

Debian jessie builds of DKMSed upstream DRBD8 Kernel Module

Motivation

History of DRBD at ITEG & Clazzes.org

For many years (since CentOS 5) we at ITEG (also the main force behind Clazzes.org) have used DRBD to mirror the partition of each virtual machine (then OpenVZ, now mostly LXC).

Kernel 4.6 hickups caused by DRBD 8.4.6 to 8.4.9-1

After shredding some data last year (probably caused by a combination of aggressive settings and unfortunate module versions), it started to make big troubles again when a VLAN/Bridging bug forced us to upgrade our Debian jessie machines to kernels from jessie-backports (kernels 4.4 to now 4.7) with DRBD module 8.4.6 up to upstream version 8.4.9-1.

For details see [Ever-rising load on Debian jessie + DRBD8 + LXC host pairs](#).

Debianizing the most-recent upstream release of the kernel modul was our way to divide (as in *divide and conquer*).

We now know that a wrong kernel 4.0+ adaption in DRBD8 was the root cause of that problem - and **we had a kernel expert fix it for us (and hence everybody)**, see [this post in the drbd-dev mailing list](#) and followups.

As of [this commit in Linus' linux.git](#) the patch we triggered (and financed) made it into the kernel mainstream, i.e. 4.9. It will also be part of all older 4.x kernels still maintained. The next LTS kernel to be, 4.10, also chosen by Debian for `stretch`, will get a rewrite of the whole section and hence also be free of that bug.

But as long as the fix does no make it into Debian `jessie-backports` (via upstream kernel upgrades) and as long as `stretch` does not become `stable` we will continue to provide the DKMS-debianized DRBD module as described below, #OpenSourceRules.

Up-to-date & FIXED DRBD 8 packages in Clazzes.org' Debian repository

Introduction

The packages below contain the most recent DRBD8 module plus the "kernel_sendmsg() usage" patch, and are in permanent and successful use on all 10 nodes we maintain.

Disclaimers

1. We have no intentions to damage Linbit's business model, but the `kernel_sendmsg()` bug (that should have been recognized and fixed much earlier) almost destroyed one of our business models, so I won't hold back on this one for as long as this is resolved for everybody.
2. Anyone using these packages does so at one's own peril.

Packages

Clazzes.org's Deb server [deb.clazzes.org](#) contains a repository "jessie-drbdpkg-8" providing 2 DRBD packages:

`drbd8-dkms 8.4.9-1clazzes6`

This packages contains the most-recently available sources of the DRBD8 module, the "kernel_sendmsg() usage" patch, and DKMS integration for Debian jessie (probably usable for jessie-based derivates too).

On installation of `drbd8-dkms_8.4.9-1clazzes6` (or later) alongside matching pairs of `linux-header-*-amd64` and `linux-image-*-amd64` packages the up-to-date DRBD module is automatically built and installed.

Remark 1: On some nodes `dkms` triggered the installation of a `linux-image-3.2.0-4-rt-amd64` which can be removed afterwards.

Remark 2: DKMS sucks at building modules for kernel versions that are installed but not yet active, and we haven't got it under perfect control yet. See bottom of this page.

We'll update this package as long as `stretch` does not become `stable`. and `jessie-backports` does not include the patch (or goes to 4.10).

The Debianization part is open sourced too, see [svn.clazzes.org/svn/drbdpkg/trunk/drbd8-dkms/](#).

`drbd-utils 8.9.9-1`

We 'cheated' a bit here: `drbd-utils_8.9.9-1` is the package from Debian's `unstable/experimental` repositories (i.e. [here](#)), re-integrated in our DRBD8-repository.

This way any Debian `jessie` installation has access to up-to-date DRBD8 packages without the need to care about compiling manually or adding Debian `unstable` to `sources.list`.

We'll update this package as long as `stretch` does not become `stable`.

Repository activation & package installation Quick Guide

```
wget -O - https://deb.clazzes.org/gpg/pba-archiver.clazzes.org.asc |apt-key add -

cd /etc/apt/sources.list.d
wget https://deb.clazzes.org/debian/sources.list.d/jessie/jessie-drbdpkg-8.list

# evtl. change jessie-drbdpkg-8.list to use https over http, if apt-transport-https is installed

apt-get update

apt-get install drbd8-dkms drbd-utils

# get secondary on all drbd resources
# restart DRBD
/etc/init.d/drbd stop
/etc/init.d/drbd start
# check version:
cat /proc/drbd |head -1
```

DKMS Problems, Solutions, hints

Problem

DKMS somehow fails to build the DRBD8 module correctly for newer kernels that are installed but not yet active, so after reboot the DRBD module can't be loaded:

```
modprobe: ERROR: could not insert 'drbd': Exec format error
```

Solution used

We solved it with ...

```
dpkg-reconfigure drbd8-dkms

# ... or ...
apt-get remove drbd8-dkms
apt-get install drbd8-dkms

# ... and finally ...
/etc/init.d/drbd start
```

Version checks

To check which `drbd8` module version is loaded:

```
cat /proc/drbd |head -1
```

To check which version the kernel will load next time it loads the `drbd8` module:

```
modinfo drbd |egrep -i "^version:"
```