ISICC | IBM SAP International Competence Center

Power Systems

Planning Guide

SAP Applications with AIX Live Update

IBM Systems
Created on August 31, 2020 – Version 1.0
Last modified on August 31, 2020 – Version 1.0
© Copyright IBM Corporation, 2020

IBM

Note: Before using this information and the product it supports, be sure to read the general information under ”Copyrights and Trademarks” on page 23 as well as “Disclaimer and Special Notices” on page 24.

Edition Notice and Version Information

© Copyright IBM Corporation 2020. All Rights Reserved.
US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp. All trademarks or registered trademarks mentioned herein are the property of their respective holders.
IBM Corporation
Systems and Technology Group
Route 100
Somers, NY 10589
Edition Notice: this is version 1.0 of this document.
Focus: AIX life update for SAP Production systems
Target:

- AIX Release 7.2 technology level 3, and higher
- SAP products based on ABAP stack, e.g. NetWeaver 7.5

Doc Version	Changes
1.0	Initial Publication

Preface

Running SAP on IBM Power Systems offers customers a consistent platform for their applications, best-in-class performance, resilience for critical workloads, and most flexible infrastructure. This holds for existing IT assets, as servers and storage, as well as skills and operation procedures.

One feature offered by the IBM AIX platform from AIX release 7.2 on is the Live Update ability, i.e. to update to a higher AIX technology level or service pack without need of reboot.

In this project the AIX feature of Live Update was validated successfully on SAP systems.

The scope of the validation included small to medium sized SAP systems with ABAP stack. Any issues found during testing were investigated and resolved. They are documented in this paper together with a set of best practices collected during the project.

About This Document

This document is intended for architects and specialists planning to use AIX Live Update for their SAP databases and application servers to minimize planned downtime.

The most recent document version can be downloaded from IBM TechDocs:
http://www.ibm.com/support/techdocs/atsmastr.nsf/Web/WhitePapers

Feel free to provide feedback and change requests for this document via email to enable.SAP@de.ibm.com.

Authors:

Ralf Kokowski is working with IBM Germany, Research and Development Lab, as a member of the SAP on IBM POWER (AIX, Linux) development team, and as team lead for the SAP on AIX and Linux on POWER support team.

Bernhard Mollner-Reinbold is working with IBM Germany, Research and Development Lab, as a member of the SAP on IBM POWER (AIX, Linux) development team.

Paul H Hernandez is an experienced IBM AIX system developer - a member of the Live Update development team.

David Sheffield is an IBM STSM in Austin Texas working as part of the AIX architecture team. David has over 30 years experience in AIX development, and AIX Live Update has been his focus in recent years.

Table of Contents

Figures

Tables

IBM Systems, © Copyright IBM Corporation, 2020

Introduction

The AIX Live Update feature allows to upgrade an AIX service pack or even technology level without need for reboot, so that any applications can continue to run. Users experience no disruptions, besides a minor freeze time of usually less than a minute.

For validation a multitude of scenarios was tested on two SAP installations, each comprised of an SAP *central instance* host (LPAR, logical partition) plus another host holding an *additional dialog instance*. A central instance setup means to have an SAP dialog instance plus all other required SAP components, namely the central services component and the global file tree. Also, the database was installed on that host. AIX Live Update tests were run on the central instance host as well as on the additional instance host.

Different load and other conditions were set up during the test runs, and any issues found were documented and evaluated, for providing guidance of best practices. The validation project was successful, finding all AIX Live Update executions succeeding, when few conditions, documented here, are ensured.

The overall project team was comprised of two IBM teams, one located in Germany in the IBM department SAP Technical Enablement, the other in US, Austin development labs. The Germany based team was responsible for setting up the SAP installations and pursuing the tests. The US team worked in the roles of AIX kernel development expert and AIX Live Update architect.

IBM Systems, © Copyright IBM Corporation, 2020

SAP Systems

The two SAP systems used for validation were both NetWeaver installations with ABAP stack, one with Oracle database, the other with Db2 database. Both SAP systems were distributed ones, comprised of two hosts (LPARs[1]) each. One host was set up as central instance with database, central services and primary application server (PAS), and global file tree ("/sapmnt"), and the other host had only an additional application server (AAS) installed.

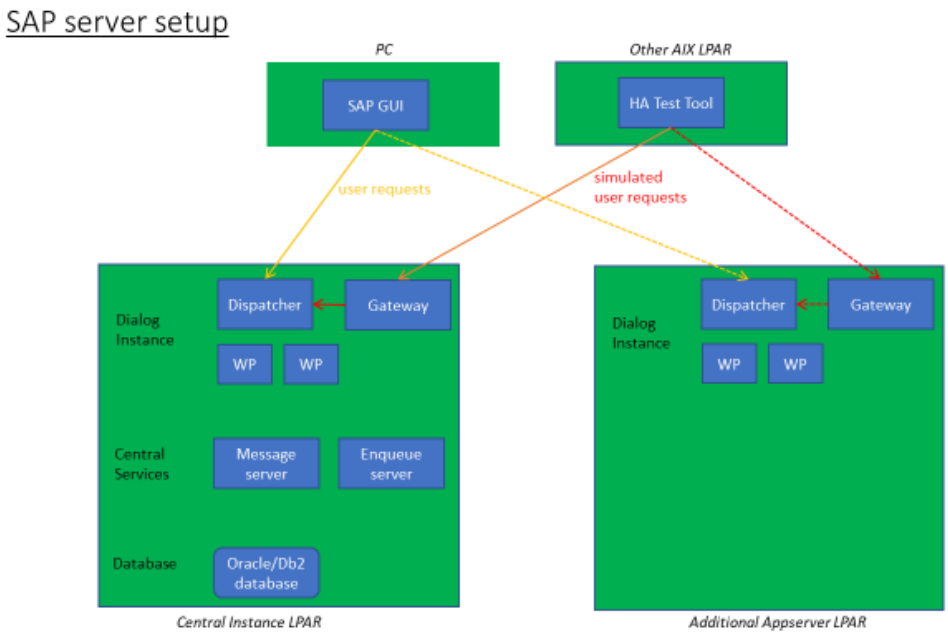

Figure 1: SAP server setup

The IBM server used to hold the LPARs for both SAP systems was a IBM Power 730 Express server (8231-E2C), with 128 GB RAM and 16 cores.

Each central instance host was configured as a shared pool LPAR with 18 GB RAM, 6 virtual CPUs, SMT-4, and an entitled capacity of 1. The additional application server hosts were configured similarly, only with reduced amount of RAM of 16 GB.

The AIX version on all LPARs was release 7.2 technology level 3 service pack 1 at the start of all Live Update tests. The target version chosen during Live Update was AIX 7.2 technology level 3 service pack 2. For easier handling some AIX Live Update testing was pursued without actual version update, focussing on the Live Update procedures.
The SAP version used on all SAP installations was NetWeaver 7.5 with SAP kernel version 7.53

[1] Logical partitions

Test scenarios

Conceptual approach

To achieve a representative test coverage, a wide variety of SAP scenarios on AIX is to be tested, while running some well specified sets of workload.

The variety is achieved conceptually by defining a set of relevant attributes, also called "dimensions", which can be varied independently, and are each represented by a small set of attribute values.

While some dimensions and representation points may be regarded as kind of natural and mandatory, as setting some sizing attribute to low/medium/high, other dimension's relevance is not as sharply defined. Selecting suitable dimensions, and related aspects, means a challenge, to create a conceptual model that well represents reality, here typical customer scenarios.

Building on such model, a full "test matrix" is then defined as the "product" of all dimensions, each represented by few value-points per dimension. Ideally all elements in such matrix, each representing a scenario, would be tested. From a practical perspective this will not be possible for an ambitious model with many dimensions, due to the effect of "combinatoric explosion".

Primary dimensions

Attribute "A" for the "architecture":
The variants are:
A1: Run the Live Update on the host with central instance setup, including a database installation.
A2: Run the Live Update on the host with the additional dialog instance.
Two attribute-values: A1=db_plus_appserver, A2=appserver_alone

Attribute "W" for the "workload":
Workload was generated by two means. The simplest method was to use the *SAP Load Generator* (“SGEN”), which is already available in the SAP GUI, and is capable to create high workload in the order of minutes to half an hour. SGEN is basically doing database reads, then some processing in the ABAP engine and then database updates and inserts.

The other method was to use the *SAP High Availability Test Tool* (Release 2.14), or in short HA Test Tool, to create SAP workload. It offers various so-called test classes which are representing different SAP workloads. The name element "High Availabilty" refers to some abilities to support high availability testing, but it can as well be used just for workload generation.

The variants used for attribute “W" are:
W1: Small workload: Run from SAP GUI the transaction code SGEN.

W2: HA Test Tool generated workload, that simulates SAP user dialogs leading to high load on the database server and the application server: "hatool_DB_load". It is parameterized by the number of users/clients and others.
W3: HA Test Tool generated workload, that can produce high CPU utilization on the central instance server: "hatool_ENQ_CPU_load". It is mainly parameterized by the number of clients.
W4: HA Test Tool generated workload, that can produce high memory utilization on the central instance server: "hatool_ENQ_MEM_load". It is mainly parameterized by the number of clients. Despite its reference to the SAP enqueue server, the major memory consumption happens in the SAP work processes.
Four attribute-values: W1=sgen_load, W2=hatool_DB_load, W3=hatool_ENQ_CPU_load, W4=hatool_ENQ_MEM_load

Attribute "B" for the "blackout time extension":
The AIX Live Update procedure does implicitly involve a so called "blackout time", also called “blackout period”, where all memory-activities are stopped, until the surrogate LPAR is ready to run. For investigation of the effects of longer lasting blackout times an internal method was used to prolong it.

The variants are:
B1: Run the Live Update without modifying the blackout time.
B2: Extend the Live Update blackout time to a value of "N" minutes that impacts the behaviour, likely by hitting a timeout given by the environment. Possibly the first impact may be experienced when 1 minute is exceeded.
Two attribute-values: B1=blackout_unchanged, B2=blackout_N_minutes

Secondary dimensions

Attribute „D“ for the „database“:
The variants are:
D1: Using the Oracle-based SAP installation
D2: Using the Db2-based SAP installation

Test matrices

The “core tests”, with most dense test matrix, were pursued on the SAP installation with Oracle database. Reason for selecting this installation as the most essential, was the interest and cooperation of one customer in this project, who had available mainly Oracle based SAP installations.

Test period 1: Core testing (Oracle database)
The test matrix is shown below.
It is comprised by the primary dimensions "A * W * B":
Two architectural variants were specified as A1 (db_plus_appserver, "2-tier") and A2 (appserver_alone, "3-tier"). Four Workload variants were used, one generated from SAP GUI “SGEN” transaction, and the other three generated from HA Test Tool. For the blackout period two variants were selected, one with unchanged blackout time, usually less than a minute, and secondly with blackout time extended by one minute. This altogether results in 2 * 4 * 2 = 16 test cases.

Secondary dimensions:
The following secondary dimensions are of fixed-values here: Database: D=oracle

Test ID	A/architecture	W/workload	B/blackout	D=oracle
A1W1B1core	db_plus_appserver	sgen_load	blackout_unchanged	x
A1W2B1core	db_plus_appserver	hatool_DB_load	blackout_unchanged	x
A1W3B1core	db_plus_appserver	hatool_ENQ_CPU_load	blackout_unchanged	x
A1W4B1core	db_plus_appserver	hatool_ENQ_MEM_load	blackout_unchanged	x
A2W1B1core	appserver_alone	sgen_load	blackout_unchanged	x
A2W2B1core	appserver_alone	hatool_DB_load	blackout_unchanged	x
A2W3B1core	appserver_alone	hatool_ENQ_CPU_load	blackout_unchanged	x
A2W4B1core	appserver_alone	hatool_ENQ_MEM_load	blackout_unchanged	x
A1W1B2core	db_plus_appserver	sgen_load	blackout_N_minutes	x
A1W2B2core	db_plus_appserver	hatool_DB_load	blackout_N_minutes	x
A1W3B2core	db_plus_appserver	hatool_ENQ_CPU_load	blackout_N_minutes	x
A1W4B2core	db_plus_appserver	hatool_ENQ_MEM_load	blackout_N_minutes	x
A2W1B2core	appserver_alone	sgen_load	blackout_N_minutes	x
A2W2B2core	appserver_alone	hatool_DB_load	blackout_N_minutes	x
A2W3B2core	appserver_alone	hatool_ENQ_CPU_load	blackout_N_minutes	x
A2W4B2core	appserver_alone	hatool_ENQ_MEM_load	blackout_N_minutes	x

Table 1: Test matrix for the SAP Oracle installation

Test period 2: Db2 testing
The secondary dimension is changed to: database D=db2.
A sub-selection of the full matrix "A * W * B" was specified as test matrix.

Test ID	A/architecture	W/workload	B/blackout	D=db2
A1**W1**B1db2	db_plus_appserver	sgen_load	blackout_unchanged	x
A2**W2**B1db2	appserver_alone	hatool_DB_load	blackout_unchanged	x
A1**W3**B1db2	db_plus_appserver	hatool_ENQ_CPU_load	blackout_unchanged	x
A2**W4**B1db2	appserver_alone	hatool_ENQ_MEM_load	blackout_unchanged	x
A1W1**B2**db2	db_plus_appserver	sgen_load	blackout_N_minutes	x

Table 2: Test set for the SAP Db2 installation

Test variants
Note that most test cases were executed not only once but multiple times, with varied parameter changes, especially if issues were found and required to be drilled down. The variations could be about increasing trace levels, or setting parameters on HA Test Tool level, SAP level or OS level for better understanding.

Preparations for AIX Live Update

Before running an AIX Live Update some preparations are required, e.g. some resources must be provided and the *rootvg* volume group should be backed up. This is outlined in the following.

Resource requirements

Some resource demands must be fulfilled to be able to run a Live Update on an AIX LPAR, respectively on the IBM POWER server holding it. The IBM knowledge center documentation is found at
https://www.ibm.com/support/knowledgecenter/en/ssw_aix_72/install/lvupdate_configuration.html

Double RAM resources, temporary
The IBM physical server must have available free memory resources of the same amount that the Original LPAR holds. This is required since a Surrogate LPAR as kind of clone is created, while the Original LPAR runs, so that double RAM resources are needed. At the end of AIX Update the Original LPAR will be decommissioned, and its RAM resources given back.

Double CPUs resources, or less, temporary
From first approach the IBM physical server must have available free CPU resources of the same amount that the "Original LPAR" holds, to provide to the "Surrogate LPAR". At the end of AIX Update the Original LPAR will be decommissioned, and its CPU resources given back.
But this demand might be reduced, when specifying an option *"cpu_reduction = yes"* in the configuration file *'lvupdate.data'* (see below). Note that the processes on the LPAR must be able to cope with the reduced amount of CPUs for some time period, until Live Update finished.

Additional Physical Volumes
At least two additional physical volumes with the size of the rootvg are required.

If it's intended to avoid future reboots, a third physical volume is required as preparation for the next Live Update[2].

[2] See section Multiple Live Updates without a reboot

Example disk and volume group setup:

```
root@XXXXXX/# lspv
hdisk0          00f86d3f3ea81468                rootvg          active
hdisk1          00f86d3fd4371009                None
hdisk2          00f86d3fd4371080                None
hdisk3          00f86d3fd43710fc                None
hdisk4          00f86d3fd4371182                sap_vg          active
root@XXXXXX/# getconf DISK_SIZE /dev/hdisk0
40960
root@XXXXXX/# getconf DISK_SIZE /dev/hdisk1
40960
root@XXXXXX/# getconf DISK_SIZE /dev/hdisk2
40960
root@XXXXXX/# getconf DISK_SIZE /dev/hdisk3
40960
root@XXXXXX/# getconf DISK_SIZE /dev/hdisk4
133120
```

Here the rootvg has size of 40 GB, provided by hdisk0. The SAP application is contained in the volume group sap_vg with size of 130 GB, provided by hdisk4. The additional disks hdisk1, hdisk2, hdisk3 are used for AIX Live Update.

General preparations

Besides the resource demands to fulfill also some generic preconditions should be checked and precautions considered.

- If an AIX Live Update is to be run the first time, the IBM knowledge center documentation should be studied well. https://www.ibm.com/support/knowledgecenter/en/ssw_aix_72/install/live_update_install.html

- Especially the section „Planning“ should be read, with its advices for „LPAR requirements for Live Update“, „Live Update restrictions“, and „Best practices for the Live Update function“.

- In the SAP context also the best practices of this document should be studied.

- Backup the rootvg of the AIX LPAR and consider to back up the LPAR profiles defined on the hardware management console (HMC), that manages the LPAR to be updated.

- Get user ID and password to access the HMC.

- Before doing the live update, do a prerequisite check provided by geninstall tool:

```
hmcauth -a <hmchost> -u <hmcuser> -p <hmcpassword>
geninstall -k -p
```

Configuration file lvupdate.data

A configuration file */var/adm/ras/liveupdate/lvupdate.data* needs to be filled with the specification of the planned AIX Live Update run. It can be created as copy of */var/adm/ras/liveupdate/lvupdate.template* and then modified appropriately. The template file contains detailed parameter descriptions. The following web link shows the newest update of the parameters and descriptions:
https://www.ibm.com/support/knowledgecenter/ssw_aix_72/install/lvupdate_configuration.html

As very minimum, two hard disks are to be specified which will contain clones of the rootvg, and the IP address of a Hardware Management Console, and a user to access it.

```
disks:
        nhdisk  = hdisk1
        mhdisk  = hdisk2
        tohdisk =
        tshdisk =

hmc:
        lpar_id =
        management_console = XX.XX.XX.XXX
        user = hscroot
```

Multiple Live Updates without a reboot

Attempting to execute a second Live Update with the same specification for *nhdisk* will result in a failure because *nhdisk* is still in use. This is used as the initial boot disk for the surrogate LPAR during Live Update and will be unavailable for other use until the next successful Live Update or a reboot. It will be part of a special volume group called *lvup_rootvg*. In order to execute another Live Update, another disk or disks must be specified for *nhdisk* in the *lvupdate.data* file. After a second Live Update, the original *nhdisk* will no longer be in use and is therefore available for another Live Update.

In order to avoid having to switch the value of *nhdisk* in the *lvupdate.data* file for each Live Update, it's possible to specify both *nhdisk* and *alt_nhdisk*. If *nhdisk* is not available, Live Update will use *alt_nhdisk*. In this way, by specifying *mhdisk*, *nhdisk*, and *alt_nhdisk*, any number of Live Updates can be performed without needing to update the *lvupdate.data* file each time.

The extended configuration for multiple live updates is recommended here. It could look as follows.

```
disks:
        nhdisk     = hdisk1
        mhdisk     = hdisk2
        alt_nhdisk = hdisk3
        tohdisk =
        tshdisk =

hmc:
        lpar_id =
        management_console = XX.XX.XX.XXX
        user = hscroot
```

Example AIX Live Update run

The AIX Live Update run described on the following pages was selected as a representative example. In this scenario the SAP High Availability Test Tool was used to create an SAP workload. After starting the program, three phases can be distinguished:

1. Initialization
2. High load phase
3. Post-processing i.e. creating of output data like log files and summaries

The working method of the tool can be roughly described as follows. The tool simulates a configurable number of clients through threads during the high load phase. Each client thread processes the following sequence of calls in a loop:

1. Login
2. Execute steps 1 – n
3. Logoff

SAP's HA Test Tool offers various so-called test classes which are representing different SAP workloads. The test class CL_HA_DB_LOAD was chosen for the sample AIX Live Update run. The selected test class puts load on the database and is customized as shown below:

- *clientnum = 50: number of simulated client users*
- *stepnum = 5: number of steps per loop i.e. insert + update + delete + rollback + check*
- *duration = 1800: duration in seconds of high load phase of the test run*

Two AIX LPARs were used to run the test. The first AIX LPAR hosted an Oracle database and an SAP application server. The AIX Live Update run has taken place on this LPAR and covered a migration from AIX operating system level 7200-03-01-1838 to 7200-03-02-1846. At the same time the SAP High Availability Test Tool was started on the other AIX LPAR simulating clients which log on to the SAP application server on the first AIX LPAR. This setting has the Test ID A1W2B1core as described in chapter 3.4.
Below are the steps for performing the AIX Live Update run:

1. At the beginning system level is 7200-03-01-1838. The `root` user authenticated to the HMC using the `hmcauth` command. This is a precondition to run AIX Live Update.

```
(0)root @ xxxx01: /
# oslevel -s
7200-03-01-1838

(0) root @ xxxx01: /
# hmcauth -a x.x.x.x -u hscroot -p xxxxxxxx
```

2. The geninstall command initiated the AIX Live Update run as well as the migration to service pack 7200-03-02-1846 which was located in /mnt update_all (the date command was used as an easy way to get a timestamp at the start and the end of the test).

```
(0)root @ xxxx01: /
# date; geninstall -k -d /mnt update_all; date
Sun Aug 16 04:21:34 CDT 2020

Validating live update input data.

Computing the estimated time for the live update operation:
-----------------------------------------------------
LPAR: xxxx01.x.x.x.x
Blackout time(in seconds): 122
Total operation time(in seconds): 1074

Checking mirror vg device size:
----------------------------------------
Required device size: 24320 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.

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.
…

        Blackout Time started.

        Blackout Time end.

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

Sun Aug 16 04:53:45 CDT 2020
(0) root @ xxxx01: /
```

3. After the AIX Live Update run had finished the AIX operating system level was 7200-03-02-1846. The actual duration of the blackout time shown by the alog command was 41.079693s compared to the estimated duration of 122 seconds.

```
(0) root @ xxxx01: /
# oslevel -s
7200-03-02-1846

(0) root @ xxxx01: /# alog -t mobte -o | tail -1
time=081620:04:53:32 pid=23363239110770689 type=Global
dep_mfreq=6039527429Hz dep_sfreq=6039527429Hz dep_cont=101
bw=220000000B/s dep_iorbw=280285910B/s dep_iowbw=196018307B/s
dep_pipebw=167773440B/s dep_rvgsz=24320MB dep_lulvsz=9980MB
dep_cpuavail=10000 dep_flags=0x0 arr_mfreq=6039527429Hz
arr_sfreq=6039527429Hz arr_cont=101 arr_iorbw=280285910B/s
arr_iowbw=196018307B/s arr_pipebw=167773440B/s arr_rvgsz=24320MB
arr_lulvsz=9980MB arr_cpuavail=10000 arr_flags=0x0 files=955 procs=163
threads=363 datasz=503543184B socks=75 mempg=655305 nmqs=1 nmsgs=0
nqbs=0 nsems=193 shmsz=248365815798 mmapsz=30568 flags=0x0
stdl_blackout=78.000000s stdl_global=1392.000000s blackout=41.079693s
global=1215.240479s
```

The measured runtimes of the test in detail:

- AIX Live Update started at Aug 16 04:21:34 CDT 2020
- AIX Live Update finished at Aug 16 04:53:45 CDT 2020
- SAP High Availability Test Tool started at Aug 16 04:21:45 CDT 2020
- SAP High Availability Test Tool finished at Aug 16 04:53:42 CDT 2020

SAP's High Availability Test Tool provides a graphical view of the test result (loadtest.html). In the given scenario it looked like shown below.
The selected example showed a typical behavior seen during the tests: SAP's High Availability Test Tool was facing long response times during the blackout period, but it completed normally. An abnormal termination of the test tool run was not detected.

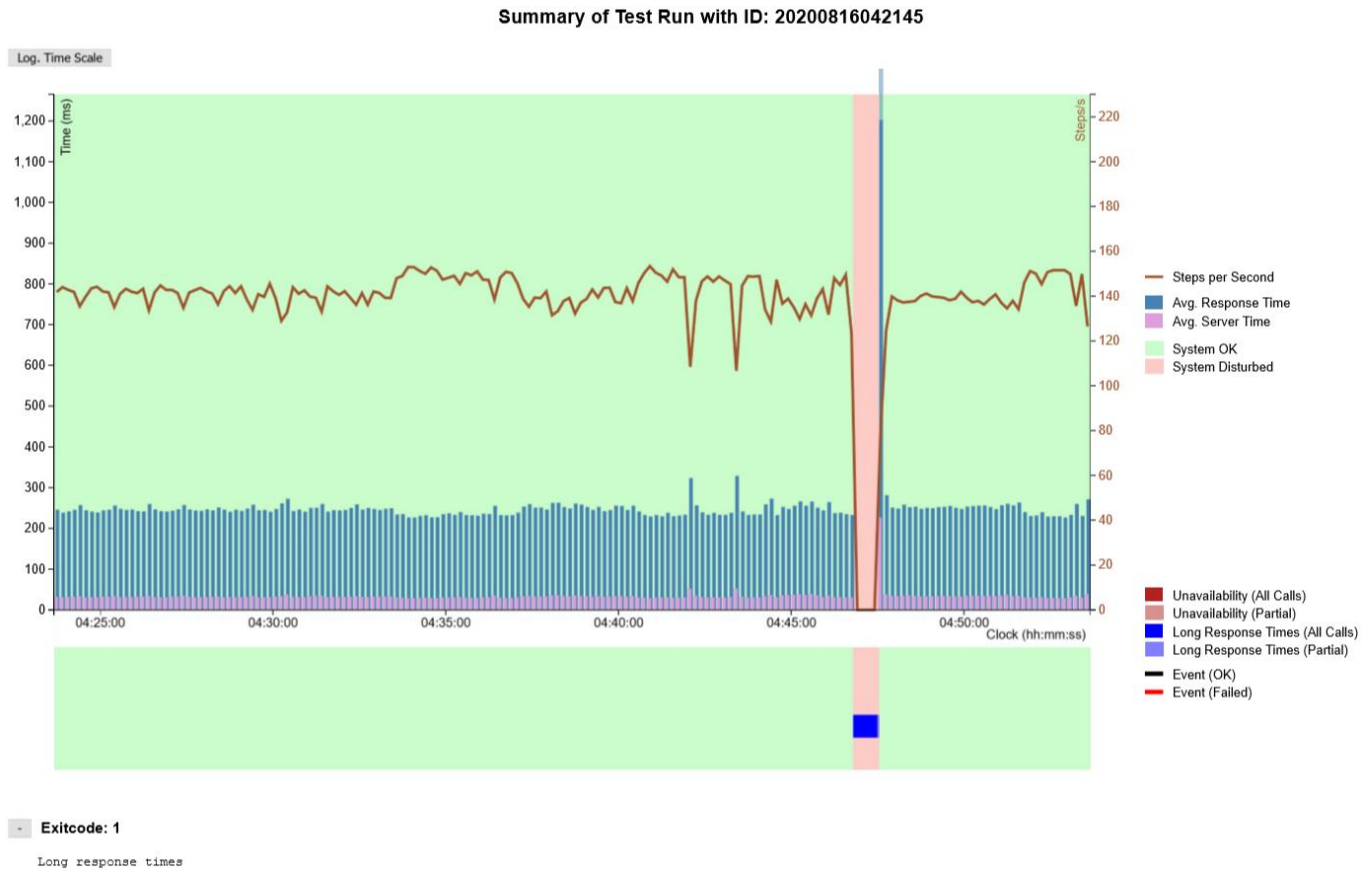

Figure 2: Summary of HA Test Tool run

The summary has the following structure:
Header: Testrunid of test run

- Green background: test run was ok
- Red background: test run finished with errors

Chart: Depicts the high load phase of the test run

- Light green background: system ok
- Light red background: system disturbed (periods of unavailability or long response times occurred)
- Upper part
  - Steps per second
  - Average response and server time
  - Logarithmic time scale offered if response time exceeds the linear time scale
- Lower part
  - Displayed if the system was disturbed for some time or if events were executed
  - Periods of unavailability & Periods of long response times
  - Events
- Additional information
- Exit code

IBM Systems, © Copyright IBM Corporation, 2020

Issues found and best practices

The issues found during test cases are documented in this paper. All were evaluated until understood, and best practices concluded, so that they could be avoided. Few issues were already known at project start, from an earlier evaluation, and also documented here.

Tape devices

Currently tape devices are not supported for AIX Live Update.
The precondition check via 'geninstall -k -p' shows following message if tape devices exist on the current LPAR: "1430-011 FAILED: local tape devices are not supported."
An Request For Enhancement had been issued, "SAN Tape Support in AIX LKU Live Kernel update", with ID 133495, and accepted by IBM development as an Uncommited Candidate for a future release.

Lightweight memory trace

In some scenarios the Live Update prerequisite check has issued following:
"1430-036 FAILED: Not enough available space for /var to dump Light weight memory Trace buffers."

Besides increasing the /var file system and providing more space, it can also help to change an RAS parameter as follows: The 'raso' tunable 'mtrc_commonbufsize' should be set low, e.g. 30.
Documentation for that parameter:
The command *raso -h mtrc_commonbufsize* gives some technical information:
Help for tunable *mtrc_commonbufsize*:
Purpose:
Specifies the memory trace buffer size for common events of Lightweight Memory Trace (LMT) which provides system trace information for First Failure Data Capture (FFDC).
Values:

Default: 1458
Range: 1 - 12211
Type: Dynamic
Unit: 4KB pages

Tuning:
The default value is based on data generation under a reference system-wide activity and HW and system characteristics. The range upper limit is based on HW and system characteristics and depends on the current value of mtrc_rarebufsize because they share the LMT resource. Recorded events are saved in system dump and/or reported through user commands.

IPSEC

The feature *IPSEC* must be deactivated during AIX Live Update.
The Live Update precondition check would show following:
"1430-202 FAILED: Live update cannot run when there is active ipsec configuration."

In SAP scenarios IPSEC may be used for SAP Clones. In that use case the IPSEC feature is only required on the SAP Clone system, as long as it is active.

Live Update memory demands

The Live Update procedures have some notable memory demands on their own. Therefore in scenarios already short on memory, the Live Update demands can lead to an out of memory situation (OOM), and start paging, and potentially run out of paging space. Investigations by AIX kernel development lead to an estimated consumption of 2.5 MB per thread of the system and following rule of thumb.

Important recommendation:
To avoid the risk of exhausting memory during Live Update, the applications and AIX must not consume more than 80% of the LPARs physical memory (RAM), so that 20% are available for Live Update operations.
Note: The available memory is comprised of not only the free memory but also the majority of the file cache. If any application or here the Live Update routines demand more memory than is available in the free list, then the file cache will be shrinked down to a very minimum, usually to 3% of RAM. The command 'svmon -G' shows the available memory explicitly.

No NFS based executables

A constraint for Live Update applicability is that all processes running on the original LPAR must be started from *local executable binaries*, and never from binaries located on NFS mounts. If any running executables were loaded from NFS the Live Update will fail.

For the SAP context the `/sapmnt` directory tree is the critical element here. In typical SAP installations this directory is NFS mounted onto all host with SAP instances of an SAP system, besides on the host where the central instance runs.
Per SAP recommendations the SAP executables should be copied from the /sapmnt tree to a local directory and started from there, but this may not always be ensured.

Recommendation:
In the SAP context, it should be checked before Live Update, whether any processes are loaded from NFS-based executables, especially for the executables located on the /sapmnt file tree.

Volume Group related failures

The Live update may not succeed, showing volume group related error messages. This would not lead to a crash situation, just the upgrade attempt not succeeding.
The error messages could look like following:

```
0516-024 lqueryvg: Unable to open physical volume.
        Either PV was not configured or could not be opened. Run
        diagnostics.

0516-1140 importvg: Unable to read the volume group descriptor area
        on specified physical volume.
```

This is a rare timing issue when Live Update performs disk health checks, and is unrelated to SAP. If the issue is hit an intermediate fix ("iFix") for APAR IJ26902 can be requested at IBM support.

SAP Syslog messages

The SAP System Log, with transaction code *SM21*, may show some messages, flagged as errors, warnings or information, induced by the AIX Live Update.
Shown below are the kind of messages that occurred during the validated runs. They were assessed by SAP experts and assessed as uncritical.
Some messages, which show the flag *RD* for "gateway reader process", are related to the SAP instances gateway process. This process is meant to receive requests from external programs outside of the SAP instance via so called *RFC* protocol, defined by SAP. The gateway process will forward valid RFC requests in some internal format to the dispatcher process which then selects a work process to handle the request.
The HA Test Tool used in the testing is one example of an external program sending request to the gateway process. The found errors related to the communication between HA Test Tool and gateway might also occur when other external programs communicate with the gateway.
Usually only few of the following kind of error messages occurred during HA Test Tool usage or none at all. The number of messages can increase to a high amount, if the so called *blackout period* of the AIX Live Update, where both the Original LPAR and the Surrogate LPARs are frozen, takes longer than 60 seconds.

write error

Example:

```
10:14:23 RD  Q0 I Operating system call writev failed (error no. 32 )

10:14:23 RD  S2 3 Connection to CPI-C client 273 was closed
```

Usually the "writev" error triggers the "Connection ... closed" error as a follow up message.

The AIX error code 32/EPIPE is documented as "Broken pipe".
Explanation: A keep alive request is sent from SAP gateway process to the HA Test Tool, but not answered. These keep alive messages belong to the underlying SAP RFC protocol and are assessed as uncritical.

recv error

Example:

```
10:14:23 RD  Q0 I Operating system call recv failed (error no. 73 )
The AIX error-code 73/ECONNRESET  is documented as "Connection reset by
peer".
```

Explanation: The SAP gateway process is waiting for further data from the HA Test Tool, but not receiving. That is also assessed as being part of the keep-alive mechanism, and so uncritical.

CPIC Communication error

A log message like the following appears, flagged as "warning", telling about "Communication error" and "CPIC", always followed by two informational messages.

```
10:55:31| DIA |007       |001|HATT  |        |R49       |
Communication error, CPIC return code 020, SAP return code 223
|
10:55:31| DIA |007       |001|HATT  |        |R5A       |
> Conversation ID: 75869720
10:55:31|XXXXXX_A75_02|DIA |007       |001|HATT  |        |R64       |
> CPI-C function: CMSEND(SAP)
```

Explanation: This is also assessed as related to the RFC protocol used by the HA Test Tool and can be avoided or at least diminished by setting the parameter jco.cpic_pong_timeout to a value that is higher than the expected blackout period.

Buffer synchronization

10:14:23 DIA 037000SAPSYSBS 5
The buffer synchronization has not been called for 234 seconds
This is assessed as an invalidation of some SAP internal caches due to too long outstanding synchronization and seen as uncritical.

Conclusion

The AIX feature of Live Update was validated successfully on SAP systems with ABAP stack. The project experiences show it to work well, if the systems are set up according to the guidelines given in the IBM knowledge center documentation, and according to the best practices given in this paper.

References

IBM knowledge center, official web documentation for AIX Live Update
https://www.ibm.com/support/knowledgecenter/ssw_aix_72/install/live_update_install.html

It contains a valuable video documentation by David Sheffield named “No Reboot Required! AIX 7.2.1 Live Update”
https://www.youtube.com/watch?v=dHvBQOXtjaY

IBM AIX 7.2 Live Update - Experiences testing AIX 7.2 Live Update feature with Oracle Database 12c Release 2
https://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102794

IBM Best practices for the Live Update function
https://www.ibm.com/support/knowledgecenter/ssw_aix_72/install/lvupdate_best_practice.html

Configuring IBM Db2 with BLU Acceleration to support IBM AIX Live Update
https://developer.ibm.com/articles/au-aix-db2-blu/

SAP Test Tool for High Availability Environments
https://blogs.sap.com/2016/01/28/test-tool-for-high-availability-environments/

Community Weblinks about Live Update best practices
http://gibsonnet.net/blog/cgaix/html/Chriss_AIX_Live_Update_Best_Practices.html

Copyrights and Trademarks

© Copyright 2020 IBM Corporation. All Rights Reserved.
IBM Corporation
Systems and Technology Group
Route 100
Somers, NY 10589

Neither this documentation nor any part of it may be copied or reproduced in any form or by any means or translated into another language, without the prior consent of the IBM Corporation.
IBM makes no warranties or representations with respect to the content hereof and specifically disclaim any implied warranties of merchantability or fitness for any particular purpose. IBM assumes no responsibility for any errors that may appear in this document. The information contained in this document is subject to change without any notice. IBM reserves the right to make any such changes without obligation to notify any person of such revision or changes. IBM makes no commitment to keep the information contained herein up to date.

Edition Notice: 2020
This is version 1.0 of this document.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. A current list of IBM trademarks is available on the web at:
Copyright and trademark information.

Adobe and PostScript are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States and/or other countries.
Intel, Intel Xeon, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc., in the United States, other countries, or both.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Linux is the registered trademark of Linus Torvalds in the United States, and other countries.
Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
SAP HANA, SAP NetWeaver are trademarks or registered trademarks of SAP Corporation in the United States, other countries, or both.
Oracle is a registered trademark of Oracle Corporation and/or its affiliates.
Other company, product or service names may be trademarks or service marks of others.

Information is provided "AS IS" without warranty of any kind.

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 home pages. 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.

Disclaimer and Special Notices

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.
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.
Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems.

Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.
Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.
This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products.
All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs
ANY INFORMATION HEREIN IS PROVIDED “AS IS” WITHOUT WARRANTY OR INDEMNIFICATION OF ANY KIND BY IBM AND DO NOT ANY EXPRESS OR IMPLIED, WARRANTIES OR CONDITIONS OF MERCHANTABILITY, FITNESS OR USAGE FOR PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT.