

1 Objectives

The objective of this project is to help customer to resolve SQL width problem by automatically updating sql column width in schema whenever there is a data truncation event for the column.

This autonomous schema update (ASU) has to be authorized by the customer. Sql will not turn on ASU by default.

ASU will enable queries to see application data in its true (actual) size in a 2 step process.

- In 1st execution of the query, sql will learn about the true size of data.
- In the 2nd execution sql will return data in its actual size. This 2 step process is not an objective but a limitations inherent in how sql applications bind to (learn about) the database.

2 Overview

The sql columns which experience data exceeding their defined column size during a query execution will be updated with the actual data size as column width in the schema.

The query will return data in actual size if it is executed again.

Schema will be updated only if data truncation is turned on and there are data truncation events.

The new startup parameter which enables Autonomous Schema Update (ASU) will automatically turn on the Authorized Data Truncation feature (released in OpenEdge 11.5.1).

3 Terminology

1. Autonomous – acting independently.
2. Autonomous Schema Update - the name for the sql feature described in this doc.
3. ASU – acronym for Autonomous Schema Update.
4. ADT – acronym for Authorized Data Truncation (feature implemented for 11.5.1 to help sql width problems).

4 New Keywords

None

5 Introduction

Autonomous Schema Update (ASU) feature will update the schema automatically, without any user intervention, at appropriate time.

In 11.5.1, a new feature, Authorized Data Truncation (ADT) was introduced. If width of the column data being operated on in a query exceeds the defined width of the column, SQL throws error. But if ADT feature is turned on, then SQL returns truncated data for such too-large column rather than throwing error.

ASU will enable sql to automatically increase the defined column size to actual data size whenever there is data truncation event for the column. Re-execution of the query will return data in actual (rather than truncated data) as ASU will take care to increase column size in schema.

6 Functional Description

6.1 Startup Parameter

A new server startup parameter “-SQLWidthUpdate” will be added to control “ASU” related operation. Value of this parameter will be remembered for the lifetime of the server and for all connections. If value of this parameter is supplied as “ON” then “ASU” functionalities will be executed. If value of this parameter is supplied as “OFF” then “ASU” functionalities will not be executed. This parameter will be treated as “OFF” if not supplied at all during database startup.

ASU parameter can only be provided during server startup. Unlike ADT, ASU parameter can't be provided in connection URL.

6.1.1 ASU and ADT Server Startup Parameter Interconnection

Turning ON ASU will turn ON ADT explicitly. But on the contrary, turning OFF ASU will not turn OFF ADT. Note that by default both ADT and ASU are OFF.

Following examples elaborate this in detail.

- If ASU is turned ON but ADT is turned OFF, then ADT will be turned ON explicitly. In this case for any data truncation event, schema update operation (ASU) operation will be performed.
- If ASU and ADT are both turned ON then for any data truncation event, schema update operation (ASU) operation will be performed.
- If ASU is turned OFF but ADT is turned ON then only authorized data truncation (ADT) will be performed but schema update (ASU) will not be performed.
- If both ASU and ADT are turned OFF then neither ADT nor ASU operation will be performed.

6.1.2 ADT Parameter in Connection URL and ASU Parameter

Other than server startup parameter, ADT parameter can also be provided in user connection URL. ADT parameter in connection URL overrides the ADT server startup parameter i.e. if ADT parameter is present in the connection URL then this value is considered ignoring the server startup parameter.

Above behavior for ADT parameter will be continued even after introduction of ASU.

When ASU server startup parameter is switched OFF, this behavior is not conflicting with the rules described in above section.

But when ASU startup parameter is turned ON, startup parameter for ADT is also turned ON. In this case, if a user connects with ADT parameter OFF in connection URL, then ADT will be turned OFF for the connection. So for that connection ASU will stay turned ON but ADT will stay turned OFF, conflicting the rules described in above section.

However this behavior is acceptable as this will enable to turn OFF ASU indirectly for that particular connection only. As ADT is turned OFF for this connection, no data truncation event will be there. So, even though ASU is turned ON, there will be nothing for schema update for this connection.

6.1.3 Examples

If ASU is ON, ADT is OFF and in connection URL ADT is OFF. In this case ASU will turn ON ADT but will be turned OFF by connection URL. So finally ASU will be ON and ADT will be OFF. So for this connection no data truncation event will be there and hence ASU will have nothing to update even though it is ON.

If ASU is ON, ADT is ON and in connection URL ADT is OFF. In this case ADT will be turned OFF by connection URL. So for this connection no data truncation event will be there and hence ASU will have nothing to update even though it is ON.

If ASU is ON, ADT is ON and in connection URL ADT is ON. In this case both ASU and ADT are ON. So for any data truncation event, schema update operation (ASU) operation will be performed.

If ASU is ON, ADT is OFF and in connection URL ADT is ON. In this case ADT will be turned ON by both ASU and the connection URL. So for any data truncation event, schema update operation (ASU) operation will be performed.

If ASU is OFF, ADT is ON and in connection URL ADT is OFF. In this case ASU and ADT will be both OFF. So neither ADT nor ASU operation will be performed.

If ASU is OFF, ADT is OFF and in connection URL ADT is ON. In this case only ADT will be ON. So in this case, only data truncation event will happen but not ASU.

6.2 Functionality – How it works

In the OE sql server, server management components organize and monitor the server framework. This includes startup and shutdown, user connection creation, and other overhead operations.

Whenever a user requests a connection to the server for query execution, the server management components creates a new user connection for the user. Execution of query is performed under the newly established user connection. The user connection runs in its own operating system thread, within the sql server process. As part of ASU, when a user query experiences data truncation events (column value larger than defined size), the user connection will accumulate information about the actual size of the columns involved.

As part of ASU, the server management components will be enhanced to gather schema update information from client connections, and periodically update the database schema appropriately.

Note that user transactions running in user connections do not themselves update the schema. All schema updates are applied by the server management components, using the server connection to the database.