

Metacat - Bug #2983

Update PostgreSQL from 7.4 to 8.2 in knb.msi

10/11/2007 02:00 PM - Jing Tao

Status:	Resolved	Start date:	10/11/2007
Priority:	Normal	Due date:	
Assignee:	Jing Tao	% Done:	0%
Category:	metacat	Estimated time:	0.00 hour
Target version:	1.8.1	Spent time:	0.00 hour
Bugzilla-Id:	2983		
Description			
<p>Here is the email from matt: Jing,</p> <p>I see no problem, so go ahead. However, you should be sure to absolutely minimize the downtime for the KNB server, and if at all possible do the upgrade when few people will need access to the server (some evening, or on a weekend). Try to limit the outage to just a minute or two if possible -- can you install both databases side by side so that the version 7 db can continue to run while you set up version 8? Also, can the version 8 db use the db structure from version 7, so that you don't have to do time-consuming table copies?</p> <p>Could you try the upgrade on a test machine first, and develop a proposed upgrade plan that outlines the steps you plan to complete, and how long each step would require to move from 7 to 8? I'd like to review that plan before you proceed.</p> <p>Thanks,</p> <p>Matt</p> <p>Jing Tao wrote:</p> <p>Hi, matt and other devs:</p> <p>Since postgresql 8 has a better performance than 7 and postgresql 8 works so far so good at LTER metacat, I am thinking to upgrade postgresql to 8 in knb.msi. For example, during % search, postgresql 7 couldn't handle one query which contains a list of all documents(because the query is too long), we have to cut the list into couple subset and run queries couple times. It is not efficient. However, postgresql 8 can run this long query.</p> <p>Do you see any problem or have some suggestion for the update?</p> <p>Thanks,</p> <p>Jing</p>			

History

#1 - 10/11/2007 02:12 PM - Jing Tao

In my test machine, I did successfully installed 8 while 7 was running. In order to do so, we need install 8 from source code (If we use binary installation, some files of version 7 will be overwritten). Since data file structure of 8 is different to 7, I have to dump the data to a sql file from 7 and import the sql file to postgresql 8.

Procedure:

```
1. Install postgresql 8.2.5 from source code
root@localhost postgresql-8.2.5]# ./configure --prefix /usr/local/pgsql/8.2.5
root@localhost postgresql-8.2.5]# gmake
root@localhost postgresql-8.2.5]# gmake install
```

2. Start postgresql at another port:

```
root@localhost postgresql-8.2.5]# cd /usr/local/pgsql/8.2.5/
```

```
[root@localhost 8.2.5]# mkdir data
[root@localhost 8.2.5]# chown postgres data
[root@localhost 8.2.5]# su - postgres
-bash-3.00$ cd /usr/local/pgsql/8.2.5/
-bash-3.00$ ./bin/initdb data
edit pg_hba.conf and it looks like:
```

```
local    all             postgres    trust
local    knb             metacat     password
local    kepler          kepler      password
host     all             all         127.0.0.1  255.255.255.255 password
```

```
-bash-3.00$ export LD_LIBRARY_PATH=/usr/local/pgsql/8.2.5/lib
-bash-3.00$ ./bin/postgres -i -p 5433 -D data >/dev/null 2>&1 &
```

3. Create knb and kepler databases, metacat and kepler role in version 8

```
-bash-3.00$ /usr/local/pgsql/8.2.5/bin/createdb -p 5433 knb
-bash-3.00$ /usr/local/pgsql/8.2.5/bin/createdb -p 5433 kepler
-bash-3.00$ ./bin/psql -p 5433 knb
knb=# CREATE ROLE metacat LOGIN PASSWORD 'password';
knb=# CREATE ROLE kepler LOGIN PASSWORD 'password';
knb=# \q
```

4. Dump knb 7.4 postgresql to a file in KNB machine

```
bash-3.00$ pg_dump knb > knbbackup (6.8G)
Start 6:00, end 6:20
pg_dump kepler > keplerbackup (45M)
couple seconds
```

5. Restore database into 8.2

```
-bash-3.00$ ./bin/psql -p 5433 knb < /tmp/knbbackup
```

before export data file to db:

Filesystem	Size	Used	Avail	Use%	Mounted on
/dev/hda1	54G	35G	18G	67%	/
none	506M	0	506M	0%	/dev/shm

started at 9:15 PM ended (?) overnight (but > 3 hours)

Filesystem	Size	Used	Avail	Use%	Mounted on
/dev/hda1	54G	51G	443M	100%	/
none	506M	0	506M	0%	/dev/shm

6. Modify the metacat.properties:

```
defaultDB=jdbc:postgresql://localhost:5433/knb
```

7. Restart tomcat (took 1 minute)

8. Test the new metacat. Since my local machine couldn't handle so huge data (I always get out of memory error for searching), I use ant "registry-schemas" for testing.

9. Shut down the 7.4 postgres

10. Test metacat again.

11. Add a file psql.sh in /etc/profile.d. The file looks like -

```
#
```

```
1. Settings for PostgreSQL database #
PSQL=/usr/local/pgsql/8.2.5
export PSQL
```

```
PATH=$PSQL/bin:$PATH
export PATH
```

```
MANPATH=$PSQL/man:$MANPATH
export MANPATH
```

```
LD_LIBRARY_PATH=$PSQL/lib
export LD_LIBRARY_PATH
```

12. Modify the postgresql in /etc/init.d

a. cp postgresql to postgresql.backup

b. Modify those parameters:

```
PGVERSION=8.2.5 # Set defaults for configuration variables
```

```
PGENGINE=/usr/local/pgsql/8.2.5/bin
PGPORT=5433
export PGDATA=/usr/local/pgsql/8.2.5/data
if [ -f $PGDATA/PG_VERSION ] && [ -d $PGDATA/base/template1 ]
then
echo "Using old-style directory structure"
else
export PGDATA=/usr/local/pgsql/8.2.5/data
fi
```

13. Restart postgresql by /etc/init.d/postgresql restart (Couple seconds)

14. Check process - ps -ef | grep post

15. Check by "netstat -plt"

16. Restart tomcat

17. Test metacat again.

18. Remove postgresql 7.4 (waiting...)

19. Remove data directory of 7.4 (waiting...)

#2 - 10/11/2007 02:38 PM - Jing Tao

Here is the time issue:

Since we can install postgresql 8 while 7 is running, installation time is not a big issue. However, in order to keep the databases identical, there shouldn't be any insert, update or delete activities in the version 7 database when we dump data from 7 to sql file and restore the sql file to version 8 database. The first activity (dump data from 7 to sql file) in knb.msi will take about 20 minutes. The second activity (restore the sql file to version 8 database) is unknown in knb.msi (but in my machine it took > 3 hours - overnight). So let's guess there would be an hour which should NOT allow any insert, update or delete activities in database. I think it is good to do this job in a weekend while shutdown metacat.

#3 - 10/11/2007 03:07 PM - Matt Jones

The disadvantage of installing pg8 from source code is that security patches will not be automatically applied. So, given that you have to shut down the metacat in order to export the db and reimport it into pg 8, I think we might as well go with an RPM upgrade of pg8, so that automatic updates will continue to work.

#4 - 10/11/2007 04:42 PM - Jing Tao

Another issue is about the install dir and the database data dir. The space needed to install postgresql is about 30 M (without data directory). So Postgresql 8 can be installed to anywhere. The default installation directory is /usr/local/pgsql (from source). I propose the installation directory will be /usr/local/pgsql/8.2.5

Let's find a place for data directory. Here is the information of disk usage in knb.msi:

Filesystem	Size	Used	Avail	Use%	Mounted on
/dev/mapper/vg01-sysroot	4.0G	1.3G	2.6G	33%	/
/dev/sda1	471M	54M	394M	12%	/boot
none	4.0G	0	4.0G	0%	/dev/shm
/dev/mapper/vg02-home	9.8G	6.8G	2.5G	74%	/home
/dev/mapper/vg02-tmp	7.9G	2.2G	5.3G	30%	/tmp
/dev/mapper/vg01-usr	29G	6.2G	22G	23%	/usr
/dev/mapper/vg02-var	7.9G	4.2G	3.4G	56%	/var
/dev/mapper/vg03-pgsql	184G	84G	91G	48%	/var/lib/pgsql
/dev/mapper/vg04-metacat	184G	61G	113G	36%	/var/metacat
/dev/mapper/vg05-www	46G	4.6G	39G	11%	/var/www

Since current size of knb database is about 16 G and we should consider future usage, it seems /var/lib/pgsql is a good place for data directory.

Here is the file structure of this dir:

```
[root@knb lib]# ls postgresql/
backups data initdb.i18n knb-backup-05-31-2006 lost+found nohup.out null
postgres-backup-05-31-2006
```

It already has a data directory used by postgresql 7. I propose to create new directory named /var/lib/pgsql/8.2.5/data as the data directory for postgresql 8.

Any suggestion and comment will be appreciated.

#5 - 10/11/2007 05:05 PM - Jing Tao

Consider auto update, RPM upgrade is better choice. No matter rpm or source installation, we have to shutdown metacat to export and import data :(

How about the data directory for 8? Do you think /var/lib/pgsql/8.2.5/data is fine?

Thanks.

#6 - 10/17/2007 04:05 PM - Jing Tao

Here is my propose for the updating by rpms.

1. Download packages :

```
postgresql-8.2.5-1PGDG.rhel4.i686.rpm
postgresql-contrib-8.2.5-1PGDG.rhel4.i686.rpm
postgresql-devel-8.2.5-1PGDG.rhel4.i686.rpm
postgresql-docs-8.2.5-1PGDG.rhel4.i686.rpm
postgresql-libs-8.2.5-1PGDG.rhel4.i686.rpm
postgresql-plperl-8.2.5-1PGDG.rhel4.i686.rpm
postgresql-plpython-8.2.5-1PGDG.rhel4.i686.rpm
postgresql-pltcl-8.2.5-1PGDG.rhel4.i686.rpm
postgresql-server-8.2.5-1PGDG.rhel4.i686.rpm
postgresql-test-8.2.5-1PGDG.rhel4.i686.rpm
```

2. Showdown tomcat

3. Dump knb and kepler database into files

```
bash-3.00$ pg_dump knb >/tmp/knbbackup (6.8G)
```

Start 6:00, end 6:20

```
pg_dump kepler > /tmp/keplerbackup (45M)
```

couple seconds

4. Showdown PostgreSQL 7.4

5. Update PostgreSQL from 7.4 to 8.2

```
rpm -Uvh postgresql*
```

6. Create a dir - /var/lib/pgsql/8.2.5/ by user postgres

7. Modify the /etc/init.d/postgresql file:

```
export PGDATA=/usr/local/pgsql/8.2.5/data
if [ -f $PGDATA/PG_VERSION ] && [ -d $PGDATA/base/template1 ]
then
echo "Using old-style directory structure"
else
export PGDATA=/usr/local/pgsql/8.2.5/data
fi
```

8. Start postgresql - /etc/init.d/postgresql start

9. Create knb and kepler databases (as unix user postgres), metacat and kepler role in version 8

```
-bash-3.00$ createdb knb
```

```
-bash-3.00$ createdb kepler
```

```
-bash-3.00$ ./bin/psql -p 5433 knb
```

```
knb=# CREATE ROLE metacat LOGIN PASSWORD 'password';
```

```
knb=# CREATE ROLE kepler LOGIN PASSWORD 'password';
```

```
knb=# \q
```

10. Restore database into 8.2

```
-bash-3.00$ psql knb < /tmp/knbbackup
```

```
-bash-3.00$ psql kepler < /tmp/keplerbackup
```

11. Edit pg_hba.conf in /usr/local/pgsql/8.2.5/data and it should look like:

local	all	postgres	trust
local	knb	metacat	password
local	kepler	kepler	password
host	all	all	127.0.0.1 255.255.255.255 password

12. Restart PostgreSQL

13. Restart tomcat

14. Test metacat

15. Delete data directory of 7.4 (/var/lib/pgsql/data) after a while (one week late?)

Notes:

1. Since my redhat box is pretty old - Fedora Core 3:(, I couldn't find proper rpms and didn't run it.
2. Could you double check the rpm files are completed? In 7.4, there are jdbc and odbc packages. But I couldn't find them in 8.2
3. Backup plan: if rpm failed. I can go to the installation from source code.

Any suggestion and comment will be appreciated.

#7 - 03/06/2008 06:39 PM - Jing Tao

Here is my new propose for the updating by rpms.

Main change:

- a. Install postgresql 8.3.0 rather than 8.2.5
- b. We should deleted the data in /var/lib/pgsql since no space for two copies.
(Partition have 184 G space, but it is used 107 G).
- c. Base on my experience in dev, I added a step to remove postgresql 7.4 before installation of 8.3.

1. Download packages :

```
compat-postgresql-libs-3-2PGDG.rhel4.i686.rpm 2008-02-05 01:11:10 54.3 KB
compat-postgresql-libs-debuginfo-3-2PGDG.rhel4.i686.rpm 2008-02-05 01:11:12 131.3 KB
postgresql-8.3.0-1PGDG.rhel4.i686.rpm 2008-02-02 02:22:00 1.7 MB
postgresql-contrib-8.3.0-1PGDG.rhel4.i686.rpm 2008-02-02 02:22:07 347.6 KB
postgresql-debuginfo-8.3.0-1PGDG.rhel4.i686.rpm 2008-02-02 02:27:52 12.3 MB
postgresql-devel-8.3.0-1PGDG.rhel4.i686.rpm 2008-02-02 02:28:22 1.3 MB
postgresql-docs-8.3.0-1PGDG.rhel4.i686.rpm 2008-02-02 02:32:50 10.2 MB
postgresql-libs-8.3.0-1PGDG.rhel4.i686.rpm 2008-02-02 02:32:54 185.3 KB
postgresql-plperl-8.3.0-1PGDG.rhel4.i686.rpm 2008-02-02 02:32:55 33.4 KB
postgresql-plpython-8.3.0-1PGDG.rhel4.i686.rpm 2008-02-02 02:32:56 24.5 KB
postgresql-pltcl-8.3.0-1PGDG.rhel4.i686.rpm 2008-02-02 02:32:57 23.3 KB
postgresql-server-8.3.0-1PGDG.rhel4.i686.rpm 2008-02-02 02:34:08 4.4 MB
postgresql-test-8.3.0-1PGDG.rhel4.i686.rpm 2008-02-02 02:34:36 1.2 MB
```

2. Showdown tomcat

3. Dump knb and kepler database into files

```
bash-3.00$ pg_dump knb >/var/lib/pgsql/knbbackup (6.8G)
```

Start 6:00, end 6:20

```
pg_dump kepler > /var/lib/pgsql/keplerbackup (45M)
```

couple seconds

4. Showdown postgresQL 7.4

5. Reomve postgresql 7.4

```
rpm -e $(rpm -qa|grep postgresql) --nodeps
```

6. Install postgresQL 8.3

```
rpm -ivh compat*.rpm
```

```
rpm -ivh postgres*.rpm
```

7. Delete /usr/lib/pgsql/data by user postgres

```
rm data -rf
```

8. Initialize a PostgreSQL database cluster

```
initdb data
```

9. Start postgresql - /etc/init.d/postgresql start

10. Create knb and kepler databases (as unix user postgres), metacat and kepler roles in version 8

```
-bash-3.00$ createdb knb
```

```
-bash-3.00$ createdb kepler
```

```
-bash-3.00$ ./bin/psql -p 5433 knb
```

```
knb=# CREATE ROLE metacat LOGIN PASSWORD 'password';
```

```
knb=# CREATE ROLE kepler LOGIN PASSWORD 'password';
```

```
knb=# \q
```

11. Restore database in 8.3

```
-bash-3.00$psql knb < /tmp/knbbackup
```

```
-bash-3.00$psql kepler < /tmp/keplerbackup
```

12. Edit pg_hba.conf in /usr/local/pgsql/8.2.5/data and it should look like:

local	all	postgres	trust
local	knb	metacat	password
local	kepler	kepler	password

host all all 127.0.0.1 255.255.255.255 password

13. Restart postgresSQL

14. Restart tomcat

15. Test metacat

Backup plan:

If 8.3 installation fails, we can use up2date to install 7.4 again. Then use saved script to restore database.

Note:

1. Metacat 1.8.0 has a sql command which isn't compatible to 8.3. However cvs head does work. Currently, dev machine is running cvs metacat with postgresql 8.3.
2. Postgresql 8.3 has the new feature of free test search.

#8 - 06/16/2008 03:04 PM - Jing Tao

In step 5, i first ran:

```
[root@knb postgres-8.3.3]# rpm -ef `rpm -qa | grep postgresql`
```

error: Failed dependencies:

libpq.so.3 is needed by (installed) perl-DBD-Pg-1.31-6.i386

libpq.so.3 is needed by (installed) mod_auth_psql-2.0.1-7.1.i386

libpq.so.3 is needed by (installed) php-pgsql-4.3.9-3.22.9.i386

Then ran:

```
[root@knb postgres-8.3.3]# rpm -ef `rpm -qa | grep postgresql` --nodeps
```

Note: the command in step has some errors. the above command is correct.

#9 - 06/16/2008 03:34 PM - Jing Tao

I have some issues about compat-postgresql-libs-3-2PGDG.rhel4.i686.rpm and compat-postgresql-libs-4-1PGDG.rhel4.i686.rpm.

Talked with nick we decided only to install compat-postgresql-libs-3-2PGDG.rhel4.i686.rpm

Here is the process:

```
rpm -ivh compat-postgresql-libs*.rpm
```

warning: compat-postgresql-libs-3-2PGDG.rhel4.i686.rpm: V3 DSA signature: NOKEY, key ID 442df0f8

warning: package compat-postgresql-libs = 3-2PGDG.rhel4 was already added, replacing with compat-postgresql-libs <= 4-1PGDG.rhel4

warning: package compat-postgresql-libs-debuginfo = 3-2PGDG.rhel4 was already added, replacing with compat-postgresql-libs-debuginfo <= 4-1PGDG.rhel4

```
Preparing... ##### [100%]
```

```
1:compat-postgresql-libs-##### [ 50%]
```

```
2:compat-postgresql-libs-##### [100%]
```

It seems only 4-1 being installed.

I uninstalled 4-1 and install 3-2.

Here are the des[

```
root@knb postgres-8.3.3]# rpm -qilp compat-postgresql-libs-3-2PGDG.rhel4.i686.rpm
```

warning: compat-postgresql-libs-3-2PGDG.rhel4.i686.rpm: V3 DSA signature: NOKEY, key ID 442df0f8

Name : compat-postgresql-libs Relocations: (not relocatable)

Version : 3 Vendor: (none)

Release : 2PGDG.rhel4 Build Date: Sun 13 Jan 2008 09:32:51 PM PST

Install Date: (not installed) Build Host: koji-centos4-i386-pg83

Group : Applications/Databases Source RPM: compat-postgresql-libs-3-2PGDG.rhel4.src.rpm

Size : 111100 License: BSD

Signature : DSA/SHA1, Fri 01 Feb 2008 04:35:49 PM PST, Key ID 1f16d2e1442df0f8

URL : <http://www.postgresql.org/>

Summary : PostgreSQL client programs and libraries

Description :

This package provides compatibility libraries for PostgreSQL

/usr/lib/libpq.so.3.1

```
[root@knb postgres-8.3.3]# rpm -qilp compat-postgresql-libs-4-1PGDG.rhel4.i686.rpm
```

Name : compat-postgresql-libs Relocations: (not relocatable)

Version : 4 Vendor: (none)

Release : 1PGDG.rhel4 Build Date: Sun 08 Jun 2008 01:44:38 PM PDT

Install Date: (not installed) Build Host: DBI-XEN0

Group : Applications/Databases Source RPM: compat-postgresql-libs-4-1PGDG.rhel4.src.rpm

Size : 115388 License: BSD

Signature : (none)

URL : <http://www.postgresql.org/>

Summary : PostgreSQL client programs and libraries

Description :

This package provides compatibility libraries for PostgreSQL
/usr/lib64/libpq.so.4.1
cription about the two files:

#10 - 06/16/2008 03:35 PM - Jing Tao

Step 7. Delete /usr/lib/pgsql/data by user postgres
rm data -rf

was changed to: mv data data-for-7.4

#11 - 06/26/2008 11:28 AM - Jing Tao

It was done and no error is reported so far. Close this bug.

#12 - 03/27/2013 02:21 PM - Redmine Admin

Original Bugzilla ID was 2983