

Mar 18 '14

Using the Live Partition Mobility Feature of PowerVM

by [Jaqui Lynch](#) in [Enterprise Tech Journal](#)

Since POWER6, IBM has provided a way to move a logical partition (LPAR) between servers while it's running, allowing you to migrate an LPAR or LPARs without any downtime. This feature is called Live Partition Mobility (LPM) and is provided by the Enterprise Edition of PowerVM, which must be installed on all servers involved.

LPM is used for multiple reasons, ranging from server consolidation, server migration and workload balancing to completely evacuating servers so they can be shut down for planned maintenance or other purposes. LPM can be used to move Linux and AIX (5.3 and up) workloads between POWER6 and POWER7 servers and back, provided certain prerequisites are met. With certain limitations, iOS LPARs can also be made mobile. However, LPM isn't disaster recovery nor is it a replacement for PowerHA or other high-availability solutions. It's designed for planned migrations only. Crashed kernels can't be migrated and partitions can't be moved from failed machines.

There are three options for partition mobility:

- Active partition mobility is the actual movement of a running LPAR from one server to another without disrupting the applications currently running in the LPAR. At the end of the migration, network applications may see a brief (less than 1 second) suspension, but connectivity isn't lost. The amount of time for the migration depends on the amount of memory in the LPAR and how busy the LPAR is.
- Inactive partition mobility is when an LPAR is transferred when it's shut down. This kind of migration only takes seconds.
- Suspended partition mobility refers to an LPAR that has been suspended and is then migrated from one server to another.

Migration Prerequisites

Several prerequisites must be met in order for LPM to work smoothly. Details on all prerequisite firmware, virtual I/O server (VIOS), etc. levels can be found online (see the "References" section).

Server: First, the servers must be compatible. Both the source and target need to be POWER6 or above and PowerVM Enterprise Edition must be installed on both servers. There are also minimum firmware requirements. Additionally, the target server must have sufficient available memory, CPU and virtual adapter slots to accommodate the LPAR being moved. Finally, the logical memory block (LMB) or memory region size for the servers must be the same. This can be checked on the server via the hardware management console (HMC). If they aren't the same, they will need to be reset, which requires powering the server off and on.

Management console: LPM servers need to be managed by a management console. An HMC can be used to manage current POWER blades as well as POWER servers and Flex servers. Integrated virtualization manager (IVM) can be used to manage POWER blades only. Flex system manager (FSM) can be used to manage Flex servers only. Since an HMC can manage them all, it's recommended that HMCs be used where possible to provide one simple interface. There are minimum software requirements for each console and these affect the options available for migration, including the number of concurrent migrations possible. As an example, remote migration (between servers controlled by different HMCs) became available at v7.3.4 of the HMC, and mobility between the HMC and an FSM is only available after v7.7.1.0. It's recommended that the HMC be kept at the current level that's available.

VIO Server: LPM was first introduced at VIOS v1.5.2.1-FP11 and v2.1. The current HMC version is v2.2.3.1, which

became available in November 2013. A VIOS on each server must be set as a mover server partition (MSP), as these are used to control the migration.

LPAR: The mobile LPAR itself also has some requirements. There are minimum operating system levels as well as configuration requirements. The first requirement is that all resources be virtualized through a VIO server. This means all storage is provided via the storage area network (SAN) either using virtual SCSI (vSCSI) or N_Port ID virtualization (NPIV), and all networking is provided through the VIO server using a shared Ethernet adapter (SEA). Adapters should also be defined as desired, not required and must not be marked as for “Any client.” While dedicated adapters can be used, they must be removed prior to the move. This also applies to any virtual CDs (vtopt) defined to take advantage of file backed optical (FBO). The LPAR must also have a name that’s unique to both the source and target servers. Additionally, the LPAR can’t be marked as the service partition, nor can it have any open consoles at the time of the move.

Inactive migration has the same requirements as active migration with a couple of exceptions. An LPAR can use inactive migration if it uses huge memory pages, has barrier synchronization register (BSR) arrays, uses redundant error reporting or has physical I/O attached. Active migrations can’t be done in those cases. It should be noted that an LPAR must have been activated at least once (even if only in system managed services [SMS] mode) in order for any migration to occur.

Other: All the LPARs that are to be mobile must be on the same network with resource monitoring and control (RMC) subsystem connections established to the HMC. If active memory sharing (AMS) is being used, then the destination must have a paging device available for the mobile LPAR. If the LPAR is using active memory expansion (AME), then AME must be supported on the destination server as well.

Not only must the storage be virtualized, but it also has to be zoned correctly. For vSCSI, both the source and target HBAs must be zoned. For NPIV, each virtual Fibre Channel adapter has two worldwide port names (WWPNs)—these both have to be zoned for LPM to work. Additionally, whole LUNs need to be passed across (if vSCSI); logical volume manager (LVM-) based disks can’t be used and all hdisks must be external and set to `reserve_policy=no_reserve`. Details on this are provided in the LPM Redbook.

Migrating Between Different Server Types

Instruction sets between POWER6, POWER6+ and POWER7/7+ differ, which causes some issues with migration back to the older technology if care isn’t taken. Migration from POWER6 to POWER6+, POWER7 or POWER7+ is very straightforward. The LPAR will run in POWER6 compatibility mode once it’s on the new server and will remain in that mode unless it’s changed and rebooted. If it’s in POWER6 compatibility mode, then it’s easy to migrate it back to the POWER6 if necessary. Migrating from POWER7 to POWER7+ and back involves no compatibility modes, as they use the same instruction set; however, it’s important to check that the LPAR being migrated is set to an entitlement of at least 0.1 (not 0.05) when migrating from POWER7+ to POWER7, as 0.1 is the minimum supported on POWER7.

Migration Phases

There are slightly different rules for inactive vs. active migrations and they both go through a slightly different set of phases. For an active migration, the phases consist of validating the configuration, creating the new LPAR, creating the new virtual resources, migrating the state of the LPAR in memory, removing the old LPAR configuration and freeing up the old resources. For an inactive migration, the phases consist of validating the configuration, creating the new LPAR, creating the new virtual resources, removing the old LPAR configuration and freeing up the old resources.

In the validation phase, the hypervisor and HMC perform checks to ensure the migration will go through with minimal

risk. Checks are performed on capabilities and compatibility (hypervisor, VIOS and MSPs), RMC connectivity, partition readiness, target system resource availability, virtual adapter mapping, operating system and application readiness as well as the uniqueness of the LPAR name and whether the number of current active migrations is less than the number of supported active migrations.

If validation passes, then the migration phase begins. The HMC creates an LPAR on the target and configures the MSPs that connect to the hypervisor to set up a private channel to transfer partition state data. The HMC creates the target virtual devices and adapters and then the MSP on the source starts sending partition state data to the target. Once almost all pages are moved, the MSP on the source has the hypervisor suspend the mobile partition while the last modified memory pages and state data are moved. The partition is resumed on the destination and any uncompleted I/O is recovered, and cleanup begins on the source.

Remote migration: Remote migration is the ability to move LPARs between two servers on different HMCs. To do this, a minimum of v7.3.4 of the HMC software is required along with network access between the two HMCs. SSH key authentication needs to be set up to the remote HMC and all involved LPARs (VIOS and mobile LPARs).

Summary

LPM is a valuable tool for any data center that regularly does migrations to new technology or that needs to do firmware and other updates without taking outages. As long as you've done proper planning, LPM can save significant time and allow for shorter maintenance windows.

References

For more information on LPM, check out the following resources:

- LPM Redbook: www.redbooks.ibm.com/redbooks/pdfs/sg247460.pdf
- LPM Prerequisites: <http://pic.dhe.ibm.com/infocenter/powersys/v3r1m5/index.jsp?topic=/p7hc3/p7hc3firmwaresupportmatrix.htm>
- LPM Demo: www.circle4.com/movies/lpm/index.html
- AIX Virtual User Group (lots of topics including LPM): www.tinyurl.com/ibmaixvug. **ETI**