

5<sup>th</sup> October 2015

## AIX Live Update

Starting with AIX Version 7.2, the AIX operating system provides the AIX Live Update function which eliminates downtime associated with patching the AIX operating system. Previous releases of AIX required systems to be rebooted after an interim fix was applied to a running system. This new feature allows workloads to remain active during a Live Update operation and the operating system can use the interim fix immediately without needing to restart the entire system. In the first release of this feature, AIX Live Update will allow customers to install interim fixes (ifixes) only. Ultimately it may be possible to use this function to install AIX Service Packs (SPs) and Technology Levels (TLs) without a reboot.

IBM delivers kernel fixes in the form of ifixes to resolve issues that are reported by customers. If a fix changes the AIX kernel or loaded kernel extensions that cannot be unloaded, the host logical partition (LPAR) must be rebooted. To address this issue, AIX Version 7.1, and earlier, provided concurrent update-enabled ifixes that allowed deployment of some limited kernel fixes to a running LPAR. Unfortunately not all ifixes could be delivered as “concurrent update-enabled”. The AIX Live Update solution is not constrained by the same limitations as in the case of concurrent update enabled ifixes. The AIX 7.2, Live Update feature will allow customers to install ifixes without needing to reboot their AIX systems, avoiding downtime for their mission critical, production workloads.

This article will discuss the high-level concepts relating to AIX Live Updates and then provide a real example of how to use the tool to patch a live AIX system. I was fortunate enough to take part in an Early Ship Program (ESP) for AIX 7.2. During the ESP I had the opportunity to test the AIX Live Update feature. I'll share my experience using this tool in the example that follows.

## AIX Live Update Concepts

Live Update is the next generation in AIX Live Update technology. The development team set out to provide an innovative tool for patching, that could leverage existing AIX maintenance models and tools, such as **emgr** and **installp**. The tool was designed to allow for non-disruptive updates for all AIX components, such as the kernel, commands and libraries. The starting point was for ifixes only but longer term the goal will be to provide non-disruptive updates for SPs and TLs.

To achieve this goal, the AIX Live Update function utilises what's known as original and surrogate AIX LPARs. An AIX Live Update operation is started on the original partition. Another LPAR is provisioned (automatically) and will become the surrogate partition. This partition is patched, live, while your workloads continue to run on the original partition. At a point in time, the workload is migrated from the original partition to the “new” patched surrogate partition. Essentially the partition undergoes a “checkpointing” process in which the workload is paused and its current state is saved (for all running processes). Once the “checkpointing” is complete the processes are migrated to (restarted/un-paused on) the new partition. The checkpoint saves and validates the status of the current workload and then starts its back up on the other LPAR in this saved state. This is similar to Workload Partition Live Application Mobility which was introduced with AIX 6.1 in 2007.

The ifix is applied on the surrogate LPAR and the running workload is transferred from the original partition to the surrogate partition. There are several critical steps in a Live Update operation, these are listed below:

- The root volume group of the original partition is cloned using standard AIX alternate disk management utilities.
- The ifix is applied on the cloned volume group that serves as the boot volume group for the surrogate partition. This disk is assigned to the new surrogate partition from which it boots a minimal AIX environment.
- After the surrogate partition is booted and while the workloads are still running on the original partition, the root volume group of the surrogate partition is mirrored.
- The workload processes are checkpointed and moved over to the surrogate partition.
- Workloads resume on the surrogate partition in a “chrooted” environment on the mirrored volume group. During this process, the workloads continue to run without being stopped, although there is a short blackout time when they are suspended.

The following diagram provides a basic overview of the components of a Live Update environment for AIX 7.2.

## Next Generation AIX Live Update Technology

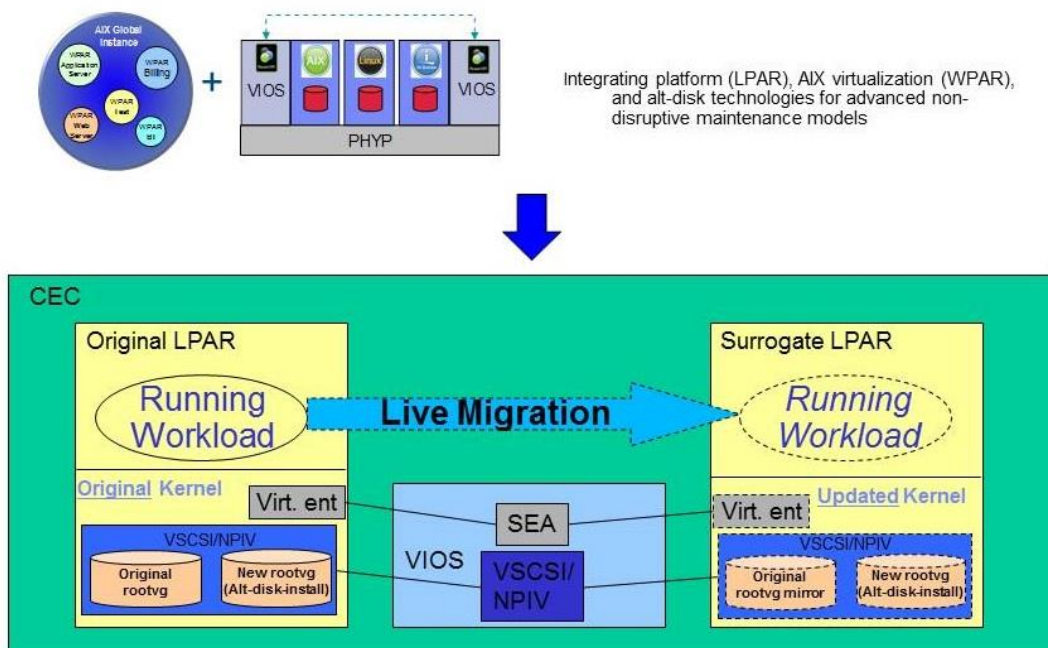


Figure 1 – AIX Live Update components

The AIX Live Update operation can be launched using the **geninstall** command with the **-k** flag or through the Network Installation Manager (**NIM**) or the System Management Interface Tool (**SMIT**). You configure AIX Live Update by modifying the stanzas in the **/var/adm/ras/liveupdate/lvupdate.data** file. A template of this file is supplied with AIX 7.2, called

`/var/adm/ras/liveupdate/lvupdate.template`. You must copy and edit this file to reflect your own configuration. The `geninstall` command uses a lock file, `/usr/lpp/.genlib.lock.check`, to guarantee that no other Live Update process can run simultaneously. The Live Update operation runs in one of the following modes:

- **Preview mode**

In preview mode, estimation of the total operation time, estimation of application blackout time, and estimation of resources such as storage and memory are provided to the user. These estimations are based on the assumption that the surrogate partition has the same resources in terms of CPU, memory and storage as the original partition. All the provided inputs are validated and the AIX Live Update limitations are checked.

- **Automated mode**

In automated mode, a surrogate partition with the same capacity as the original partition is created, and the original partition is turned off and discarded after the AIX Live Update operation completes.

The mirror copy of the original root volume group (rootvg) is retained after the AIX Live Update operation is complete. Thus, if you want to return to the state of the system before applying the ifix, the LPAR can be restarted from the disk that was specified as the mirror volume group.

The main item to consider is that there must be sufficient resources (CPU and memory) available in your environment for a second “copy” or “clone” of your partition to be created during the AIX Live Update process.

## Planning for Live Updates on AIX

If you plan to use Live Update in your AIX 7.2 environment, the following minimum requirements must be met.

- AIX Live Update is currently only supported with ifixes.
- All I/O devices must be virtualized (virtual Ethernet, Virtual Small Computer System Interface (VSCSI) or N-Port Id Virtualisation (NPIV) with AIX multipath I/O (MPIO)).
- Temporary CPU and memory is required (on the same frame).
- Two disks required:
  1. Initial boot disk for surrogate (freed after subsequent AIX Live Update or reboot).
  2. New rootvg (mirrored/split during AIX Live Update) – “old\_rootvg” can be freed after AIX Live Update.

The following system firmware, Hardware Management Console (HMC) and Virtual I/O Server (VIOS) levels must be installed for AIX Live Update to function and be supported in your environment.

### System firmware

- Ax730\_066\*
- Ax740\_043\*

- Ax770\_063
  - Ax773\_056
  - Ax780\_056
  - Ax810 or later
- \* Limitation: PowerVC cannot seamlessly manage the updated LPAR

#### HMC

- 840

#### Virtual I/O Server

- 2.2.3.50
- 2.2.4.0

#### RSCT (if required)

- 3.2.1.0

#### PowerHA (if required)

- 7.2.0

#### PowerSC (if required)

- 1.1.4.0

The following is a list of currently known requirements and limitations with AIX Live Update.

- Support for ifixes only, including kernel and kernel extension ifixes (no SPs or TLs).
- The AIX administrator must be able to authenticate with the HMC before updating. The **hmcauth** utility should be used to establish this authentication prior to the AIX Live Update process starting.
- There must be at least 2 paths to storage (half of the paths will be removed during update).
- Not intended for updates of an Oracle RAC or DB2 PureScale cluster node. RSCT cluster services will be stopped during the update.
- In a PowerHA environment the node will be “unmanaged” during the AIX Live Update operation.
- Only JFS2 and NFS file systems supported.
- Workload must be able to accommodate the “blackout” period. The blackout time is the duration when the running processes are paused during the AIX Live Update operation. The blackout time can be estimated by running the AIX Live Update operation in preview mode.
- Transmission control protocol (TCP) connections will be maintained. Protocols like TCP use a back-off retransmit timeout that allows TCP connections to remain active during the blackout time, so the blackout time is not apparent to most workloads.
- Preview mode will estimate the blackout time (in seconds).
- The lpar\_id value changes as a result of the AIX Live Update operation. You can request a specific lpar\_id value in the **lvupdate.data** file, but it cannot be the same as the original value.

- **I/O restrictions**
  - Any Coherent Accelerator Processor Interface (CAPI) device must not be open during the AIX Live Update operation.
  - No physical or virtual tape or optical device is supported. These devices must be removed before the AIX Live Update operation can proceed.
  - The mirrorvg utility can mirror up to 3 copies. If the root volume group of the original partition is already being mirrored with 3 copies, the AIX Live Update operation cannot proceed.
  - The AIX Live Update operation is not supported on diskless AIX clients.
  - The AIX Live Update operation is not supported in a multibos environment.
  - Data Management API (DMAPI) is not supported by the AIX Live Update feature.
  - VSCSI support for the AIX Live Update operation is only for those logical unit numbers (LUNs) that are backed by physical volumes, not logical volumes.
  - VSCSI disk support excludes the option where the VSCSI server adapter can be mapped to any partition or partition slot.
  - At the time of writing (September 2015), Shared Storage Pool (SSP) disks are not supported with AIX Live Update and VSCSI clients. Attempting a Live Update operation on an AIX partition with SSP hdisks will fail. This is intended to be a supported environment. In the interim NPIV storage or VSCSI disks backed by whole disks is supported.
  
- **Security restrictions**
  - The AIX Live Update operation is not supported when a process is using Kerberos authentication.
  - The AIX Live Update feature does not support PowerSC Trusted Logging.
  - The AIX Live Update feature is not supported by an active Department of Defence (DoD) security profile.
  - The AIX Live Update feature is not supported when audit is enabled for a stopped workload partition (WPAR).
  - The AIX Live Update feature does not support Public-Key Cryptography Standards # 11 (PKCS11). The security.pkcs11 fileset cannot be installed.
  - The AIX Live Update feature is not supported by any of the following Trusted Execution options in the trustchk command:
    - TEP=ON
    - TLP=ON
    - CHKSHLIB=ON and STOP\_UNTRUSTD=ON
    - TSD\_FILES\_LOCK=ON
  
- **Reliability, availability and serviceability (RAS) restrictions**
  - System trace of the AIX Live Update operation is not possible if channel 0 is already in use.
  - The AIX Live Update feature is not supported when ProbeVue is running. The ProbeVue session needs to be stopped to run the AIX Live Update operation.
  - User storage keys are not supported in the AIX Live Update environment.

- Any system dump that is present on the root volume group of the original LPAR is not available after a successful AIX Live Update operation.
- **Miscellaneous restrictions**
  - The ifix must have the LU CAPABLE attribute, which means the ifix must be compatible with the AIX Live Update operation. The **emgr** command can display this attribute. Ideally, all the ifixes can be applied with the AIX Live Update operation, but there might be some exceptions.
  - The location of the ifix files must be on the root volume group of the client partition in either /, /usr, /home, /var, /opt, or /tmp file systems.
  - Network File System (NFS)-mounted executables must not be running during a AIX Live Update operation.
  - Active WPARs must be stopped before the AIX Live Update operation.
  - RSCT Cluster Services are stopped during AIX Live Update operations, and then restarted before the AIX Live Update operation completes.
  - A configuration with 16 MB page support is not allowed. The promoted (16 MB Multiple Page Segment Size (MPSS)) pages by Dynamic System Optimizer (DSO) are supported by the AIX Live Update operation.
  - The AIX Live Update operation is supported when the DSO running, but DSO optimization is reset by the AIX Live Update operation. The optimization begins again based on workload monitoring after the AIX Live Update operation.
  - The AIX Live Update feature is not supported on a partition that participates in Active Memory Sharing (AMS).
  - The AIX Live Update feature is not supported on a remote restartable partition.
  - If an ifix is installed, without the AIX Live Update operation, that requires a restart, the restart must be completed before a subsequent AIX Live Update operation can be started.

Please refer to the AIX Knowledge Centre for the latest information on known limitations and current requirements.

In my Power Systems lab environment I had the following configuration and levels installed:

|                        |   |
|------------------------|---|
| <b>HMC</b>             | V8R8.4.0.0. Early ship code.  |
| <b>VIOS</b>            | 2.2.3.52  |
| <b>AIX</b>             | 7200-00-00-0000. Early ship code.   |
| <b>Disks</b>           | 3 disks. 1 x rootvg and 2 x spare disks. SAN LUNs presented via SAN Volume Controller (SVC) to both VIOS. |
| <b>Dual VIOS</b>       | s824vio1 and s824vio2.  |
| <b>Server</b>          | POWER8 S824.  |
| <b>System Firmware</b> | SV810_081, FW810.10.  |

## Performing AIX Live Updates on AIX

In the following example I will show you how to patch a live AIX system using AIX Live Update. I'll start with an unpatched AIX 7.2 system. There are no ifixes installed. I also checked how much spare capacity (CPU and memory) I would need available on my system before starting the process. In this case I'd need 0.1 processing units and 2GB of memory. Support for live update is provided by the **bos.liveupdate.rte** fileset.

```
# oslevel -s
7200-00-00-0000
# lslpp -L bos.liveupdate.rte
  Filesset              Level  State  Type  Description (Uninstaller)
  -----
  bos.liveupdate.rte    7.2.0.0  C    F    Live Update Runtime

# emgr -l
There is no efix data on this system.

# lparstat -i | egrep 'Online Mem|Entitled Capacity '
Entitled Capacity      : 0.10
Online Memory          : 2048 MB
```

The AIX development team provided a “dummy” ifix for me to use in my testing during the ESP. I'll use this ifix in my example. You can determine if an ifix is AIX Live Update capable by previewing the fix with the **emgr** command and checking the '**LU CAPABLE**' attribute.

```
# ls -ltr /tmp/cg/dummy/dummy.150813.epkg.Z
-rw-r--r--  1 root      system      35625 Aug 20 16:30 dummy.150813.epkg.Z

root@aix721534A / # emgr -p -e /tmp/cg/dummy/dummy.150813.epkg.Z | grep LU
LU CAPABLE:      yes
```

Three disks were assigned to my AIX partition. All were VSCSI hdisks, backed by SAN LUNs on both of my VIOS. There were two spare disks (hdisk2 and hdisk3) which I could use with the AIX Live Update process. All disks were 50GB in size.

```
# lsdev -Cc disk
hdisk1 Available  Virtual SCSI Disk Drive
hdisk2 Available  Virtual SCSI Disk Drive
hdisk3 Available  Virtual SCSI Disk Drive

# lspv
hdisk1          00f94f58cc2de28d          rootvg          active
hdisk2          00f94f58c90ae057          None
hdisk3          00f94f58c90ae0db          None

# echo cvai | kdb -script
read VSCSI_scsi_ptrs OK, ptr = 0xF1000000C0126DF0
(0)> cvai
Executing cvai command
NAME          STATE      CMDS_ACTIVE  ACTIVE_QUEUE  HOST
vscsi0        0x000007  0x0000000000 0x0           s824vio1->vhost3
vscsi1        0x000007  0x0000000000 0x0           s824vio2->vhost5

End of execution for cvai command
(0)> Executing q command

# getconf DISK_SIZE /dev/hdisk1
51200
```

```
# getconf DISK_SIZE /dev/hdisk2
51200
```

```
# getconf DISK_SIZE /dev/hdisk3
51200
```

I confirmed that there were two paths to each of my disks and that each disk was mapped across both of my VIOS.

```
root@aix721534A / # lspath
Enabled hdisk1 vscsi1
Enabled hdisk2 vscsi1
Enabled hdisk3 vscsi1
Enabled hdisk1 vscsi0
Enabled hdisk2 vscsi0
Enabled hdisk3 vscsi0
```

```
$ hostname
s824vio1
```

```
$ lsmap -vadapter vhost3
SVSA          Physloc          Client Partition ID
-----
vhost3        U8286.42A.214F58V-V1-C90      0x00000063
```

```
VTD          vtscsi0
Status       Available
LUN          0x8100000000000000
Backing device hdisk7
Physloc      U78C9.001.WZS01K8-P1-C6-T1-W5005076801404F98-L6000000000000
Mirrored     false
```

```
VTD          vtscsi5
Status       Available
LUN          0x8200000000000000
Backing device hdisk8
Physloc      U78C9.001.WZS01K8-P1-C6-T1-W5005076801404F98-L8000000000000
Mirrored     false
```

```
VTD          vtscsi6
Status       Available
LUN          0x8300000000000000
Backing device hdisk9
Physloc      U78C9.001.WZS01K8-P1-C6-T1-W5005076801405173-L7000000000000
Mirrored     false
```

```
$ hostname
s824vio2
```

```
$ lsmap -vadapter vhost5
SVSA          Physloc          Client Partition ID
-----
vhost5        U8286.42A.214F58V-V2-C91      0x00000063
```

```
VTD          vtscsi0
Status       Available
LUN          0x8100000000000000
Backing device hdisk7
Physloc      U78C9.001.WZS01K8-P1-C3-T1-W5005076801404F98-L6000000000000
Mirrored     false
```

```
VTD          vtscsi5
Status       Available
LUN          0x8200000000000000
Backing device hdisk8
```



```

Physloc          U78C9.001.WZS01K8-P1-C3-T1-W5005076801404F98-L8000000000000
Mirrored        false

VTD             vtscsi8
Status          Available
LUN             0x8300000000000000
Backing device  hdisk9
Physloc          U78C9.001.WZS01K8-P1-C3-T1-W5005076801405173-L7000000000000
Mirrored        false

```

I ensured that none of my file systems were full, particularly /var, where the AIX Live Update process would log all of its activity.

```

# df -g
Filesystem      GB blocks   Free %Used   Iused %Iused Mounted on
/dev/hd4        0.36       0.19  48%     11436  21% /
/dev/hd2        2.28       0.20  92%     41276  45% /usr
/dev/hd9var     1.34       1.13  17%     3739   2% /var
/dev/hd3        1.09       1.01   8%       766   1% /tmp
/dev/hd1        0.02       0.02   3%         7   1% /home
/dev/hd11admin  0.12       0.12   1%         5   1% /admin
/proc           -          -    -         -    - /proc
/dev/hd10opt    0.03       0.02  52%      223   6% /opt
/dev/livedump   0.25       0.25   1%         4   1% /var/adm/ras/livedump

```

The AIX Live Update process must be able to communicate with your HMC in order for it to control the original and surrogate LPARs. The root user must be able to authenticate to the HMC that manages the partition. You can authenticate to the HMC by using the **hmcauth** command or by defining an HMC object through NIM. The hmcclientliveupdate HMC role has all the privileges that are required for the AIX Live Update operation. If a user is defined on the HMC with this role, the authentication can be done with this user rather than the hscroot user. In my lab, I ran the **hmcauth** command to authenticate with my HMC as the hscroot user (as shown below).

```

# hmcauth -a hsc02 -u hscroot -p abc1234
#

# hmcauth -l
Address   : 10.1.50.30
User name: hscroot
port     : 12443
TTL      : 20:25:41 left

```

I configured the AIX Live Update data file with the appropriate information for my environment. I set the **mode** to preview. I set **nhdisk** to hdisk2; this is the name of the disk that will be used to create a copy of the original rootvg and boot the surrogate partition. I set **mhdisk** to hdisk3; this is the name of the disk used for the mirrored rootvg on the surrogate partition. I chose 98 as the new **lpar\_id** for the surrogate partition and I entered the appropriate HMC user name (**user**) and IP address (**management\_console**) for my environment.

```

# vi /var/adm/ras/liveupdate/lvupdate.data

general:
    mode = preview
    kext_check = no

disks:

```

```
nhdisk = hdisk2
mhdisk = hdisk3
tohdisk =
tshdisk =
```

hmc:

```
lpar_id = 98
management_console = 10.1.50.30
user = hscroot
```

I performed a preview of the AIX Live Update operation using the **geninstall** command with the **-k** and **-p** flags. This performed a number of “pre-flight” checks, each of which passed.

```
# geninstall -k -p -d /tmp/cg/dummy dummy.150813.epkg.Z
```

Validating live update input data.

Computing the estimated time for the live update operation:

```
-----
LPAR: aix721534A
Mode: F
Blackout_time(s): 77
Global_time(s): 624
```

Checking mirror vg device size:

```
-----
Required device size: 7216 MB
Given device size: 51200 MB
PASSED: device size is sufficient.
```

Checking new root vg device size:

```
-----
Required device size: 7216 MB
Given device size: 51200 MB
PASSED: device size is sufficient.
```

Checking temporary paging space device size:

```
-----
Required device size: 512 MB
```

Checking temporary dump device size:

```
-----
Required device size: 100 MB
```

Validating the adapters and their paths:

```
-----
PASSED: adapters can be divided into two sets so that each has paths to all disks.
```

Checking other requirements:

```
-----
PASSED: sufficient space available in /var.
PASSED: sufficient space available in /.
PASSED: no existing altinst_rootvg.
PASSED: rootvg is not part of a snapshot.
PASSED: pkcs11 is not installed.
PASSED: rootvg is not part of a snapshot.
PASSED: The trustchk Trusted Execution Policy is not on.
PASSED: The trustchk Trusted Library Policy is not on.
PASSED: The trustchk TSD_FILES_LOCK policy is not on.
PASSED: the boot disk is set to the current rootvg.
PASSED: the mirrorvg name is available.
PASSED: the rootvg is uniformly mirrored.
PASSED: the rootvg does not have the maximum number of mirror copies.
PASSED: the rootvg does not have stale logical volumes.
PASSED: all of the mounted file systems are of a supported type.
```

```

PASSED: this AIX instance is not diskless.
PASSED: no Kerberos configured for NFS mounts.
PASSED: multibos environment not present.
PASSED: Trusted Computing Base not defined.
PASSED: no local tape devices found.
PASSED: live update not executed from console.
PASSED: the execution environment is valid.
PASSED: enough available space for /var to dump Component Trace buffers.
PASSED: enough available space for /var to dump Light weight memory Trace buffers.
PASSED: all devices are virtual devices.
PASSED: No active workload partition found.
PASSED: nfs configuration supported.
PASSED: HMC token is present.
PASSED: HMC token is valid.
PASSED: HMC requests successful.
PASSED: Provided LPAR ID is available.
PASSED: A virtual slot is available.
PASSED: RSCT daemons are active.
PASSED: no Kerberos configuration.
PASSED: lpar is not remote restart capable.
PASSED: no virtual log device configured.
PASSED: lpar is not using shared memory or ams resources are available.
PASSED: the disk configuration is supported.
PASSED: no Generic Routing Encapsulation (GRE) tunnel configured.
PASSED: Firmware level is supported.
PASSED: vNIC resources available.
INFO: Any system dumps present in the current dump logical volumes will not be available after
live update is complete.

```

To perform the actual Live Update operation I first needed to change the mode from **preview** to **automated** in the **lvupdate.data** configuration file. Then I was able to start the Live Update process on my running system. The need to manually change the mode was required only during the ESP and should not be necessary when AIX 7.2 becomes generally available. I watched as the process first validated the LPAR environment was suitable for live updates, and then it created a clone of the root volume group for booting the surrogate. The surrogate LPAR was booted next and a mirror of the original rootvg created and assigned to the surrogate, followed by a migration of the running workload to the surrogate partition. After the blackout time had ended, the workload was now running on the surrogate partition. At this point the original LPAR was shut down and deleted.

```
# perl -pi -e 's/= preview/= automated/g' /var/adm/ras/liveupdate/lvupdate.data
```

```
# cat /var/adm/ras/liveupdate/lvupdate.data
```

```
general:
```

```
    mode = automated
```

```
# geninstall -k -d /tmp/cg/dummy dummy.150813.epkg.Z
```

```
Validating live update input data.
```

```
Computing the estimated time for the live update operation:
```

```
-----
LPAR: aix721534A
Mode: F
Blackout_time(s): 77
Global_time(s): 670
```

```
Checking mirror vg device size:
```

```
-----
Required device size: 7216 MB
Given device size: 51200 MB
```

PASSED: device size is sufficient.

Checking new root vg device size:

-----

Required device size: 7216 MB

Given device size: 51200 MB

PASSED: device size is sufficient.

Checking temporary paging space device size:

-----

Required device size: 512 MB

Checking temporary dump device size:

-----

Required device size: 100 MB

Validating the adapters and their paths:

-----

PASSED: adapters can be divided into two sets so that each has paths to all disks.

Checking other requirements:

-----

PASSED: sufficient space available in /var.

PASSED: sufficient space available in /.

PASSED: no existing altinst\_rootvg.

PASSED: rootvg is not part of a snapshot.

PASSED: pkcs11 is not installed.

PASSED: rootvg is not part of a snapshot.

PASSED: The trustchk Trusted Execution Policy is not on.

PASSED: The trustchk Trusted Library Policy is not on.

PASSED: The trustchk TSD\_FILES\_LOCK policy is not on.

PASSED: the boot disk is set to the current rootvg.

PASSED: the mirrorvg name is available.

PASSED: the rootvg is uniformly mirrored.

PASSED: the rootvg does not have the maximum number of mirror copies.

PASSED: the rootvg does not have stale logical volumes.

PASSED: all of the mounted file systems are of a supported type.

PASSED: this AIX instance is not diskless.

PASSED: no Kerberos configured for NFS mounts.

PASSED: multibos environment not present.

PASSED: Trusted Computing Base not defined.

PASSED: no local tape devices found.

PASSED: live update not executed from console.

PASSED: the execution environment is valid.

PASSED: enough available space for /var to dump Component Trace buffers.

PASSED: enough available space for /var to dump Light weight memory Trace buffer

PASSED: all devices are virtual devices.

PASSED: No active workload partition found.

PASSED: nfs configuration supported.

PASSED: HMC token is present.

PASSED: HMC token is valid.

PASSED: HMC requests successful.

PASSED: Provided LPAR ID is available.

PASSED: A virtual slot is available.

PASSED: RSCT daemons are active.

PASSED: no Kerberos configuration.

PASSED: lpar is not remote restart capable.

PASSED: no virtual log device configured.

PASSED: lpar is not using shared memory or ams resources are available.

PASSED: the disk configuration is supported.

PASSED: no Generic Routing Encapsulation (GRE) tunnel configured.

PASSED: Firmware level is supported.

PASSED: vNIC resources available.

INFO: Any system dumps present in the current dump logical volumes will not be a

Non-interruptable live update operation begins in 10 seconds.

Broadcast message from root@aix721534A (pts/0) at 09:28:43 ...

Live AIX update in progress.

```
.....
Initializing live update on original LPAR.

Validating original LPAR environment.

Beginning live update operation on original LPAR.

Requesting resources required for live update.
.....
Notifying applications of impending live update.

Creating rootvg for boot of surrogate.
.....
Starting the surrogate LPAR.
.....
Creating mirror of original LPAR's rootvg.
.....
Moving workload to surrogate LPAR.
.....
        Blackout Time started.
.....
        Blackout Time end.

Workload is running on surrogate LPAR.
.....
Shutting down the Original LPAR.
.....
The live update operation succeeded.
```

Broadcast message from root@aix721534A (pts/0) at 09:40:40 ...

**Live AIX update completed.**

The AIX Live Update process completed successfully and the ifix was installed, as shown by the **emgr** output below.

```
# emgr -l

ID  STATE LABEL          INSTALL TIME          UPDATED BY ABSTRACT
===  =====
1   ST   dummy          09/14/15 09:38:56          Test fix for Live Update

STATE codes:
S = STABLE
M = MOUNTED
U = UNMOUNTED
Q = REBOOT REQUIRED
B = BROKEN
I = INSTALLING
R = REMOVING
T = TESTED
P = PATCHED
N = NOT PATCHED
SP = STABLE + PATCHED
SN = STABLE + NOT PATCHED
QP = BOOT IMAGE MODIFIED + PATCHED
```

QN = BOOT IMAGE MODIFIED + NOT PATCHED  
 RQ = REMOVING + REBOOT REQUIRED

The AIX error report showed messages for the start and successful completion of the AIX Live Update operation. The entire process took approximately 11 minutes to complete, all the while my workload remained active and was not disrupted.

```
# errpt -NLVUPDATE
IDENTIFIER  TIMESTAMP    T C RESOURCE_NAME  DESCRIPTION
12295E0B   0922152915  I S LVUPDATE      Live AIX update completed successfully
9A74C7AB   0922151815  I S LVUPDATE      Live AIX update started
```

While the AIX Live Update operation was taking place, I noticed on my HMC, that a new AIX partition was created, with id 98. The original LPAR was renamed and **\_lku0** appended to its name.




Figure 2 – New surrogate partition created. Original partition renamed with **\_lku0** appended.




Figure 3 – Surrogate partition booting.




Figure 4 – AIX started on the surrogate partition.



Figure 5 – Surrogate and original partitions both running.



Figure 6 – Moving workload from original to surrogate partition. Signified by led 000e on both LPARs.




Figure 7 – Original partition shutdown after successful workload migration to surrogate partition.




Figure 8 – Original partition deleted. New partition running the workload. New partition ID.

From the AIX command line, I was able to verify the partition ID had changed, as expected, from 99 to 98, using the **uname** command (as shown below).

#### Before AIX Live Update

```
# uname -L
99 AIX72_Build_1534A
```

#### After AIX Live Update

```
# uname -L
98 AIX72_Build_1534A
```

The “black out” time was short. I noticed my **topas** sessions “paused” for a period of time (27 seconds) while the workload migration to the new partition was performed.

```
Topas Monitor for host:aix721534A
Mon Sep 14 15:13:20 2015 Interval:2

EVENTS/QUEUES FILE/TTY
Cswitch 415 Readch 97.8K
Syscall 1914 Writech 14387
CPU User% Kern% Wait% Idle% Phyc Entc% Reads 79 Rawin 0
Total 5.0 14.3 0.1 80.6 0.04 37.034 Writes 121 Ttyout 600
Network BPS I-Pkts O-Pkts B-In B-Out Execs 3 Namei 140
Total 20.1K 116.5 171.0 8.49K 11.6K Runqueue 0.50 Dirblk 0
Waitqueue 0.0
Disk Busy% BPS TPS B-Read B-Writ MEMORY
Total 1.7 177K 40.50 76.0K 101K PAGING Real,MB 2048
Faults 843 % Comp 56
FileSystem BPS TPS B-Read B-Writ Steals 0 % Noncomp 25
Total 110K 174.5 97.7K 12.4K PgpsIn 0 % Client 25
PgpsOut 0
Name PID CPU% PgSp Owner PageIn 24 PAGING SPACE
java 8913316 1.1 64.9M root PageOut 0 Size,MB 512
topas 8782148 0.5 3.80M root Sios 25 % Used 2
getty 7274990 0.3 656K root % Free 98
sshd 9044346 0.2 1.23M root
clcomd 8126756 0.1 1.73M root NFS (calls/sec)
iostat 4653484 0.1 508K root SerV2 0 WPAR Activ 0
gil 1507634 0.1 960K root CliV2 0 WPAR Total 0
sshd 5439772 0.0 1.23M root SerV3 0 Press: "h"-help
Cliv3 0 0 0 Cliv3 0 0 0 Quit

Topas Monitor for host:aix721534A
Mon Sep 14 15:13:47 2015 Interval:2

EVENTS/QUEUES FILE/TTY
Cswitch 29915 Readch 18.1M
```

```

845                               Syscall  837762 Writech  16.2M
080   User% Kern% Wait% Idle%   Phyc  Entc%  Reads   6766 Rawin    0
Total   13.5 42.6  2.9 41.0    1.38 1381.75 Writes  2769 Ttyout  752
                               Forks    155 Igets   192
Network   BPS  I-Pkts  O-Pkts   B-In  B-Out  Execs   155 Namei   2412
Total   22.6M  5.27K  5.27K  21.2M  1.46M Runqueue 0.50 Dirblk   0
                               Waitqueue 0.0

Disk   Busy%      BPS      TPS  B-Read  B-Writ
Total   6.2      21.5M   506.5  7.05M  14.4M
                               MEMORY
                               Real,MB  2048
Filesystem      BPS      TPS  B-Read  B-Writ  PAGING
Total           0        0        0        0  Faults  33279 % Comp  40
                               Steals  0 % Noncomp 11
                               Pgspln  5127 % Client 11
                               Pgspln  0
Name           PID  CPU%  PgSp  Owner  PageIn  7903  PAGING SPACE
ksh           6881644  7.5  800K  root  PageOut  3537  Size,MB  512
sbiod         4391024  7.2  896K  root  Sios    11441  % Used  2
mcr           4522026  4.6  1.75M  root  % Free  98
rtcmd         3866794  1.8  512K  root  NFS (calls/sec)
vtiol         786456  1.3  704K  root  SerV2   0  WPAR Activ  0
j2pg          1966152  0.1  3.62M  root  Cliv2   0  WPAR Total  0
topas         8782148  0.1  3.81M  root  SerV3   0  Press: "h"-help
sshd          9044346  0.0  1.23M  root  Cliv3   0  0  5126uit

```

I also noticed that during the AIX Live Update process, as expected, only a single path was configured for my disks on the original LPAR. The other path was configured on the surrogate LPAR at this time, as the AIX Live Update process booted the new LPAR.

```

# lspath
Enabled hdisk1 vscsi0
Enabled hdisk2 vscsi0
Enabled hdisk3 vscsi0

```

On the second VIOS, s824vio2, I noticed that the vhost5 adapter was now mapped to the new partition ID 62 (98 in decimal).

```

$ hostname
s824vio2

```

```

$ lsmap -vadapter vhost5
SVSA          Physloc          Client Partition ID
-----
vhost5        U8286.42A.214F58V-V2-C91  0x00000062

VTD           vtscsi0
Status        Available
LUN           0x8100000000000000
Backing device hdisk7
Physloc       U78C9.001.WZS01K8-P1-C3-T1-W5005076801404F98-L6000000000000
Mirrored      false

VTD           vtscsi5
Status        Available
LUN           0x8200000000000000
Backing device hdisk8
Physloc       U78C9.001.WZS01K8-P1-C3-T1-W5005076801404F98-L8000000000000
Mirrored      false

VTD           vtscsi8
Status        Available
LUN           0x8300000000000000
Backing device hdisk9
Physloc       U78C9.001.WZS01K8-P1-C3-T1-W5005076801405173-L7000000000000

```



Mirrored false

After the AIX Live Update process finished, the vhost adapter mapping on both VIOS showed all disks mapped to the new partition ID. The partition ID changed from 99 (63 in Hex) to 98 (62 in Hex). And all of the disk paths on the LPAR were restored to two paths per disk (as shown by the **lspath** command on the AIX partition).

```
$ hostname
s824vio1
$ lsmap -vadapter vhost3
SVSA          Physloc          Client Partition ID
-----
vhost3        U8286.42A.214F58V-V1-C90      0x00000062

VTD           vtscsi0
Status        Available
LUN           0x8100000000000000
Backing device hdisk7
Physloc       U78C9.001.WZS01K8-P1-C6-T1-W5005076801404F98-L6000000000000
Mirrored      false

VTD           vtscsi5
Status        Available
LUN           0x8200000000000000
Backing device hdisk8
Physloc       U78C9.001.WZS01K8-P1-C6-T1-W5005076801404F98-L8000000000000
Mirrored      false

VTD           vtscsi6
Status        Available
LUN           0x8300000000000000
Backing device hdisk9
Physloc       U78C9.001.WZS01K8-P1-C6-T1-W5005076801405173-L7000000000000
Mirrored      false

$ hostname
s824vio2
$ lsmap -vadapter vhost5
SVSA          Physloc          Client Partition ID
-----
vhost5        U8286.42A.214F58V-V2-C91      0x00000062

VTD           vtscsi0
Status        Available
LUN           0x8100000000000000
Backing device hdisk7
Physloc       U78C9.001.WZS01K8-P1-C3-T1-W5005076801404F98-L6000000000000
Mirrored      false

VTD           vtscsi5
Status        Available
LUN           0x8200000000000000
Backing device hdisk8
Physloc       U78C9.001.WZS01K8-P1-C3-T1-W5005076801404F98-L8000000000000
Mirrored      false

VTD           vtscsi8
Status        Available
LUN           0x8300000000000000
Backing device hdisk9
Physloc       U78C9.001.WZS01K8-P1-C3-T1-W5005076801405173-L7000000000000

# lspath
```

```

Enabled hdisk1 vscsi0
Enabled hdisk2 vscsi0
Enabled hdisk3 vscsi0
Enabled hdisk1 vscsi1
Enabled hdisk2 vscsi1
Enabled hdisk3 vscsi1

```

All AIX Live Update operations are logged to the `/var/adm/ras/liveupdate/logs` directory. You can find a full set of detailed logs in this location, useful for troubleshooting Live Update operations. The `snap` command can also be used (with the `-U` flag) to collect AIX Live Update information and save it in the `/tmp/ibmsupt/liveupdate` directory, which could be shared with IBM support if they need to assist you in troubleshooting an issue with AIX Live Update.

```

# ls -ltr
total 71512
-rw-r--r--  1 root      system      230727 Sep 12 06:23 oplhmclog.2015-09-12_06:23:04.364
-rw-r--r--  1 root      system      165165 Sep 12 06:23 lvupdlog.2015-09-12_06:23:04.364
-rw-r--r--  1 root      system      230727 Sep 12 06:25 oplhmclog.2015-09-12_06:25:21.559
-rw-r--r--  1 root      system      3043257 Sep 12 06:27 olhmclog.2015-09-12_06:25:21.559
-rw-r--r--  1 root      system        36357 Sep 12 06:27 olvupdlog.2015-09-12_06:25:21.559
-rw-r--r--  1 root      system     3310341 Sep 12 06:27 otlog.2015-09-12_06:25:21.559
-rw-r--r--  1 root      system     3490422 Sep 12 06:27 lvupdlog.2015-09-12_06:25:21.559
..etc...
-rw-r--r--  1 root      system      230769 Sep 21 15:07 oplhmclog
-rw-r--r--  1 root      system      165260 Sep 21 15:07 lvupdlog

```

If, for some reason, I needed to return the system to its previous state before the AIX Live Update operation, the fastest way to achieve this is to boot the partition from the original rootvg disk. All I need do is set the boot list to point to the original rootvg hdisk and reboot the system. The system will restart, unpatched.

```

# lspv
hdisk1      00f94f58cc2de28d      rootvg      active < Patched rootvg disk
hdisk2      00f94f58c90ae057      None        < Original rootvg disk
hdisk3      00f94f58c90ae0db      lvup_rootvg < Surrogate boot disk

# bootlist -m normal -o
hdisk1 blv=hd5 pathid=1

# bootlist -m normal hdisk2

# shutdown -Fr

# emgr -l
There is no efix data on this system.

# lspv
hdisk1      00f94f58cc2de28d      None        < Patched rootvg disk
hdisk2      00f94f58c90ae057      rootvg      active < Original rootvg disk
hdisk3      00f94f58c90ae0db      None        < Surrogate boot disk

```

## Conclusion

The aim of this article was to introduce and demonstrate the new AIX Live Update feature of the AIX 7.2 operating system. In the example above I showed you how to perform an AIX Live Update using the AIX command line tool `geninstall`. More information can be found in the AIX Knowledge Centre. AIX Live Update is a powerful new feature which will allow customers to take another step towards

their continuous availability goals for their mission critical, production AIX environments. I look forward to the future enhancements and capabilities that are yet to come with this tool.

## **Acknowledgement**

The author sincerely acknowledges David Sheffield (IBM Senior Technical Staff Member, AIX Operating System Architect) for reviewing this article and providing his valuable suggestions and feedback.

**Job Title:** AIX and Power Systems Client Technical Specialist

**Email:** cgibson@au1.ibm.com

**Bio:** Chris Gibson is a Power Systems Client Technical Specialist at IBM Systems. Located in Melbourne, Australia, he has co-authored several IBM Redbooks® on AIX. Chris contributes to the AIX community through his [AIX blog](#) and Twitter (@cgibbo).

<https://www.ibm.com/developerworks/community/blogs/cgaix?lang=en>

**Abstract:** This article describes how to use the new AIX Live Update feature of the AIX® 7.2 operating system. You will learn how to avoid downtime on your mission critical workloads when patching the AIX OS. General concepts will be discussed, followed by an example of how to use the tool.