



# Debugging GPFS

---

- Got problem? Don't panic!
  - Check for possible basic problems:
    - Is Network OK?
    - Check status of the cluster: "mmgetstate -a"
    - Check status of NSDs: "mmlsdisk fsname"
  - Take a 5 min break (take a coffee, smoke a cigarette)
    - In major cases GPFS will recover by it self without need of any intervention from the administrator
  - If not recovered
    - **Ensure that you are the only person who is doing the work!**
    - check gpfs logs (first on cluster manager, then on FS manager, then on NSD servers)
    - check syslog (/var/log/messages) for eventual errors
    - Check disks availability (mmlsdisk fsname)
    - Consult "Problem determination guide"



# Debugging GPFS (cont.)

---

- Some useful commands:
  - “mmfsadm dump waiters” will help to find long lasting processes
  - “mmdiag --network|grep pending” helps to individuate non-responsive node
  - “mmdiag --iohist” lists last 512 I/O operations performed by GPFS on current node (helps to find malfunctioning disk)
  - “gpfs.snap” will gather all logs and configurations from all nodes in the cluster
    - the first thing to send to IBM support when opening service request



# GPFS V3.4 Problem Determination Guide

---

- <http://publib.boulder.ibm.com/infocenter/clresctr/vrx/index.jsp>



# NFS stale file handle

---

- When a GPFS mount point is in the "NFS stale file handle" status, example

```
[root@diskserve-cms-2 root]# df
Filesystem 1K-blocks Used Available Use% Mounted on
/dev/gpfs_cms4 8125032448 8023801088 101231360 99% /storag/gpfs_cm
df: ` /storage/gpfs_cms5': Stale NFS file handle
```

- check if there is any NSD with status "down"

```
[root@diskserve-cms-2 root]# mmlsdisk gpfs_cms5
disk driver sector failure holds holds
name type size group metadata data status availability
-----
disk53 nsd 512 4015 yes yes ready up
disk54 nsd 512 4015 yes yes ready down
disk58 nsd 512 4015 yes yes ready down
disk59 nsd 512 4013 yes yes ready up
```

restart the NSDs (important: do it for all NSD with status "down" in one command):

```
[root@diskserve-cms-2 root]# mmchdisk gpfs_cms5 start -d "disk54;disk58"
```

re-mount the filesystem



## Recovery of GPFS configuration

---

- If a node of the cluster lost its configuration (has been re-installed) but still present as member of this cluster ("mmgetstate" lists it in "unknown" state) use this command to recover the node:

```
/usr/lpp/mmfs/bin/mmsdrrestore -p  
diskerv-san-5 -R /usr/bin/scp
```



# Checking existing NSD

---

- If get this warning while creating new nsd

Disk descriptor xxx system refers to an existing NSD

Use this command to verify if this device is actually used in one of the file systems

```
mmfsadm test readdescraw /dev/emcpowerax
```