

Dracut

A Generic Initramfs
Infrastructure

Harald Hoyer
Red Hat

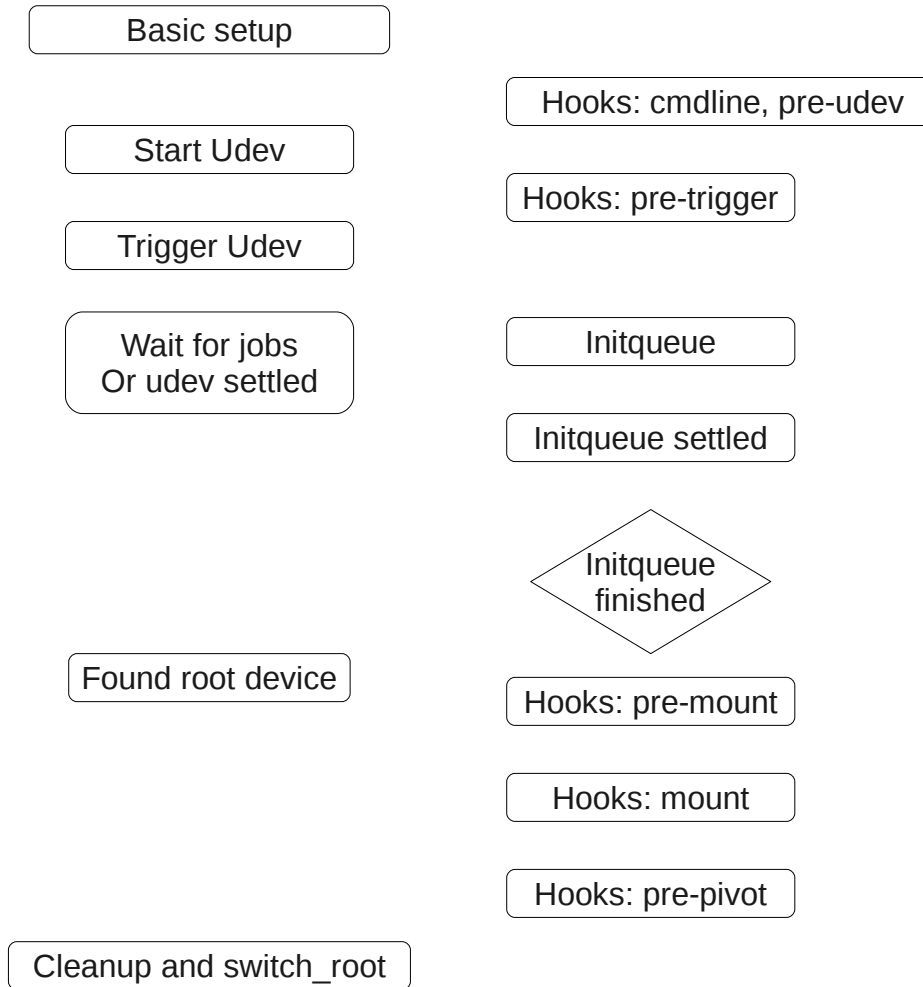
Dracut

- 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

Overview



Modules

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

Hooks

- Hooks: cmdline, pre-udev, pre-trigger, pre-mount, pre-pivot, emergency
- POSIX shell scripts ending with .sh in /\$hookname
- rdbreak=\$hookname

Queues

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

DMraid

- Udev rules match on
ENV{ID_FS_TYPE}=="*_raid_member"
- \$sys/\$devpath/holders/dm-[0-9]*
- dmraid -ay -i -p --rm_partitions "\$setname"
- GUID Partition Table (GPT):
kpartx -a -p p "/dev/mapper/\$setname"

MDraid

- `ENV{ID_FS_TYPE}=="linux_raid_member|isw_raid_member"`
- `$sys/$devpath/holders/md[0-9]*`
- `RUN+="/sbin/mdadm -I --no-degraded $env{DEVNAME}"`
- `/sbin/mdadm -As --auto=yes --run`

LVM

- `ENV{ID_FS_TYPE}=="LVM2_member"`
- Create `/etc/lvm/lvm.conf`
`filter = ["a|^/dev/DEVNAME$", "r/*/"]`
- `lvm vgscan`
- `lvm vgchange -ay`

Network

- ip={dhcp|on|any}
- ip=<interface>:{dhcp|on|any}
- ip=<client-IP-number>:[<server-id>]:<gateway-IP-number>:<netmask>:<client-hostname>:<interface>:{none|off}
- ifname=<interface>:<MAC>
- ACTION=="add", SUBSYSTEM=="net", RUN+="/sbin/ifup \$env{INTERFACE}"
- Match specific interfaces by name or MAC
- Run DHCP if needed

NFS

- root=[<server-ip>:]<root-dir>[:<nfs-options>]
- root=nfs:[<server-ip>:]<root-dir>[:<nfs-options>]
- root=nfs4:[<server-ip>:]<root-dir>[:<nfs-options>]
- Just try to mount it

NBD

- `root=nbd:<server>:<port>[:<fstype>]`
`[:<mountopts>]`
- `nbd-client`
- Wait for `/dev/nbd0`

iSCSI

- root=iscsi:
[username:password[reverse:password]@]
[<servername>]:[<protocol>]:[<port>]:
[<LUN>]:<targetname> (rfc4173)
- Iscsistart from iscsi-initiator-utils
(<http://www.open-iscsi.org>)
- SYMLINK=="disk/by-path/*-iscsi-*-<iscsi lun>"
- Problem without config files

FCoE

- `fcoe=<networkdevice>:<dcb|nodcb>`
`fcoe=<macaddress>:<dcb|nodcb>`
- `ACTION=="add", SUBSYSTEM=="net",
NAME=="<ifacename>", RUN+="/sbin/fcoe-up
$env{INTERFACE}"`
- `/sbin/ip link set "$netif" up`
- `echo -n "$netif" >
/sys/module/fcoe/parameters/create`

Plymouth

- Use `/usr/libexec/plymouth/plymouth-populate-initrd`
- Run a small `udev` trigger for graphic devices only
- Provides dialog for `/sbin/cryptsetup luksOpen`

Cryptoluks

- `ENV{ID_FS_TYPE}=="crypto_LUKS"`
- Runs in non-settled initqueue

DMSquash

- Used to find the Live CD
- Loop mounts images
- Sets up device mapper overlays

Debug

- rdinfo
- rdshell
- rdinitdebug
- rdbreak
- rdbreak={cmdline|pre-udev|pre-trigger|initqueue|pre-mount|mount|pre-pivot}
- rdudevinfo
- rdudevdebug
- rdnetdebug

Test Suite

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

Use Cases

- Cluster setups
- Rescue System
- Kexec crash dumps

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`