
Db2 Basics: Monitoring and Troubleshooting Db2
Ember Crooks
Session code:
#IDUGDb2
Db2

Agenda

- Db2 Documentation
- Error Messages
- Db2 Diagnostic Log
- Overview of Monitoring Tools:
  - db2top
  - ~~dsmtop~~
  - dmctop
  - monreport Module
  - mon_get Table Functions and Views

Troubleshooting Methodology

Take a Systemic Approach

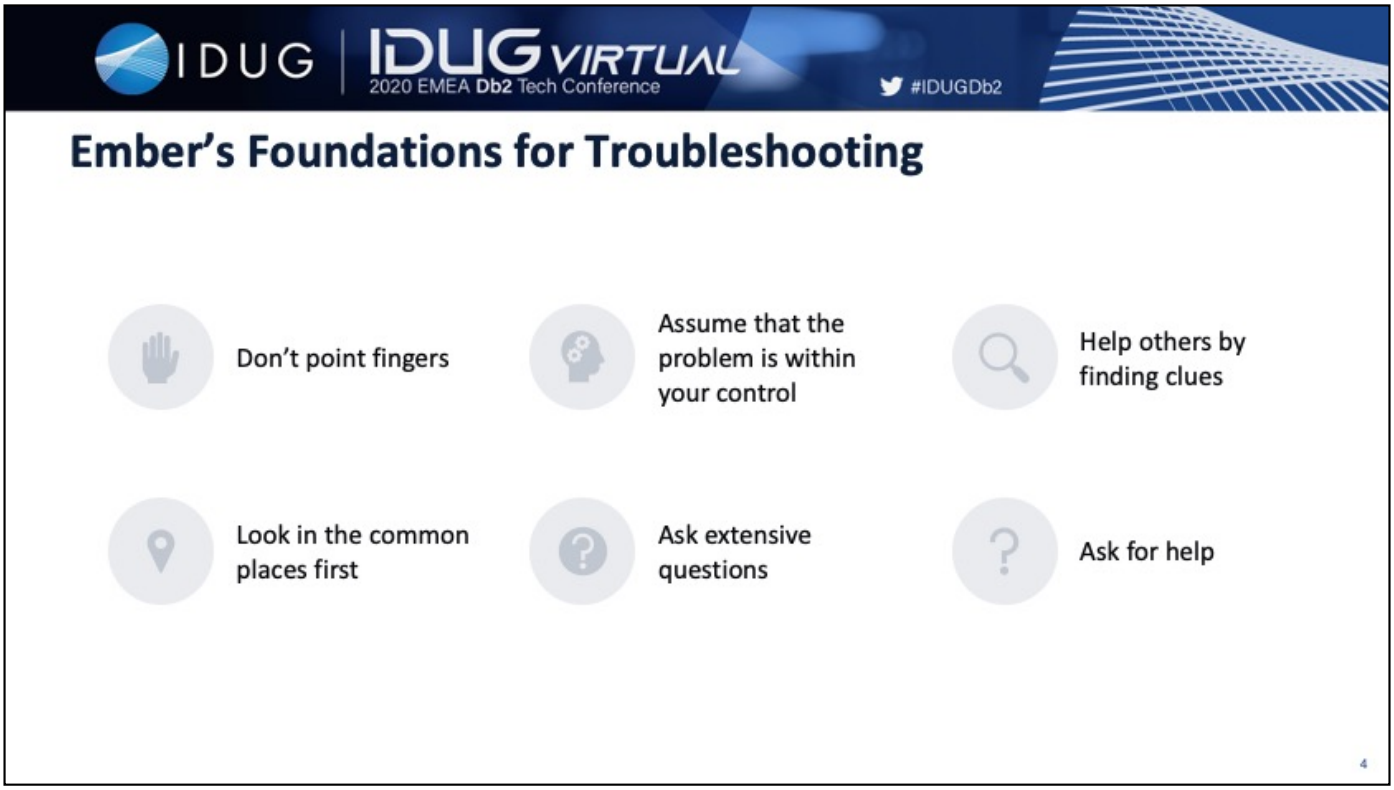

When troubleshooting you're rarely able to immediately say "this is a database problem", "this is an application problem", or "this is a network problem". The key is that if you're working with experts in other areas or even if you're trying to cover all the areas yourself, don't rule things out prematurely. If someone calls me in for a problem, I do my best not to just say "That's not the database.", even if it clearly isn't. Even if it isn't the database, often there are clues that you can find within the database that may help others locate the problem in their areas. My approach is to start from "This is in the database, but where?", and then once I have significant proof it isn't, provide others with the details that will help them in their own troubleshooting efforts.

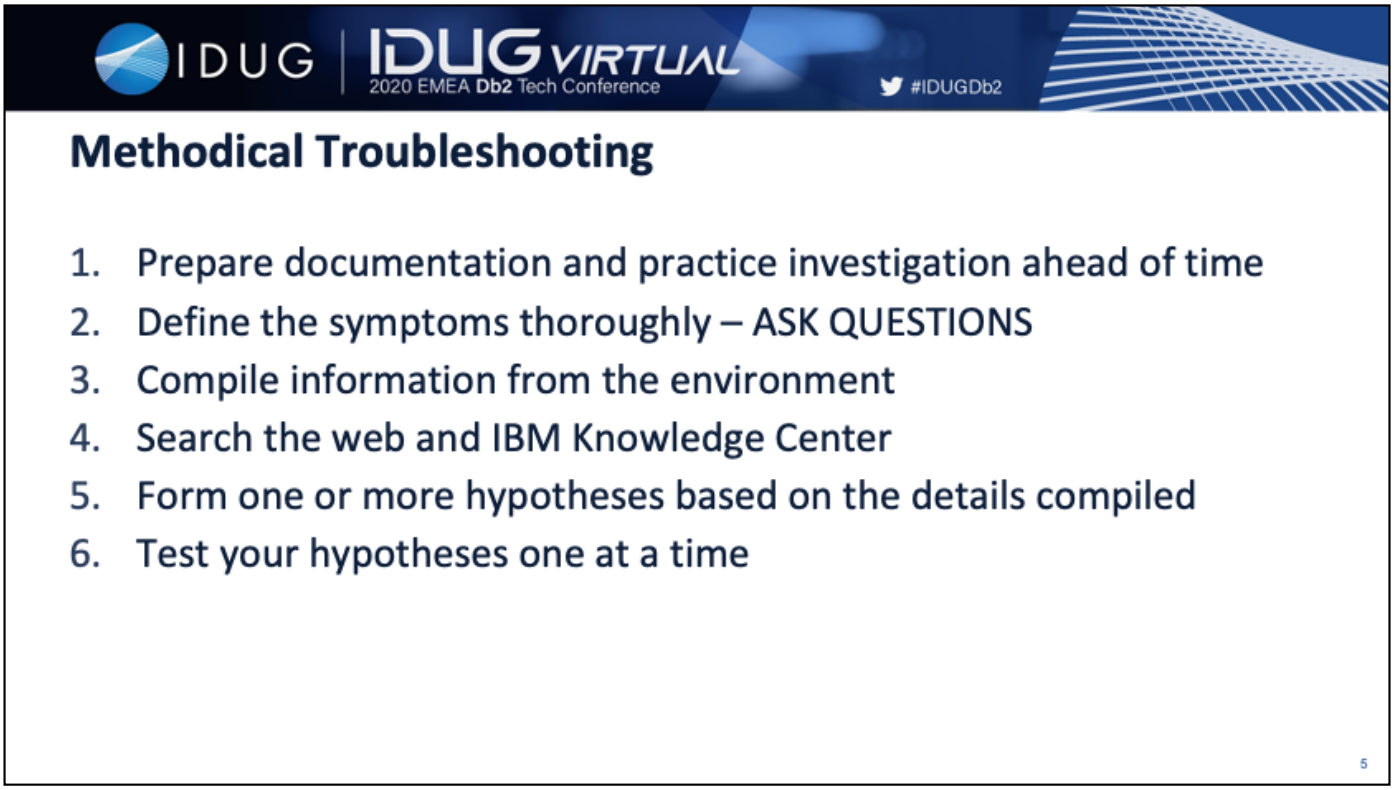

It is hard to perform proper troubleshooting if you don't have the documentation to understand what is in place, and if you haven't prepared with basic skills ahead of time. It is hard to practice troubleshooting, but you can at least practice using the tools. I worked with a junior DBA once who would pass his laptop to me and say "break my database!". He was running Db2 on a VM, and this was an excellent way for him to come across at least a few scenarios to practice on.

Asking questions is one of the most important troubleshooting skills. The first symptoms of a problem may come in through user complaints and "the database doesn't work" simply isn't enough information. Gather all the information you can, as you never know which details are going to be relevant.

It is easy in the heat of troubleshooting to start chasing what you think is the problem and overlook some basic. The two things that come to mind here are checking the diagnostic log for errors and checking to see when the last runstats was. These are both fundamental things that may be critical to the problem. Never forget to check the basics.

I well remember the days when searching a problem online didn't return many results, and we spent a lot more time with the manuals and asking our colleagues for advice and expertise. Now there is a wealth of information online, and even if someone doesn't have the exact same problem, their information may point you in the right direction to look or try things.

It is very helpful if you can formulate one or more hypotheses as to why the problem is occurring and why now. With those in hand, you can move forward to actions or information gathering that can either rule out or confirm your ideas.

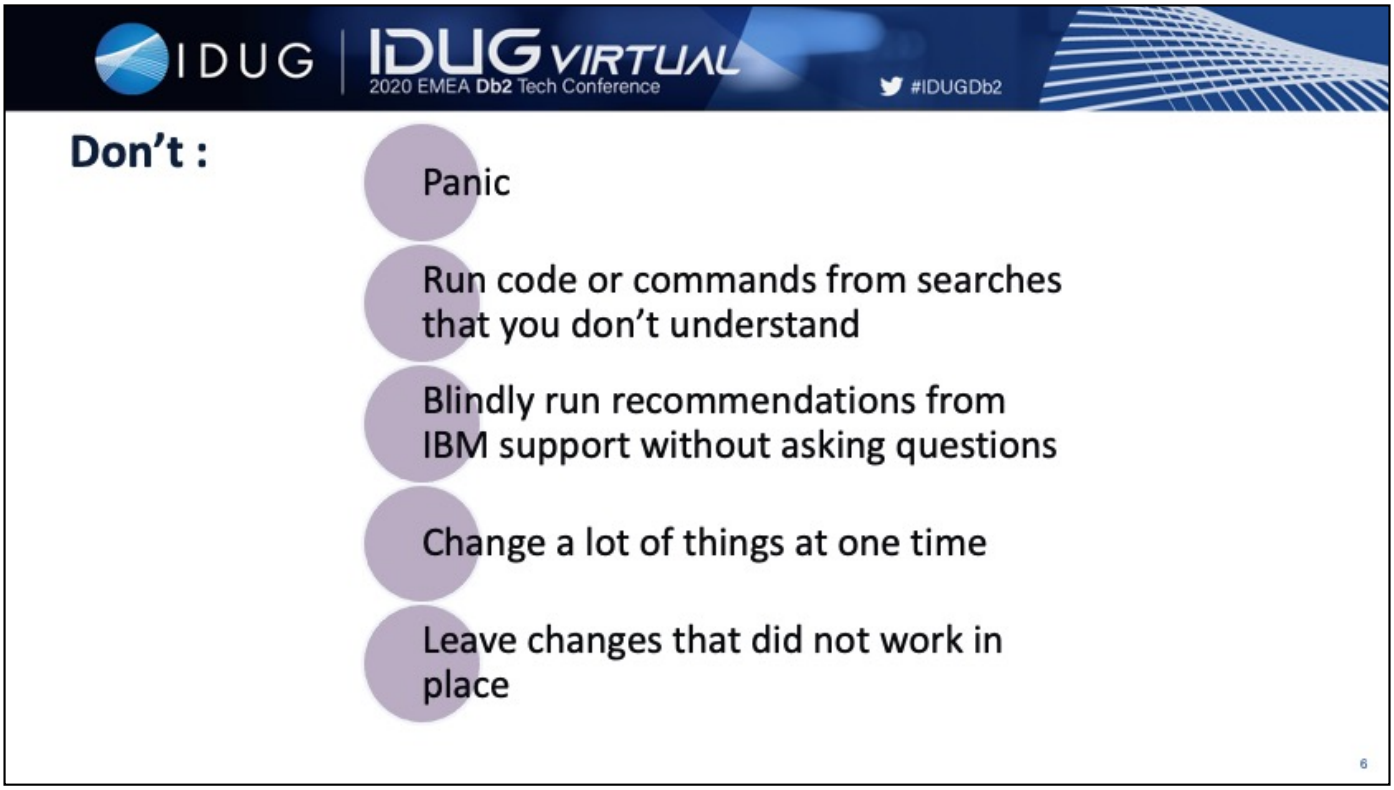

It is very important to stay calm, positive, and focused while troubleshooting. If someone called you in the middle of the night, and you had trouble logging in, and you're nervous, just taking a moment to take a deep breath can be hugely beneficial.

There's an internet legend of someone who was having query performance problems and found a post online recommending truncating the table. They truncated the table, and only then learned that truncation meant deleting all the data, and their data was now gone. Maybe that's an extreme example, but it is easy to run across commands that you don't understand while searching. Be cautious and make sure you understand the impacts before running these commands. Many commands can impact performance or availability further, and some can cause harm that is hard to undo.

As absurd as it may seem, this also applies to recommendations from IBM support. More than once, I've had support recommend something that will cause an outage without telling me it will. More than once, I've had support recommend something that was not appropriate for my environment. Ask support questions about the impact of their recommendations, and if they cannot answer those questions, either insist on answers before you run them, or research to understand the impact yourself.

It can be tempting when troubleshooting to change a number of things at once, particularly if you have more than one thing you need to try that requires an outage. Try to resist this and change only one thing at a time, so you understand what it is that really resolved the problem.

If possible, remove changes that did not fix the problem. Sometimes you try things, and it's best to remove them if they did not work.

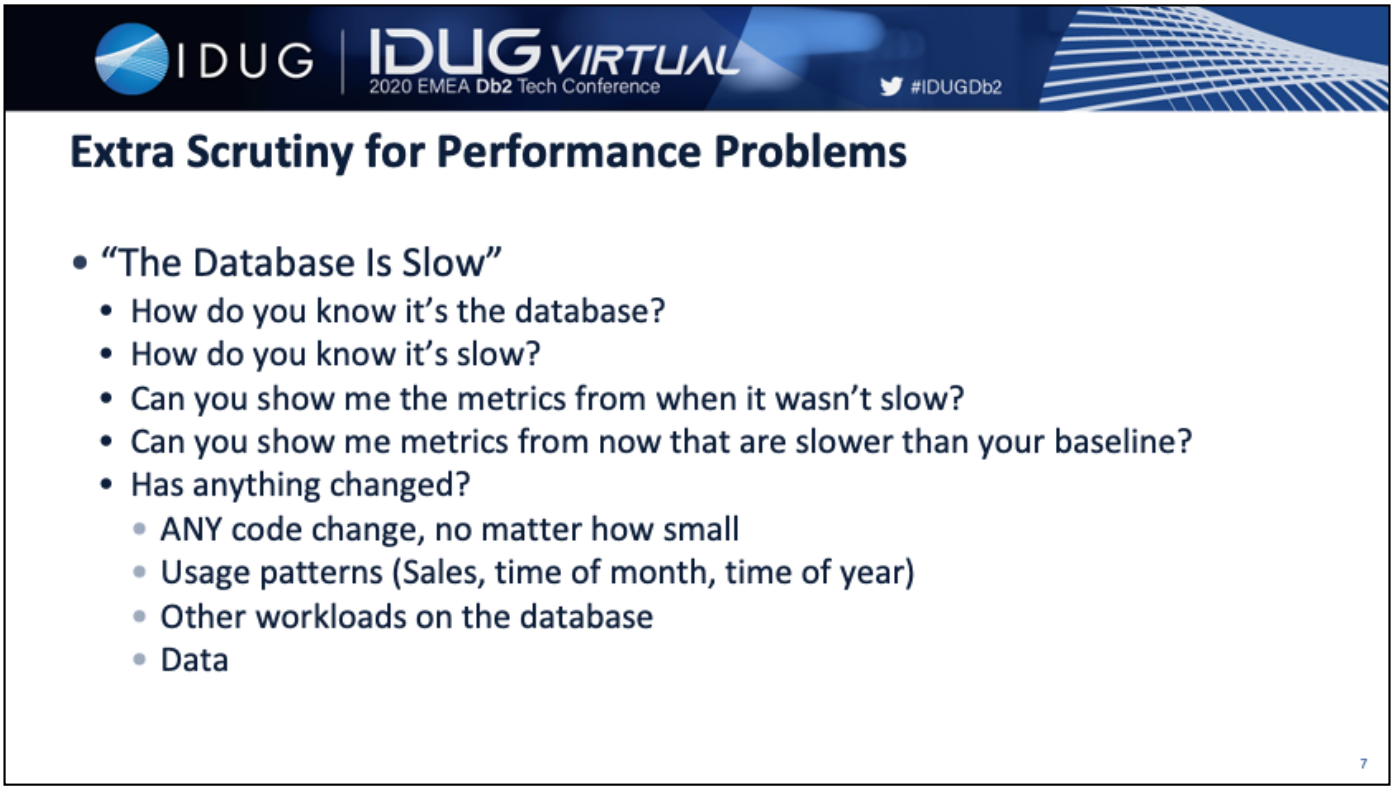

One of the keys for troubleshooting is to ask questions. This is more important for performance troubleshooting than any other area. I once had a developer come to me with a performance problem, complaining that simply returning the CURRENT TIMESTAMP from SYSIBM.SYSDUMMY1 was taking 800 milliseconds – nearly a whole second. This should return very quickly. As I started to investigate, he sent me a screen shot showing the time, and it said 800 microseconds, not milliseconds, which was much more reasonable. Always ask questions, always ask to see what someone is looking at and dig into the details

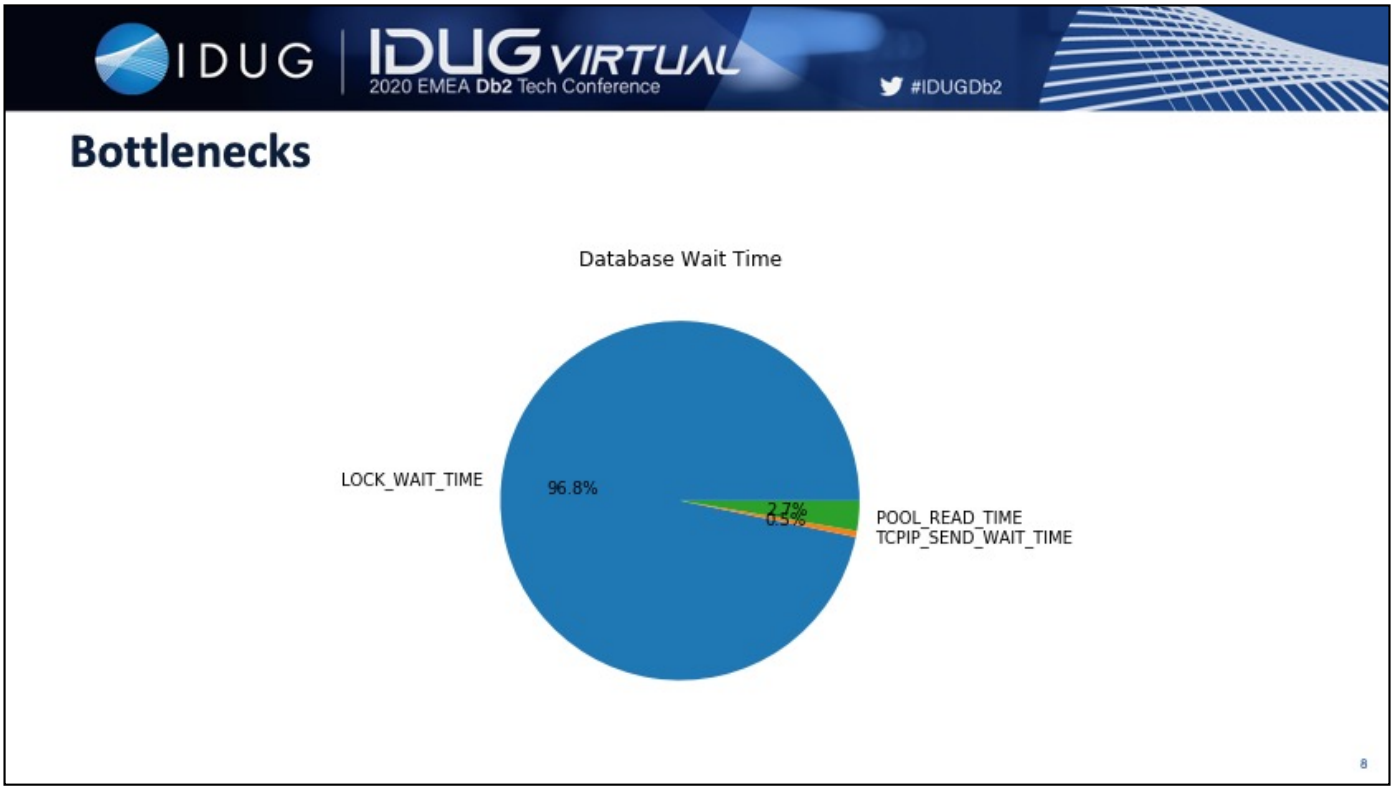

It is sometimes easy to start off investigating a problem, find a promising direction, and loose sight of the bigger picture. For example, if I have a graphic like this one representing wait time within my database, and I go off looking into a network problem, the most I can decrease database wait time by is 0.5%. In this example, my bottleneck is clearly LOCK_WAIT_TIME, and if I spend too much time in other areas, I'm probably wasting my time. Even if I completely eliminated all pool read and all network time, I'd barely touch the overall wait time. That doesn't mean that there cannot be some other area contributing to the lock wait time, just that if whatever I'm working on doesn't reduce lock wait time in some way, it is not addressing my bottleneck.

There are often bottlenecks in database systems. As soon as we eliminate one, something else shows up as the new bottleneck. I remember when a network bottleneck was eliminated once between application servers and the database server – the network went from 10M to 100M. And that made it so database performance was the new focus. Had database performance changed? No. But with the old bottleneck gone, the database was the new bottleneck.

Db2 Documentation

Navigating the IBM Db2 Documentation

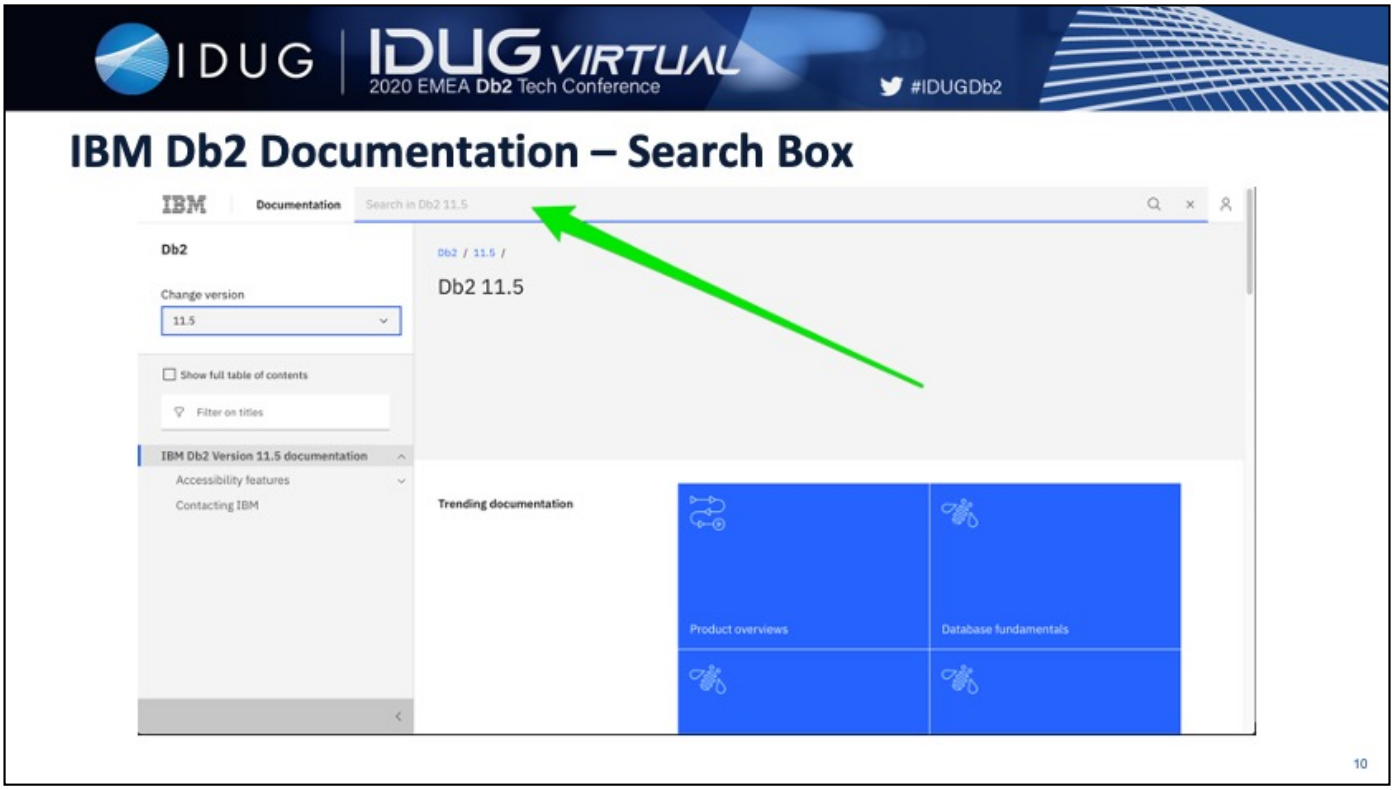

There is a search box at the upper right. The search in the IBM Db2 Knowledge Center is an order of magnitude better than it used to be. Google or your favorite search engine will also return pages from the Knowledge Center, but sometimes in different languages.

Search – Results

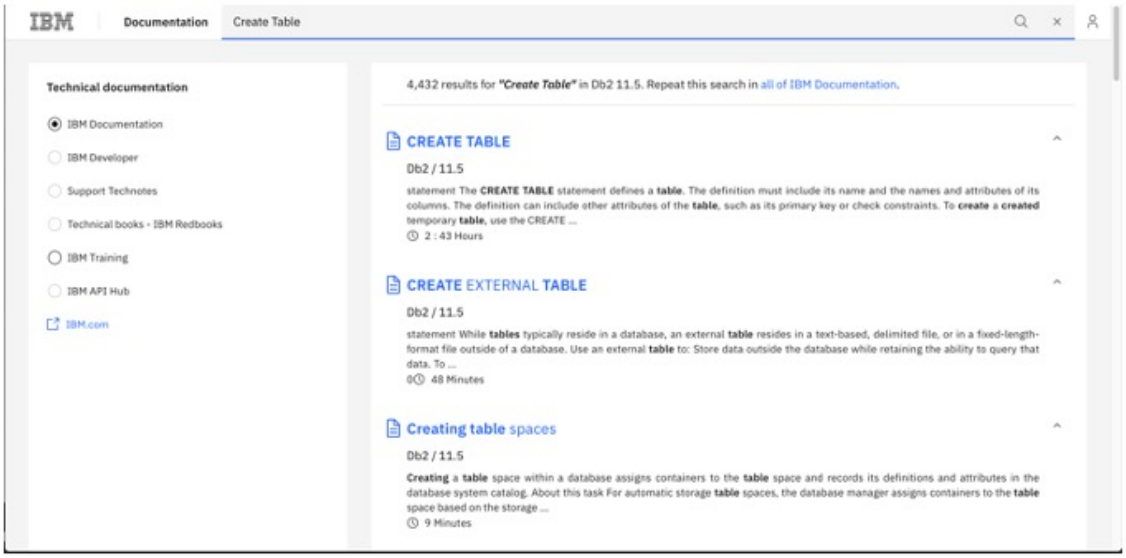

11

By clicking the plus-sign at the left of any search result, you can get an idea of what the result looks like to decide if it is the one you actually want to click on.

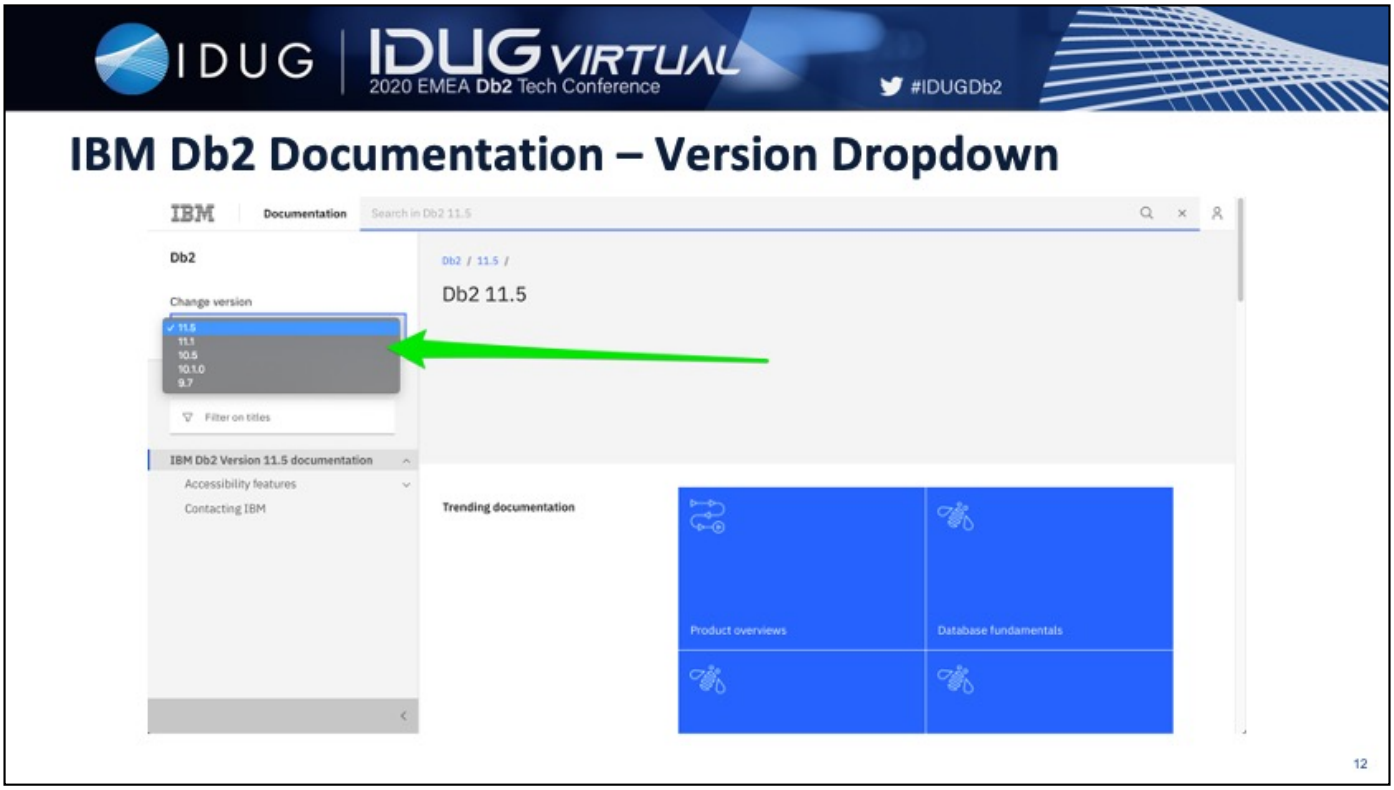

Some of the changes between versions are drastic. It is critical to make sure you're looking at the documentation for the version you're working with. On every page, there is a drop-down box that allows you to select different versions of that page. For some versions, some pages don't exist. For example, BLU functionality was added in 10.5, and so the pages related to 10.5 don't exist in 10.1 and before. Documentation is also removed for the oldest versions.

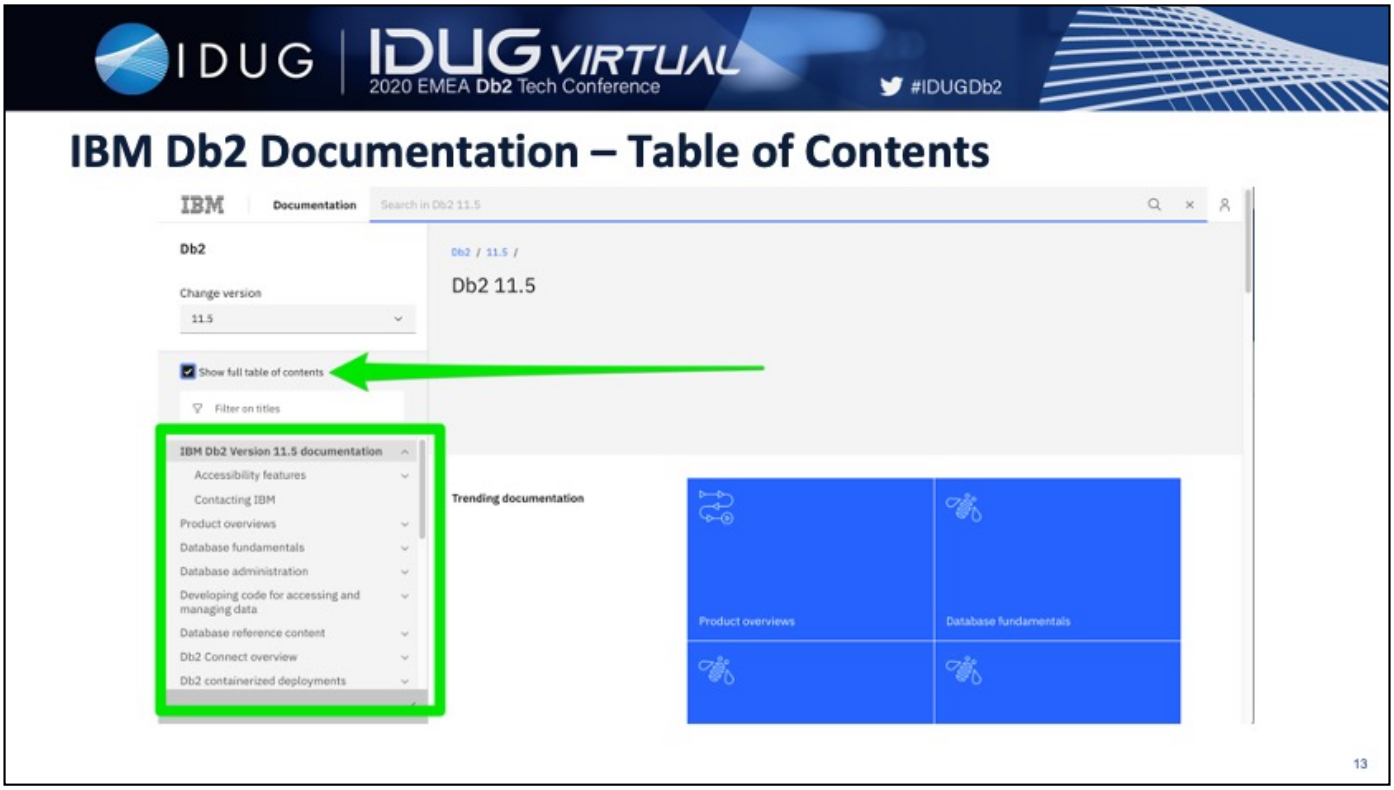

For any page you're on, you can also pull up the table of contents. This will show all the different areas, as well as where the page you currently have up is within the hierarchy. For old-school dbas, this is really easy to navigate because we used to have to use an entire shelf's worth of Db2 books to find this information.

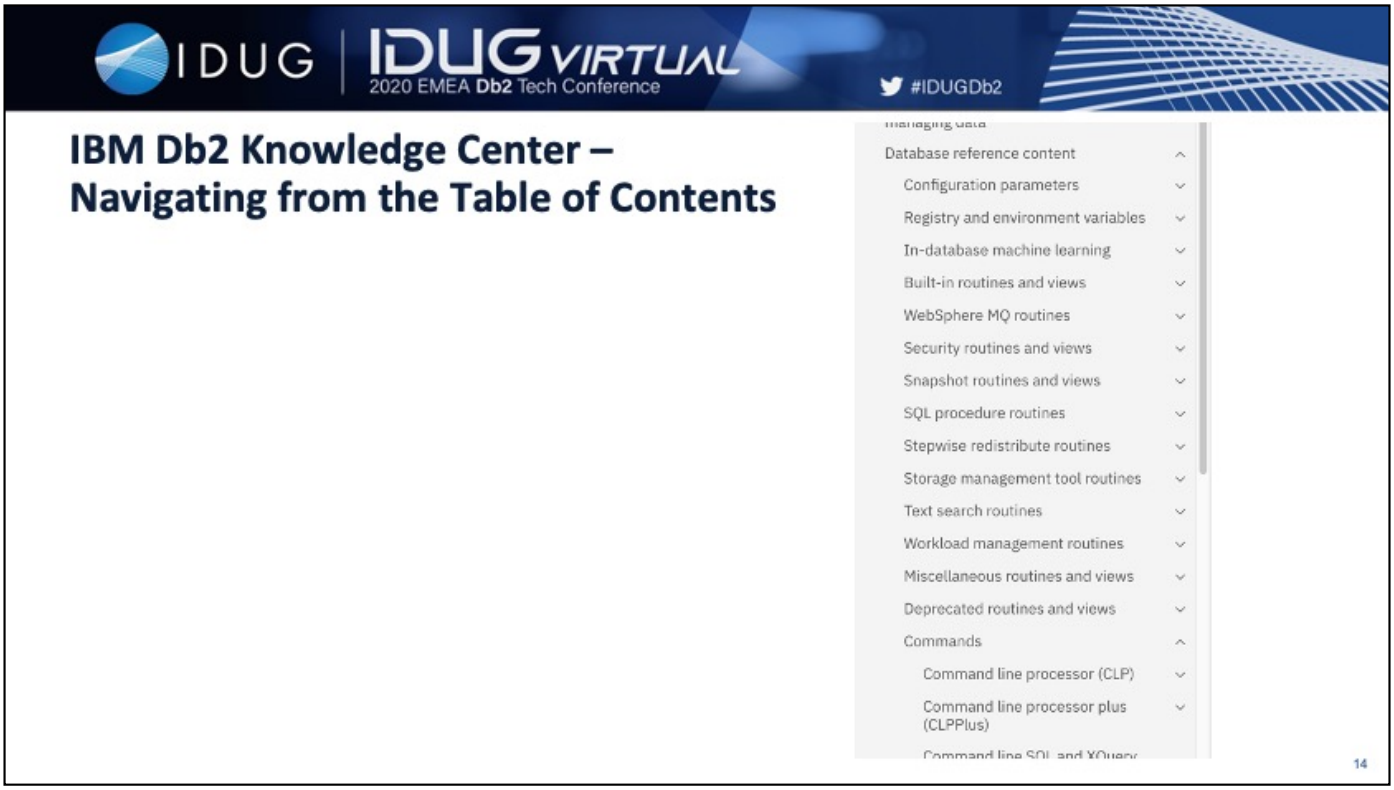

It is amazing how much is in the Db2 Knowledge Center. Even DBAs who have been working with Db2 for a while don't always realize the depth of the documentation. The "Database fundamentals" and "Database administration" sections have some more conceptual Db2 topics. The "Database Reference" is where DBAs spend most of their time, but the "SQL" section there has a lot of details for developers, too. Some of the structure here may be faster to navigate than search. If you're new to Db2, just perusing some of these topics may be interesting.

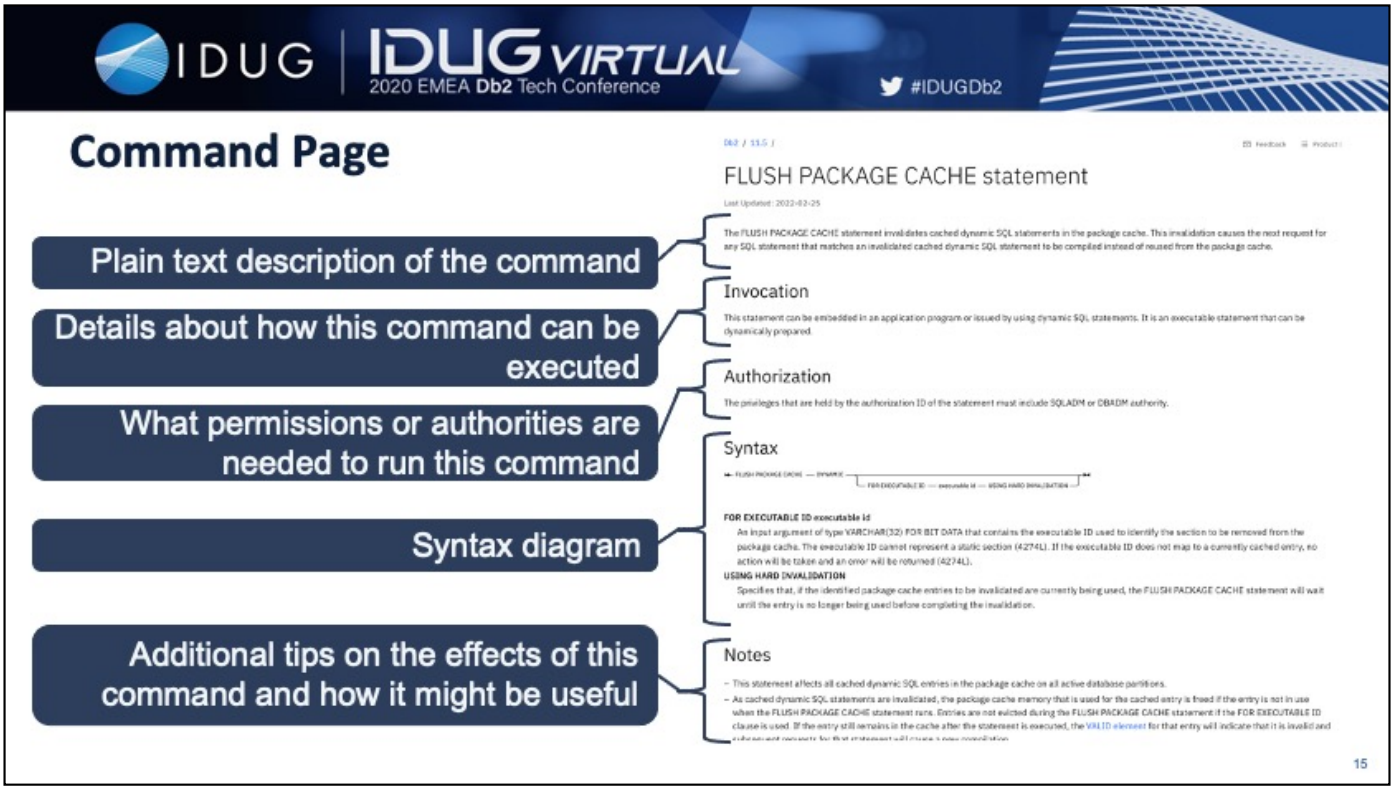

https://www.ibm.com/docs/en/db2/11.5?topic=statements-flush-package-cache

Each command is laid out in detail in the IBM Db2 Knowledge Center. Note the authorization section which specifies what authorities or permissions are needed to execute a command. The syntax diagram follows a specific format covered on future slides. At the bottom of this page, there are often examples of executing this command.

Sample Db2 Syntax Diagram

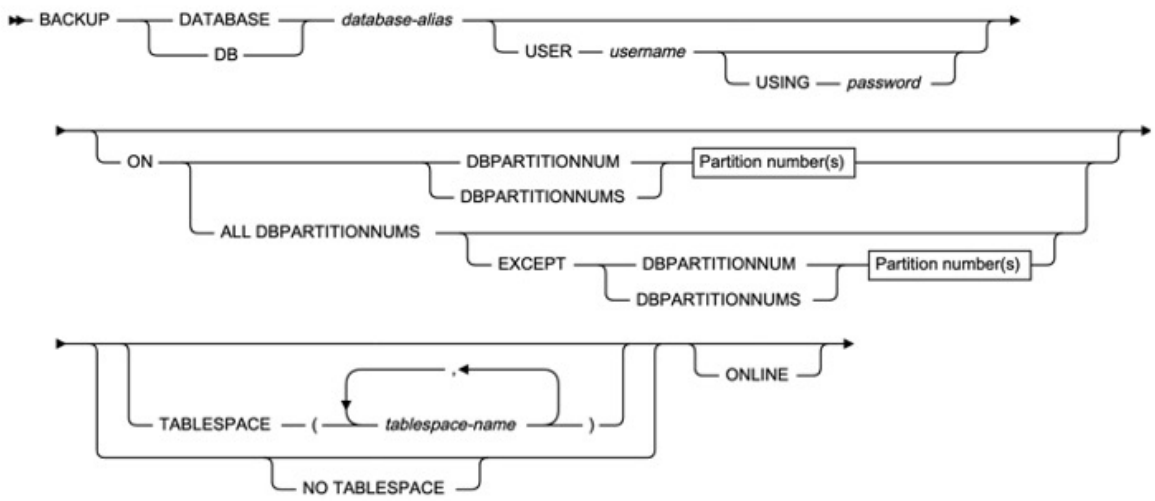

How to Read a Db2 Syntax Diagram (1|4)

- Read the syntax diagrams from left to right and top to bottom, following the path of the line.
- The ⏭ symbol indicates the beginning of a syntax diagram.
- The → symbol indicates that the syntax is continued on the next line.
- The ⏮ symbol indicates that the syntax is continued from the previous line.
- The ⏭ symbol indicates the end of a syntax diagram.
- A word or phrase in a box indicates a parameter block.

17

https://www.ibm.com/docs/en/db2/11.5?topic=sql-how-read-syntax-diagrams

How to Read a Db2 Syntax Diagram (2|4)

- If you *must* choose one of the items, one item of the stack appears on the main path.

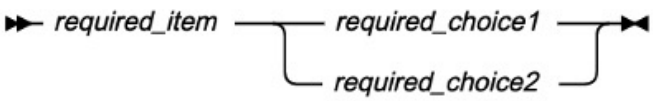

- If one of the items is the default, it will appear above the main path, and the remaining choices will be shown below.

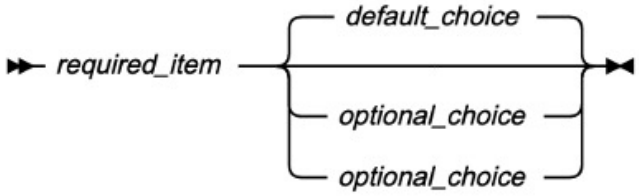

How to Read a Db2 Syntax Diagram (3|4)

- An arrow returning to the left, above the main line, indicates an item that can be repeated. In this case, repeated items must be separated by one or more blanks.

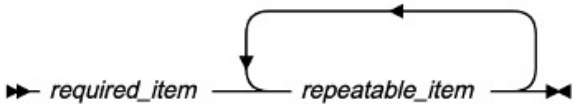

- If the repeat arrow contains a comma, you must separate repeated items with a comma.

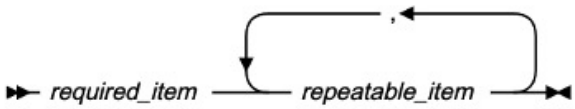

How to Read a Db2 Syntax Diagram (4|4)

- Sometimes a single variable represents a larger fragment of the syntax. For example, in the following diagram, the variable parameter-block represents the whole syntax fragment that is labeled **parameter-block**:

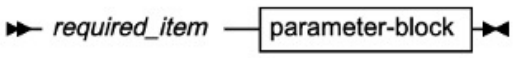

parameter-block

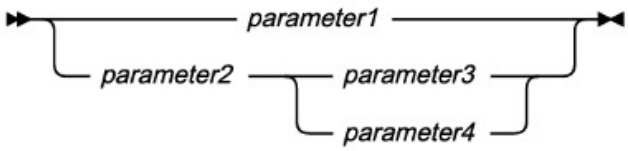

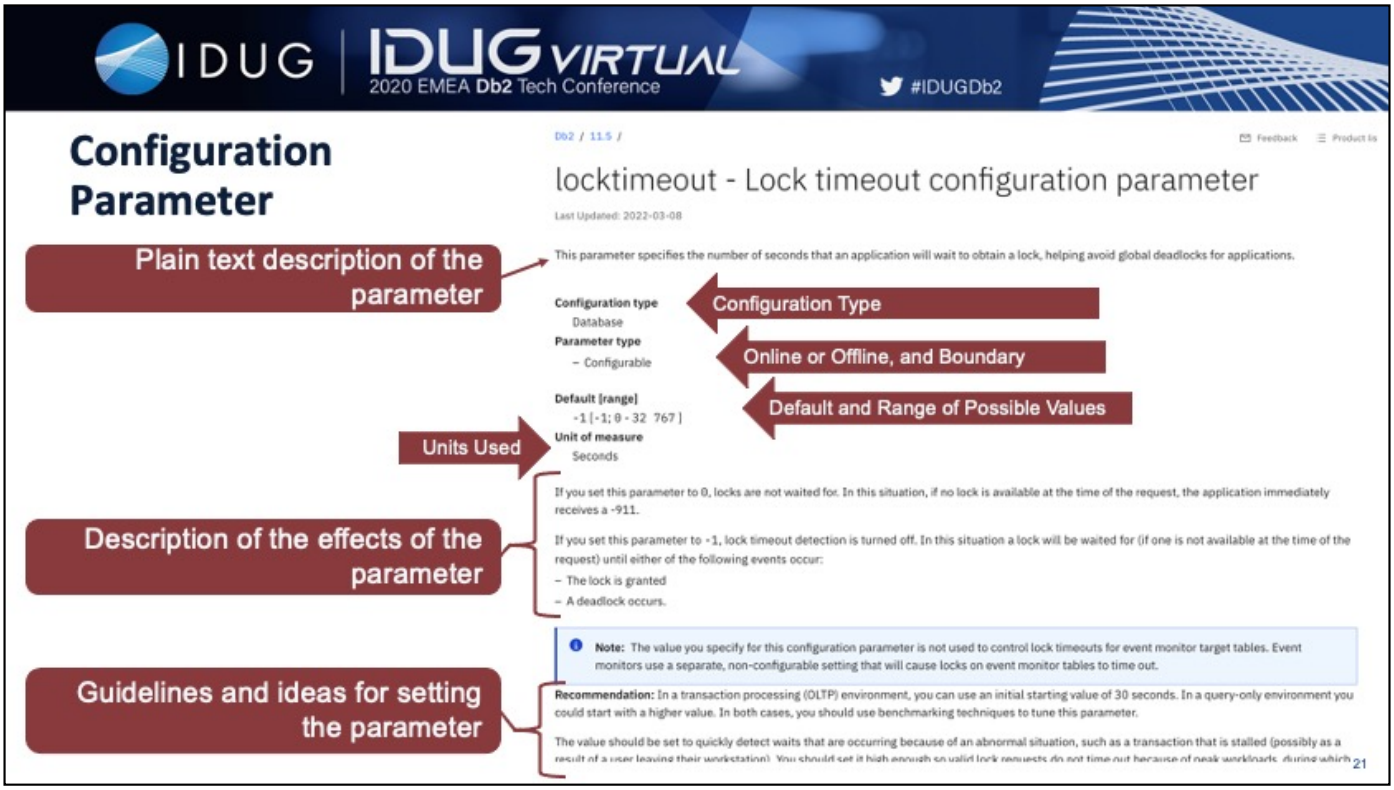

Each configuration parameter is described in great detail in the IBM Knowledge center. Note particularly the Parameter type, which tells you if a parameter can be changed online or not. The recommendation section often provides a bit of advice on things to consider when setting a parameter.

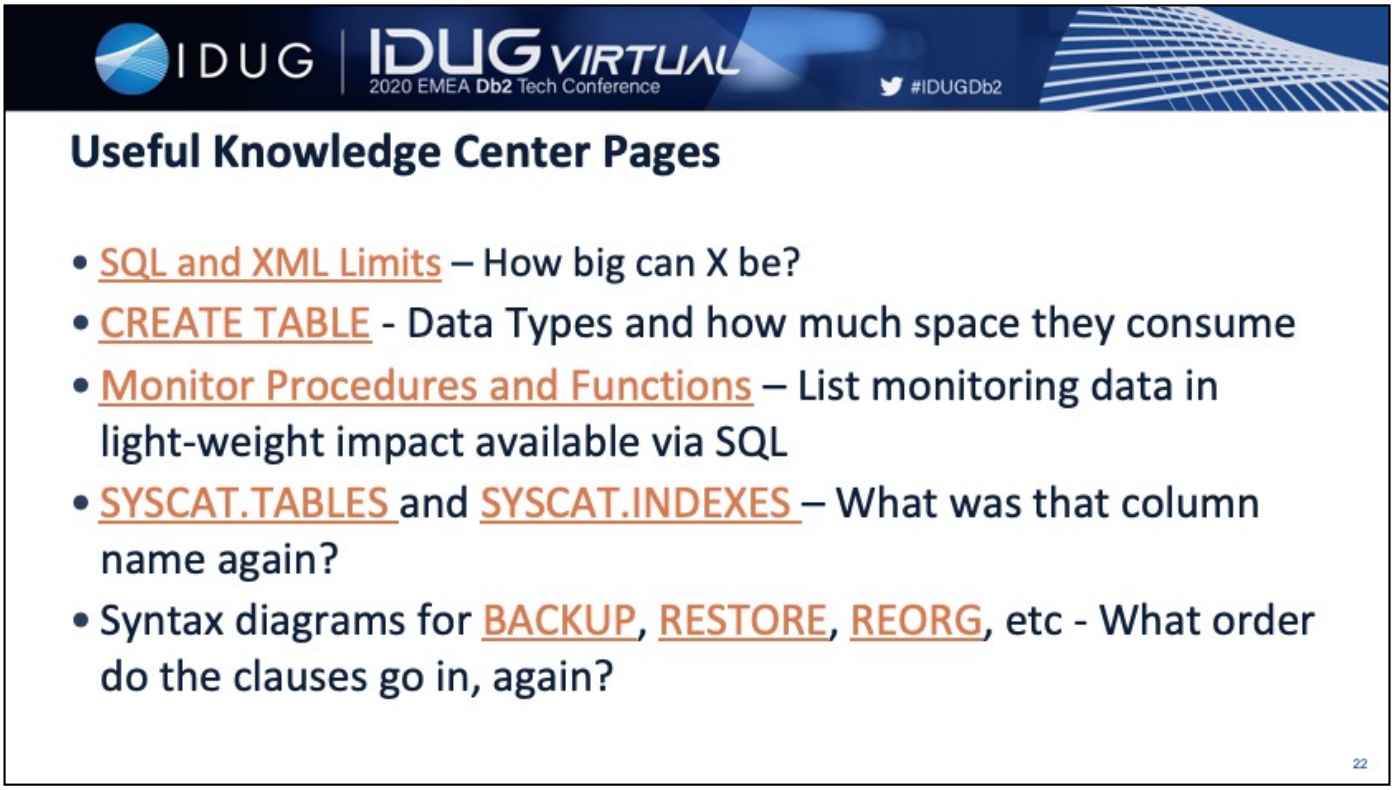

There are some pages that you find yourself visiting again and again. These are a few of them that I find myself returning to over and over again, even after 19 years of experience.

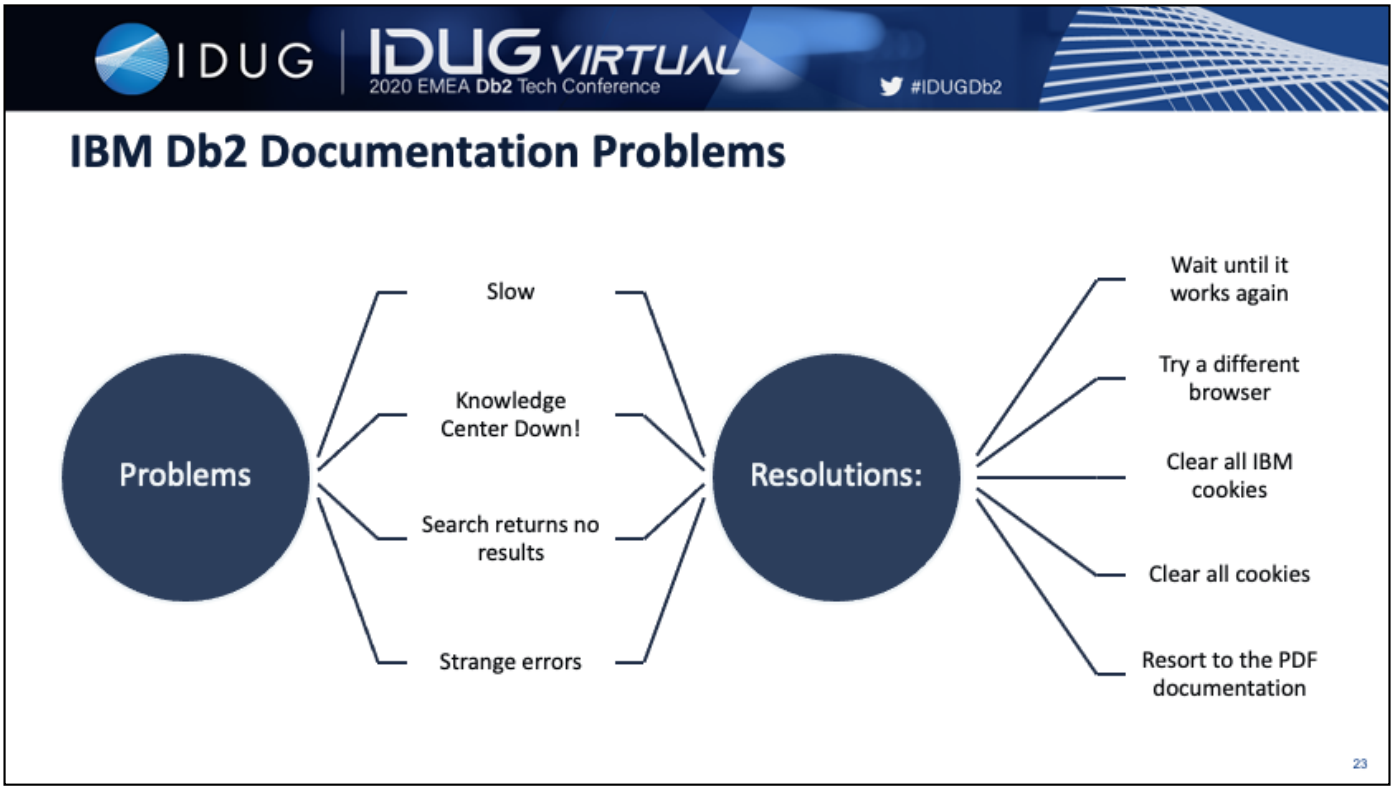

Some places do not allow access to the internet while connected to their VPN. Sometimes the IBM Db2 Knowledge Center is unavailable – even if this is for short periods of time in the middle of the night, it always seems to be when you're in the middle of a major problem or change. DBAs don't always sleep.

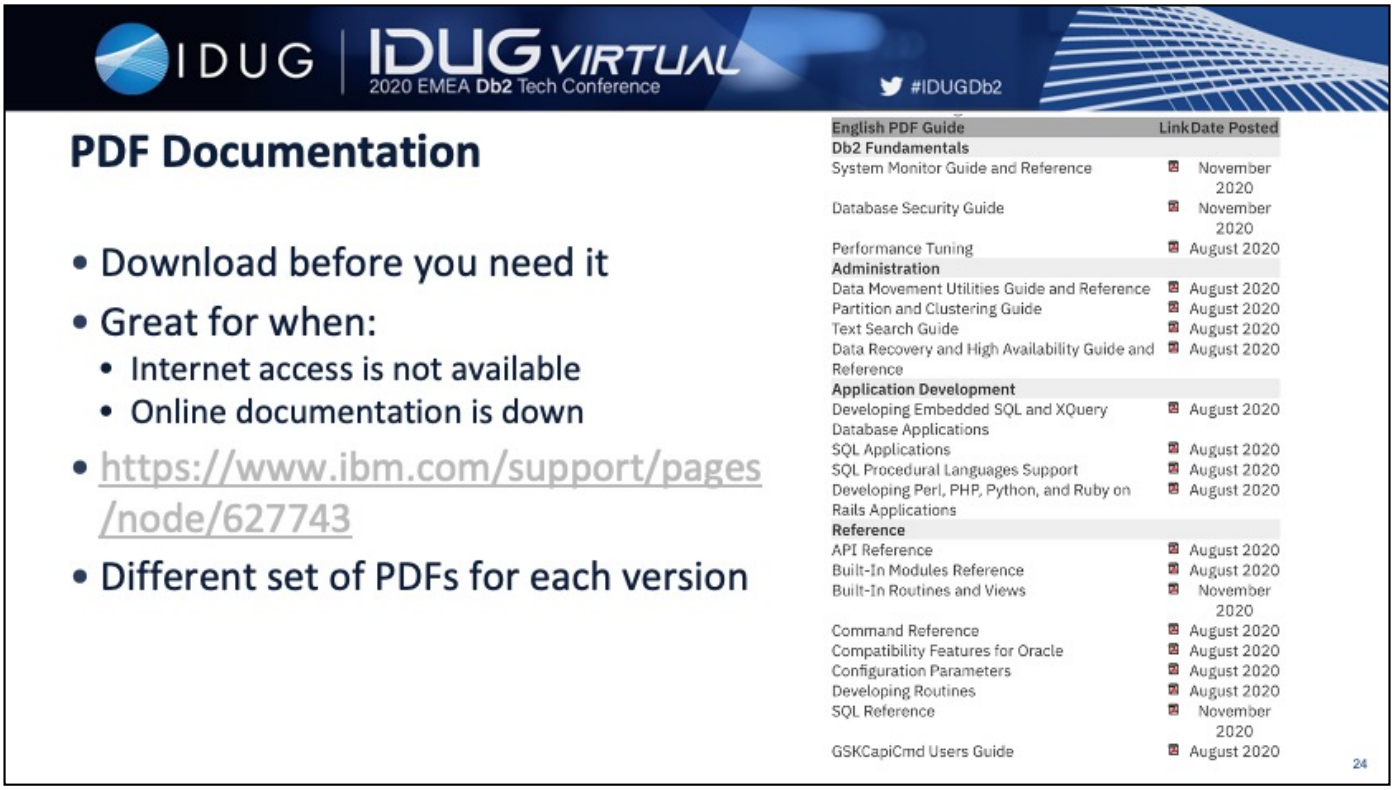

IBM was hesitant to continue offering PDF and offline forms of the Db2 documentation with 11.1. However, an RFE (Request for Enhancement) was opened, and it was the most popular one. IBM decided to start offering offline documentation again, in the form of PDFs. Some places do not allow access to the internet while connected to their VPN. Sometimes the IBM Db2 Knowledge Center is unavailable – even if this is for short periods of time in the middle of the night, it always seems to be when you're in the middle of a major problem or change. DBAs don't always sleep. It is great to have another option available. The key is to download these before you need them.

Db2 Error Messages

A Wealth of Information in a Small Package

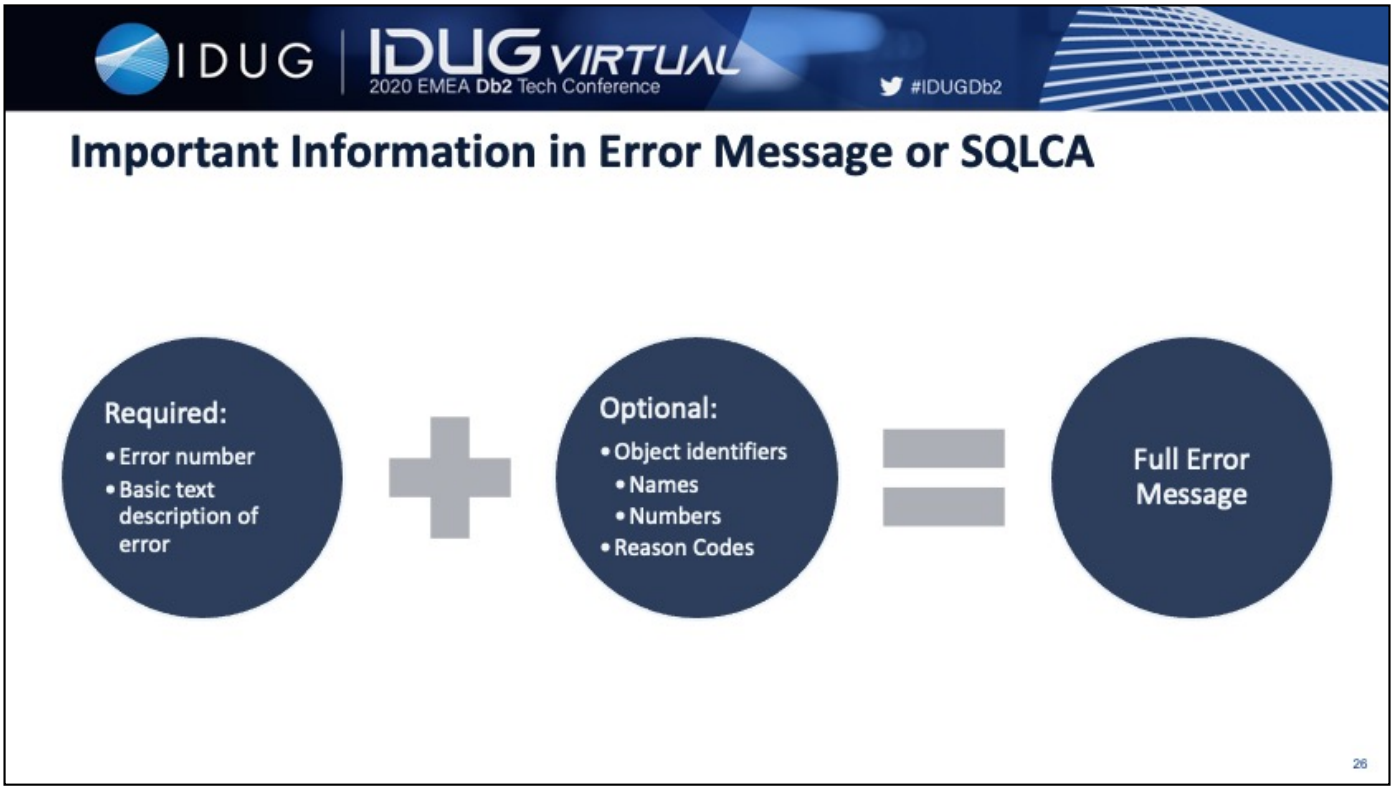

There is a wealth of information available in Db2 error messages.

Error Messages

Three letters, usually SQL, but sometimes others such as DSN or DIA

SQL 0911 N
XXX NNNN X

Usually 4 digits, sometimes 5

One Character indicating the intensity of the issue. Commonly:

- I – Informational
- W – Warning
- N – Statement Failure
- C – System Failure

- https://www.ibm.com/docs/en/db2/11.5?topic=content-messages

Getting Information on Errors – Command Line (1|2)

```
$ db2 ? SQL0911N

SQL0911N  The current transaction has been rolled back because of a
      deadlock or timeout. Reason code "<reason-code>".

Explanation:

The current unit of work was involved in an unresolved contention for
use of an object and had to be rolled back.

The reason codes are as follows:

2

         The transaction was rolled back due to a deadlock.

68

         The transaction was rolled back due to a lock timeout.

72
…
```

Getting Information on Errors – Command Line (2|2)

```
$ db2 ? SQL0551N


SQL0551N  The statement failed because the authorization ID does not
      have the required authorization or privilege to perform the
      operation. Authorization ID: "<authorization-ID>". Operation:
      "<operation>". Object: "<object-name>".

Explanation:

The operation could not be performed on the specified object. In
general, this message is returned because the authorization ID does not
have the required authority or privilege to perform the operation. In
some cases, it is returned for an object that does not allow the
operation even when the authorization ID has an administrative
authority.

…
```

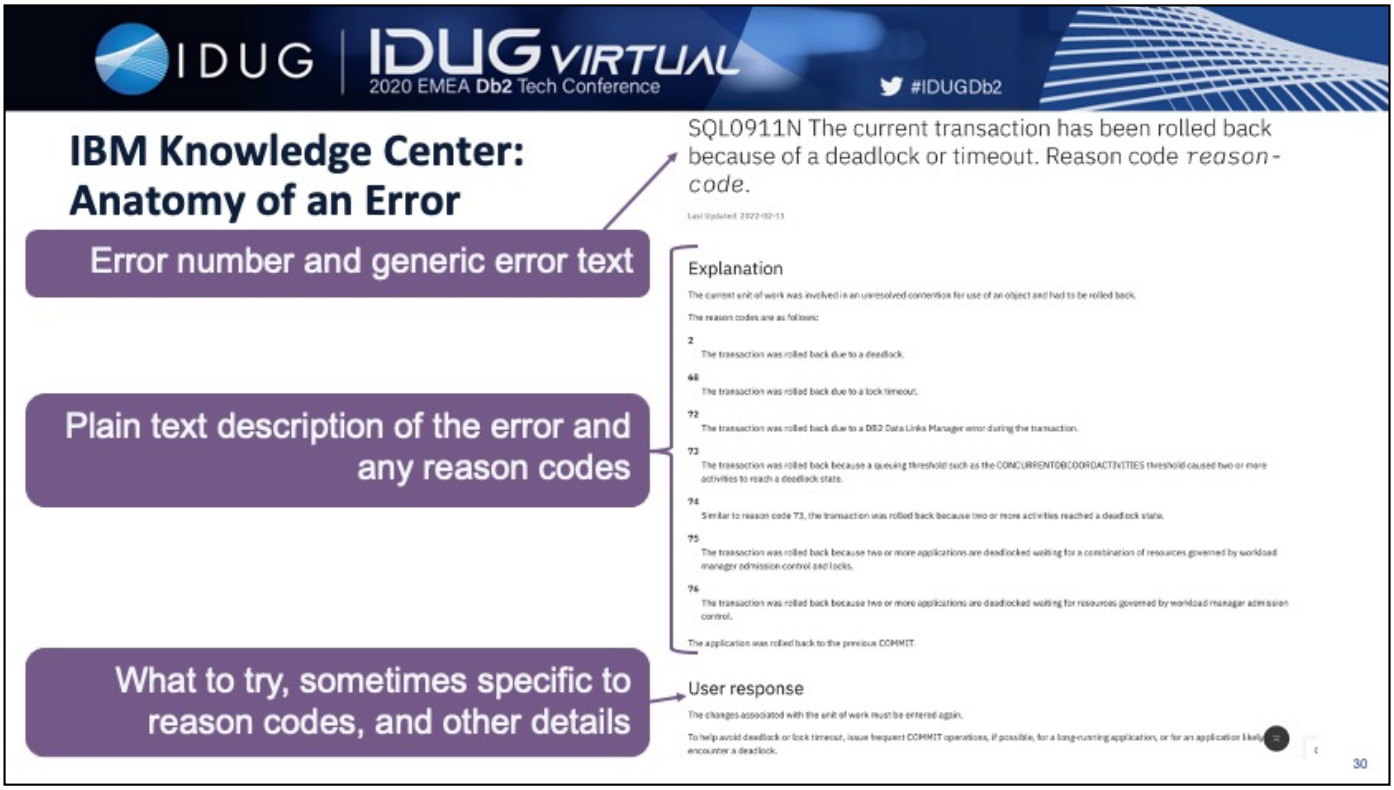

Each Section for an error on an error page has a well-defined structure. Knowing that structure can help you find the data you're looking for more quickly. The first part is the actual error message in its generic format. Next, in the explanation section, there is a plain text description of the error and descriptions of each reason code, if there are any. Finally, in the user response section, there are ideas for things to try to deal with the error, including suggestions specific to different reason codes.

Db2 Diagnostic Log

The Preeminent Location for Information About System-Level Issues

Where is the Db2 Diagnostic Log?

- The default Linux/UNIX location for the DB2 diagnostic log: $INSTHOME/sqllib/db2dump
- The default Windows (hidden!) location for the DB2 diagnostic log: C:\ProgramData\IBM\DB2\<db2 copy name>\DB2
- To get the path where the DB2 diagnostic log is stored:

```
$ db2 get dbm cfg |grep DIAGPATH
Diagnostic data directory path        (DIAGPATH) = /db2home/db2inst1/sqllib/db2dump
```

OR

```
db2 "select VALUE from SYSIBMADM.DBMCFG where NAME='diagpath'"
```

Db2 Diagnostic Log – Example of an Error

```
2017-01-18-13.37.20.290401-480 I381648021E487      LEVEL: Warning
PID     : 30986        TID : 140732262246144       PROC : db2sysc 0
INSTANCE: db2inst1     NODE : 000                  DB : SAMPLE
APPHDL  : 0-62447
HOSTNAME: prd-01.example.com
EDUID   : 11715        EDUNAME: db2agent (SAMPLE) 0
FUNCTION: DB2 UDB, bsu security, sqlexLogPluginMessage, probe:20
DATA #1 : String with size, 67 bytes
Password validation for user db2inst1 failed with rc = -2146500507
```

Timestamp

Diagnostic Level

Instance

Database Name

Db2 Error Log – Working with RCs

```
PID     : 641                  TID : 139928742651648 PROC : db2ckpwd
INSTANCE: db2inst1             NODE : 000
HOSTNAME: 90664639b1c3
EDUID   : 2                    EDUNAME: db2wdog [db2inst1]
FUNCTION: DB2 UDB, oper system services, sqloSpawnAndWaitForPasswordCheckExe, probe:130
MESSAGE : ZRC=0x800F006A=-2146500502=SQLO_BAD_USER "Bad User"
          DIA8117C Error with userid "".
DATA #1 : signed integer, 4 bytes
2
```

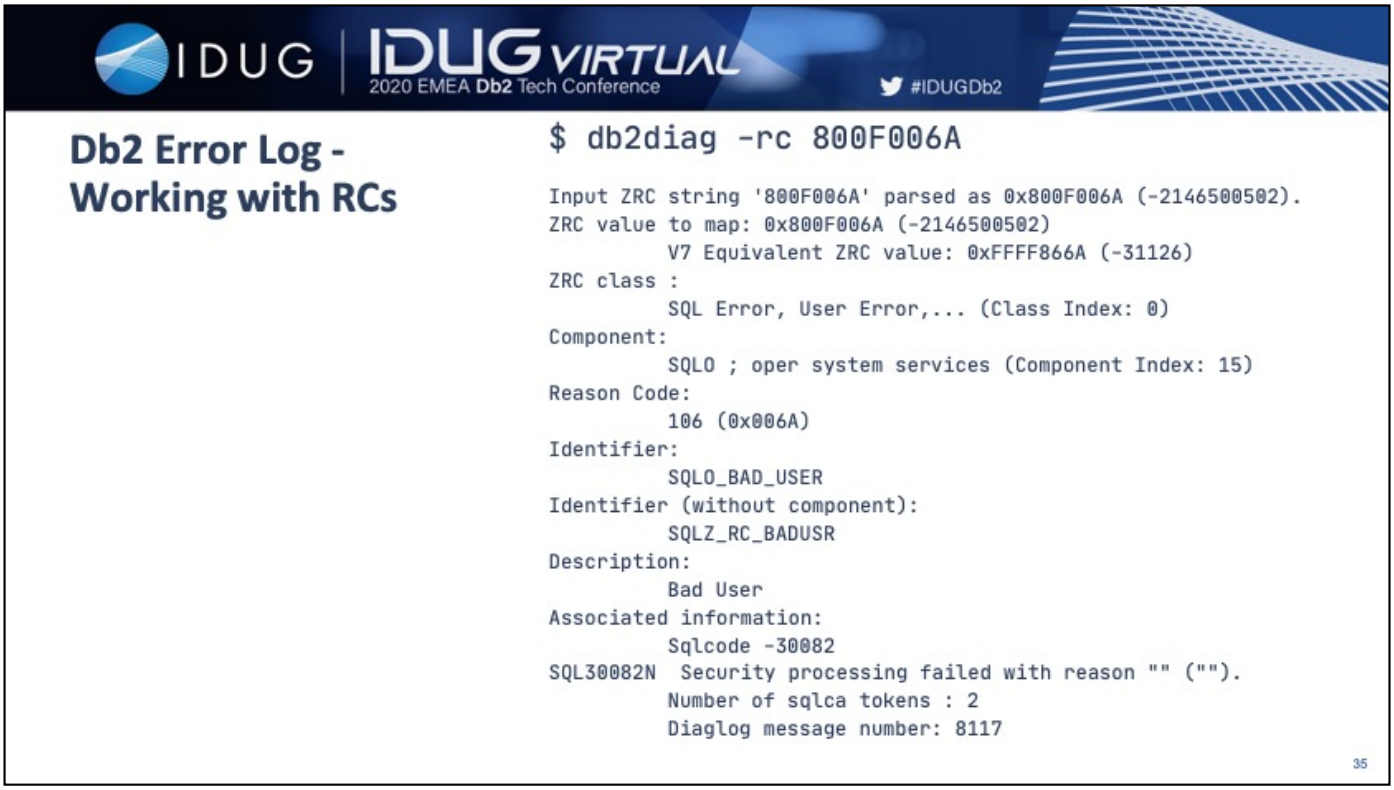

Some of the error messages in the Db2 diagnostic log can seem a bit cryptic. The db2diag system command can be used to help you translate odd error messages into something you can more easily look up. In this case, 800F006A translates to SQL30082N.

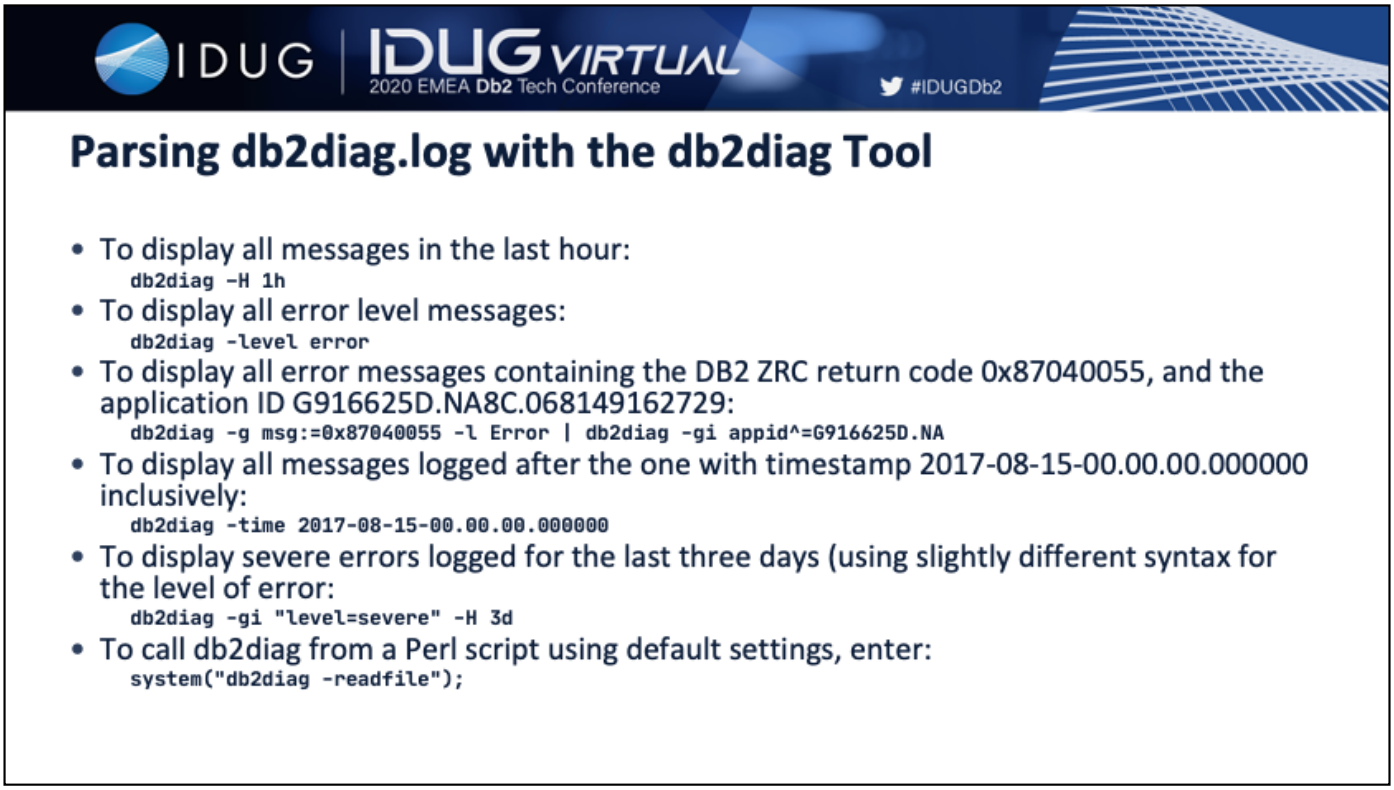

https://www.ibm.com/support/knowledgecenter/SSEPGG_11.1.0/com.ibm.db2.luw.admin.cmd.doc/doc/r0011728.html

https://www.ibm.com/support/knowledgecenter/SSEPGG_11.1.0/com.ibm.db2.luw.admin.trb.doc/doc/c0020701.html

The db2diag tool is excellent for parsing the Db2 diagnostic log and sorting through the information available there. You could also choose to query what's in the Db2 diagnostic log using the PD_GET_LOG_MSGS table function.

Searching and Formatting using db2diag

```
$ db2diag -e 4500 -fmt %dataobject
DATA #1 :
Package Cache Overflow
memory needed              : 10592
current used size (OSS)    : 134106751
maximum cache size (APM)   : 130191196
maximum logical size (OSS): 134150417
maximum used size (OSS)    : 215154688
owned size (OSS)           : 137101312
number of overflows        : 9400
```

Monitoring Tools

Many Methods to Understand what is Happening

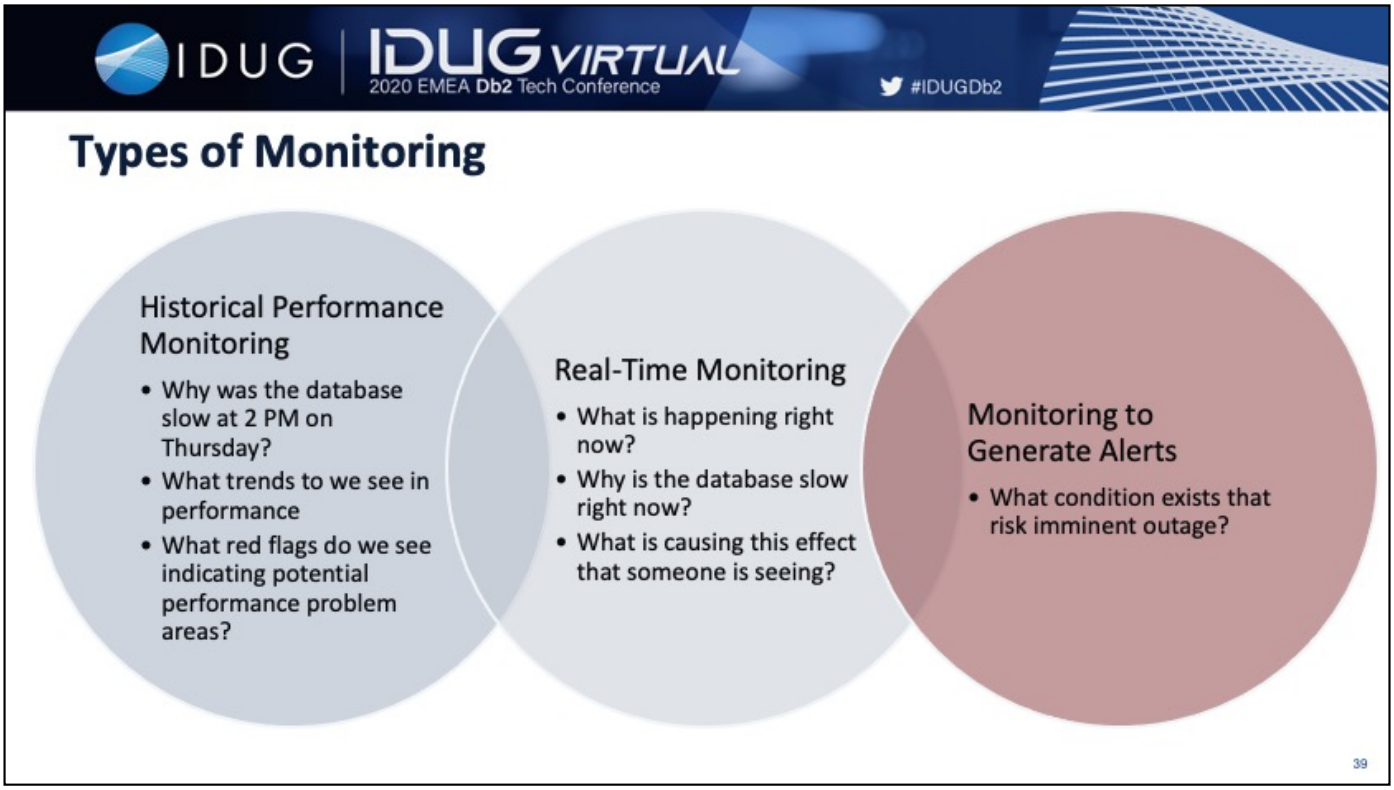

Monitoring is a general term that may include a lot of different potential reasons and actions. In the case of this presentation, we're talking about real-time monitoring to see what is going on in the database at the moment and to troubleshoot problems.

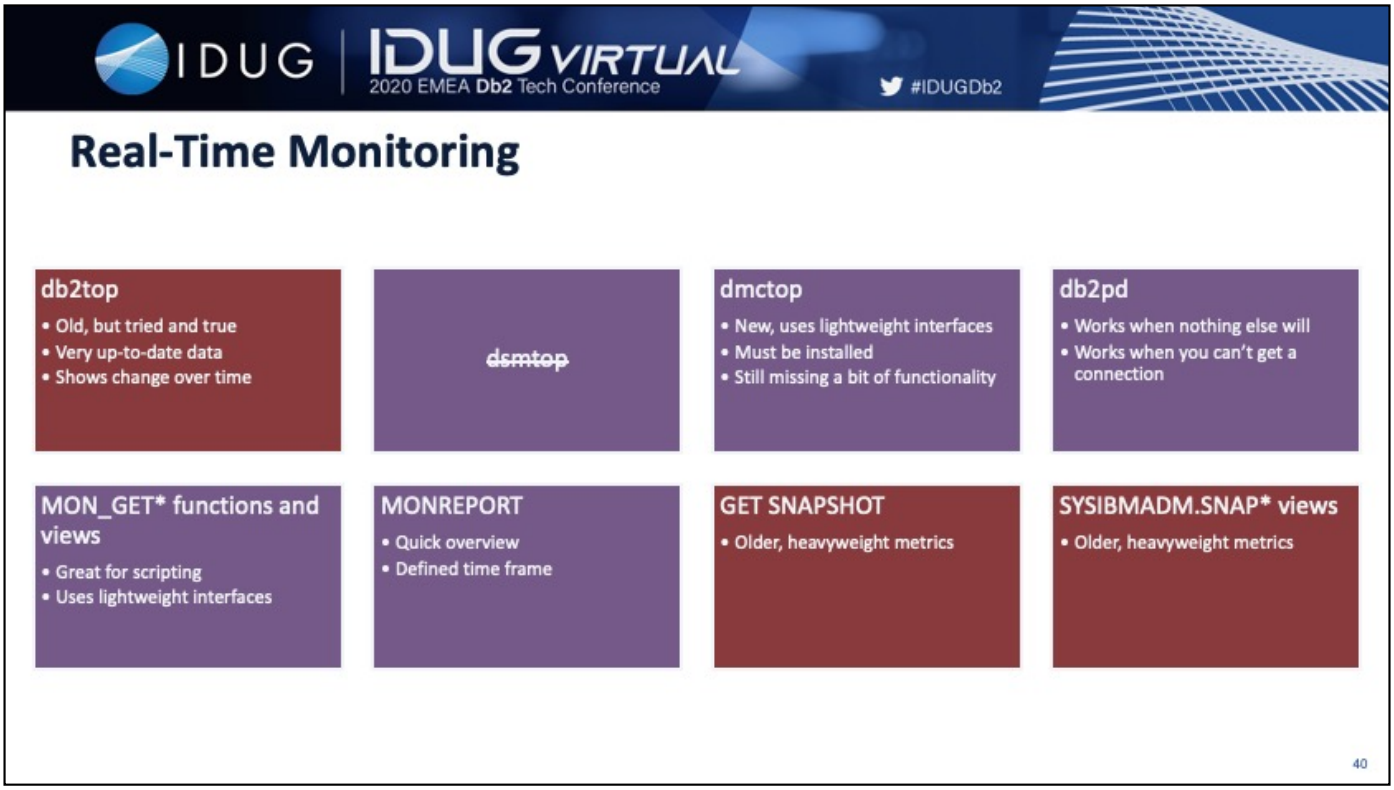

There are a lot of choices in real-time monitoring. Which you choose depends on your purposes, constraints, and preferences.

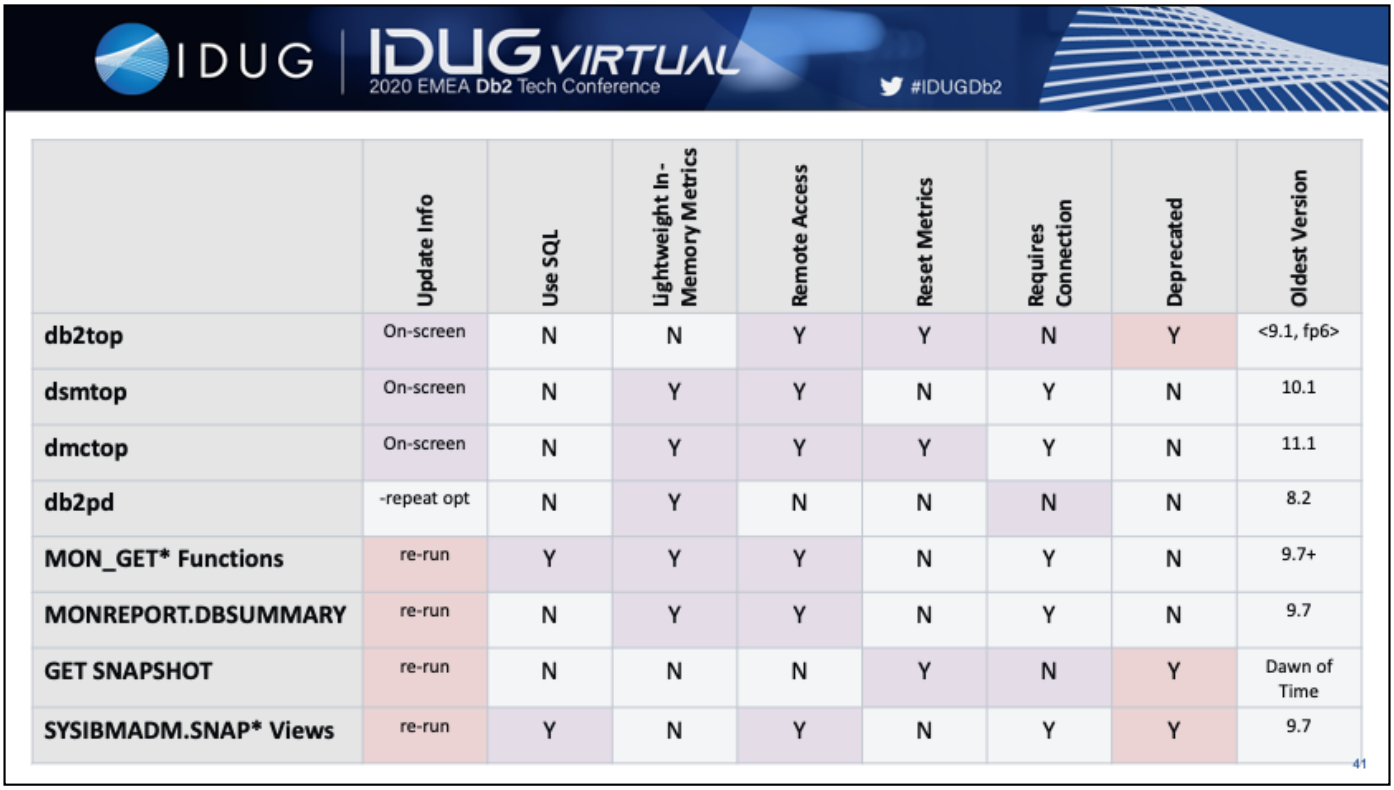

	Update Info	Use SQL	Lightweight In-Memory Metrics	Remote Access	Reset Metrics	Requires Connection	Deprecated	Oldest Version
db2top	On-screen	N	N	Y	Y	N	Y	<9.1, fp6>
dsmtop	On-screen	N	Y	Y	N	Y	N	10.1
dmctop	On-screen	N	Y	Y	Y	Y	N	11.1
db2pd	-repeat opt	N	Y	N	N	N	N	8.2
MON_GET* Functions	re-run	Y	Y	Y	N	Y	N	9.7+
MONREPORT.DBSUMMARY	re-run	N	Y	Y	N	Y	N	9.7
GET SNAPSHOT	re-run	N	N	N	Y	N	Y	Dawn of Time
SYSIBMADM.SNAP* Views	re-run	Y	N	Y	N	Y	Y	9.7

41

This is a summary of the major options for real-time monitoring and the areas in which they succeed or fail.

db2top

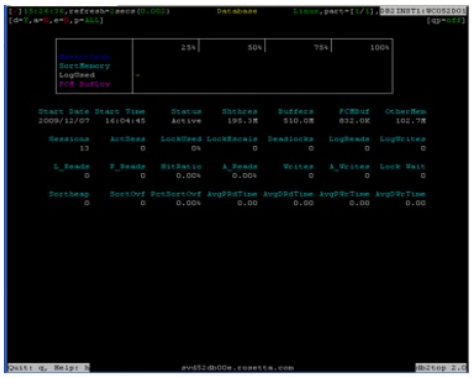
25%
50%
75%
100%
SortMemory
LogUsed
FCM Outflow
Start Date Start Time Status ShrUsed BufferX FCMBuf OtherMem
2009/12/07 16:04:45 Active 195.3M 510.0M 832.0K 102.7M
Sessions ActSess LockWait LockEscals Deadlocks LogReads LogWrites
13 0 0% 0 0 0 0
L_Reads P_Reads HitRatio A_Reads Writes A_Writes Lock Wait
0 0 0.00% 0.00% 0 0 0
SortAmap SortOvf PctSortOvf AvgPRdTime AvgSRdTime AvgPWrTime AvgSWrTime
0 0 0.00% 0.00 0.00 0.00 0.00
Quit: q, Help: h
avd52db00e.rosetta.com
db2top 2.0

db2top – Help Screen

- Pressing the h key in db2top will get you the help screen, which has a list of all the nifty options

db2top – Interval and Cumulative

- Default interval is 2 seconds
- A different interval can be specified with the –i option when db2top is started
- db2top only reports the data for the last <interval> seconds (current view)
- db2top can also report cumulative values since the last db restart or reset (cumulative view)
- Pressing k will toggle between the current view and the cumulative view
- Pressing R will reset the monitor data – for that session only

db2top – Lock Chaining

- Use U to get to the lock screen:

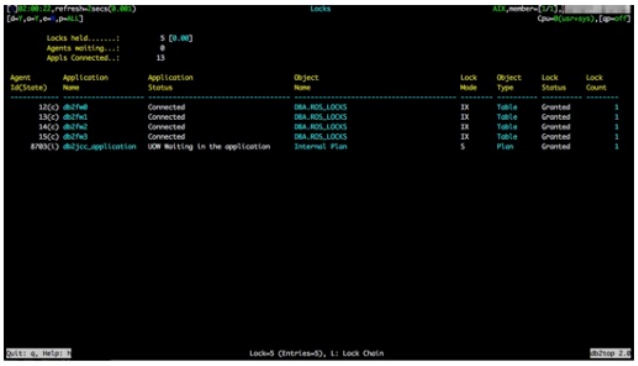

db2top – Lock Chaining

- If locks are waiting, db2top can quickly tell which connection is waiting on which connection

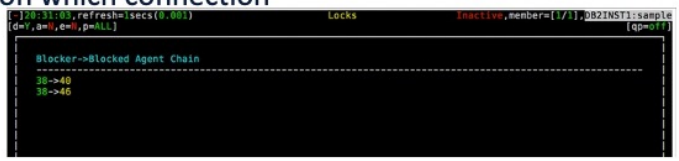

db2top - Scrolling

- Scrolling up and down is not possible
- Many screens have more columns that will fit on the screen
- Left and right scrolling can be done with > and <

dmctop

- Brand new as of Summer 2020
  - Is included in Db2, though often older versions
- Easy to install by downloading a file
- Can connect remotely, with or without TLS/SSL
- Does not yet have full feature parity with db2top
- Better than dsmtop

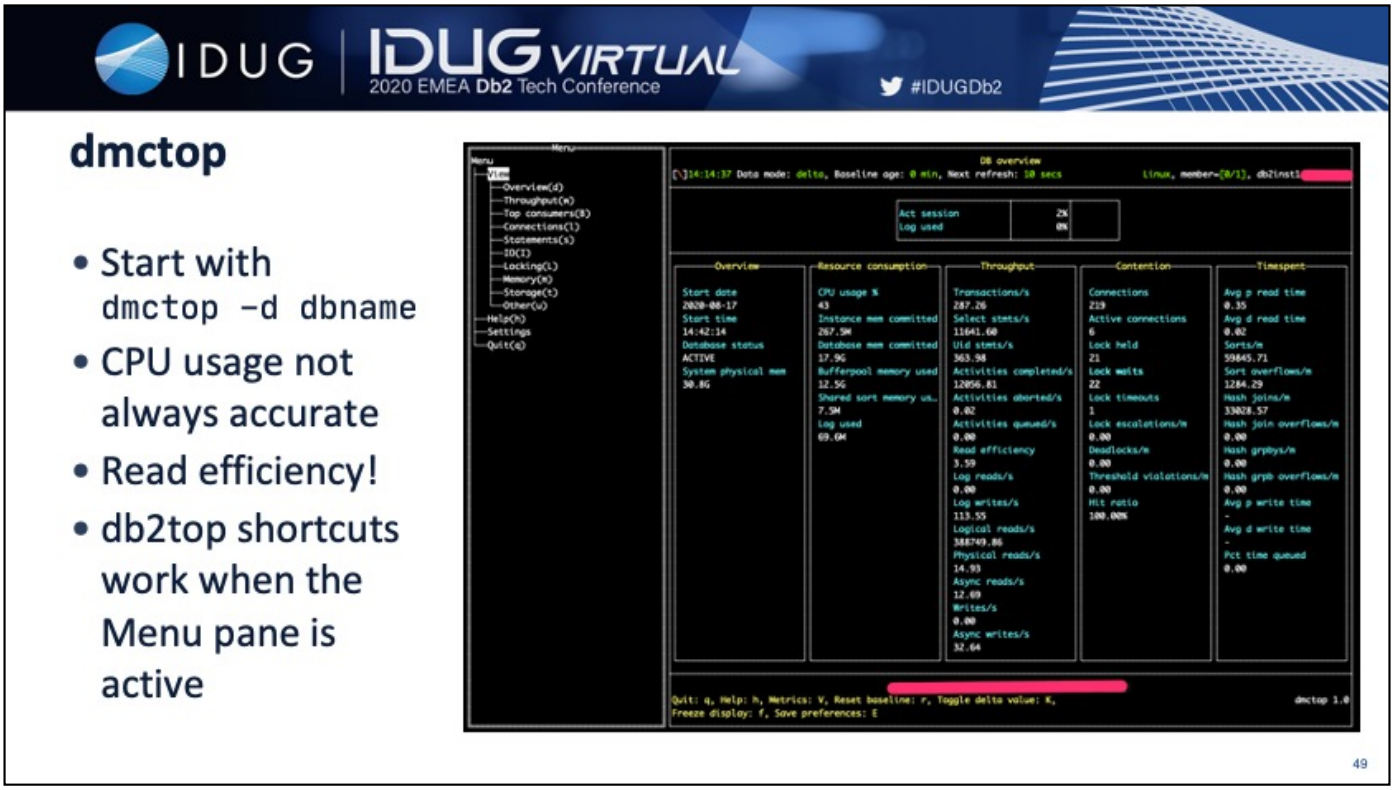

dmctop has a lot of great information on the database overview screen. A few things to note are that I and some friends have seen the CPU usage metric not be very accurate. Personally, I love the read efficiency metric, which was not available in db2top. On this screen, there are two panes – the left menu pane and the right output pane. Keys pressed currently have a different meaning depending on which pane is active, so if keys aren't doing what you think they should, try hitting ESC to put the focus back on the menu pane.

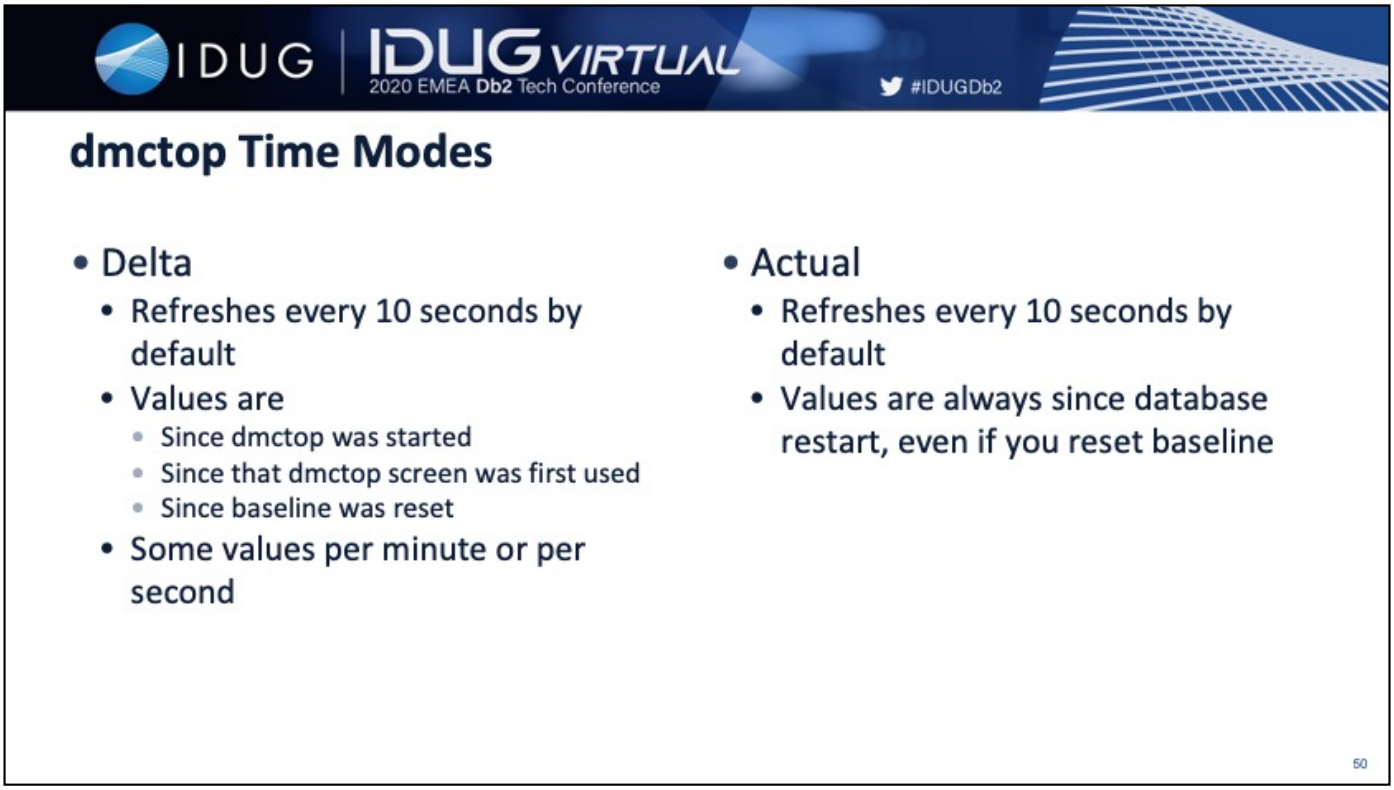

There are two different time modes for dmctop. The Delta mode allows values to be reset.

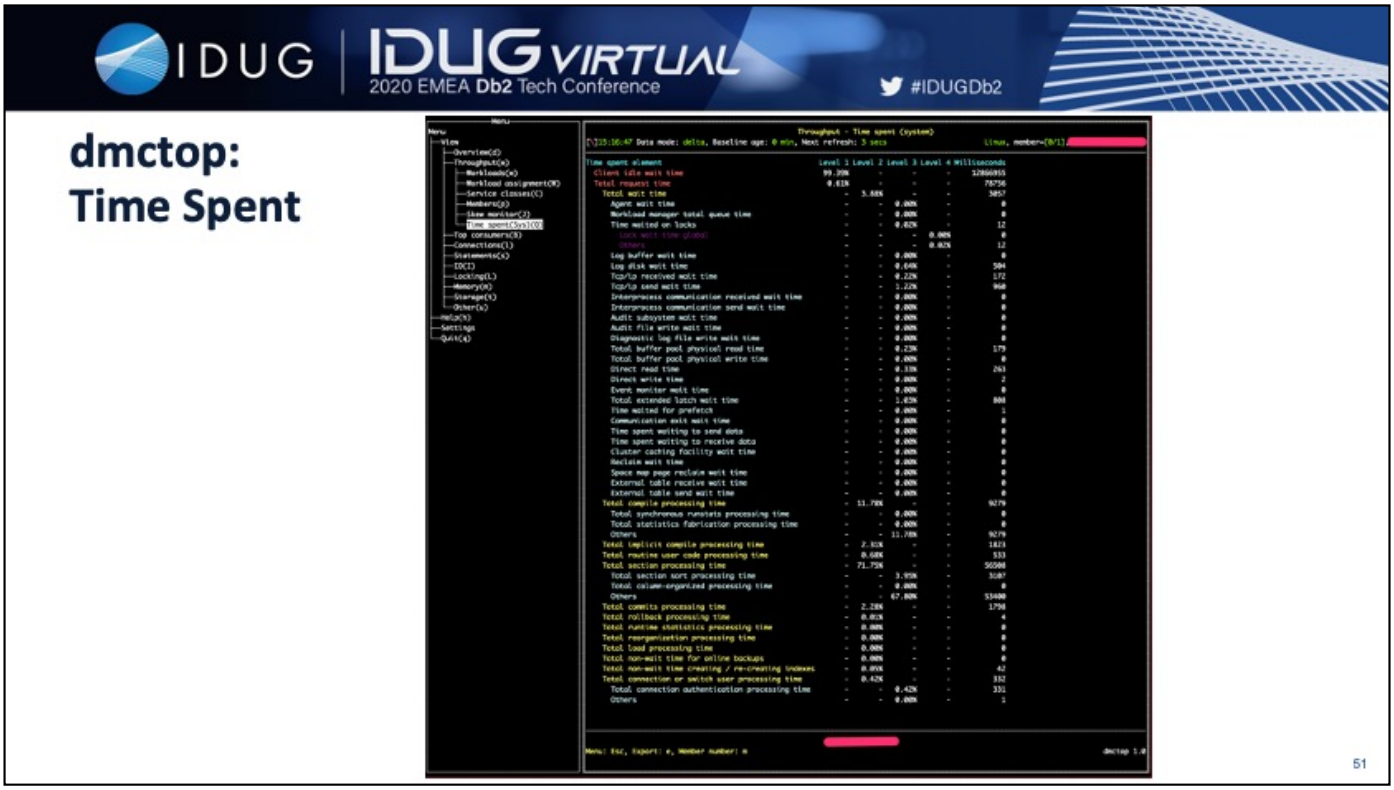

My favorite screen in dmctop is the time spent screen. It provides a nice overview of where Db2 is spending time that allows you to drill down into problem areas or quickly see where a problem might be lurking.

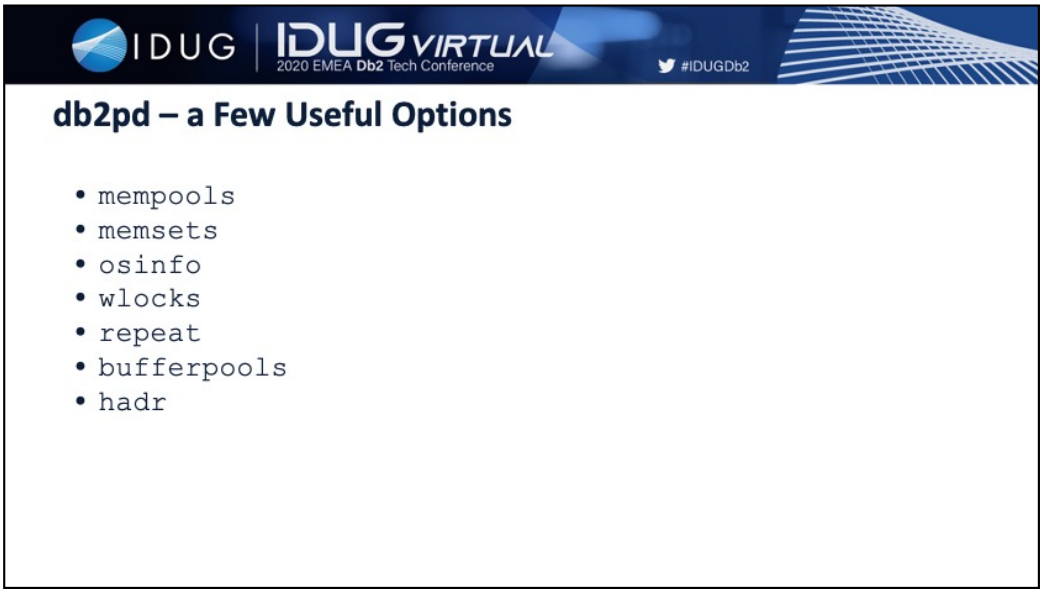

Db2pd has a few advantages. First, it can be used without a database connection, meaning that if the database is hung, you can still get information. It is also a very lightweight interface, so there is very little risk of it adding load to a busy server. I find it most useful in a few areas, listed here. I often also find the –repeat option useful to have it repeat the command every minute or so – allowing me to track the progress of something like HADR.

MONREPORT.DBSUMMARY

- Good summary of data
- Snapshot-like format that uses more recent monitoring infrastructure
- Snipets…

```
Part 1 - System performance

 Work volume and throughput
  ------------------------------------------------------------------------------
                                Per second             Total
                                ---------------------  ----------------------
  TOTAL_APP_COMMITS             0                      3
  ACT_COMPLETED_TOTAL           1                      11
  APP_RQSTS_COMPLETED_TOTAL     2                      27

  TOTAL_CPU_TIME                = 44364
  TOTAL_CPU_TIME per request    = 1643

  Row processing
    ROWS_READ/ROWS_RETURNED     = 2 (33/16)
    ROWS_MODIFIED               = 0
```

MONREPORT.DBSUMMARY Snippets

```
Buffer pool
  ------------------------------------------------------
  Buffer pool hit ratios

  Type             Ratio            Formula
  ---------------  ---------------  -------------------
  Data             100              (1-(0+0-0)/(37+0))
  Index            100              (1-(0+0-0)/(34+0))
  XDA              0                (1-(0+0-0)/(0+0))
  COL              0                (1-(0+0-0)/(0+0))
  LBP Data         100              (37-0)/(37+0)
  LBP Index        0                (0-0)/(34+0)
  LBP XDA          0                (0-0)/(0+0)
  LBP COL          0                (0-0)/(0+0)
  GBP Data         0                (0 - 0)/0
  GBP Index        0                (0 - 0)/0
  GBP XDA          0                (0 - 0)/0
  GBP COL          0                (0 - 0)/0
```

MONREPORT.DBSUMMARY Snippets

```
                        Per activity                    Total
                        ------------------------------  ----------------------
LOCK_WAIT_TIME          0                               0
LOCK_WAITS              0                               0
LOCK_TIMEOUTS           0                               0
DEADLOCKS               0                               0
LOCK_ESCALS             0                               0

<SNIP>

Sort
------------------------------------------------------------------------------
TOTAL_SORTS                         = 4
SORT_OVERFLOWS                      = 0
POST_THRESHOLD_SORTS                = 0
POST_SHRTHRESHOLD_SORTS             = 0
```

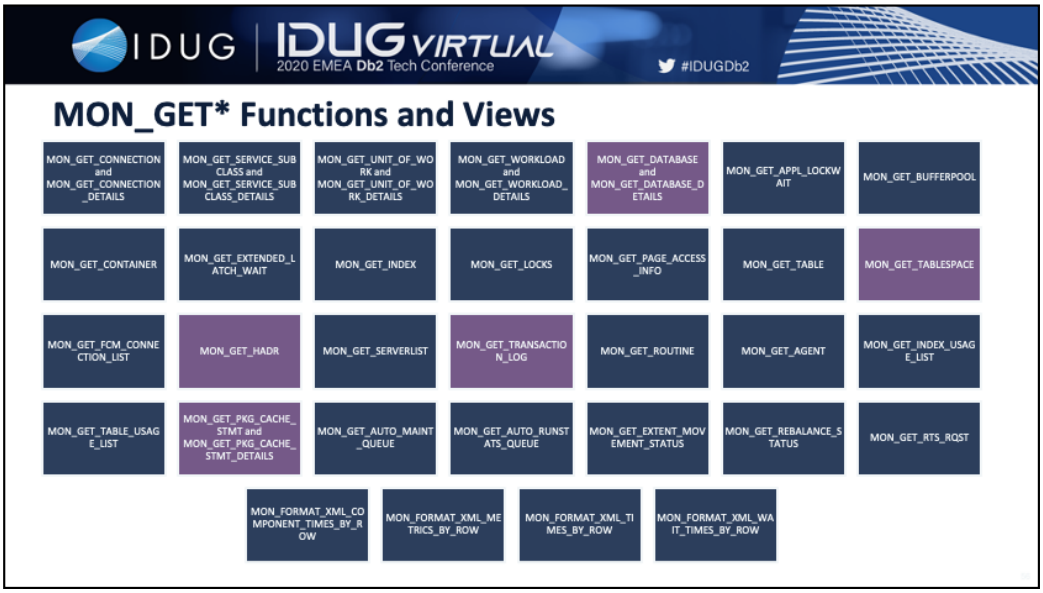

SQL interface to lightweight in-memory monitoring

Querying MON_GET* Table Functions

```
select * from table(mon_get_database(-2))
```

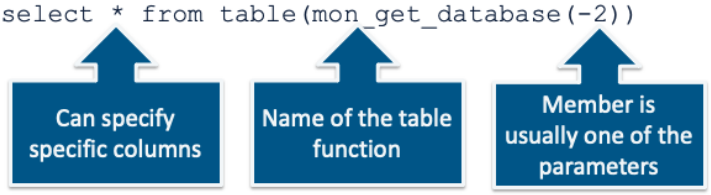

- Variable number of parameters depending on the object
- For example, mon_get_table requires tabschema and tabname
- Can use '' or NULL for many of the parameters to return data on all objects

Interacting with IBM Support

- Open a ticket via the support portal or phone
- Front-line support reps are sometimes good, but sometimes do not seem to know much about Db2.
- IBM will ask for additional information. Respond quickly.
- Usually, you must push IBM to keep working on an issue by updating the ticket regularly
- If a ticket is not getting needed attention, escalate to a duty manager (call in and ask for this)
- If you did not get reasonable answers or responses, close a ticket as "non-sat" – this gets attention, but often no further action.

58

https://datageek.blog/en/2012/12/04/getting-the-most-out-of-db2-support/

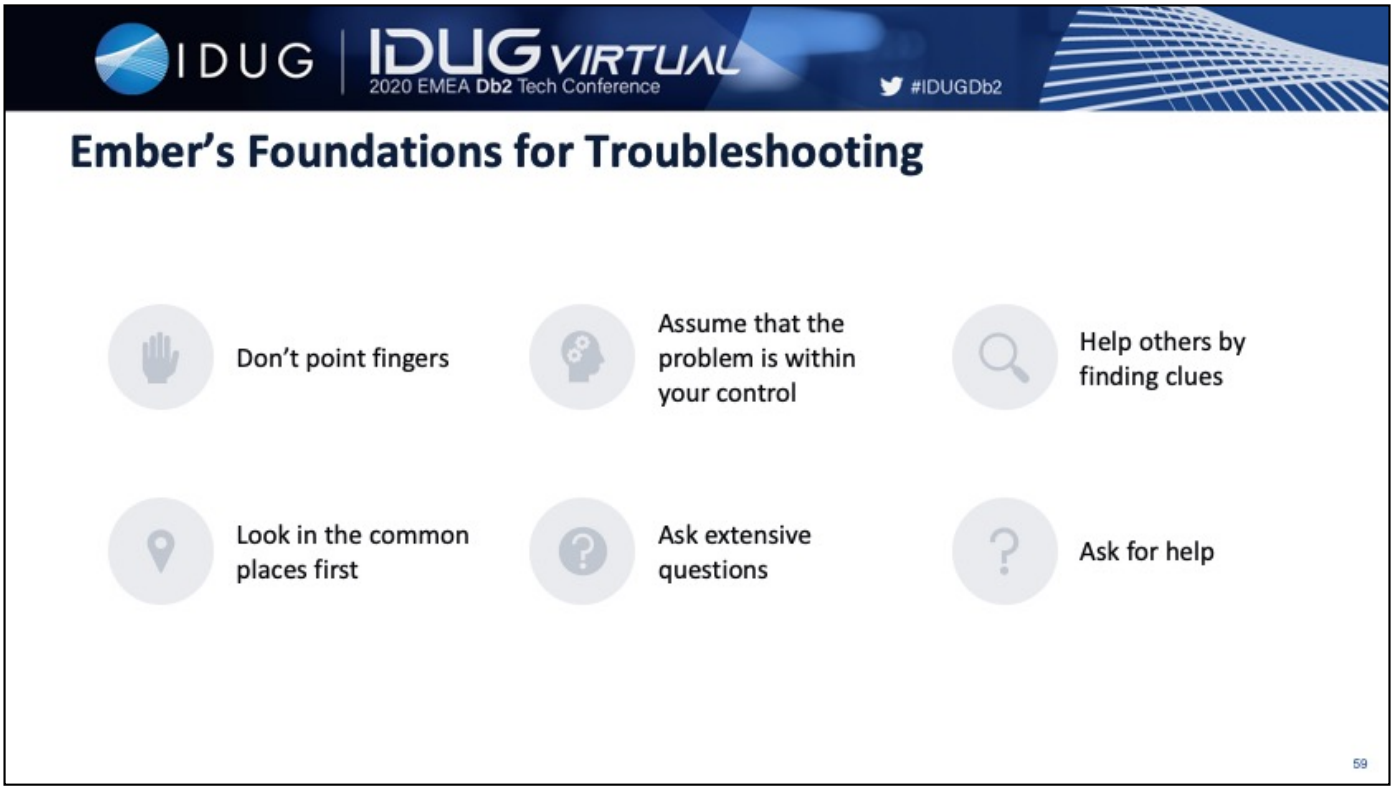

When troubleshooting you're rarely able to immediately say "this is a database problem", "this is an application problem", or "this is a network problem". The key is that if you're working with experts in other areas or even if you're trying to cover all the areas yourself, don't rule things out prematurely. If someone calls me in for a problem, I do my best not to just say "That's not the database.", even if it clearly isn't. Even if it isn't the database, often there are clues that you can find within the database that may help others locate the problem in their areas. My approach is to start from "This is in the database, but where?", and then once I have significant proof it isn't, provide others with the details that will help them in their own troubleshooting efforts.

Ember has 19+ years of experience with Db2 on Linux, Unix, and Windows platforms. She is the founder and principal author of the popular datageek.blog technical blog where she educates herself and others through example and case study. Ember is an IBM Gold Consultant and IBM Champion in Information Management.