

Power775
Drive (HDD/SSD) Service Procedure
Last Modified 9/28/2011



Power775 Drive (HDD/SSD) Service Procedure
PN: 41U8515, EC N44172

CONTENTS

1	GENERAL	4
1.1	RELEASE / REVISION HISTORY	4
1.2	WHERE TO FIND THIS DOCUMENT, AND CONTENTS OF THE PARENT PDF	4
1.3	REQUIRED DOCUMENTS	4
1.4	RELATED DOCUMENTS	4
1.5	ABBREVIATIONS	5
2	OVERVIEW	6
2.1	SAFETY NOTICES	6
2.2	ATTENTION: FILE SYSTEM RISK STATEMENT	6
2.3	CONFIRM HOW YOU GOT TO THIS POWER 775 DRIVE SERVICE PROCEDURE	6
2.4	DRIVE DESCRIPTION	7
2.5	DRIVE WEIGHT	10
2.6	BACKGROUND	11
2.6.1	Enclosure VPD label and Location Code and Recovery Group Name	13
2.7	CONCURRENCY	14
2.8	REQUIRED SSRS AND ROLES	14
2.9	ESTIMATED SERVICE TIME	14
2.10	REQUIRED TOOLS	14
2.11	PREREQUISITES FOR THIS PROCEDURE	14
2.12	SERVICE PROCEDURE OVERVIEW	14
3	SERVICE PROCEDURE	15
3.1	IDENTIFY THE FRAME <= SSR TASK	15
3.2	SUSPEND THE DRIVES IN THE DRIVE CARRIER <= CUSTOMER TASK	22
3.4	DRIVE REPLACEMENT <= SSR TASK	26
3.5	UNSUSPEND THE DRIVES <= CUSTOMER TASK	32
3.6	VERIFY THAT THE DRIVE REPLACE FIXED THE PROBLEM <= CUSTOMER TASK	33
3.7	CONFIRM WITH THE CUSTOMER <= SSR TASK	33
4	END OF POWER775 DRIVE (HDD/SSD) SERVICE PROCEDURE	33
5	APPENDIX A: POWER775 BPC FSP COMMAND LINE PROCEDURE	34
5.1	PROCEDURE TO ACCESS THE BPC FSP COMMAND LINE	34
5.2	END OF APPENDIX A: POWER775 BPC FSP COMMAND LINE PROCEDURE	36

Power775 Drive (HDD/SSD) Service Procedure
PN: 41U8515, EC N44172

Figure List:

Figure 1 Hard Drive Assembly Photo..... 7
Figure 2 SSD Assembly Photo..... 7
Figure 3 Hard Drive Assembly Drawing 7
Figure 4 SSD Assembly Drawing..... 7
Figure 5 Drive Carrier Drawing 8
Figure 6 Drive Carrier Photo - Top View 8
Figure 7 Drive Carrier Photo - Side View (Labels Show Drive Position)..... 8
Figure 8 Disk Enclosure (Front View – showing 48 Drive Carriers, 8 Port Cards, and 2 DCCAs)..... 9
Figure 9 Disk Enclosure Photo (Front View – showing 48 Drive Carriers, 8 Port Cards, and 2 DCCAs)..... 9
Figure 10 Disk Enclosure (Rear View – showing 48 Drive Carriers and 8 Port Cards) 9
Figure 11 Electro-mechanical Actuator in Drive Carrier Latch (2 Views)..... 11
Figure 12 DE VPD Label..... 13
Figure 13 Disk Enclosure HEX Cage ID 16
Figure 14 Frame Photo showing UEPO location..... 17
Figure 15 UEPO assembly showing frame Identify Amber LED..... 18
Figure 16 Door Latch..... 19
Figure 17 Door Latch handle shown extended..... 19
Figure 18 The Front DE Enclosure ID LED is located on the DCCAs..... 20
Figure 19 Rear DE Enclosure ID LED is located on the bottom right side of the disk enclosure:..... 21
Figure 20 Disk Enclosure LEDs 27
Figure 21 Port Card Slot Identify LEDs 27
Figure 22 Drive Carrier Slot Identify LED 28
Figure 23 Drive Carrier LEDs Drawing..... 28
Figure 24 Drive 1 LEDs and arrow 29
Figure 25 Drive Carrier LEDs and Drive Locations..... 30
Figure 26 Drive Removal..... 31

Table List

Table 1 Release / Revision History..... 4
Table 2 Required Documents..... 4
Table 3 Related Documents 4
Table 4 Abbreviations..... 5

1 GENERAL

1.1 Release / Revision History

Document Name	Date	PDF name	Description
Power775 Drive Service Procedure	9/28/2011	"p775_de_drive.pdf"	Initial Release

Table 1 Release / Revision History

1.2 Where to find this document, and contents of the parent PDF

The current "Power775 Drive Service Procedure" document is "p775_de_drive.pdf" which is to be downloaded from:
InfoCenter Website: <http://publib.boulder.ibm.com/infocenter/powersys/v3r1m5/topic/p7ee2/p7ee2kickoff.htm>

Click "PDF files for the IBM Power 775 (9125-F2C) removing and replacing parts"

Under "Disk Enclosure", click "Drive (HDD/SSD)" to download PDF "p775_de_drive.pdf"

This is the only valid source for the latest Power775 Drive Service Procedure.

1.3 Required Documents

Document	Doc Number	Location
Safety Notices http://publib.boulder.ibm.com/infocenter/powersys/v3r1m5/topic/p7hdx/G229-9054.pdf	G229-9054	InfoCenter

Table 2 Required Documents

1.4 Related Documents

Document	Document Number	Location
"GPFS Native RAID Administration and Programming Reference"	SA23-1354	InfoCenter *

Table 3 Related Documents

*InfoCenter Website: <http://publib.boulder.ibm.com/infocenter/powersys/v3r1m5/topic/p7ee2/p7ee2kickoff.htm>

Power775 Drive (HDD/SSD) Service Procedure
PN: 41U8515, EC N44172

1.5 Abbreviations

Abbreviation	Definition	Details
BPA	Bulk Power Assembly	
CEC	Central Electronic Complex	Also referred to as the “node”.
DCCA	Distributed Conversion and Control Assembly	The power supplies for the CEC and DE are called the CEC DCCA and DE DCCA respectively.
DE	Disk Enclosure	
GPFS	Global Parallel File System	IBM’s file system utilizing software RAID
HDD	Hard Disk Drive	This also means hard drive
HPIC	High Power Interface Cable	DCCA power cable
LED	Light Emitting Diode	
LIC	Licensed Internal Code	
PCB	Printed Circuit Board	
RAID	Redundant Array of Inexpensive Disks	
SAS	Serial Attached SCSI	Protocol used for direct attached storage
SSR	Systems Services Representative	IBM Service personnel
SSD	Solid State Drive	
UEPO	Unit Emergency Power Off	

Table 4 Abbreviations

2 OVERVIEW

This section is an overview only. Do not start the service procedure until Section 3 which contains the detailed steps.

2.1 Safety Notices

Read “Safety_Notices “ available from InfoCenter – see Section 1.3.
<http://publib.boulder.ibm.com/infocenter/powersys/v3r1m5/topic/p7hdx/G229-9054.pdf>

The following cautions apply to all Power775 service procedures:

CAUTION:

Energy hazard present. Shorting might result in system outage and possible physical injury. Remove all metallic jewelry before servicing. (C001)

CAUTION:

The doors and covers to the product are to be closed at all times except for service by trained service personnel. All covers must be replaced and doors locked at the conclusion of the service operation. (C013)

CAUTION:

Service of this product or unit is to be performed by trained service personnel only. (C032)

2.2 ATTENTION: File System Risk Statement

The Global Parallel File System (GPFS) implementation of software RAID stripes data across all the Disk Enclosures in the cluster. **It is very important to follow the procedure and only replace a single drive at a time following the replacement priority order; if too many drives are removed while GPFS is active, the file system will become unavailable. DO NOT swap drives between different locations. (There is a special case where multiple drives in the same Drive Carrier may be replaced at the same time. This is described in the procedure notes.)**

2.3 Confirm how you got to this Power 775 Drive Service Procedure

You should be performing this procedure if you have discovered through routine maintenance checks or from automated notification that the maximum limit threshold of defective drives in a declustered array has been exceeded. Continued operation under this condition may put at risk the loss of data if an additional drive failure occurs.

The P7 cluster event notification TEAL agent can be configured to send disk replacement notices when they occur to the P7 cluster EMS.

You should download this procedure from:

InfoCenter Website: <http://publib.boulder.ibm.com/infocenter/powersys/v3r1m5/topic/p7ee2/p7ee2kickoff.htm>

This is the only valid source for the latest Power775 Drive Service Procedure

2.4 Drive Description

The definition of the Drive in the context of Power775 is either of the following, and this procedure applies to both types of drives:

- a. A 2.5” Small Form Factor (SFF) Serial Attached SCSI (SAS) Hard Drive Assembly
This will be referred to as “hard drive”. See Figure 1 & 3.
- b. A 2.5” SFF SAS Solid State Drive (SSD) Assembly
This will be referred to as “SSD”. See Figure 2 & 4.



Figure 1 Hard Drive Assembly Photo



Figure 2 SSD Assembly Photo




Figure 3 Hard Drive Assembly Drawing




Figure 4 SSD Assembly Drawing

Power775 Drive (HDD/SSD) Service Procedure

PN: 41U8515, EC N44172

The Drive Carrier is used to hold up to four Drives located (Figures 5,6 and 7).
There are 96 Drive Carriers in the Disk Enclosure (DE) (see Figures 8,9 and 10).




Figure 5 Drive Carrier Drawing



Figure 6 Drive Carrier Photo - Top View



Figure 7 Drive Carrier Photo - Side View (Labels Show Drive Position)

Power775 Drive (HDD/SSD) Service Procedure
PN: 41U8515, EC N44172




Figure 8 Disk Enclosure (Front View – showing 48 Drive Carriers, 8 Port Cards, and 2 DCCAs)

A = Drive Carrier B = Port Card C = DE DCCA



Figure 9 Disk Enclosure Photo (Front View – showing 48 Drive Carriers, 8 Port Cards, and 2 DCCAs)




Figure 10 Disk Enclosure (Rear View – showing 48 Drive Carriers and 8 Port Cards)

2.5 Drive Weight

Hard Drive Assembly	0.51 lbs (0.23 kg)
SSD Assembly	0.30 lbs (0.12 kg)
Drive Carrier by itself	0.87 lbs (0.39 kg)
Drive Carrier with 4 Drive Assemblies installed	2.91 lbs (1.32 kg)

2.6 Background

The Global Parallel File System (GPFS) implementation of software RAID stripes data across all the recovery groups in a cluster. Every recovery group contains one small declustered array (DA) of 4 SSDs and from 1 to 4 DAs of 47 hard drives each. Each Disk Enclosure (DE) contains 2 recovery groups. One consists of the drives in the top Drive Carriers and the other consists of the drives in the bottom Drive Carriers. Each Drive Carrier contains from 1 to 4 drives. Concentrating on half a DE, these drives are further divided into four groups of 47 hard drives and 1 SSD. For the upper recovery group, the drives of the first DA reside in drive position 1 of the 48 upper Drive Carriers. The drives in the second DA reside in drive position 2, and so on.

The entire file system can be compromised if a single declustered array has lost enough drives to eliminate redundancy and then loses another drive. When you remove a Drive Carrier for service, you are removing 4 drives at once from 4 different DAs. Only the file system can check the health of its declustered arrays. **Thus, it is critical to interact with the file system to ensure that the drives on the Drive Carrier have been suspended and that only the correct Drive Carrier is removed when performing the service.**

An electro-mechanical actuator locks the Drive Carrier's latch as soon as one or more drives are powered on in that Drive Carrier. This was implemented as a final check that the SSR has identified the correct Drive Carrier to remove in order to do the Drive service, as only the Drive Carrier being serviced will be unlocked. The actuator pin is moved to lock the latch trigger. See Figure 7 below. (The blue and red colors are only for contrast.)

NOTE – When the Drive Carrier is removed, if the actuator pin is accidentally moved to the locked position, it will be necessary to move gently back to the unlocked position with your finger before reinstalling the Drive Carrier.




Figure 11 Electro-mechanical Actuator in Drive Carrier Latch (2 Views)

Power775 Drive (HDD/SSD) Service Procedure

PN: 41U8515, EC N44172

GPFS has a replacement threshold of 1 set for the small DAs containing the SSDs, and a replacement threshold of 2 set for the large DAs. This means that if an SSD fails in a small DA, it will be marked immediately for replacement. By contrast, when only one hard drive fails, it is not marked for replacement; when the second hard drive in that DA fails, both failed hard drives are marked for replacement.

GPFS provides several indications that disk replacement is required:

- * AIX error report and Linux syslog entries
- * The pdReplacePdisk callback can be configured to run an admin-supplied script at the moment a pdisk is marked for replacement
- * The P7 cluster event notification TEAL agent can be configured to send disk replacement notices when they occur to the P7 cluster EMS
- * The mmlsrecoverygroup command with the -L and --pdisk flags shows the states of all pdisks, which may be examined for the "replace" attribute
- * The mmlspdisk command with the --replace flag lists only those pdisks that are marked for replacement

Power775 Drive (HDD/SSD) Service Procedure PN: 41U8515, EC N44172

2.6.1 Enclosure VPD label and Location Code and Recovery Group Name

The VPD label is placed on the rear of a DE just above the drawer LEDs. The following picture shows the information contained in the label.



Figure 12 DE VPD Label

An example of a location code of Drive 4 in Drive Carrier 35 contained in the DE that this VPD label is attached to would be as follows:

U78AD.001.992001B-P1-C35-D4

By convention as explained in the GPFS D-RAID Administration and Programming Reference document, the 7 digit unique SN on the enclosure VPD label, is to be used in the recovery group name as follows:

xxxxxxxTOP
xxxxxxxBOT

where xxxxxxx is the final 7 digits of the unique SN.

2.7 Concurrency

The drive is concurrently maintainable without losing any customer function and while maintaining redundancy.

NOTE: If a Drive Carrier is out longer than the repair time-out (configurable parameter in GPFS with default of 5 minutes), the disks go from Suspended to Suspended/Missing state and GPFS will start draining the data that was on those disks to spare space. Normally, this won't really hurt anything, but in the extreme case where some affected array brought down too many disks, it can cause the server to resign control of the recovery group, causing delays and other inconvenience. The procedure out of this situation is to put the carrier back and run mmchcarrier --replace or --resume. In the resign situation, the mmchcarrier replace/resume command may take a long time to work while it waits for the backup server to take over the recovery group.

2.8 Required SSRs and Roles

This service procedure contains steps to be performed by a customer system administrator and an IBM SSR (System Service Representative).

- The customer system administrator performs the STEPs in Sections marked with “<= CUSTOMER TASK”
 - The SSR performs the STEPs in Sections marked with “<= SSR TASK”
-

2.9 Estimated Service Time

0.5 Hour – 1 Hour

2.10 Required Tools

The P7IH qualified service ladder PN is 46G5947.

2.11 Prerequisites for this Procedure

In order to perform this procedure, you will need the following information:

- 1) The location code of the FRU to be serviced
 - 2) The cage location of the FRU to be serviced
 - 3) The frame number and frame serial number of the FRU to be serviced
-

2.12 Service Procedure Overview

This is an overview of the tasks to be performed. Read this overview but do not perform any of the tasks yet.

3.1	IDENTIFY THE FRAME <= SSR TASK	15
3.2	SUSPEND THE DRIVES IN THE DRIVE CARRIER <= CUSTOMER TASK	22
3.4	DRIVE REPLACEMENT <= SSR TASK	26
3.5	UNSUSPEND THE DRIVES <= CUSTOMER TASK	32
3.6	VERIFY THAT THE DRIVE REPLACE FIXED THE PROBLEM <= CUSTOMER TASK	33
3.7	CONFIRM WITH THE CUSTOMER <= SSR TASK	33

3 SERVICE PROCEDURE

STOP – Do not proceed unless you have read “Safety Notices” which is available from InfoCenter; see Section 1.3.

These steps define the procedure to replace a drive in a P7IH cluster.

For additional detail on using GPFS, see “GPFS D-RAID Administration and Programming Reference” available on the Boulder public website. Also, manual pages for the GPFS commands will be installed as part of GPFS.

3.1 Identify the frame <= SSR TASK

- STEP 1 Use the procedure “Appendix A: Power775 BPC FSP Command Line Procedure” (included in Section 5 of this document) to access the BPC FSP command line of the frame containing the Disk Enclosure with the FRU needing replacement.
- STEP 2 Blink the Enclosure Identify LEDs of the DE with the FRU needing replacement by using the following command from the BPC Command Line:

This is the command syntax:

```
bpccmd -c 27<blank>0000<cage ID><FRU ID><LED blink state 00=off, 02=blink>
```

<blank>	= “11”
<cage ID>	= see Figure 13
<FRU ID>	= “01” for the DE Identify LED.
<LED state>	= “02” for this step

Type in the following, inserting the correct Cage ID DE chassis location:

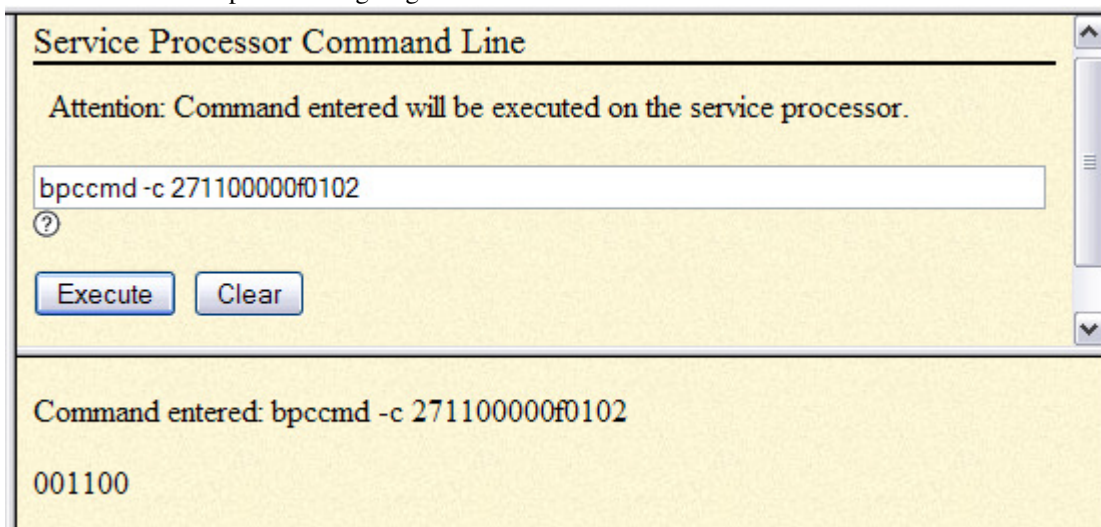
```
bpccmd -c 27110000xx0102
```

Where “xx” is the Cage ID of the DE you are servicing.

This field will be 05,07,09,0B,0D, or 0F (see Figure 13).

The correct return code is 001100. If you do not receive this, contact the next level of support.

Below is an example assuming Cage ID=0f



Power775 Drive (HDD/SSD) Service Procedure
 PN: 41U8515, EC N44172

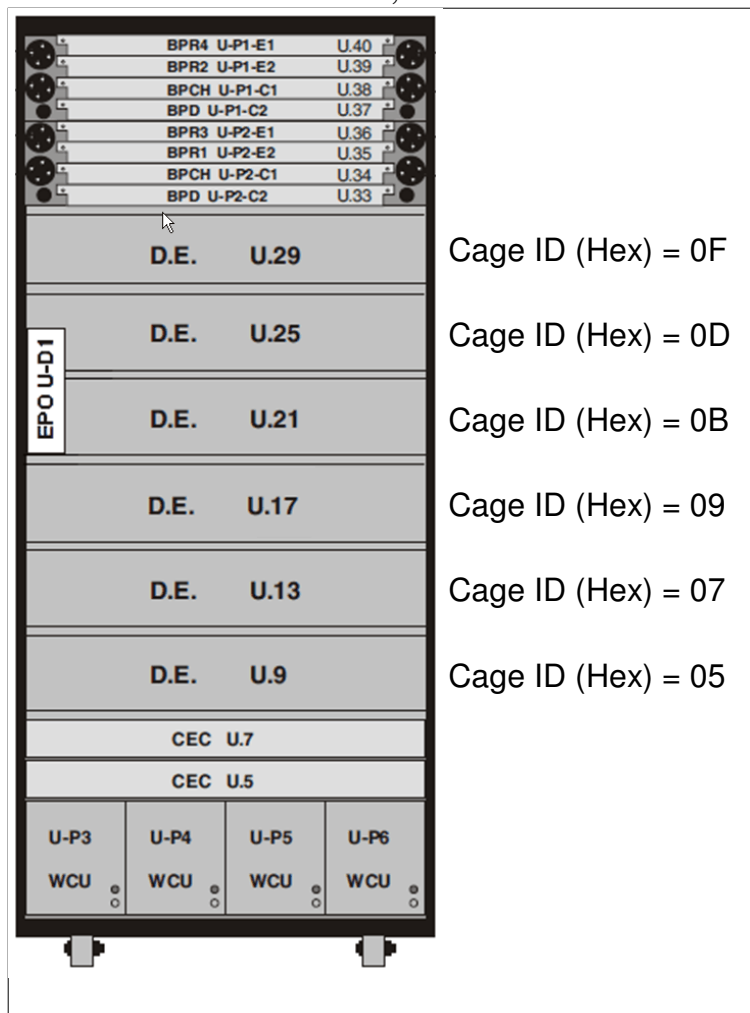


Figure 13 Disk Enclosure HEX Cage ID

Power775 Drive (HDD/SSD) Service Procedure
PN: 41U8515, EC N44172

- STEP 3 Locate the frame. Confirm that it is the correct frame by checking that its UEPO Panel Identify LED is blinking. This LED is located on the UEPO panel and it is marked by the “!” symbol (see Figure 14 and Figure 15).




Figure 14 Frame Photo showing UEPO location

Power775 Drive (HDD/SSD) Service Procedure PN: 41U8515, EC N44172




Figure 15 UEPO assembly showing frame Identify Amber LED

STEP 4 If the Drive Carrier is at location c001 to c056, then open the front door.

If the Drive Carrier is from location c057 to c112, then open the rear door.

Locate the latch along the left side of the door; push in on the X button, the latch handle will pop out, extended out from the latch pocket. Grasp the latch handle and swing the door to the right (see Figure 16 and Figure 17).

Power775 Drive (HDD/SSD) Service Procedure
PN: 41U8515, EC N44172



Figure 16 Door Latch



Figure 17 Door Latch handle shown extended

Power775 Drive (HDD/SSD) Service Procedure
PN: 41U8515, EC N44172

STEP 5 Identify the Disk Enclosure

Identify the Disk Enclosure requiring service by looking for the blinking amber Enclosure ID LEDs on the DE DCCAs on the front side of the Disk Enclosure (see Figure 18) OR on the right side of the rear of the Disk Enclosure (see Figure 19).

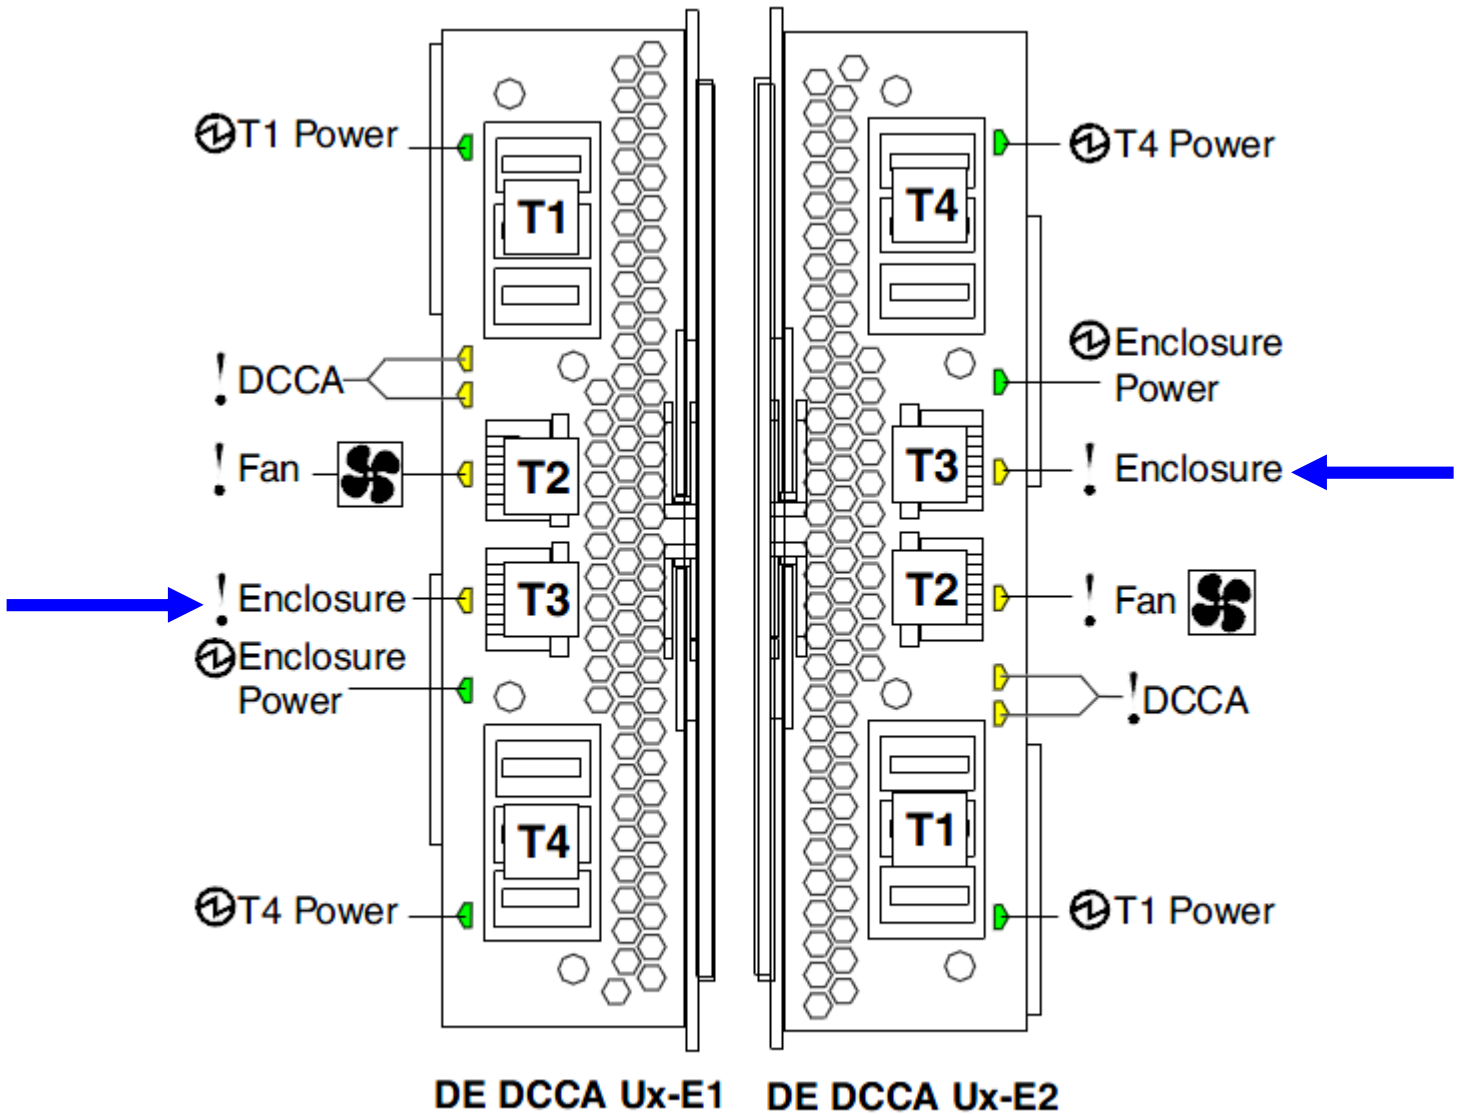


Figure 18 The Front DE Enclosure ID LED is located on the DCCAs

Power775 Drive (HDD/SSD) Service Procedure PN: 41U8515, EC N44172




Figure 19 Rear DE Enclosure ID LED is located on the bottom right side of the disk enclosure:

STEP 6 Check the Enclosure VPD Label

Verify that the location code of the Drive to be replaced matches the enclosure VPD Label of the DE you have identified (see Figure 12).

STEP 7 Bring over the Ladder if you are going to need it for service.

STEP 8 Turn off the Enclosure Identify LEDs of the DE by using the following command from the BPC FSP Command Line:

This is the command syntax:

```
bpccmd -c 27<blank>0000<cage ID><FRU ID><LED blink state 00=off, 02=blink>
```

- <blank> = "11"
- <cage ID> = see table below
- <FRU ID> = "01" for the DE Identify LED.
- <LED state> = "00" for this step

Type in the following, inserting the correct cage ID DE chassis location:

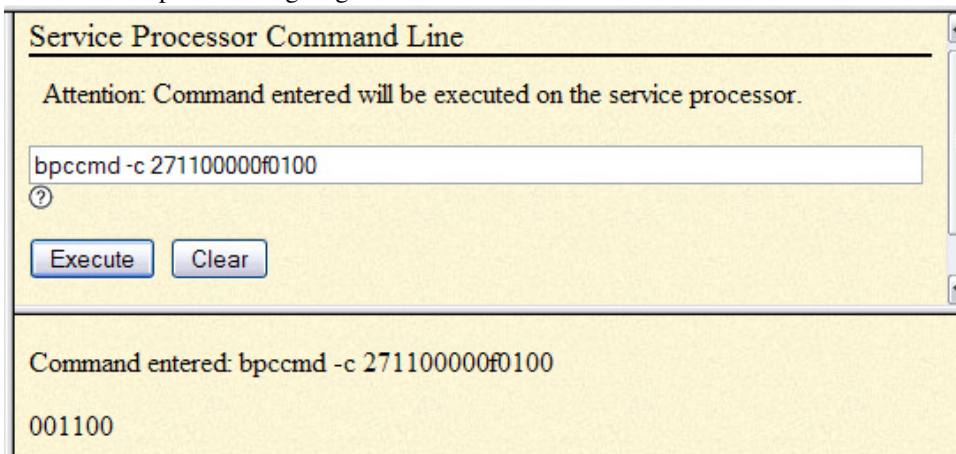
```
bpccmd -c 27110000xx0100
```

Where "xx" is the cage ID of the DE you are replacing.

This field will be 05,07,09,0B,0D, or 0F (see Figure 13).

The correct return code is 001100. If you do not receive this, contact the next level of support.

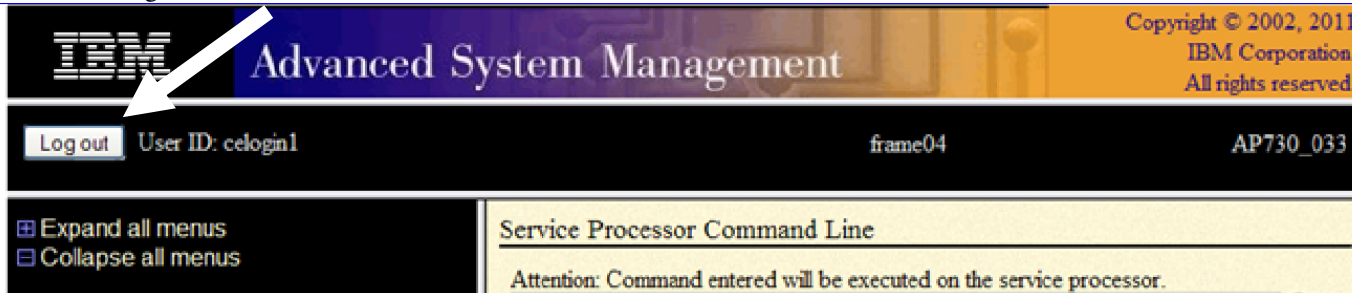
Below is an example assuming Cage ID=0f



Power775 Drive (HDD/SSD) Service Procedure

PN: 41U8515, EC N44172

STEP 9 Log out of BPC FSP ASM.



3.2 Suspend the drives in the Drive Carrier <= Customer Task

STEP 10 Log onto the Executive Management Server (EMS)

STEP 11 Log onto the Service Node.

STEP 12 SSH into a GPFS I/O Node

STEP 13 From the GPFS LPAR, issue command:

```
"/usr/lpp/mmfs/bin/mmlspdisk all --replace".
```

This will list all pdisks requiring replacement in all the DE recovery groups.

ATTENTION: You must only extract one Drive Carrier at a time.

If the command shows no pdisks requiring replacement AND if you are sure you there is a disk you want to replace, proceed to STEP 13.

Example output assuming 2 Drive fails in declustered array "DA3" in DE with enclosure VPD label "000DE37"

```
pdisk:
  replacementPriority = 1.00
  name = "c046d3"
  device = "/dev/rhdisk158,/dev/rhdisk62"
  recoveryGroup = "000DE37TOP"
  declusteredArray = "DA3"
  state = "dead/systemDrain/noRGD/noVCD/replace" ← Indicates defective
drive
  freeSpace = 513745775616
  fru = "74Y4936"
  location = "78AD.001.000DE37-C46-D3"
  WWN = "naa.5000C5001D90E17F"
  server = "server1"
  reads = 986
  writes = 293368
  bytesReadInGiB = 0.826
  bytesWrittenInGiB = 286.696
  IOErrors = 0
  IOTimeouts = 12
  mediaErrors = 0
```

Power775 Drive (HDD/SSD) Service Procedure

PN: 41U8515, EC N44172

```
checksumErrors = 0
pathErrors = 0
timeBadness = 33.278
dataBadness = 0.000
```

pdisk:

```
replacementPriority = 1.00
name = "c047d3"
device = "/dev/rhdisk630,/dev/rhdisk726"
recoveryGroup = "000DE37TOP"
declusteredArray = "DA3"
state = "dead/systemDrain/noRGD/noVCD/noData/replace" ← Indicates
```

defective drive

```
freeSpace = 599147937792
fru = "74Y4936"
location = "78AD.001.000DE37-C47-D3"
WWN = "naa.5000C5001DC5DF3F"
server = "server1"
reads = 1104
writes = 379053
bytesReadInGiB = 0.844

bytesWrittenInGiB = 370.717
IOErrors = 0
IOTimeouts = 0
mediaErrors = 0
checksumErrors = 10
pathErrors = 0
timeBadness = 0.000
dataBadness = 10.000
```

Record the following information from each “pdisk” report output:

- replacementPriority (Replacement Priority is the number displayed after the “=” character.)
- name (pdisk name is the text between the locations)
- recoveryGroup (Recovery Group is the text between the quotations.)
- location (Location Code is the text between the quotations.)

Replacement priority determines Drive Replacement Order!

You must only extract one Drive Carrier at a time. The lowest priority number must be replaced first! (It is possible to have two drives with the same priority. In that case, choose either as long as they are both the lowest priority number).

Priority numbers can be negative. A replacement order example is below:

Replacement Priority	Replace Drive Order
-1.00	1 st
1.00	2 nd
3.98	3 rd

In our example for the first “pdisk” instance (Drive #3 of Carrier #46)

```
replacementPriority = 1.00
name = "c046d3"
recoveryGroup = "000DE37TOP"
location = "78AD.001.000DE37-C46-D3"
```

Power775 Drive (HDD/SSD) Service Procedure PN: 41U8515, EC N44172

NOTE: If the particular recovery group is known (xxxxxxxTOP or xxxxxxxBOT for top or bottom drive Carriers), you can issue the command `“/usr/lpp/mmfs/bin/mmldisk xxxxxxxTOP --replace”`, where xxxxxxx is the recovery group’s number. For the remainder of the document, xxxxxxxTOP will be used to reference the recovery group in question – alter the command to refer to the exact name of the recovery group you are servicing.

NOTE: Remember that a hard drive could be failed but will not yet be marked for replacement if it is the first hard drive in that DA to fail. The command `“/usr/lpp/mmfs/bin/mmlsrecoverygroup xxxxxxxTOP --L”` will list the details for that recovery group, including the replace thresholds.

NOTE: To list all recovery groups, use command `“/usr/lpp/mmfs/bin/mmlsrecoverygroup”`.

NOTE: If there are 2,3 or 4 drives in the same Drive Carrier, all with marked for replacement, and all having the same replacement priority, it is possible to replace them all at the same time.

STEP 14 Once the drive (or drives) to be replaced is chosen and confirmed by STEP 12, the customer sys admin must notify the SSR that a drive replacement is required.

Multiple drive replaces can be identified and confirmed via STEP 12 and the SSR can plan to do more than one during the visit (**but only one at a time**).

NOTE: If a carrier is out longer than the repair time-out (configurable parameter in GPFS with default of 5 minutes), the disks go from Suspended to Suspended/Missing state and GPFS will start draining the data that was on those disks to spare space. Normally, this won't really hurt anything, but in the extreme case where some affected array brought down too many disks, it can cause the server to resign control of the recovery group, causing delays and other inconvenience. The procedure out of this situation is to put the carrier back and run `mmchcarrier --replace` or `--resume`. In the resign situation, the `mmchcarrier replace/resume` command may take a long time to work while it waits for the backup server to take over the recovery group.

NOTE ON TIME CONSTRAINT: The following step should be performed after the SSR prepares the frame for work. The door should be opened and, if necessary, the ladder should be in place to minimize repair time. If the Drive Carrier is not replaced and resumed within the repair time-out (configurable parameter in GPFS with default of 5 minutes), then it's possible that this could result a server resign, causing longer unavailability.

STEP 15 From the GPFS LPAR, issue command from either of the command formats below

`/usr/lpp/mmfs/bin/mmchcarrier xxxxxxxTOP --release --pdisk <pdisk name>`

NOTE: The command will suspend drives that will result in the loss of the availability of the file system if this is necessary in order to fix/rebuild the corrupted drives. In this case it will suspend disk operations to the unavailable DA until the replacement is complete.

NOTE: If in STEP 12 you have confirmed more than 1 drive needs to be replaced in the Drive Carrier, then use command of this format:

`/usr/lpp/mmfs/bin/mmchcarrier xxxxxxxTOP --release --pdisk ‘<pdisk name1>;<pdisk name2>’`

For example:

`/usr/lpp/mmfs/bin/mmchcarrier 000DE37TOP --release --pdisk ‘c046d1;c046d2’`

Power775 Drive (HDD/SSD) Service Procedure

PN: 41U8515, EC N44172

NOTE: If the Drive is not marked for replacement and if you are sure you want to release this Drive, add the “--force-release” option, which will override the check to see if a disk requires replacement. **ONLY USE THIS COMMAND IF ABSOLUTELY NECESSARY.** If a sufficiently large number of pdisks are down in the affected arrays, the file system can pause due to the release.

Example command using “—force-release”:

```
/usr/lpp/mmfs/bin/mmhcarrier xxxxxxTOP --release --force-release --pdisk <pdisk name>
```

The above command suspends all the drives on the Drive Carrier to be removed. This command should complete quickly and release the Drive Carrier.

Example output on success:

```
mmhcarrier 000DE37TOP --release --pdisk c046d3
```

```
[I] Suspending pdisk c046d1 of RG 000DE37TOP in location 78AD.001.000DE37-C46-D1.  
[I] Suspending pdisk c046d2 of RG 000DE37TOP in location 78AD.001.000DE37-C46-D2.  
[I] Suspending pdisk c046d3 of RG 000DE37TOP in location 78AD.001.000DE37-C46-D3.  
[I] Suspending pdisk c046d4 of RG 000DE37TOP in location 78AD.001.000DE37-C46-D4.  
[I] Carrier released.
```

- Remove carrier.
- Replace disk in location 78AD.001.000DE37-C46-D3 with FRU 74Y4936.
- Reinsert carrier.
- Issue the following command:

```
mmhcarrier 000DE37TOP --replace --pdisk 'c046d3'
```

Repair timer is running. Perform the above within 5 minutes to avoid pdisks from being reported as missing.

The command will fail if the customer sys admin attempts to release a Drive Carrier where a Drive does not need replacement.

Example output on fail:

```
[E] Pdisk c046d3 of RG xxxxxxTOP --location is not currently slated for replacement.  
mmhcarrier: Command failed. Examine previous error messages to determine cause.
```

3.4 Drive Replacement <= SSR Task

NOTE ON TIME CONSTRAINT: If the Drive Carrier is not replaced and resumed within the repair time-out (configurable parameter in GPFS with default of 5 minutes), then it's possible that this could result a server resign, causing longer unavailability.

- STEP 16 Check the LED states and the state of the latch actuator on the Drive Carrier. Remember that removing the wrong Drive Carrier may bring down the entire cluster filesystem. Before unseating the Drive Carrier, the SSR MUST:
- a. Check that the Frame, Cage Location, and FRU location code for the Drive Carrier that was deactivated by the customer sys admin all match exactly the physical Frame, Cage Location and FRU location code.
 - b. Check that the Frame (UEPO) , Enclosure ID, Port Card Slot ID, and amber Drive Carrier Slot ID LEDs are blinking (see Figure 20, Figure 21, and Figure 22).

NOTE: The Drive Carrier Slot ID and the Port Card Slot ID LED functions have 2 LEDs on the backplane going into a single light pipe. If the LED is not blinking, do not assume that the LED is broken as it is very unlikely that both are broken (see Figure 21) If this LED is not blinking, it is recommended to stop and contact next level of support.

- c. Verify the green Drive Carrier power LED is off. Do not remove the Drive Carrier if its power LED is on. If the power LED is illuminated at least one hard drive or SSD is active. Removal of the carrier while any drive is active may BRING DOWN THE FILESYSTEM. Make sure all green drive activity LEDs on the Drive Carrier are off before removal of the Drive Carrier (see **Error! Reference source not found.**Figure 23).
- d. Check that the Drive Carrier latch is unlocked. Remember that the latch for the Drive Carrier is on the left side. Make sure you have the correct latch and you are not off by one slot. When the latch is unlocked, a slight downward and outward pressure with your finger tip on the terracotta touch point will move the latch. No side to side motion is required.

NOTE: Do not attempt to circumvent the latch locking mechanism. If slight pressure does not move it, double check all of the items above. Remember that if you unseat the wrong Drive Carrier, you may bring down the entire cluster filesystem. If you are positive that you have the correct Drive Carrier and the LEDs are in the correct state, and yet the latch is locked, CONTACT THE NEXT LEVEL OF SUPPORT.

Power775 Drive (HDD/SSD) Service Procedure

PN: 41U8515, EC N44172




Figure 20 Disk Enclosure LEDs




Figure 21 Port Card Slot Identify LEDs

Power775 Drive (HDD/SSD) Service Procedure
PN: 41U8515, EC N44172




Figure 22 Drive Carrier Slot Identify LED




Figure 23 Drive Carrier LEDs Drawing

STEP 17 After the Drive Carrier to be removed has been verified, unseat the card using the terracotta lever and remove the Drive Carrier. Make sure to complete STEP 18 right away after removal.

ATTENTION: While handling the Drive Carrier, be very careful not to damage the signal pins on the 120 signal pin connector. If the pin field looks damaged, do not plug the Drive Carrier into the DE as it will likely damage the backplane connector on the chassis.

STEP 18 Confirm the Drive (or possibly multiple Drives) to be replaced by observing the Drive Identify LED. Locate the blinking amber Drive Identify LED which indicates the position of the Drive to be replaced. There are 2 LEDs for each drive position for redundancy. The amber LEDs are next to a silk screen arrow pointing to the latch of the

Power775 Drive (HDD/SSD) Service Procedure

PN: 41U8515, EC N44172

Drive to be replaced, and are labeled “Drive 1”, “Drive 2”, “Drive 3”, and “Drive 4”. (see Figure 24 and Figure 25).

NOTE: The DE Port Cards control 12 Drive Carrier cards in a redundant fashion. Each DE Port Card controls one of two DRIVE Identify LEDs for each drive position. Since there are two separate LED drivers and a low duty cycle, the LEDs will blink out of phase. If there is a broken LED, there will be only a single LED blinking for the drive position – this is not a problem and the Drive Carrier should not be replaced for this since the redundancy is working as intended.




Figure 24 Drive 1 LEDs and arrow

Power775 Drive (HDD/SSD) Service Procedure
PN: 41U8515, EC N44172




Figure 25 Drive Carrier LEDs and Drive Locations

Power775 Drive (HDD/SSD) Service Procedure

PN: 41U8515, EC N44172

STEP 19 Remove the failed

To remove the drive from the Drive Carrier:

- a. Lift the blue latch in order to slide the tab out of the slot in the Drive Carrier PCB.
- b. Slide the hard drive or SSD away from its connector (towards the latch end).
- c. Lift the latch end first out of the Drive Carrier, followed by the connector end.




Figure 26 Drive Removal

- STEP 20 Install the replacement Drive by placing the drive into the Drive Carrier making sure 4 tabs on the Drive sled align and slide into the slots in the Drive Carrier rails.

Press down and slide the Drive toward the connector until it is fully seated in the connector. The blue latch should have snapped into place once the Drive was fully seated. If not, press on the blue latch to make sure it engages with the slot in the Drive Carrier printed circuit board (PCB)..

NOTE: The amber Drive Identify LEDs will continue to blink even after installing the new Drive as long as the capacitors have store energy left. **BE CAREFUL** not to swap any Drive locations. The only Drive that should be removed is the Drive that is being replaced in the present procedure.

- STEP 21 Inspect the Drive Carrier signal connector for bent pins before installing the Drive Carrier back into the same location it was removed from.

ATTENTION: DO NOT plug a Drive Carrier with bent connector signal pins into a Disk Enclosure. It will damage the Disk Enclosure. If there are any bent pins, replace the Drive Carrier with a new Drive Carrier and be careful to move the drives from the old Drive Carrier to the same locations in the new Drive Carrier.

Power775 Drive (HDD/SSD) Service Procedure

PN: 41U8515, EC N44172

- STEP 22 After inspecting the Drive Carrier connector signal pins plug the Drive Carrier back into the chassis at the same location it was removed from. The DE Port Card will sense plug detect of the Drive Carrier and will power on the hard drives and SSD if applicable. The DE Port Card will turn off the amber Drive Carrier slot LED, turn on the green Drive Carrier power LED, and lock the electro-mechanical actuator. It should take no longer than about one minute for this power up procedure to complete.
- STEP 23 Close the doors to the frame and inform the customer sys admin that the Drive has been replaced.

3.5 Unsuspend the drives <= Customer Task

- STEP 24 The command to issue will be shown as output of the command you used to suspend the drives. Normally, you would have done a "--release" and will do a "--replace" command to unsuspend the drives.

Issue the following command to GPFS to resume the use of the suspended drives and begin using the replaced drive:

```
/usr/lpp/mmfs/bin/mmchcarrier xxxxxxTOP --replace --pdisk <pdisk name>
```

For example:

```
/usr/lpp/mmfs/bin/mmchcarrier 000DE37TOP --replace --pdisk c046d3
```

NOTE: If the Drive Carrier was out longer than the repair time-out (configurable parameter in GPFS with default of 5 minutes), the disks have gone from Suspended to Suspended/Missing state and GPFS will start draining the data that was on those disks to spare space. Use the same command above to unsuspend the drives if this happens. In the resign situation, the mmchcarrier replace/resume command may take a long time to work while it waits for the backup server to fail over the recovery group.

NOTE: If in STEP 13 you targeted more than 1 drive for replacement in the Drive Carrier, then use command of this format:

```
/usr/lpp/mmfs/bin/mmchcarrier xxxxxxTOP --replace --pdisk '<pdisk name1>;<pdisk name2>'
```

For example:

```
/usr/lpp/mmfs/bin/mmchcarrier 000DE37TOP --replace --pdisk 'c046d3;c046d4'
```

NOTE: If you had to use the "--force release" option, you will have to use the "--resume" option to resume the drives.

For example:

```
/usr/lpp/mmfs/bin/mmchcarrier 000DE37TOP --resume --pdisk c046d3
```

NOTE: After the drive is replaced or resumed, the following LEDs should turn off: the Drive Carrier Slot ID LED, the Port Card Slot ID LED, the DE Enclosure ID LED, and the UEPO Panel ID LED.

3.6 Verify that the drive replace fixed the problem <= Customer Task

STEP 25 From the GPFS LPAR, issue command:

`"/usr/lpp/mmfs/bin/mmlspdisk all --replace".`

This will list all pdisks requiring replacement in all the DE recovery groups.

Make sure the drive location that you replaced no longer shows up in the list.

3.7 Confirm with the customer <= SSR Task

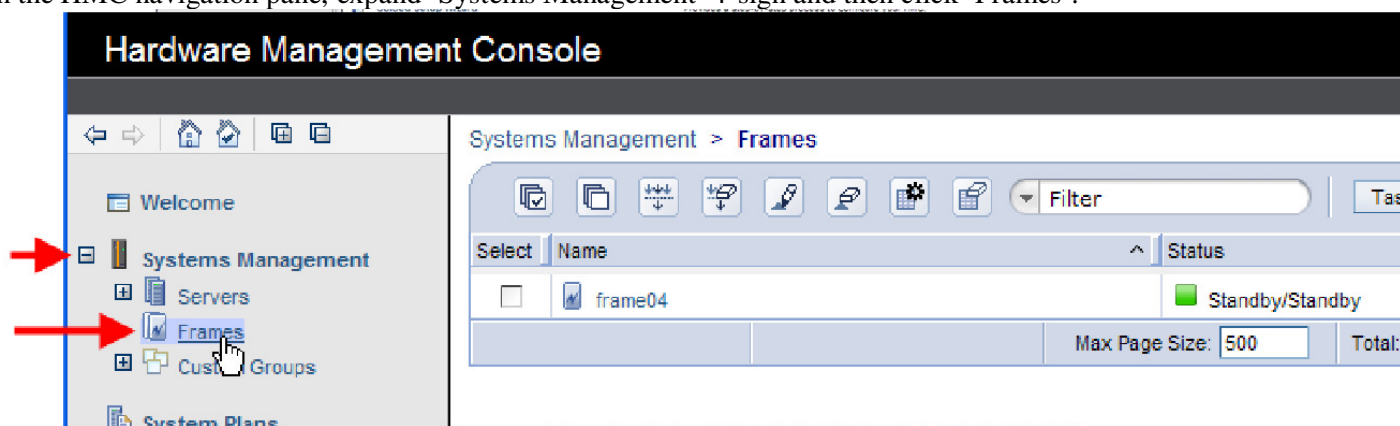
STEP 26 Confirm with the customer that the Drive Carrier replacement fixed the problem.

4 END OF POWER775 DRIVE (HDD/SSD) SERVICE PROCEDURE

5 APPENDIX A: POWER775 BPC FSP COMMAND LINE PROCEDURE

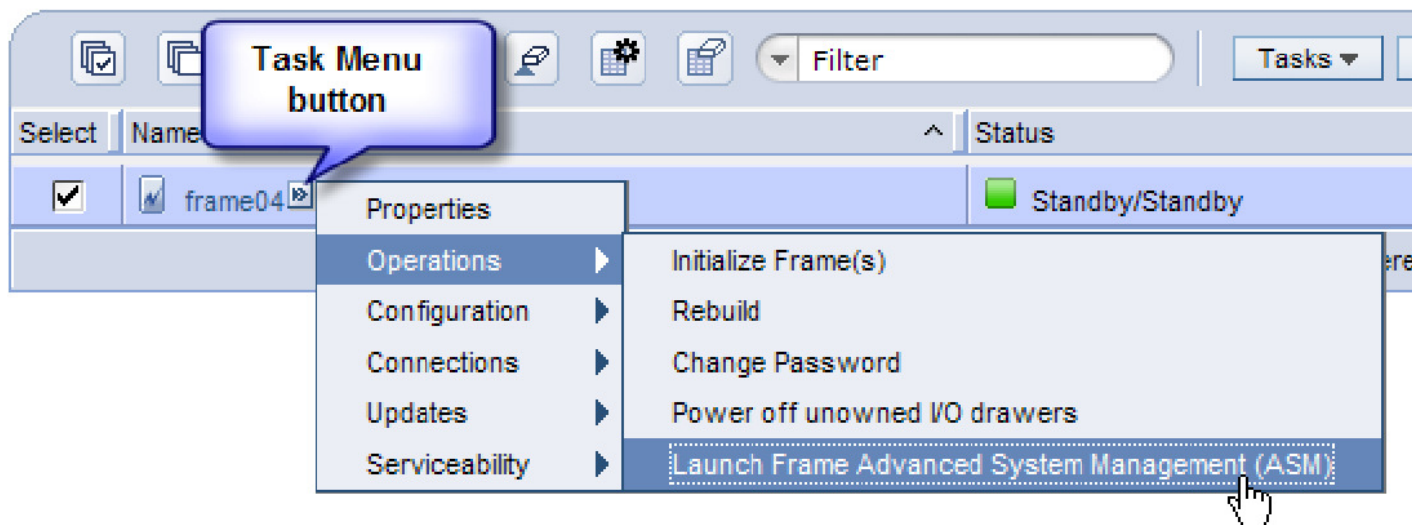
5.1 Procedure to Access the BPC FSP Command Line

1. The HMC can be accessed via the keyboard/display that resides in the network management rack.
2. Login to the HMC if not done already.
3. In the HMC navigation pane, expand 'Systems Management' + sign and then click 'Frames':



4. From the Tasks Menu right-arrow pull-down menu, click **Operations** → **Launch Frame Advanced Systems Management (ASM)**

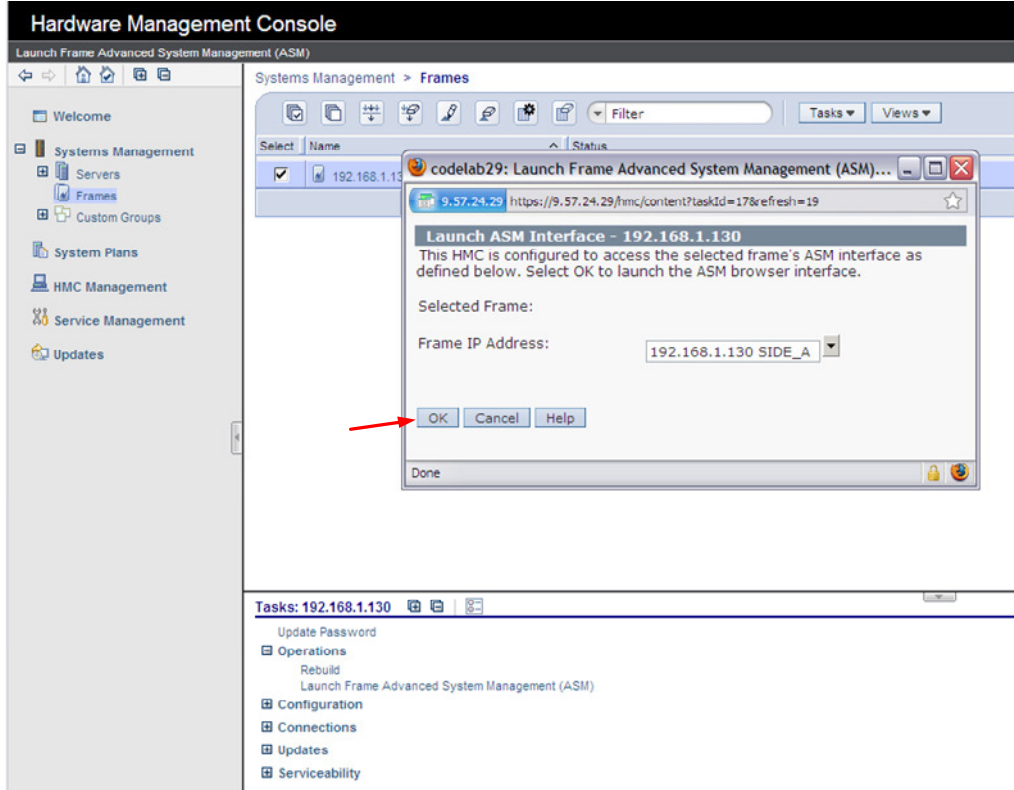
Systems Management > Frames



Power775 Drive (HDD/SSD) Service Procedure

PN: 41U8515, EC N44172

5. From the Launch ASM Interface window, select **SIDE_A** for the 'Frame IP Address' then click the **OK** button.




6. The ASM login window is presented. Acquire the necessary User ID and Password.

“celogin” requires Daily PW from the IBM Support Center.

“celogin1” might be has enabled by the customer. If so, obtain the password from the customer.


1. Enter User ID
2. Enter Password
3. Click **Log in** button



Power775 Drive (HDD/SSD) Service Procedure

PN: 41U8515, EC N44172

- Expand 'System Service Aids' + sign and Select 'Service Processor Command Line'.



- The ASM BPC FSP Command line will be presented. Enter the commands defined in the paper service procedures at this command line and press the Enter key (or click the **Execute** button).
- Return to the step of the paper service procedure that directed you to this Appendix.

5.2 End of Appendix A: Power775 BPC FSP Command Line Procedure