

sshfs auto reconnect

Scenario: using SSHFS mounts from my laptop to a central server, the SSHFS mount is broken after a longer disconnect (e.g. during suspend), cause the underlying SSH connection timed out.

Use `-o reconnect,ServerAliveInterval=15,ServerAliveCountMax=3`

The combination `ServerAliveInterval=15,ServerAliveCountMax=3` causes the I/O errors to pop out after one minute of network outage. This is important but largely undocumented. If `ServerAliveInterval` option is left at default (so without the alive check), processes which experience I/O hang seem to sleep indefinitely, even after the sshfs gets reconnected. I regard this a useless behaviour.

In other words what happens on `-o reconnect` without assigning `ServerAliveInterval` is that any I/O will either succeed, or hang the application indefinitely if the ssh reconnects underneath. A typical application becomes entirely hung as a result. If you'd wish to allow I/O to return an error and resume the application, you need `ServerAliveInterval=1` or greater.

The `ServerAliveCountMax=3` is the default anyway, but I like to specify it for readability.

fstab entry example for laptop, automount after network available:

```
nuc:/    /media/kodi    fuse.sshfs
reconnect,ServerAliveInterval=15,ServerAliveCountMax=3,auto,users,_netdev,x-
systemd.after=network-online.target    0    0
```

From:

<http://wuff.dyndns.org/> - **Wulf's Various Things**

Permanent link:

<http://wuff.dyndns.org/doku.php?id=linux:sshfs&rev=1595199204>

Last update: **2023/05/29 11:53**

