[RFC][PATCH] inotify 0.8

Source: http://linux.derkeiler.com/Mailing−Lists/Kernel/2004−07/6040.html

---

From: John McCutchan (*ttb_at_tentacle.dhs.org*)
Date: 07/30/04

```
Date:   Thu, 29 Jul 2004 18:31:31 −0400
To: linux−kernel@vger.kernel.org
```

I am resubmitting inotify for comments and review. Inotify has changed drastically from the earlier proposal that Al Viro did not approve of. There is no longer any use of (device number, inode number) pairs. Please give this version of inotify a fresh view.

Inotify is designed as a replacement for dnotify. The key difference's are that inotify does not require the file to be opened to watch it, when you are watching something with inotify it can go away (if path is unmounted) and you will be sent an event telling you it is gone and events are delivered over a fd not by using signals.

Inotify is a character device that when opened offers 2 IOCTL's. (It actually has 4 but the other 2 are used for debugging)

INOTIFY_WATCH:

> Which takes a path and event mask and returns a unique (to the instance of the driver) integer (wd [watcher descriptor] from here on) that is a 1:1 mapping to the path passed. What happens is inotify gets the inode (and ref's the inode) for the path and adds a inotify_watcher structure to the inodes list of watchers. If this instance of the driver is already watching the path, the event mask will be updated and the original wd will be returned.

INOTIFY_IGNORE:

> Which takes an integer (that you got from INOTIFY_WATCH) representing a wd that you are not interested in watching anymore. This will:
>
> send an IGNORE event to the device
> remove the inotify_watcher structure from the device and from the inode and unref the inode.

After you are watching 1 or more paths, you can read from the fd and get events. The events are struct inotify_event. If you are

watching a directory and something happens to a file in the directory
the event will contain the filename (just the filename not the full path).

KERNEL VFS CHANGES START HERE

Aside from the inotify character device driver.
The changes to the kernel are very minor.

The first change is adding calls to inotify_inode_queue_event and
inotify_dentry_parent_queue_event from the various vfs functions. This
is identical to dnotify.

The second change is more serious, it adds a call to inotify_super_block_umount
inside generic_shutdown_superblock. What inotify_super_block_umount does
is:

find all of the inodes that are on the super block being shut down,
sends each watcher on each inode the UNMOUNT and IGNORED event
removes the watcher structures from each instance of the device driver
and each inode.
unref's the inode.

I have tested this code on my system for over a week now and have not
had problems. I would appreciate design review, code review and testing.

Attached is the patch to Linux 2.6.7

John

–

- text/plain attachment: inotify.patch