


Session DDB-1176
Upgrading to DB2 With the Latest
Version of BLU Acceleration

Melanie Stopfer
mstopfer@us.ibm.com

 mstopfer1

Insight2015
Lead in the **Insight Economy**

© 2015 IBM Corporation #ibminsight

Melanie Stopfer is a IBM Consulting Learning Specialist and Developer and recognized worldwide as a DB2, IBM BLU and Analytics subject matter expert. As a Certified DB2 10.5 Advanced Technical Expert, she has provided in-depth technical customer support in recovery, performance and database upgrade best practices since 1988. In 2009, Melanie was the first DB2 LUW speaker to be inducted into the IDUG Speaker Hall of Fame and selected Best Overall Speaker at IDUG EMEA 2009, 2011 and 2012, and IDUG NA 2009, 2012 and 2014 and included in Top 10 Presenters in all years presenting. She has successfully presented at all IBM IOD and Insight conferences and been recognized as a top presenter.

Please connect with Melanie on Twitter at @mstopfer1 and LinkedIn.

Objectives

- Describe changes to DB2 10.5 packaging and software and operating system requirements.
- Understand DB2 10.5 changes to DDL, utilities, tools and monitoring.
- Configure a database and instance to implement DB2 10.5 new features.
- Analyze DB2 10.5 changes that affect applications and database runtime behavior.
- Develop a DB2 10.5 upgrade strategy



Insight2015

Lead in the Insight Economy

Abstract:

Are you responsible for upgrading your database to the latest version of DB2 with BLU Acceleration? Everyone is talking about in-memory computing and columnar databases to improve analytic query performance. Take advantage of columnar data store integrated into the core DB2 engine which has huge benefits for Business Intelligence workloads. In addition to BLU, there is additional new DB2 great functionality. Unlock the secrets. If you need a detailed step-by-step approach to implement the newest version of DB2, come learn about often overlooked but very important best practices to understand before and after upgrading. Satisfy your ever increasing demands to analyze business data quickly and get insights into your data.

Please note my other DB2 LUW Sessions



- **DDB-1178: Meet the Experts:
Power of db2pd: Monitor and Troubleshoot Your World**
Date/Time: Mon, 26-Oct, 02:30 PM-03:20 PM
Venue: Mandalay Bay South Convention Center Level 1
Room: Meet the Experts - Forum 2, EXPO
- **DDB-1179:
Enhancements and Tips for DB2 LUW and IBM BLU**
Date/Time: Wed, 28-Oct, 02:30 PM-03:30 PM
Venue: Mandalay Bay South Convention Center Level 3
Room: South Seas D
- **DDB-1177:
DB2 Memory Management: Have You Lost Your Memory?**
Date/Time: Thu, 29-Oct, 09:45 AM-10:45 AM
Venue: Mandalay Bay South Convention Center Level 3
Room: South Seas D

insight2015



Insight2015

Lead in the **Insight Economy**

2

Planning Your Upgrade to DB2 10.5 LUW



ibm/insight

DB2 10.5 LUW Upgrade from *pre-DB2 Version 10.5* copy:


- DB2 10.1-> DB2 10.5
- DB2 9.8 -> DB2 10.5
- DB2 9.7 -> DB2 10.5

If DB2 servers are running on release prior DB2 9.7:

- Upgrade to DB2 9.7 or 10.1 latest fixpack
- Then upgrade to DB2 10.5

See notes for more details/URLs:

- Upgrading to DB2 V10.5 Roadmap
- DB2 Upgrade Portal
- DB2 V10.5 English Manuals
- DB2 V10.5 Translated Manuals



SC27-5519-00 What's New for DB2 V10.5
SC27-5513-00 Upgrading to DB2 V10.5 Guide



Insight2015

Lead in the Insight Economy

7

DB2 10.5 for Linux, UNIX and Windows became generally available on June 14, 2013. In the upgrade tasks, the term *pre-DB2 Version 10.5 releases* refers to DB2 Version 10.1, DB2 Version 9.8 or DB2 Version 9.7.

Refer to the DB2 Information Center for details about the installation and upgrade process. It is highly recommended that you consult the DB2 IC while planning and executing an upgrade. You should also test the upgrade process on non-production servers first. Your first point of reference to learn how to upgrade each component of your DB2 environment and how to create a upgrade plan is the **SC27-5513-00 *Upgrading to DB2 Version 10.5 Guide***. The content of this guide is available in the *DB2 Information Center*. You may also download the DB2 Version 10.5 for Linux, UNIX, and Windows English manuals from the following website:

<http://www-01.ibm.com/support/docview.wss?uid=swg27038855>

The DB2 Version 10.5 for Linux, UNIX, and Windows translated manuals can be downloaded from <http://www-01.ibm.com/support/docview.wss?rs=71&uid=swg27038856>

Checking Whether Upgrade is Possible

Direct upgrades are supported from Version 9.7 and higher systems. Upgrading your DB2 server requires that you install a DB2 10.5 copy and then upgrade all the instances and databases to be able to run them under the DB2 10.5 copy. You can directly upgrade existing DB2 9.8, 9.7 or DB2 10.1 instances and databases to DB2 10.5. On Windows operating systems, you have an option to automatically upgrade an existing pre-V10.5 DB2 copy. If you choose to upgrade your existing DB2 copy during installation, you only need to upgrade your databases after installation.

Review also the DB2 upgrade portal at www.ibm.com/support (formerly known as DB2 migration portal) that provides access to additional resources and up-to-date information about the upgrade process as they become available. These resources include educational material, white papers, and webcasts for upgrade. The **DB2 Upgrade Portal** is located at <http://www-01.ibm.com/support/docview.wss?uid=swg21200005>

The Upgrading to DB2 Version 10.5 Roadmap is located at <http://www-01.ibm.com/support/docview.wss?uid=swg21633449>

Supported Client Access and Client Upgrade



#Domainsight

- **ALWAYS UPGRADE SERVERS before upgrading CLIENTS!**
- **DB2 10.5, 10.1 and DB2 9.7 IBM clients and drivers can access remote DB2 10.5 Server:**
 - IBM Data Server Driver Package
 - IBM Data Server Driver for JDBC and SQLJ
 - IBM Data Server Driver for ODBC and CLI
 - IBM Data Server Runtime Client
 - IBM Data Server Client
- **If use pre-10.5 clients to access DB2 10.5 servers, applications are restricted to functionality available for that client.**
- If different versions of client and DB2 server located on same system, local client-to-server connections using IPC are not supported. Establish connection as remote connection (loopback connection) using TCP/IP.
- When later-level client accesses earlier-level server, later level client functionality not available to server.
- Upgrading client involves installing DB2 10.5 client copy and then upgrading client instance.
- Current level of client installed determines way to proceed with upgrade to DB2 10.5.
 - Can directly upgrade to DB2 10.5 clients from DB2 10.1 or DB2 9.7.
 - **If have DB2 9.5 or earlier clients, migrate to any DB2 10.1 or DB2 9.7 client first.**
- Which JDBC driver (JCC) version corresponds with each DB2 release and Fix Pack level?
 - <http://www-01.ibm.com/support/docview.wss?uid=swg21363866>
 - For JDBC or SQLJ applications, if you are using the IBM DB2 Driver for SQLJ and JDBC, you can determine the level of the driver by running `db2jcc -version`



Insight2015

Lead in the **Insight Economy**

10

A client instance allows connecting application to database and keeps information about client configuration, cataloged nodes, and cataloged databases. The upgrade process consists of all the tasks that you must perform to have your DB2 environment running successfully after upgrading. These components are DB2 servers, DB2 clients, database applications, and routines. Upgrading DB2 servers to DB2 10.5 requires an understanding of upgrade concepts, upgrade restrictions, upgrade recommendations, and your DB2 server. When you have a complete understanding of what upgrading your DB2 server involves, you can create your own upgrade plan. The upgrade of each of the components in your DB2 environment requires that you perform different tasks:


- *Upgrading DB2 servers* involves upgrading your existing instances and databases so that they can run in the new release.
- *Upgrading clients* involves upgrading your client instances to keep the configuration of your existing clients.
- *Upgrading database applications and routines* involves testing them in the new release and modifying them only when you must support changes in this new release.

The IBM clients and drivers are as follows:

- IBM Data Server Driver Package
- IBM Data Server Driver for JDBC and SQLJ
- IBM Data Server Driver for ODBC and CLI
- IBM Data Server Runtime Client
- IBM Data Server Client

Features and Functions by DB2 LUW Editions

	EXPRESS c	EXPRESS Server	Workgroup Server	Advanced Workgroup	Enterprise Server	Advanced Enterprise
Storage Optimization	No	No	No	Yes	No	Yes
pureScale Cluster	No	No	IBM DB2 Business Application Continuity Offering	Yes	IBM DB2 Business Application Continuity Offering	Yes
HADR	No	Yes	Yes	Yes	Yes	Yes
BLU Acceleration In-Memory	No	No	IBM DB2 BLU Acceleration In-Memory Offering	Yes	IBM DB2 BLU Acceleration In-Memory Offering	Yes
Native Database & Native Backup Encryption	Yes	IBM DB2 Encryption Offering	IBM DB2 Encryption Offering	Yes	IBM DB2 Encryption Offering	Yes
Multi-Temperature Storage	No	No	No	Yes	Yes	Yes
Database Partitioning	No	No	No	Yes	No	Yes
Oracle Compatibility	Yes	Yes	Yes	Yes	Yes	Yes
Temporal Tables	Yes	Yes	Yes	Yes	Yes	Yes
Range Partitioned Tables	No	No	Yes	Yes	Yes	Yes
Workload Management	No	IBM DB2 Performance Management Offering	IBM DB2 BLU Acceleration In-Memory or Performance Mgt. Offering	Yes	IBM DB2 BLU Acceleration In-Memory or Performance Mgt. Offering	Yes
Data Studio	Yes	Yes	Yes	Yes	Yes	Yes


Insight2015 Lead in the **Insight Economy**

The following new DB2 offerings are available as add-ons to different DB2 product editions:

- IBM DB2 BLU Acceleration In-Memory Offering
- IBM DB2 Business Application Continuity Offering
- IBM DB2 Encryption Offering
- IBM DB2 Performance Management Offering

The DB2 Editions are:

- DB2 Advanced Enterprise Server Edition and DB2 Advanced Workgroup Server Edition include all the above options
- DB2 Enterprise Server Edition allows the above options to be added by registering an additional associated license certificate file with a minimum purchase of 25 users per 100 Processor Value Units
- DB2 Workgroup Server Edition allows the above options to be added by registering an additional associated license certificate file. You must acquire a separate user license for each authorized user of this product, with a minimum purchase of five users per server.
- DB2 Express Server Edition allows adding the IBM DB2 Encryption Offering and/or the IBM DB2 Performance Management Offering by registering an additional associated license certificate file. You must acquire a separate user license for each authorized user of this product with a minimum purchase of five users per server.
- DB2 Express-C Edition includes by default the native backup encryption and native database encryption functionality introduced through the IBM DB2 Encryption Offering. No other DB2 add-on offerings can be added to DB2 Express-C.

Product packaging enhancements & licensing



#Dominsight

- DB2 Product Editions descriptions
http://www-01.ibm.com/support/knowledgecenter/?lang=en#!/SSEPGG_10.5.0/com.ibm.db2.luw.licensing.doc/doc/c0058536.html
- Functionality in DB2 features and DB2 product editions details:
http://www-01.ibm.com/support/knowledgecenter/SSEPGG_10.5.0/com.ibm.db2.luw.licensing.doc/doc/r0053238.html?lang=en
- Which DB2 10.5 product edition is right for you?
<http://www.ibm.com/developerworks/data/library/techarticle/dm-1311whichdb2edition/>
- Compare the distributed DB2 10.5 database servers
<http://www.ibm.com/developerworks/data/library/techarticle/dm-1311db2compare/index.html?ca=drs>
- IBM DB2 Recommendation Tool
www-01.ibm.com/software/data/db2/editiontool#/?S_CMP=fromsocialmedia&utm_content=buffercf556&utm_medium=social&utm_source=twitter.com&utm_campaign=buffer?S_CMP=fromsocialmedia
- Read about products and view related licensing and marketing information:
<http://www-01.ibm.com/software/data/db2/linux-unix-windows/>
- Licensing distributed DB2 10.5 servers in HA environment
<http://www.ibm.com/developerworks/data/library/techarticle/dm-1311db2halicensing/>



Insight2015

Lead in the **Insight Economy**

15

The functionality in DB2 features and DB2 product editions is explained at http://www-01.ibm.com/support/knowledgecenter/#!/SSEPGG_10.5.0/com.ibm.db2.luw.licensing.doc/doc/r0053238.html.

IBM has updated product packaging to provide simplicity through fewer packages, and greater value through more function and features included in the base DB2 editions. To read about these products and to view the related licensing and marketing information, see the DB2 product page at <http://www-01.ibm.com/software/data/db2/linux-unix-windows/>.

Program charges:

- **DB2 Express Edition** can be deployed on pervasive SMB operating systems, such as Linux, Windows, or Solaris and includes pureXML, web services federation, DB2 Homogeneous Federation, Homogeneous SQL Replication, and backup compression. If licensed as a yearly subscription (DB2 Express Edition FTL), it also includes High Availability feature provided both primary and secondary servers in the high availability cluster are licensed. If licensed under the Limited Use Virtual Server metric, DB2 Express Edition uses up to eight cores on the server. The DB2 data server cannot use more than 64 GB of memory per server. You must acquire a separate user license for each authorized user of this product with a minimum purchase of five users per server.
- **DB2 Express-C** can be used for development and deployment at no charge, and can also be distributed with third-party solutions without any royalties to IBM. It can be installed on physical or virtual systems with any amount of CPU and RAM, and is optimized to use up to a maximum of 2 cores and 16 GB of memory. DB2 Express-C is refreshed at major release milestones and comes with online community-based assistance. Users requiring more formal support, access to fix packs, or additional capabilities such as high availability clustering and replication features, can purchase optional yearly subscription for DB2 Express Edition (FTL) or upgrade to other DB2 editions.
- **DB2 Developer Edition** - The software in this package cannot be used for production systems. You must acquire a separate user license for each Authorized User of this product.

Benchmark DB2 server performance – before and after

- Run performance tests before upgrading your DB2 server.
- Use **db2batch** benchmark tool to collect elapsed and CPU times for running queries.
`db2batch -d DBX -f input.sql -r output.fil -i complete -o e yes -isol CS`
- Record the exact environment conditions where you run your tests.
- Keep record of **db2exfmt** command output for each test query.
- Compare the results before and after upgrade.
- Helps identify and correct any performance degradation that might occur.



Insight2015

Lead in the Insight Economy

16

Run a number of performance tests before upgrading your DB2 server. The db2batch benchmark tool helps you to collect elapsed and CPU times for running queries. You can use this tool to develop performance tests. Record the exact environment conditions where you run your tests.

Also, keep a record of the db2exfmt command output for each test query. Compare the results before and after upgrade. This practice can help to identify and correct any performance degradation that might occur.

db2prereqcheck command

- Check if system meets DB2 10.5.0.0 prerequisites:

- **db2prereqcheck –v 10.5.0.0 -s**

-v <version>

-s Print prerequisite validation summary on screen

-c Checks prerequisites for thin client

-p Checks prerequisites for DB2 pureScale environment

-i Checks the non-DB2 pureScale prerequisites for the latest DB2 version that is defined in the resource XML file.

If do not specify any options, both DB2 pureScale and non-DB2 pureScale server prerequisites checked for all DB2 versions listed in DB2prereqs.xml.

The -i and -v parameters are mutually exclusive.

-o <report-file>

- Uses resource XML file that contains prerequisites. Default path of XML file located in: <DB2 installation>/cfg/DB2prereqs.xml



Insight2015

Lead in the Insight Economy

20

The supported UNIX, Linux and Windows operating systems have changed in DB2 Version 10.5. Review the installation requirements for DB2 servers and IBM data server clients to determine whether your operating system version is supported and if you need to upgrade your operating system before installing DB2 10.1. Newer versions of operating systems can also bring new hardware requirements. Performing hardware and operating system upgrades separately from DB2 database product upgrade simplifies problem determination if you encounter upgrade difficulties. If you upgrade your software or hardware before a DB2 database product upgrade, ensure that your system is operating as expected before attempting to upgrade your DB2 database product. Checks whether your system meets the prerequisites for the installation of a specific version of DB2 Database for Linux, UNIX, and Windows. By using the db2prereqcheck command, you can determine whether your system satisfies the prerequisites before you download DB2 Database for Linux, UNIX, and Windows and start the installation process. The db2prereqcheck command uses a resource XML file that contains the prerequisites. The default path of the XML file is located in DB2 installation/cfg/DB2prereqs.xml. You must have read or write permissions on the XML file. Do not modify the contents of the XML file.


- -c Checks the prerequisites for thin client. The -c, -u, -g, and -p parameters are mutually exclusive.
- -u Checks the uDAPL requirement for the DB2 pureScale environment (Linux and AIX only). The -c, -u, -g, and -p parameters are mutually exclusive. On RoCE networks, you must ensure the AIX and uDAPL software prerequisites are satisfied. The db2prereqcheck command does not automatically validate these levels.
- -g Checks the GPL compilation requirement (Linux only). The -c, -u, -g, and -p parameters are mutually exclusive.
- -p Checks the prerequisites for the DB2 pureScale environment (Linux and AIX only). The -c, -u, -g, and -p parameters are mutually exclusive.
- -t network configuration type. The -t parameter validates prerequisites requirement for a specific type of network configuration (Linux only). The network configuration type must be SINGLE_IB_PORT_CLUSTER, MULT_IB_PORT_CLUSTER, SINGLE_ROCE_PORT_CLUSTER or MULTI_ROCE_PORT_CLUSTER.
- -s Print prerequisite validation summary on screen.
- -i Checks the prerequisites for the latest DB2 version that is defined in the resource XML file.
- To display the latest version number in the XML file, run -i with -s. The -i and -v parameters are mutually exclusive.
- -v version - Checks the prerequisites for the specific DB2 version. You can use this parameter with other db2prereqcheck command parameters. The -i and -v parameters are mutually exclusive.
- -f xml-file - Specifies the name of the XML resource file. The default XML file (with file path : DB2 installation/cfg/DB2prereqs.xml) will be used if the -f parameter is not specified.
- -o report-file - Specifies a name for the report file. For example, db2prereqcheck.rpt.
- -hl list_of_hostnames - Specifies a list of hosts that are checked for passwordless root SSH access between all hosts.
- -dl list_of_shared_devices - Specifies one or more PVIDs (on AIX operating systems) or WWNs (on Linux operating systems) of the shared disks that are verified to make sure that they are accessible by all hosts. If you specify this parameter, you must also specify the -hl parameter.
- -iu instance_user - Specifies the instance user name. The UID and GID of the instance user name are verified to make sure they are the same across all hosts. If you specify this parameter, you must also specify the -hl parameter.
- -nl list_of_netnames - Specifies the list of netnames. These netnames are used to ping the RoCE & IB networks to verify that they are pingable between all hosts.
- -upgrade - Checks whether the required entry in the netmon.cf file is correct when you are upgrading from the GA level of Version 9.8 or Version 9.8 Fix Pack 1.

INSTALLATION METHODS



#Dominsight

- On Linux/UNIX, use **db2setup** command or **response file** installation method.
 - `./db2setup -l /temp/db2setup.log -t /tmp/db2setup.trc`
 - `./db2setup -r responsefile_directory/response_file`
 - **-c validates contents of response file before install fix pack. Ensures contents of response file correct before installation.**
- On Windows, use **setup** command or **response file** installation method.
 - `setup -l c:\temp/db2setup.log -t c:\tmp/db2setup.trc`
 - `setup -u c:\ responsefile_directory/response_file`



- **db2_install** is deprecated and might be removed in future release
- **GC27-5514-00 Installing DB2 Servers**
- **GC27-5515-00 Installing IBM Data Server Clients**



Insight2015

Lead in the **Insight Economy**

21

The `db2setup` command queries the operating system to determine the existing language settings. If the language setting of your operating system is supported by `db2setup`, then that language will be used when displaying the DB2 Setup wizard. If your system uses the same code pages but different locale names than those supported by the DB2 interface, you can still see the translated `db2setup` by setting your `LANG` environment variable to the appropriate value. The `db2setup` command is only available on Linux and UNIX operating systems. The command for Windows operating systems is **setup**.

To start the IBM DB2 Setup Launchpad, run the `db2setup` command from the directory where the DB2 installation image resides:

```
./db2setup -l /temp/db2setup.log -t /tmp/db2setup.trc
```

- The `-l` and `-t` recommended parameters enable DB2 logging and tracing.

To create a response file use the DB2 Setup wizard:

- Type the `db2setup` command to launch the DB2 Instance Setup Wizard.
- On the Select installation, response file creation or both panels, select the **Save my installation settings in a response file** option.

After creating a response file, you can make changes to the response file to activate or deactivate keywords.

```
./db2setup -r responsefile_directory/response_file
```

Checklist: First steps after installation



#Dominsight

1. [Deleting a Firefox browser profile](#)

After run DB2 Setup wizard or DB2 First Steps, may want to delete browser profile to prevent JavaScript from running automatically when visiting other Web pages with this profile.

2. [Applying DB2 licenses](#)

To license your product, register appropriate license key with DB2 database product.

3. [Updating DB2 licenses](#)


Can switch between editions in DB2 10.5 - Workgroup, Enterprise, and Advanced Editions common image.

4. [Post-installation tasks for DB2 database servers](#)

After installing DB2 database, you must complete a number of tasks.

5. OPTIONAL: [First Steps interface](#)

Launches First Steps interface, which contains links to functions that you need to begin learning about and using DB2 product.



Insight2015

Lead in the **Insight Economy**

30

After you install a DB2 database product, there are additional necessary and optional tasks available.


2. Applying DB2 licenses – db2licm



#Dominsight

- If want license key added automatically during product installation, copy license key to **/db2/license** directory of installation image before launching DB2 Setup wizard.
- If do not add any license files to /db2/license directory, see *"License not registered"* message from **db2licm -l** command.

- Can apply licenses manually after installation by running **db2licm -a**.



- If register valid license key, **db2licm -l** command successfully lists all products with available license information.



Insight2015

Lead in the **Insight Economy**

32

To license your product, register the appropriate license key with your DB2 database product.

If you want the license key added automatically during product installation, copy the license key to the /db2/license directory of the installation image before launching the DB2 Setup wizard.

If you do not add any license files to the /db2/license directory, you will see a "License not registered" message in the output from the db2licm -l command. The following base licenses are available at the /db2/license directory of the server image and the product-specific base license is automatically applied during DB2 product installation.

- db2ese.lic
- db2exp.lic
- db2consv.lic

You can apply the licenses manually after the installation by running the db2licm -a command. If you register a valid license key, the db2licm -l command will successfully list all the products with available license information

To use DB2 features or to be in compliance with your current product entitlement, you might be required to register additional license keys.

How to download DB2 license?



#Dominsight

- **License files for DB2 Version 10.5 are shipped separately.**
 - Done for convenience so downloading license file takes less time due to small size.
 - Download license activation key from Passport Advantage.
 - Then install it.



How do I download my DB2 license from Passport Advantage?

[http://www-01.ibm.com/support/docview.wss?](http://www-01.ibm.com/support/docview.wss?uid=swg21305417&myns=swgimgmt&mynp=OCSSEPGG&mynp=OCSSEPDU&mync=E)

[uid=swg21305417&myns=swgimgmt&mynp=OCSSEPGG&mynp=OCSSEPDU&mync=E](http://www-01.ibm.com/support/docview.wss?uid=swg21305417&myns=swgimgmt&mynp=OCSSEPGG&mynp=OCSSEPDU&mync=E)

DB2 10.5 Activation Key Part Numbers

<http://www-01.ibm.com/support/docview.wss?uid=swg21641360>



Insight2015

Lead in the **Insight Economy**

33

How do I download my DB2 license from Passport Advantage? The license files for DB2 Version 10.5 are shipped separately. This is done for convenience so that downloading the license file takes less time due to its small size. You will need to download the license activation key from Passport Advantage and then install it.

A detailed example of how to obtain the activation key for your DB2 product is shown at: <http://www-01.ibm.com/support/docview.wss?uid=swg21305417&myns=swgimgmt&mynp=OCSSEPGG&mynp=OCSSEPDU&mync=E>

The Activation key part numbers for DB2 10.5 products are located at the following website: <http://www-01.ibm.com/support/docview.wss?uid=swg21641360>

3. Updating licenses

- To switch between DB2 10.5 Workgroup, Enterprise, and Advanced Edition licenses:
 - Cannot apply DB2 pureScale license over DB2 Advanced Enterprise Server Edition or DB2 Advanced Workgroup Server Edition terabyte license. If a DSF license exists then cannot apply DB2 Advanced Enterprise Server Edition or DB2 Advanced Workgroup Server Edition terabyte licenses. DB2 pureScale license cannot coexist with DB2 Advanced Enterprise Server Edition or DB2 Advanced Workgroup Server Edition terabyte licenses.
- Run **db2licm -l** to list all DB2 product licenses registered on system.
db2licm -l product-identifier
- Run **db2licm -r** to remove existing product license, feature license, and trial license.
db2licm -r product-identifier
- Run **db2licm -a** to apply new DB2 product license on system. If existing product license was not removed by running db2licm -r, it is automatically removed.
db2licm -a product-identifier
- Run **db2licm -l** to view newly applied DB2 product license information.
db2licm -l product-identifier
- Run **db2start** to restart instance, and for license compliance report to take effect for new license.
db2start



You can switch between editions in the DB2 Version 10.5 - Workgroup, Enterprise, and Advanced Editions common image.

4. Post-installation tasks: DB2 Windows servers



#Dominsight

- Add users to DB2ADMNS and DB2USERS user groups
- Update 32-bit DB2 instances to 64-bit instances
- **Validating DB2 copy: `db2val -a`**
- Set up notification and contact lists
- List DB2 copies and change default DB2 or default IBM database client interface copy after installation:
`db2swtch -l`
`db2swtch -d DB2_copy_name`
- Check for DB2 updates
- Install SSH Secure Shell Server component after installing DB2 product
 - Provides secure platform to run commands from remote machines
 - IBM SSH for Windows required for Optim Database Administrator
http://www-01.ibm.com/support/knowledgecenter/SSEPGG_10.5.0/com.ibm.db2.luw.qb.server.doc/doc/c0059310.html?lang=en



Insight2015

Lead in the **Insight Economy**

35

- After installing DB2 database on Windows systems, you must complete a number of tasks.
- **Adding your user ID to the DB2ADMNS and DB2USERS user groups (Windows)**
 - After successfully completing a DB2 installation, you now have to add users to the DB2ADMNS or the DB2USERS groups for users that need to run local DB2 applications and tools on the machine. The DB2 installer creates two new groups. You can either use a new name or accept the default names. The default group names are DB2ADMNS and DB2USERS. This process is also applicable to instance based clients.
- **Updating your 32-bit DB2 instances to 64-bit instances (Windows)**
 - If you have more than one 32-bit DB2 copy on your system, you can convert them to 64-bit instances.
- **Validating your DB2 copy**
 - The `db2val` command ensures that your DB2 copy is functioning properly.
- **Setting up notification and contact lists**
 - For your DB2 database product to contact you or others regarding the status of your databases, set up notification and contact lists. If you did not do this during the DB2 database product installation, you can manually set up these lists.
- **Changing the default DB2 and default IBM database client interface copy after installation (Windows)**
 - The `db2swtch` command can be run from any DB2 copy or IBM data server driver copy. If no arguments, launches the utility in graphical mode. The `-l` displays a list of DB2 copies and IBM data server driver copies on the system. To switch the default DB2 copy to the name specified: `db2swtch -db2 -d DB2_copy_name-client`. The `-d` specifies the name of DB2 copy or IBM data server driver copy. To switch the default client interface copy to the name specified: `db2swtch -client -d name of DB2 copy or IBM data server driver copy-d DB2_copy_name`. To switch both default DB2 copy and client interface copy to the name specified: `db2swtch -d DB2_copy_name`.
- **IBM data server client connectivity using multiple copies**
 - When using multiple DB2 copies of DB2 database products or multiple data server driver copies, various options are available for applications to access the DB2 databases. Existing applications will continue to function properly.
- **Working with existing DB2 copies**
 - You can install and run multiple DB2 copies on the same computer. Each DB2 copy can either be at the same DB2 database product level or at a different DB2 database product level.



4. Post-installation tasks: DB2 Linux/UNIX servers

- List DB2 database products installed on system: **db2ls**
 - To list info about DB2 database products or features in particular installation path `q` parameter must be specified:
 - **db2ls -q -p -b baseInstallDirectory**
 - `q` specifies querying product or feature. Parameter is mandatory.
 - `p` specifies listing displays products rather than listing features.
 - `b` specifies installation directory of product or feature. Parameter is mandatory if not running command from installation directory.
 - Forgot which DB2 features you installed?
`/opt/IBM/db2/V10.5/install/db2ls -q -a -b /opt/IBM/db2/V10.5`
- Verify basic functions of DB2 copy by checking state of installation files, instance setup, and local database connections.
 - **Validate all instances for DB2 copy: `db2val -a`**
- Main menu entries for DB2 tools (Linux)
 - After installation, can add several DB2 tools to Main menu. See notes for details.



After installing DB2 database on Linux or UNIX systems, you must complete a number of tasks.

• **Listing DB2 database products installed on your system (Linux and UNIX)**

On supported Linux and UNIX operating systems, the `db2ls` command lists the DB2 database products and features installed on your system, including the DB2 Version 10.1 HTML documentation. With the ability to install multiple copies of DB2 database products on your system and the flexibility to install DB2 database products and features in the path of your choice, you need a tool to help you keep track of what is installed and where it is installed. On supported Linux and UNIX operating systems, the `db2ls` command lists the DB2 products and features installed on your system, including the DB2 HTML documentation. The `db2ls` command can be found both in the installation media and in a DB2 install copy on the system. The `db2ls` command can be run from either location. The `db2ls` command can be run from the installation media for all products except IBM® Data Server Driver Package. The `db2ls` command can be used to list:

- Where DB2 database products are installed on your system and list the DB2 database product level
- All or specific DB2 database products and features in a particular installation path.

The output that the `db2ls` command lists is different depending on the ID used:

- When the `db2ls` command is run with root authority, only root DB2 installations are queried.
- When the `db2ls` command is run with a non-root ID, root DB2 installations and the non-root installation owned by matching non-root ID are queried. DB2 installations owned by other non-root IDs are not queried.

The `db2ls` command is the only method to query a DB2 database product. You *cannot* query DB2 database products using Linux or UNIX operating system native utilities, such as `pkginfo`, `rpm`, `SMIT`, or `swlist`. Any existing scripts containing a native installation utility that you use to query and interface with DB2 installations must change. You *cannot* use the `db2ls` command on Windows operating systems.

• **Validating your DB2 copy**

- The `db2val` command ensures that your DB2 copy is functioning properly. The `db2val` tool verifies the core function of a DB2 copy by validating installation files, instances, database creation, connections to that database, and the state of partitioned database environments. This validation can be helpful if you have manually deployed a DB2 copy on Linux and UNIX operating systems using `tar.gz` files. The `db2val` command can quickly ensure that all the configuration has been correctly done and ensure that the DB2 copy is what you expect it to be. You can specify instances and databases or you can run `db2val` against all of the instances. The `db2val` command can be found in the `DB2-install-path\bin` and `sqllib/bin` directories

db2ckupgrade - Check database for upgrade



#ibmnsight

db2ckupgrade <database-> -l filename


- In partitioned database, db2ckupgrade will check each database partition.
- On Linux/UNIX, install new DB2 copy and run from *DB2DIR*/bin directory where *DB2DIR* is location where DB2 copy is installed
- On Windows, insert DB2 product CD and run from db2\Windows\Utilities directory on CD
- SYSADM authority

db2ckupgrade -e -l filename

- e specifies all local cataloged databases are to be scanned

db2ckupgrade <database-> -l filename

- l filename specifies log file to keep list of errors and warnings



- Confirm running correct level of db2ckupgrade command. Ensure db2ckupgrade log file contains following text:

Version of DB2CKUPGRADE being run: Version 10.5



Insight2015

Lead in the Insight Economy

41

The **db2ckupgrade** command must be issued prior to upgrading the instance to verify that your databases are ready for upgrade. Backup all databases prior to upgrade, and prior to the installation of the current version of DB2 database product on Windows operating systems. Once the codebase has been installed you should then run the db2ckupgrade before upgrading the instance or databases. The db2ckupgrade fails to run against databases which are catalogued as remote databases.

The **db2ckupgrade** command verifies that a database can be migrated. In a partitioned database environment, the db2ckupgrade command will check each database partition.

Ensure that the log file for db2ckupgrade command contains the following text:

Version of DB2CKUPGRADE being run: Version 10.5. This text confirms that you are running the correct level of the db2ckupgrade command.

SYSADM authority is required to execute. No database connection is required.

db2chkupgrade or db2iupgrade -b option

- > **PROBLEM:** db2ckupgrade or db2iupgrade take long time to complete.
 - > Additional check for tables in load pending state added to db2ckupgrade tool in the DB2 10.1 and 10.5.
 - > If large number of tables, indexes, data, check can take long time.
- > **SOLUTION:** Only use if 0 rows returned!
`SELECT tabname FROM SYSIBMADM.ADMINTABINFO
WHERE LOAD_STATUS IS NOT NULL`
0 Records returned
- > **Internal -b option added in DB2 10.5 Fixpack 3 and DB2 10.1 Fixpack 3 bypasses check for tables in load pending state:**
`db2ckupgrade <dbname> -b -l <mylog.log>
db2iupgrade -b <myinstance>`

Remember: db2iupgrade calls db2ckupgrade!!



Insight2015

Lead in the Insight Economy

45

The db2ckupgrade or db2iupgrade tools may take a long time to complete when upgrading to the DB2 v10.1 or v10.5 products. An additional check for tables in load pending state was added to the db2ckupgrade tool in the DB2 v10.1 and v10.5 products. This check can take a long time if a database has a large number of tables, defined indexes and hosts a considerable amount of data. A query against SYSIBMADM.ADMINTABINFO, may be used to obtain this information from the database.

To diagnose the problem, looking at the db2diag.log file. It may show that the db2ckupgrade tool is spending a lot of time using the sqllemCkTableLoadPending() function.

The following example from db2diag.log file shows gap of 2 hours between beginning and end message for the sqllemCkTableLoadPending function:

2013-02-16-09.44.55.802740+000 I161183A335 LEVEL: Warning

PID : 5178038 TID : 1 PROC : db2ckupgrade64_exe

INSTANCE: db2inst NODE : 000

APPID : *N0.db2inst.130216094415

EDUID : 1

FUNCTION: <0>, <0>, <0>, probe:4138

DATA #1 : <preformatted>

DB2CKUPGRADE: Begin: sqllemCkTableLoadPending

Disk space requirements (1)

• System catalog table space (SYSCATSPACE)

- Increase to twice the total of used space.
- Free pages should be \geq used pages.
- System catalog table space is required for both old and new database catalogs during upgrade.



• Temporary table space (TEMPSPACE1)

- Increasing total size to twice total size of system catalog table space.

```
SELECT
SUBSTR(TBSP_NAME,1,15) NAME, TBSP_TYPE TYPE,
TBSP_AUTO_RESIZE_ENABLED AUTO_RESIZE, TBSP_NUM_CONTAINERS
CONTAINERS, TBSP_TOTAL_PAGES TOTAL_PGS, TBSP_USED_PAGES USED_PGS,
TBSP_FREE_PAGES FREE_PGS, TBSP_MAX_SIZE MAX_SZ, TBSP_PAGE_SIZE PG_SZ
FROM SYSIBMADM.TBSP_UTILIZATION
WHERE TBSP_CONTENT_TYPE IN ('ANY','SYSTEMP')"
```



Insight2015

Lead in the **Insight Economy**

47

You must be aware that the upgrade process requires additional disk space. Ensure that you have enough free disk space to complete this process successfully. The above disk space recommendations are applicable for upgrading to DB2 10.5.

Ensure that you have sufficient free space on the system catalog and the system temporary table spaces for the databases that you are upgrading. System catalog table space is required for both old and new database catalogs during upgrade. The amount of free space required varies, depending on the complexity of the database, as well as on the number and size of database objects.

To increase the amount of free space on your automatic storage table spaces, you can increase the space on the current storage paths or add a new storage path.

To increase the amount of free space on your System Managed Space (SMS) table spaces, free sufficient disk space on the corresponding file systems or increase the size of your file systems if you are using a volume manager.

To increase the amount of free space on your Database Managed Space (DMS) table spaces, you can increase the size of existing containers. You can also add additional containers although this might trigger data rebalancing. You can reduce the size of the containers after upgrade.

Disk space requirements (2)

- **Instance Upgrade Trace File written to /tmp for LINUX and UNIX**
 - Ensure 5GB of free space in /tmp directory.
- **Database log file space**
 - Upgrade changes to system catalog objects are single UOW.
 - Set LOGSECOND to twice current value of LOGPRIMARY and LOGSECOND if file system with log files has enough disk free space.
 - If already have large log file space, might not be necessary.
 - **Only need to increase log space on catalog partition.**
 - **Must update these values before upgrade the instance to DB2 10.5**
 - Optional: Consider infinite logging (LOGSECOND=-1)

db2 GET DB CFG FOR sample | grep '(LOG[FPS])' tee logsize.txt

db2 UPDATE DB CFG FOR sample using LOGSECOND
(current_value of LOGPRIMARY + current_value of LOGSECOND)

- **Index space**
 - Each index on every populated table requires 1 additional page per index to use following functionality:
 - Real-time statistics, Deferred cleanup roll out for MDC tables, Index rebuild on a populated table.
- **Automatic storage files**
 - If enable automatic storage on database with ALTER DATABASE ... ADD STORAGE, creates SQLSGF.1 and SQLSGF.2 files required for maintaining automatic storage.



#Dominsight



Insight2015

Lead in the **Insight Economy**

48

Instance upgrade trace file is written to /tmp for Linux and UNIX. Ensure at least 5GB of free space in /tmp directory.

You must update these LOGSECOND values before you upgrade the instance to DB2 10.5, because you will not be able to update these DB CFG parameters until issue the UPGRADE DATABASE command. If this command fails because there is insufficient log file space, then you can set these database configuration parameters to higher values and then re-issue the UPGRADE DATABASE command.

The new database configuration parameter settings for log space can be restored to their original value after the upgrade is complete.

If have limited amount of free disk space for indexes, can get error message SQL0289N that indicates the table space is full. Ensure that you have enough free pages in the corresponding index table space to account for one additional page per index on populated tables before:

- Populating tables in databases created in DB2 Version 9.7 or later, real-time statistics are enabled by default in these newly created databases.
- Enabling deferred cleanup roll out by setting DB2_MDC_ROLLOUT to DEFER, or when DB2_WORKLOAD is set to SAP.
- Reorganizing or recreating indexes on populated tables.

Back up databases before AND after upgrade




#Dominsight

After upgrade instances to DB2 10.5, cannot backup databases until upgrade them!

- Perform a full offline back up for each of your local databases:

1. To disconnect all applications and users:
db2 FORCE APPLICATION ALL



2. Backup database using BACKUP DATABASE.
db2 BACKUP DATABASE database_alias TO backup-dir

- In **partitioned database, back up all database partitions.**
- If performed full online or offline database backup recently and cannot perform another one before upgrading, can perform incremental offline database backup instead if **TRACKMOD=ON**.
- If activated and configured DB2 Advanced Copy Services (ACS) on DB2 9.7 or later, use USE SNAPSHOT parameter to perform snapshot backup.
 - **Can only restore snapshot backup to instance of same version.**
 - **Cannot use snapshot backup to upgrade to new version.**



Insight2015

Lead in the **Insight Economy**

50

Before and after the upgrade process to DB2 10.5, it is strongly recommended that you perform a full *offline* database backup. If an error occurs during the upgrade process, you need full database backups to recover and upgrade your databases.

After you upgrade your instances to DB2 10.5, you cannot backup databases until you upgrade them.

Backing up DB2 server configuration and diagnostic info

1. Collect information from DB2 servers by running db2support command for all databases going to upgrade or convert to DB2 pureScale, in all your instances. Collect information about database system catalog, DB CFG and DBM CFG settings, DB2 registry variables settings, explain table data, and diagnostic information required by DB2 support in case of problems.
db2support output-directory -d database-name -cl 0

- **Keep zip for several months after complete upgrade or conversion to DB2 pureScale to help resolve performance issues with new release.**

2. Back up information about all packages for applications associated with each database.
db2 LIST PACKAGES FOR ALL SHOW DETAIL > /upgrade/sample_pckg.txt

3. If enabled audit facility, back up audit cfg for each instances:
db2audit describe > audit_instance-name.cfg

4. Back up all external routines.
cp -R \$INSTHOME/sqllib/function \$INSTHOME/routine_backup

Optional Steps: db2support HTML report includes following info.


5. ** Use GET DBM CFG to back up settings for each instance:
db2 GET DBM CFG > dbm_instname.cfg

6. ** Use GET DB CFG to back up settings for each database.
If not same settings, back up DB CFG for each database partition.
db2 CONNECT TO database_alias
db2 GET DB CFG FOR database_alias SHOW DETAIL > db_database_alias.cfg

7. ** Use db2look to save additional information.
db2look -d sample -e -o sample_tbs.db2 -l -x

8. ** Use db2set command to back up DB2 profile registry variables settings:
db2set -all > reg_instname.txt

9. If set DB2 environment variables, use appropriate system command to list environment variables values.
set | grep DB2 > env_instname.txt (UNIX example)



**** Optional: Contained in db2support zip file**



Insight2015

Lead in the Insight Economy

51

Backing up your settings for database and database manager configuration parameters before DB2 server upgrade, or conversion to DB2 pureScale, allows you to verify DB2 server behavior after upgrade, or converting to DB2 pureScale, and to re-create instances and databases. In addition, you can collect information from your DB2 servers about the database system catalogs, DB2 registry variables settings, explain table data, and diagnostic information that can help in problem determination if you encounter any post-upgrade differences in the database manager behavior or performance.

The db2support -cl 0 parameter collects the database system catalog, database and database manager configuration parameters settings, DB2 registry variables settings. The information collected is stored in the db2support.zip compressed zip file under the output directory. A summary report in HTML format is included. In the db2supp_opt.zip file that is also included, you should check the optimizer.log file to verify that the collection of information was performed successfully. Keep this zip file for several months after you complete the upgrade, or conversion to DB2 pureScale. The information in the zip file can help in quickly resolving any performance issues with the new release.

The LIST PACKAGES ... FOR SCHEMA clause allows you to list all packages for a specific schema, if your application has several schemas you need to repeat this command for each schema name or use the FOR ALL clause.

When backing up external routines, INSTHOME is set to home directory of instance owner. If you have specified a full path that is not under the default routines path when you created your external routines in the database, you must ensure the existing libraries remain on their original location.

The SHOW DETAIL clause on GET DB CFG displays the values calculated by the database manager when configuration parameters are set to AUTOMATIC.

When possible, use the output from the set command and run the db2set command to set these environment variables as registry variables in the DB2 profile registry.

Validating Instance Upgrade



Review **db2iupgrade log file** written to Windows current path or UNIX/Linux instance home directory. Review **db2diag** and **administrative notification logs**.


#ibmresgpt

- Switch to the instance owner and issue:

✓ **db2start**
db2level

DB21085I This instance or install (instance name, where applicable: "inst461") uses "64" bits and DB2 code release "SQL10050" with level identifier "0601010E". Informational tokens are "DB2 v10.5.0.0", "s130528", "LINUXAMD64105", and Fix Pack "0". Product is installed at "/opt/ibm/db2/V10.5".

- On Windows, db2level shows DB2 copy name.
- To display the same information for DB2 pureScale member, use:
db2level -localMember



Insight2015

Lead in the Insight Economy

54

The db2level output shown above was ran on a Linux 64-bit operating system with DB2 10.5 and no fixpacks installed.

The db2iupgrade log file is created in the instance home directory for Linux and UNIX operating systems or in the current directory for Windows operating systems. The instance upgrade will not continue if the db2ckupgrade command returns any errors.

After upgrade, check output of db2diag and administrative notification logs.

If the db2iupgrade command completes successfully, switch to the instance owner and issue the **db2start** to restart the instance. To confirm you are now running DB2 10.5 issue the **db2level** command.

The **-localMember** parameter specifies the current version and service level for the member where the db2level command is being issued.

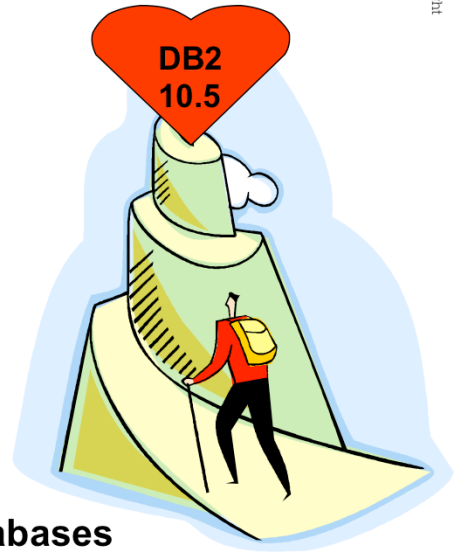
Database directory upgrade

- During first time access, database directory is implicitly upgraded if necessary.
- Database directory is accessed when issue:
 - **LIST DATABASE DIRECTORY**
 - **UPGRADE DATABASE**
- Obtain list of all local databases that must be upgraded:

db2 list db directory > db-dir-before.txt

**DBY
DBX**

**Must upgrade all databases
in upgraded instance**



#Dominsight



Insight2015

Lead in the **Insight Economy**

56

The database directory is implicitly upgraded if necessary during first time access.

After the instance upgrade has been completed, all databases need to be upgraded to the latest version. Obtain a list of all local databases using the **list database directory** command.

Smart Idea - Before upgrade, clean up diagpath

- Rename or delete db2diag log files so new files are created.
- Remove or move to another directory any existing dump files, trap files, and alert log files in diagnostic path indicated by diagpath.
- Now files only contain information about database upgrade process.
- Helps to isolate and understand any problem that might occur during database upgrade.



Insight2015

Lead in the Insight Economy

58

Rename or delete the db2diag log files so that new files are created. Also, remove or move to another directory any existing dump files, trap files, and alert log files in the directory indicated by the diagpath parameter. By doing this, the files only contain information about the upgrade process that helps you to isolate and understand any problem that might occur during database upgrade.

Upgrading Local Databases to DB2 10.5

• **db2 UPGRADE DATABASE *database_alias* -REBINDALL**

- Instance owner authority required
- **REBINDALL** rebinds all packages during upgrade and ensures step not missed and that applications not started before completed
 - Alternately, rebind all valid and invalid packages after database upgrade:
 - **db2rbind *database-name* -l *logfile* all**

db2 UPGRADE DATABASE DBY -REBINDALL

db2 UPGRADE DATABASE DBX -REBINDALL

If error occurs during upgrade, may need to issue TERMINATE



```
db2 CONNECT to <dbname>
```

```
db2 GET DB CFG FOR <dbname> SHOW DETAIL  
> <dbname>-dbcfg-after.txt
```

```
db2 LIST TABLESPACES SHOW DETAIL  
> <dbname>-tsp-after.txt
```

```
db2 LIST PACKAGES FOR ALL SHOW DETAIL  
> <dbname>-pkg-after.txt
```

```
db2 CONNECT RESET
```

```
...
```

Compare before and after txt files.

Insight2015

Lead in the Insight Economy

59

The **db2ckupgrade** command must be issued prior to upgrading the instance to verify that your databases are ready for upgrade. The **db2iupgrade** command implicitly calls the **db2ckupgrade**. Backup all databases prior to upgrade, and prior to the installation of the current version of DB2 database product on Windows operating systems.

The **upgrade database** command will only upgrade a database to a newer version, and cannot be used to convert an upgraded database to its previous version. The database must be cataloged before upgrade.

If an error occurs during upgrade, it might be necessary to issue the TERMINATE command before attempting the suggested user response. For example, if a log full error occurs during upgrade (SQL1704: Database upgrade failed. Reason code "3"), it will be necessary to issue the TERMINATE command before increasing the values of the database configuration parameters LOGPRIMARY and LOGFILSIZ. The CLP must refresh its database directory cache if the upgrade failure occurs after the database has already been relocated (which is likely to be the case when a "log full" error returns).

After each database has been upgraded repeat the statistics gathering as above and compare the two sets of information. With the exception of the system packages, the two sets of information should be identical.

The **REBINDALL** option specifies that a REBIND of all packages is performed during upgrade. Performing the REBINDs automatically during the upgrade ensures that the step is not missed and will help ensure that other applications are not started before the REBINDs are complete.

Upgrade Database Warning and Error Messages



ibm/insight

SQL1704N Database upgrade failed. Reason code "3".

Fix: Increase LOGFILSIZ

SQL1243W Database upgrade warning

Fix: Drop or rename SYSTOOLS.DB2LOOK_INFO table. Otherwise, ALTER TABLE and COPY SCHEMA statements fail to run.

- db2 "SELECT tabname, tabschema, definer FROM syscat.tables WHERE tabschema = 'SYSTOOLS' AND tabname = 'DB2LOOK_INFO'"
- If did create: db2 RENAME SYSTOOLS.DB2LOOK_INFO TO new-table-name
- If you did not create: db2 DROP TABLE SYSTOOLS.DB2LOOK_INFO

For other upgrade database warning and error messages:

http://www-01.ibm.com/support/knowledgecenter/#!/SSEPGG_10.5.0/com.ibm.db2.luw.qb.upgrade.doc/doc/t0007193.html?cp=SSEPGG_10.5.0%2F2-3-1-3-2



Insight2015

Lead in the Insight Economy

60

If the UPGRADE DATABASE command fails and returns the [SQL1704N](#) error message with a reason code that describes the cause of the failure, find this SQL error code and determine the action to take from the list of the possible solutions for each reason code. One of the most common causes of upgrade failure is that the log file space is not large enough.

If the UPGRADE DATABASE command returns the [SQL1243W](#) warning message, you must drop or rename the SYSTOOLS.DB2LOOK_INFO table. Otherwise, the ALTER TABLE and COPY SCHEMA statements will fail to run.

ENV_GET_INSTANCE_CODE_LEVELS



ibm/insight

- **db2 GET DBM CFG | grep -i CUR_EFF**
- **SELECT NAME, DEFERRED_VALUE FROM SYSIBMADM.DBMCFG**
...
cur_eff_arch_lvl V:10 R:5 M:0 F:1 I:0 SB:0
cur_eff_code_lvl V:10 R:5 M:0 F:1 I:0 SB:0
 - Informational configuration parameters
- Invoke **ENV_GET_INSTANCE_CODE_LEVELS()** function
 - can use to facilitate rolling fixpack upgrades
 - **SELECT ... FROM TABLE (SYSPROC.ENV_GET_INSTANCE_CODE_LEVELS())**

V:10 R:5 M:0 F:0 I:0 SB:0



Insight2015

Lead in the Insight Economy

62

You can see the output of instance code levels either when doing a "get dbm cfg" for the informational configuration parameters **CUR_EFF_ARCH_LVL** or **CUR_EFF_CODE_LVL**, or by invoking the **ENV_GET_INSTANCE_CODE_LEVELS()** function. The function can be used to facilitate rolling fixpack upgrades.

- **cur_eff_arch_level** - current effective architecture level configuration parameter. This parameter displays the current effective architecture level (CEAL) at what the instance is operating.
- **cur_eff_code_level** is current effective code level configuration parameter. This parameter displays the current effective code level (CECL) at what the instance is operating.

The ENV_GET_INSTANCE_CODE_LEVELS table function returns the code level for the instance, members, and cluster facilitators.

```
SELECT ... FROM TABLE (SYSPROC.ENV_GET_INSTANCE_CODE_LEVELS())  
V:10 R:5 M:0 F:0 I:0 SB:0
```

See following URL for more details:

[http://www-01.ibm.com/support/knowledgecenter/#!/SSEPGG_10.5.0/
com.ibm.db2.luw.sql.rtn.doc/doc/r0060901.html?cp=SSEPGG_10.5.0%2F3-6-1-3-7-6l](http://www-01.ibm.com/support/knowledgecenter/#!/SSEPGG_10.5.0/com.ibm.db2.luw.sql.rtn.doc/doc/r0060901.html?cp=SSEPGG_10.5.0%2F3-6-1-3-7-6l)

How to revalidate invalid DB2 SYSTEM views after upgrade?



#ibminsight

After update run query to see if invalid system views after upgrade:

```
db2 "select name, valid from sysibm.sysviews"
```

<u>NAME</u>	<u>VALID</u>
SQLFOREIGNKEYS	N

Revalidate with administrative routine.

```
db2 "CALL SYSPROC.ADMIN_REVALIDATE_DB_OBJECTS  
( 'view', 'MY_SCHEMA', 'MY_VIEW' )"
```



Insight2015

Lead in the Insight Economy

64

The ADMIN_REVALIDATE_DB_OBJECTS procedure revalidates database objects.

This procedure takes various input parameters that control the level of revalidation that is performed:

- To revalidate all of the invalid objects in the database, either specify NULL for all parameters, or call the procedure without parameters.
- To revalidate all of the invalid database objects under a specific schema, specify a value for *object_schema*, and specify NULL for *object_name* and *object_type*.
- To revalidate a specific invalid database object, specify valid values for the first parameters.
- To force the revalidation of compiled SQL PL or PL/SQL objects regardless of their state being valid or invalid, specify a value for *object_type*, *object_schema*, *object_name*, and specify 'Y' for force.
- To enable or disable debug on compiled SQL PL or PL/SQL objects being revalidated, specify 'Y' or 'N' for *debug_flag*.

This procedure will revalidate only invalid objects and regular or materialized query tables in reorg-pending state. All invalid objects can be found in SYSCAT.INVALIDOBJECTS. To find out which tables are in reorg-pending state, use the ADMIN_GET_TAB_INFO table function.

Upgrade explain tables & rerun benchmarks



#DbmInsight

- To maintain explain table info gathered from previous releases, upgrade explain tables to DB2 10.5.

db2exmig -d dbname -e explain_schema

OR

CALL SYSPROC.SYSINSTALLOBJECTS ('EXPLAIN', 'M', CAST (NULL AS VARCHAR(128)), CAST (NULL AS VARCHAR(128)))

- Renames original explain tables, creates new set of tables by using EXPLAIN.DDL file, and copies contents of original explain tables to new tables. Finally, the tool drops original explain tables and preserves any user-added columns in explain tables.
- After upgrade database, update statistics, rerun benchmarks, compare and examine explain output if necessary.
 - Use **db2batch** benchmark tool to collect elapsed and CPU times
db2batch -d DBX -f input.sql -r output.fil -i complete -o e yes -isol CS
 - Keep **db2exfmt** output for each test query



Insight2015

Lead in the **Insight Economy**

67

Run a number of performance tests before upgrading your DB2 server. The **db2batch** benchmark tool helps you to collect elapsed and CPU times for running queries. You can use this tool to develop performance tests. Record the exact environment conditions where you run your tests.

Also, keep a record of the **db2exfmt** command output for each test query. Compare the results before and after upgrade. This practice can help to identify and correct any performance degradation that might occur.

The SYSINSTALLOBJECTS procedure can migrate explain tables. The M option is only valid when used with the tool name EXPLAIN. This option migrates explain tables to be compatible with the current version.

The db2batch tool reads SQL statements and XQuery statements from a flat file, dynamically prepares and describes the statements, returns an answer set, populates explain tables, and measure the elapsed time to run each statement where the prepare, execute, and fetch times are reported separately. The default isolation level is RR.

The db2exfmt tool can be used to analyze explain output stored in explain tables.

Upgrade event monitor output tables



#Dominsight

- As event monitors enhanced, tables require new columns.
- **EVMON_UPGRADE_TABLES** stored procedure
 - upgrades definitions of existing tables along with all data
 - automatically deactivate event monitor before upgrade
 - automatically reactivates after upgrade completed

call **EVMON_UPGRADE_TABLES** (*evmon_name*,
evmon_type, *options*, *num_evmons_evaluated*,
num_evmons_to_upgrade, *num_evmons_upgraded*)

or

call **evmon_upgrade_tables** (null, null, null, ?, ?, ?)

- Can also use ALTER EVENT MONITOR statement to add new logical groups to event monitor as alternative



Insight2015

Lead in the **Insight Economy**

68

Beginning in DB2 10.5, you can upgrade event monitor output tables after you upgrade the DB2 product. This capability lets you retain any data that might exist in event monitor tables that you had before you upgraded. As event monitors are enhanced in the DB2 product, the tables they produce might change. For example, new columns might be added to a table for reporting new monitor elements. Before 10.5, if you had existing event monitors that wrote to tables that contained data that you wanted to retain, and you wanted to collect the data in the newly-added columns, you were required to manually alter them after upgrading to the new release. This alteration involved adding any of the new columns that you might want to use. If you did not add the new columns, the event monitor would work as it had in the previous release, capturing only the data supported by that the event monitor in that release. Unformatted event tables that had changed could not be upgraded at all; you were required to drop them and then re-create them.

The EVMON_UPGRADE_TABLES stored procedure upgrades the definitions of existing event monitor tables to match those produced by the current level of the DB2 product. This feature lets you keep any existing tables that you might have, along with all the data they contain, eliminating the need to manually alter, or to drop and re-create tables.

The EVMON_UPGRADE_TABLES procedure works with both regular and UE tables. For regular tables, the procedure adds any new columns needed, drops old columns that are no longer required, and alters any columns as needed. For UE tables, the procedure adds new columns and modifies existing columns as needed to allow the UE table to be processed by the db2evmonfmt tool, or the EVMON_FORMAT_UE_TO_TABLES or EVMON_FORMAT_UE_TO_XML routines.

Important: Any active event monitors must be deactivated for the upgrade process to work properly. The EVMON_UPGRADE_TABLES procedure automatically deactivates any active event monitors before it begins upgrading tables. Do not reactivate any event monitors with tables being processed by EVMON_UPGRADE_TABLES, or the upgrade process will fail. Any event monitors that were active before the upgrade are activated again after the upgrade has completed.

Implications of not upgrading event monitor tables

As in past releases, you can choose to not upgrade your event monitor tables. However, any new columns that have been added to the event monitor in the new release will not be populated with data, and will not be available for queries. Also, the values for any monitor elements that previously existed in the old release and that increased in size in the new release might be truncated. For example, if a monitor element increases in size from VARCHAR(20) to VARCHAR(128) in the new release, and you do not upgrade the previously-existing tables, the column that contains the monitor element values will continue to store only 20 characters of data, even though the system may be sending 128-bytes of data for that monitor element to the event monitor.

Migrating Data to IBM DB2 with BLU Acceleration (1)



ibmresearch

- **db2set WORKLOAD=ANALYTICS** automatically sets configuration parameters when create new database unless disable Configuration Advisor.
- **TIP: Running AUTOCONFIGURE against existing database when DB2_WORKLOAD=ANALYTICS has the same result.**
 - dft_table_org (default table organization for user tables) db cfg parm set to COLUMN.
 - dft_degree (default degree) db cfg parm set to ANY.
 - dft_extent_sz (default extent size) db cfg parm set to 4.
 - catalogcache_sz (catalog cache) db cfg parm set to value > value for non-analytics workload.
 - sortheap (sort heap) and sheapthres_shr (sort heap threshold for shared sorts) db cfg parm are calculated specifically for analytics workload. Settings take into account additional memory requirements for processing column-organized data.
 - util_heap_sz (utility heap size) db cfg parm set to value that takes into account additional memory required to load data into column-organized tables.
 - auto_reorg (automatic reorganization) db cfg parm set to ON.
- The following extra choices are made automatically:
 - Default database page size for newly created database is set to 32 KB.
 - Intraquery parallelism enabled for any workload (including SYSDEFAULTUSERWORKLOAD) that inherits intrapartition parallelism setting from intra_parallel (enable intrapartition parallelism) dbm cfg parm, even if intra_parallel set to NO.
 - Concurrency control enabled on SYSDEFAULTMANAGEDSUBCLASS.
- Automatic table maintenance performs space reclamation for column-organized tables by default. For other considerations see
http://www-01.ibm.com/support/knowledgecenter/SSEPGG_10.5.0/com.ibm.db2.luw.admin.dboj.doc/doc/c0060592.html?lang=en
<http://www.ibm.com/developerworks/data/library/techarticle/dm-1309db2bluaccel/index.html?ce=ism0070&ct=is&cmp=ibmsocial&cm=h&cr=crossbrand&ccy=us#Scenario2>



Insight2015

Lead in the Insight Economy

72

DB2_WORKLOAD registry variable is available to be set on all operating systems. The default is that it is not set. Each value for DB2_WORKLOAD represents a specific grouping of several registry variables with predefined settings. Only one value for DB2_WORKLOAD can be in effect at any given time. Use ANALYTICS setting before creating the database to establish an optimal default configuration for analytic workloads. The ANALYTICS option ensures that configuration parameters are automatically set as follows unless you disable the Configuration Advisor:


- The dft_table_org (default table organization for user tables) database configuration parameter is set to COLUMN.
- The dft_degree (default degree) database configuration parameter is set to ANY.
- The dft_extent_sz (default extent size) database configuration parameter is set to 4.
- The catalogcache_sz (catalog cache) database configuration parameter is set to a value that is higher than the value for a non-analytics workload.
- The values of the sortheap (sort heap) and sheapthres_shr (sort heap threshold for shared sorts) database configuration parameters are calculated specifically for an analytics workload. These settings take into account the additional memory requirements for processing column-organized data.
- The util_heap_sz (utility heap size) database configuration parameter is set to a value that takes into account the additional memory that is required to load the data into column-organized tables.
- The auto_reorg (automatic reorganization) database configuration parameter is set to ON.

Tip: Running the AUTOCONFIGURE command against an existing database when DB2_WORKLOAD is set to ANALYTICS has the same result.

The following extra choices are made automatically:

- The default database page size for a newly created database is set to 32 KB.
- A larger database shared sort heap is allocated.
- Intraquery parallelism is enabled for any workload (including SYSDEFAULTUSERWORKLOAD) that inherits the intrapartition parallelism setting from the intra_parallel (enable intrapartition parallelism) database manager configuration parameter, even if intra_parallel is set to NO.
- Concurrency control is enabled on SYSDEFAULTMANAGEDSUBCLASS.
- Automatic table maintenance performs space reclamation for column-organized tables by default.

Migrating Data to IBM DB2 with BLU Acceleration (2)



- **BACKUP YOUR DATABASE!!!**
- Set the workload to ANALYTICS and autoconfigure
 - **db2set DB2_WORKLOAD=ANALYTICS**
 - **db2stop force**
 - **db2start**
 - **db2 CONNECT TO <dbalias>**
 - **db2 AUTOCONFIGURE APPLY DB AND DBM**
or display changes and apply yourself with:
db2 AUTOCONFIGURE APPLY NONE
- If not already using create Automatic Storage:
 - **db2 CREATE STOGROUP <storage group name> ON <path(s)> SET AS DEFAULT**
- Convert DMS data and index table spaces to AUTOMATIC STORAGE.
 - **db2 ALTER TABLESPACE <tablespace name> MANAGED BY AUTOMATIC STORAGE**
 - **db2 ALTER TABLESPACE <tablespace name> REBALANCE**
 - **db2 CREATE TABLESPACE <tablespacename> PAGESIZE 32K EXTENTSIZE 4**
- Convert tables from row to column organization:
 - **db2convert -d mydatabase -z user -t table1 -ts coltbsp**
or convert all tables
db2convert -d mydatabase

Insight2015 Lead in the Insight Economy 73

The above steps can be used to perform the data migration task to BLU acceleration.

The db2convert utility converts one or all row-organized user tables in a specified database into column-organized tables. The row-organized tables remain online during command processing. For monitoring purposes, the command displays statistics about the conversion. You must have SQLADM or DBADM authority to invoke the ADMIN_MOVE_TABLE stored procedure, on which the db2convert command depends. You must also have the appropriate object creation authorities, including the authority to issue the SELECT statement on the source table.

Convert the tables from row organization to column organization using the db2convert utility:

db2convert -d mydatabase -z user -t table1
or to convert all tables
db2convert -d mydatabase

Also if creating a new database, see *5 Steps for Migrating Data to IBM DB2 with BLU Acceleration* by Mohankumar Saraswatipura at: http://ibmdatamag.com/2013/06/5-steps-for-migrating-data-to-ibm-db2-with-blu-acceleration/?utm_content=buffer430d1&utm_source=buffer&utm_medium=facebook&utm_campaign=Buffer

After setting DB2_WORKLOAD=ANALYTICS, run the AUTOCONFIGURE command to get most of the settings configured. During table space creation, explicitly specify 32K page size and 4 extent size since the database level settings are likely different. To avoid having to specify the extent size on every table space creation, you could update DFT_EXTENT_SZ to 4 instead.

A reasonable starting point is to set SHEAPTHRES_SHR to the size of the buffer pool and SORTHEAP to 1/20th of SHEAPTHRES_SHR.

Workload Table Organization Advisor (OQWT)



ibm.com

- **OQWT (Optim Query Workload Tuner)** can identify tables that are candidates for being column-organized tables.
 - Bases advice on table statistics and query workload characteristics.
 - Indicates expected performance improvement for queries in workload.
 - Can decide if want to execute script generated by tool to implement recommendations.
 - Make sure to back up database prior to making these changes.
- Should be able to measure compression savings and start running queries without any change to SQL.
 - No need for indexes or MQTs to be created.
 - Only indexes for column-organized tables are ones used for enforced primary keys or unique constraints.
 - If analytic data is already cleansed during ETL processing, create non-enforced constraints to allow DB2 Optimizer to leverage this information.



Insight2015

Lead in the Insight Economy

76

For additional information about OQWT see:

<http://pic.dhe.ibm.com/infocenter/dstudio/v4r1/index.jsp?topic=%2Fcom.ibm.datatools.qrytune.workloadtunedb2luw.doc%2Ftopics%2Fgenrecswtoa.html>

IBM InfoSphere Optim Query Workload Tuner for DB2 for LUW estimates benefits for column-organized tables



Advisor identifies candidate tables for conversion to columnar format.

Analyzes SQL workload and estimates execution cost on row- and column-organized tables.

Table	Creator	Current Organization	Recommended Organization	Conversion Warning	Finality	References to Table	Cumulative Total Cost
HOUSEHOLD_DEMOG...	TPCDS	ROW	COLUMN	Indexes will be removed	300.00	2	846,761,862.00
DATE_DIM	TPCDS	ROW	COLUMN	Indexes will be removed	349.00	51	15,408,311,462.00
WEB_SALES	TPCDS	ROW	COLUMN	Indexes will be removed	388.00	15	13,587,838,304.00
STORE	TPCDS	ROW	COLUMN	Indexes will be removed	302.00	1	973,711,296.00
STORE_SALES	TPCDS	ROW	COLUMN	Indexes will be removed	388.00	19	14,261,145,494.00
CUSTOMER_ADDRESS	TPCDS	ROW	COLUMN	Indexes will be removed	300.00	3	1,022,691,606.00
STORE_RETURNS	TPCDS	ROW	COLUMN	Indexes will be removed	324.00	1	973,711,296.00
ITEM	TPCDS	ROW	COLUMN	Indexes will be removed	736.00	25	11,926,096,848.00
CUSTOMER	TPCDS	ROW	COLUMN	Indexes will be removed	316.00	4	2,635,452,752.00

Execution Count	Weight	Estimated Performance Gain(%)	Cost Before	Cost After	Statement Text
1	0.00	-202.20	22,203,01...	67,097.3...	SELECT Q5.'SS_ADDR_SK' AS Q5C6, Q5.'SS_CDEMO_SK' AS Q5C4, Q5.'SS_CUSTOMER_SK' AS Q5C3, Q5.'SS...
1	0.00	78.99	438,398,0...	92,125.0...	select iss_i_brand_id as brand_id ,iss_i_class_id class_id ,iss_i_category_id category_id from tpcds.store_sal...
1	0.00	81.84	824,558,8...	149,764...	select avg(ss_quantity), avg(ss_ext_sales_price), avg(ss_ext_wholesale_cost), sum(ss_ext_wholesale_cost) from ...
1	0.00	77.40	2,085,428...	471,247...	with cross_items as (select i_item_sk ss_item_sk from tpcds.item, (select iss_i_brand_id brand_id ,iss_i_cles...
1	0.00	48.12	175,929.7...	91,275.6...	select cd_gender, cd_marital_status, cd_education_status, count(*) , cd_purchase_estimate, count(*) , c...

IBM® InfoSphere® Optim™ Query Workload Tuner Version 4.1.1 includes the Workload Table Organization Advisor for DB2 LUW V10.5 Fixpack 4, which examines all of the tables that are referenced by the statements in a query workload. Its recommendations lead to the best estimated performance improvement for the query workload as a whole. The advisor presents its analysis and rationales so that you can see the tables that are recommended for conversion from row to column organization.

IBM® InfoSphere® Optim™ Query Workload Tuner Version 4.1.0 included the Workload Table Organization Advisor for DB2 LUW V10.5 Fixpack 1, 2 and 3.

Sample recommendations for workload using Optim Query Workload Tuner



Review Workload Advisor Recommendations

This page shows the recommendations from the advisors that you ran.

Database connection: TPCSDANv10.2hotel67 (DB2 for Linux, UNIX, and Windows V10.5.0)

► Status/Description

Statements | Summary | **Table organization** | Candidate Table Organization

Estimated performance improvement: 83.44 %
Number of tables referenced in the workload: 11 Number of tables recommended for conversion: 11

Show DDL Script Test Candidate Table Organization Filter by Tables to be converted

Table	Creator	Current Organization	Recommended Organization	Conversion Warning
HOUSEHOLD_DEMOG...	TPCDS	ROW	COLUMN	Indexes will be remove
DATE_DIM	TPCDS	ROW	COLUMN	Indexes will be remove
WEB_SALES	TPCDS	ROW	COLUMN	Indexes will be remove
STORE	TPCDS	ROW	COLUMN	Indexes will be remove
STORE SALES	TPCDS	ROW	COLUMN	Indexes will be remove



Insight2015

Lead in the **Insight Economy**

78

For additional information about OQWT see:

<http://pic.dhe.ibm.com/infocenter/dstudio/v4r1/index.jsp?topic=%2Fcom.ibm.datatools.qrytune.workloadtunedb2luw.doc%2Ftopics%2Fgenrecswtoa.htm>

Database Configuration Parameter Changes



#Dominsight

- **SOFTMAX** db cfg is **deprecated**. Replaced with **page_age_trgt_mcr** and **page_age_trgt_gcr** which control saving changed buffer pool pages.
 - **page_age_trgt_mcr** - Page age target member crash recovery
 - Specifies target duration (in seconds) for changed pages kept in local buffer pool
 - Default 120 for pureScale, 240 for ESE and partitioned databases
 - **page_age_trgt_gcr** - Page age target group crash recovery
 - Specifies target duration (in seconds) for changed pages kept in group buffer pool before pages are persisted to disk
 - PureScale only, default is 240
 - `page_age_trgt_gcr` parameter must be \geq `page_age_trgt_mcr`
 - Existing upgraded databases continue to use softmax parameter.
 - Switch from using softmax to these new parameters by setting softmax to 0.
 - New DB2 10.5 databases created with value of softmax set to 0 by default.
- **buffpage** db cfg is **deprecated**. Use SIZE parameter on CREATE/ALTER BUFFERPOOL for explicit value or AUTOMATIC with self tuning memory.
- **auto_stats_prof** and **auto_prof_upd** db cfgs are **discontinued** because automatic statistics profiling is discontinued.



Insight2015

Lead in the **Insight Economy**

79

The `page_age_trgt_mcr` parameter is used when the softmax parameter is configured to a value of 0. Migrated databases retain the previous value for softmax and ignore the `page_age_trgt_mcr` parameter if this value is not 0. The value of softmax is set to 0 in new databases. Increasing the value of this configuration parameter keeps changed pages in memory for a longer time, allowing more page updates to be buffered before the pages are persisted to disk. This behavior can help to improve performance but also increases recovery time. The `page_age_trgt_gcr` parameter must be greater than or equal to `page_age_trgt_mcr` database configuration parameter. The `page_age_trgt_gcr` parameter is used when the softmax parameter is configured to a value of 0. Migrated databases retain the previous value for softmax and ignore the `page_age_trgt_gcr` parameter if this value is not 0. The value of softmax is set to 0 in new databases. Increasing the value of this configuration parameter keeps changed pages in memory for a longer time, allowing more page updates to be buffered before the pages are persisted to disk. This behavior can help to improve performance but also increases recovery time.

The **buffpage** is deprecated. This parameter indicated the default buffer pool size for the ALTER BUFFERPOOL or CREATE BUFFERPOOL statements that indicated the value -1 as the size. Use the SIZE parameter for both statements to indicate an explicit value or use the AUTOMATIC value to enable self tuning of the buffer pool.

The `auto_stats_prof` is discontinued because automatic statistics profiling is discontinued. The `auto_prof_upd` is discontinued because automatic statistics profiling is discontinued. Automatic statistics profiling is discontinued because of its noticeable performance overhead and restrictions.

DB2 command changes for DB2 10.5 (1)



#DbmInsight

DB2 Command	Command Option
db2pd	<ul style="list-style-type: none">- apinfo displays more info about current/past UOW activities- edus displays agents processing column-organized data- locks showlocks displays table name (TableNm) and schema name (SchemaNm) of locks being held- tablespace displays if reclaimable space enabled (RSE)- transactions displays total number of application commits and rollbacks- ruStatus displays status of online fix pack update in DB2 pureScale instance- wlocks ub option displays TableNm, SchemaNm, and AppNode columns for application locks being waited on- extentmovement displays extent movement status of database- membersubsetstatus displays state of member subsets- subsetid identifies subsets



Insight2015

Lead in the **Insight Economy**

88

The changes in DB2 10.5 to CLP and system commands can impact your existing applications and scripts after you upgrade to DB2 10.5. Modify any existing scripts and applications to discontinue using deprecated or discontinued options and parameters.

db2pd

- The **-apinfo** parameter now displays more information about current and past activities of the unit of work (UOW).
- The **-edus** parameter now also displays the agents that are processing column-organized data.
- The **showlocks** parameter now displays the TableNm and SchemaNm columns, which indicates the table name and schema name of locks that are being held by applications. To display this information, you must use the showlocks parameter with the **-locks** parameter.
- The **-tablespace** parameter now displays the RSE column to indicate whether the reclaimable space feature is enabled.
- The **-transactions** parameter now displays the total number of application commits and the total number of application rollbacks.
- The new **-ruStatus** parameter displays the status of an online fix pack update in a DB2 pureScale instance.
- The new details **uboption** for the **-wlocks** parameter displays the TableNm, SchemaNm, and AppNode columns for application locks that are being waited on.
- The new **-extentmovement** parameter displays the extent movement status of a database.
- The new **-membersubsetstatus** parameter displays the state of member subsets.
- The new **-subsetid** parameter identifies subsets.

DB2 command changes for DB2 10.5 (3)



#DbInsight

DB2 Command	Changes
db2support	<ul style="list-style-type: none">-d now supports collection of information from multiple databases. Separate database names with comma.-alldatabases specifies that information about all databases in database directory is collected.-system_group collects system-related information about specified system group.-system_user collects system-related information about specified system user.-wlm collects information related to WLM issues as part of optimizer mode with collection level 0 and above.



Insight2015

Lead in the Insight Economy

90

db2support

The **-d** parameter now supports collection of information from multiple databases. To specify multiple databases, separate the database names with a comma.

The new **-alldatabases** parameter specifies that information about all databases in the database directory is collected.

The new **-system_group** parameter collects system-related information about the specified system group.

The new **-system_user** parameter collects system-related information about the specified system user.

The new **-wlm** parameter collects information that is related to WLM issues as part of the optimizer mode with collection level 0 and above.

LOAD command STATISTICS parameter changes



#Dunnisght

- **STATISTICS YES** is discontinued.
- **LOAD ... STATISTICS USE PROFILE**
 - Instructs load to collect statistics during load according to profile defined for table.
 - Default for column-organized tables during LOAD REPLACE.
 - Also during LOAD REPLACE against column-organized table, column compression dictionaries replaced by default.
 - To explicitly disable automatic statistics collection, specify STATISTICS NO parameter.
 - Profile must be created and stored in SYSCAT.TABLES.STATISTICS_PROFILE before issuing LOAD command.
**RUNSTATS ON TABLE mel.t1
WITH DISTRIBUTION AND DETAILED INDEXES ALL
SET PROFILE [ONLY]**
 - For row-organized tables, if profile does not exist and STATISTICS USE PROFILE specified, warning returned and no statistics collected.
 - For column-organized tables, if profile does not exist and STATISTICS USE PROFILE specified, load utility uses same default RUNSTATS options as those used during automatic runstats.
 - During load, distribution statistics are not collected for XML column type.
- **LOAD ... STATISTICS NO**
 - Specifies that no statistics are to be collected and statistics in catalogs are not to be altered.
 - Default for row-organized tables.



Insight2015 Lead in the Insight Economy

93

The **STATISTICS YES** parameter of the LOAD command is discontinued.

The **STATISTICS USE PROFILE** instructs load to collect statistics during the load according to the profile defined for this table. This is the default for column-organized tables. The profile must be created before issuing the LOAD command. The profile is created by the RUNSTATS command. For row-organized tables, if the profile does not exist and this parameter is specified, a warning is returned and no statistics are collected. For column-organized tables, if the profile does not exist and this parameter is specified, the load utility uses the same default RUNSTATS command options as those that are used during an automatic runstats operation. During load, distribution statistics are not collected for columns of type XML.

The **STATISTICS NO** specifies that no statistics are to be collected, and that the statistics in the catalogs are not to be altered. This option is the default for row-organized tables. **LOAD** For column-organized tables, automatic statistics collection occurs by default during the LOAD REPLACE command. To explicitly disable automatic statistics collection, specify the STATISTICS NO parameter. Also, during a LOAD REPLACE operation against a column-organized table, the column compression dictionaries are replaced by default.

Online Table Reorganization Enhancements



#Dbinsight

- Reorganization capabilities enhanced in following ways:
 - INPLACE (online) table reorganization now supported for tables that use adaptive compression.
 - CLEANUP OVERFLOWS suboption of INPLACE reorg removes pointer overflow records but bypasses other row movement processing
 - Table size will not be reduced
 - Overflow records may be introduced in compressed tables when rows are updated
- CLASSIC can be specified to indicate offline table reorg processing mode

```
db2 REORG TABLE T1 INPLACE CLEANUP OVERFLOWS
vs
db2 REORG TABLE T1 INPLACE FULL
```

```
db2 REORG TABLE T1 CLASSIC ALLOW READ ACCESS USE TEMP1
```



Insight2015

Lead in the **Insight Economy**

94

CLEANUP OVERFLOWS: An INPLACE CLEANUP OVERFLOWS reorganization traverses the table and searches for pointer or overflow records. Any record found is converted to a normal record by the operation. This operation improves performance for tables that have a significant number of pointer or overflow records. The operation does not result in a reduction of the size of the table.

FULL: The table is reorganized to fill pages while respecting PCTFREE for the table. Optionally, when the INDEX clause is specified, row data is moved within the table to recluster the data. Overflow records are also converted to normal records as part of this process. This behavior is the default behavior.

Additional INPLACE options

```
.-ALLOW WRITE ACCESS-.
|-----+-----|
>
'-ALLOW READ ACCESS--'

.-FULL-.                                .-TRUNCATE TABLE---.  .-
START--.
>-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+---|
|               '-INDEX--index-name-' '-NOTRUNCATE TABLE-' | '-
RESUME-'
'-CLEANUP OVERFLOWS-----'
```

NOT ENFORCED primary key & unique constraints support



#Dominsight

- Informational constraints help query optimizer to select optimal data access plans when index access to the data provides no additional benefit.
- If know that data conforms to these constraints, can use NOT ENFORCED capability to help achieve two goals:
 1. Improve performance, primarily in insert, update, and delete operations on table
 2. Reduce space requirements associated with enforcing primary key or unique constraint
- Cannot reference not-enforced primary key constraints in any enforced referential integrity definitions.
- **Prior DB2 versions supported not-enforced foreign keys and not-enforced check constraints.**



Insight2015

Lead in the **Insight Economy**

95

You can create a not-enforced primary key or unique constraint for either a column-organized or row-organized table. Unlike an enforced primary key or unique constraint, a not-enforced primary key or unique constraint does not create an index on the data.

Specify an informational constraint only if the table data is independently known to conform to the constraint. Because the DB2 database manager does not enforce uniqueness for these constraints, if the table data violates the not-enforced constraint, incorrect results can occur. You cannot reference not-enforced primary key constraints in any enforced referential integrity definitions.

Expression-based indexes improve query performance



#Dominsight

- CREATE INDEX statement extension to create index that contains expression-based key.
- Can create index on data that is not stored in table.
- When create index with expression-based key, results of expression are stored in index.
- SQL statement that contains predicate that includes expression can benefit from index whose definition includes expression-based keys. .
- Expression in index useful if query references expression in other ways:
 - query might return that expression in its results
 - query might order result set by that expression
- If index created with UNIQUE option, uniqueness enforced against values stored in index.
 - Uniqueness not enforced against original column values
- Fix Pack 1: LOB types (BLOB, CLOB, and DBCLOB) supported as input to expression-based index keys.
- Create index with expression-based key using upper case of employee's name and ID:

```
CREATE INDEX EMP_UPPERNAME ON EMPLOYEE (UPPER(NAME) , ID)
```



Insight2015

Lead in the Insight Economy

96

Starting with DB2 10.5, you can use an extension to the CREATE INDEX statement to create an index that contains an expression-based key. Using this expression-based key capability, you can create an index on data that is not stored in a table. When you create an index with an expression-based key, the results of the expression are stored in the index. An SQL statement that contains a predicate that includes an expression can benefit from an index whose definition includes expression-based keys. With this feature, the database manager can choose from many highly efficient index-based methods to access the data and evaluate the predicate. As a result, performance is improved for queries that contain expressions. An expression in an index is also useful if the query references the expression in other ways. For example, the query might return that expression in its results, or it might order the result set by that expression. To benefit from efficient evaluation of expressions that are specified at run time, use an expression-based key, rather than other types of indexes whose keys only consist of one or more table columns.

Expression-based indexes are best suited when you want an efficient evaluation of queries that involve a column expression. Simple index keys consist of a concatenation of one or more specified table columns. Compared to simple indexes, the index key values of expression-based indexes are not the same as values in the table columns. The values are transformed by the expressions that you specify. You can create the index with the CREATE INDEX statement. If an index is created with the UNIQUE option, the uniqueness is enforced against the values that are stored in the index. The uniqueness is not enforced against the original column values.

EXCLUDE NULL KEYS from Indexes



ibm.com

- **EXCLUDE NULL KEYS** on **CREATE INDEX** reduces size of sparse indexes.
 - does not insert key into index object when all columns in key are null.
 - Can result in improved storage and performance optimization when do not want queries to access data associated with null keys.
 - For unique indexes, enforcement of uniqueness of table data ignores rows where the index key is null.
- **INCLUDE NULL KEYS** specifies that index entry is created when all parts of index key contain null value. Default setting.
- Cannot specify **EXCLUDE NULL KEYS** with following:
 - Nickname
 - **GENERATE KEY USING XMLPATTERN** clause
 - **EXTEND USING** clause.



Insight2015

Lead in the Insight Economy

97

Starting with DB2 10.5 the new EXCLUDE NULL KEYS clause can be used on the CREATE INDEX statement to reduce the size of sparse indexes.

An index created with the EXCLUDE NULL KEYS clause does not insert a key into the index object when all the columns in the key are null. Excluding null keys can result in improved storage and performance optimization for cases where you do not want queries to access data associated with null keys. For unique indexes, the enforcement of the uniqueness of table data ignores rows where the index key is null.

Unique indexes are indexes that help maintain data integrity by ensuring that no two rows of data in a table have identical key values.

When you create a unique index for an existing table with data, values in the columns or expressions that comprise the index key are checked for uniqueness. If the table contains rows with duplicate key values, the index creation process fails. When a unique index is defined for a table, uniqueness is enforced whenever keys are added or changed within the index. This enforcement includes insert, update, load, import, and set integrity, to name a few. In addition to enforcing the uniqueness of data values, a unique index can also be used to improve data retrieval performance during query processing.

Non-unique indexes are not used to enforce constraints on the tables with which they are associated. Instead, non-unique indexes are used solely to improve query performance by maintaining a sorted order of data values that are used frequently.

It is important to understand that there is no significant difference between a primary key or unique key constraint and a unique index. To implement the concept of primary and unique key constraints, the database manager uses a combination of a unique index and the NOT NULL constraint. Therefore, unique indexes do not enforce primary key constraints by themselves because they allow null values. Although null values represent unknown values, when it comes to indexing, a null value is treated as being equal to other null values.

Therefore, if a unique index consists of a single column, only one null value is allowed-more than one null value would violate the unique constraint. Similarly, if a unique index consists of multiple columns, a specific combination of values and nulls can be used only one time.

INCLUDE NULL KEYS - Specifies that an index entry is created when all parts of the index key contain the null value. This is the default setting.

EXCLUDE NULL KEYS - Specifies that an index entry is not created when all parts of the index key contain the null value. When any part of the index key is not a null value, an index entry is created. You cannot specify EXCLUDE NULL KEYS with the following syntax elements:

- A nickname
- The **GENERATE KEY USING XMLPATTERN** clause
- The **EXTEND USING** clause.

Large row size support – Oracle Compatibility



#Dominsight

- Row size support extended allowing to create table where row length can exceed maximum record length for page size of table space
- **EXTENDED_ROW_SZ=ENABLE** db cfg parameter must be set
 - Migrate tables that have row sizes exceeding 32K
 - Improve performance where majority of data rows can fit on smaller page
 - Create tables with more VARCHAR or VARGRAPHIC columns
 - Maximum number of columns is not changing
 - Ability to exceed maximum record length for page size allows for more columns
- After upgrade, alter existing tables to take advantage of extended row size support.
CREATE TABLE T1 (C1 INT, C2 VARCHAR(3995)) in TS1
 - Table space TS1 has 4K page size and T1 created without need to store any data out of row because its byte count is 4005.
ALTER TABLE T1 ADD C3 CLOB(1M)
 - Byte count for T1 now exceeds maximum record length, but ALTER TABLE successful if **EXTENDED_ROW_SZ=ENABLE**.
- When row data exceeds maximum row length, moves subset of VARCHAR or VARGRAPHIC data out of row and stores as LOB data on different data page(s).
 - **data access might be slower because additional I/O resources might be required to fetch, insert, or update the LOB data**



Insight2015

Lead in the **Insight Economy**

101

Starting with DB2 10.5 row size support is extended allowing you to create a table where its row length can exceed the maximum record length for the page size of the table space.

Rows of table data are organized into blocks called pages which can be four sizes: 4, 8, 16, and 32 kilobytes. All tables created within a table space of a particular size have a matching page size. In previous releases, the maximum number of bytes allowed in a table row was dependant on the page size of the table space. Any attempt to create a table whose row length exceeded the maximum record length for the page size would result in an error (SQLSTATE 54010). For example, in previous releases the following table could not be created in a 4K page size table space because of its large row size. **CREATE TABLE T1 (C1 INTEGER, C2 VARCHAR(5000))** By extending row size support to allow for the creation of tables containing large rows that exceed the maximum record length for the page size of the table space, the table T1 can be created in a 4K page size table space.

When extended_row_sz is set to ENABLE, extended row size is enabled. During an INSERT or DELETE table operation, row data that exceeds the maximum row length moves a subset of VARCHAR or VARGRAPHIC data out of the row and stores it as LOB data. While the data is in LOB format, data access might be slower because additional I/O resources might be required to fetch, insert, or update the LOB data.

For further details see: http://www-01.ibm.com/support/knowledgecenter/#/SSEPGG_10.5.0/com.ibm.db2.luw.admin.dboj.doc/doc/c0061057.html

New Database and Instance Monitoring Functions



#Dominsight

• **MON_GET_DATABASE, MON_GET_DATABASE_DETAILS**

- Returns database statistics previously only available using DATABASE snapshot monitoring
- Determine activation time, activation state, total connections, and connection high water mark for current database:

```
SELECT DB_CONN_TIME, DB_ACTIVATION_STATE, TOTAL_CONS,  
CONNECTIONS_TOP FROM TABLE (MON_GET_DATABASE(-2));
```

• **MON_GET_INSTANCE**

- Returns DB2 instance level information previously only available using DATABASE MANAGER snapshot monitoring

- Determine start time and total number of current connections:

```
SELECT DB2START_TIME, TOTAL_CONNECTIONS  
FROM TABLE (MON_GET_INSTANCE(-2));
```



Insight2015

Lead in the **Insight Economy**

102

The MON_GET_DATABASE table function returns database level information within the monitor infrastructure.

The MON_GET_DATABASE_DETAILS table function retrieves database information metrics in an XML document. The schema for the XML document that is returned in the DETAILS column is available in the file sqllib/misc/DB2MonRoutines.xsd. Further details can be found in the file sqllib/misc/DB2MonCommon.xsd.

The MON_GET_INSTANCE table function returns DB2 instance level information within the monitor infrastructure.

Active utility monitoring - MON_GET_UTILITY



#DbmInsight

```
SELECT UTILITY_ID, UTILITY_START_TIME, UTILITY_TYPE,  
  VARCHAR(OBJECT_SCHEMA,12) AS SCHEMA,  
  VARCHAR(OBJECT_NAME,20) AS OBJECT_NAME,  
  SUBSTR(UTILITY_DETAIL,1,200) AS CMD  
FROM TABLE(MON_GET_UTILITY(-1)) AS T1
```

```
UTILITY_ID  UTILITY_START_TIME          UTILITY_TYPE  SCHEMA  OBJECT_NAME  
-----  
          2 2013-06-25-14.02.22.325849  LOAD                TEST2    ACCT  
CMD
```

```
-----  
[LOADID: 21.2013-06-25-14.02.22.266904.0 (6;8)] [*LOCAL.inst20.130625175129]  
OFFLINE LOAD DEL  
AUTOMATIC INDEXING REPLACE COPY NO TEST2    .ACCT
```

```
1 record(s) selected.
```

- SNAPUTIL administrative view and SNAP_GET_UTIL table function deprecated.
- Alternative to db2 list utilities show detail and db2pd -util



Insight2015

Lead in the Insight Economy

105

The MON_GET_UTILITY table function returns information about utilities currently running on the database.

One of the following authorizations is required:

- EXECUTE privilege on the routine
- DATAACCESS authority
- DBADM authority
- SQLADM authority

Important: The SNAPUTIL administrative view and the SNAP_GET_UTIL table function are deprecated and have been replaced by the MON_GET_UTILITY table function.

DISCONTINUED SQL administrative routines and views (1)

UPGRADE DATABASE command drops discontinued administrative routines and marks dependent objects inoperative or invalid.

DISCONTINUED in DB2 10.5	REPLACEMENT ROUTINE/VIEW
SNAPSHOT_AGENT	MON_GET_AGENT, MON_GET_CONNECTION
SNAPSHOT_APPL	MON_GET_CONNECTION
SNAPSHOT_APPL_INFO	MON_GET_CONNECTION
SNAPSHOT_BP	MON_GET_BUFFERPOOL
SNAPSHOT_CONTAINER	MON_GET_CONTAINER
SNAPSHOT_DATABASE	MON_GET_DATABASE MON_GET_TRANSACTION_LOG
SNAPSHOT_DBM	MON_GET_INSTANCE
SNAPSHOT_DYN_SQL	MON_GET_PKG_CACHE_STMT
SNAPSHOT_FCMNODE	MON_GET_FCM
SNAPSHOT_FILEW	SNAP_WRITE_FILE




Insight2015

Lead in the Insight Economy

113

A number of SQL administrative routines are discontinued in Version 10.5. You must use the replacement routines. The following shows the SQL administrative routines that are discontinued, their replacement routines or views, and the version when the replacement views or routines became available:

- Discontinued SNAPSHOT_AGENT table function replaced with MON_GET_AGENT table function and MON_GET_CONNECTION table function in Version 10.5 and 9.7.
- Discontinued SNAPSHOT_APPL table function replaced with MON_GET_CONNECTION table function and MON_GET_UNIT_OF_WORK table function in Version 9.7.
- Discontinued SNAPSHOT_APPL_INFO table function replaced with MON_GET_CONNECTION table function and MON_GET_UNIT_OF_WORK table function in Version 9.7.
- Discontinued SNAPSHOT_BP table function replaced with MON_GET_BUFFERPOOL table function in Version 9.7.
- Discontinued SNAPSHOT_CONTAINER table function replaced with MON_GET_CONTAINER table function in Version 9.7.
- Discontinued SNAPSHOT_DATABASE table function replaced with MON_GET_DATABASE table function and MON_GET_TRANSACTION_LOG table function in Version 10.5.
- Discontinued SNAPSHOT_DBM table function replaced with MON_GET_INSTANCE table function in Version 10.5.
- Discontinued SNAPSHOT_DYN_SQL table function replaced with MON_GET_PKG_CACHE_STMT table function in Version 9.7.
- Discontinued SNAPSHOT_FCM table function replaced with MON_GET_FCM table function in Version 9.7 Fix Pack 2.
- Discontinued SNAPSHOT_FCMNODE table function replaced with MON_GET_FCM_CONNECTION_LIST table function in Version 9.7 Fix Pack 2.
- Discontinued SNAPSHOT_FILEW procedure replaced with SNAP_WRITE_FILE procedure in Version 9.1.
- Discontinued SNAPSHOT_LOCK table function replaced with MON_GET_APPL_LOCKWAIT table function, MON_GET_LOCKS table function, MON_FORMAT_LOCK_NAME table function in Version 9.7 Fix Pack 1.
- Discontinued SNAPSHOT_LOCKWAIT table function replaced with MON_GET_APPL_LOCKWAIT table function, MON_GET_LOCKS table function, and MON_FORMAT_LOCK_NAME table function in Version 9.7 Fix Pack 1.
- Discontinued SNAPSHOT QUIESCERS table function replaced with SNAP_GET_TBSP QUIESCER table function in Version 9.1.
- Discontinued SNAPSHOT_RANGES table function replaced with SNAP_GET_TBSP_RANGE table function in Version 9.1.



DISCONTINUED SQL administrative routines and views (2)

DISCONTINUED in DB2 10.5	REPLACEMENT ROUTINE/VIEW
SNAPSHOT_LOCK, SNAPSHOT_LOCKWAIT	MON_GET_APPL_LOCKWAIT, MON_GET_LOCKS, MON_FORMAT_LOCK_NAME
SNAPSHOT QUIESCERS	SNAP_GET_TBSP QUIESCER
SNAPSHOT_STATEMENT	MON_GET_ACTIVITY, MON_CURRENT_SQL
SNAPSHOT_SUBSECT	SNAP_GET_SUBSECTION
SNAPSHOT_SWITCHES	SNAP_GET_SWITCHES
SNAPSHOT_TABLE	MON_GET_TABLE
SNAPSHOT_TABREORG	MON_GET_UTILITY
TAB_REORG	MON_GET_TABLESPACE
SNAPSHOT_UTIL_PROG	SNAPUTIL_PROGRESS, SNAP_GET_UTIL_PROGRESS

115

Insight2015 Lead in the Insight Economy

Modify all user-defined objects that are dependent on the discontinued routines. Re-create these objects using the replacement routines or views indicated. If you upgrade a database that has dependent objects, the UPGRADE DATABASE command drops the discontinued administrative routines and marks the dependent objects inoperative or invalid.

Modify all your applications and scripts and remove all references to these routines or use the replacement routines or views indicated.

- Discontinued SNAPSHOT_LOCK table function replaced with MON_GET_APPL_LOCKWAIT table function, MON_GET_LOCKS table function, MON_FORMAT_LOCK_NAME table function in Version 9.7 Fix Pack 1.
- Discontinued SNAPSHOT_LOCKWAIT table function replaced with MON_GET_APPL_LOCKWAIT table function, MON_GET_LOCKS table function, and MON_FORMAT_LOCK_NAME table function in Version 9.7 Fix Pack 1.
- Discontinued SNAPSHOT QUIESCERS table function replaced with SNAP_GET_TBSP QUIESCER table function in Version 9.1.
- Discontinued SNAPSHOT_RANGES table function replaced with SNAP_GET_TBSP_RANGE table function in Version 9.1.
- Discontinued SNAPSHOT_STATEMENT table function replaced with MON_GET_ACTIVITY table function and MON_CURRENT_SQL administrative view in Version 10.5 and Version 9.7 Fix Pack 1.
- Discontinued SNAPSHOT_SUBSECT table function replaced with SNAP_GET_SUBSECTION table function in Version 9.1.
- Discontinued SNAPSHOT_SWITCHES table function replaced with SNAP_GET_SWITCHES table function in Version 9.1.

Deprecated SQL administrative routines and views (1)

DEPRECATED in DB2 10.5	REPLACEMENT ROUTINE/VIEW
APPL_PERFORMANCE	MON_CONNECTION_SUMMARY
APPLICATIONS	MON_GET_CONNECTION
BP_HITRATIO, BP_READ_IO,BP_WRITE_IO	MON_BP_UTILIZATION
CONTAINER_UTILIZATION	MON_GET_CONTAINER
LOG_UTILIZATION	MON_TRANSACTION_LOG_UTILIZATION
SNAPAGENT, SNAP_GET_AGENT	MON_GET_CONNECTION, MON_GET_AGENT
SNAPAPPL, SNAP_GET_AGENT, SNAPAPPL, SNAP_GET_APPL	MON_GET_CONNECTION
SNAPBP, SNAP_GET_BP, SNAPBP_PART,SNAP_GET_BP_PART	MON_GET_BUFFERPOOL
SNAPCONTAINER,SNAP_GET_CONTA INER	MON_GET_CONTAINER
SNAPDB, SNAP_GET_DB	MON_GET_DATABASE
SNAPDBM, SNAP_GET_DBM	MON_GET_DATABASE



Insight2015

Lead in the Insight Economy

117

Expanded support is provided for the existing administrative routines and views by replacing them with new, more comprehensive routines or views. Starting with 10.1, routines with version-suffixed names were deprecated. The names of the routines do not have a version suffix so that the names remain consistent across releases. The replacement routines without a version-suffix might include modifications such as: columns added or removed, new data types for existing columns, or new values for existing columns. Use good practices when issuing queries on SQL administrative routines to minimize the impact from changes to these routines.

- Deprecated APPL_PERFORMANCE administrative view - Retrieve percentage of rows selected for an application replaced with MON_CONNECTION_SUMMARY - Retrieve metrics for all connections.
- Deprecated APPLICATIONS administrative view - Retrieve connected database application information replaced with MON_GET_CONNECTION table function - Get connection metrics.
- Deprecated BP_HITRATIO administrative view - Retrieve bufferpool hit ratio information replaced with MON_BP_UTILIZATION administrative view - Retrieve metrics for bufferpools.
- Deprecated BP_READ_IO administrative view - Retrieve bufferpool read performance information replaced with MON_BP_UTILIZATION administrative view - Retrieve metrics for bufferpools.
- Deprecated BP_WRITE_IO administrative view - Retrieve bufferpool write performance information replaced with MON_BP_UTILIZATION administrative view - Retrieve metrics for bufferpools.
- Deprecated CONTAINER_UTILIZATION administrative view - Retrieve table space container and utilization information replaced with MON_GET_CONTAINER table function - Get table space container metrics.
- Deprecated LOG_UTILIZATION administrative view - Retrieve log utilization information replaced with MON_TRANSACTION_LOG_UTILIZATION administrative view - Retrieve log utilization information.
- Deprecated LONG_RUNNING_SQL administrative view replaced with MON_CURRENT_SQL - Retrieve key metrics for all activities on all members.
- Deprecated SNAP_WRITE_FILE procedure – not available.
- Deprecated SNAPAGENT administrative view and SNAP_GET_AGENT table function - Retrieve agent logical data group application snapshot information replaced with MON_GET_CONNECTION table function - Get connection metrics and MON_GET_AGENT table function - List agents running in a service class.
- Deprecated SNAPAPPL administrative view and SNAP_GET_APPL table function - Retrieve appl logical data group snapshot information replaced with MON_GET_CONNECTION table function - Get connection metrics.
- Deprecated SNAPAPPL_INFO administrative view and SNAP_GET_APPL_INFO table function - Retrieve appl_info logical data group snapshot information replaced with MON_GET_CONNECTION table function - Get connection metrics.
- Deprecated SNAPBP administrative view and SNAP_GET_BP table function - Retrieve bufferpool logical group snapshot information replaced with MON_GET_BUFFERPOOL table function - Get buffer pool metrics.
- Deprecated SNAPBP_PART administrative view and SNAP_GET_BP_PART table function - Retrieve bufferpool_nodeinfo logical data group snapshot information replaced with MON_GET_BUFFERPOOL table function - Get buffer pool metrics.
- Deprecated SNAPCONTAINER administrative view and SNAP_GET_CONTAINER table function - Retrieve tablespace_container logical data group snapshot information replaced with MON_GET_CONTAINER table function - Get table space container metrics.
- Deprecated SNAPDB administrative view and SNAP_GET_DB table function - Retrieve snapshot information from the dbase logical group replaced with MON_GET_DATABASE table function - Get database level information.
- Deprecated SNAPDBM administrative view and SNAP_GET_DBM table function - Retrieve the dbm logical grouping snapshot information replaced with MON_GET_INSTANCE table function - Get instance level information.

Deprecated SQL administrative routines and views (2)

DEPRECATED in DB2 10.5	REPLACEMENT ROUTINE/VIEW
SNAPFCM, SNAP_GET_FCM	MON_GET_FCM
SNAPDETAILLOG, SNAP_GET_DETAILLOG	MON_GET_TRANSACTION_LOG
SNAPFCM_PART, SNAP_GET_FCM_PART	MON_GET_FCM
SNAPSTMT, SNAP_GET_TBSP	MON_GET_TABLESPACE
SNAPTbsp, SNAP_GET_TBSP	MON_GET_TABLESPACE
SNAPTbsp_PART, SNAP_GET_TBSP_PART	MON_GET_TABLESPACE
SNAPUTIL, SNAP_GET_UTIL	MON_GET_UTILITY
TBSP_UTILIZATION	MON_TBSP_UTILIZATION



Insight2015

Lead in the Insight Economy

118

- Deprecated SNAPFCM administrative view and SNAP_GET_FCM table function - Retrieve the fcm logical data group snapshot information replaced with MON_GET_FCM - Get FCM metrics.
- Deprecated SNAPDETAILLOG administrative view and SNAP_GET_DETAILLOG table function - Retrieve snapshot information from the detail_log logical data group replaced with MON_GET_TRANSACTION_LOG table function - Get log information.
- Deprecated SNAPFCM_PART administrative view and SNAP_GET_FCM_PART table function - Retrieve the fcm_node logical data group snapshot information replaced with MON_GET_FCM - Get FCM metrics.
- Deprecated SNAPSTMT administrative view and SNAP_GET_STMT table function - Retrieve statement snapshot information replaced with MON_CURRENT_SQL - Retrieve key metrics for all activities on all members and MON_GET_CONNECTION table function - Get connection metrics.
- Deprecated SNAPTbsp administrative view and SNAP_GET_TBSP table function - Retrieve table space logical data group snapshot information replaced with MON_GET_TABLESPACE table function - Get table space metrics.
- Deprecated SNAPTbsp_PART administrative view and SNAP_GET_TBSP_PART table function - Retrieve tablespace_nodeinfo logical data group snapshot information replaced with MON_GET_TABLESPACE table function - Get table space metrics.
- Deprecated SNAPUTIL administrative view and SNAP_GET_UTIL table function - Retrieve utility_info logical data group snapshot information replaced with MON_GET_UTILITY table function - Get utilities running on the database.
- Deprecated TBSP_UTILIZATION administrative view - Retrieve table space configuration and utilization information replaced with MON_TBSP_UTILIZATION - Retrieve monitoring metrics for all table spaces and all database partitions.

Administrative Views and Routines Changes in DB2 10.5.

- **New columns, changed descriptions, changed column data types, and increased column lengths**

- ADMINTABINFO and ADMIN_GET_TAB_INFO
- ENV_SYS_INFO
- MON_BP_UTILIZATION
- MON_CONNECTION_SUMMARY
- MON_CURRENT_SQL
- MON_DB_SUMMARY
- MON_FORMAT_XML_COMPONENT_TIMES_BY_ROW
- MON_FORMAT_XML_METRICS_BY_ROW
- MON_FORMAT_XML_TIMES_BY_ROW
- MON_GET_APPL_LOCKWAIT
- MON_GET_BUFFERPOOL
- MON_GET_CONNECTION
- MON_GET_CONNECTION_DETAILS
- MON_GET_PKG_CACHE_STMT
- MON_GET_SERVICE_SUBCLASS
- MON_GET_SERVERLIST
- MON_GET_TABLE
- MON_GET_TABLESPACE
- MON_GET_TABLE_USAGE_LIST
- MON_TBSP_UTILIZATION
- MON_GET_UNIT_OF_WORK
- MON_GET_UNIT_OF_WORK_DETAILS
- MON_GET_WORKLOAD
- MON_GET_WORKLOAD_DETAILS
- MON_WORKLOAD_SUMMARY

For complete list of changes and deprecated table functions and views in DB2 9.7, 10.1 and 10.5 see URLs in notes.


Insight2015 Lead in the Insight Economy

119

The above administrative views and routines have changed in DB2 10.5. Most modifications consist of new columns, new values, changed column data types, and increased column lengths.

Review the list of the deprecated administrative routines and their replacement routines or views to determine additional changes that might impact your applications and scripts.

http://www-01.ibm.com/support/knowledgecenter/#!/SSEPGG_10.5.0/com.ibm.db2.luw.sql.rtn.doc/doc/r0023171.html?cp=SSEPGG_10.5.0%2F3-6-1-4




New Routines and Views in DB2 10.5

- **New monitor routines and views:**
 - MON_CAPTURE_ACTIVITY_IN_PROGRESS procedure
 - MON_COLLECT_STATS procedure
 - MON_GET_ACTIVITY table function
 - MON_GET_AGENT table function
 - MON_GET_DATABASE table function
 - MON_GET_DATABASE_DETAILS table function
 - MON_GET_INSTANCE table function
 - MON_GET_QUEUE_STATS table function
 - MON_GET_SERVICE_SUBCLASS_STATS table function
 - MON_GET_SERVICE_SUPERCLASS_STATS table function
 - MON_GET_UTILITY table function
 - MON_GET_WORKLOAD_STATS table function
 - MON_GET_WORK_ACTION_SET_STATS table function
 - MON_TRANSACTION_LOG_UTILIZATION administrative view

- **New environment routines:**
 - ENV_GET_DB2_EDU_SYSTEM_RESOURCES table function
 - ENV_GET_INSTANCE_CODE_LEVELS table function

- **New workload management routines:**
 - WLM_ALTER_MEMBER_SUBSET procedure
 - WLM_CREATE_MEMBER_SUBSET procedure
 - WLM_DROP_MEMBER_SUBSET procedure

#ibminights



Lead in the Insight Economy

120

Adjust your applications to the changes in existing system catalog objects such as new columns, or modified data type in columns. Also, start using new system catalog objects that provide more comprehensive views or routines, or access to information about new functionality.

To minimize the impact of changes to built-in routines and views, review Best practices for calling built-in routines and views in applications.

http://www-01.ibm.com/support/knowledgecenter/#/SSEPGG_10.5.0/com.ibm.db2.luw.sql.rtn.doc/doc/c0056435.html?cp=SSEPGG_10.5.0%2F3-6-1-0

For a list of the data dictionary-compatible views, see Data dictionary-compatible views at::

http://www-01.ibm.com/support/knowledgecenter/#/SSEPGG_10.5.0/com.ibm.db2.luw.wn.doc/doc/i0058717.html

System Catalog Views – New and Changed



ibm/insight

- **New System Catalog Views**
 - SYSCAT.MEMBERSUBSETATTRS
 - SYSCAT.MEMBERSUBSETMEMBERS
 - SYSCAT.MEMBERSUBSETS
- **New columns, changed descriptions, changed column data types, and increased column lengths**
 - SYSCAT.ATTRIBUTES catalog view
 - SYSCAT.CHECKS catalog view
 - SYSCAT.COLUMNS catalog view
 - SYSCAT.CONTROLS catalog view
 - SYSCAT.DATATYPES catalog view
 - SYSCAT.INDEXCOLUSE catalog view
 - SYSCAT.INDEXES catalog view
 - SYSCAT.PACKAGES catalog view
 - SYSCAT.ROUTINEPARMS catalog view
 - SYSCAT.ROUTINES
 - SYSCAT.ROWFIELDS
 - SYSCAT.SERVICECLASSES
 - SYSCAT.STOGROUPS
 - SYSCAT.TABDEP
 - SYSCAT.TABLES-new column called TABLEORG indicates table organization.
 - SYSCAT.TABLESPACES
 - SYSCAT.TRIGGERS
 - SYSCAT.VARIABLES
 - SYSCAT.VIEWS
 - SYSSTAT.COLUMNS
 - SYSSTAT.INDEXES
 - SYSSTAT.TABLES



Insight2015

Lead in the Insight Economy

121

To support new features in DB2 Version 10.5, database catalog objects such as system catalog views, built-in functions and global variables, built-in administrative routines and views have been added and modified.

Most modifications to catalog views consist of new columns, changed descriptions, changed column data types, and increased column lengths.

HADR Changes

- As DB2 10.5, **including environments with one standby database**, always set **HADR_TARGET_LIST** as part of initializing HADR databases. Initialization of HADR without setting HADR_TARGET_LIST is deprecated.
- HADR log spooling now enabled by default. **HADR_SPOOL_LIMIT** has default of AUTOMATIC for new databases. For upgraded databases, value of `hadr_spool_limit` is not changed.
- HADR now supported in DB2 pureScale.
 - In DB2 pureScale environments, must set HADR_TARGET_LIST.

For details see:

http://www-01.ibm.com/support/knowledgecenter/SSEPGG_10.5.0/com.ibm.db2.luw.admin.ha.doc/doc/t0011725.html?lang=en



Insight2015

Lead in the Insight Economy

131

DB2 pureScale environments now support DB2 high availability disaster recovery (HADR). This enhancement combines the continuous availability of the DB2 pureScale Feature with the robust disaster recovery capabilities of HADR.

HADR log spooling is now enabled by default. The `hadr_spool_limit` database configuration parameter, which specifies log spooling behavior, now has a default setting of AUTOMATIC. For DB2 pureScale databases, the `hadr_spool_limit` parameter is set to AUTOMATIC. Because HADR was previously not supported with the DB2 pureScale Feature, this change has no impact.

For all other DB2 databases, the value of the `hadr_spool_limit` parameter is not changed. Introduced in 10.1, HADR log spooling allows transactions on the primary to make progress without having to wait for the log replay on the standby. Testing has shown that spooling can significantly reduce any impact that HADR might have on the primary's workload but has a negligible impact on a standby's replay performance.

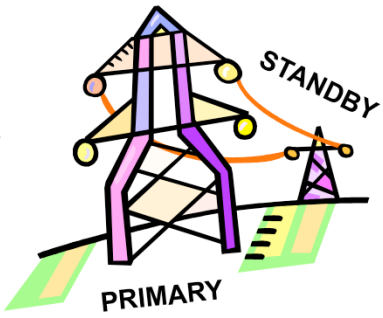
As of Version 10.5, the steps for initializing your HADR databases have changed. You should now use the `hadr_target_list` configuration parameter even if you are only configuring one standby. The `hadr_target_list` database configuration parameter was introduced in Version 10.1 to configure multiple HADR standby databases. You could use this parameter to specify up to three standby databases. This parameter was not required for HADR setups that had only one standby database.

As of Version 10.5, in DB2 environments that do not use the DB2 pureScale Feature, you should always set the `hadr_target_list` database configuration parameter as part of initializing your HADR databases, including environments with one standby database. Initialization of HADR without setting the `hadr_target_list` database configuration parameter is deprecated and might be discontinued in a future release. In DB2 pureScale environments, you must set the `hadr_target_list` database configuration parameter to initialize HADR.

If you are creating a new HADR setup, perform all the steps in the initializing HADR task, including setting the `hadr_target_list` database configuration parameter. If your HADR is already set up using the old method (that is, the `hadr_target_list` parameter is empty), see:
http://www-01.ibm.com/support/knowledgecenter/#!/SSEPGG_10.1.0/com.ibm.db2.luw.admin.ha.doc/doc/t0060257.html


Upgrading DB2 servers in automated HADR environment

1. As instance owner, stop HADR on both standby and primary servers:
db2 stop hadr on database sample
2. Verify database role has changed to standard on both servers using:
db2 get db cfg for sample | grep -i role
3. Ensure snapshot of resource model captured via **Issam** command.
Optionally, capture resource and policy info with **Issamctrl** and **sampolicy**.
4. Delete cluster domain using **db2haicu** command:
db2haicu -delete
5. Deactivate database on both standby and primary servers:
db2 deactivate db sample
6. Stop DB2 instance: **db2stop**
7. As root, install DB2 V10.x in new directory path.
8. As root, upgrade DB2 instance: **DB2DIR/instance/db2iupgrade db2inst1**
Where DB2DIR represents installation location of new release of DB2 database system installed.
9. As instance owner, start DB2 instance: **db2start**
10. As instance owner, upgrade database:
db2 upgrade database sample
11. Reinitialize HADR by following steps documented on this link:
http://www-01.ibm.com/support/knowledgecenter/SSEPGG_10.1.0/com.ibm.db2.luw.admin.ha.doc/doc/t0011725.html
12. Once HADR reestablished, recreate cluster domain resources. See "Automated HADR configuration setup" whitepaper:
<http://www.ibm.com/developerworks/data/library/long/dm-0907hadrb2haicu/>



#ibmresilient

134


Lead in the Insight Economy

In an automated High Availability Disaster Recovery (HADR) environment, rolling upgrade procedure will not work when upgrading from an earlier to a later version of a DB2 database system.

Note : Deleting the cluster domain is an irreversible step. Ensure that a snapshot of the resource model is captured via the **Issam** command before deleting the cluster domain. The resource model will need to be recreated once the DB2 upgrade has completed.

Upgrading DB2 servers in an automated HADR environment:

https://www.ibm.com/developerworks/community/blogs/IMSupport/entry/75_ways_to_demystify_db2_31_upgrading_db2_servers_in_an_automated_hadr_environment?lang=en_us&hootPostID=259ffd5591284888fab1b8a481941758

DB2 10.5 for LUW New Features and Database Upgrade Considerations



#IBMinsight

- Course Title: **DB2 10.5 for LUW New Features and Database Upgrade Considerations**

Course Code: **CL314** (classroom)

3L314 (ILO- Instructor Led Online)

Course number may vary depending upon language and Global Training Partner.



- This course can be customized to meet your needs.
- Checkout IBM Analytics Certification and Skills at:
<http://www.ibm.com/software/data/education/>



Insight2015

Lead in the **Insight Economy**

141

Course Code CL314 provides training on DB2 10.5 for LUW New Features and Database Upgrade Considerations. See the above URL for more details.

DB2 10.5 Additional Information



#IBMinsight



- **DB2 10.5 with BLU Acceleration - New Dynamic In-Memory Analytics for the Era of Big Data**
<http://bit.ly/DB2BLUFlashbook>

- **Technical Talks:**
<http://www-01.ibm.com/software/data/db2/tech-talk.html>



- **BLU HUB: IBM DB2 with BLU Acceleration**
<http://ibmbluhub.com/>

- **SG24-8212 Architecting and Deploying DB2 with BLU Acceleration in your Analytical Environment**



Insight2015

Lead in the **Insight Economy**

142

Download additional information about DB2 10.5 for free from the above urls.

DB2 10.5 for LUW Certification Exams




#IBMinsight

- IBM Certified Database Associate
- **NEW!** Test C2090-615: IBM DB2 10.5 Fundamentals for LUW

- IBM Certified Database Administrator
 - Test C2090-311: IBM DB2 10.5 DBA for LUW Upgrade from DB2 10.1
 - For details: <http://www-03.ibm.com/certify/tests/ovrC2090-311.shtml>

- **DB2 10.1/10.5 for Linux, UNIX, and Windows Database Administration: Certification Study Guide**

- ISBN-10: 1583473750 or ISBN-13: 978-1583473757



Insight2015

Lead in the **Insight Economy**

143

DB2 10.5 for Linux, Unix and Windows Upgrade certification examination, preparation suggestions and sample assessment tests are available.

If you are knowledgeable with DB2 10.5 and are capable of performing the intermediate to advanced skills required in the day-to-day administration of DB2 instances and databases, you may benefit from this certification role.

A resource for Exam C2090-311 (DB2 10.5 DBA for LUW Upgrade from 10.1) is the Knowledge Center.

There is one physical book for preparation of both exams: [DB2 10.1/10.5 for Linux, UNIX, and Windows Database Administration: Certification Study Guide](#)

You can order from Amazon at

<http://www.amazon.com/10-1-Linux-Windows-Database-Administration/dp/1583473750>

Summary

- Describe changes to DB2 10.5 packaging and software and operating system requirements.
- Understand DB2 10.5 changes to DDL, utilities, tools and monitoring.
- Configure a database and instance to implement DB2 10.5 new features.
- Analyze DB2 10.5 changes that affect applications and database runtime behavior.
- Develop a DB2 10.5 upgrade strategy



Insight2015

Lead in the Insight Economy

144

BIO

Melanie Stopfer is a Consulting Learning Specialist and Developer for IBM Software Group. She is recognized worldwide as a database specialist and DB2 LUW technical expert. As an IBM Certified DB2 10.5 DBA, Advanced Technical Expert and Learning Facilitation Specialist, she has provided in-depth technical support to IM customers specializing in recovery performance and database upgrade and migration best practices since 1988. In 2009, Melanie was the first DB2 LUW speaker to be inducted into the IDUG Speaker Hall of Fame and was again selected Best Overall Speaker at IDUG EMEA 2011 and 2012 and IDUG NA 2012 and 2014.

We Value Your Feedback!



insight2015

Don't forget to submit your Insight session and speaker feedback! Your feedback is very important to us – we use it to continually improve the conference.

Access your surveys at insight2015survey.com to quickly submit your surveys from your smartphone, laptop or conference kiosk.

Session DDB-1176
Upgrading to DB2 With the Latest Version of BLU Acceleration



Insight2015

Lead in the **Insight Economy**

145



Thank You

Melanie Stopfer
mstopfer@us.ibm.com



Insight2015
Lead in the **Insight Economy**

© 2015 IBM Corporation

#ibminisight

