: Sat, 14 Feb 2004 21:36:42 +0600

Hi all,

I am starting to write some code to add a feature which I think would be very useful, and like to get comments and suggestions from LKML. Please ignore this mail if it sounds like nonsense ;-)

Also, if this is already happening somewhere, please enlighten me.

Short version: This feature will add a "table" file type and SQL execution primitives to the kernel, and also relevant userspace programs.

Still reading? Ok, here are the glory details:

Longer version

Both the two popular open source database systems, Postgres and Mysql are going through the filesystem layer to manage databases, as opposed to some commercial counterparts that talk to the block IO layer directly. Also, most databases implement user management separate from the UNIX user database. And the database is separate from the rest of the filesystem.

The feature I am planning to implement will overcome these limitations by implementing a "table" filetype, and the primitive SQL operations in the kernel.

Parsing SQL, optimizing etc. will happen at the user space.

But here comes the interesting part. I am planning to do this without breaking VFS and traditional UNIX syscalls. To clarify, consider this table:

Name Version

Linux-Kernel: Implementing SQL on files

David Weinehall 2.0
Alan Cox 2.2
Marcelo Tosatti 2.4

Using the userspace tools, one can create a "table" file (say maintainers), and insert the data to that file. Each file (or may be filesystem) has two characters (or strings) associated with them: field separator and record separator. Say, colon and newline. If I cat the file:

```
% cat maintainers
David Weinehall:2.0
Alan Cox:2.2
Marcelo Tosatti:2.4
%
```

Now, if I want to add something to the table, either I can use the relevant userspace tools, but the following also will work.

```
% echo 'Linus Torvalds:2.6' > maintainers
%
```

Even if one uses vi to edit the file instead of a simple echo, it won't get that complicated.

The nice thing about this is that now we don't need a separate authentication mechanism. Also, the database no longer needs to be a separate entity (it need not exist at first place!).

As far as I can see, the best way to implement this is as a ReiserFS 4 module.

Like to hear some comments and suggestions.

Thanks for reading all the way here ;-)

Anuradha

--

<http://www.linux.lk/~anuradha/>

-

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@vger.kernel.org

More majordomo info at <http://vger.kernel.org/majordomo-info.html>

Please read the FAQ at <http://www.tux.org/lkml/>