Informix Tech Talks by the IIUG

IIUG Tech Talks Everything You Need to Know About Statistics and Data Distributions

by Art S. Kagel

We will get started shortly.

Informix Tech Talks by the IIUG

IIUG Tech Talks Everything You Need to Know About Statistics and Data Distributions

by Art S. Kagel

Date: Thursday, May 22, 2025, 1:00pm EST

Description: A deep dive into UPDATE STATISTICS and important dostats options that you may not know enough about to understand how to use them.

Update Statistics and the Optimizer

Let’s start with some basics.

What does the update statistics command do?

- Updates the system catalog tables with basic statistical detail and stores them in:
  - Systables
  - Syscolumns
  - Sysindices
- Collects data distributions and stores them in:
  - Sysdistrib
  - Sysfragdist
- Re-optimizes (compiles) stored procedures
- Triggers index rebuilds after a server update that changes the index layout!

Update Statistics and Data Distribution

Update statistics [LOW|MEDIUM|HIGH] for table tablename (columns);

LOW- least amount of data gathered; updates systables, syscolumns, sysindexes; does NOT update sysdistrib

HIGH - most amount of data gathered; requires one or more full table scans
updates sysdistrib and possibly sysfragdist. Default resolution of 0.5 % - (200 bins)

MEDIUM - data obtained by sampling; not full table scans, faster than HIGH. Updates sysdistrib and possibly sysfragdist. Default resolution is 2.5% - (40 bins)

Drop Distributions:

UPDATE STATISTICS LOW FOR TABLE tablename (columns) DROP DISTRIBUTIONS;

Stored Procedures:

UPDATE STATISTICS FOR ROUTINE stored_procedure_name;

Statistics Collected During Update Statistics Low

- Systables
  - Number of rows
  - Number of pages to store the data
  - Time LOW stats were last collected for the table
- Syscolumns
  - Second largest value for a column
  - Second smallest value for a column
- Sysindexes
  - Number of unique values for the lead key
  - Width and depth of the index structure
  - How highly clustered the values are for the lead key

Statistics Collected During Update Statistics Low

- This is the basic level of statistics used by the deepest core of the optimizer if there are no data distributions for the table's columns.
- Causes the optimizer to revert to the behavior of the v5.xx optimizer.

UPDATE STATISTICS LOW FOR TABLE <tabname>;

Data Distributions

- Created with update statistics MEDIUM or HIGH
- Distributions are a mapping of the data in the column into a carefully chosen set of column value ranges.
- The data in the column is examined and divided into bins, which represent a percentage of data. For example, if the distributions are gathered into 50 bins then each bin will hold 2% of the data (assuming an even distribution of values)..
- RESOLUTION is the percent of data that is held in each bin. Defaults:
  - MEDIUM = 2.5% (40 bins)
  - HIGH = 0.5% (200 bins)
- The Optimizer uses distributions for columns referenced in the ON and WHERE clauses to estimate the effect on the cost of each possible query path.
- Distributions are also considered when choosing an index for a table.
- You must have DBA privileges, or be the owner of the table, in order to create HIGH or MEDIUM distributions.

Use Dbschema to View Distributions

- dbschema –d database –hd table
  - dbschema –d stores_demo –hd orders
- Displays histograms of the distribution for columns of a specified table, or "all" for all tables.

Distributions Output

```
Distribution for art.extents.size

Constructed on 2025-05-05 17:18:01.99874

Medium Mode, 2.500000 Resolution, 0.950000 Confidence


--- DISTRIBUTION ---

      (                                    1)
  1: (   5208,      10,                    32)


--- OVERFLOW ---

  1: (   1540,                             4)
  2: (   1540,                             5)
  3: (   3814,                             6)
  4: (   9755,                             7)
  5: (  15256,                             8)
  6: (  17603,                             9)
  7: (  22958,                            10)
  8: (  47896,                            11)
  9: (  20244,                            12)
 10: (  22077,                            13)
 11: (  17383,                            14)
 12: (   8802,                            15)
 13: (   9022,                            16)
 14: (   3961,                            17)
 15: (   3301,                            18)
 16: (   1540,                            21)
 17: (   3961,                            22)
```

The report shows the table and column names, the time that the distribution was collected, the level (Medium or High), the resolution, and (for Medium level) the confidence that the data is as good as High level distribution data would be.

This section shows the # of row represented in each Bin, the # of different values in the rows in the Bin, and the highest value in the Bin.

When one value represents more than 25% of the rows in a Bin it will get placed in the Overflow Bin which shows the overflow value and the number of rows containing that value.

Process for Collecting Distributions

- If AUTO mode is not set for the table, nor as the default in the ONCONFIG file, or if the FORCE option was specified then continue. Else:
- Determine if the table should be processed in AUTO mode:
  - If the table is set to AUTO distributions mode, or is set as the default to the ONCONFIG setting determine the number of inserts, updates, and deletes that the table has experienced since distributions were last gathered.
  - If these changes exceed the configured STATLEVEL set for the table (or the default set in the ONCONFIG file) continue.
  - If the threshold is not exceeded, no-op the request and return success.

Note: AUTO mode is ignored when changing sampling size, resolution, or confidence

Process for Collecting Distributions

- Develop scan plan based on available resources
  - This may involve only scanning one or more indexes or performing a table scan or multiple table scans
- Scan table under Isolation Dirty Read
  - High: Read all rows
  - Medium: Sample of percentage of rows
- Sort each column for which distributions are requested
- Collect distributions into bins
  - Identify outliers and rebalance bins
- Begin transaction
  - Delete old column distributions
  - Insert new column distributions
- Commit transaction
- Mark existing distributions cache entries 'stale'

Notes about distributions at the partition level

Partition level distributions are set using the ALTER TABLE statement to set the STATLEVEL attribute of the table. Options are:

- TABLE – Distributions are calculated for the table as a whole
- FRAGMENT – Distributions are calculated separately for each partition of the table
- AUTO – the default. If the table is partitioned via EXPRESSION, LIST, or INTERVAL/RANGE and has at least 1 million rows partition distributions are calculated, otherwise table level distributions. Note that prior to v11.70 partition level distributions were not available.

Note: Partition level distributions are stored in the SYSSBSPACE smart blobspace. If one is not configured and set in the ONCONFIG file, UPDATE STATISTICS will return an error for eligible tables!

Update Statistics Medium Resolution Clause

- Use the Resolution clause to adjust the number of distribution bins. Designate whether or not to avoid calculating data on indexes.
- In MEDIUM mode also can adjust the confidence level or the SAMPLING size.
  - confidence_level – A factor of the likelihood the that sampling in MEDIUM mode will produce the same results as the exact HIGH mode. Default level is 0.95. Must be within the range from 0.80 (minimum) to 0.99 (maximum).
  - resolution - Percentage of sample in each bin of distribution. Default is 2.5% (40 bins) for MEDIUM and 0.5 for HIGH (200 bins).
    - Example
      - 100,000 rows in the table
      - Resolution of 2%
      - Each of the 50 bin will represent approximately 2,000 rows
  - Default Medium with confidence_level = .95 and percent = 2.5 will sample up to 2,963 rows!.

Update Statistics Medium SAMPLING SIZE Clause

- Use the SAMPLING SIZE clause to specify the maximum number of rows to be sampled in MEDIUM mode (100 -> nrows)
  - Example:

```
UPDATE STATISTICS MEDIUM FOR TABLE mytable
    SAMPLING SIZE 120000;
```

- SAMPLING SIZE values less than or equal to 1 represent a maximum percentage of rows to process (not documented)!
- Values less than 0.01 (ie 1%) will also set the CONFIDENCE to set a different maximum number of rows sampled.

Note: First available in v11.10

Recommendations for Update Statistics

- Execute UPDATE STATISTICS (LOW) for the database on a regular basis.
- Execute UPDATE STATISTICS MEDIUM DISTRIBUTIONS ONLY for selected tables, or for the entire database if time is available.
  - The DISTRIBUTIONS ONLY keyword prevents re-updating the LOW data and speeds up the process.
- Execute UPDATE STATISTICS HIGH for:
  - all columns that are listed first in an index key
  - if multiple indexes start with the same columns also do HIGH for the first column that is different from each key
  - Non-index leading columns used in joins
  - Non-index leading columns used in WHERE clause as filters
  - Optionally columns used in partitioning expressions

- **The goal is to balance the time required to execute update statistics versus improving query performance. HIGH with many Bins on every column in a table will always give the best performance. However, it will take long to run and the recommendations here will usually do as well!**

Improving Performance for Update Statistics

- SET PDQPRIORITY when running update statistics, but only for tables
- Set PDQPRIORITY to zero when updating statistics for SPL routines. (If an SPL routine is compiled with positive PDQPRIORITY it will always run with that priority which will limit concurrent sessions!)
- When running HIGH or MEDIUM, increase the memory update statistics has to work with:
  - DS_TOTAL_MEMORY
  - DS_NONPDQ_QUERY_MEM (if PDQ 0)
- Enable parallel sorting (i.e. PDQPRIORITY, PSORT_NPROCS)
- Enable parallel temp space (PSORT_DBTEMP, DBSPACETEMP)

Get my dostats utility

Synopsis:

Dostats automatically generates optimal UPDATE STATISTICS statements for a table(s) or database(s) per the latest Performance Guide manual, and John Miller III's paper on improvements to the way that UPDATE STATISTICS works internally and how to take advantage of those improvements. Options control what databases and tables are affected, whether commands are executed, output to a script, or scheduled to run later using the task scheduler, whether and how stored procedures are handled, the level of verbosity, tweaking the granularity of the statistical distributions captured, specifying criteria for selecting tables to update, and much more. Go to:

https://www.askdbmgt.com/my-utilities

And download the utils2_ak package. You can build it just by unzipping it and running “make” on Linux (instructions for other environments are included in the BUILDING file and in notes in the Makefile).

Dostats options - Basics

-? (or no args) – Print usage with paging
-h <servername> - Note using a shared memory connection requires adding the -s option!
-s – handle shared memory connections. Only one database can be processed.
-d <database spec> - MATCHES specification of databases to process. Asterisk (‘*’) or “ALL” processes all databases.
-t <table spec> - MATCHES specification of tables to process. Omitting -t or passing -t ‘*’ processes all tables in each database.
-p – enable compiling stored procedures. Defaults to enabled if none of -t, -i, or -x are specified, disabled otherwise.
-w <secs> - sets the lock wait time. Default: 10 seconds.

Dostats options - PDQPRIORITY

Setting a positive PDQPRIORITY when processing distributions for tables enables parallel sorting and increases the amount of memory available for sorting without writing to disk.
-Q <PDQ> - the PDQPRIORITY to use for processing tables. Defaults to the value of the PDQPRORITY environment variable if set, zero if not set.
-P <PDQ> - the PDQPRIORITY to use when compiling stored procedures. Defaults to zero. CAUTION: SPL routines compiled with a non-zero PDQPRORITY will always execute at that PDQ level! As noted above, that will limit the number of concurrent sessions processing any SPL routine compiled with any positive PDQ!

Dostats options – Table handling

-i – Various ways to provide a specific list of tables to process. Multiple -i options can be included:

-i <tabname> - similar to -t but does not support wildcards

-i @file – provide a file that contains a list of table names

-i @ - read table names from stdin until EOF

-i !<where clause to select tables>

-i !select tabname from …

-i >procname – name a specific routine to recompile

--as-listed – process tables provided by -i options in the order they were specified. Default process in tabid order.

Dostats options – Table handling

-x – Various ways to provide a specific list of tables to exclude from processing. Multiple -x options can be included:

-x <tabname> - similar to -t but does not support wildcards

-x @file – provide a file that contains a list of table names

-x @ - read table names from stdin until EOF

-x !<where clause to select tables>

-x <database>: - specify a database to exclude from processing when -d ‘*’ or -d all was specified.

Dostats options – Table handling

--small-tables-high – specifies that tables with fewer than a specified number of rows process all columns HIGH

--small-tables-threshold – specifies the row count threshold for tables to be considered a small table for –small-tables-high.

Note: The threshold can also be set in sysadmin:ph_thresholds in the “AUS_SMALL_TABLES” threshold record and enabled by also passing –aus-thresholds

Dostats options – output & processing options

By default, dostats processes the update statistics commands that it generates immediately. Other options are:

-f <file> - output the commands to an SQL file for later processing
--proc-local – create a table of commands in each database along with a stored procedure to execute the contents of the table.
--procedure=<procname> - name the procedure that –proc-local uses
--execute-local – similar to –proc-local but executes the procedure connected to the local database. Required instead of –proc-local for ANSI and unlogged databases.

Dostats options – Clear Old Distributions

In some cases, especially after a version upgrade, it is recommended to remove existing distributions before running update statistics. Dostats provides two options:
-- drop-distributions – Can only be included by user informix. Manually deletes all distributions in the sysdistrib and sysfragdist tables before other processing. Cannot be used if -t, -i, or -x are provided.
--clean-distributions – Manually drops distributions for each table before processing that table further. Can be used with -t, -i, and/or -x.

Dostats options – Adjusting distributions

HIGH

-R <resolution> - Set the resolution (% of rows in each bin).

--distributions-high=<file> - specifies a file containing a list of columns that are to be processed HIGH despite other criteria. Each line in the file contains three space separated fields: database, table, & column.

Dostats options – Adjusting distributions

MEDIUM

-r <resolution> - Set the resolution (% of rows in each bin).

-c <confidence> - Set the confidence level

-Z <sampling> - Set the SAMPLING SIZE.

A fraction <= 1 & >0 representing the percent of rows to sample.

An integer >= 100 representing the maximum number of rows to sample.

Including -Z together with -c may adjust the actual sampling rate.

Dostats options – Misc. Options

Reporting

-e - Expand reporting to output elapsed time of each stage in seconds and fractions.
-E - Expand reporting to output elapsed time of each stage in hours, minutes, seconds and fractions.
--display-time Display starting time for each command and block of commands in UTC timezone. --time-display is the same
--display-local - Display starting time for each command and block of commands in the local timezone.
--display-subsec - display start times to 1/1000000th second. Enables –display-time. Can be used with --display-local

Dostats options – Misc Options

-X – enable SET EXPLAIN during processing. Use to determine settings and environment variables that might improve runtime.
-V Print current version and copyright.
-F Disable single column index optimization. By default index keys for singleton indexes are processed without the DISTRIBUTIONS ONLY allowing LOW and HIGH stats to be calculated in a single command. On some systems it is faster to separate the two commands.
-N – do not use the SPECIFIC keyword when processing stored routines.
-n errlimit – Do not exit until “errlimit” errors have been encountered.
--error-count – if -n is specified, returns the actual number of errors encountered.

Dostats options – Misc. Options

--aus-thresholds – Use the relevant thresholds set in the sysadmin:ph_thresholds table for aging, browsing, and small table processing. Dostats options for thresholds override specific AUS thresholds if included.

--isolation=<lvl> where “lvl” is one of:

- c – COMMITTED READ
- d – DIRTY READ
- l – COMMITTED READ READ COMMITTED

--force-run – override the AUTO_STAT_MODE setting for all tables if set. Always build distributions.

--auto-run – override the AUTO_STAT_MODE setting if clear. Operate is if it is set.

Dostats options – Misc. Options

-m – Process catalog tables in addition to any other tables selected for processing. Can improve query performance in databases with many many objects or permissions.

-g – Process columns used in FRAGMENT BY expressions for HIGH distributions.

The environment variables INFORMIXUSER and INFORMIXPASS can to used to specify a username and password (respectively) to be used to connect to the database rather than the current user's id and password.

Getting it done even faster using drive_dostats!

The utils2_ak package also includes a script, drive_dostats, that can partition the tables in your database and run multiple copies of dostats to complete the job faster by processing multiple tables in parallel.

Getting it done even faster using drive_dostats!

$ drive_dostats

Usage:
  drive_dostats nprocs dbs [tablespec] [-x@excl] [-xexctbl] [dostats options]
                                [-a] [-i@incl] [-iinctbl]

  Driver script to run 'nprocs' copies of dostats each working on a subset of the requested tables.

Getting it done even faster using drive_dostats!

nprocs – number of parallel copies of dostats to run

dbs – the single database to process

tablespec - a MATCHES style wildcard to select tables to include

Getting it done even faster using drive_dostats!

-a - Process smallest tables first. (Default: Largest tables first)
-x@excl - excl is a file containing tablenames to ignore
-xexctbl - exctbl is a tablename to ignore
Multiple -x and -x@ options are accepted and can be mixed
-i@incl - incl is a file containing tablenames to process
-iinctbl - inctbl is a tablename to process
Multiple -i and -i@ options are accepted and can be mixed

dostats options - Most dostats options are passed on and are valid. However, some have no meaning in context of using drive_dostats and may cause the run to fail.

When mixing -x/-x@ and -i/-i@ options only tables which appear in the include list but do not appear in the exclude list will be processed. Mixing these options should be carefully considered and planned to avoid unexpected results.

Compare dostats with & without PDQ versus drive_dostats

Base run:
$ time dostats -d art –force-run
…
Processing completed

real 0m0.828s
user 0m0.041s
sys 0m0.041s

Compare dostats with & without PDQ versus drive_dostats

Base run:
$ time dostats -d art –force-run -Q 100
…

real 0m1.434s
user 0m0.024s
sys 0m0.046s

Compare dostats with & without PDQ versus drive_dostats

drive_dostats run:

```
$ time drive_dostats 7 art --force-run
…
Dostats complete logfiles are:
        /tmp/drive_dostats.80574.0.log
        /tmp/drive_dostats.80574.1.log
        /tmp/drive_dostats.80574.2.log
        /tmp/drive_dostats.80574.3.log
        /tmp/drive_dostats.80574.procs.log

real    0m0.412s
user    0m0.061s
sys     0m0.063s
```

ASK Database Management

Informix Tech Talks by the IIUG on YouTube

Please subscribe to the IIUG TechTalks channel on YouTube to stay informed. This is THE place for Informix how-to videos.

Subscribe at: https://www.youtube.com/c/InformixTechTalksbytheIIUG

www.iiug.org

ASK Database Management

Thank You

Art S. Kagel, President ASK Database Management Corp.

art.kagel@gmail.com
https://www.askdbmgt.com

Download my utilities at: https://www.askdbmgt.com/my-utilities.html