

## Servers - Bug #1666

### [parabolaweb]: fails to update

2018-03-18 10:26 PM - bill-auger

<b>Status:</b>	fixed	<b>% Done:</b>	100%
<b>Priority:</b>	broken		
<b>Assignee:</b>	lukeshu		
<b>Category:</b>			
<b>Description</b>			
the parabolaweb website has not updated in almost a month			
it looks like all of the services are running normally on proton - the 'parabolaweb-reporead-rsync' systemd cron job syncs the repo with rsync without error every hour but throws many parsing errors of various types from the python script /srv/http/www.parabola.nu/web/devel/management/commands/reporead.py - these could be the source of the problem			
<pre>\$ cp /usr/bin/parabolaweb-reporead-rsync ~ \$ ~/parabolaweb-reporead-rsync</pre>			
there are general classes or error noted below - the entire console output is attached			
-----			
<pre>2018-03-18 22:04:57 -&gt; INFO: adding 21 files for package python2-jsonlines Traceback (most recent call last):   File "/srv/http/www.parabola.nu/web/manage.py", line 10, in &lt;module&gt;     execute_from_command_line(sys.argv)   File "/usr/lib/python2.7/site-packages/django/core/management/__init__.py", line 385, in execute_from_command_line     utility.execute()   File "/usr/lib/python2.7/site-packages/django/core/management/__init__.py", line 377, in execute     self.fetch_command(subcommand).run_from_argv(self.argv)   File "/usr/lib/python2.7/site-packages/django/core/management/base.py", line 288, in run_from_argv     self.execute(*args, **options.__dict__)   File "/usr/lib/python2.7/site-packages/django/core/management/base.py", line 338, in execute     output = self.handle(*args, **options)   File "/srv/http/www.parabola.nu/web/devel/management/commands/reporead.py", line 76, in handle     return read_repo(arch, filename, options)   File "/srv/http/www.parabola.nu/web/devel/management/commands/reporead.py", line 600, in read_repo     db_update(arch, repo, packages_arches[arch], force)   File "/srv/http/www.parabola.nu/web/devel/management/commands/reporead.py", line 402, in db_update     populate_pkg(dbpkg, pkg, timestamp=timestamp)   File "/srv/http/www.parabola.nu/web/devel/management/commands/reporead.py", line 238, in populate_pkg     populate_files(dbpkg, repopkg, force=force)   File "/srv/http/www.parabola.nu/web/devel/management/commands/reporead.py", line 329, in populate_files     batched_bulk_create(PackageFile, pkg_files)   File "/srv/http/www.parabola.nu/web/devel/management/commands/reporead.py", line 278, in batched_bulk_create     return model.objects.bulk_create(all_objects)   File "/usr/lib/python2.7/site-packages/django/db/models/manager.py", line 92, in manager_method     return getattr(self.get_queryset(), name)(*args, **kwargs)   File "/usr/lib/python2.7/site-packages/django/db/models/query.py", line 409, in bulk_create     self._batched_insert(objs_without_pk, fields, batch_size)   File "/usr/lib/python2.7/site-packages/django/db/models/query.py", line 938, in _batched_insert     using=self.db)   File "/usr/lib/python2.7/site-packages/django/db/models/manager.py", line 92, in manager_method     return getattr(self.get_queryset(), name)(*args, **kwargs)   File "/usr/lib/python2.7/site-packages/django/db/models/query.py", line 921, in _insert     return query.get_compiler(using=using).execute_sql(return_id)</pre>			

```
File "/usr/lib/python2.7/site-packages/django/db/models/sql/compiler.py", line 921, in execute_s
ql
    cursor.execute(sql, params)
File "/usr/lib/python2.7/site-packages/django/db/backends/utils.py", line 81, in execute
    return super(CursorDebugWrapper, self).execute(sql, params)
File "/usr/lib/python2.7/site-packages/django/db/backends/utils.py", line 65, in execute
    return self.cursor.execute(sql, params)
File "/usr/lib/python2.7/site-packages/django/db/utils.py", line 94, in __exit__
    six.reraise(dj_exc_type, dj_exc_value, traceback)
File "/usr/lib/python2.7/site-packages/django/db/backends/utils.py", line 65, in execute
    return self.cursor.execute(sql, params)
django.db.utils.DataError: integer out of range
```

-----

2018-03-18 22:05:45 -> INFO: Updating ~fauno (any)

Traceback (most recent call last):

```
File "/srv/http/www.parabola.nu/web/manage.py", line 10, in <module>
    execute_from_command_line(sys.argv)
File "/usr/lib/python2.7/site-packages/django/core/management/__init__.py", line 385, in execute
_from_command_line
    utility.execute()
File "/usr/lib/python2.7/site-packages/django/core/management/__init__.py", line 377, in execute
    self.fetch_command(subcommand).run_from_argv(self.argv)
File "/usr/lib/python2.7/site-packages/django/core/management/base.py", line 288, in run_from_ar
gv
    self.execute(*args, **options.__dict__)
File "/usr/lib/python2.7/site-packages/django/core/management/base.py", line 338, in execute
    output = self.handle(*args, **options)
File "/srv/http/www.parabola.nu/web/devel/management/commands/reporead.py", line 76, in handle
    return read_repo(arch, filename, options)
File "/srv/http/www.parabola.nu/web/devel/management/commands/reporead.py", line 600, in read_re
po
    db_update(arch, repo, packages_arches[arch], force)
File "/srv/http/www.parabola.nu/web/devel/management/commands/reporead.py", line 380, in db_upda
te
    dbpkgs = update_common(archname, reponame, pkgs, True)
File "/srv/http/www.parabola.nu/web/devel/management/commands/reporead.py", line 342, in update_
common
    repository = Repo.objects.get(name__iexact=reponame)
File "/usr/lib/python2.7/site-packages/django/db/models/manager.py", line 92, in manager_method
    return getattr(self.get_queryset(), name)(*args, **kwargs)
File "/usr/lib/python2.7/site-packages/django/db/models/query.py", line 357, in get
    self.model._meta.object_name)
main.models.DoesNotExist: Repo matching query does not exist.
```

-----

2018-03-18 22:05:46 -> INFO: Adding package firestarter

Traceback (most recent call last):

```
File "/srv/http/www.parabola.nu/web/manage.py", line 10, in <module>
    execute_from_command_line(sys.argv)
File "/usr/lib/python2.7/site-packages/django/core/management/__init__.py", line 385, in execute
_from_command_line
    utility.execute()
File "/usr/lib/python2.7/site-packages/django/core/management/__init__.py", line 377, in execute
    self.fetch_command(subcommand).run_from_argv(self.argv)
File "/usr/lib/python2.7/site-packages/django/core/management/base.py", line 288, in run_from_ar
gv
    self.execute(*args, **options.__dict__)
File "/usr/lib/python2.7/site-packages/django/core/management/base.py", line 338, in execute
    output = self.handle(*args, **options)
File "/srv/http/www.parabola.nu/web/devel/management/commands/reporead.py", line 76, in handle
    return read_repo(arch, filename, options)
File "/srv/http/www.parabola.nu/web/devel/management/commands/reporead.py", line 600, in read_re
po
    db_update(arch, repo, packages_arches[arch], force)
```

```
File "/srv/http/www.parabola.nu/web/devel/management/commands/reporead.py", line 402, in db_upda
te
    populate_pkg(dbpkg, pkg, timestamp=timestamp)
File "/srv/http/www.parabola.nu/web/devel/management/commands/reporead.py", line 232, in populat
e_pkg
    dbpkg.signature_bytes = b64decode(repokg.pgpsig.encode('utf-8'))
AttributeError: 'NoneType' object has no attribute 'encode'
```

#### Related issues:

Has duplicate Servers - Bug #1205: [repo] Several user repos cause reporead e...

not-a-bug

#### History

##### #1 - 2018-03-18 10:27 PM - bill-auger

- File *parabolaweb-reporead-rsync-errors* added

##### #2 - 2018-03-22 04:29 PM - bill-auger

- Description updated

##### #3 - 2018-03-27 03:57 AM - lukeshu

I didn't notice this because the script was exiting with a 0 exit code. I've pushed an update to *parabolaweb-utils* so that errors get noticed sooner (and the systemd service fails), but IDK what's causing the error.

##### #4 - 2018-03-28 03:30 AM - lukeshu

Of the 67 (wow, that's a lot!) total repos, there are errors from *reporead* on 36 of them. Not all the same error:

(25) `django.db.utils.DataError: integer out of range`

```
community/os/armv7h/community.files.tar.gz
community/os/i686/community.files.tar.gz
community/os/x86_64/community.files.tar.gz
core/os/armv7h/core.files.tar.gz
core/os/i686/core.files.tar.gz
core/os/x86_64/core.files.tar.gz
extra/os/armv7h/extra.files.tar.gz
extra/os/i686/extra.files.tar.gz
extra/os/x86_64/extra.files.tar.gz
java/os/i686/java.files.tar.gz
java/os/x86_64/java.files.tar.gz
libre/os/armv7h/libre.files.tar.gz
libre/os/i686/libre.files.tar.gz
libre/os/x86_64/libre.files.tar.gz
nonprism/os/i686/nonprism.files.tar.gz
pcr/os/armv7h/pcr.files.tar.gz
pcr/os/i686/pcr.files.tar.gz
pcr/os/x86_64/pcr.files.tar.gz
testing/os/armv7h/testing.files.tar.gz
testing/os/i686/testing.files.tar.gz
testing/os/x86_64/testing.files.tar.gz
~lukeshu/os/i686/~lukeshu.files.tar.gz
~lukeshu/os/x86_64/~lukeshu.files.tar.gz
multilib-testing/os/x86_64/multilib-testing.files.tar.gz
multilib/os/x86_64/multilib.files.tar.gz
```

(8) `AttributeError: 'NoneType' object has no attribute 'encode'`

```
libre-testing/os/i686/libre-testing.files.tar.gz
libre-testing/os/x86_64/libre-testing.files.tar.gz
~jorginho/os/i686/~jorginho.files.tar.gz
~jorginho/os/x86_64/~jorginho.files.tar.gz
~smv/os/i686/~smv.files.tar.gz
~smv/os/x86_64/~smv.files.tar.gz
~xihh/os/i686/~xihh.files.tar.gz
~xihh/os/x86_64/~xihh.files.tar.gz
```

(2) `main.models.DoesNotExist: Repo matching query does not exist.`

```
~fauno/os/i686/~fauno.files.tar.gz
```

```
~fauno/os/x86_64/~fauno.files.tar.gz
```

```
(1) tarfile.ReadError: file could not be opened successfully
```

```
extra/os/x86_64/community.files.tar.gz
```

The tarfile.ReadError is probably transient because of db-import-pkg not updating the file atomically.

TBH, I'm not too worried about The AttributeError (except maybe on [libre-testing]). I'm pretty sure it's caused by old packages that don't have PGP signatures; which makes sense--it's mostly on personal repos of hackers I haven't seen around in quite a while.

Now, django.db.utils.DataError: integer out of range is the target of real concern. I have no idea about what's causing it.

#### #5 - 2018-03-31 12:49 AM - lukeshu

```
parabolitsa=> select max(id) from package_files;
```

```
max
```

```
2147483413
```

```
(1 row)
```

```
parabolitsa=> \d+ package_files
```

```
Table "public.package_files"
```

Column	Type	Modifiers	Storage
Stats target	Description		
id	integer	not null default nextval('package_files_id_seq'::regclass)	plain

```
...
```

```
integer 4 bytes typical choice for integer -2147483648 to +2147483647
```

<https://www.postgresql.org/docs/9.1/static/datatype-numeric.html>

```
$ bc
```

```
2147483647 - 2147483413
```

```
234
```

#### #6 - 2018-03-31 12:53 AM - lukeshu

```
parabolitsa=> select * from package_files_id_seq;
```

```
sequence_name | last_value | start_value | increment_by | max_value | min_value | cache_valu  
e | log_cnt | is_cycled | is_called
```

```
-----+-----+-----+-----+-----+-----+-----  
package_files_id_seq | 2147504862 | 1 | 1 | 9223372036854775807 | 1 |  
1 | 32 | f | t  
(1 row)
```

#### #7 - 2018-03-31 12:54 AM - lukeshu

```
bc
```

```
2147483647 - 2147504862
```

```
-21215
```

#### #8 - 2018-03-31 01:02 AM - lukeshu

currently running:

```
parabolitsa=> ALTER SEQUENCE package_files_id_seq RESTART;
```

```
ALTER SEQUENCE
```

```
parabolitsa=> UPDATE package_files SET id = DEFAULT;
```

**#9 - 2018-03-31 01:12 AM - lukeshu**

Both CPUs are pegged on iowait, iotop tells me postgresql is dominating, although XFS-related kernel tasks also pop up. That makes sense, just wait it out.

systemd-journald is appearing too, which doesn't make too much sense to me.

**#10 - 2018-03-31 03:43 PM - lukeshu**

Followed that with VACUUM ANALYZE package\_files, and things seem mostly OK now.

But I've temporarily disabled parabolaweb-reporead-rsync.timer anyway. It seems that for some reason, some packages are deleted/recreated each run. Include a few packages with tens of thousands of files (like rocksndiamonds-data) in that, and that would explain the churn that lead to us overflowing an int32.

So I'm looking in to why that's happening. I suspect it's related to the signature errors that oaken-source is seeing.

**#11 - 2018-03-31 03:47 PM - lukeshu**

Oh, I should verify that SELECT last\_value FROM package\_files\_id\_seq is >= SELECT max(id) FROM package\_files.

**#12 - 2018-03-31 04:32 PM - lukeshu**

They are in fact ==, which is what is expected.

**#13 - 2018-03-31 04:34 PM - lukeshu**

With 3 consecutive parabolaweb-reporead-rsync runs with the rsync bit disabled (so we're reporeading over static non-changing data), we get the same output for each, with slight variation in package order. Each run takes 10-20 minutes.

**#14 - 2018-03-31 04:37 PM - lukeshu**

It looks like [core] is a good place to debug this--there's some churn on core, but it's just a handful (<10) packages, each with a small number of files.

**#15 - 2018-04-01 11:34 PM - isacdaavid**

I just want to point out that db-import-pkg hasn't been importing anything anyway since the introduction of [this commit](#). the script halts as soon as it fails to find \${LOCAL\_CONFIG}. is that behaviour part of the ongoing debugging efforts?

**#16 - 2018-04-02 12:58 AM - lukeshu**

No, I don't think that's affecting this.

However... whuuuu? The [[ -f ]] guard is meant to handle that. I know the overall status is != 0 for that expression, but set -e doesn't trip for commands with && in them. Is trap ... ERR what's tripping?

**#17 - 2018-04-02 01:17 AM - lukeshu**

Oh, but the non-0 \$? is also taken to be the return code of the source command in db-import-pkg:9.

**#18 - 2018-04-02 01:46 AM - lukeshu**

FYI, I've just pushed a fix for db-import-pkg crashing (dbscripts 20180401).

**#19 - 2018-04-02 02:16 AM - isacdaavid**

lukeshu wrote:

No, I don't think that's affecting this.

i understand that. i was rather asking if your investigations of #note-14 would be disrupted by my appending of || : to db-import-pkg:9 or something along those lines.

many thanks to both of you and everyone who remains active.

**#20 - 2018-05-13 10:49 PM - bill-auger**

this seems to be working now at least for the upstream imports - is it completely fixed for all repos ?

**#21 - 2018-06-18 05:20 AM - lukeshu**

Until the "churn" mentioned in [#1666-10](#), [#1666-13](#), [#1666-14](#) is fixed: Even when we "fix" the PostgreSQL DB (as in [#1666-8](#), [#1666-10](#)), the django.db.utils.DataError: integer out of range issue will keep eventually coming back.

It's worked it's way back up and it's having that issue again now.

**#22 - 2018-07-16 10:45 PM - lukeshu**

I'm trying to do the sequence again:

```
ALTER SEQUENCE package_files_id_seq RESTART;  
UPDATE package_files SET id = DEFAULT;  
VACUUM ANALYZE package_files;
```

But because some of the low ids are still used, I'm having to do the UPDATE in batches.

**#23 - 2018-07-17 06:31 AM - lukeshu**

The churn appears to be **very** bad. I've sudo systemctl stop parabolaweb-reporead-rsync.timer in the mean time.

**#24 - 2018-07-18 02:58 AM - lukeshu**

It's being caused by arch=(any) packages that differ between arches of the same repo, because of db-import-pkg.

1. You're saying that shouldn't happen with current db-import-pkg? Well it is! Between ALARM and Arch.
2. But they're in different pools, surely there's some isolation going on? Well, the assumption in the Arch tools (dbscripts, and archweb) is that all arches of a given reponame share the same pool; eg [extra] and [community] don't need to share the same pool, but [community] (armv7h) and [community] (i686) do.
3. Even if we fix the pool selection bug in db-import-pkg (item 1), that doesn't fix the problem that some arch=(any) packages are present in a repo on one arch but not others, like archlinux32-keyring getting imported to [core] on i686, but [libre] on other arches. This is almost certainly exacerbated by Arch32 holding packages in testing repos for longer than us/Arch/ALARM.

**#25 - 2018-07-18 03:03 AM - lukeshu**

A hack that I'm tempted to do is to have reporead override the package's architecture, and pretend that all packages have the architecture of the repo they're in.

**#26 - 2018-07-18 03:08 AM - lukeshu**

Yeah, until we make some long-term changes to db-import, that's what we'll have to do.

1. Still, it's a short-term bug that db-import-pkg is allowing the same filename in multiple pools.
2. db-update should forbid releasing a file that exists in another pool, unless the new file is a bit-for-bit copy of the old file (as has happened with packages released to [nonsystemd] having the same filename as packages in the packages pool.

**#27 - 2018-07-18 04:10 AM - lukeshu**

I've modified reporead and added fixtures so now any packages show up as "any\_\$(tarch)".

I'm now waiting on it to complete Package.objects.filter(arch\_id=1).delete() to remove the packages with any as the arch, as reporead refuses to do it because it thinks something is wrong.

**#28 - 2018-07-18 04:14 AM - lukeshu**

Given that this bug involves multiple edge cases across multiple different pieces of infrastructure, I think it's fitting that the issue number is [#1666](#).

**#29 - 2018-07-18 04:27 AM - lukeshu**

Oh, I should have just updated the arch\_id to be the any\_x86\_64 id, it would have been faster, and the subsequent reporead would have been faster too. Oh well.

**#30 - 2018-07-18 05:09 AM - lukeshu**

I'm not sure what's wrong, but the any\_ arches aren't working, I'm going to go straight to tarch.

**#31 - 2018-07-18 05:11 AM - brettgilio**

lukeshu wrote:

Given that this bug involves multiple edge cases across multiple different pieces of infrastructure, I think it's fitting that the issue number is [#1666](#).

This means it will reoccur at 2666.

**#32 - 2018-07-18 05:17 AM - lukeshu**

Oh, wait, it was this bit in reporead.py:

```
arch in Arch.objects.filter(agnostic=True):
```

```
packages_arches[arch.name] = []
```

### #33 - 2018-07-18 06:08 AM - lukeshu

Well, it was already done: no more "any" arch in parabolaweb.

Remaining errors from reporead:

```
(10) AttributeError: 'NoneType' object has no attribute 'encode'
```

```
Traceback (most recent call last):
...
File "/srv/http/www.parabola.nu/web/devel/management/commands/reporead.py", line 232, in populate_
pkg
    dbpkg.signature_bytes = b64decode(repopkg.pgpsig.encode('utf-8'))
AttributeError: 'NoneType' object has no attribute 'encode'
```

```
libre-testing/os/i686/libre-testing.files.tar.gz
libre-testing/os/x86_64/libre-testing.files.tar.gz
~jorginho/os/i686/~jorginho.files.tar.gz
~jorginho/os/x86_64/~jorginho.files.tar.gz
~lukeshu/os/i686/~lukeshu.files.tar.gz
~lukeshu/os/x86_64/~lukeshu.files.tar.gz
~smv/os/i686/~smv.files.tar.gz
~smv/os/x86_64/~smv.files.tar.gz
~xihh/os/i686/~xihh.files.tar.gz
~xihh/os/x86_64/~xihh.files.tar.gz
```

```
(6) main.models.DoesNotExist: Repo matching query does not exist.
```

```
nonsystemd-multilib/os/x86_64/nonsystemd-multilib.files.tar.gz
nonsystemd/os/armv7h/nonsystemd.files.tar.gz
nonsystemd/os/i686/nonsystemd.files.tar.gz
nonsystemd/os/x86_64/nonsystemd.files.tar.gz
~fauno/os/i686/~fauno.files.tar.gz
~fauno/os/x86_64/~fauno.files.tar.gz
```

A no-op run of reporead now takes about 2.5 minutes.

### #34 - 2018-07-18 06:14 AM - lukeshu

I'm going to go ahead and re-enable the timer.

### #35 - 2018-07-18 04:17 PM - lukeshu

The missing repos have been added to parabolaweb. That just leaves unsigned packages.

Unsigned packages:

Filename	Build date	Packager
libre-testing/os/i686/mozilla-searchplugins-1.0-1-any.pkg.tar.xz	2011-12-27 02:14:04	Unknown Packager
libre-testing/os/x86_64/iceweasel-libre-9.0.1.1-2-x86_64.pkg.tar.gz	2012-01-04 01:08:56	Luke Shumaker <lukeshu@sbcglobal.net>
libre-testing/os/x86_64/mozilla-searchplugins-1.0-1-any.pkg.tar.xz	2011-12-27 02:14:04	Unknown Packager
~fauno/os/i686/argyllcms-1.3.3-1-i686.pkg.tar.xz	2011-07-28 17:41:35	Nicolás Reynolds <fauno@kiwwwi.com.ar>
~fauno/os/i686/distcc-zeroconf-3.1-1-i686.pkg.tar.xz	2011-06-21 15:35:05	Nicolás Reynolds <fauno@kiwwwi.com.ar>
~fauno/os/i686/emerillon-0.1.90-2-i686.pkg.tar.xz	2011-09-24 04:54:46	Nicolás Reynolds <fauno@kiwwwi.com.ar>
~fauno/os/i686/ethos-0.2.2-6-i686.pkg.tar.xz	2011-09-24 04:19:04	Nicolás Reynolds <fauno@kiwwwi.com.ar>
~fauno/os/i686/faenza-icon-theme-1.2-2-any.pkg.tar.xz	2012-08-31 19:13:05	Nicolás Reynolds <fauno@kiwwwi.com.ar>
~fauno/os/i686/geoclue-0.12-2-i686.pkg.tar.xz	2011-09-24 03:56:56	Nicolás Reynolds <fauno@kiwwwi.com.ar>
~fauno/os/i686/get_flash_videos-1.24-2-i686.pkg.tar.xz	2011-07-20 22:22:18	Nicolás Reynolds <fauno@kiwwwi.com.ar>
~fauno/os/i686/gnome-schedule-2.1.5-1-any.pkg.tar.xz	2011-10-04 00:38:46	Nicolás Reynolds <fauno@kiwwwi.com.ar>

~fauno/os/i686/haskell-base64-bytestring-0.1.0.3-1-i686.pkg.tar.xz 2011-09-26 01:56:13	Nicolás Reynolds <fauno@kiwwwi.com.ar>
~fauno/os/i686/haskell-citeproc-hs-0.3.3-1-i686.pkg.tar.xz 2011-09-26 18:47:44	Nicolás Reynolds <fauno@kiwwwi.com.ar>
~fauno/os/i686/haskell-digest-0.0.1.0-1-i686.pkg.tar.xz 2011-09-26 01:04:09	Nicolás Reynolds <fauno@kiwwwi.com.ar>
~fauno/os/i686/haskell-dlist-0.5-18-i686.pkg.tar.xz 2011-09-26 18:51:15	Nicolás Reynolds <fauno@kiwwwi.com.ar>
~fauno/os/i686/haskell-hs-bibutils-4.12-4-i686.pkg.tar.xz 2011-09-26 02:49:02	Nicolás Reynolds <fauno@kiwwwi.com.ar>
~fauno/os/i686/haskell-json-0.4.4-18-i686.pkg.tar.xz 2011-09-26 17:00:09	Nicolás Reynolds <fauno@kiwwwi.com.ar>
~fauno/os/i686/haskell-pandoc-types-1.8.2-1-i686.pkg.tar.xz 2011-09-26 18:21:41	Nicolás Reynolds <fauno@kiwwwi.com.ar>
~fauno/os/i686/haskell-tagsoup-0.12.3-1-i686.pkg.tar.xz 2011-09-26 18:58:19	Nicolás Reynolds <fauno@kiwwwi.com.ar>
~fauno/os/i686/haskell-texmath-0.5.0.1-19-i686.pkg.tar.xz 2011-09-26 02:08:32	Nicolás Reynolds <fauno@kiwwwi.com.ar>
~fauno/os/i686/haskell-xml-1.3.9-1-i686.pkg.tar.xz 2011-09-26 02:03:15	Nicolás Reynolds <fauno@kiwwwi.com.ar>
~fauno/os/i686/haskell-zip-archive-0.1.1.7-19-i686.pkg.tar.xz 2011-09-26 01:09:18	Nicolás Reynolds <fauno@kiwwwi.com.ar>
~fauno/os/i686/libchamplain-0.12.0-1-i686.pkg.tar.xz 2011-09-24 04:44:23	Nicolás Reynolds <fauno@kiwwwi.com.ar>
~fauno/os/i686/librest-0.7.10-1-i686.pkg.tar.xz 2011-09-24 04:32:55	Nicolás Reynolds <fauno@kiwwwi.com.ar>
~fauno/os/i686/python-html2text-3.02-1-any.pkg.tar.xz 2011-07-27 15:35:55	Nicolás Reynolds <fauno@kiwwwi.com.ar>
~fauno/os/i686/rss2email-2.71-2-any.pkg.tar.xz 2011-07-27 15:42:43	Nicolás Reynolds <fauno@kiwwwi.com.ar>
~fauno/os/i686/ruby-enterprise-1.8.7-12-i686.pkg.tar.xz 2012-02-11 14:51:58	Nicolás Reynolds <fauno@kiwwwi.com.ar>
~fauno/os/i686/ruby-minitar-0.5.2-1-any.pkg.tar.xz 2012-01-10 16:26:34	Nicolás Reynolds <fauno@kiwwwi.com.ar>
~fauno/os/i686/subtle-hg-3171-1-i686.pkg.tar.xz 2012-01-10 16:30:32	Nicolás Reynolds <fauno@kiwwwi.com.ar>
~fauno/os/i686/transmission-remote-cli-0.10.4-1-any.pkg.tar.xz 2011-11-11 20:53:01	Nicolás Reynolds <fauno@kiwwwi.com.ar>
~fauno/os/i686/transmission-remote-cli-git-20110623-1-any.pkg.tar.xz 2011-06-23 05:18:00	fauno <fauno@kiwwwi.com.ar>
~fauno/os/i686/vala-notmuch-20120112-1-i686.pkg.tar.xz 2012-01-12 14:44:44	Nicolás Reynolds <fauno@kiwwwi.com.ar>
~fauno/os/i686/vim-conque-2.3-1-any.pkg.tar.xz 2011-09-28 17:37:29	Nicolás Reynolds <fauno@kiwwwi.com.ar>
~fauno/os/x86_64/faenza-icon-theme-1.2-2-any.pkg.tar.xz 2012-08-31 19:13:05	Nicolás Reynolds <fauno@kiwwwi.com.ar>
~fauno/os/x86_64/gnome-schedule-2.1.5-1-any.pkg.tar.xz 2011-10-04 00:38:46	Nicolás Reynolds <fauno@kiwwwi.com.ar>
~fauno/os/x86_64/python-html2text-3.02-1-any.pkg.tar.xz 2011-07-27 15:35:55	Nicolás Reynolds <fauno@kiwwwi.com.ar>
~fauno/os/x86_64/rss2email-2.71-2-any.pkg.tar.xz 2011-07-27 15:42:43	Nicolás Reynolds <fauno@kiwwwi.com.ar>
~fauno/os/x86_64/ruby-minitar-0.5.2-1-any.pkg.tar.xz 2012-01-10 16:26:34	Nicolás Reynolds <fauno@kiwwwi.com.ar>
~fauno/os/x86_64/transmission-remote-cli-0.10.4-1-any.pkg.tar.xz 2011-11-11 20:53:01	Nicolás Reynolds <fauno@kiwwwi.com.ar>
~fauno/os/x86_64/transmission-remote-cli-git-20110623-1-any.pkg.tar.xz 2011-06-23 05:18:00	fauno <fauno@kiwwwi.com.ar>
~fauno/os/x86_64/vim-conque-2.3-1-any.pkg.tar.xz 2011-09-28 17:37:29	Nicolás Reynolds <fauno@kiwwwi.com.ar>
~jorginho/os/i686/firestarter-1.0.3-12-i686.pkg.tar.xz 2012-04-16 16:58:23	Jorge Lopez <jorginho@adinet.com.uy>
~jorginho/os/x86_64/firestarter-1.0.3-12-x86_64.pkg.tar.xz 2012-04-16 22:49:46	Jorge Lopez <jorginho@adinet.com.uy>
~lukeshu/os/i686/conkeror-git-20120228-1-any.pkg.tar.xz 2012-02-28 07:21:02	Luke Shumaker <lukeshu@sbcglobal.net>
~lukeshu/os/x86_64/conkeror-git-20120228-1-any.pkg.tar.xz 2012-02-28 07:21:02	Luke Shumaker <lukeshu@sbcglobal.net>
~smv/os/i686/eggdrop-1.8-1-i686.pkg.tar.xz g> 2011-07-28 16:47:44	Sorin-Mihai Vărgolici <smv@yobicore.or>
~smv/os/i686/gitosis-git-20110728-1-any.pkg.tar.xz g> 2011-07-28 06:16:09	Sorin-Mihai Vărgolici <smv@yobicore.or>
~smv/os/i686/lxc-0.7.4.1-1-i686.pkg.tar.xz g> 2011-07-25 09:30:18	Sorin-Mihai Vărgolici <smv@yobicore.or>
~smv/os/i686/pan-0.135-1-i686.pkg.tar.xz	Sorin-Mihai Vărgolici <smv@yobicore.or>



```

g> 2011-07-25 10:02:13
~smv/os/i686/planner-0.14.5-1-i686.pkg.tar.xz Sorin-Mihai Vârgolici <smv@yobicore.or
g> 2011-07-25 10:24:44
~smv/os/x86_64/eggdrop-1.8-1-x86_64.pkg.tar.xz Sorin-Mihai Vârgolici <smv@yobicore.or
g> 2011-07-28 16:44:05
~smv/os/x86_64/gitosis-git-20110728-1-any.pkg.tar.xz Sorin-Mihai Vârgolici <smv@yobicore.or
g> 2011-07-28 06:16:09
~smv/os/x86_64/lxc-0.7.4.1-1-x86_64.pkg.tar.xz Sorin-Mihai Vârgolici <smv@yobicore.or
g> 2011-07-25 09:28:55
~smv/os/x86_64/pan-0.135-1-x86_64.pkg.tar.xz Sorin-Mihai Vârgolici <smv@yobicore.or
g> 2011-07-25 09:47:52
~smv/os/x86_64/planner-0.14.5-1-x86_64.pkg.tar.xz Sorin-Mihai Vârgolici <smv@yobicore.or
g> 2011-07-25 10:15:27
~xihh/os/i686/couchapp-git-20120413-1-any.pkg.tar.xz Joshua I. Haase H. <hahj87@gmail.com>
2012-04-13 09:05:35
~xihh/os/i686/python2-argh-0.15.1-1-any.pkg.tar.xz Joshua I. Haase H. <hahj87@gmail.com>
2012-04-13 08:53:26
~xihh/os/i686/python2-pathtools-20120413-1-any.pkg.tar.xz Joshua I. Haase H. <hahj87@gmail.com>
2012-04-13 08:48:59
~xihh/os/i686/python2-restkit-4.0.0-1-any.pkg.tar.xz Joshua I. Haase H. <hahj87@gmail.com>
2012-04-13 08:57:41
~xihh/os/i686/python2-socketpool-0.2.3-1-any.pkg.tar.xz Joshua I. Haase H. <hahj87@gmail.com>
2012-04-13 08:47:45
~xihh/os/i686/python2-watchdog-0.6.0-1-any.pkg.tar.xz Joshua I. Haase H. <hahj87@gmail.com>
2012-04-13 08:54:54
~xihh/os/x86_64/couchapp-git-20120413-1-any.pkg.tar.xz Joshua I. Haase H. <hahj87@gmail.com>
2012-04-13 09:05:35
~xihh/os/x86_64/python2-argh-0.15.1-1-any.pkg.tar.xz Joshua I. Haase H. <hahj87@gmail.com>
2012-04-13 08:53:26
~xihh/os/x86_64/python2-pathtools-20120413-1-any.pkg.tar.xz Joshua I. Haase H. <hahj87@gmail.com>
2012-04-13 08:48:59
~xihh/os/x86_64/python2-restkit-4.0.0-1-any.pkg.tar.xz Joshua I. Haase H. <hahj87@gmail.com>
2012-04-13 08:57:41
~xihh/os/x86_64/python2-socketpool-0.2.3-1-any.pkg.tar.xz Joshua I. Haase H. <hahj87@gmail.com>
2012-04-13 08:47:45
~xihh/os/x86_64/python2-watchdog-0.6.0-1-any.pkg.tar.xz Joshua I. Haase H. <hahj87@gmail.com>
2012-04-13 08:54:54

```

All of those are from either 2011 or 2012.

### #36 - 2018-07-18 04:58 PM - lukeshu

```

[repo@winston ~]$ db-remove libre-testing any mozilla-searchplugins
==> Removing packages from [libre-testing]...
-> mozilla-searchplugins (x86_64)
(none)
-> mozilla-searchplugins (i686)
(none)
-> mozilla-searchplugins (armv7h)
(none)
==> WARNING: No packages remain, creating empty database.

```

Oh look, it's a phantom package.

### #37 - 2018-07-18 05:28 PM - bill-auger

lukeshu wrote:

That just leaves unsigned packages.

is that [#1652](#) ?

### #38 - 2018-07-18 05:35 PM - bill-auger

there is also the issue of the "maintainer" filter [#1897](#)

### #39 - 2018-07-18 05:40 PM - lukeshu

Discrepancies between desc files in the .files and .db databases:

```

[lukeshu@beefcake repo-sanity]$ diff -r db files
Only in files/libre-testing/os/i686: mozilla-searchplugins-1.0-1
Only in files/libre-testing/os/x86_64: iceweasel-libre-9.0.1.1-2
Only in files/libre-testing/os/x86_64: mozilla-searchplugins-1.0-1

```

```
Only in files/pcr/os/x86_64: sfml-git-20120812-1
Only in files/~fauno/os/i686: argyllcms-1.3.3-1
Only in files/~fauno/os/i686: get_flash_videos-1.24-2
Only in files/~fauno/os/i686: python-html2text-3.02-1
Only in files/~fauno/os/i686: rss2email-2.71-2
Only in files/~fauno/os/x86_64: python-html2text-3.02-1
Only in files/~fauno/os/x86_64: rss2email-2.71-2
Only in files/~smv/os/i686: eggdrop-1.8-1
Only in files/~smv/os/i686: gitosis-git-20110728-1
Only in files/~smv/os/i686: lxc-0.7.4.1-1
Only in files/~smv/os/i686: pan-0.135-1
Only in files/~smv/os/i686: planner-0.14.5-1
Only in files/~smv/os/x86_64: eggdrop-1.8-1
Only in files/~smv/os/x86_64: gitosis-git-20110728-1
Only in files/~smv/os/x86_64: lxc-0.7.4.1-1
Only in files/~smv/os/x86_64: pan-0.135-1
Only in files/~smv/os/x86_64: planner-0.14.5-1
```

#### #40 - 2018-07-18 05:43 PM - lukeshu

bill-auger wrote:

lukeshu wrote:

That just leaves unsigned packages.

is that [#1652](#) ?

No, that's something else (mis-configured makepkg.conf?) The unsigned packages thing is packages that were uploaded without a PGP signature. Consider <https://www.parabola.nu/packages/testing/i686/python-astroid/> it says "Signed By: Unknown (0xA952AC0F)" because Erich's PGP key isn't set up in parabolaweb, but it at least has the key fingerprint there.

#### #41 - 2018-07-18 08:23 PM - lukeshu

- % Done changed from 0 to 70

If you were wondering, my criteria for closing this is: `systemctl --failed` no longer says `parabolaweb-reporead-rsync.service`.

#### #42 - 2018-08-23 03:34 AM - lukeshu

I just ran:

```
ssh repo@repo.parabola.nu
cd /srv/repo
(set -euE -o pipefail; find main/~{fauno,jorginho,smv,xihh} -not -type d|while read -r src; do dst=http/${src#main/}; mkdir -p -- "${dst%/*}"; mv -Tv -- "$src" "$dst"; done)
find main/~{fauno,jorginho,smv,xihh} -depth -type d -exec rmdir -- {} \;
```

#### #43 - 2018-08-23 04:35 AM - lukeshu

I manually deleted these packages from parabolaweb:

- `~xihh/i686/django-merengue`
- `~xihh-x86_64/django-merengue`

Then I deleted these repos from parabolaweb:

- `~fauno`
- `~jorginho`
- `~smv`
- `~xihh`

#### #44 - 2018-08-23 04:56 AM - lukeshu

- Has duplicate Bug #1205: [repo] Several user repos cause reporead errors. added

#### #45 - 2018-09-22 01:22 AM - lukeshu

I've tricked `repo-remove` in to removing the packages from the repos with discrepancies by renaming the `repo.files.tar.gz` to `repo.db.tar.gz` in a temporary directory.

It was also complaining that [build-support] for i686 had only 45% of the packages in the web DB. So I've marked it as a staging repo in parabolaweb to skip that check.

**#46 - 2018-09-22 01:37 AM - lukeshu**

- Assignee set to lukeshu

The only remaining errors are coming from [~lukeshu].

**#47 - 2018-09-22 02:31 AM - lukeshu**

I've now run:

```
db-remove '~lukeshu' x86_64 conkeror-git  
db-remove '~lukeshu' i686 conkeror-git
```

**#48 - 2018-09-22 03:04 AM - lukeshu**

- % Done changed from 70 to 100

- Status changed from open to fixed

And it ran cleanly!

**#49 - 2018-09-26 08:11 PM - lukeshu**

When I'd manually edited the .files databases, they lost g+w permissions, so only I could push to [pcr] or [libre-testing].

## Files

---

parabolaweb-reporead-rsync-errors	163 KB	2018-03-18	bill-auger
-----------------------------------	--------	------------	------------