

# TruCluster V5 disk manipulation commands fail with I/O errors

Last Updated: 2002-04-26 Solution-id: 253.0.4518925.2847710

## Environment:

TruCluster V5.0A

TruCluster V5.1

disklabel

mkfdmn

HSZ80

HSG80

HSV110

## Problem Description:

disk manipulation commands fail with I/O errors

persistent reservation

## Solution:

Persistent reservations are maintained as lists of reservation keys. The reservations can be listed using the show keys command of scu:

```
# scu -f /dev/rdisk/dsk10c
scu> show keys
Persistent Reservation Header:
Generation Value: 2
Additional Length: 16
Reservation Key List:
Key Entry 0: 0x10001
Key Entry 1: 0x10001
```

To remove the persistent reservations on a device you can use scu to register a new key and then clear it which will remove all reservations:

```
scu> preserve register key 0 skey 1
scu> preserve clear key 1
scu> show keys
Persistent Reservation Header:
Generation Value: 4
Additional Length: 0
Reservation Key List:
```

## Cause:

Compaq's TruCluster software uses SCSI bus reservations to promote device sharing whilst maintaining data integrity. The TruCluster V1.x products used the SCSI-2 reserve/release commands to protect shared disk access but these reservations are lost during bus reset or power downs. The cluster software had to monitor the reservations and

if they were lost the reserves had to be taken out again. If a cluster became partitioned it was possible for the devices to be accessed simultaneously from both partitions and data corruption could result.

The TruCluster Server V5.x products overcome this problem by using the SCSI-3 persistent reservation commands. With persistent reserves the devices retain reservation information after bus resets or power loss.

During normal cluster use the system manager will not be aware of the existence of persistent reservations but if cluster maintenance needs to be performed by booting from the Tru64 UNIX install disk then the cluster disks may be unreachable. Here is an example of a disklabel command being run against a cluster disk when booted from the standalone Tru64 UNIX disk:

```
# disklabel -r dsk10
disklabel: dsk10: I/O error
```

If the I/O error is examined using DECEvent or Compaq Analyze then the cause will be seen as a "Unit Reserved" error from the device. If the device is on a HSZ80 or HSG80 RAID controller then a show <unit> command on the console will give a state of "Persistent reserved".

#### **With the release of PK1 for V5.1a cleanPR is now supported on HSV110:**

```
HSG80> show d2
LUN Uses Used by
-----
D2 BOOTMIR (partition)
LUN ID: 6000-1FE1-0008-8380-0009-0240-3165-0025
IDENTIFIER = 2
Switches:
RUN NOWRITE_PROTECT READ_CACHE
READAHEAD_CACHE WRITEBACK_CACHE
MAXIMUM_CACHED_TRANSFER_SIZE = 256
Access:
ALL
State:
ONLINE to this controller
Persistent reserved
Size: 8884583 blocks
Geometry (C/H/S): ( 2629 / 20 / 169 )
```