

# IBM i Live Partition Mobility in a multichassis IBM PureFlex System environment

## Preparation and running

Luca Comparini

February 28, 2013

This article explains the main steps to configure Live Partition Mobility (LPM) for IBM® i in an IBM PureFlex™ System multichassis environment. Running LPM is also documented, according to a video caption taken during the IBM Flex System Manager™ GA3 release testing.

## What is (and is not) Live Partition Mobility

LPM is an IBM PowerVM® feature capable of migrating a running partition from one IBM Power Systems™ server to another; migration is performed without disrupting the transactions and the applications that are running on the partition, in other words transparently for the business.

LPM is not a business continuity feature. LPM improves the serviceability and maintainability of an environment running on Power Systems servers. In other words, this feature is designed to make the IT manager's life easier, when a planned maintenance is scheduled because it is capable of reducing the downtime to zero for a planned maintenance activity. Reasons for using LPM also include workload consolidation (from many servers to one), workload balancing (distribute workloads across a pool of servers), and workload resilience in response to a Predictive Failure Analysis (PFA) event.

## IBM i and LPM requirements

Describing LPM requirements is not the goal of this article as the support statements are reported in IBM Redbooks®. General support statements for IBM i LPM are documented in the [IBM i Technology Updates - Live Partition Mobility](#) article.

## The building blocks

The following list of building blocks is generally applicable to any LPM scenario:

- The IBM **Power® server defined as *source***, running the IBM i partition before running LPM

- The **Power server defined as *target***, running the IBM i partition running LPM
- The **storage system hosting the IBM i logical unit numbers (LUNs)**: LUNs must be visible to both the source and destination servers
- The **LPM *orchestrator***, validating and coordinating the migration operations
- The **IBM i partition**, running the workload

## The hardware

### The source server

The source server is an IBM Flex System™ p460 Compute Node, in an IBM Flex System Enterprise Chassis called **Chassis 1**.

**Figure 1. The source server for LPM occupies slots 3 and 4 in Chassis 1**



The Flex System p460 Compute Node, where the IBM i workload is running, occupies slots 3 and 4, as it is a full-wide server.

The IBM i partition is installed with the IBM i 7.1 TR4 release and the latest program temporary fix (PTF) groups, with 4 GB of RAM and 0.2 processor uncapped units. The IBM i partition is fully virtual, with two Virtual I/O Server (VIOS) partitions managing the disks and network virtualization. Storage is provided by two 25 GB LUNs, visible to both the VIOS and mapped to the IBM i logical partition (LPAR) through virtual Small Computer System Interface (VSCSI) adapters.

### The destination server

The destination server is another IBM Flex System p460 Compute Node, in a different IBM Flex System Enterprise Chassis called **Chassis 3**.

**Figure 2. The destination server for LPM occupies slots 1 and 2 in Chassis 3**



The Flex System p460 Compute Node where IBM i will be migrated, occupies slots 1 and 2, as it is a full-wide server, too. Also in this server, two VIOS partitions are running and the same IBM i storage LUNs are already visible to them.

### The storage system

All the IBM i storage space is provided by two 25 GB LUNs on an IBM Storwize® V7000 storage system. As the IBM i disks are made visible to the partition through virtual SCSI adapters, LUNs are mapped to the Fibre Channel (FC) cards owned by the VIOS.

**Figure 3. Storwize V7000 host mappings view - the two IBM i LUNs are mapped to each VIOS**

v7k01 > Hosts > Host Mappings ▾

⋮ Actions ▾

Host Name	SCSI ID	Volume Name	Volume Unique Identifier
SN1054A5B_VIOS2	1	ibmi-SN1054A3B-01	600507680280876500000000000000A7
SN1054A3B_VIOS1	4	ibmi-SN1054A3B-01	600507680280876500000000000000A7
SN1054A3B_VIOS2	4	ibmi-SN1054A3B-01	600507680280876500000000000000A7
SN1054A5B_VIOS1	1	ibmi-SN1054A3B-01	600507680280876500000000000000A7
SN1054A5B_VIOS2	2	ibmi-SN1054A3B-02	600507680280876500000000000000A8
SN1054A3B_VIOS1	5	ibmi-SN1054A3B-02	600507680280876500000000000000A8
SN1054A3B_VIOS2	5	ibmi-SN1054A3B-02	600507680280876500000000000000A8
SN1054A5B_VIOS1	2	ibmi-SN1054A3B-02	600507680280876500000000000000A8

In this context, on each VIOS on the *source* server:

- The two IBM i LUNs must be visible as hdisk devices.

- The two IBM i LUNs must be associated with two virtual destination devices (VTD) associated to the logical representation of the virtual SCSI adapter pointing to the IBM i partition (vhost).

### Listing 1. View of the VIOS 1 running on the source server: LUNs are visible and mapped to the IBM i partition

```
$ hostname
sn1054a3bvios1

$ lspv -size | grep 'SIZE\|25600'
NAME          PVID          SIZE(megabytes)
hdisk5        00054a3b868044e2 25600
hdisk6        00054a3b868072a7 25600

$ lsmmap -vadapter vhost2
SVSA          Physloc          Client Partition ID
-----
vhost2        U7895.42X.1054A3B-V1-C7 0x000000005

VTD           vtscsi2
Status        Available
LUN           0x8100000000000000
Backing device hdisk5
Physloc       U78AF.001.WZS0224-P1-C35-L1-T1-W5005076802158969-L4000000000000
Mirrored      false

VTD           vtscsi3
Status        Available
LUN           0x8200000000000000
Backing device hdisk6
Physloc       U78AF.001.WZS0224-P1-C35-L1-T1-W5005076802158969-L50000000000000
Mirrored      false
```

### Listing 2. View of the VIOS 2 running on the source server: LUNs are visible and mapped to the IBM i partition

```
$ hostname
sn1054a3bvios2

$ lspv -size | grep 'SIZE\|25600'
NAME          PVID          SIZE(megabytes)
hdisk5        00054a3b868044e2 25600
hdisk6        00054a3b868072a7 25600

$ lsmmap -vadapter vhost2
SVSA          Physloc          Client Partition ID
-----
vhost2        U7895.42X.1054A3B-V2-C7 0x000000005

VTD           vtscsi2
Status        Available
LUN           0x8100000000000000
Backing device hdisk5
Physloc       U78AF.001.WZS0224-P1-C37-L1-T1-W5005076802158969-L4000000000000
Mirrored      false

VTD           vtscsi3
Status        Available
LUN           0x8200000000000000
Backing device hdisk6
Physloc       U78AF.001.WZS0224-P1-C37-L1-T1-W5005076802158969-L50000000000000
Mirrored      false
```

On the *destination* server:

- The two IBM i LUNs must be visible as hdisk devices.

### Listing 3. View of the VIOS 1 running on the destination server: IBM i LUNs are visible

```
$ hostname
sn1054a5bvios1

$ lspv -size | grep 'SIZE\|25600'
NAME                PVID                SIZE(megabytes)
hdisk1              00054a3b868044e2   25600
hdisk2              00054a3b868072a7   25600

$ lsmmap -all
$
```

### Listing 4. View of the VIOS 2 running on the destination server: IBM i LUNs are visible

```
$ hostname
sn1054a5bvios2

$ lspv -size | grep 'SIZE\|25600'
NAME                PVID                SIZE(megabytes)
hdisk1              00054a3b868044e2   25600
hdisk2              00054a3b868072a7   25600

$ lsmmap -all
$
```

### Setting the `reserve_policy` attribute

```
chdev -dev hdiskX -attr reserve_policy=no_reserve
```

In some configuration, the *reserve\_policy* attribute of the disks must be changed. For instance, in dual-VIOS environments, any disk device that is presented to a client partition through a virtual SCSI adapter must be set to *no\_reserve*.

For partition mobility, the *reserve\_policy* on the destination VIOS must be the same as on the source VIOS.

## IBM Flex System Manager

In an IBM PureFlex System environment, Flex System Manager (FSM) is an intelligent appliance designed for management-based tasks for PureFlex System hardware, storage, and networking. FSM is also oriented to the virtualization management and OS provisioning. When managing Power Systems compute nodes, FSM is irreplaceable because the expertise proper to Hardware Management Console (HMC) (including the capability to orchestrate the LPM), merged into the FSM.

## Double-check the FSM release

To check the FSM release, use Secure Shell (SSH) to log into the FSM and enter the `smcli lsconfig` command. FSM must be at 1.2 release or later.

## Double-check that servers are capable of IBM i Live Partition Mobility

Perform the following steps to double-check that servers are capable of IBM i LPM.

1. On the FSM Resource Explorer page, identify the source server.
2. Right-click the source server to get the navigation menu, and then click **System Configuration**.
3. Click **Edit Host**.
4. Click the **Migration** tab.
5. Double-check that IBM i Partition Mobility capable is true.

**Figure 4. FSM - Migration tab**

**Migration Capabilities**

---

Type	Capable	Number of Supported
Inactive	True	16
Active	True	16
Total: 2, Displayed: 2		

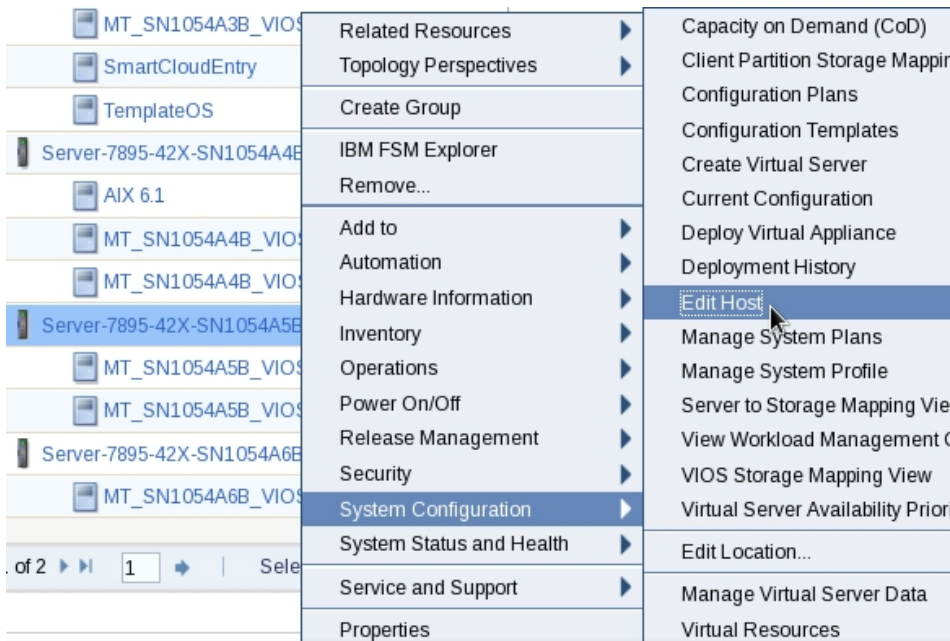
Repeat the above steps also for the destination server.

## Define the service mover partition on both source and destination

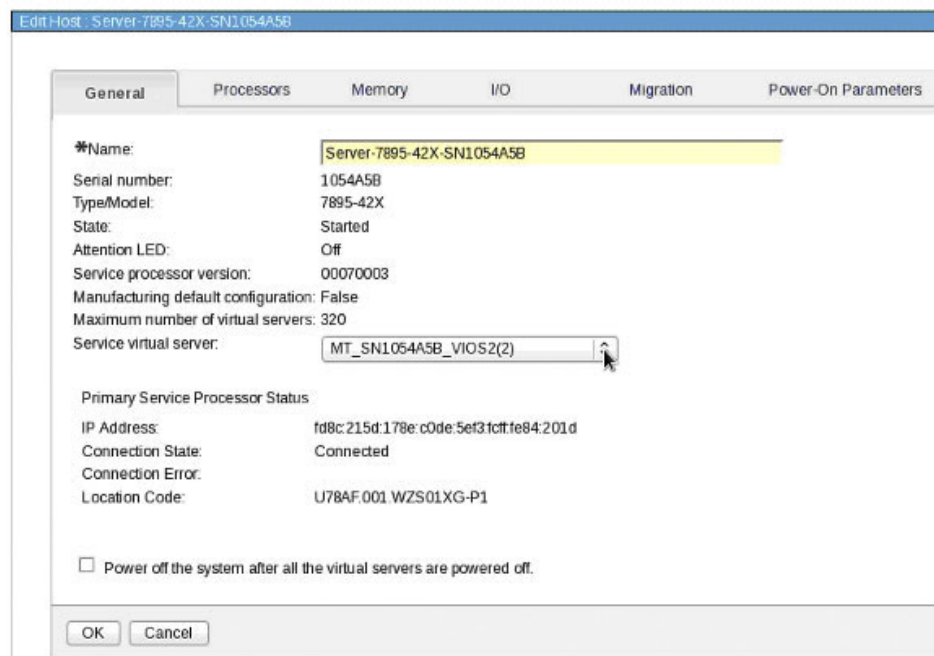
FSM is instrumental to set up the VIOS partition as the service mover, on both source and destination servers.

1. On the FSM Resource Explorer page, identify the source server.
2. Right-click the source server to get the navigation menu and then click **System Configuration**.
3. Click **Edit Host**.
4. Select one of the VIOS as the service virtual server.

## Figure 5. FSM - Selecting the Edit Host option



## Figure 6. FSM - Selecting one of the VIOS as the service virtual server



Repeat the steps above also for the destination server.

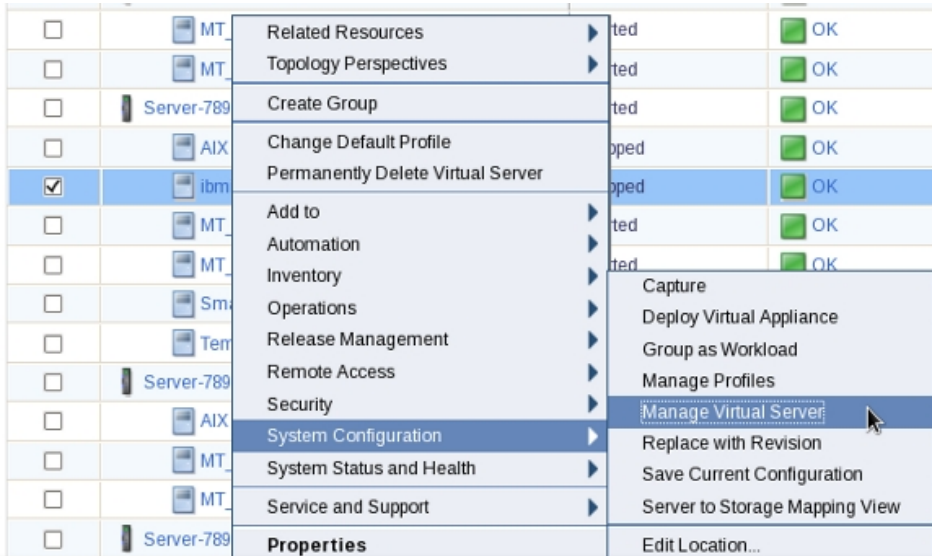
## The IBM i partition

IBM i must be installed with the 7.1 TR4 code (or later) and the latest PTF groups available have to be installed and applied.

1. On the FSM Resource Explorer page, identify the source server.

2. Right-click the IBM i partition that is running on it to get the navigation menu and then click **System Configuration**.
3. Click **Manage Virtual Server**.

**Figure 7. FSM - Manage Virtual Server option**



IBM i partition must be fully virtual; in other words, no physical slot must be assigned to it, as any I/O resource has to be accessible through the VIOS virtualization features. The **IBM i restricted I/O mode** check box must be selected.

**Figure 8. IBM i restricted I/O mode**

Virtual server name:

OS installed : IBM i Licensed Internal Code 7.1.0 410 0

IP address: Not Available

Processors: 0.2

Memory 5.50 GB

---

**General Configuration**

Maximum virtual adapters:

Remote restartable: No

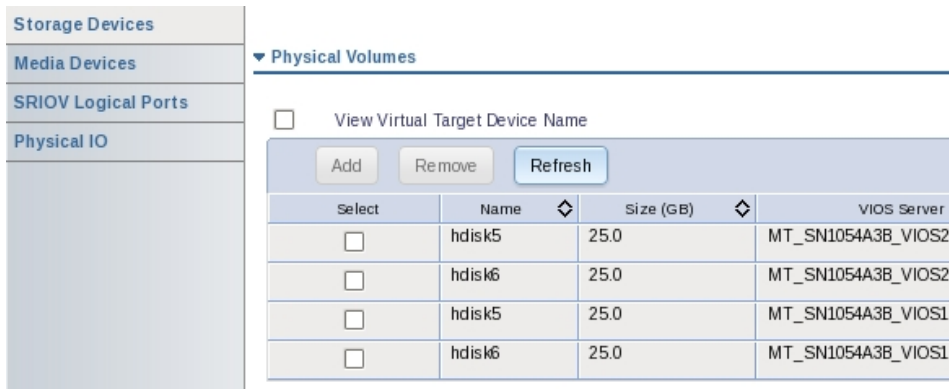
Suspend capable

IBM i restricted I/O mode

From a disk perspective, LUNs are mapped to the VIOS and therefore assigned to the IBM i partition through virtual SCSI adapters. This refers to the sections in Listing 1 and Listing 2. The IBM i partition profile must be consistent too.



**Figure 9. IBM i - Storage Devices page**




**Figure 10. IBM i - Storage Adapters page**

Available Virtual Slots: 5




As described in the previous sections, in this scenario, IBM i has two 25 GB disks available and part of the System ASP (this is not a mandatory condition, just the way I decided to test it).

**Figure 11. IBM i - WRKDSKSTS**



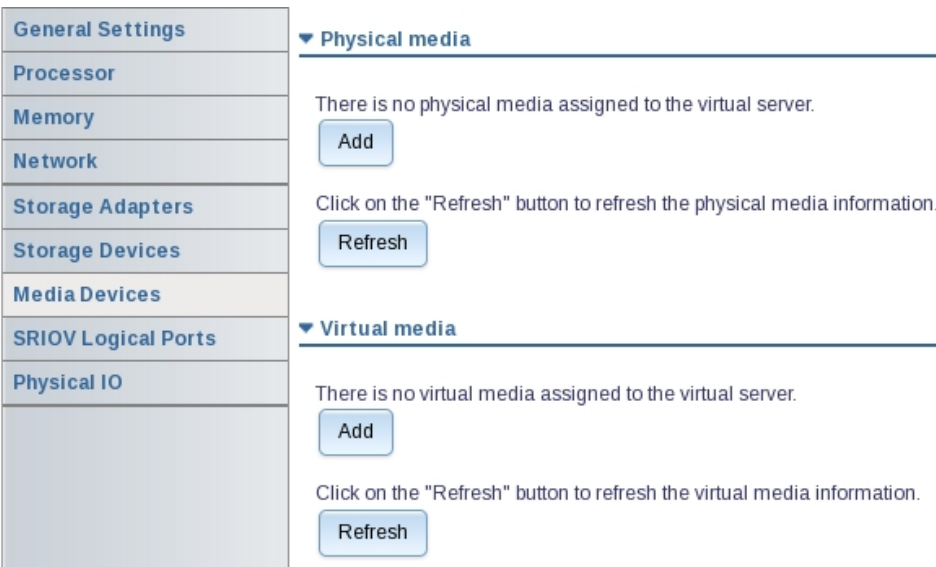
From a network perspective, it is not important how many virtual Ethernet adapters are associated to the IBM i partition; the only important condition is that any bridge (SEA) and VLAN configured on the source server is consistently there, also on the destination server.

**Figure 12. IBM i – Network page**



IBM i partition must not have any physical or virtual media device.

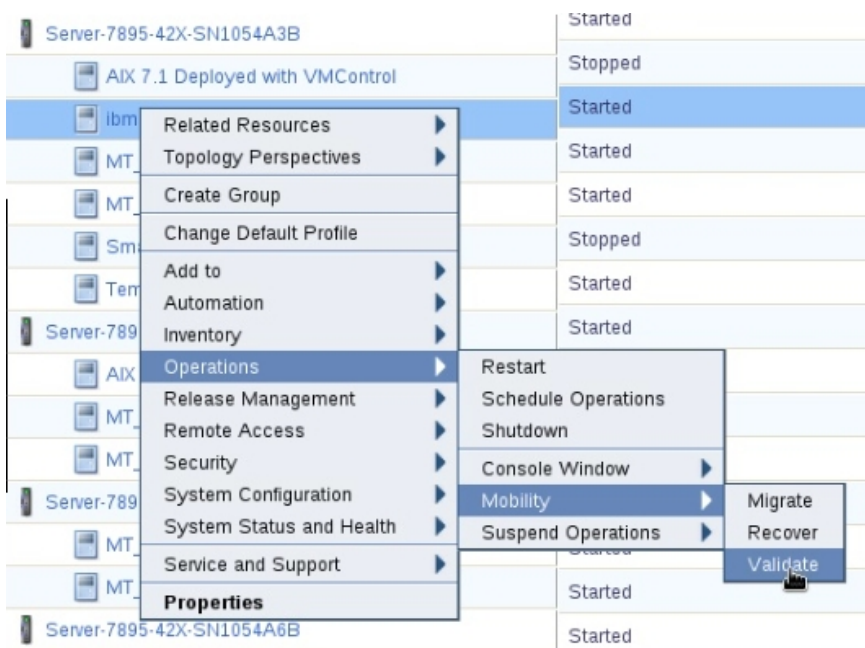
**Figure 13. IBM i - Media Devices page**



## Running LPM

My preferred method to run the LPM in a PureFlex System environment consists of using the **Validation** feature and then **Migrating** afterwards. This can be done by right-clicking the IBM i partition to be migrated, and then clicking **Operations >Mobility >Validate**.

**Figure 14. FSM - Validate option**



**Figure 15. FSM - LPM validation**

Migrating virtual server: ibmi

Remote management console:

Remote user:

\*Destination system: Server-7895-42X-SN1054A5B

Destination profile name: DefaultProfile

Destination shared processor pool: DefaultPool (0)

Source mover service partition: MT\_SN1054A3B\_VIOS1

Destination mover service partition: MT\_SN1054A5B\_VIOS2

Wait time (in min):

Override virtual network errors when possible:

Override virtual storage errors when possible:

Virtual Storage assignments :

Select	Source Slot ID	Slot Type	Destination VIOS
<input type="checkbox"/>	4		SCSI
<input checked="" type="checkbox"/>	4		SCSI
<input checked="" type="checkbox"/>	3		SCSI
<input type="checkbox"/>	3		SCSI

The validation task runs all the needed tasks to check whether a migration would happen correctly. After the validation happens, simply click **Migrate** to initiate the LPM of the IBM i partition.

## Resources

[Video: IBM i Live Partition Mobility in a multi-chassis PureFlex environment](#)

© Copyright IBM Corporation 2013

([www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml))

[Trademarks](#)

([www.ibm.com/developerworks/ibm/trademarks/](http://www.ibm.com/developerworks/ibm/trademarks/))