



IBM Advanced Technical Sales Proof of Concept:

Power 7 Servers Live Partition Mobility Using FlashSystems Storage

IBM Power Advanced Technical Sales
IBM FlashSystems Center of Competency

Table of Contents

I.	Executive Summary	2
	Highlights	
	Objectives	
	Mission Statement	
	Keys to Success	
II.	Description of Proof of Concept	3
	Architecture	
III.	Proof of Concept Implementation	5
	Execution	
	Confirm Attached Disks on Virtual Server are IBM FlashSystem Disks	
	Confirm AIX & VIOS Levels	
	Confirm DB2 is UP and Running with an Active Database	
	Confirm IBM ISD is Up and Running	
	Starting ISD Health Monitor Dashboard to Monitor CPU & Adapter Speeds	
	Confirm Source Virtual Server Serial Number Prior to Migration	
	Start Active Migration of Virtual Server "m1aix3" from Power 7 Server "m1p740" to "m2p740" (GUI or CLI)	
	During Active Migration, Execute Inventory Collection on ISD to Drive up Workload	
	View Status of Inventory Collection as Migration Runs	
	Connect to "DIRDB01" as Migration Runs	
	Confirm Serial Number Changes Following Successful Migration	
	Confirm IBM ISD & DB2 Access (GUI & CLI) Following Successful Active Migration	
	Reboot Virtual Server "m1aix3" on Target System "m2p740"	
	Confirm Both IBM ISD & DB2 Started Correctly Following Virtual Server Reboot	
	Conclusion	
IV.	Appendix	22
	Acknowledgements	
	Trademarks	
	POC Reference Material	

Executive Summary


The purpose of this proof of concept (PoC) is to demonstrate the ability of Power 7 servers to perform active Live Partition Mobility (LPM) operations while utilizing IBM FlashSystems storage. This PoC is designed to highlight the interoperability between Power 7 servers and IBM FlashSystems storage for those clients considering integrating this next generation storage into their Power 7 environments. This PoC was executed on October 10, 2013 at the IBM FlashSystems Center of Competency in Coppel, TX. Below you will find the details of the activities performed that day.

Highlights

Successful execution of one or more Live Partition Mobility (LPM) operations between two Power 7 servers using IBM FlashSystems storage.

Migration of active virtual server running IBM Systems Director & DB2 database software executing workload.

Migration of virtual server took under 5 minutes.



Objectives

Demonstrate the active migration of a Power 7 virtual server using IBM FlashSystems

Demonstrate the active migration of an AIX LPAR while actively processing IBM Systems Director & DB2 workloads.

Demonstrate virtual server stability after LPM by rebooting the virtual server and confirming virtual server & applications start back up.

Mission Statement

Highlight interoperability and synergy between IBM Power and IBM FlashSystems by performing an active migration of a virtual machine using FlashSystems storage.

Keys to Success

Successful active LPM of a running virtual server using IBM FlashSystems storage

Ability for enterprise-class applications IBM Systems Director and IBM DB2 to remaining running and processing workload during the migration

Virtual server reboot and application start-up success

Description of Proof of Concept

The Proof of Concept (POC) will demonstrate the ability to actively migrate a virtual server running enterprise-class software between a pair of Power 7 servers while actively processing workload.

Architecture

The environment used for this PoC is illustrated below in Figure 1:

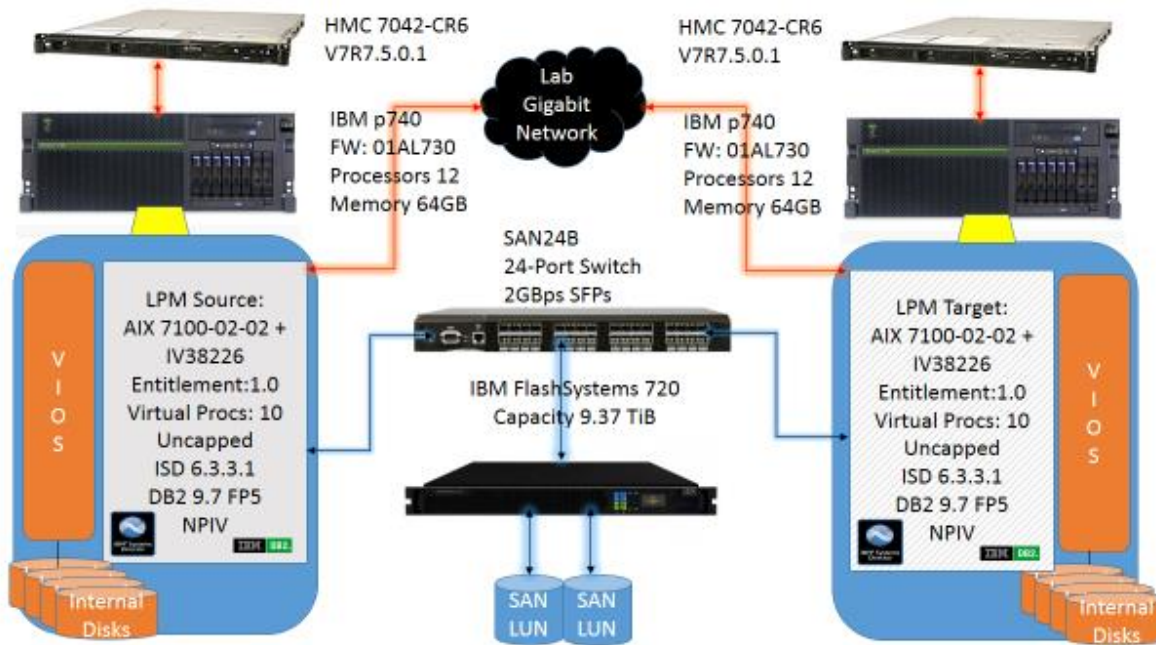


Figure 1. Hardware and Software Used for AIX on FlashSystems LPM

Implementation of Proof of Concept

Execution

The LPM PoC will execute the following procedures:

1. Confirm attached disks on the virtual server are FlashSystem disks
2. Confirm AIX and VIOS operating system levels
3. Confirm IBM DB2 is up and running with an active database
4. Confirm IBM ISD is up and running and connected to its DB2 hosted database
5. Start Health Summary Dashboard in ISD to monitor CPU % and Network Recv/Send
6. Confirm source virtual server serial number prior to migration
7. Start active migration of virtual server 'm1aix3' from Power 7 server 'm1p740' to 'm2p740' (HMC GUI or CLI)
8. During migration, execute inventory collection on ISD to drive up workload
9. View status of inventory collection as migration runs
10. Connect to DIRBD01 DB as migration runs
11. After migration completes, confirm serial number has changed
12. Confirm ISD and DB2 could still be accessed (GUI & CLI)
13. Reboot virtual server and confirm it comes back up successfully
14. Check status of IBM ISD & DB2

Confirm Attached Disks on Virtual Server are Flash Systems Disks



For **AIX 61 TL8 SP3** and **AIX 71 TL2 SP3**, or newer release, you don't need to do any additional config for FlashSystem, the new logical unit (LUN) can be recognized by MPIO without any problem.

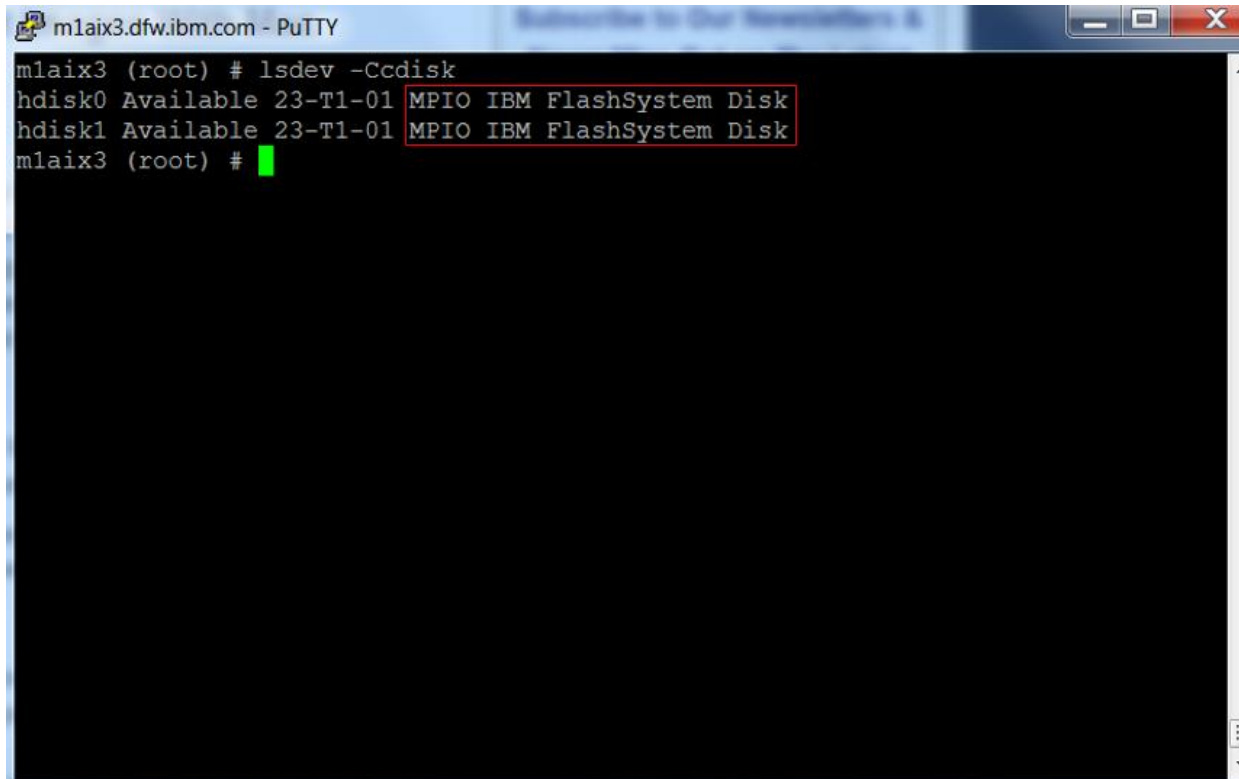
For **AIX 61 TL8 SP2** and **AIX 71 TL2 SP2**, you will need an APAR to properly set the device type and queue depth. The APAR description and fix locations are:

FC attached IBM FlashSystem Storage is recognized by AIX as "Other FC Disk Drive" and the disk is configured as non-MPIO and has a queue_depth of "1".

APAR for **AIX 61 TL8 SP2**: <http://www-01.ibm.com/support/docview.wss?uid=isg1IV38191>

APAR for **AIX 71 TL2 SP2**: <http://www-01.ibm.com/support/docview.wss?uid=isg1IV38226>

We confirmed that FlashSystems disks were being used by running the lsdev command as shown in Figure 2:

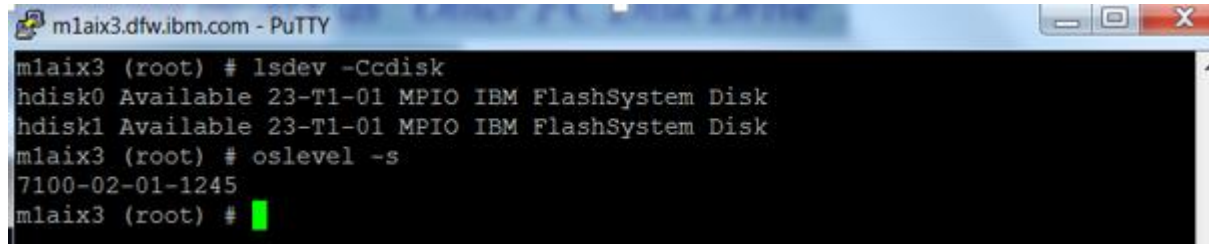
A terminal window titled 'mlaix3.dfw.ibm.com - PuTTY' showing the execution of the 'lsdev -Cdisk' command. The output lists two disks: 'hdisk0' and 'hdisk1', both available and identified as 'MPIO IBM FlashSystem Disk'. The text 'MPIO IBM FlashSystem Disk' is highlighted with a red box. The prompt 'mlaix3 (root) #' is visible at the end of the output.

```
mlaix3 (root) # lsdev -Cdisk
hdisk0 Available 23-T1-01 MPIO IBM FlashSystem Disk
hdisk1 Available 23-T1-01 MPIO IBM FlashSystem Disk
mlaix3 (root) #
```

Figure 2. Confirming FlashStorage disks are connected

Confirm AIX and VIOS Levels

Figure 3 below shows the “oslevel” command being run on the target AIX virtual server “m1aix3”:

A terminal window titled "m1aix3.dfw.ibm.com - PuTTY" showing the execution of the "oslevel" command. The prompt is "m1aix3 (root) #". The user enters "lsdev -Ccdisk" and the output is "hdisk0 Available 23-T1-01 MPIO IBM FlashSystem Disk" and "hdisk1 Available 23-T1-01 MPIO IBM FlashSystem Disk". Then the user enters "oslevel -s" and the output is "7100-02-01-1245". The prompt returns to "m1aix3 (root) #".

```
m1aix3 (root) # lsdev -Ccdisk
hdisk0 Available 23-T1-01 MPIO IBM FlashSystem Disk
hdisk1 Available 23-T1-01 MPIO IBM FlashSystem Disk
m1aix3 (root) # oslevel -s
7100-02-01-1245
m1aix3 (root) #
```

Figure 3. AIX Operating System Level Confirmation

Figure 4 below shows the “ioslevel” command being run against the source VIOS virtual server ‘m1demo9’:

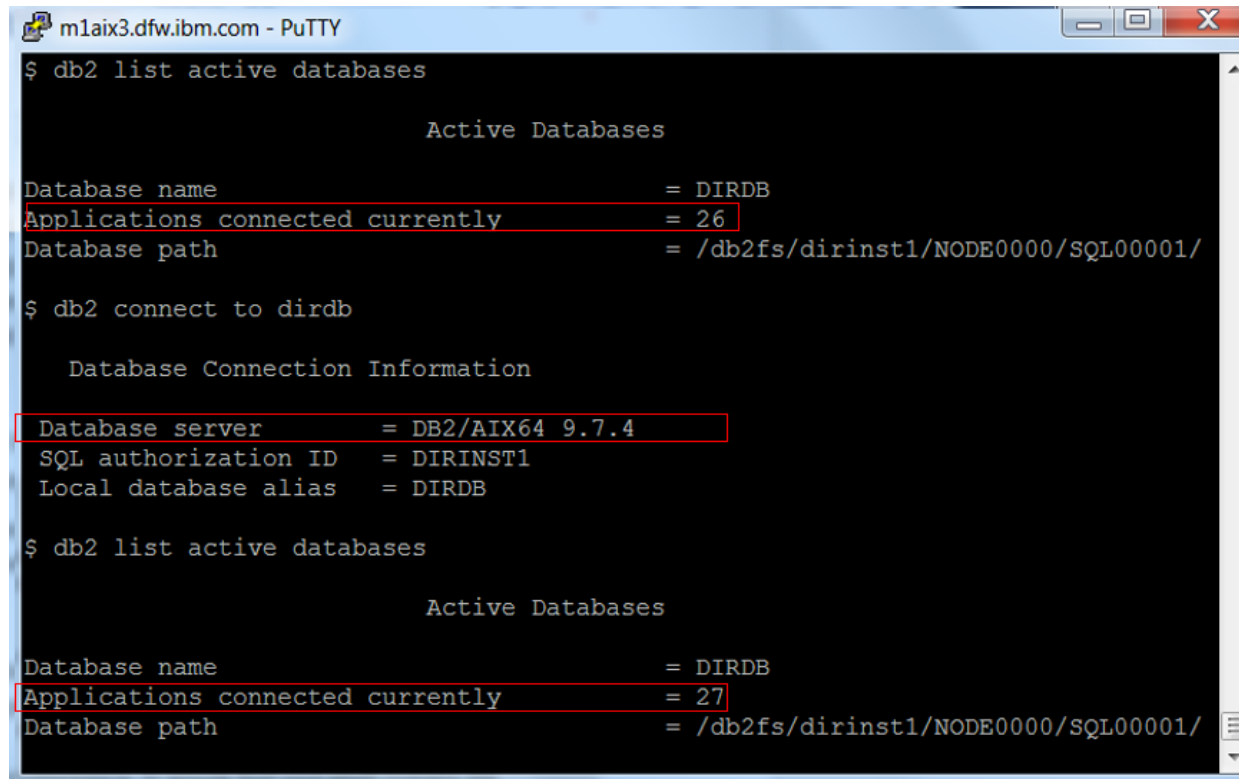
A terminal window titled "m1demo1.dfw.ibm.com - PuTTY" showing the execution of the "ioslevel" command. The prompt is "\$". The user enters "ioslevel" and the output is "2.2.2.2". The prompt returns to "\$".

```
$ ioslevel
2.2.2.2
$
```

Figure 4. VIOS Operating System Level Confirmation

Confirm IBM DB2 is Up and Running with an Active Database

Log on as DB2 instance owner on LPM source virtual server “m1aix3” and run the following commands to confirm that DB2 is up and running with an active database and applications connected to it as shown in Figure 5 below:



```
m1aix3.dfw.ibm.com - PuTTY
$ db2 list active databases

                Active Databases

Database name           = DIRDB
Applications connected  = 26
Database path           = /db2fs/dirinst1/NODE0000/SQL00001/

$ db2 connect to dirdb

Database Connection Information

Database server         = DB2/AIX64 9.7.4
SQL authorization ID   = DIRINST1
Local database alias   = DIRDB

$ db2 list active databases

                Active Databases

Database name           = DIRDB
Applications connected  = 27
Database path           = /db2fs/dirinst1/NODE0000/SQL00001/
```

Figure 5. Confirming Database Active & Accepting Connections

Confirm IBM ISD is Up and Running

As root user, run the “smstatus” command as shown in Figure 6. below to confirm the level & state of the ISD server:



```
m1aix3.dfw.ibm.com - PuTTY
m1aix3 (root) # smcli lsver
6.3.3.1
m1aix3 (root) # smstatus
Active
m1aix3 (root) #
```

Figure 6. Confirming ISD Level & Status

Starting Health Monitor Dashboard to Monitor CPU & Adapter Speeds

Logged into the IBM ISD web interface and configured monitors for the ‘m1aix3’ endpoint to allow ISD to monitor its operating system CPU utilization %, network adapter throughput, and Memory usage:

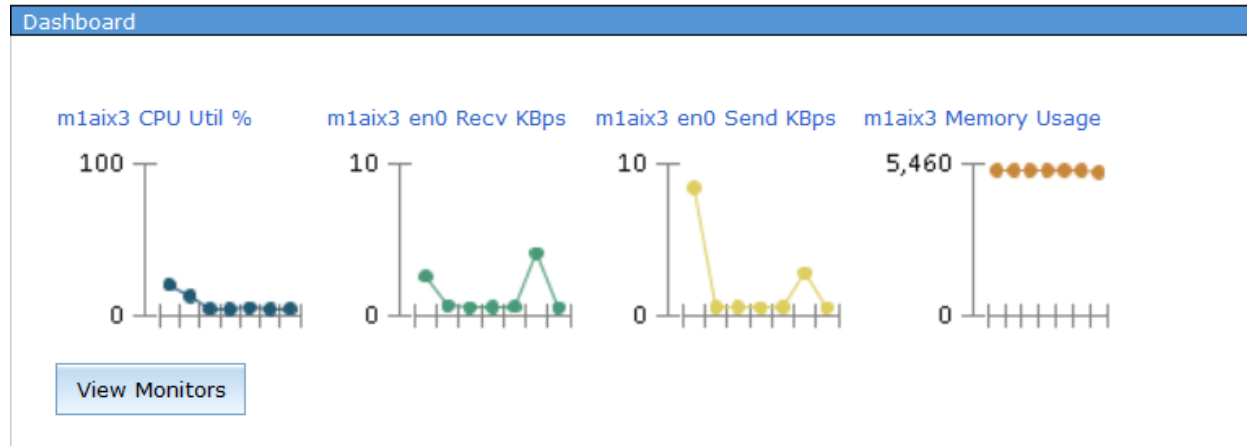
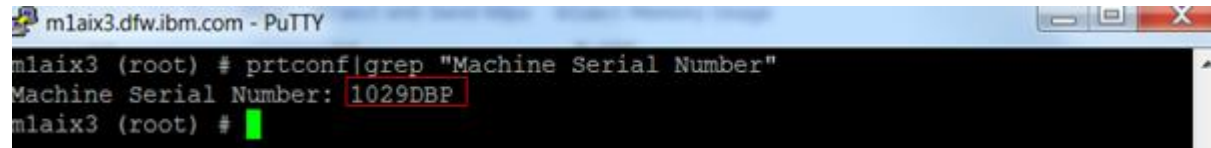


Figure 7. Starting ISD Health Monitor Dashboard

The above will be kept running during the active migration to demonstrate IBM ISD can continue to function and monitor metrics while its underlying virtual server is being migrated from m1p740 to m2p740.

Confirm Source Virtual Server Serial Number Prior to Migration

Before we begin the active migration we need to confirm the serial number on the source virtual server so that we may confirm if the machine serial number has in fact changed after the migration completes. This will allow us to confirm that the virtual server was in fact successfully migrated from system “m1p740” to “m2p740” as noted in Figure 8:



```
m1aix3.dfw.ibm.com - PuTTY
m1aix3 (root) # prtconf|grep "Machine Serial Number"
Machine Serial Number: 1029DBP
m1aix3 (root) #
```

Figure 8. Confirming Serial Number Change Prior to Migration

Start Active Migration of Virtual Server “m1aix3” from Power 7 Server “m1p740” to “m2p740” (GUI or CLI)

The migration was executed both through the HMC GUI and CLI. The CLI does provide a bit more flexibility in controlling elements of LPM operations as well as lends itself nicely to scripting. Below we show you both methods of execution starting with the GUI method in Figure 9:

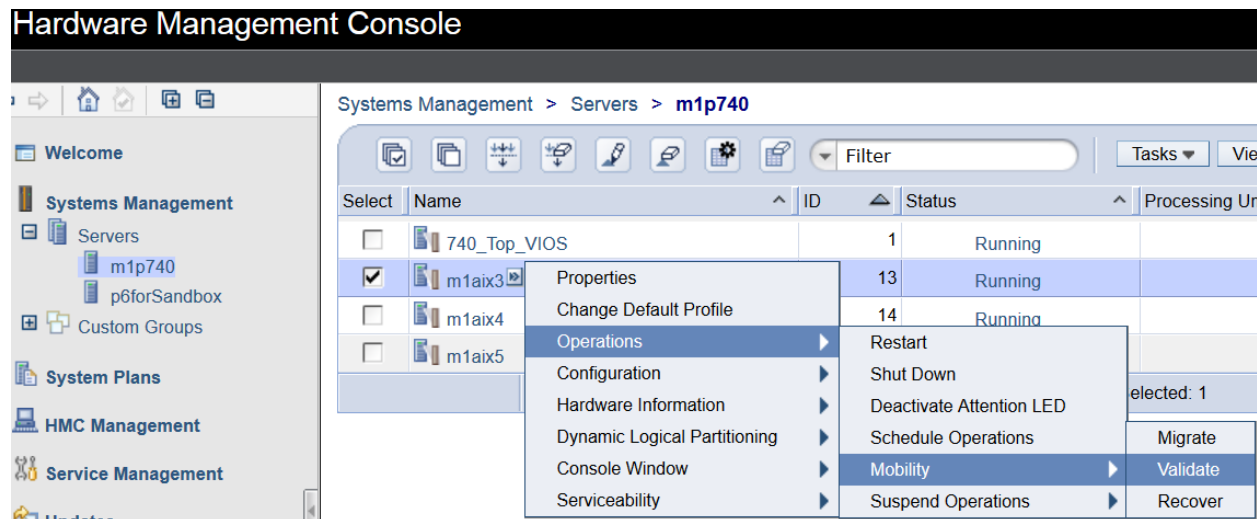


Figure 9. Executing LPM via HMC GUI

Figure 10. below shows the screen populated with the remote HMC, destination system, MSP pairing, and the virtual slot that will be assigned on the destination VIOS:

Partition Migration Validation - m1p740 - m1aix3

Fill in the following information to set up a migration of the partition to a different managed system. Click Validate to ensure that all requirements are met for this migration. You cannot migrate until the migration set up has been verified.

Source system : m1p740
Migrating partition: m1aix3
Remote HMC: m2demo9
Remote User: hscroot
Destination system: m2p740
Destination profile name: m1aix3_npiv
Destination shared processor pool: DefaultPool (0)
Source mover service partition: 740_Top_VIOS
Destination mover service partition: 740_Bottom_VIOS
Wait time (in min): 5
Override virtual network errors when possible:
Override virtual storage errors when possible:
Virtual Storage assignments :

Select	Source Slot ID	Slot Type	Destination VIOS
<input checked="" type="checkbox"/>	23	Fibre	740_Bottom_VIOS

Figure 10. LPM Validation Screen

We are now ready to execute the migration by clicking on the migrate button listed in Figure 10. above. We are then presented the Partition Migration Status window as noted in Figure 11. below:




Figure 11. Migration Status

As the migration progresses, we can view the ISD Health Summary Dashboard to confirm ISD is still working correctly and continues to monitor key performance metrics associated with the m1p740 frame and VIOS virtual:

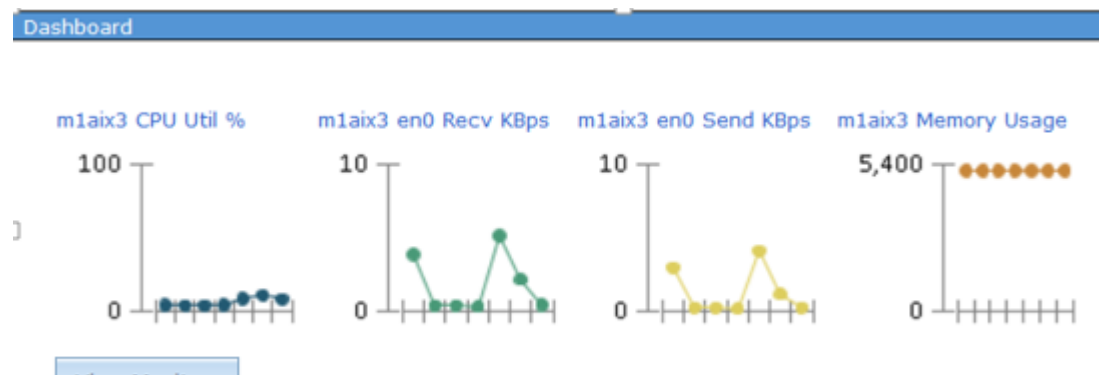


Figure 12. ISD Health Summary Dashboard

We continue to monitor the migration progress a few minutes later as noted below:

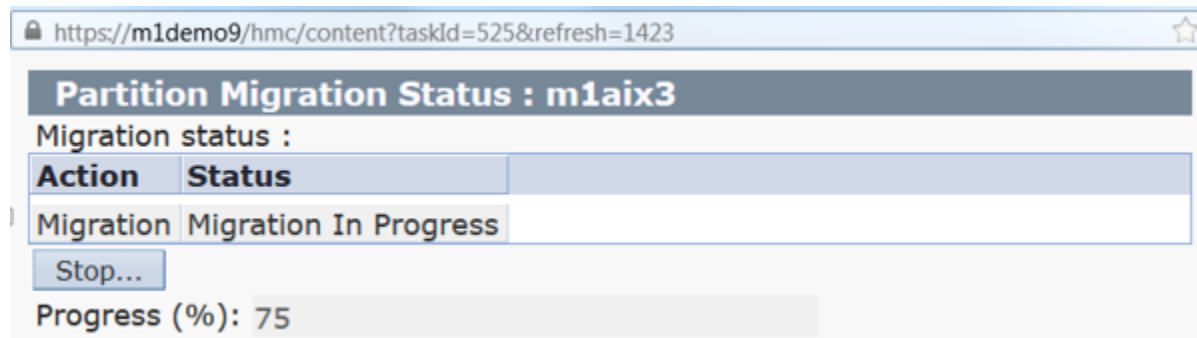


Figure 13. LPM Status Window

HMC CLI Active Migration Method

The above migration task was also performed via the HMC CLI by issuing the commands listed below to start the migration in the background and check on the status:

```
m1demo9.dfw.ibm.com - PuTTY
hscroot@m1demo9:~> migrlpar -o m -m mlp740 -t m2p740 -p m1aix3 --ip m2demo9 -u h
hscroot &
[1] 31856
hscroot@m1demo9:~> lsiparmigr -r sys -m mlp740 -F "num_active_migrations_in_prog
ress" --header
num_active_migrations_in_progress
1
hscroot@m1demo9:~>
```

Figure 14. HMC CLI LPM & Status Check Commands

During Migration, Execute Inventory Collection on ISD to Drive Up Workload

We had a previously run inventory collection of all endpoints and simply restarted that job to drive up workload on the ISD server as the migration continues to further test application functionality while migration continues to run:

Select	Name	Status	Progress	Last Run St...	Description	Next Run	Last Run	Task	Created By
<input type="checkbox"/>	Service and Support Manager Su...	Active	100%	Complete (view l...	Repeat job Ever...	10/4/13 at 10:0...	10/3/13 at 10:0...	Allocate space f...	System
<input type="checkbox"/>	Send System Service Information...	Active	100%	Complete (view l...	Repeat job Ever...	11/3/13 at 10:0...	10/3/13 at 10:0...	Send System Se...	System
<input type="checkbox"/>	Collect Performance Data	Active	100%	Complete (view l...	Repeat job Ever...	10/4/13 at 10:0...	10/3/13 at 10:0...	Collect Performa...	System
<input type="checkbox"/>	Update Compliance - October 3, ...	Complete	100%	Complete (view l...	Run once on 10/...		10/3/13 at 9:53...	Update Complian...	System
<input checked="" type="checkbox"/>	Collect Inventory LOAD TEST	Active	100%	Complete (view l...	Run once on 10/...		10/3/13 at 9:50...	Collect Inventory	root
<input type="checkbox"/>	Update Compliance - Octobe	Complete	100%	Complete (view l...	Run once on 10/...		10/2/13 at 5:39...	Update Complian...	System
<input type="checkbox"/>	Collect Inventory ALL TEST	Complete	100%	Complete (view l...	Run once on 10/...		10/2/13 at 5:37...	Collect Inventory	root
<input type="checkbox"/>	Update Compliance - Octobe	Complete	100%	Complete (view l...	Run once on 10/...		10/2/13 at 5:17...	Update Complian...	System
<input type="checkbox"/>	System Discovery - m1aix3	Complete	100%	Complete (view l...	Run once on 10/...		10/2/13 at 5:13...	System Discovery	root
<input type="checkbox"/>	Update Compliance - Octobe	Complete	100%	Complete (view l...	Run once on 10/...		10/2/13 at 5:12...	Update Complian...	System
<input type="checkbox"/>	Update Compliance - Octobe	Complete	100%	Complete (view l...	Run once on 10/...		10/2/13 at 5:10...	Update Complian...	System
<input type="checkbox"/>	System Discovery - m1aix3	Complete	100%	Complete (view l...	Run once on 10/...		10/2/13 at 3:41...	System Discovery	root

Figure 15. Executing ISD Inventory Collection to Drive Load

View Status of Inventory Collection as Migration Runs

Select	Name	Status	Progress	Last Run St...	Description	Next Run	Last Run	Task	Created B
<input checked="" type="checkbox"/>	Collect Inventory LOAD TEST	Active	52%	Running	Run once on 10/...		10/3/13 at 10:21 AM	Collect Inventory	root

Figure 16. ISD Inventory Collection Job Status

Select	Name	Status	Progress	Last Run St...	Description	Next Run	Last Run	Task	Created B
<input checked="" type="checkbox"/>	Collect Inventory LOAD TEST	Active	94%	Running	Run once on 10/...		10/3/13 at 10:21 AM	Collect Inventory	root

Figure 17. ISD Inventory Collection Job Status

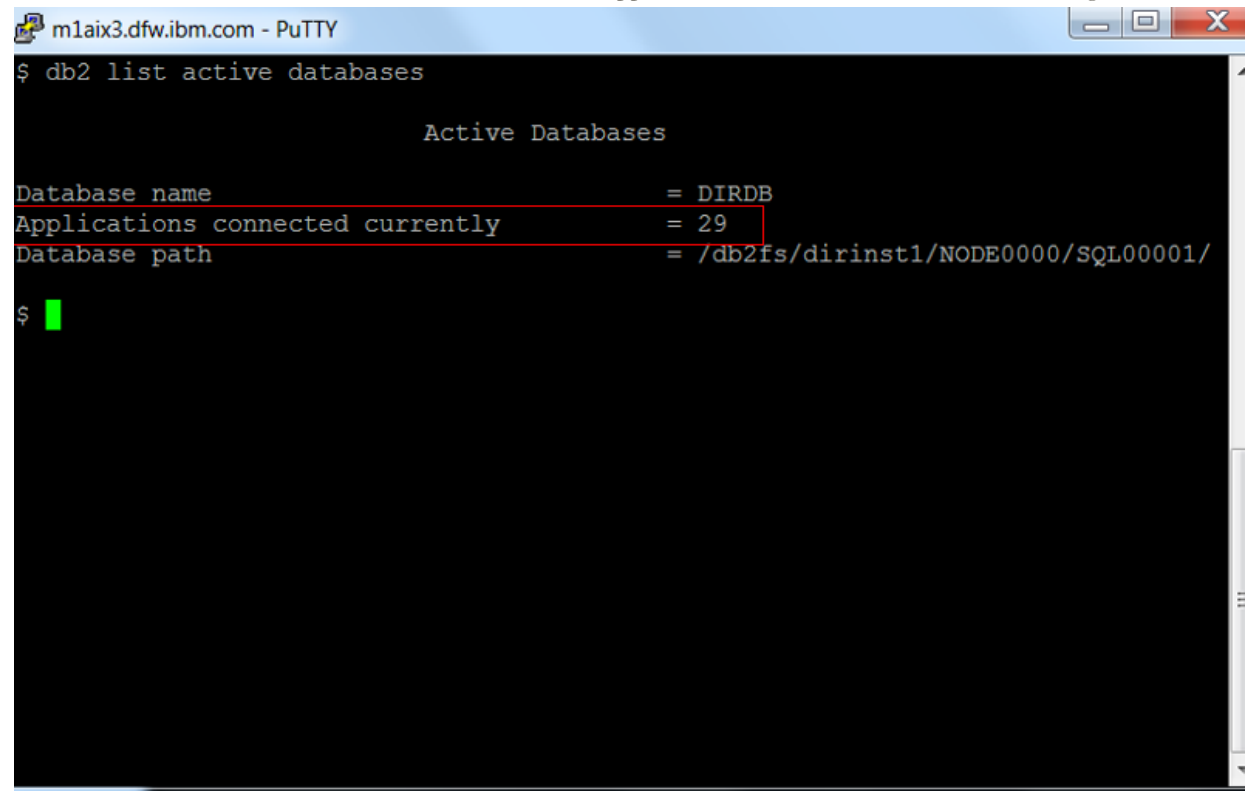
Based on the below snippets from the log, we can see that the job took just under 2 minutes to run and completed without errors:

October 3, 2013 10:21:24 AM CDT-Level:1-MEID:0--MSG: Job "Collect Inventory LOAD TEST" activated.

October 3, 2013 10:23:14 AM CDT-Level:1-MEID:0--MSG: Job activation status changed to "Complete".

Connect to DIRBD01 DB as Migration Runs

We first confirmed that IBM DB2 is still active with ISD applications attached to it's DB as shown in Figure 18. below:



```
m1aix3.dfw.ibm.com - PuTTY
$ db2 list active databases

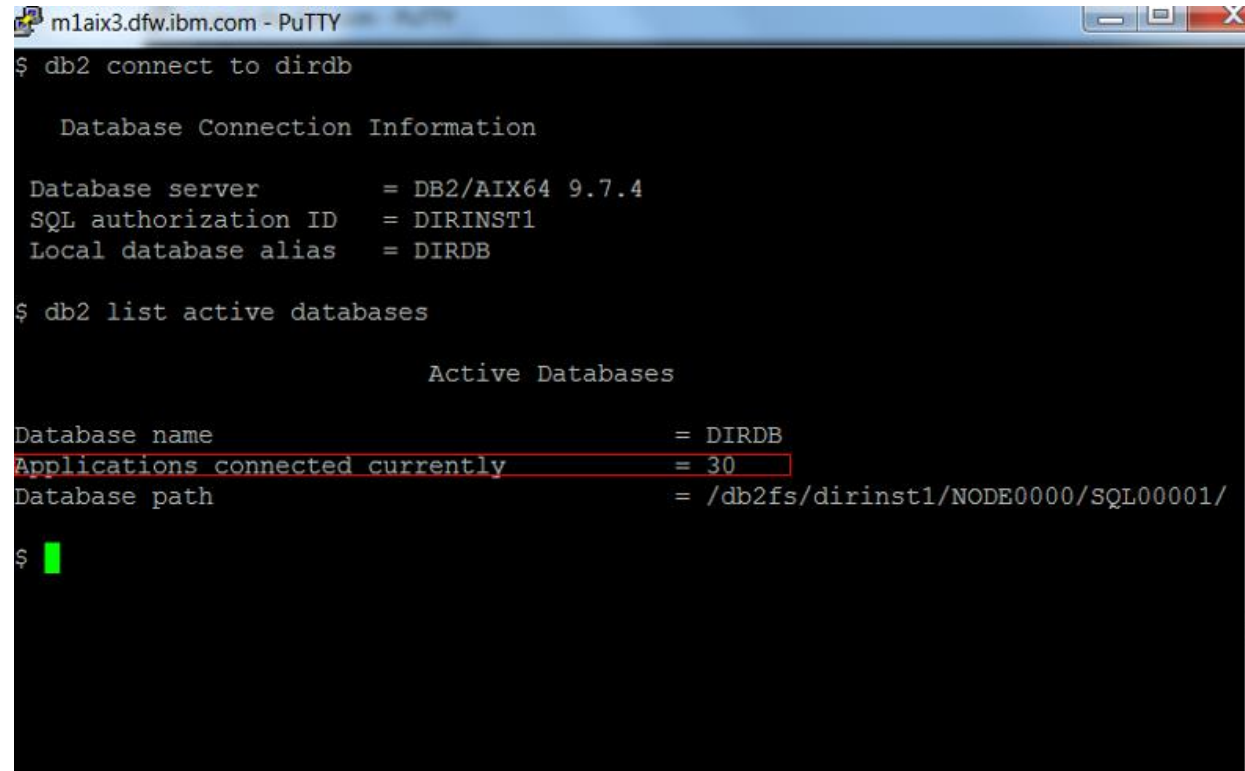
                Active Databases

Database name           = DIRDB
Applications connected currently = 29
Database path          = /db2fs/dirinst1/NODE0000/SQL00001/

$ █
```

Figure 18. Confirm IBM DB2 DB "DIRDB" is Active

We then confirmed we could still directly connect to the ISD database hosted by DB2 as shown below:



```
m1aix3.dfw.ibm.com - PuTTY
$ db2 connect to dirdb

Database Connection Information

Database server          = DB2/AIX64 9.7.4
SQL authorization ID    = DIRINST1
Local database alias    = DIRDB

$ db2 list active databases

Active Databases

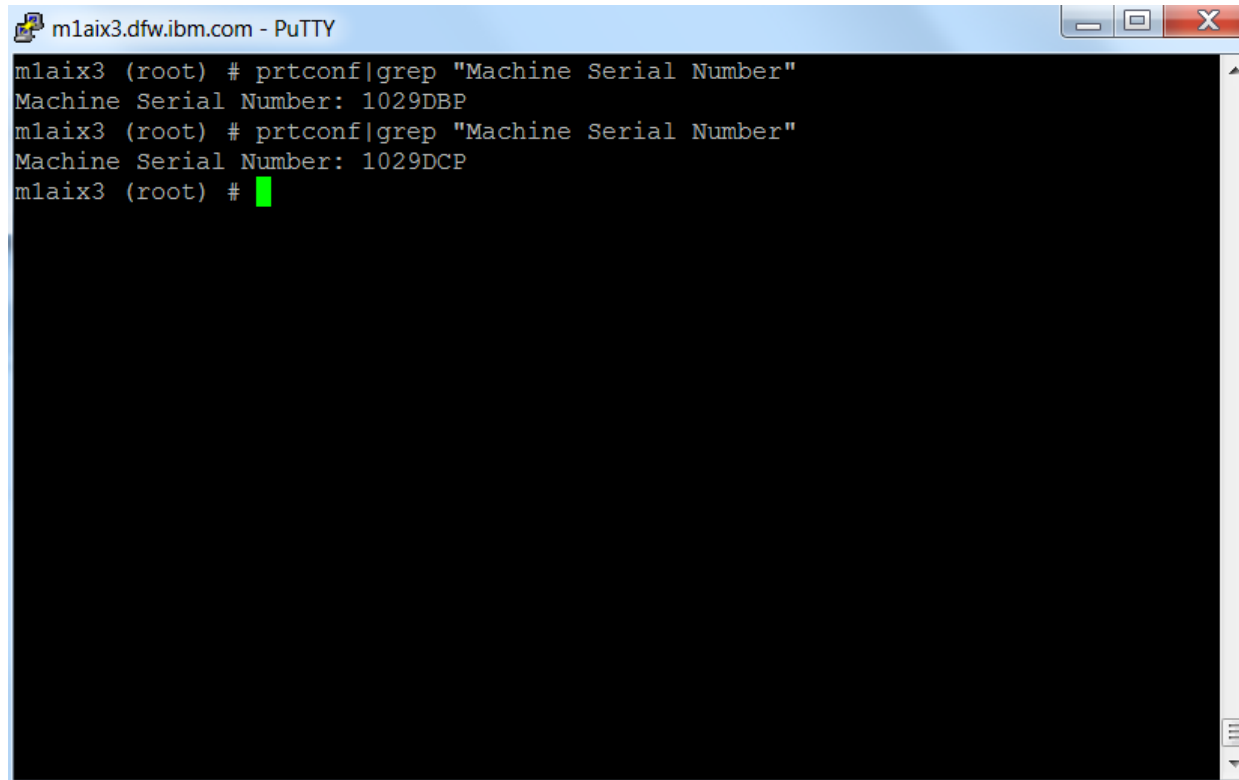
Database name           = DIRDB
Applications connected currently = 30
Database path           = /db2fs/dirinst1/NODE0000/SQL00001/

$ █
```

Figure 19. Connecting to DIRDB DB2 Database

Confirm Serial Number Change Following Successful Active Migration

Once the migration completed we confirmed the system serial number has changed on virtual server:

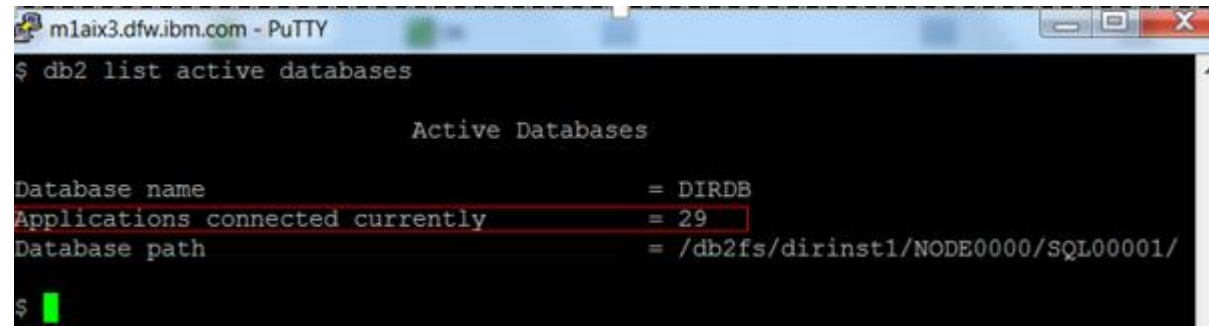


```
m1aix3.dfw.ibm.com - PuTTY
m1aix3 (root) # prtconf|grep "Machine Serial Number"
Machine Serial Number: 1029DBP
m1aix3 (root) # prtconf|grep "Machine Serial Number"
Machine Serial Number: 1029DCP
m1aix3 (root) # █
```

Figure 20. Post LPM Serial Number Change

Confirm ISD and DB2 Access (GUI & CLI) Following Successful Active Migration

We also confirm that IBM DB2 is still active with ISD applications connected to it:



```
m1aix3.dfw.ibm.com - PuTTY
$ db2 list active databases

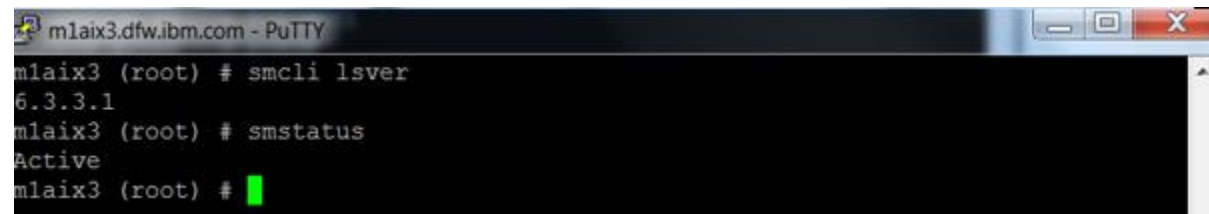
Active Databases

Database name           = DIRDB
Applications connected currently = 29
Database path          = /db2fs/dirinst1/NODE0000/SQL00001/

$
```

Figure 21. Confirming DB2 Active Post LPM

We also confirm that IBM ISD is still active and thereby connected to its IBM DB2 DB:



```
m1aix3.dfw.ibm.com - PuTTY
m1aix3 (root) # smcli lsver
6.3.3.1
m1aix3 (root) # smstatus
Active
m1aix3 (root) #
```

Figure 22. Confirming ISD Active Post LPM

We can see that ISD continues to collect dashboard data in the GUI and monitor m1p740 utilization. Note, that m1aix3 is no longer listed under m1p740 in ISD due to it being migrated to m2p740 which is not monitored by ISD:

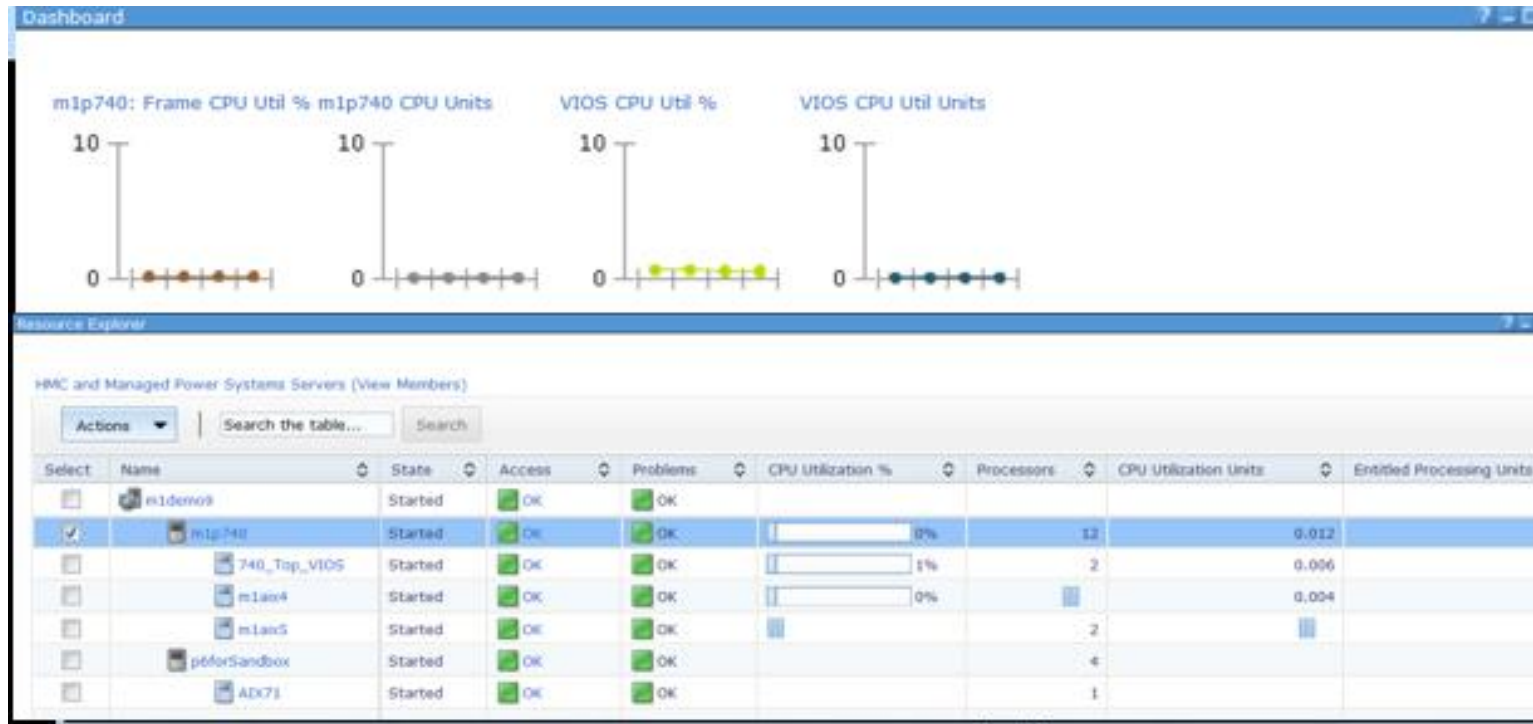
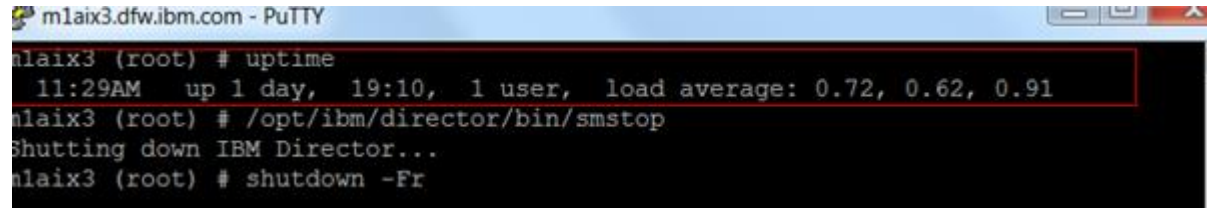


Figure 23. ISD Health Summary Dashboard & HMC Resource Explorer View

Reboot Virtual Server 'mlaix3' on Target System m2p740

We will now reboot the virtual server to confirm it comes backup following the active migration:



```
mmlaix3.dfw.ibm.com - PuTTY
mlaix3 (root) # uptime
11:29AM up 1 day, 19:10, 1 user, load average: 0.72, 0.62, 0.91
mlaix3 (root) # /opt/ibm/director/bin/smstop
Shutting down IBM Director...
mlaix3 (root) # shutdown -Fr
```

Figure 24. Restarting Virtual Server Post LPM

After severel minutes, the virtual server was back up and running and we ran several commands to confirm the disks are avilable and needed volume groups online:




```
mmlaix3 (root) # uptime
11:34AM up 2 mins, 1 user, load average: 1.90, 0.60, 0.22
mlaix3 (root) # lsdev -Cdisk
hdisk0 Available 23-T1-01 MPIO IBM FlashSystem Disk
hdisk1 Available 23-T1-01 MPIO IBM FlashSystem Disk
mlaix3 (root) # lsvg -o
datavg
rootvg
mlaix3 (root) # lspv
hdisk0 00f629db70b86547 rootvg active
hdisk1 00f629db71747c3c datavg active
mlaix3 (root) #
```

Figure 25. Virtual Server Backup After Reboot with VGs & Disks Online

Confirming Both IBM ISD & DB2 Started Correctly Following Virtual Server Reboot

We can see below that both IBM DB2 and ISD started correctly (ISD can take up to 15-20 minutes to fully activate following a reboot):



```
m1aix3.dfw.ibm.com - PuTTY
$ db2 list active databases

           Active Databases

Database name           = DIRDB
Applications connected currently = 26
Database path           = /db2fs/dirinst1/NODE0000/SQL00001/

$
```

Figure 26. IBM DB2 Status Following Reboot



```
m1aix3.dfw.ibm.com - PuTTY
m1aix3 (root) # smcli lsver
6.3.3.1
m1aix3 (root) # smstatus
Active
m1aix3 (root) #
```

Figure 27. IBM ISD Status Following Reboot

Conclusion

The preceding POC effectively demonstrated the interoperability and synergy of Power 7 servers used in conjunction with IBM FlashSystems storage to perform Live Partition Mobility operations. Clients looking to incorporate IBM FlashSystems in their Power 7 environment which utilize Live Partition Mobility can rest assured that it is supported and will work.

Appendix

Acknowledgements

The table below list the team that approved, designed, implemented, and documented the preceding POC:

James Nash	IBM ATS
Franklin Almonte	IBM ATS
Eric Hess	IBM ATS
Sumanta Banerji	IBM STG

Trademarks

Power Systems®, IBM DB2®, IBM Systems Director Server®, are registered trademarks of IBM Corporation.

POC Reference Material

Flash resources: <http://www-03.ibm.com/systems/storage/flash/resources.html>

LPM resources: <http://www.redbooks.ibm.com/abstracts/sg247460.html>