

OCFS2 with openSUSE 12.3

I know the cluster-file system OCFS2 from my workplace. With this file system, multiple nodes (in the office: more than 20) can access the same volume. And when this is possible with SLES 11, it should work with openSUSE 12.3 as well.

That's what I thought until I tried.

Installation

That was the simple part:

```
zypper se ocfs2
zypper in ocfs2-tools-o2cb ocfs2console
```

and all packages are installed, together with the other packages that are required to work. The packages come from the openSUSE distro or from the OSS-update repository, respectively.

Create a file system

I used iSCSI to access a volume on my NAS. I created a partition on the volume and initialized the file system:

```
mkfs.ocfs2 -L wernasocfs2 -T mail --fs-feature-level=default --global-
heartbeat \
  --cluster-stack o2cb --cluster-name OCFS2WERNAS -F /dev/sdc1

mkfs.ocfs2 1.8.2
Cluster stack: o2cb
Cluster name: OCFS2WERNAS
Stack Flags: 0x1
NOTE: Feature extended slot map may be enabled
Overwriting existing ocfs2 partition.
WARNING: Cluster check disabled.
Proceed (y/N): y
Filesystem Type of mail
Label: wernasocfs2
Features: sparse extended-slotmap backup-super unwritten inline-data strict-
journal-super xattr indexed-dirs refcount discontig-bg
Block size: 4096 (12 bits)
Cluster size: 4096 (12 bits)
Volume size: 1610611687424 (393215744 clusters) (393215744 blocks)
Cluster groups: 12191 (tail covers 15104 clusters, rest cover 32256
clusters)
Extent allocator size: 100663296 (24 groups)
```

```
Journal size: 268435456
Node slots: 16
Creating bitmaps: done
Initializing superblock: done
Writing system files: done
Writing superblock: done
Writing backup superblock: 6 block(s)
Formatting Journals: done
Growing extent allocator: done
Formatting slot map: done
Formatting quota files: done
Writing lost+found: done
mkfs.ocfs2 successful
```

Configuration

This was a lot more work to do than on my SLES 11 boxes. There, I created der /etc/ocfs2/cluster.conf manually, entered "rco2cb start", "rcocfs2 start" and everything was fine. But here...

```
o2cb list-nodes OCFS2WERNAS
o2cb: Unknown code ____ 251 while loading cluster configuration file
'/etc/ocfs2/cluster.conf'
```

Not a problem, let's use the graphical tool to create a working cluster.conf:

```
# ocfs2console
Traceback (most recent call last):
  File "/usr/sbin/ocfs2console", line 33, in <module>
    from ocfs2interface.console import main
  File "/usr/lib64/python2.7/site-packages/ocfs2interface/console.py", line
26, in <module>
    from mount import mount, unmount
  File "/usr/lib64/python2.7/site-packages/ocfs2interface/mount.py", line
20, in <module>
    import ocfs2
ImportError: /usr/lib64/python2.7/site-
packages/ocfs2interface/ocfs2module.so: undefined symbol: io_queue_release
```

Hum, with SLES 11 I never had those problems. Why here? So I opened a bug for ocfs2console: https://bugzilla.novell.com/show_bug.cgi?id=827504

So, everything has to be entered manually:

```
klops:~ # o2cb add-cluster OCFS2WERNAS
klops:~ # o2cb heartbeat-mode OCFS2WERNAS global
klops:~ # o2cb add-heartbeat OCFS2WERNAS /dev/disk/by-label/wernasocfs2
klops:~ # o2cb start-heartbeat OCFS2WERNAS
Global heartbeat started
klops:~ # o2cb add-node --ip 192.168.2.55 --number 0 OCFS2WERNAS klops
```

```
klops:~ # o2cb add-node --ip 192.168.2.77 --number 1 OCFS2WERNAS sieben
klops:~ # o2cb add-node --ip 192.168.2.44 --number 2 OCFS2WERNAS thinkt51
klops:~ # o2cb list-cluster OCFS2WERNAS
heartbeat:
    region = 5F1672E500D94138BF3C4F517813A336
    cluster = OCFS2WERNAS

node:
    number = 0
    name = klops
    ip_address = 192.168.2.55
    ip_port = 7777
    cluster = OCFS2WERNAS

cluster:
    node_count = 3
    heartbeat_mode = global
    name = OCFS2WERNAS
```

Pardon? Oh, really, in the file `/etc/ocfs2/cluster.conf` there is only one node. The “node count” was increased correctly. Four weeks ago, the bugzilla entry https://bugzilla.novell.com/show_bug.cgi?id=822190 was created, but is still showing “new”. I added https://bugzilla.novell.com/show_bug.cgi?id=827479, maybe this helps.

As an alternative there is `o2cb_ctl`, but:

```
klops:~ # o2cb_ctl -C -n sieben -t node -a number=1 -a
ip_address=192.168.2.77 -a ip_port=7777 -a cluster=OCFS2WERNAS
o2cb_ctl: Unable to access cluster service Cannot initialize cluster
```

OK, after starting `o2cb` gestartet wurde, the command worked – but again only the “node count” is increased, no additional node entry in file `/etc/ocfs2/cluster.conf`.

The commands I used:

```
klops:~ # o2cb heartbeat-mode OCFS2WERNAS global
klops:~ # o2cb list-heartbeats OCFS2WERNAS
heartbeat:
    region = 5F1672E500D94138BF3C4F517813A336
    cluster = OCFS2WERNAS

klops:~ # o2cb list-cluster OCFS2WERNAS
heartbeat:
    region = 5F1672E500D94138BF3C4F517813A336
    cluster = OCFS2WERNAS

cluster:
    node_count = 0
    heartbeat_mode = global
    name = OCFS2WERNAS
```

```
klops:~ # o2cb add-node --ip 192.168.2.55 --number 0 OCFS2WERNAS klops
klops:~ # o2cb list-cluster OCFS2WERNAS
heartbeat:
  region = 5F1672E500D94138BF3C4F517813A336
  cluster = OCFS2WERNAS
```

```
node:
  number = 0
  name = klops
  ip_address = 192.168.2.55
  ip_port = 7777
  cluster = OCFS2WERNAS
```

```
cluster:
  node_count = 1
  heartbeat_mode = global
  name = OCFS2WERNAS
```

```
klops:~ # o2cb add-node --ip 192.168.2.77 --number 1 OCFS2WERNAS sieben
klops:~ # o2cb list-cluster OCFS2WERNAS
heartbeat:
  region = 5F1672E500D94138BF3C4F517813A336
  cluster = OCFS2WERNAS
```

```
node:
  number = 0
  name = klops
  ip_address = 192.168.2.55
  ip_port = 7777
  cluster = OCFS2WERNAS
```

```
cluster:
  node_count = 2
  heartbeat_mode = global
  name = OCFS2WERNAS
```

```
klops:~ # o2cb add-node --ip 192.168.2.44 --number 2 OCFS2WERNAS thinkt51
klops:~ # o2cb list-cluster OCFS2WERNAS
heartbeat:
  region = 5F1672E500D94138BF3C4F517813A336
  cluster = OCFS2WERNAS
```

```
node:
  number = 0
  name = klops
  ip_address = 192.168.2.55
  ip_port = 7777
  cluster = OCFS2WERNAS
```

```
cluster:
  node_count = 3
```

```
heartbeat_mode = global  
name = OCFS2WERNAS
```

Mounting

Well, sometimes I can mount the volume:

```
# mount /dev/disk/by-label/wernasocfs2 /zdisk/ocfs2  
# mount | grep ocfs2  
ocfs2_dlmfs on /dlm type ocfs2_dlmfs (rw,relatime)  
/dev/sdc1 on /zdisk/ocfs2 type ocfs2  
(rw,relatime,_netdev,heartbeat=global,nointr,data=ordered,errors=remount-  
ro,atime_quantum=60,cluster_stack=o2cb,coherency=full,user_xattr,acl)
```

But then again:

```
klops:~ # mount /zdisk/ocfs2/  
mount.ocfs2: Invalid argument while mounting /dev/sdc1 on /zdisk/ocfs2.  
Check 'dmesg' for more information on this error.
```

And dmesg tells me:

```
[ 2907.104719] OCFS2 1.5.0  
[ 2907.147382] (mount.ocfs2,8078,0):ocfs2_verify_userspace_stack:876 ERROR:  
cluster stack passed to mount, but this filesystem does not support it  
[ 2907.147440] ocfs2: Unmounting device (8,33) on (node 0)  
[ 2907.147457] (mount.ocfs2,8078,0):ocfs2_fill_super:1230 ERROR: status =  
-22
```

What's going on here?

From:
<http://www.wernerflamme.net/> - **Werners Wiki**

Permanent link:
<http://www.wernerflamme.net/doku.php?id=comp:en:comp:ocfs2os123>

Last update: **2013-06-29 2233**

