

SGI UNICOS DCE/DFS Administrative Aids

The Distributed Computing Environment (DCE) and Distributed File Service (DFS) are very powerful products, encompassing many different distributed computing functions. Because of this complexity, it may be hard for a site analyst to know where to start when isolating a suspected DCE or DFS problem or looking for usage information. The following document is a list of basic information which may be of use in this area. It consists of 4 sections; General Information, Quick and Easy Sanity Checks and Information Resources, Specific Problems/Resolutions, and Miscellaneous Information.

The intended audience is SGI analysts and system administrators knowledgeable with UNICOS, but less experienced with DCE and DFS. This document was written for the DCE/DFS product supported on the UNICOS operating system, but in most cases the information is also relevant for the UNICOS/Mk operating system.

Table of Contents

1. General Information	2
1.1 The DCE/DFS daemons	2
1.2 DFS Kernel Threads	2
1.3 SGI's DCE/DFS World Wide Web location	3
1.4 Other World Wide Web pages of interest	3
1.5 DFS Tuning Options section in the Release Overview	3
1.6 Tuning options documented in man pages	3
2. Quick and Easy Sanity Checks and Information Resources	4
2.1 Sanity checks which indicate DCE core services are working	4
2.2 Verifying that the session has DCE credentials	4
2.3 The dcestat command	5
2.4 Logs	5
2.5 UNICOS crash utility dfsstat command	6
2.6 mbuf usage	6
2.7 The dce_shutdown script	6
3. Specific Problems/Resolutions	7
3.1 System console is receiving request denial messages from m_gets	7
3.2 The DFS client is working, but not the server	8
3.3 User receives an "I/O error" when doing a change directory (cd)	8
3.4 User receives a "Connection timed out" message	9
3.5 User receives a "No such device" message	11
3.6 User receives a "permission denied" error message	12
3.7 User receives a "Unable to set context: internal error in sec_login (dce / sec)"	13
3.8 User receives "permission denied" when accessing DFS file system with MLS	14
3.9 Unable get to the junction after CDS (Cell Directory Server) is restored	18
4. Miscellaneous Information	19
4.1 Requirements for a Cray only cell	19
4.2 DFS use with NQE	19
4.3 Changing your dce login password	19

1. General Information

1.1 The DCE/DFS daemons.

- dced(8) - The DCE Host daemon, who's services include -
- > Endpoint Mapper: maintains a database which allows DCE clients to find servers, individual services provided by servers, and objects managed by services on the host.
 - > Host Data Management: maintains local files of host data that include the host_name, cell_name, cell_aliases, and a post_processor file.
 - > Server Management: maintains data that describes the startup configuration and execution state for each server, as well as the ability to start or stop particular servers.
 - > Security Validation: acts as the client side of the security server.
 - > Key Table Management: allows for remote maintenance of server's key tables.
- cdsadv(8) - Cell Directory Service advertisement and solicitation daemon.
- dfsbind(8) - Provides user-space information to the Cache Manager and File Exporter. (DFS Client and Server)
- ftserver(8) - Fileset Server, handles file administration operations. (DFS Server)

Further information on all of the above daemons may be obtained by viewing the associated man pages distributed with the UNICOS DCE/DFS release package.

1.2 DFS Kernel Threads.

The following threads may or may not be present depending if the system is running DFS Client, DFS server or both. Also if viewing via the "ps" command the number of threads of each type will vary depending on system activity, configuration options, etc.

	<u>Thread Name</u>	<u>Used By DFS</u>	<u>Function</u>
DCEspawn		Client/ Server	The parent thread of all other kernel threads. Created during system startup and always present if DFS is built into the kernel.
DCEcthrd	Call Threads	Client/ Server	Executes the server end of RPC calls.
DCEdfsd	Cache manager	Client	Make RPC calls on behalf of users of the DFS client.

DCEfxd	File Exporter	Server	Performs periodic housekeeping tasks in DFS server.
DCElsn	Network Listener Service	Client/ Server	Reads all incoming RPC network traffic and hands it on to the correct thread.
DCEpool		Client/ Server	Used for timer events, for when an RPC call must be made from the DFS server to multiple clients at once, and for assorted housekeeping tasks.
DCEtimer		Client/ Server	Processes timer events within the RPC runtime library.
DCEconvm	Conversation manager*	Client/ Server	Used internally within RPC runtime code to manage the state of ongoing RPC calls.

* Not always present.

1.3 SGI's DCE/DFS World Wide Web location - (may change without warning)

<http://www.cray.com/products/software/misc/dceindex.html>

provides basic information about SGI's DCE/DFS products, on-line SGI DCE/DFS documentation, including this document, information on problem avoidance and resolutions, future releases, as well as pointers to other DCE/DFS Web Sites.

1.4 Other World Wide Web pages of interest - (may change without warning)

Open Group

<http://www.opengroup.org/tech/dce/>

Transarc Corporation

<http://www.transarc.com/>

1.5 The DFS Tuning Options section of the Cray DCE Client Services 1.1.0/Cray DCE DFS Server 1.1.0 Release Overview, (RO-5225 1.1.0), supplies information on tuning DFS for the type of usage DFS has at your site.

1.6 Options to tune the DFS daemons may be modified in your rc.dfs file or through the dce_config utility. These options are documented in the associated man pages (i.e. fxd, dfsd).

2. Quick and Easy Sanity Checks and Information Resources

2.1 Sanity checks which indicate DCE core services are working:

If the following 2 commands work without error, it is a good indication that the DCE core services are functioning properly.

```
% /opt/dcelocal/bin/dcecp -c directory show /:
% /opt/dcelocal/bin/dce_login {valid username}
```

If either of the above commands indicate a problem, try them from another host in the cell to ensure the server(s) are functioning correctly.

The DCE core services may be restarted by doing a -

```
# /opt/dcelocal/etc/dce_shutdown
Stopping DCE daemons...Stopped all DCE daemons.

# /etc/rc.dce
Starting DCE daemons from rc.dce
/opt/dcelocal/bin/dced -rpcd -f -sec_clientd -v -d -runpipe /opt/dcelocal/etc/CMD
/opt/dcelocal/bin/cdsadv -d -s
DCE daemon start complete: Mon Oct 24 09:54:51 CDT 1994
```

2.2 To verify that the session has DCE credentials -

Verify that there is an environment variable called KRB5CCNAME and it is set to something similar to the following (the file name will vary) -

```
$ echo $KRB5CCNAME
FILE:/opt/dcelocal/var/security/creds/dcecred_41ffffc
```

The “klist” command returns DCE Identity Information -

Note - make sure you are using the correct klist command, /opt/dcelocal/bin/klist, **not** /usr/bin/klist

Valid DCE Credentials (only part of output listed)

```
$ /opt/dcelocal/bin/klist
DCE Identity Information:
Warning: Identity information is not certified
Global Principal: /.../solar.universe.com/galileo
Cell: 00690ba4-1182-1de1-87e4-80a2155aaa77 /.../solar.universe.com
Principal: 000024ad-b28a-2e50-ba00-80a2155aaa77 galileo
Group: 00000408-b314-2de3-b901-80a2155aaa77 astronomer
```

No Valid DCE Credentials

```
$ /opt/dcelocal/bin/klist
No DCE identity available: No currently established network identity for which context exists (dce / sec)
```

Kerberos Ticket Information:

klist: Unknown code DCE:krb 191 (336761023) while setting cache flags (ticket cache /tmp/krb5cc_9389)

2.3 The `dcestat` command can give you an indication which daemons are running. The following is sample output from the `dcestat` command -

% `/opt/dcelocal/bin/dcestat`

DCE CORE SERVICES and DFS DAEMONS

```
-----  
THRDS  SECS   PID  PPID    SIZE   TTY   NAME  
    10  2.92  1470    1    1712   ?    dced  
    26  9.87  1821    1    2092   ?    cdsadv  
    13  0.70  1667    1    1596   ?    dfsbind  
     8  0.64 13465    1    1736   ?    ftserver
```

DFS KERNEL THREADS

```
-----  
THRDS  SECS   PID  PPID    SIZE   TTY   NAME  
     9  0.00  1715    1     8     ?    DCEcthrd  
     8  0.00  1703    1     8     ?    DCEdfsd  
     1  0.00  1954    1     8     ?    DCEfxd  
     1  0.07  1705    1     8     ?    DCElsn  
    16  0.78  1887    1     8     ?    DCEpool  
     1  0.03     8     0     8     ?    DCEspawn  
     1  0.01  1644    1     8     ?    DCEtimer
```

The kernel threads will vary depending on the type of DCE/DFS activity on the system. If you intend to run either DFS Client or Server the DCEspawn thread must be present, even if DFS has not been brought up. The DCEspawn thread indicates that DFS is built into the kernel. If it is not there see the Kernel configuration, build and install section of the Release Overview for DCE/DFS, RO-5225 1.1.0.

Note the column labeled "threads" in the above listing. If you do a "ps -ef" on any of the processes with more than one thread it will appear as several processes, but is actually just multiple threads from one process. If for some reason you need to kill one of the DCE core services or DFS daemons, the `dcestat` command is an easy way to determine which PID to use for the process.

2.4 There are several logs which may be of assistance in debugging DCE/DFS problems -

`/usr/spool/dce.log` -> Configuration, startup, shutdown log.
`/usr/adm/syslog`

There may be logfiles or core dumps in the following directories -

```
${DCELOCAL}/etc/logs/  
ls -lR ${DCELOCAL}/var | grep core
```

- 2.5 The UNICOS crash(8) utility has a `dfsstat` command which will print various cache, DFS, and RPC usage statistics.

```
> dfsstat -?  
usage: dfsstat [-g] [-c] [-p] [-r]  
-g = Grep Format  
-c = Print CM stats  
-p = Print PX stats  
-r = Print RPC stats
```

The above data may be useful for researching problems and for general usage information.

- 2.6 When analyzing problems, one easy thing to check is mbuf usage by executing the “`/usr/ucb/netstat -m`” command. Pay special attention to the first line of output, mbufs in use, as well as the last 3 lines; mbuf queue flushes, protocol mbuf queue drains, and requests for memory denied.


This subject is covered in more detail in the Specific Problems section of this document.

- 2.7 DCE core services can be taken down and brought back using the `dce_shutdown` script in `/opt/dcelocal/etc` and `/etc/rc.dce`. This capability does not exist with DFS. A host can be disabled as a DFS server by detaching all exported files on that system. It can be disabled as a DFS client by unmounting `/...`, but you will not be able to remount `/...` without a reboot.

3. Specific Problems/Resolutions

The following identifiers will be used for examples in this section.

Cell Name - solar.universe.com
Hosts - sun moon
Exported file systems - /earth /mars
DCE FS Directory structure -



3.1 Symptoms:

-> System console is receiving request denial messages from m_gets:

```
m_gets: WARNING: %d denials for %d byte requests
```

-> The “netstat -m” command shows flushes, drains, or denials.

-> Overall sluggish performance.

-> Hung interactive sessions, including dce_login.

-> System does not respond to a ping.

Possible Reason:

System may be low or out of allocatable mbufs. Check “netstat -m” on the running system or system dump. Below is an example of netstat output. Note the high number of flushes, drains, and requests denied.

netstat -m ./unicos ./dump

14874/15000 mbufs in use (15000 maximum), allocated as 8642 mbuf clusters:

```
428 mbufs allocated to data (1350 maximum)
1 mbufs allocated to packet headers (650 maximum)
388 mbufs allocated to socket structures (486 maximum)
376 mbufs allocated to protocol control blocks (484 maximum)
27 mbufs allocated to routing table entries (28 maximum)
0 mbufs allocated to fragment reassembly queue headers (14 maximum)
6 mbufs allocated to address resolution tables (6 maximum)
5 mbufs allocated to socket names and addresses (86 maximum)
0 mbufs allocated to socket options (4 maximum)
5 mbufs allocated to interface addresses (5 maximum)
1 mbufs allocated to NFS notify-when-free data areas (12 maximum)
4594 mbufs allocated to NFS static data (4867 maximum)
38 mbufs allocated to NFS dynamic data (107 maximum)
13 mbufs allocated to system buffer headers (48 maximum)
16 mbufs allocated to interface structures (16 maximum)
0 mbufs allocated to ancillary data (14 maximum)
```

63 mbufs allocated to network security data (69 maximum)
1 mbufs allocated to network security CIPSO map data (5 maximum)
8848 mbufs allocated to DCE dynamic memory allocation (8871 maximum)
54 mbufs allocated to path buffer used to get vnodes (54 maximum)
4 mbufs allocated to allocated for credentials structure (8 maximum)
2 mbufs allocated to internet multicast options (3 maximum)
4 mbufs allocated to internet multicast address (4 maximum)
856368 mbuf queue flushes
856086 protocol mbuf queue drains
14616 requests for memory denied

Resolution:

Allocate more mbufs. See the Altering kernel definitions for DFS section in the Installation and Configuration chapter of the DCE/DFS Release Overview, for further information.

3.2 **Symptom:** The DFS client is working, but not the server.

Possible Reason 1:

There may be a problem with the ftserver(8) daemon.

Resolution 1:

```
# kill {ftserver process} (-4 for a dump, -9 no dump)

# cd /opt/dcelocal/etc/cores/ftserver
# nohup /opt/dcelocal/bin/ftserver >> /opt/dcelocal/etc/logs/ftserver 2>&1 &
```

The dump will be stored in /opt/dcelocal/etc/cores/ftserver.

Possible Reason 2:

The process fxd(8) may not have completed startup successfully.

Resolution 2:

The “ps -ef | egrep fxd” should only show DCEfxd. If the process fxd is still present it is either hung or waiting for something. Check the logs documented in note 2.4, especially in the directory /opt/dcelocal/etc/logs/fxd to attempt to determine the problem.

3.3 **Symptom:** User received an “I/O error” when doing a change directory (cd).

```
$ cd /:/milkyway/sun/earth
chdir: /:/milkyway/sun/earth: I/O error
```

Possible Reason:

There may be a problem with dfsbind, the process which provides user-space services to the Cache Manager on a DFS client machine or the File Exporter on a DFS File Server machine. Use the dcestat utility to assure dfsbind is still running. Check logs for possibly problems, especially the /opt/dcelocal/etc/logs/dfsbind log.

Resolution:

Kill the dfsbind daemon (if still up, but not functioning correctly) and restart. To bring dfsbind back up -

```
# cd /opt/dcelocal/etc/cores/dfsbind
# nohup /opt/dcelocal/bin/dfsbind >> /opt/dcelocal/etc/logs/dfsbind 2>&1 &
```

3.4 **Symptom:** User gets a 'Connection timed out' message when attempting to access a file or directory within DFS.

Possible Reason 1:

The DFS cache manager may have failed to completely start up. Check for dfds(8) processes running on the system. When start-up is complete there will only be DCEdfsd processes present, the process named dfds will no longer be in the system. The logfile /opt/dcelocal/etc/logs/dfds may give a reason why it did not complete successfully, one of which would be a busy mount point. The following is the output from a ps performed on a system which is waiting on a busy mount point.

```
% ps -ef | egrep dfs
root 2256 1 ? 0:00 DCEdfsd
root 2266 1 ? 0:00 DCEdfsd
root 2268 1 ? 0:00 DCEdfsd
root 2270 1 ? 0:00 DCEdfsd
root 2352 1 p002 0:00 dfds
```

Resolution 1:

Finding the user causing a busy mount point -

Output from the /opt/dcelocal/etc/logs/dfds

```
--- Wed Oct 19 09:59:33 CDT 1994 ---
dfds: start sweeping disk cache files ....
dfds: All DFS daemons started.
dfds: can't mount DFS on /... (code 16). Retrying
```

a) Find the inode number of the mount point.

```
# ls -ld /...
5376 /...
```

b) In crash, find the inode slot number of the above inode.

```
# crash
> nc1inode | egrep 5376
325 34 9 5376 0 3 2 0 0 4096 d---555 - -
```

c) Redirect crash's output to a file and enter the "user" command.

```
> redirect crash.output
> user
```

d) From another screen - view the output from crash and search for the cdir entry with the above inode slot number. This will give you the process slot number and uid of the process which has the mount point.

```
# vi crash.output
```

```
<325>
```

```
-----
```

```
Process 9 User Area:
```

```
User ID's:   uid: 9389, gid: -1, real uid: 9389
```

```
Process Times: user: 1, sys: 0, child user: 12, child sys: 10
```

```
Process Misc:  proc slot: 9 pcomm addr: 0240553
```

```
Size:       uc_tsize: 62456, uc_dsize: 64178
```

```
IPC:       locks: unlocked
```

```
Directories: <file system> <dependent inode slot>
```

```
rdir: NULL pdir: NULL cdir: <nc1fs><325>
```

```
-----
```

e) Back in crash, you can find the process id by viewing the above process slot number.

```
> proc 9
```

```
SLT ST PID PPID PGRP UID EUID PRI CPU EVENT NAME FLAGS
 9 s 2280  1 2280 9389 9389 39 - 00000561 csh Ld rwt oldmask
```

Kill this process, or get this user to cd out of the /... directory. If you still can not access /..., and it does not appear in the output from /etc/mount, repeat the above steps. Someone else may now have the inode. /... will mount automatically once the inode is released.

Possible Reason 2:

The DFS file server is unreachable. Try the '/opt/dcelocal/bin/cm statservers' command on the client, which will attempt to access all servers. This could return immediately if the server is back up, or could take a couple of minutes before reporting which servers are down.

```
% cd ../../solar.universe.com/fs/milkyway/sun/earth/mydir
```

```
../../solar.universe.com/fs/milkyway/sun/earth/mydir: Connection timed out
```

```
% cm statservers
```

```
These servers are still down: sun.
```

Resolution 2:

a) Perform the following steps to check the network route -

-> Find the fileset associated with the mount point of the file or directory you are trying to access.

```
# fts lsmount -dir /:/milkyway/sun/earth
```

```
:/:/milkyway/sun/earth' is a mount point for fileset '#aggr.earth.disk0'
```

-> The fldb entry of the above fileset will give the name of the server.

```
# fts lsfldb -fileset aggr.earth.disk0
```

```
readWrite ID 0,,1249 valid
readOnly ID 0,,1250 invalid
backup ID 0,,1251 invalid
number of sites: 1
server flags aggr siteAge principal owner
sun-fca1 RW /earth 0:00:00 hosts/sun <nil>
```

-> Ping the server to verify the network routes.

```
# ping sun-fca1
```

- b) Go through the sanity checks, (described in Section II), on the server. This includes checking core services, logs, mbufs, etc.
- c) Check dfsbind as described in Problem/Resolution #3.3 of this document.

Possible Reason 3:

Default Machine Context has been lost. Site should see a console or log message of “dfs: ticket has expired; running unauthenticated”. See Symptom/Resolution #3.6 which is a different manifestation of the same problem.

3.5 **Symptom:** When doing a cd to a DFS mount point you receive the message - “No such device”.

```
$ cd /:/milkyway/sun/earth
chdir: /:/milkyway/sun/earth: No such device
```

Possible Reason:

You might be attempting to traverse a DFS mount point which is no longer valid. There are several DFS related components which must be in sync for DFS to access exported filesets.

The command “fts crmount” creates a DFS mount point for a fileset. This command is either entered manually as part of the fileset export procedure or automatically when using Cray’s dfsmkfs(8) command. The “fts lsmount” command displays the fileset associated with the mount point.

```
# fts lsmount -dir /:/milkyway/sun/earth
'/:/milkyway/sun/earth' is a mount point for fileset '#aggr.earth.disk0'
```

The fileset must be successfully exported, accomplished through the dfsexport(8) command issued on the DFS file server. The dfsexport command reads the /opt/dcelocal/var/dfs/dfstab file for exporting information. The dfstab(4) file information may be overridden by command line arguments. See the associated man page for further information. Issuing the dfsexport command on the DFS file server with no options lists the filesets currently exported from the local disk to the DCE namespace.

```
# dfsexport
dfsexport: /dev/dsk/disk0, ufs, 1, 0,,1321
```

The fileset location database (fldb) is used to locate the server which exports the fileset. Use the “fts lsfldb -fileset {fileset}” command with the fileset identified from the output of the lsmount

performed above to identify the server and display information about the fileset in question. Verify that the fileset ID's shown in the output from the dfsexport and the lsflldb match.

```
# fts lsflldb -fileset aggr.earth.disk0
```

```
aggr.earth.disk0
  readWrite ID 0,,1321 valid
  readOnly  ID 0,,1322 invalid
  backup    ID 0,,1323 invalid
number of sites: 1
  server    flags  aggr  siteAge principal  owner
sun-fca1   RW    /earth  0:00:00 hosts/sun <nil>
```

Resolution:

Fix the component which is out of sync. This might include doing a dfsrmfs(8) to clean up the flldb, mount point, etc., then a dfskmfs(8) to set everything up again, or reissuing the dfsexport, possibly doing a "dfsexport -detach" on the fileset first. See the associated man pages for command line options.

3.6 Symptom:

Client access to other DFS server filesystems produces a "permission denied" error message. Files created in these filesystems end up owned by nobody instead of the correct DCE identity. The user identity seems to be ignored and access seems as if the user is running unauthenticated, even though they have done a successful dce_login.

Possible Reason:

Default Machine Context has been lost. This can happen when sec_clientd attempts to refresh its context. To determine if this is the problem you can login in as root (do **not** do a dce_login) and view credentials:

```
# /opt/dcelocal/bin/klist
DCE Identity Information:
  Global Principal: /.../solar.universe.com/hosts/moon/self
  Cell: 00690ba4-1182-1de1-87e4-80a2155aaa77 /.../solar.universe.com
  Principal: 00001a49-e506-2e92-a500-80a2155aaa77 hosts/moon/self
  Group: 0000000c-1183-2de1-8701-80a2155aaa77 <group name unknown>
Local Groups:
  0000000c-1183-2de1-8701-80a2155aaa77 <group name unknown>

Identity Info Expires: 94/11/10:02:29:26
Account Expires: never
Passwd Expires: never

Kerberos Ticket Information:
Ticket cache: /opt/dcelocal/var/security/creds/dcecred_fffffffffffffff
Default principal: hosts/moon/self@solar.universe.com
```

Notice that groups is listed as **unknown**, and the Kerberos Ticket Information is incomplete. You can further check the dced log file, (/opt/dcelocal/etc/logs/dced), and look for error messages such as:

```
sec_clientd 11/02/94 04:51:13 - Refreshing machine context
```

```
sec_clientd 11/02/94 04:51:13 - Warning - unable to refresh identity ... will idle and retry  
... Context has not been validated (dce / sec)
```

Resolution:

Remove old credential files for principal **self**:

```
# rm /opt/dcelocal/var/security/creds/dccred_fffffffffffffff*
```

Kill dced and restart dced.

```
# /opt/dcelocal/bin/dcestat
```

DCE CORE SERVICES and DFS DAEMONS

```
-----  
THRDS  SECS   PID  PPID    SIZE   TTY   NAME  
    10  2.92  1470    1    1712   ?    dced  
    26  9.87  1821    1    2092   ?    cdsadv  
    13  0.70  1667    1    1596   ?    dfsbind  
     8  0.64  13465   1    1736   ?    ftserver
```

DFS KERNEL THREADS

```
-----  
THRDS  SECS   PID  PPID    SIZE   TTY   NAME  
     9  0.00  1715    1     8     ?    DCEcthrd  
     8  0.00  1703    1     8     ?    DCEdfsd  
     1  0.00  1954    1     8     ?    DCEfxd  
     1  0.07  1705    1     8     ?    DCElsn  
    16  0.78  1887    1     8     ?    DCEpool  
     1  0.03     8     0     8     ?    DCEspawn  
     1  0.01  1644    1     8     ?    DCEtimer
```

```
#  
# kill 1470  
# cd /opt/dcelocal/etc/cores/dced  
# nohup /opt/dcelocal/bin/dced -f >> /opt/dcelocal/etc/logs/dced 2>&1 &
```

3.7 **Symptom:** When doing a dce login you receive the following error -

```
$ dce_login galileo  
Unable to set context: internal error in sec_login (dce / sec)
```

Possible Reason:

The filesystem where the creds directory resides may be out of space.

```
$ pwd  
/opt/dcelocal/var/security/creds  
$ ls -ltr
```

```
total 320
-rw----- 1 root          bin      497      Dec 12 01:02 dcecred_fffffffffffffff.data
-rw----- 1 root          bin     7072      Dec 12 01:02 dcecred_fffffffffffffff
-rw-r--r-- 1 root          bin     106      Dec 12 01:02 dcecred_fffffffffffffff.nc
-rw----- 1 root          bin     574      Dec 12 01:02 dcecred_41ffffff.data
-rw----- 2 root          bin    4958      Dec 12 01:02 dcecred_07ad7a01
-rw-r--r-- 1 root          bin     220      Dec 12 01:02 dcecred_07ad7a01.nc
-rw----- 2 root          bin    4958      Dec 12 01:02 dcecred_41ffffff
-rw-r--r-- 1 root          bin     106      Dec 12 01:02 dcecred_41ffffff.nc
-rw----- 1 root          bin    1363      Dec 12 08:02 dcecred_06d67209
-rw----- 1 galileo       bin    1259      Dec 12 09:25 dcecred_a4267a01
-rw----- 1 galileo       bin      0      Dec 12 09:25 dcecred_a4267a02
```

Note that dcecred_a4267a02 has a size of zero. At the next attempt to login you get the following response -

```
$ dce_login galileo
Sorry.
User Identification Failure. - Credentials cache I/O operation failed XXX (dce / krb)
```

Resolution:

Need to free up space on the above device. Do not do this by deleting cred files.

3.8 Symptom: User receives "permission denied" when accessing DFS file system with MLS.

Possible Reason 1:

User's process is not at the correct label.

```
# /bin/spget

permits equal 00
  none
security level is 10
  level10          <----- Current security level
maximum level is 16
  level16
minimum level is 0
  level0
authorized compartments are 0400000004000040000000
  comp24
  comp39
  comp63
active compartments are 00
  none
integrity class is 0
  class0
maximum class is 0
  class0
active categories are 00
  none
authorized categories are 07000000
  secadm
  sysadm
  sysops
```

Possible Reason 2:

NAL may be configured incorrectly. Assure the level the user is running at is within the NAL label range on both systems.

```
# /etc/spnet resolve -nal hosta
ip host "hosta" {
    name = "%NAC";
    class = B1;
    min label = level10;
    max label = level10;
    ipso = cipso;
    domain = 1;
}

# /etc/spnet resolve -nal hostb
ip host "hostb" {
    name = "%NAA";
    class = B1;
    min label = level10;
    max label = level10;
    ipso = cipso;
    domain = 1;
}
```

Possible Reason 3:

Does the DFS server machine have NFS id map entries for the client machine and the user attempting to access the file system? Are their labels correct in the map and in the UDB?

```
# /etc/uidmaps/nfsuid -m hosta hosta galileo
Local name of user   : hosta
Kernel User ID map entry address: 0000116210245, 20517029, 0x0013910a5
Account ID          : 8383
Minimum Security Level : 0
Maximum Security Level : 16
Minimum Compartments : 0x0
Maximum Compartments : 0x4000004000800000
No Validated Kerberos Addresses.
No Client AUTH_KERB validations.
No Server AUTH_KERB validations.
User Entry type 1   Max 2 groups in groups list
Local uid   : 7094 Remote uid   : 7094
Local gid(0) : 1015 Remote gid(0) : 1015
Local gid(1) : -1 Remote gid(1) : -1

Local name of user   : galileo
Kernel User ID map entry address: 0000116216333, 20520155, 0x001391cdb
Account ID          : 9389
Minimum Security Level : 0
Maximum Security Level : 16
Minimum Compartments : 0x0
Maximum Compartments : 0x4000004000800000
No Validated Kerberos Addresses.
No Client AUTH_KERB validations.
No Server AUTH_KERB validations.
User Entry type 4   Max 5 groups in groups list
Local uid   : 24253 Remote uid   : 24253
Local gid(0) : 11500 Remote gid(0) : 1015
Local gid(1) : 1007 Remote gid(1) : 168
Local gid(2) : 1015 Remote gid(2) : 1007
Local gid(3) : 1026 Remote gid(3) : 12584
Local gid(4) : 11121 Remote gid(4) : -1
```

Possible Reason 4:

Are the packets going out from the client machine correctly labelled?

Step 1 - Assure you can get to the DFS junction -

```
$ cd /:/milkyway/sun/earth
```

Step 2 - In one window start the UNICOS utility /etc/crash. Type in the ccall command, don't hit carriage return:

```
# /etc/crash  
> ccall <---- No carriage control at this time
```

Step 3 - In the user window create DFS traffic from the process in question. Immediately in the crash window hit carriage return to run ccall this ccall shows the label of packets going to the server. You should see an RPC_DG_CCALL_T entry at the label the process is at.

User Window

```
$ ls -le  
cmd-1025 ls: ./file_label4: Permission denied <----- User's current label is not allowed to see label 4  
total 72t  
drwxr-xr-x    10  2 galileo  4294967294  4096      Dec 13 17:00 dir1  
-rw-----    10  1 galileo  4294967294   910      Jan  8 13:48 file1  
drwxr-xr-x    10  6 galileo  4294967294  4096      Feb  8 18:09 dir2  
-rw-r--r--    10  1 galileo  4294967294   50      Dec 11 15:22 file2
```

Crash Window

```
> ccall  
RPC_DG_CCALL_T AT ADDR=021507604  
  fault_rqe=0  
  reject_status=0 cbk_start=0 response_info_updated=1  
  server_bound=1 cbk_scall=0  
  ccte_ref.ccte=025651004 ccte_ref.gc_count=6  
  h=025503004 timeout_stamp=0101471  
  ping.next_time=073613 ping.start_time=0  
  ping.count=0 ping.pinging=0  
  quit.next_time=0 quit.quack_rcvd=0  
  cancel.next_time=0 cancel.local_count=0  
  cancel.server_count=0 cancel.timeout_time=0  
  cancel.server_is_accepting=01 cancel.server_had_pending=0  
RPC_DG_CALL_T AT ADDR=021507604  
  next=0  
  state=3 rpc_e_dg_cs_final (delayed ack pending)  
  status=0 state_timestamp=073601 cv=021507632  
  sock_ref=025722444 actid_hash=2834813747 key_info=025503205  
  current socket level: 0, min/max: 0/16  
  compartment: 0, valcmp: 07777777777777777777  
  auth_epv=04063571 timer=021475122  
  last_rcv_timestamp=073601 start_time=073601 high_seq=010  
  pkt_chain=0 com_timeout_knob=07 refcnt=03  
  stop_timer=0 is_cbk=0  
  is_in_pkt_chain=0  
  current call level: 10, min/max: 0/16 <----- Check current call level  
  compartment: 0, valcmp: 07777777777777777777  
  addr at 025730505:
```



```
149 Swaprtc:0 Swapttime:0 Swpaddr:0 Swpby:NULL
149 Swpcat_in:0 Swpcat_out:0 Swplnk:NULL Swpri_in:0.000000
149 Swpri_out:0.000000 Swpriset:0 Swpskip:0/0
149 Swpsize:0 Swqb:NULL Swqf:NULL Tccount:0 Textp:0 Time:18647
149 Ucsiz:4 Usoff:8 Wasyn:0 Writedepth:0 Xstat:0 Xswct:0
```

Level:0 <= means this process is at level 0. <= shows active levels
Level:10 <= means this process is at level 10.

Step 5 - A large value for maxFileSize from the crash command fshost indicates this system has at least talked to a Cray to negotiate the file size.

> **fshost -p**

```
FS_HOST ADDR=0120401652
HS_HOST ADDR=0120401652
  lruq=0120401652 sstates=04 HADTOKENS
  refCount=2 privatep=0 type=0 hostops=(fshs_ops)
SOCKADDR_IN: s_addr=128.162.245.5 (gust)
flags=0x720 INCALL NEW_IF NEW_ACL_IF
priIndex=0326 secIndex=0343
lastCall=Mon Apr 29 08:39:18 1996
activeCall=0
clntEpoch=Mon Apr 29 08:39:18 1996
hostLife=0170 RPCGuarante=0170 deadSrvTmout=0264
registerTokens=01 cbBinding=0120401716 maxFileSize=0x7fffffff <---- Has negotiated with a Cray
FHS_PRINCIPAL ADDR=0120324604 next=0
refCount=0 hostp=0120401652 lastCall=Thu May 2 16:12:35 1996
states (0)=NONE
XCRED AT ADDR 0120350224:
  lruq.next=010507011 lruq.prev=010507011 refcount=01
  changeCount=02 assocPag=01017777776
  flags (0)=*NONE*
  uflags (0)=*NONE*
  propListP=0120350264
UCRED AT ADDR 0116261201:
  cr_ref=1 cr_uid=-2 cr_gid=-2
  cr_sid=0 cr_ruid=-2
sec_id_pac_t addr=0
MLS min/max: 0/0 valcmp: 0
```

Resolution:

Correct the label or NAL label range.

3.9 **Symptom:** Unable get to the junction after CDS (Cell Directory Server) is restored. Receive a host timed out message.

Resolution

1. Stop (kill -9) and restart dfsbind. To restart -
cd /opt/dcelocal/etc/cores/dfsbind
nohup /opt/dcelocal/bin/dfsbind >> /opt/dcelocal/etc/logs/dfsbind 2>&1 &
2. Perform the following cm command to invalidate the mount point and force a refresh from the fldb.
/opt/dcelocal/bin/cm checkfilesets

4. Miscellaneous Information

- 4.1 Requirements for a UNICOS only cell. All DCE cells must include at least one copy of each of the servers listed below. Because systems running the UNICOS or UNICOS/MK operating system do not support these servers, they must be available on a non-UNICOS platform in your cell.

Security Server - provides security services for the cell, including the interface to the registry database.

Cell directory Service (CDS) - provides naming services for the cell.

Fileset Location Server - tracks the location of data in the cell (DFS use only).

- 4.2 DFS use with NQE. With the release of NQE version 3.0 for UNICOS NQE provides support for access to DFS filespace.

- 4.2 Changing your DCE password from the SGI UNICOS or UNICOS/MK system. With the release of UNICOS/UMK DCE 1.1 the `chpass` command is supported. Users may change their DCE registry password as shown below. You must have DCE credentials before initiating this command.

```
$ /opt/dcelocal/bin/chpass -p  
Changing registry password for galileo.  
New password:  
Re-enter new password:  
Enter your previous password for verification:
```

If the UNICOS system has the integrated login feature turned on and both the user's current UDB and DCE registry passwords match, the `/bin/passwd` command will change the user's password in both locations. If the user only wanted to change the UDB password, they can accomplish this by using the `/bin/passwd.udb` command.