



How to Ensure Statistics Are Not Obsolete

Introduction

This whitepaper provides an overview of how Autonomous Update Statistics (AUS) enables businesses to ensure that the statistics for applications using Progress® OpenEdge® SQL engine are up to date. AUS maintains current statistics used by the OpenEdge SQL cost-based optimizer to determine the most efficient query execution plan.

Are My Statistics Obsolete?

It is a known fact that statistics play a vital role in identifying the most efficient way to execute a SQL statement after considering many factors related to the tables referenced and the conditions specified in the query.

So, it's not atypical when you are working with SQL to encounter these questions:

- Is there a way to identify if the statistics are stale?
- Is there a way to identify if the data in a table has changed by X%?

This is an important step in processing any SQL statement, as it can greatly affect the query execution time. A suboptimal plan can be responsible for excessive I/O, inflated CPU, memory pressure, decreased throughput and reduced overall concurrency.

Autonomous Update Statistics (AUS), a new OpenEdge SQL feature in the 12.0 release, resolves this problem for users. Let's see how AUS keeps the statistics up to date to provide the best query execution plan.

Introducing AUS

AUS updates statistics automatically and keeps them current so that queries always perform at their best, thereby significantly reducing a DBA's maintenance work and minimizing performance problems upfront. When enabled, AUS starts an AUS server and an AUS client to execute the AUS-specific services. The AUS server performs the update statistics and the client identifies the tables with obsolete statistics.

AUS enables you to:

- Improve query performance by automatically updating SQL statistics
- Configure a Maintenance Window, during which statistics can be gathered
- Manage and monitor automatic statistics generation
- Specify the number of parallel tasks that can be spawned to update the statistics in parallel

How Can I Enable AUS?

Online

Enabling AUS is an online operation. You can enable AUS at any time after starting the database in any one of the following ways:

1. Using the AUS configuration command

```
Autoupdate statistics on;
```

2. Updating the VST pub.”_dbparams”

```
Update pub.”_dbparams” set “_DbParams-Value”=’On’ where  
“_DbParams-Name”=’-SQLAutoStats’;
```

3. Using the Promon utility

```
R&D -> Administrative Options -> Server Options ->  
Autonomous Update statistics
```

When you enable AUS, it allocates a default maintenance window and starts a dedicated server named AUSSRV, which performs the update statistics during the maintenance window. The default maintenance window configuration can be viewed using the enhanced **SHOW** command.

```

SHOW AUTOUPDATE STATISTICS ALL;
AUS Property          Value
-----
AUS Status / AUS Log  Enabled / Disabled
Start Date            xx-xx-xxxx
Start Time            xx:xx:xx
Periodicity           Week
Interval              3
Duration              5 Hours
Change Threshold      20 %
Confidence Threshold  20 %
Parallel Tasks        3

```

Default values for
start date = Following Sunday
start time = 00:00:00

However, DBAs can configure a maintenance window based on their preferences when the database load is minimal. To configure the maintenance window, the following SQL command can be used:

```

AUTOUPDATE STATISTICS ON
MAINTENANCE WINDOW EVERY (DAY | WEEK | MONTH) [INTERVAL <N>]
FOR <N> HOUR STARTS AT <TIME> [FROM NOW | <DATE>]
[CHANGE THRESHOLD <change_perc>]
[CONFIDENCE THREHSOLD <conf_perc>]
[PARALLEL TASKS <N>];

```

Where...

- Every **Day** Interval 3: Indicates once every three days
- Every **Week** Interval 2: Indicates twice in a week
- Every **Month** Interval 1: Indicates once in a month
- **Change Threshold:** Indicates the percentage difference between the predicted rows and the number of rows returned to the query
- **Confidence Threshold:** Indicates the percentage difference between the number of times a table is accessed and the number of times the table exceeds the configured change threshold
- **Parallel Tasks:** Indicates the number of parallel tasks spawned by the AUS server, wherein each task updates statistics for one table

Example:

```
AUTOUPDATE STATISTICS ON MAINTENANCE WINDOW EVERY DAY INTERVAL 15 FOR  
5 HOUR STARTS AT "10:00" FROM "01-01-2019";
```

Impact of the command:

The **AUS Client** collects datapoints for all the Select queries once the above command is executed and continues throughout the lifetime of the database or until AUS is disabled.

The **AUS Server** is active and updates the stats during the maintenance window, i.e.,

01-01-2019 10:00 AM to 01-01-2019 03:00 PM

15-01-2019 10:00 AM to 15-01-2019 03:00 PM

30-01-2019 10:00 AM to 30-01-2019 03:00 PM and so on.

Startup Parameter

You can also enable AUS using the database startup parameter **"-SQLAutoStats ON"**.

Example:

```
proserve testdb -S 6789 -SQLAutoStats On -Mn 4 -Ma 3
```

Impact of the Command:

Database **testdb** is started with

- **SQLAutoStats On** indicates AUS services on.
- **Mn 4** - Maximum number of servers that can be spawned by broker is 4(Mn). Apart from these four servers, one additional server named AUSSRV is reserved for AUS. The broker spawns AUSSRV to perform the AUS services when the service is enabled.
- **Ma 3** - For each of the SQL servers, the maximum number of allowed user connections is three. Out of these three connections, one user connection will be consumed by AUS to perform AUS services.

Summary: Total number of user connections after excluding AUS connections will be $Mn * Ma = 4 * 2 = 8$. The Mn parameter should be ≥ 2 when AUS is enabled.

Note: The above Mn Ma parameters are just examples. DBAs can configure these parameters based on their database requirements.

AUS Logging

AUS Logging can be enabled using the following SQL command

Example:

```
AUTOUPDATE STATISTICS LOG ON;
```

Enabling AUS Log results in one server log for the AUS server and another AUS client log for each of the SQL servers. If the user specifies -Mn parameter as 3, one AUS client log will be created for each of the SQL servers respectively.

How Does AUS Work?

The AUS Client

When you enable AUS, an AUS client becomes active on each of the servers. The client tracks all the Select queries executed by the server and identifies the state of statistics based on the predicted rows to the number of rows returned to a query. If there are Inserts/Updates/Deletes on table X that makes the existing statistics obsolete, the AUS client identifies the table only after a Select is performed on the corresponding table (X). The client tracks all Select queries executed by SQL connections.

The AUS Server

When you enable AUS, the broker spawns an AUS server (AUSSRV). The server waits for the maintenance window to begin and updates the statistics of all tables with stale statistics. The server remains active throughout the duration of the maintenance window and becomes idle after the configured duration is completed. The server resumes the services when the subsequent maintenance window begins and so on. The AUS server updates statistics for all the tables in the database during the first maintenance window.

A Glance at the Feature

1. Enable the feature using the proserve command.

```
linuxx86_64:120$ proserve testdb -S 2111 -SQLAutoStats On
```

2. For example, if the feature is enabled on February 14 2019, which is a Thursday, the default maintenance window provided by AUS is available on the following Sunday.

```
SQLExplorer > SHOW AUTOUPDATE STATISTICS ALL;
AUS Property                Value
-----
AUS Status / AUS Log        Enabled / Disabled
Start Date                   2-17-2019
Start Time                   00:00:00
Periodicity                  Week
Interval                     3
Duration                     5 Hours
Change Threshold             20 %
Confidence Threshold         20 %
Parallel Tasks               3
```

Note: First Maintenance window is 2-17-2019 00:00:00 to 2-17-2019 05:00:00
Second Maintenance window is 3-10-2019 00:00:00 to 3-10-2019 05:00:00, ...

3. The AUS client starts tracking all the Select queries executed by the users and keeps a list of the tables that require an update during the maintenance window. The AUS server remains idle until the maintenance window begins. Based on the database load, a DBA can customize the maintenance window configuration.

```
SQLExplorer>AUTOUPDATE STATISTICS ON MAINTENANCE WINDOW
EVERY DAY INTERVAL 15 FOR 5 HOURS STARTS AT "02:03" FROM
"02-14-2019";
```

- You can view the tables with updated statistics during the maintenance window from the table sysprogress.systblstat. The value for column “val_ts” contains the latest time stamp matching the latest maintenance window.

```
SQLExplorer>Select * from Sysprogress.Systblstat where VAL_TS is not Null order by 5 desc;
```

TBLID	PROPERTY	ATTRIBUTE	VALUE	VAL_TS	PARTITIONID
10	2			2019-02-14 02:05:21.163	0
9	2			2019-02-14 02:05:16.725	0
6	2			2019-02-14 02:05:16.1	0
7	2			2019-02-14 02:05:16.034	0
8	2			2019-02-14 02:05:11.507	0
5	2			2019-02-14 02:05:10.846	0
4	2			2019-02-14 02:05:10.84	0
1	2			2019-02-14 02:05:06.273	0
3	2			2019-02-14 02:05:05.599	0
2	2			2019-02-14 02:05:00.745	0

To summarize, you can enable AUS using either the SQL command or the startup parameter. An AUS server and an AUS client are then spawned to perform AUS services. The AUS client tracks the SQL Select queries and identifies the tables with obsolete statistics. The AUS server updates statistics for tables that the AUS client has identified during the maintenance window.

To configure the maintenance window, DBAs can use the following:

```
AUTOUPDATE STATISTICS ON MAINTENANCE WINDOW EVERY DAY INTERVAL 15 FOR 5 HOUR STARTS AT “10:00” FROM “02-12-2019”;
```

and **SHOW** command to view the configured maintenance window. Users can query the systblstat to get a list of tables with statistics updated during the maintenance window.

Conclusion



About the Author

Jyothi Sundaragiri

Jyothi Sundaragiri is the Principal Engineer at Progress. Jyothi has extensive experience on database concepts and growing technology. An active participant in developing blogs and white papers, a quick learner and proactive in her role. Prior to joining Progress, Jyothi worked in InMage Systems for 2 years.

OpenEdge SQL 12.0 provides a mechanism for identifying whether OpenEdge SQL statistics are stale and need to be updated for tables that are involved in **SELECT** statements. With the autonomous update statistics (AUS) service, you can now automatically update and collect statistics during a maintenance window. AUS offers the benefit of using up-to-date statistics while saving the time and effort involved in manually updating statistics.

OpenEdge SQL Update Statistics is no more a manual task. Enable AUS feature to update the statistics hassle free and to improve the query performance.



Learn more

About Progress

Progress (NASDAQ: PRGS) offers the leading platform for developing and deploying strategic business applications. We enable customers and partners to deliver modern, high-impact digital experiences with a fraction of the effort, time and cost. Progress offers powerful tools for easily building adaptive user experiences across any type of device or touchpoint, award-winning machine learning that enables cognitive capabilities to be a part of any application, the flexibility of a serverless cloud to deploy modern apps, business rules, web content management, plus leading data connectivity technology. Over 1,700 independent software vendors, 100,000 enterprise customers, and two million developers rely on Progress to power their applications. Learn about Progress at www.progress.com or +1-800-477-6473.

© 2019 Progress Software Corporation and/or its subsidiaries or affiliates. All rights reserved.

Rev 2019/01 | RITM0034642

