

Db2 Reference Guide
for
IBM z/OS
Remote Access Programs

January 12, 2023

International Business Machines Corporation
IBM Z Dallas ISV Center
Dallas, TX USA

Table of Contents

1 Db2 for z/OS Installation and Configuration	3
2 Dataset Information	3
3 Instructions to START/STOP the Db2 subsystem	3
4 Autostart the Db2 Subsystem.....	5
5 Accessing Db2 SPUFI interface and Db2 Products from the host.....	5
6 Accessing Db2 on the guest z/OS system remotely	6
7 Working with Db2 on z/OS.....	7
8 Db2 Archive Log Files.....	7
8.1 Non-SMS Managed Archive Logs.....	8
8.2 SMS managed Archive logs.....	9
9 Deleting Db2 Archive Log Files	9
9.1 Deleting Non-SMS Managed Archive Logs	9
9.2 Deleting SMS-Managed Archive Logs	9
9.3 Archive Logs Not Allocated.....	11
10 Db2 Maintenance	12
11 Db2 Utilities	13
12 Db2 Product Tools	14
12.1 Default Db2 product tools.....	14
12.2 All available Db2 product tools	14
12.3 How to access Db2 product tools.....	15
13 Continuous Delivery for Db2 V13	15
14 Technical Support	16

Db2 for z/OS is an optional product available upon request for the z/OS remote access environment. Please contact the IBM Z Dallas ISV Center if you wish to add Db2 to your guest z/OS system.

1 Db2 for z/OS Installation and Configuration

When requested to be connected to the guest z/OS system, Db2 will be installed and configured as documented in the Db2 for z/OS program directory and installation guide, with some minor modifications to conform to IBM Z Dallas ISV Center installation guidelines.

2 Dataset Information

Db2 for z/OS datasets will have the following format and contain information as described below:

- The high-level qualifier is DSNxxx, with xxx representing the version and release of Db2, e.g. DSN10, DSN10. If you have multiple Db2 subsystems of the same version, the high level qualifier could also be DSNx20.* where x represents the version of Db2.
- The sample installation and IVP jobs are found in library DSNxxx.NEW.SDSNSAMP.
- Installation, maintenance information, and configuration samples are provided in datasets named DSNxxx.SVSC.*.
- The default DSNZPARMS information is contained in dataset DSNxxx.NEW.SDSNSAMP(DSNTIJUZ).

3 Instructions to START/STOP the Db2 subsystem

Db2 started tasks are found in dataset VENDOR.PROCLIB and follow a naming standard of DBxG*, where x is either C (V12) or D (V13) representing the version of Db2.

The following commands can be issued from the LOG display in SDSF. From the TSO Primary Option Menu, type SD.LOG and press ENTER. In the following examples, Db2 V13 is being started.

To START Db2, issue the following command:

```
/-DBDG START DB2
```

```
DSN9022I  -DBDG DSNYASCP 'START DB2' NORMAL COMPLETION
IXL014I IXLCONN REQUEST FOR STRUCTURE DSN10_GBP8K0 096
WAS SUCCESSFUL.  JOBNAME: DBDGDBM1 ASID: 0041
CONNECTOR NAME: DB2_DBDG CFNAME: SVCF
DSNA672I  DSNA6MAI START COMMAND FOR ADMIN SCHEDULER DBDGADMT NORMAL
COMPLETION
DSNL519I  -DBDG DSNLIRSY TCP/IP SERVICES AVAILABLE 966
          FOR DOMAIN S0W1.DAL-EBIS.IHOST.COM AND PORT 5047
```

Db2 Starting Normally

This will start the following tasks associated with Db2:

DBDGMSTR
DBDGDBM1
DBDGIRLM
DBDGDIST
DBDGADMT

To STOP Db2 and all associated started tasks, issue the following commands:

```
/F DBDGADMT,APPL=SHUTDOWN
/-DBDG DIS THD(*)
    • Use this command to look for active threads
/-DBDG CANCEL THREAD(xx)
    • Where xx is the TOKEN from the display if there are any active threads
/-DBDG STOP DB2
```

```
DSN9022I  -DBDG DSNYASCP 'STOP DB2' NORMAL COMPLETION
IEF352I ADDRESS SPACE UNAVAILABLE
$HASP395 DBDGMSTR ENDED - RC=0000
IEA989I SLIP TRAP ID=X33E MATCHED.  JOBNAME=*UNAVAIL, ASID=003C.
DSN3104I -DBDG DSN3EC0X - TERMINATION COMPLETE
DSN3100I -DBDG DSN3EC0X - SUBSYSTEM DBDG READY FOR START COMMAND
DXR110I DJCG001 STOP COMMAND ACCEPTED
```

Db2 Stopping Normally

4 Autostart the Db2 Subsystem

If you would like your system to automatically start Db2 after an IPL, you will need to edit `VENDOR.PARMLIB(ALLJ2)` and add the following line to the ALLJ2 member and save it:

```
/*$VS, '-DBxG START DB2'
```

Where *x* is either C for V12 or D for V13.

Once changes have been made to the member, on the next IPL and continuing forward, your Db2 subsystem should automatically start.

```
EDIT          VENDOR.PARMLIB(ALLJ2) - 01.02
Command ==>
***** Top of Data *
000001 /*$VS, 'D ASM'
000002 /*$VS, 'S SSHD'
000003 /*$VS, '-DBDG START DB2'
***** Bottom of Data
```

VENDOR.PARMLIB(ALLJ2)

5 Accessing Db2 SPUFI interface and Db2 Products from the host

The datasets required to access the Db2 interface and Db2 product panels get allocated to your TSO session when you first logon to TSO. Several different TSO logon procs are available to use, depending on what version of Db2 you will be using. When logging on to TSO, you must specify the correct proc to get the required datasets allocated.

Once logged on to TSO, you can access the DB2I (DB2 Interactive) Primary Option menu by selecting D2 from the ISPF Primary Option Menu (the first screen you see).

Below are the correct logon procs to use. These procs vary depending on the level of z/OS you are running and the version of Db2 you wish to access.

For guest z/OS release 1.13 systems and higher:

DBPROCns - '*n*' represents the version of Db2 you want to use. '*s*' represents the Db2 subsystem identifier (usually G but could be different if you have multiple Db2 subsystems at the same Db2 version level). For example, **DBPROCCG** is used to access DBCG subsystem and product panels and **DBPROC DG** is used to access DBDG subsystem and product panels.

You can verify the logon procs available on your system by looking in the following PROCLIB datasets (in this order):

1. `VENDOR.PROCLIB`

2. SVTSC.PROCLIB
3. LVL0.PROCLIB

6 Accessing Db2 on the guest z/OS system remotely

Information needed to access the different levels of Db2:

SSID	Location	TCPIP Port	IP Address
DBC G	DALLASC	5040	xx.xx.xx.xx
DBDG	DALLASD	5045	xx.xx.xx.xx

Following is one way to setup Db2 Connect to Db2 for z/OS. To set up connectivity to Db2 for z/OS from a distributed platform, execute the following commands from the Db2

Command Line Processor:

```
db2 CATALOG TCPIP NODE znode REMOTE hostname SERVER db2port
db2 CATALOG DB dbname AT NODE znode
```

Test the connection by issuing the following command:

```
db2 CONNECT TO dbname USER uid USING pwd
```

Notes:

znode - your choice of node name for the zSeries on the client

hostname - the existing hostname or IP address of the z/OS system

db2port - tcpip port number assigned to this Db2 subsystem on z/OS (From delivery email)

dbname - the LOCATION of the subsystem (From delivery email)

uid - RACF userid defined with the appropriate authority (Cannot start with SQL or SYS)

pwd - password defined for the userid. The USING keyword and password can be omitted from the command and in that case Db2 will prompt for the password.

List commands on the client to verify the catalogs:

```
db2 LIST NODE DIRECTORY
db2 LIST DB DIRECTORY
```

If you receive error like the following:

Invocation of routine "SYSPROC.DSNUTILS "failed due to reason "00E79002".

SQLSTATE=55023

It means, that the stored procedure cannot be executed, possibly, because of lack of WLM configuration or the WLM is stopped.

To verify that WLM is active, navigate to SDSF and enter the display command `/d wlm` in the command line.

If WLM is missing, please contact the Dallas ISV Center via the IBM Z Dallas ISV Center Community user group: <https://ibm.biz/DallasISVCommunity>

If you are not a member, please refer to **Technical Support** on page 16 to join.

7 Working with Db2 on z/OS

Following are some words of advice when creating Db2 for z/OS databases and tables. The proper approach when designing databases, tablespace and tables in z/OS is to follow these steps:

1. Create your own STOGROUP and specify a VCAT name (high level qualifier of your dataset) and also specify the volume(s) you want to use for your application data. You can find volume name information in the delivery email that was sent to the system contact.
2. Create the DATABASE using the above STOGROUP. You do NOT want to create a DATABASE without specifying the STOGROUP, otherwise you will use the default STOGROUP SYSDEFLT. It is not recommended to use STOGROUP SYSDEFLT.
3. Create separate tablespaces in the above DATABASE for **EACH** table you want to create. This will help Db2 to perform much better for you.
4. Create the table in the above tablespace.
5. If a database or tablespace is not specified, Db2 will create one for you using the DEFAULT database DSNDB04 and default STOGROUP SYSDFELT. This is not recommended.
6. Note if more space is needed for application tablespaces, volumes can be added. ALTER the STOGROUP to add work volume(s) VPWRKx. Use SPUFI to execute the following SQL:

```
ALTER STOGROUP stogroupname ADD VOLUMES (VPWRKx)
```

8 Db2 Archive Log Files

Db2 has active log datasets that archive when they get full. The archive task produces archive log datasets that follow a naming standard of DSNx10.DBxG.ARCLOG* (where x represents the version of Db2).

Db2 creates two data sets for one archive log. In the archive logs name, the fourth qualifier contains the date in Julian format and the fifth qualifier includes the time.

- DSNX10.DBXG.ARCLOG1.D20186.T0003059.A0000444
- DSNX10.DBXG.ARCLOG1.D20186.T0003059.B0000444

It is much easier to manage archive logs if the time stamp option is turned on. If your archive logs do not have the date/time stamp, change the TSTAMP=NO option in DSNx10.NEW.SDSNSAMP(DSNTIJUZ) to TSTAMP=YES.

8.1 Non-SMS Managed Archive Logs

Archive logs that are NON-SMS managed need enough space to be created on a DASD volume. Ensure there is enough space for archive logs to be created. Without enough space, Db2 can stop working. Follow the steps below to mount a DASD volume if more space is needed:

Check your delivery email for available volumes on your system and the amount of space on them. Below is a sample list of possible available volumes:

```
VPWRKA (0D30)  xxxx cyls
VPWRKB (0D31)  xxxx cyls
VPWRKC (0D32)  xxxx cyls
VPWRKD (0D33)  xxxx cyls
```

Volumes must be mounted as STORAGE to use for non-sms managed Db2 archive logs. For example, if you want to mount a work volume, VPWRKA, as use=storage, you would perform the following steps:

1. Logon to TSO
2. Go into SDSF - option SD from TSO ISPF main menu
3. Type in LOG on the command line and press ENTER
4. On the command line enter **/D U,,D30,10** (to display these volumes)
5. On the command line enter **/M D30,VOL=(SL,VPWRKA),USE=STORAGE**
6. You can then scroll through the LOG output to see the command and the results. Scroll up or down using the F7 or F8 keys
7. To ensure the volume is mounted as storage, enter the following command: **/D U,,D30,1**

You will see output from the above command that looks like:

UNIT	TYPE	STATUS	VOLSER	VOLSTATE
0D30	3390	A	VPWRKA	STRG/RSDNT

When you issue these commands manually, they stay in effect for as long as your system is up. When you IPL (SVXLOG) your system, it resets everything. To make this change permanent, you will need to add an entry to member VATLST00 in VENDOR.PARMLIB.

Follow these steps:

1. Look in VENDOR.PARMLIB for member VATLST00. If you don't see it there, look in SVTSC.PARMLIB then LVL0.PARMLIB.
2. Copy VATLST00 from SVTSC.PARMLIB or LVL0.PARMLIB into VENDOR.PARMLIB.
3. Add an entry to the bottom of this member that looks like this:
 %PWRKA,0,0,* ,Y WORK VOLUME (% ADDED FOR 3380/3390)
4. Press F3 key to save the member.

8.2 SMS managed Archive logs

SMS managed archive logs are directed to the VPARC1 volume. If you need more ARCHIVE log space, please contact the Dallas ISV Center via the IBM Z Dallas ISV Center Community user group: <https://ibm.biz/DallasISVCommunity>
If you are not a member, please refer to **Technical Support** on page 16 to join.

9 Deleting Db2 Archive Log Files

The most important thing about deleting Db2 archive data sets is to ensure Db2 no longer needs the archive data set you are going to delete. If an archive data set is deleted and Db2 still needs it to start, you are going to chance corrupting the Db2 data and a Db2 coldstart may be required.

9.1 Deleting Non-SMS Managed Archive Logs

If the Db2 locks up because it cannot archive its active logs, it is best to mount a DASD volume (see page 6) as "STORAGE" to allow Db2 to finish its archive process before you attempt to delete the archive logs.

9.2 Deleting SMS-Managed Archive Logs

There are one of two ways to delete SMS-managed Db2 archive logs. If there are obsolete Db2 archive log data sets, they can be deleted. To determine if the archive logs are obsolete, look at the Julian date in the dataset name. It is best practice to keep archive data sets that were created after the last two system recycles.

To delete all archive logs, issue the command `/-DBDG ARCHIVE LOG` at least **TWO** times if you have three Db2 active logs (this is default on all RDP systems). This command will force Db2 to cut an archive log. Performing this command twice will clear out the active logs. Next, shut down Db2 and ensure it comes down cleanly. After a clean shutdown, start Db2. If

Db2 makes no reference to an archive log, then all archive logs can be deleted. Make sure between the clean startup of Db2 and deletion of the archive logs, no one makes changes to Db2.

Once you have determined it is OK to delete an archive log, you can go to ISPF 3.4 and take the following action:

Select all the archive logs for DBDG by entering DSND10.DBDG.ARC* in the following display.

```

                                Data Set List Utility
Option ===>

      blank Display data set list          P Print data set list
      V Display VTOC information          PV Print VTOC information

Enter one or both of the parameters below:
  Dsname Level . . . DSND10.DBDG.ARC*
  Volume serial . .

```

Data Set List Utility

This should show you a list similar to the following display:

```

DSLISIT - Data Sets Matching DSND10.DBDG.ARC*          Row 1 of 4

Command ===>          Scroll ===> CSR

Command - Enter "/" to select action          Message          Volume
-----
      DSND10.DBDG.ARCLOG1.D09181.T1359114.A0003011          VPARC1
      DSND10.DBDG.ARCLOG1.D09181.T1359114.B0003011          VPARC1
      DSND10.DBDG.ARCLOG1.D09181.T1425367.A0003012          VPARC1
      DSND10.DBDG.ARCLOG1.D09181.T1425367.B0003012          VPARC1

```

DSLISIT - List of ARCLOG datasets

Then enter a line command "del / purge " followed by = on each following line and press the ENTER key. See example:

```

DSLIST - Data Sets Matching DSND10.DBDG.ARC*                               Row 1 of 4

Command ===>                               Scroll  ===> CSR

Command - Enter "/" to select action      Message                               Volume
-----
del / purgeND10.DBDG.ARCLOG1.D09181.T1359114.A0003011          VPARC1
=          DSND10.DBDG.ARCLOG1.D09181.T1359114.B0003011          VPARC1
=          DSND10.DBDG.ARCLOG1.D09181.T1425367.A0003012          VPARC1
=          DSND10.DBDG.ARCLOG1.D09181.T1425367.B0003012          VPARC1

```

Example of delete purge command

9.3 Archive Logs Not Allocated

If you receive an error similar to the one below, attempt deleting obsolete archive logs first to kick start the logging process. If deleting obsolete archive logs corrects the problem, it is recommended to clear out all archive logs. Refer to section 8 Deleting Db2 Archive Log Files. If deleting obsolete archive logs did not fix the error or there are no obsolete archive logs to delete, please contact the Dallas ISV Center via the IBM Z Dallas ISV Center Community user group: <https://ibm.biz/DallasISVCommunity>

If you are not a member, please refer to **Technical Support** on page 16 to join.

```

IKJ56893I DATA SET DSND10.DBDG.ARCLOG1.D20308.T1225363.A0000734 NOT
ALLOCATED+
IGD17273I ALLOCATION HAS FAILED FOR ALL VOLUMES SELECTED FOR DATA SET
325
DSND10.DBDG.ARCLOG1.D20308.T1225363.A0000734
IGD17277I THERE ARE (2) CANDIDATE VOLUMES OF WHICH (2) ARE ENABLED OR
QUIESCED
IGD17290I THERE WERE 1 CANDIDATE STORAGE GROUPS OF WHICH THE FIRST 1
327
WERE ELIGIBLE FOR VOLUME SELECTION.
THE CANDIDATE STORAGE GROUPS WERE:DB2ARCG
IGD17279I 2 VOLUMES WERE REJECTED BECAUSE THEY DID NOT HAVE SUFFICIENT
SPACE (041A041D)

```

Example of archive logs not allocating

10 Db2 Maintenance

Every month the IBM Z Dallas ISV Center applies maintenance to the Db2 systems. Most times, this maintenance requires actions to be taken after this maintenance is rolled in via an IPL of your system. **It is the responsibility of each solution developer to perform these actions on their remote development system.**

We implement RSU level maintenance monthly and PUT level maintenance on a quarterly basis.

PUT level maintenance will be implemented as follows:

January = PUTyy12

April = PUTyy03

July = PUTyy06

October = PUTyy09

Information about the PTFs that were applied and the HOLD ACTIONS required for each PUT/RSU level of maintenance (ie. PUTyy03, RSUyy04, RSUyy05....) can be found in PUTxxxx or RSUxxxx members in the following MVS datasets that are on your system:

DSNC10.SVSC.HOLDDATA

DSND10.SVSC.HOLDDATA

JCL to execute the jobs to resolve the HOLD ACTIONS is found in:

DSNC10.SVSC.HOLDCNTL

DSND10.SVSC.HOLDCNTL

The actions you must take are typically running a z/OS batch job that will perform the required task for you. The members in DSN*10.SVSC.HOLDDATA will give you instructions for JOBS that you must execute from a dataset named DSN*10.SVSC.HOLDCNTL.

If the instructions tell you to run DSN*10.SVSC.HOLDCNTL(REBIND), this is how you would do that:

1. Logon to TSO with the IBMUSER userid. This userid has Db2 SYSADM authority, which is needed to run the batch jobs.
2. Go to the dataset list utility - from TSO main menu, option 3 then 4 - or type 3.4 on the TSO main menu command line
3. Enter the dataset name DSN*10.SVSC.HOLDCNTL
4. Press the enter Key
5. Edit the member (place an E beside the member name and press the enter key)
6. Either remove the userid and password in the jobcard (1st or 2nd line in this job) or

update the password for IBMUSER to match what it is on your system.

7. On the command line, type in SUB and press the enter key.
8. If you made an update to the JCL, you will need to type in CANCEL on the command line and press the ENTER key to back out of that member.

Note:

The REBIND job produces many log records and causes archiving. You can delete these archive log datasets that are being produced by the REBIND job. Follow the instructions in the 'Deleting Db2 Archive Log files' above.

To view the job output:

1. Go to the ISPF main menu (the menu shown when you first log on to TSO)
2. Type "SD" on the command line - press ENTER
3. Type "OWNER IBMUSER" on the command line and press ENTER
4. Type "ST" (for status) on the command line and press ENTER
5. You will see a display of the jobs you ran. Type a '?' beside the job that was just run
6. You will next see a list of job output files. Type an 'S' beside the JESMSG LG file and press enter.

11 Db2 Utilities

DB2 comes with a standard set of utilities that are used to perform such tasks as COPY, UNLOAD, LOAD, REORG as well as others. All Db2 utilities are documented in the appropriate Db2 Utility Guide and Reference which can be found at the Db2 for z/OS - Technical Resources website.

Following are locations where you can find sample JCL to perform some of the more common utilities:

- DSNx10.NEW.SDSNSAMP(DSNTEJ2A) certain steps in this JCL will provide an example of UNLOAD and LOAD
- DSNx10.NEW.SDSNSAMP(DSNTEJ1) contains some sample JCL for REORG using LISTDEF
- DSNx10.NEW.SDSNSAMP(DSNTIJIC) contains sample JCL for performing an IMAGECOPY

Please note that NONE of this JCL can be run AS IS and would require heavy modification by the solution developer. Please be sure to reference the correct version of the Db2 Utility Guide and Reference for instructions on running each utility.

12 Db2 Product Tools

The IBM Z Dallas ISV Center Db2 systems support team configures certain Db2 product tools on each system that request Db2. Below are the Db2 product tools that are configured by default as well as a list of available products that can be configured upon request (and the high-level qualifier).

12.1 Default Db2 product tools

Db2 for z/OS Version 12.1:

- IBM Db2 QMF CLASSIC EDITION V12.01.00
- IBM Db2 ADMINISTRATION TOOL FOR z/OS V12.01.00

Db2 for z/OS Version 13.1:

- IBM Db2 QMF CLASSIC EDITION V13.01.00
- IBM Db2 ADMINISTRATION TOOL FOR z/OS V13.01.00

12.2 All available Db2 product tools

Db2 for z/OS Version 12.1:

HLQ	PRODUCT
DDA211	IBM APPLICATION CONNECTIVITY TO DB2 FOR z/OS V12
ADBC10	IBM DB2 ADMINISTRATION TOOL FOR z/OS V12.01.00
GOCB10	IBM DB2 OBJECT COMPARISON TOOL FOR Z/OS V11.01.00
QMFC10	IBM DB2 QMF CLASSIC EDITION V12.01.00

Db2 for z/OS Version 12.1 products

Db2 for z/OS Version 13.1:

HLQ	PRODUCT
DDD211	IBM APPLICATION CONNECTIVITY TO DB2 FOR z/OS V13
ADBD10	IBM DB2 ADMINISTRATION TOOL FOR z/OS V13.01.00
QMFD10	IBM DB2 QMF CLASSIC EDITION V13.01.00

Db2 for z/OS Version 13.1 products

12.3 How to access Db2 product tools

Db2 product tools can be accessed from the TSO ISPF primary option menu. When on the ISPF primary option menu, press the F8 key to scroll down to see more options. Seeing an option on this menu does not guarantee that it is available on the guest z/OS system. The Db2 Object Comparison Tool is not customized on these systems.

13 Continuous Delivery for Db2 V13

Continuous delivery is a method of getting new function into the market in a timelier manner than the three-year release cycle of a new version of Db2. It allows Db2 users to adopt new functions in small, consumable units.

New function levels are sent through the maintenance stream. Db2 will not update to a new function level until an `ACTIVATE` command is used. Activating a new function level does not do anything other than allow you to use the new function. Applications will not be affected by activating a new function level because their behavior is controlled by `APPLCOMPAT`, not the function level.

To allow your applications to use the new functionality brought in with the function levels, the `APPLCOMPAT` variable will need to be set. The `APPLCOMPAT` can be explicitly defined in the application or fall to the default `APPLCOMPAT` in Db2 13 ZPARMS. The explicit definition of `APPLCOMPAT` in an application overrides the default `APPLCOMPAT` ZPARM. It is the same process for `BINDing`.

Function levels and `APPLCOMPAT`s can be reverted to a lower level if need be.

To determine what function level is activated on your Db2 13 subsystem, use the command `/-DBDG DISPLAY GROUP DETAIL` from `SD.LOG`.

```

-DBDG DISPLAY GROUP DETAIL
DSN7100I  -DBDG DSN7GCMD
*** BEGIN DISPLAY OF GROUP(.....) CATALOG LEVEL (V13R1M501)
      CURRENT FUNCTION LEVEL (V13R1M501)
      HIGHEST ACTIVATED FUNCTION LEVEL (V13R1M501)
      HIGHEST POSSIBLE FUNCTION LEVEL (V13R1M502)
      PROTOCOL LEVEL (2)
      GROUP ATTACH NAME (.....)
-----
DB2          SUB          DB2          SYSTEM          IRLM
MEMBER      ID   SYS   CMDPREF   STATUS   LVL      NAME      SUBSYS  IRLMPROC
-----
.....  0   DBDG -DBDG   ACTIVE  131502  S0W1      DJDG     DBDGIRLM
-----
SPT01 INLINE LENGTH:      32138
*** END DISPLAY OF GROUP(.....)
DSN9022I  -DBDG DSN7GCMD 'DISPLAY GROUP ' NORMAL COMPLETION

```

Example of DISPLAY GROUP

Note:

CATALOG LEVEL – When activating a new function level, the catalog may also need to be updated to a higher function level. Changing the catalog level has no effect on your applications. Catalog level cannot be reverted back to a lower function level.

CURRENT FUNCTION LEVEL – The function level your Db2 subsystem is currently operating at.

HIGHEST ACTIVATED FUNCTION LEVEL – The highest function level that has been activated on your subsystem.

HIGHEST POSSIBLE FUNCTION LEVEL – The highest level you can activate your Db2 subsystem to.

DB2 LVL – Also known as code level. This changes when a new function level is released and applied to the system via maintenance. This attribute is there to show when a new function level release is available on the system.

14 Technical Support

Technical Support is provided via our User Group (IBM Z Dallas ISV Center) at this link: <https://ibm.biz/DallasISVCommunity>. To join this User Group please follow these instructions:

Join the IBM Z Dallas ISV Community Website

This secure user group is a place for partners to collaborate, ask questions, trade tips and tricks. Your membership will not be approved if you are not enrolled in an active program in the IBM Z Dallas ISV environment.

To join, you will first need to join the IBM Z and LinuxONE Community. You can use the following invite link to do this. <https://ibm.com/community/z/join>

- Sign in with an IBMid
If you do not have an IBMid, you may create one by clicking on the link below the login box that says Don't have an account? Create an IBM id.
- Select your preference for discussion notifications
- Click Yes join the group

Once you have joined the IBM Z and LinuxONE community, please use this form to request access to the IBM Z Dallas ISV Center Community: <https://www.ibm.com/community/z/ibm-z-dallas-isv-center-request-membership/>

Once approved, you will receive an email containing a unique link to join the group.

Hours of attended operation are 07:00 to 19:00 US Central time Monday through Friday except IBM USA holidays. Emails received after hours, on weekends, or on holidays will be processed the next business day. If you require technical support, please follow the instructions on <http://dtsc.dfw.ibm.com/contact.html>

If confidential support is needed, please e-mail us at ztech@us.ibm.com.