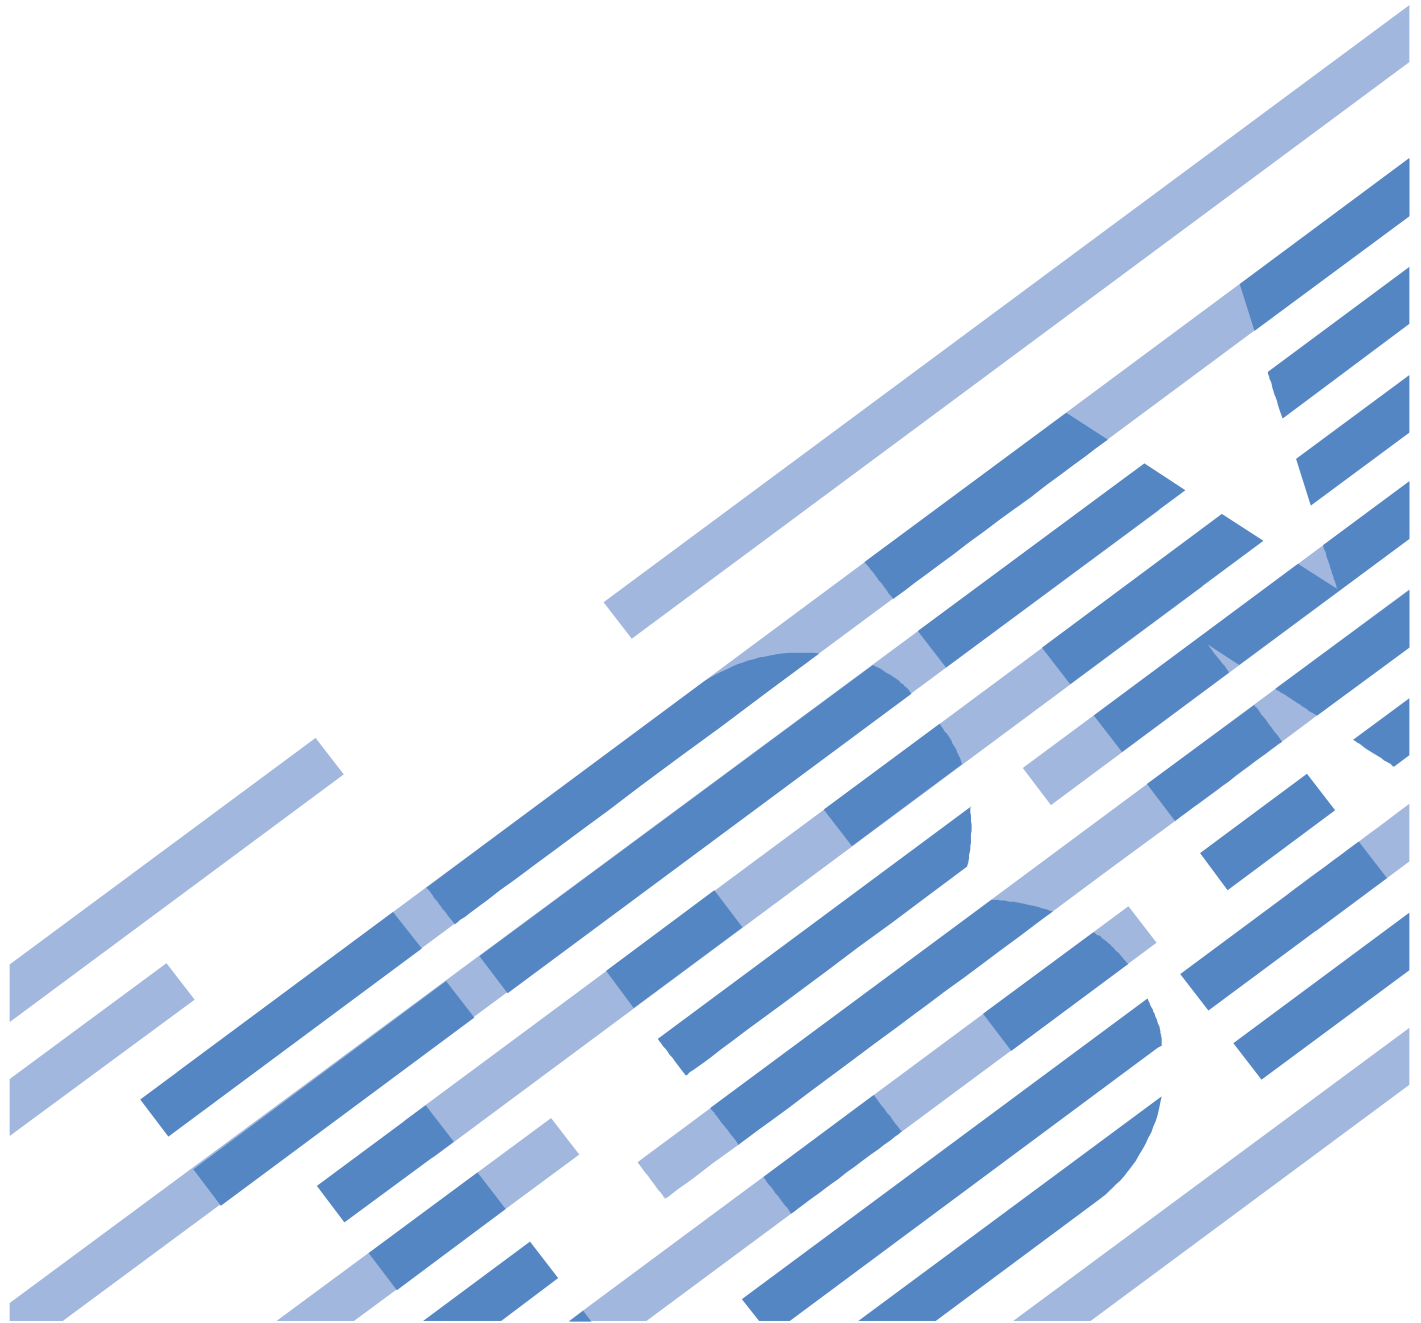# Power775
Utility Non-Volatile Storage Assembly and SAS Jumper Cable (UNVSA) Service Procedure
Last Modified 09/14/2012

CONTENTS

Figure List

Table List

1 GENERAL

1.1 Release / Revision History

Document Name	Date	PDF name	Description
Power775 Utility Non-Volatile Storage Assembly and SAS Jumper Cable (UNVSA) Service Procedure	2/13/2012	"p775_unvsa_cable.pdf"	Initial Release
"	9/14/2012	"p775_unvsa_cable.pdf"	Update

Table 1 Release / Revision History

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

The current Power775 Utility Non-Volatile Storage Assembly and SAS Jumper Cable (UNVSA) Service Procedure document is "p775_unvsa_cable.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 "CEC Drawer", click "Power775 Utility Non-Volatile Storage Assembly and SAS Jumper Cable (UNVSA) Service
      Procedure" to download PDF "p775_unvsa_cable.pdf"

This is the only valid source for the latest Power775 Utility Non-Volatile Storage Assembly and SAS Jumper Cable (UNVSA) 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	Doc# G229-9054	InfoCenter *

Table 2 Required Documents

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

1.4 Related Documents

Document	Doc Number	Location
"GPFS Native RAID Administration and Programming Reference"	SA23-1354	InfoCenter *
Managing System Dump Devices	T1000169	www.ibm.com

Table 3 Related Documents

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

1.5   Abbreviations

Abbreviation	Definition	Details
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
LED	Light Emitting Diode	
PCB	Printed Circuit Board	
RAID	Redundant Array of Inexpensive Disks	
SAS	Serial Attached SCSI	Protocol used for direct attached storage
SSR	System Service Representative	IBM Service personnel
SSD	Solid State Drive	
UEPO	Unit Emergency Power Off	
UNVSA	Utility Non-Volatile Storage Assembly	Harddrive function in PCIe cassette

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 at this link:
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:
Servicing of this product or unit is to be performed by trained service personnel only. (C032)

2.2 File System Risk Statement

The Global Parallel File System (GPFS) implementation of software RAID stripes data across all the Disk Enclosures in the cluster. **If the Utility Non-Volatile Storage Assembly and SAS Jumper Cable is powered off without a backup Utility Non-Volatile Storage Assembly and SAS Jumper Cable assembly or without all drawers being supported by the Service Node being reassigned to the backup Service Node, then all Disk Enclosures running off of this Service Node will go into panic and become unavailable.**

2.3 Confirm how you got to this Power 775 Utility Non-Volatile Storage Assembly and SAS Jumper Cable Service Procedure

You should be performing this procedure if

An SRC directed you to replace the Utility Non-Volatile Storage Assembly and SAS Jumper Cable

You should have downloaded 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 Utility Non-Volatile Storage Assembly and SAS Jumper Cable Service Procedure

2.4　Utility Non-Volatile Storage Assembly and SAS Jumper Cable Description

The Utility Non-Volatile Storage Assembly and SAS Jumper Cable units are located at the rear of a CEC Drawer that contains a Service Node. Figure 1 shows a properly configured Utility Non-Volatile Storage Assembly and SAS Jumper Cable set up.
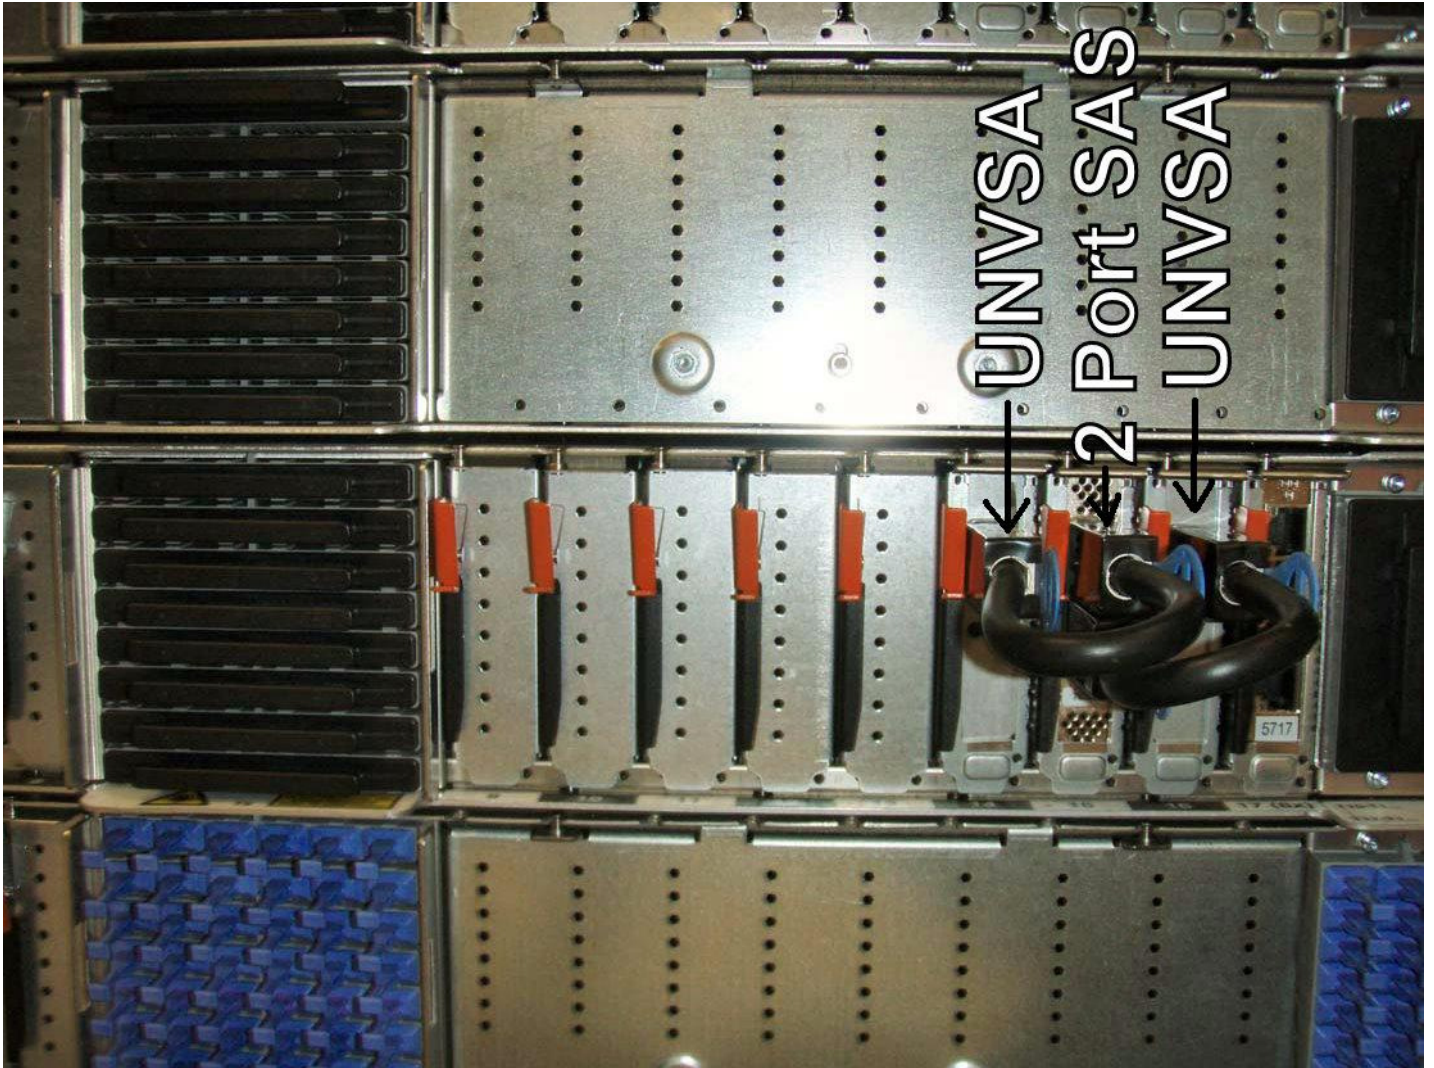
Figure 1 Utility Non-Volatile Storage Assembly and SAS Jumper Cable Locations in Service Node

2.5　Background

A Utility Non-Volatile Storage Assembly and SAS Jumper Cable, which is included in a Service Node, provides diskless boot. The Utility Non-Volatile Storage Assembly and SAS Jumper Cable unit includes a 600GB HDD along with power conversion and SAS connectivity within a PCIe form factor cassette.

2.6　Concurrency

The Utility Non-Volatile Storage Assembly and SAS Jumper Cable is concurrently maintainable. This procedure can be completed while the CEC is powered on.

2.7 Utility Non-Volatile Storage Assembly and SAS Jumper Cable Weight

1.2 lbs (0.54 kg)

2.8 Required SSRs and Roles

This service procedure requires 1 SSR.

This service procedure contains steps to be performed by a customer system administrator and an IBM System Service Representative (SSR)
- The customer system administrator and SSR tasks are separately called out in the procedure.

2.9 Estimated Service Time

90 minutes

2.10 P7IH Hand Tool Kit Required Tools

None

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 of the FRU to be serviced

2.12 Overview of Procedure

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

3 SERVICE PROCEDURE

3.1 Obtain location code of the UNVSA and SAS Jumper Cable to be replaced <= Customer Task

STEP 1   Obtain location code information for the Utility Non-Volatile Storage Assembly (UNVSA) and SAS Jumper Cable. The location code will be similar to the following example:

U78A9.001.312N003-P1-C15-T1-L3   or   U78A9.001.312N003-P1-C15-T1-L7

AIX: One method to obtain this is from the output of "lscfg –v" on the Service Node, under hdisk0 or hdisk1, whichever is associated with the UNVSA and SAS Jumper Cable to be replaced.

RHEL6: The following script will create and output that lists all of the drive names used by Red Hat with a corresponding location code like above. If the drive is no longer communicating with the OS, it will be listed as "phy not attached to end device".

Below is an example of the output of the script. In this case the drive attached to U78A9.001.9998887-P1-C15-T1-L5 and L7 (top cable) is no longer communicating with the OS.

```
SAS Adapter host0
    PCI Location=U78A9.001.9998887-P1-C15-T1
    Drive attached to U78A9.001.9998887-P1-C15-T1-L1
        attached_sas_address 0x5000c5001db6f2b6
        drive_generic_name   /dev/sg0
        drive_device_name    /dev/sda
        multipath_dev_name   dm-0
        drive_sas_address    0x5000c5001db6f2b7
        other_names          ../../dm-0 ../../dm-2 ../../dm-3 ../../dm-4 ../../dm-5 ../../dm-6
    Drive attached to U78A9.001.9998887-P1-C15-T1-L3
        attached_sas_address 0x5000c5001db6f2b5
        drive_generic_name   /dev/sg1
        drive_device_name    /dev/sdb
        multipath_dev_name   dm-0
        drive_sas_address    0x5000c5001db6f2b7
        other_names          ../../dm-0 ../../dm-2 ../../dm-3 ../../dm-4 ../../dm-5 ../../dm-6
    Drive attached to U78A9.001.9998887-P1-C15-T1-L5
        phy not attached to end device
    Drive attached to U78A9.001.9998887-P1-C15-T1-L7
        phy not attached to end device
```

Below is the script used to gather the prior data:

```
#!/bin/bash

scsi_host="/sys/class/scsi_host"
cd ${scsi_host}
subfolders=`ls -1`
for i in ${subfolders};  do
   cd ${i};
   if [ `cat proc_name` == "mpt2sas" ]; then
      host_dir=`pwd -P`
      pci_bus=`ls -1 ${host_dir}/../../../.. | grep 0`
      location=`lscfg -l ${pci_bus}|grep ${pci_bus}|awk '{print $3}'`
      echo -e SAS Adapter ${i}
      echo -e \\tPCI Location=${location}

      scsi_host_phy="/sys/class/scsi_host/${i}/device"
      cd ${scsi_host_phy}
      subfolders="1 3 5 7"
      for j in ${subfolders};  do
         cd phy-0\:${j}/sas_phy/phy*;
         echo -e \\tDrive attached to ${location}-L${j}
         if [ -f ../../port/end_device*/target*/*/sas_address ]; then
            att_phy=`cat ../../port/end_device*/target*/*/sas_address`
            echo -e \\t\\tattached_sas_address\\t${att_phy}
            sg_name=`ls -1 ../../port/end_device*/target*/*/scsi_generic`
            echo -e \\t\\tdrive_generic_name\\t/dev/${sg_name}
            dev_name=`lsscsi -g | grep ${sg_name} | awk '{print $6}'`
            echo -e \\t\\tdrive_device_name\\t${dev_name}
            dm_name=`multipath -ll ${dev_name} | grep mpath | awk '{print $3}'`
            echo -e \\t\\tmultipath_dev_name\\t${dm_name}
            mpath_name=`cat /sys/devices/virtual/block/${dm_name}/dm/uuid`
            echo -e \\t\\tdrive_sas_address\\t0x${mpath_name:7}
            other_names=`ls -l /dev/disk/by-id/* | grep ${mpath_name} | awk '{print $11}'`
            echo -e \\t\\tother_names\\t\\t${other_names}
         else
            echo -e \\t\\tphy not attached to end device
         fi;
         cd ${scsi_host_phy}
      done;
   fi;
   cd ${scsi_host}
done;
```

Note: the lsscsi rpm must be installed for this script to work. To verify it is installed, run "rpm –qa | grep lsscsi". If it is not installed, it is usually in the /install/rhel*/ppc64/Packages directory.

3.2   Identify the frame and CEC <= SSR Task

STEP 2   Access the HMC via the keyboard/display that resides in the management rack.

STEP 3   Place a checkmark in the Select column of the CEC Drawer containing the UNVSA and SAS Jumper Cable requiring service.

STEP 4   *Click* the Tasks Menu button (right-arrow) and *select* **Operations -> LED Status -> Identify LED**. See Figure 2.
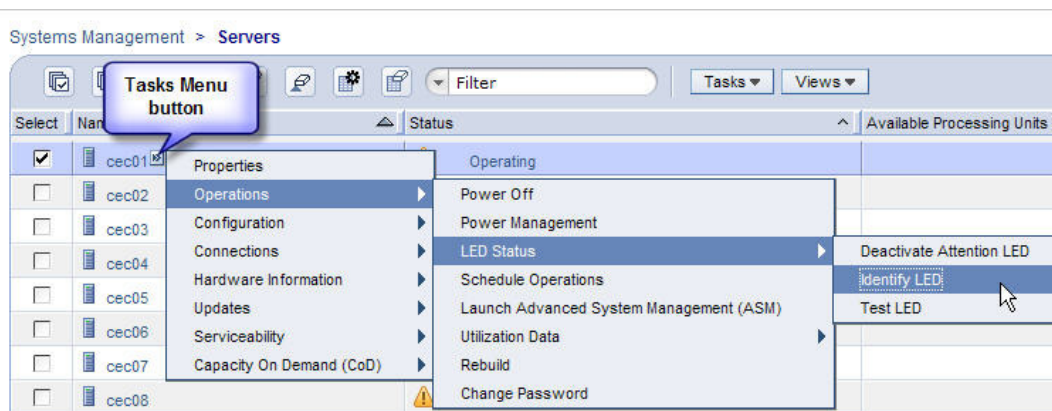
Figure 2 Identify LED Selection

STEP 5   In the window titled **Identify LED, Select Enclosure** *select* **System Unit, Model F2C** then *click* the **Activate LED** button (see Figure 3). This will cause the CEC ID LED to flash, as well as the UEPO Panel of the frame.
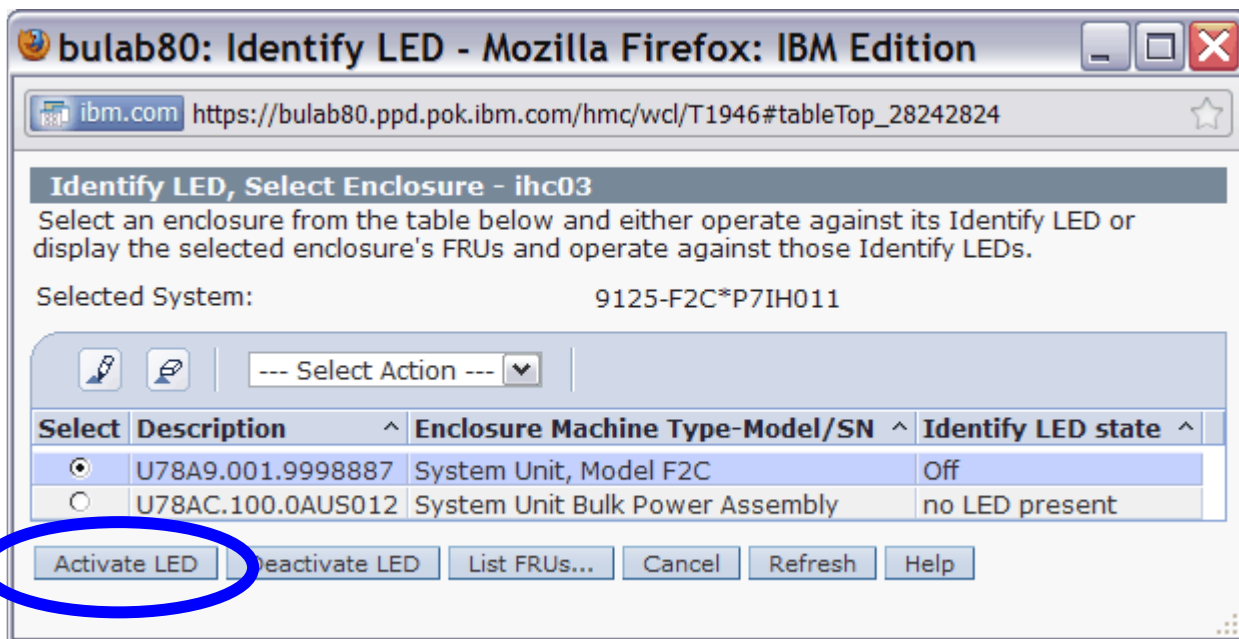
Figure 3 System Unit, Model F2C Selection – Activate LED

STEP 6   Wait for LED activation to completed and *click* the **OK** button.
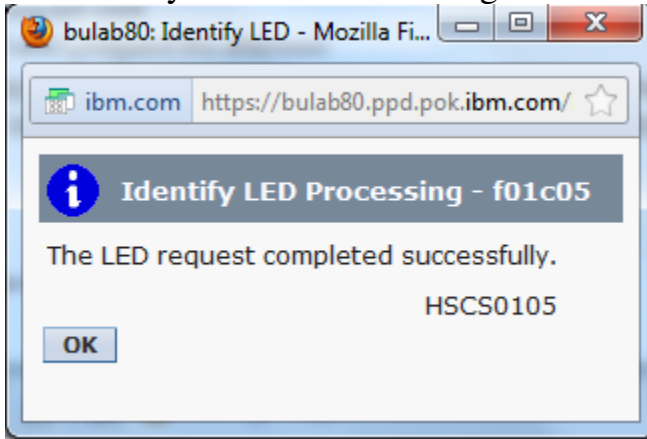
Figure 4 Activate LED Completed

STEP 7   Locate the frame. The UEPO Identify Amber Led (see Figure 5) should be flashing. If the LEDs are not flashing, verify the FRU via MTMS (machine type model serial number). The rack level MTMS is located on the face of the UEPO (see Figure 5). Open the front door of the frame.
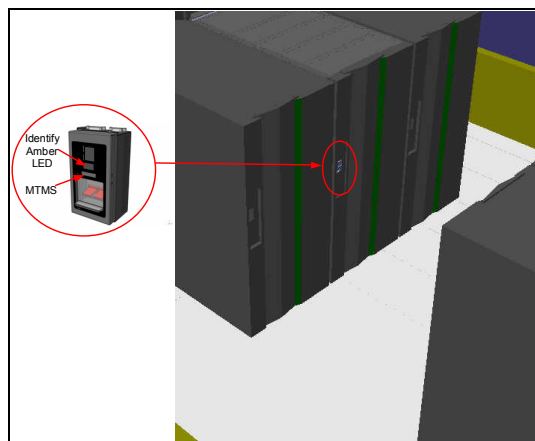
Figure 5 UEPO Identify LED location

STEP 8   Locate the CEC with the UNVSA requiring service. The CEC Drawer Identify LEDs (see Figure 6 and Figure 7) will be flashing. If the LEDs are not flashing, verify the FRU via MTMS (machine type model serial number).
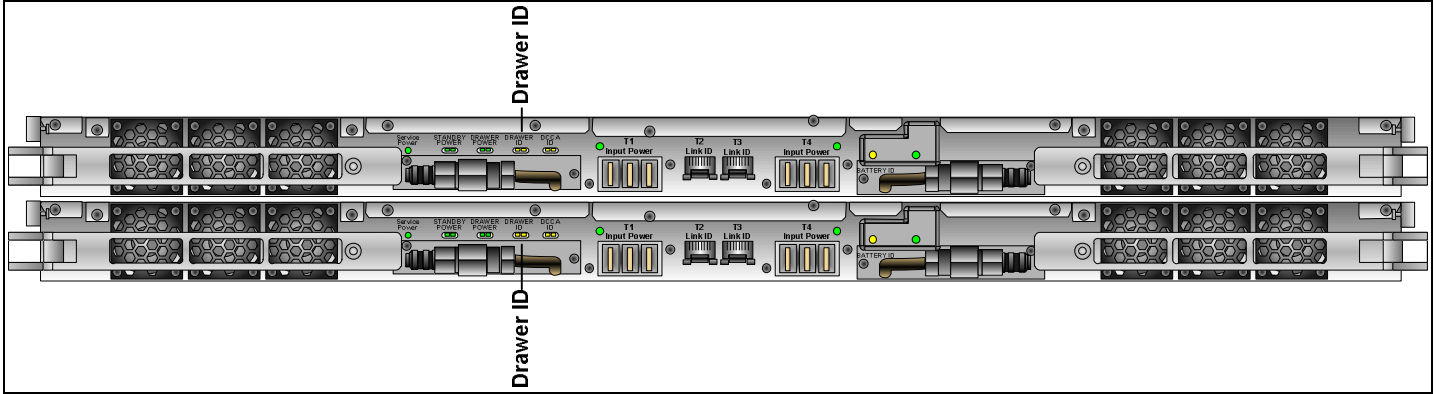
Figure 6 CEC Drawer ID – Front
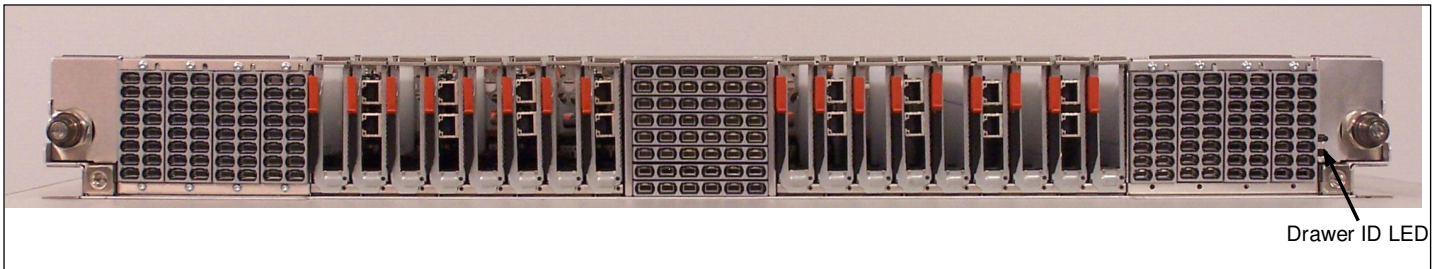
Figure 7 CEC Drawer ID - Rear

STEP 9   Verify that the serial number on the CEC Enclosure VPD label (CEC Node serial # label) matches the location code of the UNVSA and SAS Jumper cable requiring service (see Figure 8). This is the 7 digit serial number in the location code, show in bold below for example:

U78A9.001.**312N003**-P1-C15-T1-L3
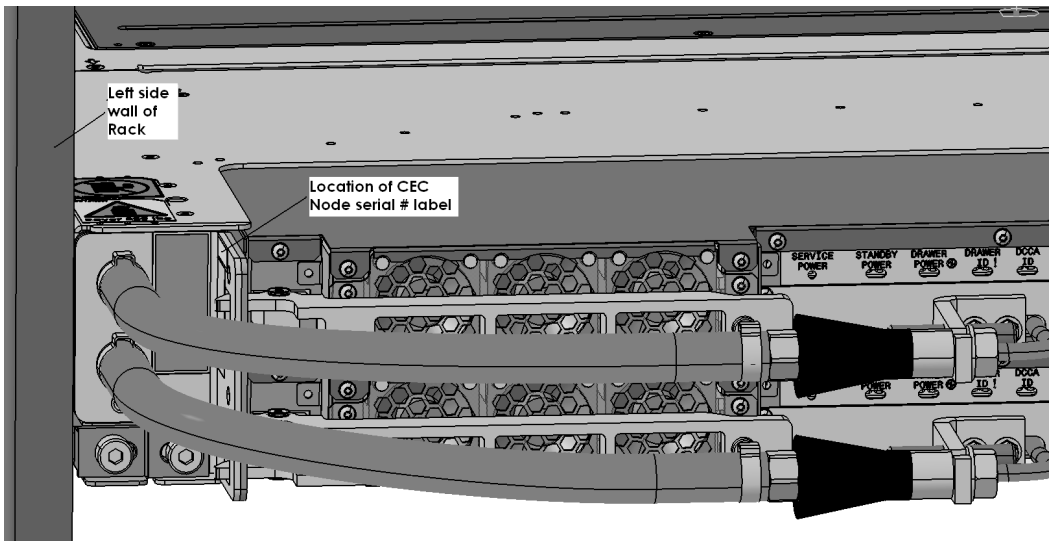
Figure 8 FRU Serial number location

3.3 Identify the UNVSA and SAS Jumper Cable to be replaced <= SSR Task

STEP 10 Open the rear door of the frame, and tracing the SAS Jumper Cable, locate the slot of the UNVSA requiring service.

The key piece of the location code is the "-Lx". The "-Lx" defines which SAS port controls the device being serviced. The "-Lx" is a reference to the two SAS ports on the PCI adapter that is typically installed in PCI slot C15, see Figure 1.

L0-L3 corresponds to the lower SAS cable port on the SAS adapter. (typically slot C16)

L4-L7 corresponds to the upper SAS cable port on the SAS adapter. (typically slot C14)

o If the location code is U78A9.001.312N003-P1-C15-T1-L3 (or L0, L1 or L2), this will refer to the UNVSA that is attached to the lower SAS port of the SAS PCI adapter. You **MUST** follow the cable from the lower SAS port to the attached UNVSA to be sure of the slot you are servicing. Using Figure 1 as an example, this would mean that slot C16 would require service.

o If the location code is U78A9.001.312N003-P1-C15-T1-L7 (or L4,L5 or L6), this will refer to the UNVSA that is attached to the upper SAS port of the SAS PCI adapter. You **MUST** follow the cable from the upper SAS port to the attached UNVSA to be sure of the slot you are servicing. Using Figure 1 as an example, this would mean that slot C14 would require service.

▪ It is very important to follow the SAS Jumper Cable to make sure you are servicing the correct slot.

   Note: Figure 1 is typical. It is possible that the Service Node is cabled differently than Figure 1.

ATTENTION: Be sure to review the location code carefully and after determining the correct port of the SAS PCI adapter, follow the SAS Jumper Cable from that port to the correct UNVSA to be replaced. Failure to follow this correctly will result in turning off the wrong PCI slot of the CEC, thus powering off the remaining good UNVSA and causing Service Node failure, which leads to OS crashes of the compute nodes and any GPFS IO nodes using that Service Node.

STEP 11   In the window titled **Identify LED, Select Enclosure** *select* **System Unit, Model F2C** then *click* the **List FRUs…** button (see **Figure 9**). This brings up the window **Identify LED, Select Location** which you will use in the next step.
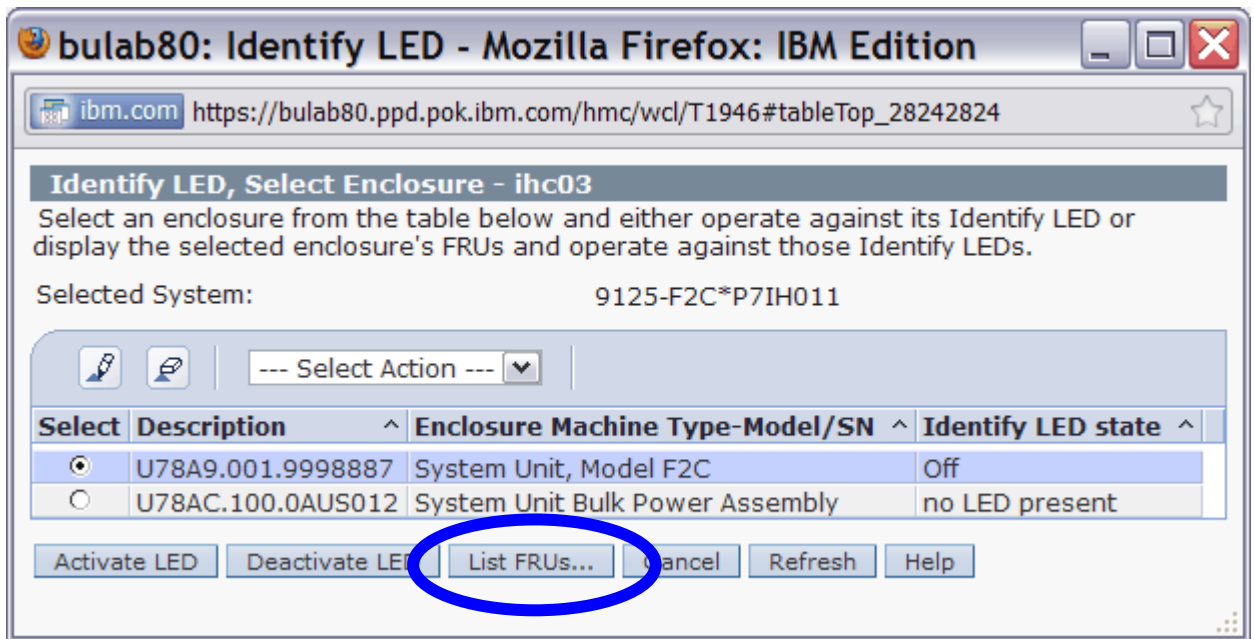
Figure 9 System Unit, Model F2C Selection – List FRUs…

STEP 12   In the window **Identify LED, Select Location**, place a checkmark in the Select column for the UNVSA(it will be shown as a PCI Adapter Card) Location Code requiring service and then *click* the **Activate LED** button. **Figure 10** shows the example of selecting the UNVSA Identify LED in **U78A9.001.[Serial#]-P1-C16** : this is one of two typical locations for a UNVSA, the other being C14. After *clicking* the **Activate LED** button, wait for the screen saying the activate was successful and *click* the **OK** button.
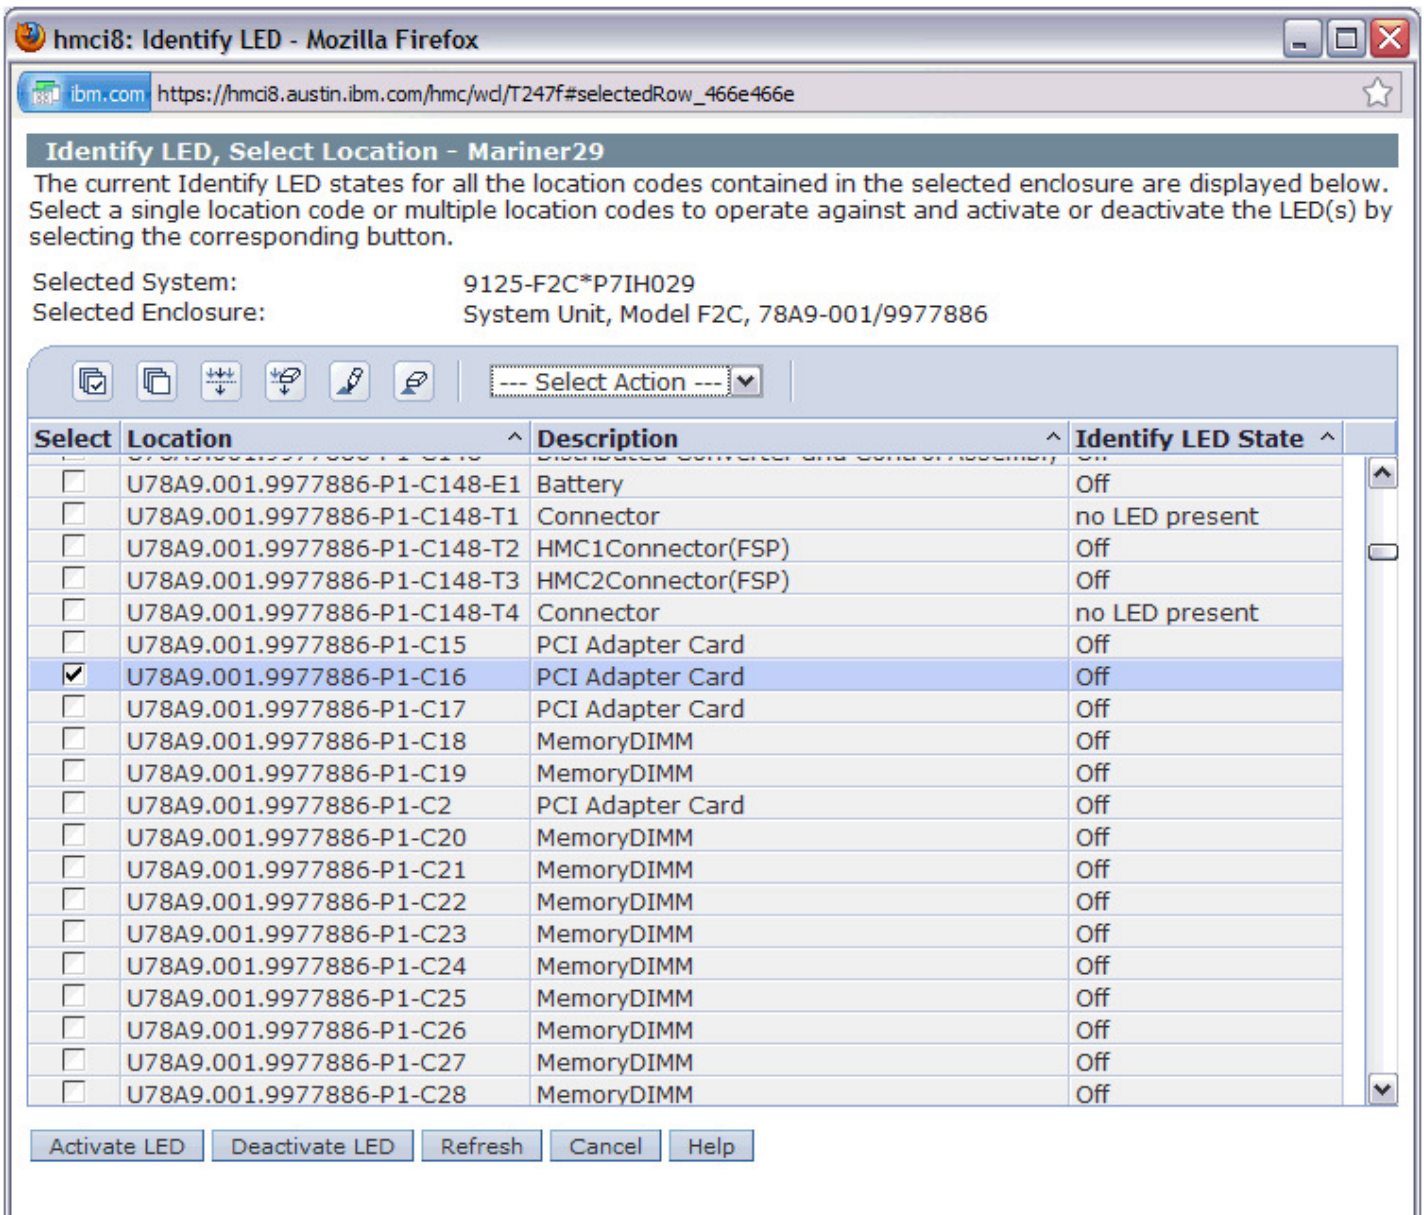
Figure 10 UNVSA Identify LED Selection

STEP 13   Verify that the Identify LED is now flashing for the PCI slot containing the UNVSA that requires service.

3.4   Unmirror the UNVSA <= Customer TASK

If the system is running RHEL6, skip to STEP 22.

AIX:

STEP 14   Ensure that the Utility Non-Volatile Storage Assembly (UNVSA) and SAS Jumper Cable being serviced has no jobs running. This procedure is a concurrent process. The procedure does not affect the other CEC Drawers in the rack.

STEP 15   If the UNVSA and SAS Jumper Cable being serviced is the only UNVSA and SAS Jumper Cable unit in the Service Node(typically there are two), then the CEC Drawers being supported by the Service Node must be reassigned to the backup Service Node in the building block before powering off the Utility Non-Volatile Storage Assembly. If there is only one Utility Non-Volatile Storage Assembly and SAS Jumper Cable, reassign the CEC Drawers to the backup Service Node and then skip to section 3.3.

STEP 16   Check the name of the volume group for hdisk0 and hdisk1. This procedure will assume rootvg. See Figure 11 for sample output

Command: `lspv hdisk0`

`lspv hdisk1`

```
[c250f52c08ap01][/]> lspv hdisk0
PHYSICAL VOLUME:     hdisk0                VOLUME GROUP:     rootvg
PV IDENTIFIER:       00c8b5f625fd4c12 VG IDENTIFIER     00c8b5f600004c000000013125fd4c91
PV STATE:            active
STALE PARTITIONS:    0                     ALLOCATABLE:      yes
PP SIZE:             1024 megabyte(s)      LOGICAL VOLUMES:  15
TOTAL PPs:           558 (571392 megabytes)  VG DESCRIPTORS:   2
FREE PPs:            437 (447488 megabytes)  HOT SPARE:        no
USED PPs:            121 (123904 megabytes)  MAX REQUEST:      1 megabyte
FREE DISTRIBUTION:   111..103..00..111..112
USED DISTRIBUTION:   01..09..111..00..00
MIRROR POOL:         None
[c250f52c08ap01][/]> lspv hdisk1
PHYSICAL VOLUME:     hdisk1                VOLUME GROUP:     rootvg
PV IDENTIFIER:       00c8b5f6e1f7dc4e VG IDENTIFIER     00c8b5f600004c000000013125fd4c91
PV STATE:            active
STALE PARTITIONS:    0                     ALLOCATABLE:      yes
PP SIZE:             1024 megabyte(s)      LOGICAL VOLUMES:  14
TOTAL PPs:           558 (571392 megabytes)  VG DESCRIPTORS:   1
FREE PPs:            441 (451584 megabytes)  HOT SPARE:        no
USED PPs:            117 (119808 megabytes)  MAX REQUEST:      1 megabyte
FREE DISTRIBUTION:   111..103..04..111..112
USED DISTRIBUTION:   01..09..107..00..00
MIRROR POOL:         None
```

Figure 11 Sample output message from lspv command

STEP 17 If there is a dedicated dump device on the UNVSA being serviced, then the dump device must be moved to a good UNVSA before completing STEP 18.

See "Managing System Dump Devices" documentation at www.ibm.com, reference # T1000169. https://www-304.ibm.com/support/docview.wss?uid=isg3T1000169

STEP 18 Remove the failed drive from the mirror. We will assume for this procedure that hdisk1 has failed.

Command: `unmirrorvg <volume group name> <old disk name>`

STEP 19 If the PV failure occurred on rootvg, remove hdisk0 from the boot list by running the following command:

Note: If your configuration uses boot devices, add them to the command syntax.   Be sure to leave out the name of the old disk.

Command: `bootlist –om normal <boot device names>`

STEP 20 Undefine the failed physical volume by entering the commands:

Command: a) `reducevg –df <volume group name> <old disk name>`

b) `rmdev –dl <old disk name>`

STEP 21 Check what disk is the most recently used boot device.

Command: `bootinfo –b`

If the failed disk is the most recently used boot device, change the link from the failed disk to the working disk.

Command: `ln /dev/r<working disk name> /dev/ipldevice`

Note: the 'r' prefix to the PV name

Verify the /dev/ipldevice hard link has been created

Command: `ls /dev/ipldevice`

If the system is running AIX, skip to STEP 27.

RHEL6: (STEPs 22 through 26 are done from Service Node)

STEP 22   Get list of RAID devices and associated physical device names.

Command: `cat /proc/mdstat`

Example output:

```
Personalities : [raid1]
md2 : active raid1 dm-6[0] dm-11[1](F)
      584796028 blocks super 1.1 [2/1] [U_]
      bitmap: 1/1 pages [64KB], 65536KB chunk

md1 : active raid1 dm-3[0] dm-8[1]
      1048568 blocks super 1.1 [2/2] [UU]

md0 : active raid1 dm-4[0] dm-9[1]
      204788 blocks super 1.0 [2/2] [UU]

unused devices: <none>
```

This example output shows 3 RAID partition names: md0, md1, and md2. In this example, dm-11 in md2 is marked as faulty; so the script run in section 3.1 (STEP 1) either listed dm-11 in the other_names list for the failing drive or the failing drive was listed as "phy not attached to end device" and dm-11 was not listed. In this example, dm-8, dm-9, and dm-11 are all associated with the failing drive. However, RAID devices md0 and md1 did not automatically detect that the drive has failed.

STEP 23   Make sure all references to the failing drive are marked as faulty (F) in all RAID devices

Command: `mdadm --manage /dev/mdX --fail /dev/dm-X`

Note: mdX is md0, md1, and md2. The command only needs to be issued for partitions that are not already marked faulty (F). The dm-X values come from STEP 22 and the output of the script in STEP 1. The dm-X values in STEP 22 show the relationship between the mdX and appropriate dm-X partitions. The output of the script in STEP 1 shows which dm-X values are associated with the drive that will and will not be replaced.

Example: Based on the example above, the following 2 commands would be run

```
mdadm --manage /dev/md0 --fail /dev/dm-9
mdadm --manage /dev/md1 --fail /dev/dm-8
```

Below is an example of the first command being run with its output:

```
# mdadm --manage /dev/md0 --fail /dev/dm-9
mdadm: set /dev/dm-9 faulty in /dev/md0
```

STEP 24   Verify all RAID references to device are marked faulty (F)

Command: `cat /proc/mdstat`

Example output:

```
Personalities : [raid1]
md2 : active raid1 dm-6[0] dm-11[1](F)
      584796028 blocks super 1.1 [2/1] [U_]
      bitmap: 1/1 pages [64KB], 65536KB chunk

md1 : active raid1 dm-3[0] dm-8[1](F)
      1048568 blocks super 1.1 [2/1] [U_]

md0 : active raid1 dm-4[0] dm-9[1](F)
      204788 blocks super 1.0 [2/1] [U_]

unused devices: <none>
```

If all references to failed drive are marked faulty (F), continue with STEP 25; otherwise, retry STEP 23 verifying that the device names used are correct.

STEP 25   Remove failed drive from RAID devices

Command: `mdadm --manage /dev/mdX --remove /dev/dm-X`

Example: Based on the example above, the following 3 commands would be run

```
mdadm --manage /dev/md0 --remove/dev/dm-9
mdadm --manage /dev/md1 --remove/dev/dm-8
mdadm --manage /dev/md2 --remove/dev/dm-11
```

Below is an example of the first command being run with its output:

```
# mdadm --manage /dev/md0 --remove /dev/dm-9
mdadm: hot removed /dev/dm-9 from /dev/md0
```

STEP 26   Verify RAID devices running off one physical drive (mirroring is disabled)

Command: `cat /proc/mdstat`

Example output:

```
Personalities : [raid1]
md2 : active raid1 dm-6[0]
      584796028 blocks super 1.1 [2/1] [U_]
      bitmap: 1/1 pages [64KB], 65536KB chunk

md1 : active raid1 dm-3[0]
      1048568 blocks super 1.1 [2/1] [U_]

md0 : active raid1 dm-4[0]
      204788 blocks super 1.0 [2/1] [U_]

unused devices: <none>
```

If all RAID devices (mdX) show only a single device (dm-X), the remove commands were successful.

3.5 Physical Replacement of UNVSA and SAS Jumper Cable <= SSR TASK

STEP 27 Return to the HMC and place a checkmark in the Select column of the CEC Drawer on which you want to power off the UNVSA requiring service.

STEP 28 From the **Tasks** menu *select* **Serviceability -> Hardware -> Power On/Off Unit**. See Figure 12.
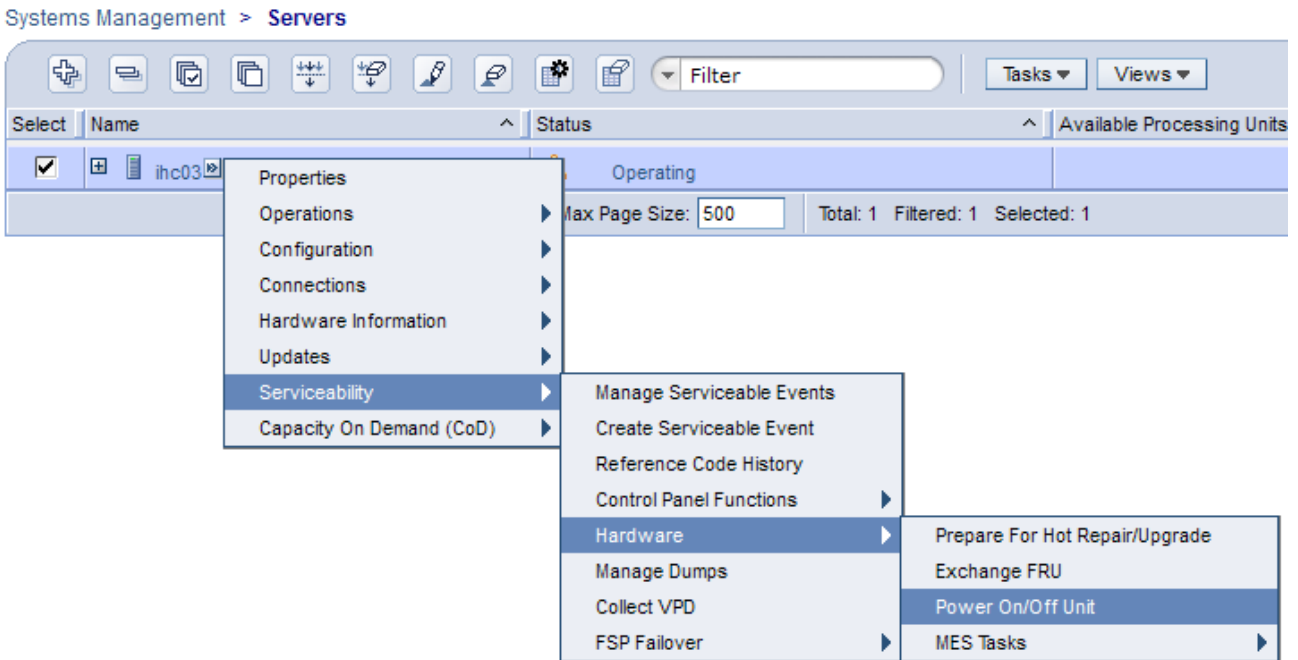
Figure 12 Power On/Off Unit

STEP 29 Power Off the UNVSA as identified in STEP 10. *Click* the CEC Drawer triangle (to expand and see busses), *click* the bus that corresponds to the UNVSA requiring service (triangle to expand), *select* the slot number that requires service, and *click* on the Power Off button. In the example in Figure 13, we show service of location code C16. C16 can be found under Bus 512; for Power Off of the C14 location code you can find this under Bus 520. See Figure 13 for example of selecting C16 location.

Figure 13 UNVSA Location Power Off/On

STEP 30 Use HMC menu's to power off slot. The HMC will go through several windows before the slot is powered off. If all conditions are met and the slot is powered off, a screen like Figure 14 is displayed. *Click* the **Next** button to get back to the Power On/Off Unit screen (Figure 13).
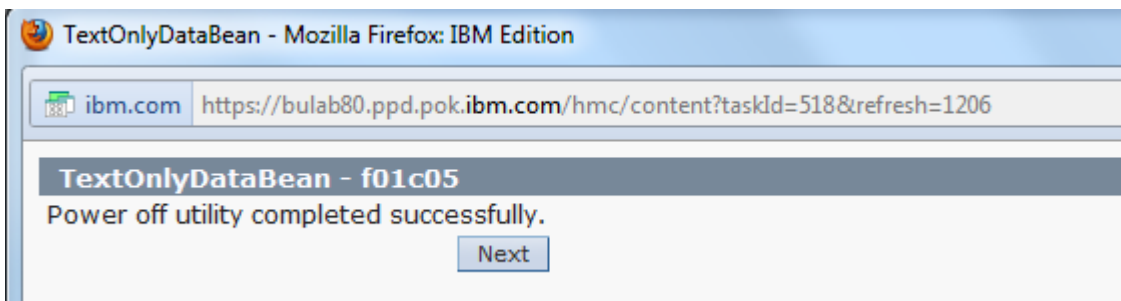
Figure 14 UNVSA Power Off Complete

STEP 31   Remove the Utility Non-Volatile Storage Assembly and SAS Jumper Cable – note use the location as identified in STEP 10, screen captures in this document are for reference only. Only remove the SAS Jumper Cable going to the drive identified in STEP 10 to be replaced. Removing the other SAS Jumper Cable can cause the Service Node to crash. See Figure 15.

   a.   Disconnect the SAS Jumper cable from the Utility Non-Volatile Storage Assembly being serviced and from the SAS adapter(typically located in C15).
   b.   Push the clasp(A) to release the black locking lever
   c.   Rotate the black locking lever(B) in the direction shown to unseat the UNVSA
   d.   Pull the UNVSA straight out of the slot
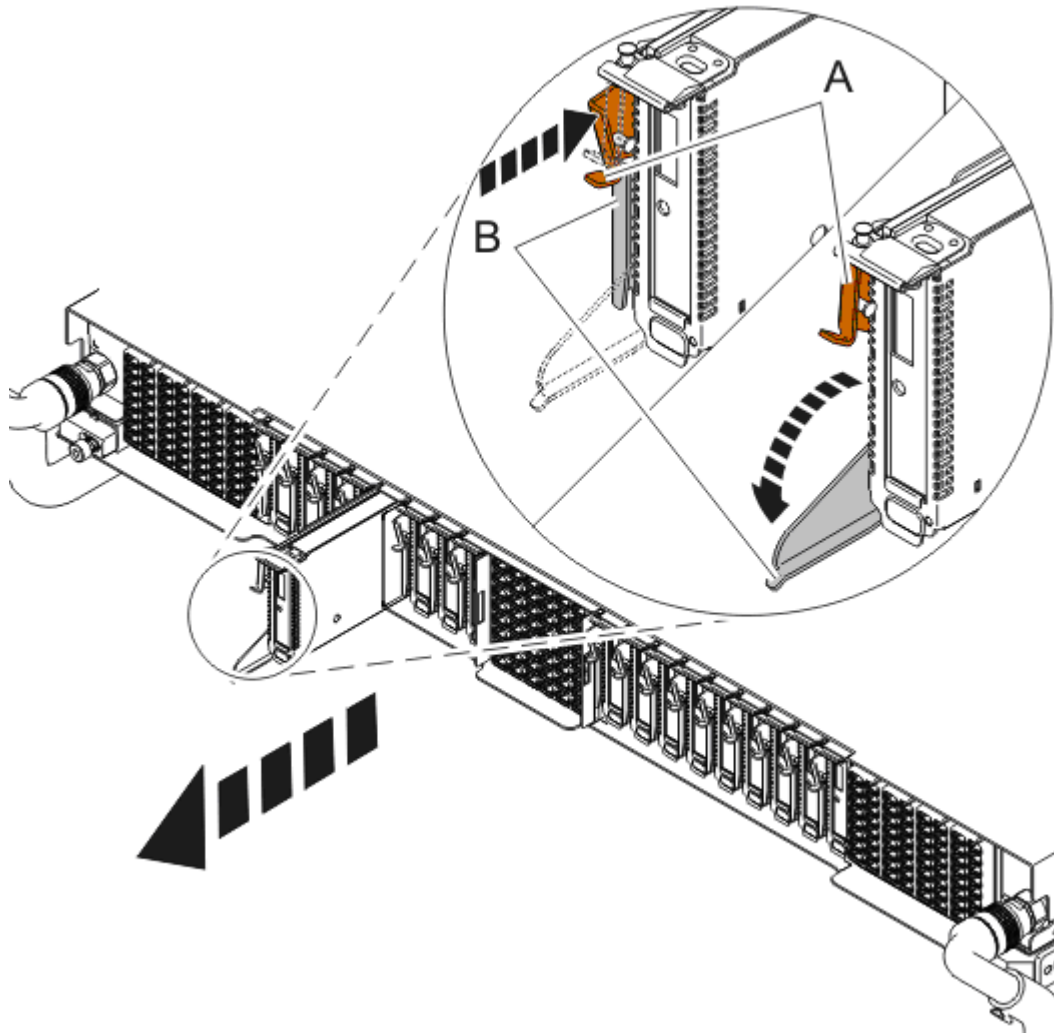
Figure 15 Removal of Utility Non-Volatile Storage Assembly from slot

STEP 32   Install the replacement Utility Non-Volatile Storage Assembly and SAS Jumper Cable. See Figure 16 for UNVSA installation.

    a. Align the guide pins on bottom of UNVSA with the PCIe slot as shown in the lower left inset of Figure 16. Bottom of UNVSA shown for reference to guide pins only.
    b. Push the UNVSA straight into the slot
    c. Rotate the black locking lever(A) in the direction shown to seat the UNVSA.
           Note: Make sure the clasp(B) closes over the locking lever.
    d. Connect the SAS cable from the UNVSA to the SAS adapter(the reverse of STEP 31 a)
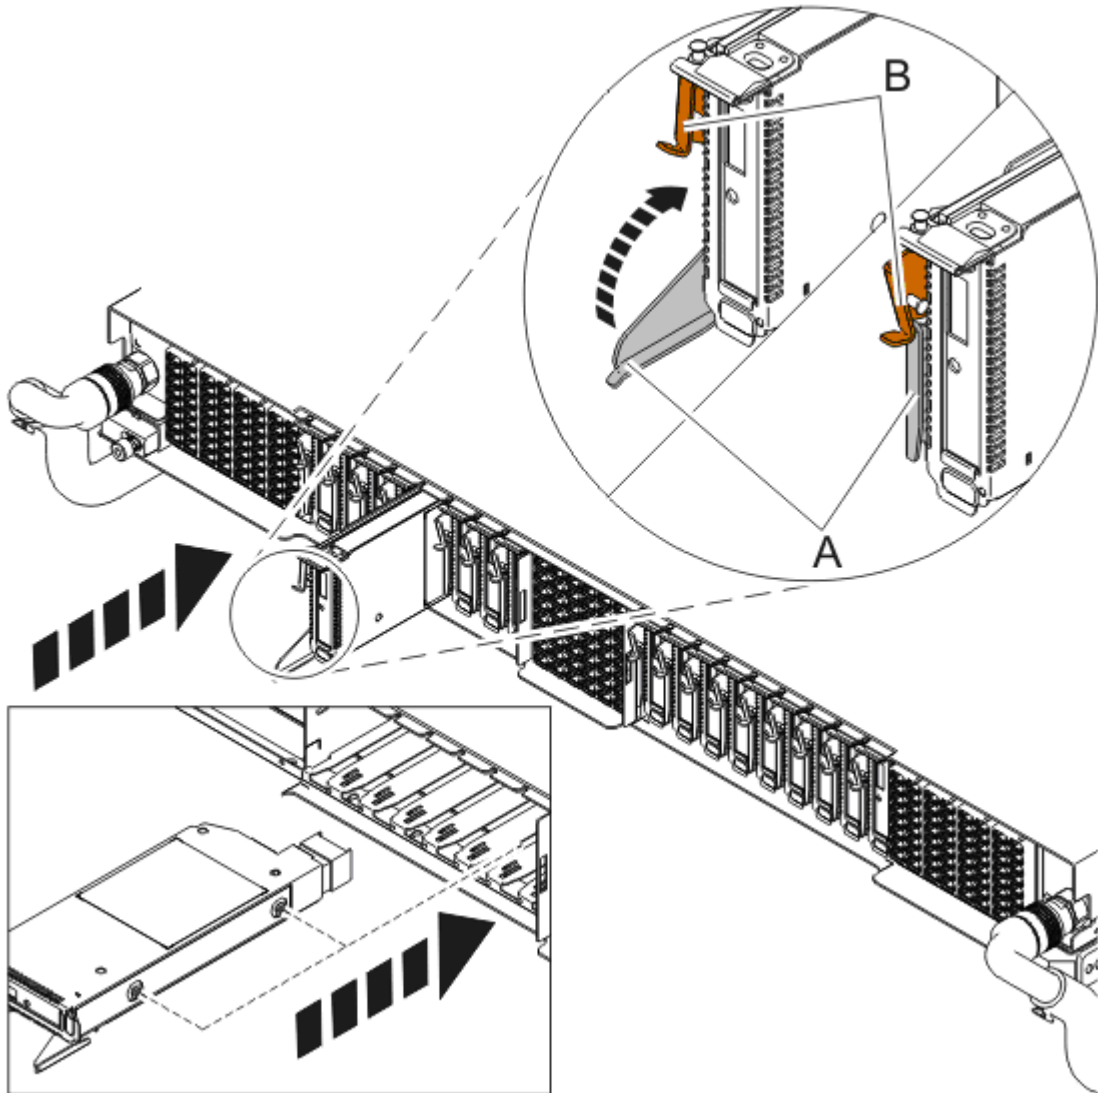
Figure 16 Install Utility Non-Volatile Storage Assembly

STEP 33   Return to the HMC screen for Power On/Off Unit. Power On the UNVSA as identified in STEP 10. Select the slot that was powered off in STEP 29 and *click* the **Power On** button. See Figure 13.
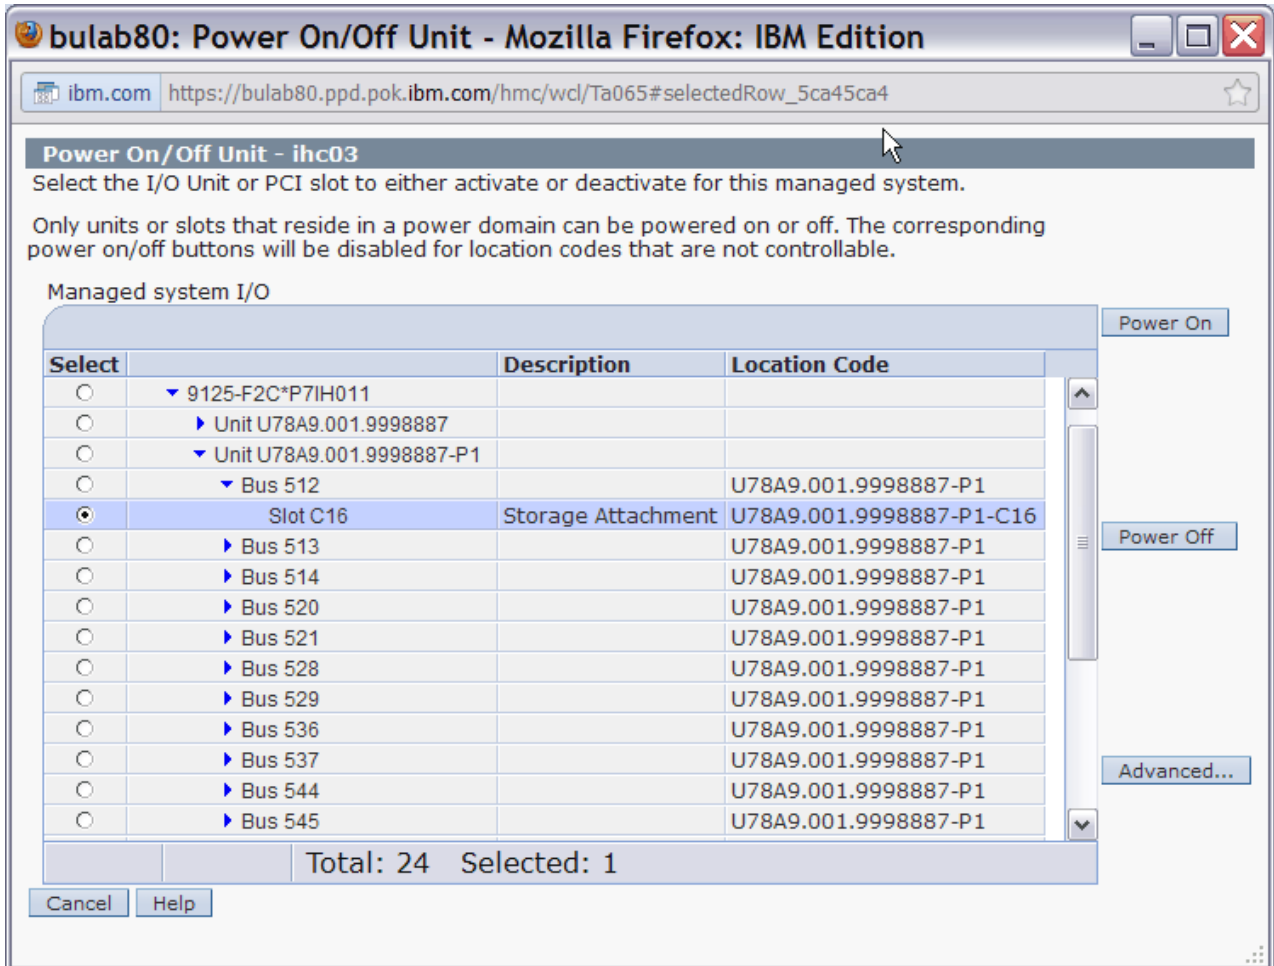
STEP 34 Use HMC menu's to power on slot. The HMC will go through several windows before the slot is powered on. If all conditions are met and the slot is powered on, a screen will be presented that says "Power on utility completed successfully". *Click* the **Next** button.

STEP 35 Verify that green LED indicating Power On for UNVSA slot has turned on solid. If it does not turn on solid it is possible that the UNVSA is not properly seated(Complete STEP 29 to STEP 34 again, if green LED does not turn on contact next level of support).

STEP 36 Click Cancel button on Power On/Off Unit

3.6 Add the UNVSA to rootvg <= Customer TASK

If the system is running RHEL6, skip to STEP 40.

AIX:

STEP 37 To define the new Physical Volume, run the following command:

Command: `cfgmgr`

STEP 38 To add the new PV to the volume group, run the following command:

Command: `extendvg rootvg <new disk name>`

NOTE: There is a force option ("-f") for the extendvg command that may be used if necessary.

STEP 39 Verify hdisk1 is now in rootvg, run the following command, see Figure 17 for sample output:

Command: `lsvg -p rootvg`

```
# lsvg -p rootvg
rootvg:
PV_NAME          PV STATE          TOTAL PPs    FREE PPs    FREE DISTRIBUTION
hdisk0           active            558          163         19..17..00..15..112
hdisk1           active            558          166         19..20..00..15..112
```

Figure 17 lsvg sample output

If the system is running AIX, skip to STEP 46.

RHEL6:

STEP 40   Find device names of new drive. If the drive was not replaced, skip to STEP 40.

In many cases it will be the same /dev/sdX names that were used to represent the drive that was replaced. One what to verify this is to run the script defined in STEP 1 and look for the "drive_device_names" that returned "cat: /sys/devices/virtual/block//dm/uuid: No such file or directory" instead of a "multipath_dev_name". This is caused by the multipath code "blacklisting" the new drive. If both drives in this RAID1 group show up with the same number of "other_names", then continue with STEP 40.

STEP 41   Find SAS WWID for new drive

Command: `multipath -d -v 3 /dev/sdX`

Example output:

```
Aug 15 15:01:59 | sda: not found in pathvec
Aug 15 15:01:59 | sda: mask = 0x11
Aug 15 15:01:59 | sda: dev_t = 8:0
Aug 15 15:01:59 | sda: size = 1172123568
Aug 15 15:01:59 | sda: subsystem = scsi
Aug 15 15:01:59 | sda: vendor = IBM
Aug 15 15:01:59 | sda: product = MBF2600RC
Aug 15 15:01:59 | sda: rev = 610A
Aug 15 15:01:59 | sda: h:b:t:l = 0:0:4:0
Aug 15 15:01:59 | sda: getuid = /lib/udev/scsi_id --whitelisted --
device=/dev/%n (config file default)
Aug 15 15:01:59 | sda: uid = 350000393681bfdf4 (callout)
Aug 15 15:01:59 | scope limited to 350000393681bfdf4
Aug 15 15:01:59 | (null): (350000393681bfdf4) wwid blacklisted
```

The key item is the SAS WWID. In this example it is "350000393681bfdf4". This value needs to be copied into the /etc/multipath.conf file. So, copy it down.

STEP 42   Change /etc/multipath.conf adding the new drive WWID

Edit the /etc/multipath.conf file adding lines for the new drive WWID in blacklist_exceptions list and in the multipaths list. Note that some versions of xCAT include the "alias mpathX" line and some do not. The recommendation is to copy what is there for the other drives and change the WWID.

Example of (part of) the /etc/multipath.conf prior to changing:

```
blacklist_exceptions {
    wwid "35000c5001db6f2b7"
    wwid "3500003930801b79c"
} multipaths {
        multipath {
                uid 0
                alias mpatha
                gid 0
                wwid "35000c5001db6f2b7"
                mode 0600
        }
        multipath {
                uid 0
                alias mpathb
                gid 0
                wwid "3500003930801b79c"
                mode 0600
        }
}
```

Example of (part of) the /etc/multipath.conf after changing:

```
blacklist_exceptions {
    wwid "35000c5001db6f2b7"
    wwid "3500003930801b79c"
    wwid "350000393681bfdf4"
}
multipaths {
        multipath {
                uid 0
                alias mpatha
                gid 0
                wwid "35000c5001db6f2b7"
                mode 0600
        }
        multipath {
                uid 0
                alias mpathb
                gid 0
                wwid "3500003930801b79c"
                mode 0600
        } multipath {
                uid 0
                alias mpathc
                gid 0
                wwid "350000393681bfdf4"
                mode 0600
        }
}
```

STEP 43   Replicate partitioning of remaining "good disk" on new disk.
Command: `sfdisk -d /dev/dm-Y | sfdisk /dev/sdX --force`

Example based on STEP 25 and STEP 38 examples:
```
sfdisk -d /dev/dm-0 | sfdisk /dev/sda --force
```

Note: if the cable is just being replaced, this command can be skipped (the partitioning data was not changed).

STEP 44   Verify partition of drives.
Command: `sfdisk -d /dev/sdX`
> Note: run this on both disks and verify both have the same number of partitions and each partition has the same size.

Example output:

```
# sfdisk -d /dev/sda
Warning: extended partition does not start at a cylinder boundary.
DOS and Linux will interpret the contents differently.
# partition table of /dev/sda
unit: sectors

/dev/sdc1 : start=      2048, size=     16384, Id=41, bootable
/dev/sdc2 : start=     18432, size=  2097152, Id=fd
/dev/sdc3 : start=  2115584, size=   409600, Id=fd
/dev/sdc4 : start=  2525184, size=1169596416, Id= 5
/dev/sdc5 : start=  2527232, size=1169594368, Id=fd
```

STEP 45   Create multipath device names for new drive

Command: `multipath -r`

Example output:

```
# multipath -r
create: mpathc (350000393681bfdf4) undef IBM,MBF2600RC
size=559G features='0' hwhandler='0' wp=undef
|-+- policy='round-robin 0' prio=1 status=undef
| `- 0:0:4:0 sda 8:0  undef ready running
`-+- policy='round-robin 0' prio=1 status=undef
  `- 0:0:5:0 sdb 8:16 undef ready running
reload: mpathb (350003930801b79c) undef IBM,MBF2600RC
size=559G features='0' hwhandler='0' wp=undef
|-+- policy='round-robin 0' prio=1 status=undef
| `- 0:0:2:0 sdc 8:32 active ready running
`-+- policy='round-robin 0' prio=1 status=undef
  `- 0:0:3:0 sdd 8:48 active ready running
```

To verify the multipath device and other names for the drive have not changed from the drive that was removed, run the script defined in STEP 1.

3.7 Mirror the Utility Non-Volatile Storage Assemblies <= Customer TASK

If the system is running RHEL6, skip to STEP 52.

AIX:

STEP 46   To build a Mirror between UNVSA devices, run the following command:

Note: The Mirroring a Volume Group may take a significant amount of time before completing because of complex error checking, the amount of logical volumes to mirror in a volume group, and the time is takes to synchronize the new mirrored logical volumes.

Command: `mirrorvg -m rootvg <new disk name>`

See Figure 18 for sample initial output from mirrorvg command, see Figure 19 for sample output after mirrorvg is completed.

```
(0) root @ f01c03sn01: 7.1.0.0: /
# mirrorvg -m rootvg hdisk1
0516-1804 chvg: The quorum change takes effect immediately.
```

Figure 18 Quorum change

```
(0) root @ f01c03sn01: 7.1.0.0: /
# mirrorvg -m rootvg hdisk1
0516-1804 chvg: The quorum change takes effect immediately.
0516-1126 mirrorvg: rootvg successfully mirrored, user should perform
        bosboot of system to initialize boot records.  Then, user must modify
        bootlist to include:  hdisk1 hdisk0.
```

Figure 19 Mirror Volume Group Complete

STEP 47   Initialize the boot record by running the following command, see Figure 20 for sample output:

Command: `bosboot -a`

```
(0) root @ f01c03sn01: 7.1.0.0: /
# bosboot -a

A previous bosdebug command has changed characteristics of this
        boot image.  Use bosdebug -L to display what these changes are.

bosboot: Boot image is 47423 512 byte blocks.
```

Figure 20 bosboot –a sample output

STEP 48   If the PV failure occurred on rootvg, update the boot list by running the following command:

Note: If your configuration uses boot devices other than your new disk, add them to the command.

Command: `bootlist -om normal <new disk name>`

STEP 49   Verify that the procedure was successful.

* To verify that all logical volumes are mirrored to the new PV, run the following command:

Command: `lslv <logical volume name>`

* To verify that all the logical volume partitions are synchronized, check that there are no stale partitions by running the following command:

Command: `lspv <new disk name>`

* If there are stale partitions, synchronize

Command: `syncvg`

* To verify the boot list, run the following command:

Command: `bootlist -m normal -o`

STEP 50   If newly installed disk is to be the boot device for the Service Node run the following command:

Command: `ln /dev/r<new disk name> /dev/ipldevice`

STEP 51   Verify the /dev/ipldevice hard link has been created by STEP 50

Command: `ls /dev/ipldevice`

If the system is running AIX, skip to STEP 57.

RHEL6:

STEP 52   Create bootstrap image in partition 1 of new drive.

Command: `mkofboot -b /dev/dm-X`

Example: `mkofboot -b /dev/dm-6`

`mkofboot: Overwrite contents of /dev/dm-7 with /usr/lib/yaboot/yaboot? [y/N] y`

Note: answer "y" to question

Note: Partition 1 is listed in the script defined in STEP 1 in the "other_names" list as the second entry. The first entry is the multipath device name for the entire drive.

Note: If mkofboot is not installed, the yaboot*.rpm will need to be installed. It is part of the Red Hat install.

STEP 53   Update initramfs file to contain information on new drive. This is done to get the new drive to be seen as a multipath device after a reboot.

Commands: `mv /boot/initramfs-$(uname -r).img /boot/initramfs-$(uname -r).old`
           `mkinitrd /boot/initramfs-$(uname -r).img $(uname -r)`

STEP 54   Add new disk into the RAID1 array.

Command: `mdadm --manage /dev/mdX --add /dev/dm-X`

Example based on STEP 40 commands:

```
mdadm --manage /dev/md0 --add /dev/dm-9
mdadm --manage /dev/md1 --add /dev/dm-8
mdadm --manage /dev/md2 --add /dev/dm-11
```

Notes: To get the correct values for the devices (/dev/mdX and /dev/dm-X), run "`cat /proc/mdstat`" to get the RAID device names (/dev/mdX) and the multipath partition names from the good drive. Run the script defined in STEP 1 to list the partition names for both drives (listed under "other_names"). Add the partition name for the new drive that is in the same position as the partition name for the good drive.

STEP 55   Update boot list order using 4 current drive device names

Running the script defined in STEP 1 will list the 4 drive paths with 4 unique drive_device_name values. For the bootlist command, only the sdX part of the name is required. Normally these values will be sda, sdb, sdc, and sdd (with sda and sdb associated with one drive; sdc and sdd associated with the other drive). The recommendation is to make sure all 4 paths are in the bootlist. The order is not critical.

Command: `bootlist -r -m normal sda sdc sdd sdb`

Example output:

```
# bootlist -r -m normal sda sdc sdd sdb
/pci@800000020000201/pci1000,72@0/sas/disk@50000393681bfdf2
/pci@800000020000201/pci1000,72@0/sas/disk@500003930801b79f
/pci@800000020000201/pci1000,72@0/sas/disk@500003930801b79e
/pci@800000020000201/pci1000,72@0/sas/disk@50000393681bfdf3
```

STEP 56   Verify reconstruction

Command: `cat /proc/mdstat`

Example output:

```
Personalities : [raid1]
md2 : active raid1 dm-11[1] dm-6[0]
      584796028 blocks super 1.1 [2/1] [U_]
      [===>................]  recovery = 18.7% (109379840/584796028)
finish=100.7min speed=78660K/sec
      bitmap: 1/1 pages [64KB], 65536KB chunk

md1 : active raid1 dm-8[1] dm-3[0]
      1048568 blocks super 1.1 [2/2] [UU]

md0 : active raid1 dm-9[1] dm-4[0]
      204788 blocks super 1.0 [2/2] [UU]

unused devices: <none>
```

This example output shows that md0 and md1 are reconstructed and md2 is 18.7% completed. Once md2 reconstruction is completed, it will look like md0 and md1.

3.8 Clean Up <= SSR Task

STEP 57   Verify with the customer that the new UNVSA and SAS jumper cable has fixed the problem.

STEP 58   In the window **Identify LED, Select Location**, *click* the **Deactivate LED** button and then *click* the Cancel button when it returns to the screen below.
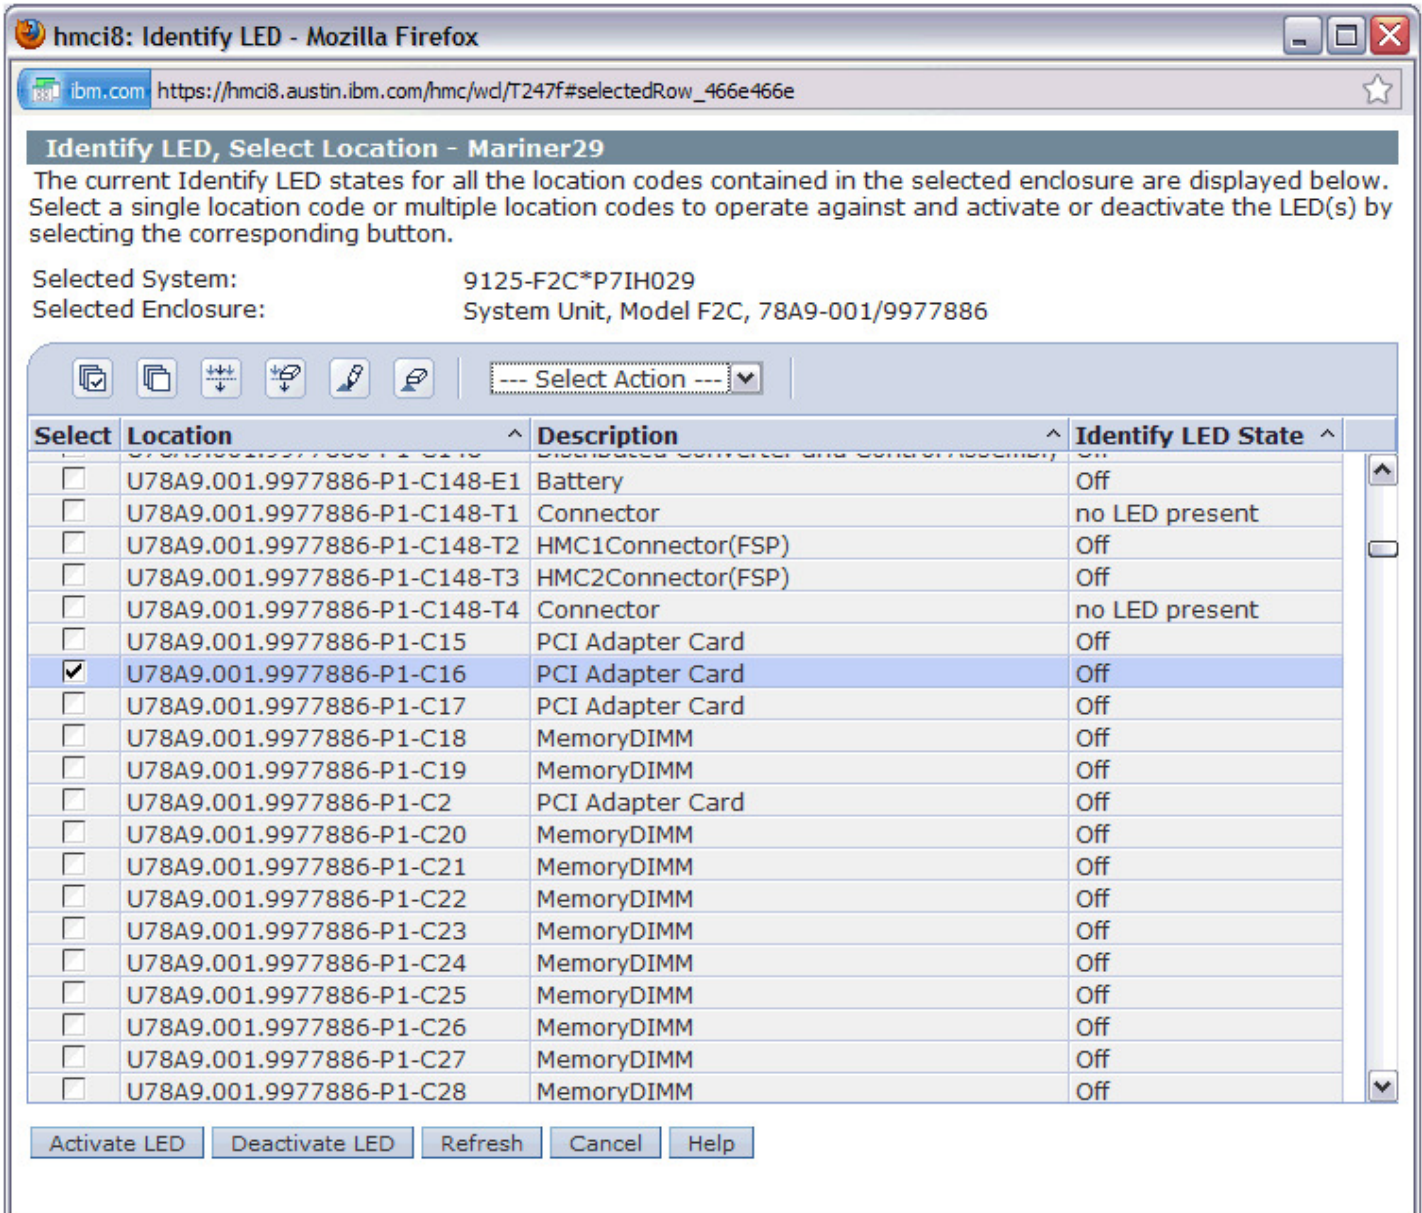

Figure 21 UNVSA Identify LED Selection

STEP 59   In the window titled **Identify LED, Select Enclosure** , *click* the **Deactivate LED** button (see **Figure 22**) and then click Cancel when it returns to the window titled **Identify LED, Select Enclosure**.
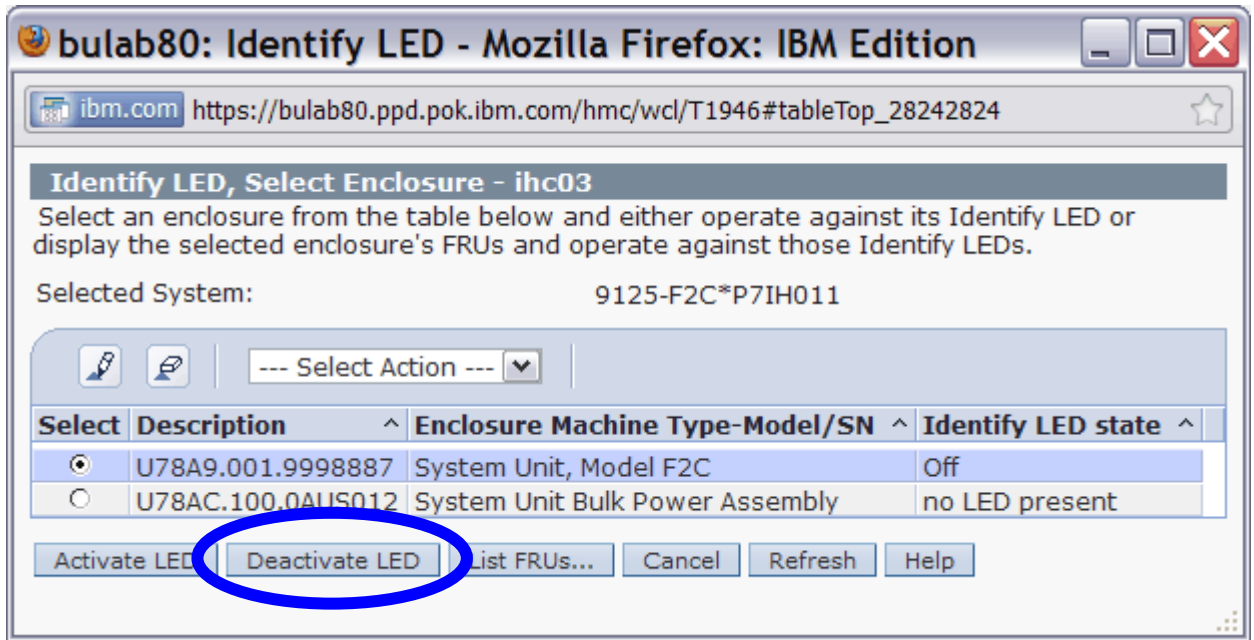
Figure 22 System Unit, Model F2C Selection

STEP 60   Verify all identify LEDs are Deactivated and close the front and rear doors.

4   END OF POWER775 UNVSA AND SAS JUMPER CABLE SERVICE PROCEDURE