

Dracut

A Generic Initramfs Infrastructure

Harald Hoyer
Red Hat

Dracut

- Tools used on the system
- Event based with udev
- Generic across distributions
- Generic across hardware
- Nothing hardcoded

Targets

- Boot any system configuration on any hardware with the same initramfs image
- Easy to extend and customize

Modules

- l18n
- Network
- Plymouth
- Crypt, DMraid, LVM, Mdrraid
- dmsquash-live
- FCoE, iSCSI, NBD, NFS
- resume, rootfs-block
- s390

Synopsis

```
dracut [OPTION]... <image> <kernel-version>
```

Options

-f, --force

overwrite existing initramfs file.

-m, --modules LIST

specify a space-separated list of dracut modules to call when building the initramfs. Modules are located in `/usr/share/dracut/modules.d`.

-o, --omit LIST

omit a space-separated list of dracut modules.

-a, --add LIST

add a space-separated list of dracut modules.

Options

-d, --drivers LIST

specify a space-separated list of kernel modules to exclusively include in the initramfs.

--add-drivers LIST

specify a space-separated list of kernel modules to add to the initramfs.

--filesystems LIST

specify a space-separated list of kernel filesystem modules to exclusively include in the generic initramfs.

-k, --kmoddir {DIR}

specify the directory, where to look for kernel modules

Options

- `--fwdir {DIR}`
specify additional directory, where to look for firmwares
- `--kernel-only`
only install kernel drivers and firmware files
- `--no-kernel`
do not install kernel drivers and firmware files
- `--mdadmconf`
include local `/etc/mdadm.conf`
- `--nomdadmconf`
do not include local `/etc/mdadm.conf`

Options

`--strip`

strip binaries in the initramfs (default)

`--nostrip`

do not strip binaries in the initramfs

`-c, --conf FILE`

specify configuration file to use.

Default: `/etc/dracut.conf`

`-l, --local`

local mode. Use modules from the current working directory instead of the system-wide installed in `/usr/share/dracut/modules.d`. Useful when running dracut from a git checkout.

Options

-H, --hostonly

Host-Only mode: Install only what is needed for booting the local host instead of a generic host.

-i, --include SOURCE TARGET

include the files in the SOURCE directory into the target directory in the final initramfs.

-I, --install LIST

install the space separated list of files into the initramfs.

dracut.conf

dracutmodules=""

omit_dracutmodules=""

add_drivers=""

filesystems=""

hostonly="yes"

mdadmconf="yes"

lvmconf="yes"

Kernel Command Line

- `root=/dev/sda1`
- `root=LABEL=rootPart`
- `root=UUID=uuidnumber`
- `root=/dev/disk/by-path/pci-0000:00:1f.1-scsi-0:0:1:0-part1`
- `rootfstype=<filesystem type>`
- `rootflags=<mount options>`

Kernel Command Line

- `rdblacklist=<drivername>`
- `rdloaddriver=<drivername>`

Kernel Command Line

- `rdinfo`
print informational output though "quiet" is set
- `rdshell`
allow dropping to a shell, if root mounting fails
- `rdinitdebug`
set `-x` for the dracut shell and logs to `dmesg`,
console and `/init.log`
- `rdbreak`
drop to a shell at the end

Kernel Command Line

- `rdbreak={cmdline|pre-udev|pre-trigger|initqueue|pre-mount|mount|pre-pivot}`
drop to a shell on defined breakpoint
- `rdudevinfo`
set udev to loglevel info
- `rdudevdebug`
set udev to loglevel debug
- `rdnetdebug`
debug network scripts in dracut. Output is written to `/tmp/`

LVM

- `rd_LVM_VG=<volume group name>`
only activate the volume groups with the given name `rd_LVM_VG` can be specified multiple times on the kernel command line.
- `rd_LVM_LV=<logical volume name>`
only activate the volume groups with the given name `rd_LVM_LV` can be specified multiple times on the kernel command line.

crypto LUKS

- rd_NO_LUKS
- rd_LUKS_UUID=<luks uuid>
- rd_NO_CRYPT-TAB

MD

- rd_NO_MD
- rd_NO_MDIMSM
- rd_NO_MDADMCONF
- rd_MD_UUID=<md uuid>

DMRAID

- rd_NO_DM
- rd_DM_UUID=<dmraid uuid>

Network

- ip={dhcp|on|any|dhcp6|auto6}
- ip=<interface>:{dhcp|on|any|dhcp6|auto6}
- ip=<client-IP>:[<server-id>]:<gateway-IP>:<netmask>:<clienthostname>:<interface>:
{none|off}
- ifname=<interface>:<MAC>
- bootdev=<interface>
- nameserver=<IP> [nameserver=<IP> ...]

NFS

- root=[<server-ip>:]<root-dir>[:<nfs-options>]
- root=nfs:[<server-ip>:]<root-dir>[:<nfs-options>]
- root=nfs4:[<server-ip>:]<root-dir>[:<nfs-options>]
- root=dhcp|dhcp6
root-path=<server-ip>:<root-dir>[,<nfs-options>]

iSCSI

- root=iscsi:[username:password[:reverse:password]@]
[<servername>]:[<protocol>]:[<port>]:
[<LUN>]:<targetname>
- root=??? iscsi_initiator=<initiator>
iscsi_target_name=<target name>
iscsi_target_ip=<target ip> iscsi_target_port=<target
port> iscsi_target_group=<target group>
iscsi_username=<username>
iscsi_password=<password> iscsi_in_username=<in
username> iscsi_in_password=<in password>
- root=??? iscsi_firmware

FCoE & NBD

- root=fcoe:<interface|MAC>:<dcb|nodcb>
- root=nbd:<server>:<port>[:<fstype>]
[:<mountopts>]
- root=dhcp

Basic setup

Hooks: cmdline, pre-udev

Start Udev

Hooks: pre-trigger

Trigger Udev

Initqueue

Wait for jobs or udev settled

Initqueue settled

Initqueue finished

Found root device

Hooks: pre-mount

Hooks: mount

Hooks: pre-pivot

Cleanup and switch_root

Initqueues

- Initqueue
- Initqueue-settled
- Initqueue-finished
- /sbin/initqueue
 - -onetime
 - -settled
 - --finished
 - -unique -name

Modules

- install
- installkernel
- check

Install

```
dracut_install ip dhclient brctl arping
```

```
inst "$moddir/ifup" "/sbin/ifup"
```

```
inst "$moddir/netroot" "/sbin/netroot"
```

```
inst "$moddir/dhclient-script" "/sbin/dhclient-script"
```

```
inst "$moddir/dhclient.conf" "/etc/dhclient.conf"
```

```
inst_hook pre-udev 50 "$moddir/iface-name-genrules.sh"
```

```
inst_hook cmdline 91 "$moddir/dhcp-root.sh"
```

```
inst_hook cmdline 99 "$moddir/parse-ip-opts.sh"
```

```
inst_hook pre-pivot 10 "$moddir/kill-dhclient.sh"
```

Installkernel

```
net_module_test() {
    local net_drivers='eth_type_trans|register_virtio_device'
    local unwanted_drivers='/(wireless|isdn|uwb)/'
    nm -uPA "$1" | egrep -q $net_drivers && \
        nm -uPA "$1" | egrep -qv 'iw_handler_get_spy' && \
        [[ ! $1 =~ $unwanted_drivers ]]
}
instmods $(filter_kernel_modules net_module_test)
instmods ecb arc4
```

Check

```
[ "$1" = "-d" ] && echo network
```

```
[ "$1" = "-h" ] && ! egrep -q '/ nfs[34 ]' /proc/mounts &&  
exit 1
```

```
which rpcbind >/dev/null 2>&1 || which portmap  
>/dev/null 2>&1 || exit 1
```

```
which rpc.statd mount.nfs mount.nfs4 umount >/dev/null  
2>&1 || exit 1
```

```
exit 0
```

Test Suite

- Uses qemu
- Covers basic modules like LVM, Dmraid, Mdrraid, NBD, iSCSI, NFS
- Client, Server communication over qemu's internal network

Other Use Cases

- Cluster setups
- Rescue System
- Kexec crash dumps
- Special Hardware, Disk Layout (OLPC)

Participate

`irc://irc.freenode.net/dracut`

`git://dracut.git.sourceforge.net/gitroot/dracut/dracut`

<http://sourceforge.net/projects/dracut>

<http://sourceforge.net/apps/trac/dracut>

Initramfs Mailing List on kernel.org

Thanks for Listening!

Questions?