

# IBM AIX 7.2 Live Update

*Experiences testing AIX 7.2 Live Update feature with Oracle Database  
12c Release 2*



*Bhargavaram Akula  
Sougata Sarkar  
IBM Systems Group  
April 2019*



## Table of contents

<b>Abstract</b> .....	<b>1</b>
<b>Introduction</b> .....	<b>1</b>
<b>Terminology</b> .....	<b>1</b>
<b>Prerequisites</b> .....	<b>2</b>
<b>Live Update restrictions</b> .....	<b>2</b>
<b>Best practices for Live Update</b> .....	<b>2</b>
<b>System topology</b> .....	<b>4</b>
<b>Live Update steps</b> .....	<b>5</b>
<b>Test description – stress / destructive tests</b> .....	<b>9</b>
<b>Performance measurements</b> .....	<b>12</b>
Database and Workload Considerations .....	12
<b>Conclusion</b> .....	<b>14</b>
<b>Resources</b> .....	<b>15</b>
<b>About the authors</b> .....	<b>16</b>
<b>Trademarks and special notices</b> .....	<b>17</b>



## Abstract

---

*This paper describes the findings of certification testing with Oracle Database (single instance) in conjunction with the AIX® Live Update (LU) capability. The objective of this certification was to validate the compatibility of the AIX 7.2 Live Update feature with Oracle Database 12c Release 2 RDBMS and Oracle ASM. This paper will introduce the concepts used to test the stability of such a solution and lists the best practices. We also observe that all tests and the certification were done with the single instance version of the Oracle DB. Use of Live Update with the Oracle RAC feature is not supported.*

## Introduction

---

IBM® delivers kernel fixes in the form of interim fixes 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 restarted. To address this issue, AIX Version 7.1, and earlier, provided concurrent update-enabled interim fixes that allow deployment of some limited kernel fixes to a running LPAR. All fixes cannot be delivered as concurrent update-enabled interim fixes.

Starting with AIX Version 7.2, IBM offers Live Update function which eliminates the downtime associated with patching the operating system. The workloads on the system are not stopped in a Live Update operation, yet the workloads can use the interim fixes after the Live Update operation.

- AIX Version 7.2 Service Pack 1 contains significant fixes for AIX Live Update.
- With AIX Version 7.2 at the 7200-01 Technology Level or later, you can use the Live Update function to update service packs and technology levels for the AIX operating system.
- With AIX Version 7.2 at the 7200-02 Technology Level or later, partitions that are managed by IBM Power® Virtualization Center (PowerVC) can also use the Live Update function.

We have used AIX 7.2 TL01 SP2 operating system in our testing.

## Terminology

---

**Original partition:** The logical partition (LPAR) on which the Live Update operation is initiated.

**Surrogate partition:** The LPAR that gets created automatically during a Live Update operation and which is identical to the original LPAR.

**Checkpointing:** A process in which the workload is paused and its current state is saved. Once the checkpointing is complete, all processes are migrated and restarted (unpaused) on the surrogate partition.

**Mobility:** Process of moving workload/application from one partition to another.

**Blackout Time or Blackout Period:** The duration of time in which the workload is paused on the original LPAR and then restarted on the surrogate LPAR.



## Prerequisites

---

Before performing the live update operation, please review the following list of hardware and software requirements, Live Update prerequisites are listed in the URL below:

[https://www.ibm.com/support/knowledgecenter/en/ssw\\_aix\\_72/com.ibm.aix.install/lvupdate\\_requisite.htm](https://www.ibm.com/support/knowledgecenter/en/ssw_aix_72/com.ibm.aix.install/lvupdate_requisite.htm)

The LPAR requirements for performing Live Update operations are listed in the URL below:

[https://www.ibm.com/support/knowledgecenter/ssw\\_aix\\_72/com.ibm.aix.install/lvupdate\\_limitations.htm](https://www.ibm.com/support/knowledgecenter/ssw_aix_72/com.ibm.aix.install/lvupdate_limitations.htm)

## Live Update restrictions

---

- The Live update operation does not support Oracle Real Application Clusters (RAC).
- Even though the Live Update operation supports Oracle ASM, it does not support Oracle ACFS (Automatic Storage Management Cluster File System).
- Live Update does not support the kernel extension “pw-syscall” which is a part of legacy Oracle Database versions 9i and 10g.
- The Live Update feature does not support Public-Key Cryptography Standards # 11 (PKCS11). The security.pkcs11 fileset cannot be installed. This fileset is not an Oracle Database prerequisite when installing on AIX 7.2.

The detailed list of Live Update restrictions are documented at the following URL:

[https://www.ibm.com/support/knowledgecenter/ssw\\_aix\\_72/com.ibm.aix.install/lvupdate\\_detail\\_restrict.htm](https://www.ibm.com/support/knowledgecenter/ssw_aix_72/com.ibm.aix.install/lvupdate_detail_restrict.htm)

## Best practices for Live Update

---

- Take a backup of rootvg using utilities alt\_disk\_copy or mksysb before performing Live Update operation.
- Make a copy of the current partition profiles on the Hardware Management Console (HMC) before performing a Live Update operation.
- Before starting an actual Live Update operation do a preview check to validate the user inputs and make sure all checks are passed. Preview mode helps in detecting unsupported kernel extensions.
- Remove the missing disk paths using the rmpath command.
- Unload the Oracle ACFS drivers before performing the Live Update operation. Sometimes ACFS drivers might get loaded unexpectedly as part of quarterly Release Updates (RU)/PSU patch.

1. If ACFS drivers are loaded, we will see the below output:

```
# genkex |grep oracle
f1000000c06ff000    3cb000 /usr/lib/drivers/oracleacfs.ext
f1000000c0619000    e6000  /usr/lib/drivers/oracleadvn.ext
6610000            f0000  /usr/lib/drivers/oracleoks.ext
```



```
$ORACLE_HOME/bin/acfsdriverstate loaded  
ACFS-9203: true
```

2. To unload ACFS drivers, run the following command:

```
# $ORACLE_HOME/bin/acfsload stop
```

3. If ACFS drivers are unloaded, we will see the below output:

```
$ORACLE_HOME/bin/acfsdriverstate loaded  
ACFS-9204: false
```

- For failed live update operations use the clvupdate command for automatic cleanup of the system.
- You can refer the below best practices URL for more details:  
[https://www.ibm.com/support/knowledgecenter/ssw\\_aix\\_72/com.ibm.aix.install/lvupdate\\_best\\_practice.htm](https://www.ibm.com/support/knowledgecenter/ssw_aix_72/com.ibm.aix.install/lvupdate_best_practice.htm)

## System topology

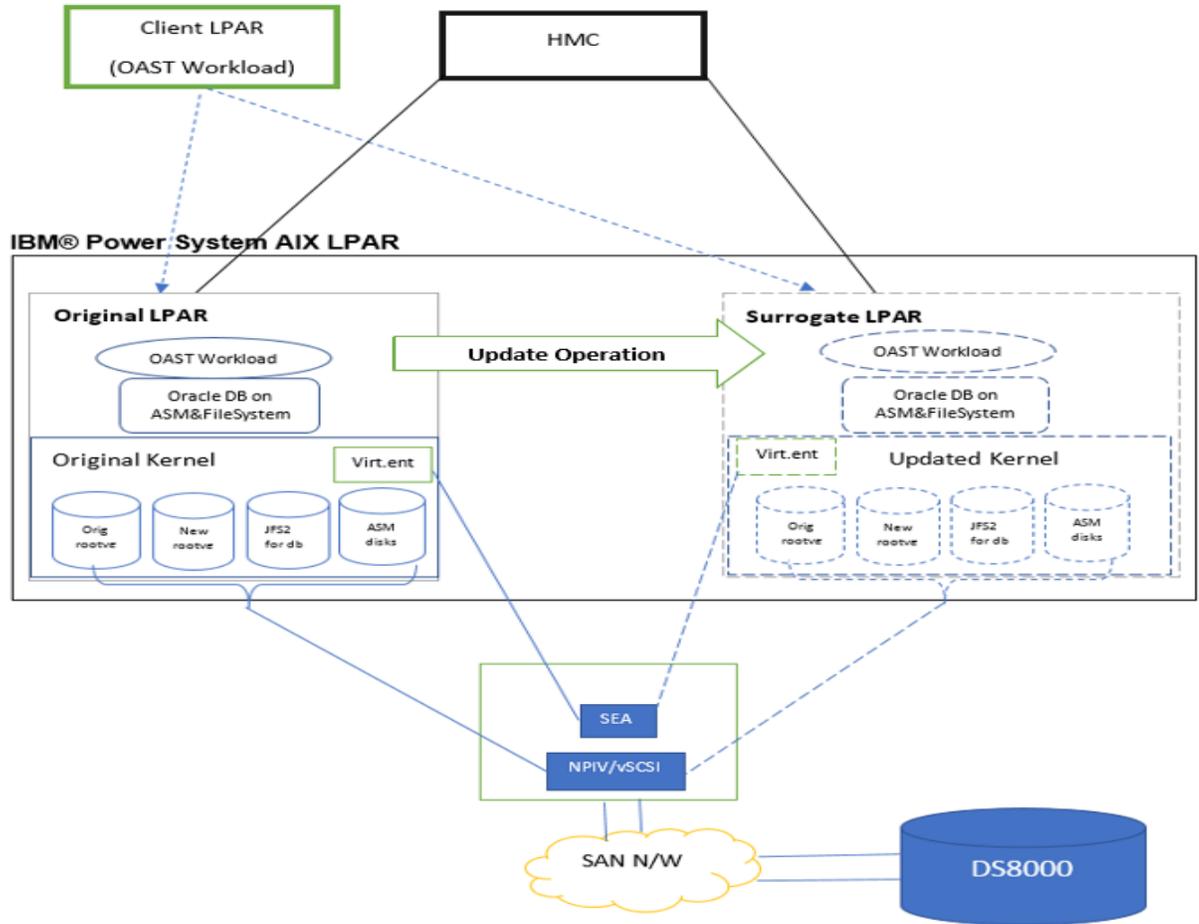


Figure 1. Sample System Configuration

An IBM Power E870 (9119-MME) system was used for our testing. An AIX LPAR with 4 cores and 48 GB memory is configured to be the original LPAR. The table below shows the details of Hardware and Software configuration for the various components used for this testing.



Component	Specifications
Server	IBM Power E870 (9119-MME)
System firmware	IBM, FW860.50 (SC860_160)
Memory and swap	RAM 48GB and swap 16G/+
HMC	8.6.0 SP3
VIOS	Dual VIOS level 2.2.5.20
Operating system	AIX 7.2 TL1 SP02 (7200-01-02-1717) 64-bit kernel
Oracle software	Standalone Grid Infrastructure (GI) 12.2 Oracle Database 12.2
Software patches	AIX – Below fixes were configured* APAR #IV89963, APAR IJ10045, APAR IJ10082 Oracle – OCT Release Update(RU) 2018
Storage configuration	All I/O devices must be virtualized (virtual Ethernet, Virtual Small Computer System Interface (VSCSI) or N-Port Id Virtualization (NPIV) with AIX multipath I/O (MPIO))  Oracle Datafiles are staged on both File System (JFS2) and ASM

\*Note: All above APARs are available in AIX 7.2 TL01 SP4

## Live Update steps

Below are the steps for performing the Live Update operation:

1. The AIX Live Update process must be able to communicate with the HMC in-order for it to control the original and the surrogate LPARs. The root user must be able to authenticate to the HMC and this can be done using the hmcauth command. HMC user should have the hmclientliveupdate role which has all the privileges that are required for the AIX Live Update operation. Port 12443 must also be opened between the original LPAR and the HMC.

```

hmcauth -a <hmchost> -u <hmcuser> -p <passwd>

Eg. # hmcauth -u hscroot -a X.X.X.X -p xxxxxx

+ Check HMC token status

# hmcauth -l
  Address   : X.X.X.X
  User name : hscroot
  port      : 12443

```

2. Configure Live Update datafile “/var/adm/ras/liveupdate/lvupdate.data”. The Live Update operation requires at least two additional disks. The first disk is required for the initial boot disk of the surrogate partition. The second disk is required to create an additional mirror of the root volume group.

nhdisk -> copy of the original rootvg which will be used to boot the Surrogate (surr-boot-rootvg)

mhdisk -> used for the mirrored rootvg (surr-mir-rootvg) on the Surrogate

tohdisk & tshdisk -> used for paging space or dump devices if they are not part of rootvg

```

Eg. /var/adm/ras/liveupdate/lvupdate.data

general:
kext_check = yes
disks:
      nhdisk = hdisk76
      mhdisk = hdisk77
      tohdisk =
      tshdisk =

hmc:
      lpar_id =
      management_console = X.X.X.X
      user = hscroot

```

3. LU preview check can be done using geninstall utility. All checks should be passed before performing the live update operation. In preview mode all configuration checks are validated without applying the update.

```

#geninstall -k -p

  Validating live update input data.
  Computing the estimated time for the live update operation:
  -----
  LPAR: XXXX.XXXX.XXXX

```

```

Blackout_time(s): 31
Global_time(s): 3453
Checking mirror vg device size:
-----
Required device size: 40896 MB
Given device size: 40959 MB
PASSED: device size is sufficient.
Checking new root vg device size:
-----
Required device size: 40896 MB
Given device size: 40959 MB
PASSED: device size is sufficient.
...
PASSED: Managed System state is operating.
INFO: Any system dumps present in the current dump logical volumes
will not be available after live update is complete.

```

- The Live Update operation first validates the LPAR requirements and then creates a clone of the root volume group for booting the surrogate. The surrogate LPAR is booted with cloned rootvg and then a mirror of the original rootvg created and is assigned to the surrogate. The workload is migrated to surrogate partition. After the blackout time ends the workload will be running on the surrogate partition. Last, the original LPAR is shut down and deleted. The actual Live Update operation can be done with the geninstall utility. Although Live Update could be initiated from a NIM client or master, in our tests we initiated it from the Original LPAR itself using the geninstall command.

Here is an example of applying a testfix using geninstall:

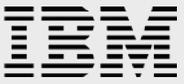
```

# time geninstall -k -d . testfix.160615.epkg.Z 2>&1 | tee
lkutest.log

  Computing the estimated time for the live update operation:
  -----
  LPAR: LKU2.isst.aus.stglabs.ibm.com
  Blackout_time(s): 31
  Global_time(s): 3453

  Checking mirror vg device size:
  -----
  Required device size: 40896 MB
  Given device size: 40959 MB
    PASSED: device size is sufficient.
    ...
  PASSED: Managed System state is operating.
  INFO: Any system dumps present in the current dump logical
  volumes will not be available after live update is complete.
  Validating live update input data.

```



```

Non-interruptable live update operation begins in 10 seconds.

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.

File /etc/inittab has been modified.

One or more of the files listed in /etc/check_config.files have
changed.
    See /var/adm/ras/config.diff for details.

```

The ifix installation can be verified using emgr command:

```

# emgr -l

ID  STATE LABEL          INSTALL TIME          UPDATED BY ABSTRACT
=== =====
=====
=====

```

1	S	lkusock	04/25/18 13:37:57	Patch for socket ENOMEM crash
2	S	tbsync	05/14/18 14:04:36	Timebase synchronization fix (1052788)
3	S	mcrtimeout	08/09/18 01:18:22	Time out for surrogate wait (1064095)
4	S	testifix	04/19/19 05:11:38	LU

capable testifix

- If the Live Update operation fails, it logs information in the `/var/adm/ras/liveupdate/logs` directory. The current log information is recorded in the file `/var/adm/ras/liveupdate/logs/lvupdlog`.
- The LU blackout time can be measured using `alog` command or using “lvupdlog” file:

```
# alog -t mobte -o
time=041919:05:22:53 pid=22236703548768257 type=Global
dep_mfreq=16341000000Hz dep_sfreq=5447000000Hz dep_cont=468
bw=120000000B/s dep_iorbw=153129588B/s dep_iowbw=103367822B/s
dep_pipebw=45512431B/s dep_rvgsz=40896MB dep_lulvsz=30486MB
dep_cpuavail=10000 dep_flags=0x1 arr_mfreq=16341000000Hz
arr_sfreq=5447000000Hz arr_cont=468 arr_iorbw=153129588B/s
arr_iowbw=103367822B/s arr_pipebw=45512431B/s arr_rvgsz=40896MB
arr_lulvsz=30486MB arr_cpuavail=10000 arr_flags=0x1 files=1234
procs=237 threads=427 datasz=135556016B socks=94 mempg=1118560
nmqs=1 nmsgs=0 nqbs=0 nsems=8 shmsz=35513574816 mmapsz=55557
flags=0x0 stdl_blackout=218.000000s stdl_global=2718.000000s
blackout=21.918161s global=6034.241211s
```

## Test description – stress / destructive tests

Oracle provided workloads are used for stressing the system. All tests are performed using a combination of Oracle CPU, IOP and DSS workloads. During the tests CPU utilization is maintained above 90% and memory utilization above 80%. In our testing we have applied testifix using the AIX Live Update feature during the workload run. The test cases cover both destructive and stress tests. The test cases are listed in the following table.



Test Scenario	Test Procedure	Observed Test Outcome and Recovery Actions
Insufficient CPUs for Surrogate	1) Start the workload, run it stable for 30mins 2) Start LU operation 3) Make no. of cpus to minimal via HMC so that surrogate LPAR creation fails 4) Wait for at least 20 minutes or until the workloads have stabilized.	LU operation failed and Oracle software stack remained up and running. The AIX workload went back to the level it was before the start of LU.
Insufficient Memory for Surrogate	1) Start the workload, run it stable for 30mins 2) Start LU operation 3) Reduce Free memory to minimal via HMC so that surrogate LPAR creation fails 4) Wait for at least 20 minutes or until the workloads have stabilized.	
Instance Recovery during LU	1) Start the workload, run it stable for 30mins 2) Start LU operation 3) Kill pmon process or do "shutdown abort" during LU 4) Start database up, instance recovery will occur	Workload run got stopped. Automatic instance recovery took place during instance startup. There were no ORA errors or database corruption.  LU operation completed successfully and ifix got applied
Original LPAR failure during LU	1) Start the workload, run it stable for 30mins 2) Start LU operation 3) Powerdown / shutdown the Original LPAR ungracefully BEFORE/DURING checkpoint phase	Workload run got stopped. Once the Original lpar came up, Oracle Database instance restarted automatically.  LU operation got failed and recovered the system using "clvupdate"
Surrogate LPAR failure during LU	1) Start the workload, run it stable for 20mins 2) Start LU operation 3) Powerdown / shutdown the surrogate lpar ungracefully BEFORE/DURING checkpoint phase 4) Wait for at least 20 minutes or until the workloads have stabilized.	LU operation got failed and workload continued to run on original lpar. There were no ORA errors or database corruption.



Storage Disk failure at Original lpar	<ol style="list-style-type: none"> <li>1) Start the workload, run it stable for 30mins</li> <li>2) Start LU operation</li> <li>3) At SAN switch, Unplug the storage disk(new rootvg disk) before creation of surrogate lpar</li> <li>4) Wait for at least 20 minutes or until the workloads have stabilized.</li> </ol>	
Storage Disk failure at Surrogate	<ol style="list-style-type: none"> <li>1) Start the workload, run it stable for 30mins</li> <li>2) Start LU operation</li> <li>3) Unplug the storage disk(new rootvg disk) which belongs to surrogate</li> <li>4) Wait for at least 20 minutes or until the workloads have stabilized.</li> </ol>	
Network failure, bring down ethernet interface at surrogate	<ol style="list-style-type: none"> <li>1)Start the workload, run it stable for 30mins</li> <li>2) Start LU operation</li> <li>3) Bring down the virtual ethernet interface at surrogate lpar</li> <li>4) Wait for at least 20 minutes or until the workloads have stabilized.</li> </ol>	
Reboot VIOS	<ol style="list-style-type: none"> <li>1) Start the workload, run it stable for 30mins</li> <li>2) Start LU operation</li> <li>3) Reboot VIOS that surrogate uses via the HMC</li> <li>4) Wait for at least 20 minutes or until the workloads have stabilized.</li> </ol>	
PIN SGA	<ol style="list-style-type: none"> <li>1) Startup the database with pinned SGA (using the lock-sga parameter)</li> <li>2) Start the workload, run it stable for 30mins</li> <li>3) Start LU operation</li> <li>4) Monitor the blackout time</li> </ol>	
Mixed Workload Stress run	<ol style="list-style-type: none"> <li>1. Let the system stabilize by running the workloads simultaneously for more than 20 minutes</li> <li>2. Perform the LU operation</li> </ol>	



	3. Wait for at least 20 minutes or until the workloads have stabilized on Surrogate lpar	
Multiple LUs run sequentially with Stress workload	<ol style="list-style-type: none"><li>1. Let the system stabilize by running the workloads simultaneously for more than 20 minutes</li><li>2. Execute a script that runs multiple LUs sequentially</li><li>3. Wait for at least 20 minutes or until the workloads have stabilized on Surrogate LPAR</li></ol>	
Database functional test	<ol style="list-style-type: none"><li>1. Let the system stabilize by running the workloads simultaneously for more than 20 minutes</li><li>2. Perform any one operation below<ul style="list-style-type: none"><li>· Add flashback database</li><li>· Take RMAN hot backup including control file</li><li>· Add a very large datafile to a tablespace</li></ul></li><li>3. Perform the LU operation</li><li>4. Wait for at least 30 minutes or until the workloads have stabilized on surrogate</li></ol>	

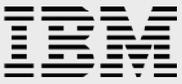
## Performance measurements

For testing purposes three Oracle generated workloads, CPU, IOP and DSS, were used. These workloads were designed to stress the system by generating database workloads with high CPU, I/O random reads and Decision Support work with high sequential reads. No optimization work was done at any level – not on the system level or the I/O subsystem level. For configuring the Oracle Single instances refer to “Oracle Database 12c Tips and Considerations” white paper.

### Database and Workload Considerations

A Container Database (CDB) was created, consisting of three PDBs (CPU, IOP, DSS). CPU PDB was built with 25 warehouses and is 5GB in size. IOP PDB was built with 50 warehouses and it is 10G. DSS was built with 250 warehouses and is 50GB. Oracle Database instance Shared Global Area (SGA) was configured as 32GB. Oracle Database workload datafiles were placed on both filesystem and Oracle ASM. The test runs were performed with 14 CPU users, 24 IOP users and 3 DSS users.

A high number of DSS workload users will introduce IO contention on the storage system and this will impact the “Creating rootvg for boot of surrogate” phase and the “Creating mirror of original LPAR’s rootvg” phases of the Live Update. These two phases will take more time than expected since they are IO



intensive. Customers that are planning to run such configurations are advised to test thoroughly before implementing it in production.

The Oracle provided workloads are tuned to stress the LPAR and during the run CPU utilization was maintained around 95-100%. The graph below shows the CPU and IOPs during the Live Update operation. The LU operation was started at 03:44 using geninstall. The blackout time started at 05:11:06 and lasted for 22s. The system data is collected for every two minutes. Since the blackout time is minimal we see slight drop in CPU utilization at 05:11 in below chart.

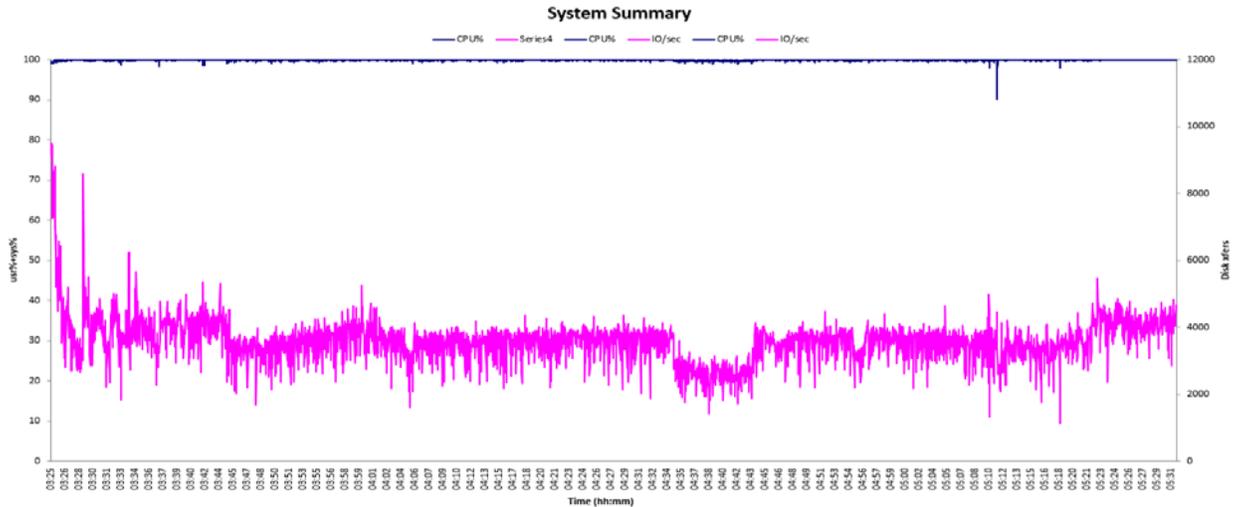


Figure 2. System CPU and IO utilization during LU

The IOP workload does many db\_file\_sequential reads and since we are using a high number of IOP workload users db\_file\_sequential read is seen as the top wait event in the AWR report, as shown in Figure 3 below.

### Top 10 Foreground Events by Total Wait Time

Event	Waits	Total Wait Time (sec)	Avg Wait	% DB time	Wait Class
db file sequential read	6,782,836	77.8K	11.47ms	27.1	User I/O
direct path read	1,218,438	42.8K	35.10ms	14.9	User I/O
DB CPU		29.5K		10.2	
enq: TX - row lock contention	474,225	11.5K	24.34ms	4.0	Application
write complete waits	101,419	7638.8	75.32ms	2.7	Configuration
latch: In memory undo latch	2,666,245	3491.8	1.31ms	1.2	Concurrency
log file sync	58,270	1778	30.51ms	.6	Commit
db file scattered read	101,590	1344.3	13.23ms	.5	User I/O
buffer busy waits	813,711	975	1.20ms	.3	Concurrency
log file switch (private strand flush incomplete)	2,598	302.2	116.33ms	.1	Configuration

Figure 3. AWR Top 10 wait events for workload run

The LU operation will introduce some blackout time and the application workload should be able to accommodate this blackout time. The LU blackout time is based on “checkpoint and restart” technology.



At the beginning of blackout time the workload is paused and system current state is saved. Once “checkpointing” completes all processes will get restarted on the surrogate LPAR. This is similar to Workload Partition Live Application Mobility(WPAR) which was introduced with AIX 6.1. With no workload on the system we had noticed the blackout time as ~15 seconds. In our testing we have noticed an average blackout time as 23 seconds. Figure 4 below shows the throughput on the system when we ran all three OAST workloads and throughput is measured in-terms of transaction-per-second(TPS).

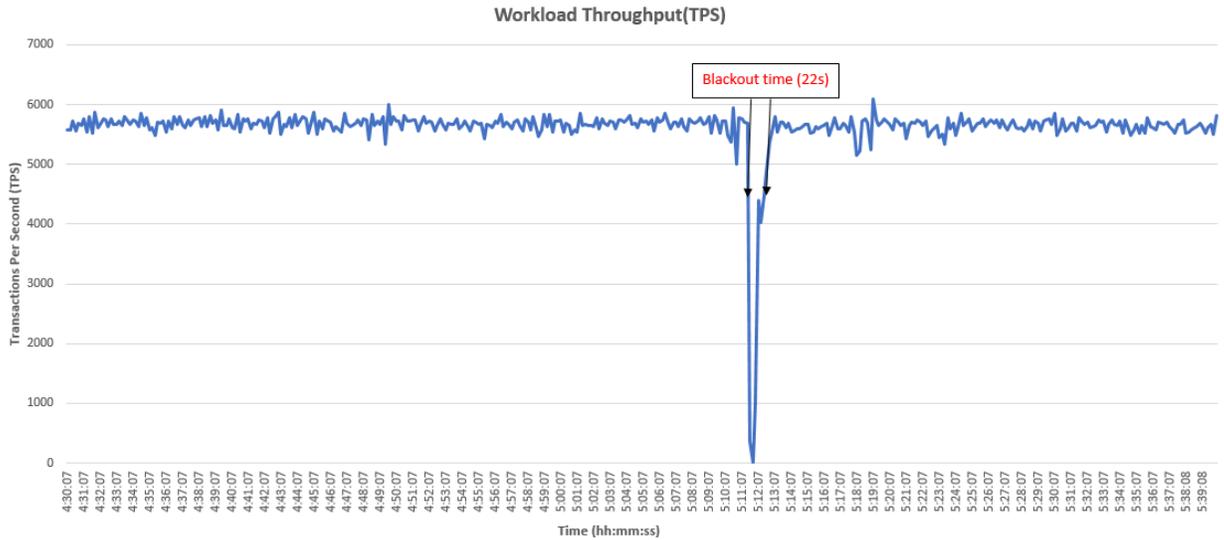


Figure 4. Workload Throughput measured in TPS

During the blackout time a VKTM warning message is seen in the database alert log and the forward time drift is equal to the blackout time of LU operation.

```
alert_iopdb.log
=====
2019-04-19T05:11:34.956742-05:00
Warning: VKTM detected a forward time drift.
Time drifts can result in unexpected behavior such as time-outs.
Please see the VKTM trace file for more details:
      /u01/diag/rdbms/iopdb/iopdb/trace/iopdb_vktml_20644226.trc
```

## Conclusion

In this document we have tested AIX 7.2 Live Update feature with Oracle Database by performing various stress and destructive tests, which were done as part of certification testing. We also listed the best practices that need to be followed when doing the LU operation.



## Resources

---

These Web sites provide useful references to supplement the information contained in this document:

- IBM Knowledgecenter – AIX Live Update  
[https://www.ibm.com/support/knowledgecenter/en/ssw\\_aix\\_72/com.ibm.aix.install/live\\_update\\_in\\_stall.htm](https://www.ibm.com/support/knowledgecenter/en/ssw_aix_72/com.ibm.aix.install/live_update_in_stall.htm)
- IBM Knowledgecenter – geinstall  
[https://www.ibm.com/support/knowledgecenter/en/ssw\\_aix\\_72/com.ibm.aix.cmds2/geninstall.htm](https://www.ibm.com/support/knowledgecenter/en/ssw_aix_72/com.ibm.aix.cmds2/geninstall.htm)
- AIX Live Update Concepts  
<https://www.ibm.com/developerworks/community/blogs/cgaix/resource/AIXLiveUpdateblog.pdf?lang=en>
- Oracle Database 12cR2 on AIX  
<https://docs.oracle.com/en/database/oracle/oracle-database/12.2/axdbi/index.html>
- Oracle DB & RAC 12c on IBM AIX : Tips and Considerations  
<https://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102425>
- IBM eServer pSeries [System p] Information Center  
<http://publib.boulder.ibm.com/infocenter/pseries/index.jsp>
- IBM Publications Center  
<http://www.elink.ibmink.ibm.com/public/applications/publications/cgibin/pbi.cgi?CTY=US>
- IBM Redbooks  
<http://www.redbooks.ibm.com/>



## About the authors

---

**Bhargavaram Akula** is a Technical Consultant with IBM India, Hyderabad. He collaborates with the specialists at the IBM Oracle International Competency Center based in Foster City, and Redwood Shores, California, US, working on Oracle product certifications on IBM Power Systems™ running AIX. He has extensive experience with Oracle Products.

**Sougata Sarkar** is Senior Staff Software Engineer working for the Integrated Software Systems Test Team for 10 years. He is working as the ISST AIX/VIOS lead . His areas of expertise are Shared Storage Pools, Virtual IO Server, PowerVM® virtualization concepts like NPIV, vSCSI, Remote Restart, Logical Partition Mobility, Oracle Database and DB2®. He has good knowledge on AIX features like Live Kernel Update. He can be reached at [sougsark@in.ibm.com](mailto:sougsark@in.ibm.com)



## Trademarks and special notices

---

© Copyright. IBM Corporation 2019. All rights reserved.

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

IBM, the IBM logo and AIX, DB2, Power, Power Systems, and PowerVM are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

The information provided in this document is distributed "AS IS" without any warranty, either express or implied.

The information in this document may include technical inaccuracies or typographical errors.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Contact your local IBM office or IBM authorized reseller for the full text of the specific Statement of Direction.

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in IBM product announcements. The information is presented here to communicate IBM's current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Photographs shown are of engineering prototypes. Changes may be incorporated in production models.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.