



MySQL and OpenSolaris

Martin 'MC' Brown
Technical Writer
Database Group

About me

- Technical Writer at MySQL
- Long time Solaris user
- OpenSolaris Contributor

Solaris Distributions

- Solaris
 - > Enterprise, Business Critical Applications
 - > Full Support
 - > Quarterly Updates with 3-5 release cycle
- OpenSolaris
 - > Developers and Early Adopters
 - > Target for Latest Innovations
 - > Short Release Cycles (weekly builds, two full releases/year)
- Solaris Express Community|Developer Edition
 - > OpenSolaris + non-open source components

Solaris Distribution Similarities

- Same codebase
- Really, same codebase
- Binary Compatibility between Distributions
 - > Build on Solaris, deploy on OpenSolaris
- Binary Compatibility between Releases
 - > Build for Solaris 10, deploy on Solaris 10u6
- Binary Compatibility between Platforms
 - > Compile on a SPARC desktop, execute on a CMT

MySQL in OpenSolaris

- MySQL 5.0.45 (32-bit) in build 79
 - > SXDE 01/08
- MySQL 5.0.45 (64-bit) in build 87
 - > SXCE 04/08
- MySQL 5.0.45 32/64-bit
 - > OpenSolaris 2008.05
 - > See coalface.mcslp.com for 32-bit/64-bit performance comparison
 - > Built using SunStudio
 - > No further upgrades; 5.1 will be next release
- MySQL 4.x in /usr/sfw/mysql is being removed

Installing MySQL on the command line

- MySQL 5.0 available as a standard package:

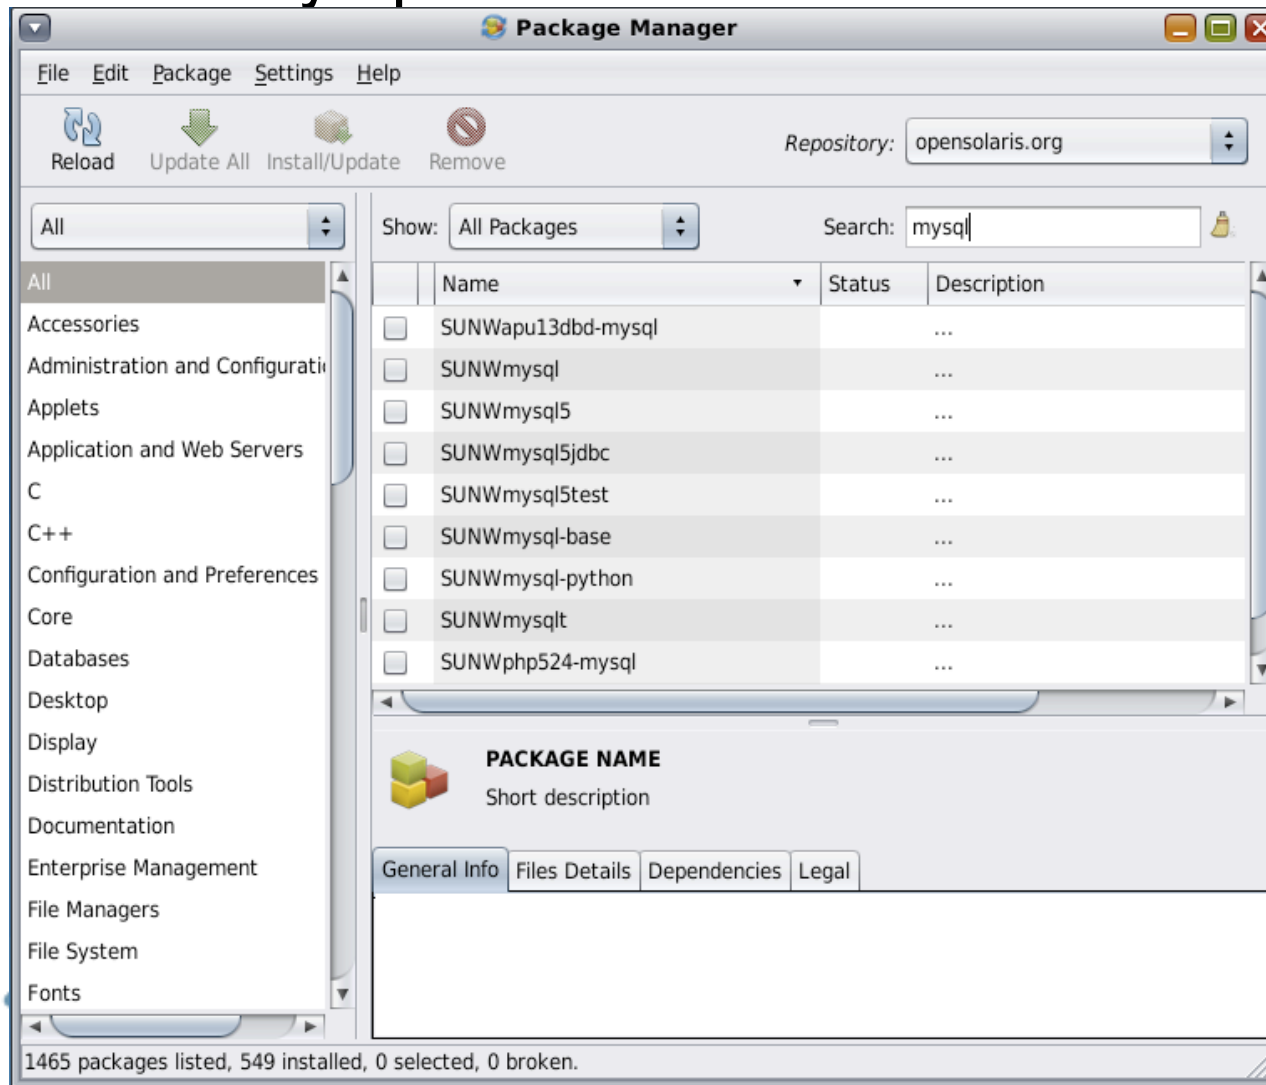
> Login as root/use pfexec

```
# pkg install SUNWmysql5
```

DOWNLOAD	PKGS
FILES XFER (MB)	
Completed	2/2
293/293 123.39/123.39	
PHASE	ACTIONS
Update Phase	2/2
Install Phase	378/378

Installing MySQL with Package Manager

- Select SUNWmysql5



Installation Layout Original Aims

- Support multiple versions
- Allow co-existence between versions
- Support independent starting/stopping of multiple versions
- Flexible data file structure

Installation Layout

- MySQL application structure

```
$ ls -al /usr/mysql
total 17
drwxr-xr-x   3 root   bin      9 Nov 12 08:18 .
drwxr-xr-x  32 root   sys     46 Nov 12 08:17 ..
drwxr-xr-x   8 root   bin     8 Nov 12 08:17 5.0
lrwxrwxrwx   1 root   root    7 Nov 12 08:18 bin -> 5.0/bin
lrwxrwxrwx   1 root   root    8 Nov 12 08:18 docs -> 5.0/docs
lrwxrwxrwx   1 root   root   11 Nov 12 08:18 include -> 5.0/include
lrwxrwxrwx   1 root   root    7 Nov 12 08:18 lib -> 5.0/lib
lrwxrwxrwx   1 root   root    7 Nov 12 08:18 man -> 5.0/man
lrwxrwxrwx   1 root   root    9 Nov 12 08:18 share -> 5.0/share
```

- MySQL data structure

```
$ ls -al /var/mysql/
total 5
drwx-----  3 mysql mysql  4 2008-11-12 09:45 .
drwxr-xr-x  34 root  sys  34 2008-11-12 09:45 ..
drwx-----  3 mysql mysql  3 2008-11-12 09:45 5.0
lrwxrwxrwx   1 root  root  8 2008-11-12 09:45 data -> 5.0/data
```

Installation Layout Updated Theory

- Multiple versions not vital
- Non-standard install location requires adding to PATH
- Confusing layout for migrating users
- Library linkage gets complicated
- Better, standardized, layout for 5.1 already planned

MySQL Integrated into SMF

- SMF provides intelligent service administration
- Self-healing and recovering
- Separate logs per service
- Allows configuration and control
- Import the manifest:

```
# svccfg import /var/svc/manifest/application/database/mysql.xml
```

- Enable/Disable:

```
# svcs mysql
```

```
STATE          STIME          FMRI
```

```
disabled       8:20:41      svc:/application/database/mysql:version_50
```

```
# svcadm enable mysql:version_50
```

Monitoring MySQL with SMF

- Check the mysql service log:

```
# tail /var/svc/log/application-database-mysql\:version_50.log
081112  8:21:17  InnoDB: Log file ./ib_logfile1 did not exist: new to be created
InnoDB: Setting log file ./ib_logfile1 size to 5 MB
InnoDB: Database physically writes the file full: wait...
InnoDB: Doublewrite buffer not found: creating new
InnoDB: Doublewrite buffer created
InnoDB: Creating foreign key constraint system tables
InnoDB: Foreign key constraint system tables created
081112  8:21:18  InnoDB: Started; log sequence number 0 0
081112  8:21:18 [Note] /usr/mysql/5.0/bin/mysqld: ready for connections.
Version: '5.0.45'  socket: '/tmp/mysql.sock'  port: 3306  Source distribution
```

Configuring MySQL with SMF

- Can configure:
 - Data directory
 - 32-bit or 64-bit binaries

- Use svccfg:

```
$ svccfg
```

```
svc:> select mysql:version_50
```

```
svc:/application/database/mysql:version_50> setprop mysql/data=/data0/mysql
```

```
svc:/application/database/mysql:version_50> setprop mysql/enable_64bit=1
```

- Then refresh and restart:

```
# svcadm disable mysql
```

```
# svcadm refresh mysql
```

```
# svcadm enable mysql
```

Quick Performance Tips

- Use the multi-threaded malloc:
 - set LD_PRELOAD environment variable:
 - LD_PRELOAD=/usr/lib/libmtmalloc.so
 - LD_PRELOAD_64=/usr/lib/amd64/libmtmalloc.so
- Use 64-bit on 64-bit architectures
- On CMT, consider using the Coolstack package:
 - <http://cooltools.sunsource.net/coolstack/>
- On standard, consider using the Webstack packages:
 - <http://opensolaris.org/os/project/webstack/>

Quick Performance Tips: ZFS

- Use ZFS
 - Typically faster and easier to manage
 - Limit the ARC cache size to limit RAM use on ZFS
 - When using InnoDB:
 - Eliminate InnoDB tablespace management
 - Set zfs recordsize to match InnoDB block size
 - `# zfs set recordsize=8K zp1/data`
 - ZFS transactions eliminate need for double write buffer
 - Use a separate zpool for logs with default recordsize (128K)
 - When using MyISAM:
 - Create a separate intent log for write-ahead
 - See MySQL section in:
 - http://www.solarisinternals.com/wiki/index.php/ZFS_Best_Practices_Guide

More Performance Tips

- MySQL InnoDB performance tuning for the Solaris 10 OS
 - http://developers.sun.com/solaris/articles/mysql_perf_tune.html
- Neelakanth Nadgir's blog
 - <http://blogs.sun.com/realneel/category/MySQL>
- Sun Studio compiler options for MySQL on Solaris 10 x64 OS
 - http://blogs.sun.com/krishs/entry/sun_studio_compiler_options_for1
- My Blogs
 - <http://coalface.mcslp.com>
 - <http://laptopsolaris.com>

Connectors and Applications

- Connector/J available now:
 - Install SUNWmysql5jdbc, based on C/J 5.1.5
- PHP -> SUNWphp52-mysql
- Python -> SUNWpython-mysql
- Or use Coolstack/Webstack
 - Optimized builds
 - Includes Apache, Lighttpd, Squid
 - PHP, Perl+mod_perl, Python, Ruby+Rails
 - Drupal, MediaWiki, Joomla, Wordpress
 - Glassfish

Quick tips on Compiling MySQL

- For GNU toolchain (gcc+gdb):
 - Install the gcc-dev package
 - Some issues with compiles and dynamic libs
 - Generally works fine
- Better, use SunStudio 12:
 - Make sure your path is correct:
 - `PATH=/opt/SUNWspro/bin:/usr/sfw/bin:...:/usr/ccs/bin`
 - Build:
 - configure `CC=cc CXX=CC [options]`
 - Use dbx, not gdb
 - See krish or my blog for flags for better performance

DTrace Benefits

- DTrace is passive
- Always available
- Low overhead to get the information
- No overhead if you don't use it
- No need to restart or change MySQL or application
- Allows you to trace and monitor any time
- Or monitor constantly and highlight

Using DTrace with MySQL

- Check out for how to use today ...
 - > DTrace and MySQL by Ben Rockwood (Joyent Inc)
 - > <http://en.oreilly.com/mysql2008/public/schedule/detail/1193>
 - > Requires knowledge about MySQL internals
 - > Best Practices for Deploying MySQL on the Solaris Platform by Ritu and Luoja Chen (Sun)
 - > <http://en.oreilly.com/mysql2008/public/schedule/detail/806>
- Making use of DTrace will soon become much easier

DTrace in MySQL

- Instrumenting MySQL with DTrace statically defined probes
- DTrace will allow performance and trace monitoring
 - > Monitor execution time for
 - > SELECT, INSERT, UPDATE, DELETE and variants
 - > Total query execution time
 - > Time in locks
 - > Lock triggers/types
 - > Filesort execution/time
 - > Query Cache hits/misses
- MySQL 6.0 will include an extended set of probes (in 6.0.8)
- OpenSolaris MySQL 5.1 will identical probes (2009.05)

DTrace Example

- Comparing execution for SELECT, 50 million rows

QC	Dur	ms	Matched	Changed	Query
0		0	0	0	select @@version_comment limit 1
0		0	0	0	SELECT DATABASE() ↵
0		0	0	0	show databases
0		0	0	0	show tables
		8289	0		Filesort for select * from s where i > 10 order by i limit 20
0		8290	20	0	select * from s where i > 10 order by i limit 20
0		331594	0	0	create index si on s (i) ↵
0		0	20	0	select * from s where i > 10 order by i limit 20



MySQL and OpenSolaris

Martin 'MC' Brown
Technical Writer
Database Group