

From: G lineau, Samuel - Montreal
Date: Thu, 1 Nov 2007 17:31:54 -0400

Hi,

Your svn2cvs testsuite failed on my machine, and I had to fix both the testsuite and the perl script. I assume it must have worked for you the last time you ran your testsuite, and I have no clue why. I'm using CVS 1.12.9.

```
> cd svn2cvs/src
```

Running ./test.sh, the first abnormal behaviour I notice is at svn commit 14, "remove everything":

```
...  
+ svn rm $svnco/dir/*  
D /dev/shm/test-svn-co/dir/bar  
D /dev/shm/test-svn-co/dir/baz  
D /dev/shm/test-svn-co/dir/keep  
D /dev/shm/test-svn-co/dir/keep-dir  
D /dev/shm/test-svn-co/dir/l1/bar  
D /dev/shm/test-svn-co/dir/l1/baz  
D /dev/shm/test-svn-co/dir/l1/l2/bar  
D /dev/shm/test-svn-co/dir/l1/l2/baz  
D /dev/shm/test-svn-co/dir/l1/l2/l3/bar  
D /dev/shm/test-svn-co/dir/l1/l2/l3/baz  
D /dev/shm/test-svn-co/dir/l1/l2/l3  
D /dev/shm/test-svn-co/dir/l1/l2  
D /dev/shm/test-svn-co/dir/l1  
D /dev/shm/test-svn-co/dir/skip_add  
D /dev/shm/test-svn-co/dir/with space  
+ svn revert $svnco/dir/keep $svnco/dir/keep-dir/keep  
Reverted '/dev/shm/test-svn-co/dir/keep'  
Skipped '/dev/shm/test-svn-co/dir/keep-dir/keep'  
+ svn commit -m 'remove everything' $svnco  
Deleting /dev/shm/test-svn-co/dir/bar  
Deleting /dev/shm/test-svn-co/dir/baz  
Deleting /dev/shm/test-svn-co/dir/keep-dir  
Deleting /dev/shm/test-svn-co/dir/l1  
Deleting /dev/shm/test-svn-co/dir/skip_add  
Deleting /dev/shm/test-svn-co/dir/with space  
  
Committed revision 14.  
...
```

The interesting lines are:

```
Skipped '/dev/shm/test-svn-co/dir/keep-dir/keep'  
...  
Deleting /dev/shm/test-svn-co/dir/keep-dir
```

It's obvious from the "svn revert" line that you never intended keep-dir to be deleted, yet it does get deleted. I'm using svn 1.3.1 (r19032).

This time, the fix is in the testsuite itself:

```

--- svn2cvs/src/test.sh 2007-11-01 19:41:19.000000000 +0000
+++ svn2cvs/src/test.sh 2007-11-01 19:49:37.000000000 +0000
@@ -145,7 +145,7 @@
test

svn rm $svn_co/dir/* || exit
-svn revert $svn_co/dir/keep $svn_co/dir/keep-dir/keep
+svn revert $svn_co/dir/keep $svn_co/dir/keep-dir{,/keep}
svn commit -m "remove everything" $svn_co || exit

test

```

The relevant commit in the new testsuite now runs as follows:

```

...
+ svn rm $svnco/dir/*
D /dev/shm/test-svn-co/dir/bar
D /dev/shm/test-svn-co/dir/baz
D /dev/shm/test-svn-co/dir/keep
D /dev/shm/test-svn-co/dir/keep-dir
D /dev/shm/test-svn-co/dir/l1/bar
D /dev/shm/test-svn-co/dir/l1/baz
D /dev/shm/test-svn-co/dir/l1/l2/bar
D /dev/shm/test-svn-co/dir/l1/l2/baz
D /dev/shm/test-svn-co/dir/l1/l2/l3/bar
D /dev/shm/test-svn-co/dir/l1/l2/l3/baz
D /dev/shm/test-svn-co/dir/l1/l2/l3
D /dev/shm/test-svn-co/dir/l1/l2
D /dev/shm/test-svn-co/dir/l1
D /dev/shm/test-svn-co/dir/skip_add
D /dev/shm/test-svn-co/dir/with space
+ svn revert $svnco/dir/keep $svnco/dir/keep-dir $svnco/dir/keep-dir/keep
Reverted '/dev/shm/test-svn-co/dir/keep'
Reverted '/dev/shm/test-svn-co/dir/keep-dir'
Skipped '/dev/shm/test-svn-co/dir/keep-dir/keep'
+ svn commit -m 'remove everything' $svnco
Deleting /dev/shm/test-svn-co/dir/bar
Deleting /dev/shm/test-svn-co/dir/baz
Deleting /dev/shm/test-svn-co/dir/l1
Deleting /dev/shm/test-svn-co/dir/skip_add
Deleting /dev/shm/test-svn-co/dir/with space

Committed revision 14.
...

```

This time, the "skip" is legitimate. The testsuite still fails, though:

```

...
+ diff -x '.svn*' -x CVS -urw $svnco/dir/ $cvsco/dir/
Only in /dev/shm/test-cvs-co//dir/: l1
Only in /dev/shm/test-cvs-co//dir/: with space
+ exit

```

It again failed at this problematic last commit, "remove everything":

```

...
Starting after revision 13

```

```
## svn export --force -q -r 14 file:///dev/shm/test-svn-rep//dir /tmp/checkoutGGpuy/dir
NOTICE: using /dir as directory for svn
```

```
-----
r 14| gelineaus | 2007-11-01T20:22:39.115442Z
```

```
remove everything
svn2cvs: D /dir/bar
#### remove file: bar at ./svn2cvs.pl line 402.
## cvs -f -d /dev/shm/test-cvs-rep/ delete 'bar'
cvs remove: scheduling `bar' for removal
cvs remove: use `cvs commit' to remove this file permanently
## cvs -f -d /dev/shm/test-cvs-rep/ commit -m 'remove everything' 'bar'
/dev/shm/test-cvs-rep/dir/bar,v <-- bar
new revision: delete; previous revision: 1.1
svn2cvs: D /dir/with space
WARNING: with space is not present in CVS, skipping...
svn2cvs: D /dir/baz
#### remove file: baz at ./svn2cvs.pl line 402.
## cvs -f -d /dev/shm/test-cvs-rep/ delete 'baz'
cvs remove: scheduling `baz' for removal
cvs remove: use `cvs commit' to remove this file permanently
## cvs -f -d /dev/shm/test-cvs-rep/ commit -m 'remove everything' 'baz'
/dev/shm/test-cvs-rep/dir/baz,v <-- baz
new revision: delete; previous revision: 1.1
svn2cvs: D /dir/l1
WARNING: l1 is not present in CVS, skipping...
svn2cvs: D /dir/skip_add
#### remove file: skip_add at ./svn2cvs.pl line 402.
## cvs -f -d /dev/shm/test-cvs-rep/ delete 'skip_add'
cvs remove: scheduling `skip_add' for removal
cvs remove: use `cvs commit' to remove this file permanently
## cvs -f -d /dev/shm/test-cvs-rep/ commit -m 'remove everything' 'skip_add'
/dev/shm/test-cvs-rep/dir/skip_add,v <-- skip_add
new revision: delete; previous revision: 1.1
commit ignored, no files
subversion revision 14 committed to CVS
/dev/shm/test-cvs-rep/dir/.svnrev,v <-- .svnrev
new revision: 1.15; previous revision: 1.14
+ update_all
+ update_svn
+ svn update /dev/shm/test-svn-co/
At revision 14.
+ update_cvs
+ cd /dev/shm/test-cvs-co/
+ cvs -f update -d dir
cvs update: Updating dir
U dir/.svnrev
cvs update: `dir/bar' is no longer in the repository
cvs update: `dir/baz' is no longer in the repository
cvs update: `dir/skip_add' is no longer in the repository
cvs update: Updating dir/keep-dir
cvs update: Updating dir/l1
cvs update: Updating dir/l1/l2
cvs update: Updating dir/l1/l2/l3
cvs update: Updating dir/with space
+ cd-
+ diff -x '*.svn*' -x CVS -urw /dev/shm/test-svn-co//dir/ /dev/shm/test-cvs-co//dir/
Only in /dev/shm/test-cvs-co//dir/: l1
Only in /dev/shm/test-cvs-co//dir/: with space
+ exit
```

The relevant lines are:

```

...
svn2cvs: D /dir/with space
WARNING: with space is not present in CVS, skipping...
...
svn2cvs: D /dir/l1
WARNING: l1 is not present in CVS, skipping...
...
Only in /dev/shm/test-cvs-co//dir/: l1
Only in /dev/shm/test-cvs-co//dir/: with space
+ exit

```

In other words, `svn2cvs.pl` cannot see directories and refuses to remove them. This is because of the `in_entries()` function, which does not consider lines describing directories. Here's the fix:

```

--- svn2cvs/src/svn2cvs.pl 2007-11-01 20:21:22.000000000 +0000
+++ svn2cvs/src/svn2cvs.pl 2007-11-01 20:36:52.000000000 +0000
@@ -263,6 +263,7 @@ sub in_entries($) {
         || return 0; #die "no entries file: $dir/CVS/Entries";
         while (<$fh>) {
             return 1 if (m{^/$file/});
+
             return 1 if (m{^D/$file/});
         }
         close($fh);
         return 0;

```

But the problematic last commit *still* fails:

```

Starting after revision 13
## svn export --force -q -r 14 file:///dev/shm/test-svn-rep//dir /tmp/checkout1jkCl/dir
NOTICE: using /dir as directory for svn

-----
r 14| gelineaus | 2007-11-01T20:53:32.113268Z

remove everything
svn2cvs: D /dir/bar
#### remove file: bar at ./svn2cvs.pl line 403.
## cvs -f -d /dev/shm/test-cvs-rep/ delete 'bar'
cvs remove: scheduling `bar' for removal
cvs remove: use `cvs commit' to remove this file permanently
## cvs -f -d /dev/shm/test-cvs-rep/ commit -m 'remove everything' 'bar'
/dev/shm/test-cvs-rep/dir/bar,v <-- bar
new revision: delete; previous revision: 1.1
svn2cvs: D /dir/with space
#### remove directory: with space at ./svn2cvs.pl line 383.
#### entries(with space) => at ./svn2cvs.pl line 249.
## cvs -f -d /dev/shm/test-cvs-rep/ delete 'with space'
cvs remove: Removing with space
## cvs -f -d /dev/shm/test-cvs-rep/ commit -m 'remove everything' 'with space'
cvs commit: Examining with space
## cvs -f -d /dev/shm/test-cvs-rep/ update -dP .
cvs update: Updating .
cvs update: Updating keep-dir
cvs update: Updating l1
cvs update: Updating l1/l2
cvs update: Updating l1/l2/l3
cvs update: Updating with space
svn2cvs: D /dir/baz
#### remove file: baz at ./svn2cvs.pl line 403.
## cvs -f -d /dev/shm/test-cvs-rep/ delete 'baz'

```

```
cvcs remove: scheduling `baz' for removal
cvcs remove: use `cvcs commit' to remove this file permanently
## cvcs -f -d /dev/shm/test-cvcs-rep/ commit -m 'remove everything' 'baz'
/dev/shm/test-cvcs-rep/dir/baz,v <-- baz
new revision: delete; previous revision: 1.1
svn2cvcs: D /dir/l1
#### remove directory: l1 at ./svn2cvcs.pl line 383.
#### entries recurse into: l1/l2 at ./svn2cvcs.pl line 239, <$fh> line 3.
#### entries recurse into: l1/l2/l3 at ./svn2cvcs.pl line 239, <> line 3.
#### entries(l1/l2/l3) => bar|baz at ./svn2cvcs.pl line 249.
#### entries(l1/l2) => bar|baz|l3/bar|l3/baz|l3 at ./svn2cvcs.pl line 249.
#### entries(l1) => bar|baz|l2/bar|l2/baz|l2/l3/bar|l2/l3/baz|l2/l3|l2 at ./svn2cvcs.pl line 249.
## cvcs -f -d /dev/shm/test-cvcs-rep/ delete 'l1/bar'
cvcs remove: scheduling `l1/bar' for removal
cvcs remove: use `cvcs commit' to remove this file permanently
## cvcs -f -d /dev/shm/test-cvcs-rep/ commit -m 'remove everything' 'l1/bar'
/dev/shm/test-cvcs-rep/dir/l1/bar,v <-- bar
new revision: delete; previous revision: 1.1
## cvcs -f -d /dev/shm/test-cvcs-rep/ delete 'l1/baz'
cvcs remove: scheduling `l1/baz' for removal
cvcs remove: use `cvcs commit' to remove this file permanently
## cvcs -f -d /dev/shm/test-cvcs-rep/ commit -m 'remove everything' 'l1/baz'
/dev/shm/test-cvcs-rep/dir/l1/baz,v <-- baz
new revision: delete; previous revision: 1.1
## cvcs -f -d /dev/shm/test-cvcs-rep/ delete 'l1/l2/bar'
cvcs remove: scheduling `l1/l2/bar' for removal
cvcs remove: use `cvcs commit' to remove this file permanently
## cvcs -f -d /dev/shm/test-cvcs-rep/ commit -m 'remove everything' 'l1/l2/bar'
/dev/shm/test-cvcs-rep/dir/l1/l2/bar,v <-- bar
new revision: delete; previous revision: 1.2
## cvcs -f -d /dev/shm/test-cvcs-rep/ delete 'l1/l2/baz'
cvcs remove: scheduling `l1/l2/baz' for removal
cvcs remove: use `cvcs commit' to remove this file permanently
## cvcs -f -d /dev/shm/test-cvcs-rep/ commit -m 'remove everything' 'l1/l2/baz'
/dev/shm/test-cvcs-rep/dir/l1/l2/baz,v <-- baz
new revision: delete; previous revision: 1.2
## cvcs -f -d /dev/shm/test-cvcs-rep/ delete 'l1/l2/l3/bar'
cvcs remove: scheduling `l1/l2/l3/bar' for removal
cvcs remove: use `cvcs commit' to remove this file permanently
## cvcs -f -d /dev/shm/test-cvcs-rep/ commit -m 'remove everything' 'l1/l2/l3/bar'
/dev/shm/test-cvcs-rep/dir/l1/l2/l3/bar,v <-- bar
new revision: delete; previous revision: 1.3
## cvcs -f -d /dev/shm/test-cvcs-rep/ delete 'l1/l2/l3/baz'
cvcs remove: scheduling `l1/l2/l3/baz' for removal
cvcs remove: use `cvcs commit' to remove this file permanently
## cvcs -f -d /dev/shm/test-cvcs-rep/ commit -m 'remove everything' 'l1/l2/l3/baz'
/dev/shm/test-cvcs-rep/dir/l1/l2/l3/baz,v <-- baz
new revision: delete; previous revision: 1.3
## cvcs -f -d /dev/shm/test-cvcs-rep/ delete 'l1/l2/l3'
cvcs remove: Removing l1/l2/l3
## cvcs -f -d /dev/shm/test-cvcs-rep/ commit -m 'remove everything' 'l1/l2/l3'
cvcs commit: Examining l1/l2/l3
## cvcs -f -d /dev/shm/test-cvcs-rep/ delete 'l1/l2'
cvcs remove: Removing l1/l2
cvcs remove: Removing l1/l2/l3
## cvcs -f -d /dev/shm/test-cvcs-rep/ commit -m 'remove everything' 'l1/l2'
cvcs commit: Examining l1/l2
cvcs commit: Examining l1/l2/l3
## cvcs -f -d /dev/shm/test-cvcs-rep/ delete 'l1'
cvcs remove: Removing l1
cvcs remove: Removing l1/l2
cvcs remove: Removing l1/l2/l3
## cvcs -f -d /dev/shm/test-cvcs-rep/ commit -m 'remove everything' 'l1'
cvcs commit: Examining l1
cvcs commit: Examining l1/l2
cvcs commit: Examining l1/l2/l3
```

```

## cvs -f -d /dev/shm/test-cvs-rep/ update -dP .
cvs update: Updating .
cvs update: Updating keep-dir
cvs update: Updating l1
cvs update: Updating l1/l2
cvs update: Updating l1/l2/l3
cvs update: Updating with space
svn2cvs: D /dir/skip_add
#### remove file: skip_add at ./svn2cvs.pl line 403.
## cvs -f -d /dev/shm/test-cvs-rep/ delete 'skip_add'
cvs remove: scheduling `skip_add' for removal
cvs remove: use `cvs commit' to remove this file permanently
## cvs -f -d /dev/shm/test-cvs-rep/ commit -m 'remove everything' 'skip_add'
/dev/shm/test-cvs-rep/dir/skip_add,v <-- skip_add
new revision: delete; previous revision: 1.1
commit ignored, no files
subversion revision 14 committed to CVS
/dev/shm/test-cvs-rep/dir/.svnrev,v <-- .svnrev
new revision: 1.15; previous revision: 1.14
+ update_all
+ update_svn
+ svn update /dev/shm/test-svn-co/
At revision 14.
+ update_cvs
+ cd /dev/shm/test-cvs-co/
+ cvs -f update -d dir
cvs update: Updating dir
U dir/.svnrev
cvs update: `dir/bar' is no longer in the repository
cvs update: `dir/baz' is no longer in the repository
cvs update: `dir/skip_add' is no longer in the repository
cvs update: Updating dir/keep-dir
cvs update: Updating dir/l1
cvs update: `dir/l1/bar' is no longer in the repository
cvs update: `dir/l1/baz' is no longer in the repository
cvs update: Updating dir/l1/l2
cvs update: `dir/l1/l2/bar' is no longer in the repository
cvs update: `dir/l1/l2/baz' is no longer in the repository
cvs update: Updating dir/l1/l2/l3
cvs update: `dir/l1/l2/l3/bar' is no longer in the repository
cvs update: `dir/l1/l2/l3/baz' is no longer in the repository
cvs update: Updating dir/with space
+ cd-
+ diff -x '*.svn*' -x CVS -urw /dev/shm/test-svn-co//dir/ /dev/shm/test-cvs-co//dir/
Only in /dev/shm/test-cvs-co//dir/: l1
Only in /dev/shm/test-cvs-co//dir/: with space
+ exit

```

The relevant lines are:

```

...
svn2cvs: D /dir/with space
## cvs -f -d /dev/shm/test-cvs-rep/ delete 'with space'
cvs remove: Removing with space
## cvs -f -d /dev/shm/test-cvs-rep/ commit -m 'remove everything' 'with space'
...
## cvs -f -d /dev/shm/test-cvs-rep/ delete 'l1'
cvs remove: Removing l1
cvs remove: Removing l1/l2
cvs remove: Removing l1/l2/l3
## cvs -f -d /dev/shm/test-cvs-rep/ commit -m 'remove everything' 'l1'
...
Only in /dev/shm/test-cvs-co//dir/: l1

```

```
Only in /dev/shm/test-cvs-co//dir/: with space
+ exit
```

Thus cvs was asked to delete the directories, but didn't. This is no surprise, since CVS was *designed* not to remove directories. From the CVS manual:

7.3 Removing directories

In concept, removing directories is somewhat similar to removing files--you want the directory to not exist in your current working directories, but you also want to be able to retrieve old releases in which the directory existed.

The way that you remove a directory is to remove all the files in it. You don't remove the directory itself; there is no way to do that. Instead you specify the '-P' option to cvs update or cvs checkout, which will cause CVS to remove empty directories from working directories. (Note that cvs export always removes empty directories.) Probably the best way to do this is to always specify '-P'; if you want an empty directory then put a dummy file (for example '.keepme') in it to prevent '-P' from removing it.

Note that '-P' is implied by the '-r' or '-D' options of checkout. This way, CVS will be able to correctly create the directory or not depending on whether the particular version you are checking out contains any files in that directory.

Thus the fix would be to automatically add a ".keepme" file upon directory creation, and to removed it (along with the other files in the directory) upon directory deletion. The testsuite also needs to be updated to use the "-P" flag. Thus, my last patch:

```
diff -r -uprN svn2cvs/src/svn2cvs.pl svn2cvs/src/svn2cvs.pl
--- svn2cvs/src/svn2cvs.pl 2007-11-01 20:36:52.000000000 +0000
+++ svn2cvs/src/svn2cvs.pl 2007-11-01 21:06:04.000000000 +0000
@@ -104,7 +104,8 @@ sub add_dir($$) {
     next if in_entries($curr_dir);
     next if ( -e "$curr_dir/CVS" );

+         log_system( "touch '$curr_dir/.keepme'", "creation of .keepme file (to keep $curr_dir)" );
-         log_system( "$cvs add '$curr_dir'", "cvs add of $curr_dir failed" );
+         log_system( "$cvs add '$curr_dir' '$curr_dir/.keepme'", "cvs add of $curr_dir failed" );
     }
}

diff -r -uprN svn2cvs/src/test.sh svn2cvs/src/test.sh
--- svn2cvs/src/test.sh 2007-11-01 20:36:32.000000000 +0000
+++ svn2cvs/src/test.sh 2007-11-01 21:06:39.000000000 +0000
@@ -52,13 +52,13 @@ cd- || exit

rm -Rf $cvs_co || exit
mkdir $cvs_co || exit
```

```

-cd $cvs_co && cvs -f co dir && cd- || exit
+cd $cvs_co && cvs -f co -P dir && cd- || exit

function svn2cvs() {
    ./svn2cvs.pl file://$svn_rep/dir $cvs_rep dir || exit
}
function update_cvs() {
-   cd $cvs_co && cvs -f update -d dir && cd- || exit
+   cd $cvs_co && cvs -f update -P -d dir && cd- || exit
}
function update_svn() {
    svn update $svn_co || exit
@@ -70,7 +70,7 @@ function update_all() {
function test() {
    svn2cvs
    update_all
-   diff -x .svn\* -x CVS -urw $svn_co/dir/ $cvs_co/dir/ || exit
+   diff -x .svn\* -x CVS -x '\.keep' -urw $svn_co/dir/ $cvs_co/dir/ || exit
}

svn2cvs

```

And now the problematic test finally passes.

Hoping to be useful,

- Samuel Gélineau

This is one of best bug reports I ever received or actually read :-). It seems that other projects also appreciate bug reports like:

- https://bugs.eclipse.org/bugs/show_bug.cgi?id=96111
- https://bugs.eclipse.org/bugs/show_bug.cgi?id=98320
- https://bugs.eclipse.org/bugs/show_bug.cgi?id=98592
- https://bugs.eclipse.org/bugs/show_bug.cgi?id=99168
- https://bugs.eclipse.org/bugs/show_bug.cgi?id=99228
- https://bugs.eclipse.org/bugs/show_bug.cgi?id=100227
- https://bugs.eclipse.org/bugs/show_bug.cgi?id=100260
- https://bugs.eclipse.org/bugs/show_bug.cgi?id=102212
- https://bugs.eclipse.org/bugs/show_bug.cgi?id=102357
- https://bugs.eclipse.org/bugs/show_bug.cgi?id=104024
- https://bugs.eclipse.org/bugs/show_bug.cgi?id=105181
- https://bugs.eclipse.org/bugs/show_bug.cgi?id=106630
- https://bugs.eclipse.org/bugs/show_bug.cgi?id=106634
- https://bugs.eclipse.org/bugs/show_bug.cgi?id=106874
- https://bugs.eclipse.org/bugs/show_bug.cgi?id=107059
- https://bugs.eclipse.org/bugs/show_bug.cgi?id=107858
- https://bugs.eclipse.org/bugs/show_bug.cgi?id=107898
- https://bugs.eclipse.org/bugs/show_bug.cgi?id=108014
- https://bugs.eclipse.org/bugs/show_bug.cgi?id=104046