

db2top - Db2 Performance Monitor

```
db2top [-d dbname] [-n nodename] [-u username]
[-p password] [-V schema] [-i interval]
[-P <partition>] [-b option] [-a] [-B] [-k] [-R]
[-x] [-f file </HH:MM:SS> <+offset> ]
[-D delimiter] [-C <option>] [-m duration]
[-o out file]
db2top [-H host ] [-N port ] [-h]
```

Description

db2top is a monitor for DB2 UDB specifically designed for DPF environments. It provides a unified, 'single system view' of a multi-partition DB2 database that can be used to quickly identify either partition-specific or global problems in the system. The user interface is character-based and built using the curses library and is similar in nature to many UNIX monitoring utilities. db2top uses the DB2 snapshot monitoring APIs to retrieve data. It uses both global as well as partition-specific monitoring information to provide aggregation as well as quick drill-down capabilities. It can be run interactively or in batch mode. Main features can be selected by an interactive command or by specifying the option in the personal or system-wide configuration file. See below .db2toprc for more information.

Command Options

n specifies the node to attach to.
d specifies the database to monitor.
u specifies the DB2 username used to access the database.
p specifies the DB2 password.
V specifies the default schema used in explains.
i specifies the delay between screen updates.
k specifies whether to display actual or delta values.
R Reset snapshot at startup.
P <no> Snapshot issued on current or partition number.
x specifies whether to display additional counters on session and appl screens (might run slower on session).
b tells db2top to run in background mode.
a specifies only active queries will be displayed.
B enable bold for active objects.

C runs db2top in snapshot collector mode, raw snapshot data is saved in <db2snap-Machine>.bin by default (unless -f is specified).
f <file> </HH:MM:SS> <+offset>. Run db2top in replay mode when snapshot data has been previously collected in <file>. offset will jump to a certain point in time in the file. It can be expressed in seconds (+10s), minutes (+10m) or hours (+10h). /HH:MM:SS will skip entries until the specified point in time.
m Will limit duration of db2top in minutes for -b and -C.
o <outfile> Outfile for -b option.
D <delimiter>. Separator for -b option.
h short help.

When using the -b option, db2top will display information in CSV format. db2top can be run in background mode in combination with reading snapshot data from collection file, in this case, use the -f <file> option.

Valid suboptions for -b and -C are:

d : database
1 : sessions
t : tablespaces
b : bufferpools
T : Tables
D : Dynamic SQL
s : Statements
U : Locks
u : Utilities
F : Federation
m : Memory pools

When using the -X switch with the -b option, the output will be displayed in XML format. When using the -L switch with the -b option, in session mode (sub-option 1), all queries will be written to 'ALL.sql'. When using the -A switch with the -b option, db2top will produce a top twenty performance report. When using the -s <sample> switch with the -b option, db2top will only display n samples and exit.

Interactive Commands

d	Goto database screen.
l	Goto sessions screen.
a	Goto application details for agent (or restrict on agent on statement screen) . db2top will prompt for the agent-id.
G	Toggle between all partitions and current partitions.
P	Select db partition on which to issue snapshot.
t	Goto tablespaces screen.
T	Goto tables screen.
b	Goto bufferpools screen.
D	Goto the dynamic SQL screen.
m	Display memory pools.
s	Goto the statements screen.
U	Goto the locks screen.
P	Goto the partitions screen.
u	Goto the utilities screen.
A	Goto the HADR screen.
F	Goto the federation screen.
B	Goto the bottleneck analysis screen.
H	Goto the history screen (experimental)
f	Freeze screen.
w	Watch mode for agentid, osuser, dbuser, application or netname. Statements returned by the session snapshot (option 1) will be written to agent.sql, osuser-agent.sql, dbuser-agent.sql, application-agent .sql or netname-agent.sql. The Dump DB Struct field will create scripts in the current directory to recreate the database. When issued from the dynamic SQL screen (option D), statements will be written to db2adv.sql in a format compatible with db2advsql.
/	Enter expression to filter data. \Expression must conform to regular expression. Each function (screen) can be filtered differently, regex check is applied to the whole row.
< >	Move to left or right of screen.
z Z	Sort on ascending or descending order.
c	This option allows to change the order of the columns displayed on the screen. The syntax is in the form: 1,2,3,... where 1,2,3 correspond

respectively to the 1st, 2nd and 3rd columns displayed. These are the column numbers to use when specifying a sort criteria.

S	Run native DB2 snapshot.
L	Allows to display the complete query text from the SQL screen. Regular DB2 explain can then be run using e or X.
R	Reset snapshot data.
i	Toggle idle sessions on/off.
I	Reset refresh interval.
k	Toggle actual vs delta values.
9	Toggle graph on/off.
X	Toggle extended mode on/off.
c	Toggle snapshot data collector on/off.
V	Set default explains schema.
O	Display session setup.
w	Write session settings to .db2toprc.
q	Quit db2top.

Interactive commands (applications screen only)

r	Return to previous function.
R	Toggle automatic refresh.
9	Toggle graph on/off.
X	Toggle extended mode on/off.
d	Display agents.

Snapshot data collector

db2top can be run in replay mode, which means it can be run against a saved copy of the raw binary snapshot data. In order to do so, db2top must be run first in data collector mode, either in batch mode by running db2top from the command line with the [-C] switch, or by activating/deactivating data collection from an interactive session by pressing C. This will create a file <db2snap-Machine> in the current directory. Afterwards, db2top can be ran against <db2snap-Machine> by using the [-f <db2snap-Machine>] arguments. db2top does not need to attach to DB2 in replay mode, which makes convenient to do remote monitoring. It is possible to limit the content of the stream file (and its size) by specifying any of the suboptions dlbtDUsUF.

Network data collector (experimental)

db2top can be run in client/server mode. This means that one instance of db2top can issue DB2 snapshots in data collection mode and send the output stream via the network to other instances of db2top. To enable this, db2top needs to be started in collection mode with -C and -N (network). Eventually, the -D will set db2top to daemon mode; in this case, messages are sent to syslog(3). Then, other instances of db2top will connect to the server using the -H [host] and -N <port> switches.

Network commands: --stop, --restart, -ping.

File db2toprc

.db2toprc is a configuration file used to setup. parameters at initialization time. db2top will search for the location of .db2toprc using the environment variable \$DB2TOPRC. If the variable is not set, db2top will then search for .db2toprc in the current directory and finally in the home directory.

Environment

DB2TOPRC	see above.
DB2DBDFT	If non null and nothing has been specified on the command line or in .db2toprc, db2top will honor the registry variable content as the default database.
EDITOR	if EDITOR is set, results of explain or native snapshots will be displayed using this command, otherwise, it will default to vi(l).

db2toprc structure

The various sections of \$DB2TOPRC are described here after.

cpu=command

This entry allows the user to specify a command which result will be displayed on the 2nd line at the right of the screen, for example:

```
cpu=vmstat 2 2 | tail -1 | awk '{printf ("%d(usr+sys)", $14 + $15);}'
```

will display 'Cpu=2 (usr+sys)' on the right of the screen.

io=command

This entry allows the user to specify a command which result will be displayed on the 2nd line at the left of the screen, for example:

```
io=vmstat 2 2 | tail -1 | awk '{printf ("%d(bi+bo)", $10+$11);}'
```

will display 'Disk=76(bi+bo)' on the right of the screen.

Both commands are run in a background process and the fields on the screen are updated asynchronously.

shell alias=command

The shell entry allows to specify a user defined command, for example:

shell M=top

will spawn top from a db2Cop session when entering M. It will return to the current screen upon exit.

function alias=command

The function entry allows to specify a user defined command, for example:

function N=netstat

will create a new function called N that will display repeatedly the output of netstat.

Sample .db2toprc file

```
#
# db2top configuration file
# On unix, should be located in $HOME/.db2toprc
# File generated by db2top-1.0h
node= # [-n] node name
database=sample # [-d] database name
user= # [-u] database user
password= # [-p] user password (crypte)
schema= # [-V] default schema for explains
intervals=2 # [-i] sampling interval
active=OFF # [-a] display active sessions (only on/off)
reset=OFF # [-R] Reset snapshot at startup (on/off)
delta=ON # [-k] Toggle display of delta/cumulative
vals
graphic=ON # True if terminal supports semi graphical chars
colors=ON # True if terminal supports colors
gauge=ON # Display graph on sessions list (on/off)
```

```

port=          #      Port for network collection
streamsize=size #      Max collection size per hr eg. 1024/1K: K|M|G
# Command to get cpu usage information from OS
cpu=vmstat 2 2 | tail -1 | awk '{printf("%d(usr+sys)0,$14+$15);}'
# Command to get IO usage information from OS
io=vmstat 2 2 | tail -1 | awk '{printf("%d(bi+bo)0,$10+$11);}'
# Ordering of information in sessions screen
sessions=0,1,18,2,3,4,5,6,7,8,9,10,11,12,13/14,15,16,17
,19,20,21,22,23
tables=0,1,2,4,3,5,6,7
tablespaces=0,1,18,2,3,4,5,6,7,8
bufferpools=0,1,18,2,3,4,5,6,7,8,9,10
dynsql=0,1,18,2,3,4,5,6,7,8,9
statements=0,1
locks=0,1
utilities=0
federation=0 ,2,4
# User defined commands
shell P=top
function N=date && netstat -t tcp

```

The columns order screen

After pressing c, a screen will be shown specifying the order of the columns displayed on screen. The left part of the screen displays the default order and column numbers whereas the right part displays the current ordering. To change the order of the columns, enter in the text field at the bottom of the screen the new column order: enter the relative column positions as displayed on the left

separated by comma. Not all columns need to be specified.

This column ordering can be saved in \$DB2TOPRC for subsequent db2top sessions by pressing w. It is possible to define the default sort order for this function by using the sort=num<a\> syntax, where num is the column number and <a\> means ascending or descending.

Example: tablespaces=0,1,18,2,3,4,5,6,7,8,sort=22a

Valid keywords for column ordering in db2toprc.

```

sessions=
tablespaces=

```

```

tables=
bufferpools=
dynsql=
statements=
locks=
utilities=
federation=

```

Bugs

Occasionally strange information may appear when the snapshot API returns unexpected very large values. db2top will try to display * instead of these values.

Alternatives

The **dsmtop** command is a simple text-based tool for monitoring, similar to the now deprecated db2top command. The dsmtop command can monitor Db2 version 10.1 and above. It is intended primarily for use on Linux and AIX®, but will also run on, Windows although with some limitations.

The **dmctop** utility is a simple text-based tool for monitoring, similar to the dsmtop. The dmctop utility can monitor Db2® version 11.1.0 and later fix packs. It is intended primarily for use on Linux® and Windows systems. Beginning with IBM® Db2 Data Management Console v3.1.5, the monitoring utility dmctop v1.0.2 is bundled with Db2 v11.5.6.

More information on internet

- db2top: <https://www.ibm.com/docs/en/db2/11.5?topic=commands-db2top-db2-monitoring-tool>
- db2toprc: <https://www.ibm.com/docs/en/db2/11.5?topic=mode-db2top-configuration-file#r0025224>
- dsmtop: <https://www.ibm.com/docs/en/db2/11.5?topic=commands-dsmtop-db2-text-based-monitoring-tool-command>
- dmctop: <https://www.ibm.com/docs/en/db2-data-mgr-console/3.1.x?topic=monitoring-utility-dmctop>