

Teiid - Scalable Information Integration

1

Teiid Reference Documentation

6.2.0

Preface	vii
1. SQL Support	1
1.1. Identifiers	1
1.2. Expressions	2
1.2.1. Column Identifiers	2
1.2.2. Literals	2
1.2.3. Aggregate Functions	3
1.2.4. Case and searched case	4
1.2.5. Scalar subqueries	4
1.2.6. Parameter references	4
1.3. Criteria	4
1.4. SQL Commands	5
1.4.1. SELECT Command	5
1.4.2. INSERT Command	7
1.4.3. UPDATE Command	7
1.4.4. DELETE Command	7
1.4.5. EXECUTE Command	7
1.4.6. Procedural Relational Command	8
1.5. Temp Tables	9
1.6. SQL Clauses	10
1.6.1. SELECT Clause	10
1.6.2. FROM Clause	10
1.6.3. WHERE Clause	11
1.6.4. GROUP BY Clause	11
1.6.5. HAVING Clause	12
1.6.6. ORDER BY Clause	12
1.6.7. LIMIT Clause	12
1.6.8. INTO Clause	13
1.6.9. OPTION Clause	13
1.7. Set Operations	14
1.8. Subqueries	14
1.8.1. Inline views	15
1.8.2. Subqueries in the WHERE and HAVING clauses	15
2. XML SELECT Command	17
2.1. Overview	17
2.2. Query Structure	17
2.2.1. FROM Clause	17
2.2.2. SELECT Clause	17
2.2.3. WHERE Clause	18
2.2.4. ORDER BY Clause	20
2.3. Document Generation	20
2.3.1. Document Validation	20
3. Datatypes	21
3.1. Supported Types	21

3.2. Type Conversions	22
3.3. Special Conversion Cases	24
3.3.1. Conversion of String Literals	24
3.3.2. Converting to Boolean	24
3.3.3. Date/Time/Timestamp Type Conversions	24
3.4. Escaped Literal Syntax	25
4. Scalar Functions	27
4.1. Numeric Functions	27
4.1.1. Parsing Numeric Datatypes from Strings	29
4.1.2. Formatting Numeric Datatypes as Strings	30
4.2. String Functions	30
4.3. Date/Time Functions	32
4.3.1. Parsing Date Datatypes from Strings	35
4.3.2. Specifying Time Zones	36
4.4. Type Conversion Functions	36
4.5. Choice Functions	36
4.6. Decode Functions	37
4.7. Lookup Function	38
4.7.1. Clearing the Cache	39
4.8. System Functions	40
4.9. XML Functions	40
4.10. Security Functions	40
4.11. User Defined Functions	41
4.11.1. UDF Definition	41
4.11.2. Source Supported UDF	42
4.11.3. Non-pushdown Support for User-Defined Functions	42
4.11.4. Installing user-defined functions	43
4.12. Nondeterministic Function Handling	44
5. Procedures	45
5.1. Procedure Language	45
5.1.1. Command Statement	45
5.1.2. Dynamic SQL Command	45
5.1.3. Declaration Statement	48
5.1.4. Assignment Statement	49
5.1.5. If Statement	49
5.1.6. Loop Statement	50
5.1.7. While Statement	50
5.1.8. Continue Statement	50
5.1.9. Break Statement	50
5.1.10. Error Statement	50
5.2. Virtual Procedures	51
5.2.1. Virtual Procedure Definition	51
5.2.2. Procedure Input Parameters	51
5.2.3. Examlpe Virtual Procedures	52

5.2.4. Executing Virtual Procedures	53
5.3. Update Procedures	53
5.3.1. Update Procedure Definition	54
5.3.2. Special Variables	54
5.3.3. Update Procedure Command Criteria	55
5.3.4. Update Procedure Processing	57
6. Transaction Support	59
6.1. AutoWrap Execution Property	59
6.2. Updating Model Count	60
6.3. JDBC and Transactions	60
6.3.1. JDBC API Functionality	60
6.3.2. J2EE Usage Models	61
6.4. Limitations and Workarounds	61
7. System Tables	63
7.1. VDB and Model Metadata	63
7.1.1. System.VirtualDatabases	63
7.1.2. System.Models	63
7.1.3. System.ModelProperties	64
7.2. Table Metadata	64
7.2.1. System.Groups	64
7.2.2. System.GroupProperties	64
7.2.3. System.Elements	65
7.2.4. System.ElementProperties	66
7.2.5. System.Keys	66
7.2.6. System.KeyProperties	67
7.2.7. System.KeyElements	67
7.3. Procedure Metadata	68
7.3.1. System.Procedures	68
7.3.2. System.ProcedureProperties	68
7.3.3. System.ProcedureParams	69
7.4. Datatype Metadata	69
7.4.1. System.DataTypes	69
7.4.2. System.DataTypeProperties	70
8. Connectors	71
8.1. Common Connector Properties	71
8.2. Source Security	72
8.2.1. Specific Source Credentials	72
8.2.2. Using Session Credentials	73
8.2.3. Session and Specific Source Credentials	73
8.2.4. Configuring Connectors For Source Security	73
9. Federated Planning	75
9.1. Overview	75
9.2. Federated Optimizations	77
9.2.1. Access Patterns	77

9.2.2. Pushdown	77
9.2.3. Dependent Joins	77
9.2.4. Copy Criteria	78
9.2.5. Projection Minimization	78
9.2.6. Partial Aggregate Pushdown	78
9.2.7. Optional Join	78
9.2.8. Standard Relational Techniques	79
9.3. Federated Failure Modes	79
9.3.1. Partial Results	79
9.4. Query Plans	81
9.4.1. Getting a Query Plan	81
9.4.2. Analyzing a Query Plan	82
9.4.3. Relational Plans	82
9.5. Query Planner	84
9.5.1. Relational Planner	84
9.5.2. Procedure Planner	90
9.5.3. XML Planner	90
9.5.4. XQuery Planner	90
10. Architecture	91
10.1. Terminology	91
10.2. Data Management	91
10.2.1. Cursoring and Batching	91
10.2.2. Buffer Management	92
10.2.3. Cleanup	92
10.3. Query Termination	93
10.3.1. Canceling Queries	93
10.3.2. Timeouts	93
10.4. Processing	93
10.4.1. Join Algorithms	93
10.4.2. Sort Based Algorithms	94
A. BNF Grammar	95
A.1. Terminals	95
A.2. Non-Terminals	96

Preface

Teiid offers a highly scalable and high performance solution to information integration. By allowing integrated and enriched data to be consumed relationally or as XML over multiple protocols, Teiid simplifies data access for developers and consuming applications.

Commercial development support, production support, and training for Teiid is available through JBoss Inc. Teiid is a Professional Open Source project and a critical component of the JBoss Enterprise Data Services Platform.

SQL Support

Teiid supports SQL for issuing queries and for defining view transformations; see also Procedure Language for how SQL is used in virtual procedures and update procedures.

Teiid provides nearly all of the functionality of SQL-92 DML. SQL-99 and later features are constantly being added based upon community need. The following does not attempt to cover SQL exhaustively, but rather highlights SQL's usage within Teiid. See the [grammar](#) for the exact form of SQL accepted by Teiid.

1.1. Identifiers

SQL commands contain references to tables and columns. These references are in the form of identifiers, which uniquely identify the tables and columns in the context of the command. All queries are processed in the context of a virtual database, or VDB. Because information can be federated across multiple sources, tables and columns must be scoped in some manner to avoid conflicts. This scoping is provided by models, which contain the information for each data source or set of views.

Fully-qualified table and column names are of the following form, where the separate 'parts' of the identifier are delimited by periods.

- TABLE: <model_name>.<table_spec>
- COLUMN: <model_name>.<table_spec>.<column_name>

Syntax Rules:

- Identifiers can consist of alphanumeric characters, or the underscore (`_`) character, and must begin with an alphabetic character. Any Unicode character may be used in an identifier.
- Because different data sources organize tables in different ways, some prepending catalog or schema or user information, Teiid allows the 'table_spec' to be a dot-delimited construct.
- Identifiers are not case-sensitive in Teiid.
- The separate parts of an identifier can be quoted, with double quotes. This is not required, but some tools do this automatically. Quotes establish another level of grouping, in addition to the dot delimiters. Quotes should not be used in such a way that the table specification, which may itself have multiple parts, is split between two quoted sections.

Some examples of valid fully-qualified table identifiers are:

- MyModel.MySchema.Portfolios
- "MyModel"."MySchema.Portfolios"

- MyModel.Inventory
- MyModel.MyCatalog.dbo.Authors

Some examples of valid fully-qualified column identifiers are:

- MyModel.MySchema.Portfolios.portfolioID
- "MyModel"."MySchema.Portfolios"."portfolioID"
- MyModel.Inventory.totalPallets
- MyModel.MyCatalog.dbo.Authors.lastName

Fully-qualified identifiers can always be used in SQL commands. Partially- or unqualified forms can also be used, as long as the resulting names are unambiguous in the context of the command. Different forms of qualification can be mixed in the same query.

1.2. Expressions

Identifiers, literals, and functions can be combined into expressions. Expressions can be used almost anywhere in a query -- SELECT, FROM (if specifying join criteria, WHERE, GROUP BY, HAVING. However you currently cannot use expressions in an ORDER BY clause.

Teiid supports the following types of expressions:

- *Column identifiers*
- *Literals*
- *Scalar functions*
- *Aggregate functions*
- *Case and searched case*
- *Scalar subqueries*
- *Parameter references*

1.2.1. Column Identifiers

Column identifiers are used to specify the output columns in SELECT statements, the columns and their values for INSERT and UPDATE statements, and criteria used in WHERE and FROM clauses. They are also used in GROUP BY, HAVING, and ORDER BY clauses. The syntax for column identifiers was defined in the *Identifiers* section above.

1.2.2. Literals

Literal values represent fixed values. These can any of the 'standard' *data types*.

Syntax Rules:

- Integer values will be assigned an integral data type big enough to hold the value (integer, long, or bigint).
- Floating point values will always be parsed as a double.
- The keyword 'null' is used to represent an absent or unknown value and is inherently untyped. In many cases, a null literal value will be assigned an implied type based on context. For example, in the function '5 + null', the null value will be assigned the type 'integer' to match the type of the value '5'. A null literal used in the SELECT clause of a query with no implied context will be assigned to type 'string'.

Some examples of simple literal values are:

- `'abc'`
- `'isn't true'` - use an extra single tick to escape a tick in a string with single ticks
- `5`
- `-37.75e01` - scientific notation
- `100.0` - parsed as double
- `true`
- `false`
- `'\u0027'` - unicode character

1.2.3. Aggregate Functions

Aggregate functions take sets of values from a group produced by an explicit or implicit GROUP BY and return a single scalar value computed from the group.

Teiid supports the following aggregate functions:

- `COUNT(*)` – count the number of values (including nulls and duplicates) in a group
- `COUNT(expression)` – count the number of values (excluding nulls) in a group
- `SUM(expression)` – sum of the values (excluding nulls) in a group
- `AVG(expression)` – average of the values (excluding nulls) in a group
- `MIN(expression)` – minimum value in a group (excluding null)
- `MAX(expression)` – maximum value in a group (excluding null)

Syntax Rules:

- Some aggregate functions may contain a keyword 'DISTINCT' before the expression, indicating that duplicate expression values should be ignored. DISTINCT is not allowed in COUNT(*) and is not meaningful in MIN or MAX (result would be unchanged), so it can be used in COUNT, SUM, and AVG.
- Aggregate functions may only be used in the HAVING or SELECT clauses and may not be nested within another aggregate function.
- Aggregate functions may be nested inside other functions.

For more information on aggregates, see the sections on GROUP BY or HAVING.

1.2.4. Case and searched case

Teiid supports two forms of the CASE expression which allows conditional logic in a scalar expression.

Supported forms:

- CASE <expr> (WHEN <expr> THEN <expr>)+ [ELSE expr] END
- CASE (WHEN <criteria> THEN <expr>)+ [ELSE expr] END

Each form allows for an output based on conditional logic. The first form starts with an initial expression and evaluates WHEN expressions until the values match, and outputs the THEN expression. If no WHEN is matched, the ELSE expression is output. If no WHEN is matched and no ELSE is specified, a null literal value is output. The second form (the searched case expression) searches the WHEN clauses, which specify an arbitrary criteria to evaluate. If any criteria evaluates to true, the THEN expression is evaluated and output. If no WHEN is true, the ELSE is evaluated or NULL is output if none exists.

1.2.5. Scalar subqueries

Subqueries can be used to produce a single scalar value in the SELECT, WHERE, or HAVING clauses only. A scalar subquery must have a single column in the SELECT clause and should return either 0 or 1 row. If no rows are returned, null will be returned as the scalar subquery value. For other types of subqueries, see the [Subqueries](#) section below.

1.2.6. Parameter references

Parameters are specified using a '?' symbol. Parameters may only be used with PreparedStatement or CallableStatements in JDBC. Each parameter is linked to a value specified by 1-based index in the JDBC API.

1.3. Criteria

Criteria are of two basic forms:

- Predicates that evaluate to true or false
- Logical criteria that combine predicates (AND, OR, NOT)

Syntax Rules:

- expression ($=|<|>|!<|>|<=|>=$) (expression|((ANY|ALL|SOME) subquery))
- expression [NOT] IS NULL
- expression [NOT] IN (expression[,expression]*)|subquery
- expression [NOT] LIKE expression [ESCAPE char]
- EXISTS(subquery)
- expression BETWEEN minExpression AND maxExpression
- criteria AND|OR criteria
- NOT criteria
- Criteria may be nested using parenthesis.

Some examples of valid criteria are:

- (balance > 2500.0)
- $100*(50 - x)/(25 - y) > z$
- concat(areaCode,concat('-',phone)) LIKE '314%1'



Comparing null Values

Null values represent an unknown value. Comparison with a null value will evaluate to 'unknown', which can never be true even if 'not' is used.

1.4. SQL Commands

There are 4 basic commands for manipulating data in SQL, corresponding to the CRUD create, read, update, and delete operations: INSERT, SELECT, UPDATE, and DELETE. In addition, procedures can be executed using the EXECUTE command or through a *procedural relational command*.

1.4.1. SELECT Command

The SELECT command is used to retrieve records any number of relations.

A SELECT command has a number of clauses:

- *SELECT ...*
- *[FROM ...]*
- *[WHERE ...]*
- *[GROUP BY ...]*
- *[HAVING ...]*
- *[ORDER BY ...]*
- *[LIMIT [offset,] limit]*
- *[OPTION ...]*

All of these clauses other than OPTION are defined by the SQL specification. The specification also specifies the order that these clauses will be logically processed. Below is the processing order where each stage passes a set of rows to the following stage. Note that this processing model is logical and does not represent the way any actual database engine performs the processing, although it is a useful model for understanding questions about SQL.

- FROM stage - gathers all rows from all tables involved in the query and logically joins them with a Cartesian product, producing a single large table with all columns from all tables. Joins and join criteria are then applied to filter rows that do not match the join structure.
- WHERE stage - applies a criteria to every output row from the FROM stage, further reducing the number of rows.
- GROUP BY stage - groups sets of rows with matching values in the group by columns.
- HAVING stage - applies criteria to each group of rows. Criteria can only be applied to columns that will have constant values within a group (those in the grouping columns or aggregate functions applied across the group).
- SELECT stage - specifies the column expressions that should be returned from the query. Expressions are evaluated, including aggregate functions based on the groups of rows, which will no longer exist after this point. The output columns are named using either column aliases or an implicit name determined by the engine. If SELECT DISTINCT is specified, duplicate removal will be performed on the rows being returned from the SELECT stage.
- ORDER BY stage - sorts the rows returned from the SELECT stage as desired. Supports sorting on multiple columns in specified order, ascending or descending. The output columns will be identical to those columns returned from the SELECT stage and will have the same name.
- LIMIT stage - returns only the specified rows (with skip and limit values).

This model can be used to understand many questions about SQL. For example, columns aliased in the SELECT clause can only be referenced by alias in the ORDER BY clause. Without

knowledge of the processing model, this can be somewhat confusing. Seen in light of the model, it is clear that the ORDER BY stage is the only stage occurring after the SELECT stage, which is where the columns are named. Because the WHERE clause is processed before the SELECT, the columns have not yet been named and the aliases are not yet known.

1.4.2. INSERT Command

The INSERT command is used to add a record to a table.

Example Syntax

- INSERT INTO table (column,...) VALUES (value,...)
- INSERT INTO table (column,...) query

1.4.3. UPDATE Command

The UPDATE command is used to modify records in a table. The operation may result in 1 or more records being updated, or in no records being updated if none match the criteria.

Example Syntax

- UPDATE table SET (column=value,...) [WHERE criteria]

1.4.4. DELETE Command

The DELETE command is used to remove records from a table. The operation may result in 1 or more records being deleted, or in no records being deleted if none match the criteria.

Example Syntax

- DELETE FROM table [WHERE criteria]

1.4.5. EXECUTE Command

The EXECUTE command is used to execute a procedure, such as a virtual procedure or a stored procedure. Procedures may have zero or more scalar input parameters. The return value from a procedure is a result set, the same as is returned from a SELECT. Note that EXEC can be used as a short form of this command.

Example Syntax

- EXECUTE proc()
- EXECUTE proc(value, ...)
- EXECUTE proc(name1=value1,name4=param4, ...) - named parameter syntax

Syntax Rules:

- The default order of parameter specification is the same as how they are defined in the procedure definition.
- You can specify the parameters in any order by name. Parameters that have default values and/or are nullable in the metadata, can be omitted from the named parameter call and will have the appropriate value passed at runtime.
- If the procedure does not return a result set, the values from the RETURN, OUT, and IN_OUT parameters will be returned as a single row.

1.4.6. Procedural Relational Command

Procedural relational commands use the syntax of a SELECT to emulate an EXEC. In a procedural relational command a procedure group name is used in a FROM clause in place of a table. That procedure will be executed in place of a normal table access if all of the necessary input values can be found in criteria against the procedure. Each combination of input values found in the criteria results in an execution of the procedure.

Example Syntax

- `select * from proc`
- `select output_param1, output_param2 from proc where input_param1 = 'x'`
- `select output_param1, output_param2 from proc, table where input_param1 = table.col1 and input_param2 = table.col2`

Syntax Rules:

- The procedure as a table projects the same columns as an exec with the addition of the input parameters. For procedures that do not return a result set, IN_OUT columns will be projected as two columns, one that represents the output value and one named {column name}_IN that represents the input of the parameter.
- Input values are passed via criteria. Values can be passed by '=', 'is null', or 'in' predicates.
- The procedure virtual group automatically has an access pattern on its IN and IN_OUT parameters which allows it to be planned correctly as a dependent join when necessary or fail when sufficient criteria cannot be found.
- Procedures containing duplicate names between the parameters (IN, IN_OUT, OUT, RETURN) and result set columns cannot be used in a procedural relational command.
- Default values for IN, IN_OUT parameters are not used if there is no criteria present for a given input. Default values are only valid for [named procedure syntax](#).



Multiple Execution

The usage of 'in' or join criteria can result in the procedure being executed multiple times.

1.5. Temp Tables

Teiid supports creating temporary, or "temp", tables. Temp tables are dynamically created, but are treated as any other physical table.

Temp tables can be defined implicitly by referencing them in a SELECT INTO or in an INSERT statement or explicitly with a CREATE TABLE statement. Implicitly created temp tables must have a name that starts with '#'.

Creation syntax:

- CREATE LOCAL TEMPORARY TABLE<temporary table name> (<column name> <data type>,...)
- SELECT <element name>,...INTO <temporary table name> FROM <table name>
- INSERT INTO <temporary table name> ((<column name>,...)VALUES (<value>,...)

Drop syntax:

- DROP TABLE <temporary table name>

Limitations:

- With the CREATE TABLE syntax only basic table definition (column name and type information) is supported.
- The "ON COMMIT" clause is not supported in the CREATE TABLE statement.
- "drop behavior" option is not supported in the drop statement.
- Only local temporary tables are supported. This implies that the scope of temp table will be either to the session or the block of a virtual procedure that creates it.
- Session level temp tables are not fail-over safe.
- temp tables are non-transactional.
- Temp tables do not support update or delete operations.

The following example is a series of statements that loads a temporary table with data from 2 sources, and with a manually inserted record, and then uses that temp table in a subsequent query.

```
...  
CREATE LOCAL TEMPORARY TABLE TEMP (a integer, b integer, c integer);  
SELECT * INTO temp FROM Src1; SELECT * INTO temp FROM Src2;  
INSERT INTO temp VALUES (1,2,3);  
SELECT a,b,c FROM Src3, temp WHERE Src3.a = temp.b;  
...
```

See [virtual procedures](#) for more on temp table usage.

1.6. SQL Clauses

This section describes the clauses that are used in the various [SQL commands](#) described in the previous section. Nearly all these features follow standard SQL syntax and functionality, so any SQL reference can be used for more information.

1.6.1. SELECT Clause

SQL queries start with the SELECT keyword and are often referred to as "SELECT statements". Teiid supports most of the standard SQL query constructs.

Usage:

```
SELECT [DISTINCT|ALL] ((expression [[AS] name])|(group  
identifier.STAR))*|STAR ...
```

Syntax Rules:

- Aliased expressions are only used as the output column names and in the ORDER BY clause. They cannot be used in other clauses of the query.
- DISTINCT may only be specified if the SELECT symbols are comparable.

1.6.2. FROM Clause

The FROM clause specifies the target table(s) for SELECT, UPDATE, and DELETE statements.

Example Syntax:

- FROM {table [AS alias]}
- FROM table1 [INNER|LEFT OUTER|RIGHT OUTER|FULL OUTER] JOIN table1 ON join-criteria
- FROM table1 CROSS JOIN table1
- FROM (subquery) [AS alias]

- FROM table1 JOIN table2 MAKEDEP ON join-criteria
- FROM table1 JOIN table2 MAKENOTDEP ON join-criteria
- FROM table1 left outer join /* *optional* */table2 ON join-criteria



DEP Hints

MAKEDEP and MAKENOTDEP are hints used to control *dependent join* behavior. They should only be used in situations where the optimizer does not chose the most optimal plan based upon query structure, metadata, and costing information.

1.6.3. WHERE Clause

The WHERE clause defines the criteria to limit the records affected by SELECT, UPDATE, and DELETE statements.

The general form of the WHERE is:

- WHERE *criteria*

1.6.4. GROUP BY Clause

The GROUP BY clause denotes that rows should be grouped according to the specified expression values. One row will be returned for each group, after optionally filtering those aggregate rows based on a HAVING clause.

The general form of the GROUP BY is:

- GROUP BY expression (,expression)*

Syntax Rules:

- Column references in the group by clause must be to unaliased output columns.
- Expressions used in the group by must appear in the select clause.
- Column references and expressions in the select clause that are not used in the group by clause must appear in aggregate functions.
- If an aggregate function is used in the SELECT clause and no GROUP BY is specified, an implicit GROUP BY will be performed with the entire result set as a single group. In this case, every column in the SELECT must be an aggregate function as no other column value will be fixed across the entire group.
- The group by columns must be of a comparable type.

1.6.5. HAVING Clause

The HAVING clause operates exactly as a WHERE clause although it operates on the output of a GROUP BY. It supports the same syntax as the WHERE clause.

Syntax Rules:

- Expressions used in the group by clause must either contain an aggregate function: COUNT, AVG, SUM, MIN, MAX. or be one of the grouping expressions.

1.6.6. ORDER BY Clause

The ORDER BY clause specifies how the returned records from a SELECT should be sorted. The options are ASC (ascending) and DESC (descending).

Usage:

```
ORDER BY column1 [ASC|DESC], ...
```

Syntax Rules:

- Sort columns may be specified positionally by a 1 based integer or string literal, by SELECT clause alias name, or by a column reference.
- Column references may appear in the SELECT clause as the expression for an aliased column or may reference columns from tables in the FROM clause. If the column reference is not in the SELECT clause the query must not be a set operation, specify SELECT DISTINCT, or contain a GROUP BY clause.
- The ORDER BY columns must be of a comparable type.
- If an ORDER BY is used in an inline view or view definition without a limit clause, it will be removed by the Teiid optimizer.



Warning

The use of positional ordering is no longer supported by the ANSI SQL standard and is a deprecated feature in Teiid. It is preferable to use alias names in the order by clause.

1.6.7. LIMIT Clause

The LIMIT clause specifies a limit on the number of records returned from the SELECT command. An optional offset (the number of rows to skip) can be specified.

Usage:

```
LIMIT [offset,] limit
```

Examples:

- LIMIT 100 - returns the first 100 records (rows 1-100)
- LIMIT 500, 100 - skips 500 records and returns the next 100 records (rows 501-600)

1.6.8. INTO Clause



Warning

Usage of the INTO Clause for inserting into a table has been deprecated. An *INSERT* with a query command should be used instead.

When the into clause is specified with a SELECT, the results of the query are inserted into the specified table. This is often used to insert records into a temporary table. The INTO clause immediately precedes the FROM clause.

Usage:

```
INTO table FROM ...
```

Syntax Rules:

- The INTO clause is logically applied last in processing, after the ORDER BY and LIMIT clauses.
- Teiid's support for SELECT INTO is similar to MS SQL Server. The target of the INTO clause is a table where the result of the rest select command will be inserted. SELECT INTO should not be used UNION query.

1.6.9. OPTION Clause

The OPTION keyword denotes options the user can pass in with the command. These options are Teiid-specific and not covered by any SQL specification.

Usage:

```
OPTION option, (option)*
```

Supported options:

- SHOWPLAN - returns the query plan along with the results
- PLANONLY - returns the query plan, but does not execute the command

- MAKEDEP [table, (table)*] - specifies source tables that should be made dependent in the join
- MAKENOTDEP [table, (table)*] - prevents a dependent join from being used
- DEBUG - prints query planner debug information in the log and returns it through the JDBC API

Examples:

- LIMIT 100 - returns the first 100 records (rows 1-100)
- LIMIT 500, 100 - skips 500 records and returns the next 100 records (rows 501-600)

1.7. Set Operations

Teiid supports the UNION, UNION ALL, INTERSECT, EXCEPT set operation as a way of combining the results of commands.

Usage:

```
command (UNION|INTERSECT|EXCEPT) [ALL] command [ORDER BY...]
```

Syntax Rules:

- The output columns will be named by the output columns of the first set operation branch.
- Each SELECT must have the same number of output columns and compatible data types for each relative column. Data type conversion will be performed if data types are inconsistent and implicit conversions exist.
- If UNION, INTERSECT, or EXCEPT is specified without all, then the output columns must be comparable types.
- INTERSECT ALL, and EXCEPT ALL are currently not supported.

1.8. Subqueries

A subquery is a SQL query embedded within another SQL query. The query containing the subquery is the outer query.

Supported subquery types:

- Scalar subquery - a subquery that returns only a single column with a single value. Scalar subqueries are a type of expression and can be used where single valued expressions are expected.
- Correlated subquery - a subquery that contains a column reference to from the outer query.
- Uncorrelated subquery - a subquery that contains no references to the outer sub-query.

Supported subquery locations:

- [Subqueries in the FROM clause](#)
- [Subqueries in the WHERE/HAVING Clauses](#)
- Subqueries may be used in any expression or CASE CRITERIA in the SELECT clause.

1.8.1. Inline views

Subqueries in the FROM clause of the outer query (also known as "inline views") can return any number of rows and columns. This type of subquery must always be given an alias.

Example 1.1. Example Subquery in FROM Clause (Inline View)

```
SELECT a FROM (SELECT Y.b, Y.c FROM Y WHERE Y.d = '3') AS X WHERE a = X.c AND
b = X.b
```

1.8.2. Subqueries in the WHERE and HAVING clauses

Subqueries supported in the criteria of the outer query include subqueries in an IN clause, subqueries using the ANY/SOME or ALL predicate quantifier, and subqueries using the EXISTS predicate.

Example 1.2. Example Subquery in WHERE Using EXISTS

```
SELECT a FROM X WHERE EXISTS (SELECT b, c FROM Y WHERE c=3)
```

The following usages of subqueries must each select only one column, but can return any number of rows.

Example 1.3. Example Subqueries in WHERE Clause

```
SELECT a FROM X WHERE a IN (SELECT b FROM Y WHERE c=3)
SELECT a FROM X WHERE a >= ANY (SELECT b FROM Y WHERE c=3)
SELECT a FROM X WHERE a < SOME (SELECT b FROM Y WHERE c=4)
SELECT a FROM X WHERE a = ALL (SELECT b FROM Y WHERE c=2)
```


XML SELECT Command

2.1. Overview

XML documents can be dynamically constructed by Teiid. The structure of the document is defined by a document model, which is generally created from a schema. The document model is bound to relevant SQL statements through mapping classes. See the Designer guide for more on creating document models.

Querying XML documents is similar to querying relational tables. An idiomatic SQL variant with special scalar functions gives control over which parts of a given document to return.

2.2. Query Structure

A valid XML SELECT Command against a document model is of the form *SELECT ... FROM ... [WHERE ...] [ORDER BY ...]*. The use of any other SELECT command clause is not allowed.

The fully qualified name for an XML element is: "model"."document name".[path to element]."element name".

The fully qualified name for an attribute is: "model"."document name".[path to element]."element name".["@]"attribute name"

Partially qualified names for elements and attributes can be used as long as the partial name is unique.

2.2.1. FROM Clause

Specifies the document to generate. Document names resemble other virtual groups - "model"."document name".

Syntax Rules:

- The from may only contain one unary clause specifying the desired document.

2.2.2. SELECT Clause

The select clause determines which parts of the XML document are generated for output.

Example Syntax:

- select * from model.doc
- select model.doc.root.parent.element.* from model.doc
- select element, element1.@attribute from model.doc

Syntax Rules:

- `SELECT *` and `SELECT "xml"` are equivalent and specify that every element and attribute of the document should be output.
- The `SELECT` clause of an XML Query may only contain `*`, `"xml"`, or element and attribute references from the specified document. Any other expressions are not allowed.
- If the `SELECT` clause contains an element or attribute reference (other than `*` or `"xml"`) then only the specified elements, attributes, and their ancestor elements will be in the generated document.
- `element.*` specifies that the element, its attribute, and all child content should be output.

2.2.3. WHERE Clause

The where clause specifies how to filter content from the generated document based upon values contained in the underlying mapping classes. Most predicates are valid in an XML SELECT Command, however combining value references from different parts of the document may not always be allowed.

Criteria is logically applied to a context which is directly related to a mapping class. Starting with the root mapping class, there is a root context that describes all of the top level repeated elements that will be in the output document. Criteria applied to the root or any other context will change the related mapping class query to apply the affects of the criteria, which can include checking values from any of the descendant mapping classes.

Example Syntax:

- `select element, element1.@attribute from model.doc where element1.@attribute = 1`
- `select element, element1.@attribute from model.doc where context(element1, element1.@attribute) = 1`

Syntax Rules:

- Each criteria conjunct must refer to a single context and can be criteria that applies to a mapping class, contain a *rowlimit* function, or contain *rowlimitexception* function.
- Criteria that applies to a mapping class is associated to that mapping class via the *context* function. The absence of a context function implies the criteria applies to the root context.
- At a given context the criteria can span multiple mapping classes provided that all mapping classes involved are either parents of the context, the context itself, or a descendant of the context.



Sibling Root Mapping Classes

Implied root context user criteria against a document model with sibling root mapping classes is not generally semantically correct. It is applied as if each of the

conjuncts is applied to only a single root mapping class. This behavior is the same as prior releases but may be fixed in a future release.

2.2.3.1. XML SELECT Command Specific Functions

XML SELECT Command functions resemble scalar functions, but act as hints in the WHERE clause. These functions are only valid in an XML SELECT Command.

2.2.3.1.1. Context Function

`CONTEXT(arg1, arg2)`

Select the context for the containing conjunct.

Syntax Rules:

- Context functions apply to the whole conjunct.
- The first argument must be an element or attribute reference from the mapping class whose context the criteria conjunct will apply to.
- The second parameter is the return value for the function.

2.2.3.1.2. Rowlimit Function

`ROWLIMIT(arg)`

Limits the rows processed for the given context.

Syntax Rules:

- The first argument must be an element or attribute reference from the mapping class whose context the row limit applies.
- The rowlimit function must be used in equality comparison criteria with the right hand expression equal to a positive integer number or rows to limit.
- Only one row limit or row limit exception may apply to a given context.

2.2.3.1.3. Rowlimitexception Function

Limits the rows processed for the given context and throws an exception if the given number of rows is exceeded.

`ROWLIMITEXCEPTION(arg)`

Syntax Rules:

- The first argument must be an element or attribute reference from the mapping class whose context the row limit exception applies.
- The `rowlimitexception` function must be used in equality comparison criteria with the right hand expression equal to a positive integer number or rows to limit.
- Only one row limit or row limit exception may apply to a given context.

2.2.4. ORDER BY Clause

The XML SELECT Command ORDER BY Clause specifies ordering for the referenced mapping class queries.

Syntax Rules:

- Each order by item must be an element or attribute reference tied a output value from a mapping class.
- The order or the order by items is the relative order they will be applied to their respective mapping classes.

2.3. Document Generation

Document generation starts with the root mapping class and proceeds iteratively and hierarchically over all of the child mapping classes. This can result in a large number of query executions. For example if a document has a root mapping class with 3 child mapping classes. Then for each row selected by the root mapping class after the application of the root context criteria, each of the child mapping classes queries will also be executed.



Document Correctness

By default XML generated XML documents are not checked for correctness vs. the relevant schema. It is possible that the mapping class queries, the usage of specific SELECT or WHERE clause values will generated a document that is not valid with respect to the schema. See [document validation](#) on how to ensure correctness.

Sibling or cousin elements defined by the same mapping class that do not have a common parent in that mapping class will be treated as independent mapping classes during planning and execution. This allows for a more document centric approach to applying criteria and order bys to mapping classes.

2.3.1. Document Validation

The execution property `XMLValidation` should be set to 'true' to indicate that generated documents should be checked for correctness. Correctness checking will not prevent invalid documents from being generated, since correctness is checked after generation and not during.

Datatypes

3.1. Supported Types

Teiid supports a core set of runtime types. Runtime types can be different than semantic types defined in type fields at design time. The runtime type can also be specified at design time or it will be automatically chosen as the closest base type to the semantic type.

Table 3.1. Teiid Runtime Types

Type	Description	Java Runtime Class	JDBC Type	ODBC Type
string	variable length character string with a maximum length of 4000	java.lang.String	VARCHAR	VARCHAR
char	a single Unicode character	java.lang.Character	CHAR	CHAR
boolean	a single bit, or Boolean, that can be true, false, or null (unknown)	java.lang.Boolean	BIT	SMALLINT
byte	numeric, integral type, signed 8-bit	java.lang.Byte	TINYINT	SMALLINT
short	numeric, integral type, signed 16-bit	java.lang.Short	SMALLINT	SMALLINT
integer	numeric, integral type, signed 32-bit	java.lang.Integer	INTEGER	INTEGER
long	numeric, integral type, signed 64-bit	java.lang.Long	BIGINT	NUMERIC
bigint	numeric, integral type, arbitrary precision of up to 1000 digits	java.lang.BigInteger	NUMERIC	NUMERIC
float	numeric, floating point type, 32-bit IEEE 754 floating-point numbers	java.lang.Float	REAL	FLOAT
double	numeric, floating point type, 64-bit IEEE 754 floating-point numbers	java.lang.Double	DOUBLE	DOUBLE
bigdecimal	numeric, floating point type, arbitrary precision of up to 1000 digits	java.math.BigDecimal	NUMERIC	NUMERIC

Type	Description	Java Runtime Class	JDBC Type	ODBC Type
date	datetime, representing a single day (year, month, day)	java.sql.Date	DATE	DATE
time	datetime, representing a single time (hours, minutes, seconds, milliseconds)	java.sql.Time	TIME	TIME
timestamp	datetime, representing a single date and time (year, month, day, hours, minutes, seconds, milliseconds, nanoseconds)	java.sql.Timestamp	TIMESTAMP	TIMESTAMP
object	any arbitrary Java object, must implement java.lang.Serializable	Any	JAVA_OBJECT	VARCHAR
blob	binary large object, representing a stream of bytes	java.sql.Blob ^a	BLOB	VARCHAR
clob	character large object, representing a stream of characters	java.sql.Clob ^b	CLOB	VARCHAR
xml	XML document	java.sql.SQLXML ^c	JAVA_OBJECT	VARCHAR

^aThe concrete type is expected to be com.metamatrix.common.types.BlobType

^bThe concrete type is expected to be com.metamatrix.common.types.ClobType

^cThe concrete type is expected to be com.metamatrix.common.types.XMLType

3.2. Type Conversions

Data types may be converted from one form to another either explicitly or implicitly. Implicit conversions automatically occur in criteria and expressions to ease development. Explicit datatype conversions require the use of the `CONVERT` function or `CAST` keyword.

Type Conversion Considerations

- Any type may be implicitly converted to the `OBJECT` type.
- The `OBJECT` type may be explicitly converted to any other type.
- The `NULL` value may be converted to any type.

- Any valid implicit conversion is also a valid explicit conversion.
- Situations involving literal values that would normally require explicit conversions may have the explicit conversion applied implicitly if no loss of information occurs.
- When Teiid detects that an explicit conversion can not be applied implicitly in criteria, the criteria will be treated as false. For example:

```
SELECT * FROM my.group WHERE created_by = 'not a date'
```

Given that `created_by` is typed as `date`, rather than converting `'not a date'` to a `date` value, the criteria will remain as a string comparison and therefore be false.

- Explicit conversions that are not allowed between two types will result in an exception before execution. Allowed explicit conversions may still fail during processing if the runtime values are not actually convertible.

Table 3.2. Type Conversions

Source Type	Valid Implicit Target Types	Valid Explicit Target Types
string	clob	char, boolean, byte, short, integer, long, biginteger, float, double, bigdecimal, xml
char	string	
boolean	string, byte, short, integer, long, biginteger, float, double, bigdecimal	
byte	string, short, integer, long, biginteger, float, double, bigdecimal	boolean
short	string, integer, long, biginteger, float, double, bigdecimal	boolean, byte
integer	string, long, biginteger, double, bigdecimal	boolean, byte, short, float
long	string, biginteger, bigdecimal	boolean, byte, short, integer, float, double
biginteger	string, bigdecimal	boolean, byte, short, integer, long, float, double
bigdecimal	string	boolean, byte, short, integer, long, biginteger, float, double
date	string, timestamp	
time	string, timestamp	
timestamp	string	date, time

Source Type	Valid Implicit Target Types	Valid Explicit Target Types
clob		string
xml		string

3.3. Special Conversion Cases

3.3.1. Conversion of String Literals

Teiid automatically converts string literals within a SQL statement to their implied types. This typically occurs in a criteria comparison where an element with a different datatype is compared to a literal string:

```
SELECT * FROM my.group WHERE created_by = '2003-01-02'
```

Here if the created_by element has the datatype of date, Teiid automatically converts the string literal to a date datatype as well.

3.3.2. Converting to Boolean

Teiid can automatically convert literal strings and numeric type values to Boolean values as follows:

Type	Literal Value	Boolean Value
String	'false'	false
	'unknown'	null
	other	true
Numeric	0	false
	other	true

3.3.3. Date/Time/Timestamp Type Conversions

Teiid can implicitly convert properly formatted literal strings to their associated date-related datatypes as follows:

String Literal Format	Possible Implicit Conversion Type
yyyy-mm-dd	DATE
hh:mm:ss	TIME
yyyy-mm-dd hh:mm:ss.ffffff ^a	TIMESTAMP

^afractional seconds are optional

The formats above are those expected by the JDBC date types. To use other formats see the functions PARSEDATE , PARSETIME , PARSETIMESTAMP .

3.4. Escaped Literal Syntax

Rather than relying on implicit conversion, datatype values may be expressed directly in SQL using escape syntax to define the type. Note that the supplied string value must match the expected format exactly or an exception will occur.

Table 3.3. Escaped Literal Syntax

Datatype	Escaped Syntax
BOOLEAN	{b'true'} or {b'false'}
DATE	{d'yyyy-mm-dd'}
TIME	{t'hh-mm-ss'}
TIMESTAMP	{ts'yyyy-mm-dd hh:mm:ss.fffffff'} ^a

^afractional seconds are optional

Scalar Functions

Teiid provides an extensive set of built-in scalar functions. See also [SQL Support](#) and [Datatypes](#). In addition, Teiid provides the capability for [user defined functions or UDFs](#).

4.1. Numeric Functions

Numeric functions return numeric values (integer, long, float, double, bigint, bigdecimal). They generally take numeric values as inputs, though some take strings.

Function	Definition	Datatype Constraint
+ - * /	Standard numeric operators	x in {integer, long, float, double, bigint, bigdecimal}, return type is same as x
ABS(x)	Absolute value of x	See standard numeric operators above
ACOS(x)	Arc cosine of x	x in {double, bigdecimal}, return type is double
ASIN(x)	Arc sine of x	x in {double, bigdecimal}, return type is double
ATAN(x)	Arc tangent of x	x in {double, bigdecimal}, return type is double
ATAN2(x,y)	Arc tangent of x and y	x, y in {double, bigdecimal}, return type is double
CEILING(x)	Ceiling of x	x in {double, float}, return type is double
COS(x)	Cosine of x	x in {double, bigdecimal}, return type is double
COT(x)	Cotangent of x	x in {double, bigdecimal}, return type is double
DEGREES(x)	Convert x degrees to radians	x in {double, bigdecimal}, return type is double
EXP(x)	e^x	x in {double, float}, return type is double
FLOOR(x)	Floor of x	x in {double, float}, return type is double
FORMATBIGDECIMAL(x, y)	Formats x using format y	x is bigdecimal, y is string, returns string
FORMATBIGINTEGER(x, y)	Formats x using format y	

Function	Definition	Datatype Constraint
		x is biginteger, y is string, returns string
FORMATDOUBLE(x, y)	Formats x using format y	x is double, y is string, returns string
FORMATFLOAT(x, y)	Formats x using format y	x is float, y is string, returns string
FORMATINTEGER(x, y)	Formats x using format y	x is integer, y is string, returns string
FORMATLONG(x, y)	Formats x using format y	x is long, y is string, returns string
LOG(x)	Natural log of x (base e)	x in {double, float}, return type is double
LOG10(x)	Log of x (base 10)	x in {double, float}, return type is double
MOD(x, y)	Modulus (remainder of x / y)	x in {integer, long, float, double, biginteger, bigdecimal}, return type is same as x
PARSEBIGDECIMAL(x, y)	Parses x using format y	x, y are strings, returns bigdecimal
PARSEBIGINTEGER(x, y)	Parses x using format y	x, y are strings, returns biginteger
PARSEDOUBLE(x, y)	Parses x using format y	x, y are strings, returns double
PARSEFLOAT(x, y)	Parses x using format y	x, y are strings, returns float
PARSEINTEGER(x, y)	Parses x using format y	x, y are strings, returns integer
PARSELONG(x, y)	Parses x using format y	x, y are strings, returns long
PI()	Value of Pi	return is double
POWER(x,y)	x to the y power	x in {double, bigdecimal, biginteger}, return is the same type as x
RADIANS(x)	Convert x radians to degrees	x in {double, bigdecimal}, return type is double
RAND()	Returns a random number, using generator established so far in the query or initializing with system clock if necessary.	Returns double.

Function	Definition	Datatype Constraint
RAND(x)	Returns a random number, using new generator seeded with x.	x is integer, returns double.
ROUND(x,y)	Round x to y places; negative values of y indicate places to the left of the decimal point	x in {integer, float, double, bigdecimal} y is integer, return is same type as x
SIGN(x)	1 if x > 0, 0 if x = 0, -1 if x < 0	x in {integer, long, float, double, biginteger, bigdecimal}, return type is integer
SIN(x)	Sine value of x	x in {double, bigdecimal}, return type is double
SQRT(x)	Square root of x	x in {long, double, bigdecimal}, return type is double
TAN(x)	Tangent of x	x in {double, bigdecimal}, return type is double
BITAND(x, y)	Bitwise AND of x and y	x, y in {integer}, return type is integer
BITOR(x, y)	Bitwise OR of x and y	x, y in {integer}, return type is integer
BITXOR(x, y)	Bitwise XOR of x and y	x, y in {integer}, return type is integer
BITNOT(x)	Bitwise NOT of x	x in {integer}, return type is integer

4.1.1. Parsing Numeric Datatypes from Strings

Teiid offers a set of functions you can use to parse numbers from strings. For each string, you need to provide the formatting of the string. These functions use the convention established by the `java.text.DecimalFormat` class to define the formats you can use with these functions. You can learn more about how this class defines numeric string formats by visiting the Sun Java Web site at the following [URL for Sun Java](http://java.sun.com/javase/6/docs/api/java/text/DecimalFormat.html) [http://java.sun.com/javase/6/docs/api/java/text/DecimalFormat.html].

For example, you could use these function calls, with the formatting string that adheres to the `java.text.DecimalFormat` convention, to parse strings and return the datatype you need:

Input String	Function Call to Format String	Output Value	Output Datatype
'\$25.30'	<code>parseDouble(cost, '\$#,##0.00;(\$#,##0.00)')</code>	25.3	double

Input String	Function Call to Output Value Format String	Output Value	Output Datatype
'25%'	parseFloat(percent, '#,##0%')	25	float
'2,534.1'	parseFloat(total, '#,##0.###;-#,##0.###')	2534.1	float
'1.234E3'	parseLong(amt, '0.###E0')	1234	long
'1,234,567'	parseInteger(total, '#,##0;-#,##0')	1234567	integer

4.1.2. Formatting Numeric Datatypes as Strings

Teiid offers a set of functions you can use to convert numeric datatypes into strings. For each string, you need to provide the formatting. These functions use the convention established within the `java.text.DecimalFormat` class to define the formats you can use with these functions. You can learn more about how this class defines numeric string formats by visiting the Sun Java Web site at the following [URL for Sun Java](http://java.sun.com/javase/6/docs/api/java/text/DecimalFormat.html) [http://java.sun.com/javase/6/docs/api/java/text/DecimalFormat.html] .

For example, you could use these function calls, with the formatting string that adheres to the `java.text.DecimalFormat` convention, to format the numeric datatypes into strings:

Input Value	Input Datatype	Function Call to Output String Format String	Output String
25.3	double	formatDouble(cost, '\$#,##0.00;(\$#,##0.00)')	'\$25.30'
25	float	formatFloat(percent, '#,##0%')	'25%'
2534.1	float	formatFloat(total, '#,##0.###;-#,##0.###')	'2,534.1'
1234	long	formatLong(amt, '0.###E0')	'1.234E3'
1234567	integer	formatInteger(total, '#,##0;-#,##0')	'1,234,567'

4.2. String Functions

String functions generally take strings as inputs and return strings as outputs.

Unless specified, all of the arguments and return types in the following table are strings and all indexes are 1-based. The 0 index is considered to be before the start of the string.

Function	Definition	Datatype Constraint
<code>x y</code>	Concatenation operator	<code>x,y</code> in {string}, return type is string
<code>ASCII(x)</code>	Provide ASCII value of the left most character in <code>x</code> . The empty string will as input will return null.	return type is integer
<code>CHR(x) CHAR(x)</code>	Provide the character for ASCII value <code>x</code>	<code>x</code> in {integer}
<code>CONCAT(x, y)</code>	Concatenates <code>x</code> and <code>y</code> with ANSI semantics. If <code>x</code> and/or <code>y</code> is null, returns null.	<code>x, y</code> in {string}
<code>CONCAT2(x, y)</code>	Concatenates <code>x</code> and <code>y</code> with non-ANSI null semantics. If <code>x</code> and <code>y</code> is null, returns null. If only <code>x</code> or <code>y</code> is null, returns the other value.	<code>x, y</code> in {string}
<code>INITCAP(x)</code>	Make first letter of each word in string <code>x</code> capital and all others lowercase	<code>x</code> in {string}
<code>INSERT(str1, start, length, str2)</code>	Insert <code>string2</code> into <code>string1</code>	<code>str1</code> in {string}, <code>start</code> in {integer}, <code>length</code> in {integer}, <code>str2</code> in {string}
<code>LCASE(x)</code>	Lowercase of <code>x</code>	<code>x</code> in {string}
<code>LEFT(x, y)</code>	Get left <code>y</code> characters of <code>x</code>	<code>x</code> in {string}, <code>y</code> in {string}, return string
<code>LENGTH(x)</code>	Length of <code>x</code>	return type is integer
<code>LOCATE(x, y)</code>	Find position of <code>x</code> in <code>y</code> starting at beginning of <code>y</code>	<code>x</code> in {string}, <code>y</code> in {string}, return integer
<code>LOCATE(x, y, z)</code>	Find position of <code>x</code> in <code>y</code> starting at <code>z</code>	<code>x</code> in {string}, <code>y</code> in {string}, <code>z</code> in {integer}, return integer
<code>LPAD(x, y)</code>	Pad input string <code>x</code> with spaces on the left to the length of <code>y</code>	<code>x</code> in {string}, <code>y</code> in {integer}, return string
<code>LPAD(x, y, z)</code>	Pad input string <code>x</code> on the left to the length of <code>y</code> using character <code>z</code>	<code>x</code> in {string}, <code>y</code> in {string}, <code>z</code> in {character}, return string
<code>LTRIM(x)</code>	Left trim <code>x</code> of white space	<code>x</code> in {string}, return string
<code>REPEAT(str1,instances)</code>	Repeat <code>string1</code> a specified number of times	<code>str1</code> in {string}, <code>instances</code> in {integer} return string

Function	Definition	Datatype Constraint
REPLACE(x, y, z)	Replace all y in x with z	x,y,z in {string}, return string
RIGHT(x, y)	Get right y characters of x	x in {string}, y in {string}, return string
RPAD(input string x, pad length y)	Pad input string x with spaces on the right to the length of y	x in {string}, y in {integer}, return string
RPAD(x, y, z)	Pad input string x on the right to the length of y using character z	x in {string}, y in {string}, z in {character}, return string
RTRIM(x)	Right trim x of white space	x is string, return string
SUBSTRING(x, y)	Get substring from x, from position y to the end of x	y in {integer}
SUBSTRING(x, y, z)	Get substring from x from position y with length z	y, z in {integer}
TRANSLATE(x, y, z)	Translate string x by replacing each character in y with the character in z at the same position	x in {string}
UCASE(x)	Uppercase of x	x in {string}

4.3. Date/Time Functions

Date and time functions return or operate on dates, times, or timestamps.

Parse and format Date/Time functions use the convention established within the `java.text.SimpleDateFormat` class to define the formats you can use with these functions. You can learn more about how this class defines formats by visiting the Sun Java Web site at the following [URL for Sun Java](http://java.sun.com/javase/6/docs/api/java/text/SimpleDateFormat.html) [http://java.sun.com/javase/6/docs/api/java/text/SimpleDateFormat.html].

Function	Definition	Datatype Constraint
CURDATE()	Return current date	returns date
CURTIME()	Return current time	returns time
NOW()	Return current timestamp (date and time)	returns timestamp
DAYNAME(x)	Return name of day	x in {date, timestamp}, returns string
DAYOFMONTH(x)	Return day of month	x in {date, timestamp}, returns integer
DAYOFWEEK(x)	Return day of week (Sunday=1)	x in {date, timestamp}, returns integer
DAYOFYEAR(x)	Return Julian day number	

Function	Definition	Datatype Constraint
		x in {date, timestamp}, returns integer
FORMATDATE(x, y)	Format date x using format y	x is date, y is string, returns string
FORMATTIME(x, y)	Format time x using format y	x is time, y is string, returns string
FORMATTIMESTAMP(x, y)	Format timestamp x using format y	x is timestamp, y is string, returns string
HOUR(x)	Return hour (in military 24-hour format)	x in {time, timestamp}, returns integer
MINUTE(x)	Return minute	x in {time, timestamp}, returns integer
MONTH(x)	Return month	x in {date, timestamp}, returns integer
MONTHNAME(x)	Return name of month	x in {date, timestamp}, returns string
QUARTER(x)	Return quarter	x in {date, timestamp}, returns integer
PARSEDATE(x, y)	Parse date from x using format y	x, y in {string}, returns date
PARSETIME(x, y)	Parse time from x using format y	x, y in {string}, returns time
PARSETIMESTAMP(x,y)	Parse timestamp from x using format y	x, y in {string}, returns timestamp
SECOND(x)	Return seconds	x in {time, timestamp}, returns integer
TIMESTAMPCREATE(date, time)	Create a timestamp from a date and time	date in {date}, time in {time}, returns timestamp
TIMESTAMPADD(interval, count, timestamp)	Add a specified interval (hour, day of week, month) to the timestamp, where intervals can have the following definition: 1. SQL_TSI_FRAC_SECOND - fractional seconds (billionths of a second)	The interval constant may be specified either as a string literal or a constant value. Interval in {string}, count in {integer}, timestamp in {date, time, timestamp}

Function	Definition	Datatype Constraint
	2. SQL_TSI_SECOND - seconds 3. SQL_TSI_MINUTE - minutes 4. SQL_TSI_HOUR - hours 5. SQL_TSI_DAY - days 6. SQL_TSI_WEEK - weeks 7. SQL_TSI_MONTH - months 8. SQL_TSI_QUARTER - quarters (3 months) 9. SQL_TSI_YEAR - years	
TIMESTAMPDIFF(interval, startTime, endTime)	Calculate the approximate number of whole intervals in (endTime - startTime) using a specific interval type (as defined by the constants in TIMESTAMPADD). If (endTime > startTime), a positive number will be returned. If (endTime < startTime), a negative number will be returned. Calculations are approximate and may be less accurate over longer time spans.	Interval in {string}; startTime, endTime in {date, time, timestamp}, returns a long.
WEEK(x)	Return week in year	x in {date, timestamp}, returns integer
YEAR(x)	Return four-digit year	x in {date, timestamp}, returns integer
MODIFYTIMEZONE(timestamp, startTimeZone, endTimeZone)	Returns a timestamp based upon the incoming timestamp adjusted for the differential between the start and end time zones. i.e. if the server is in GMT-6, then modifytimezone({ts '2006-	startTimeZone and endTimeZone are strings, returns a timestamp

Function	Definition	Datatype Constraint
	01-10 04:00:00.0'}, 'GMT-7', 'GMT-8') will return the timestamp {ts '2006-01-10 05:00:00.0'} as read in GMT-6. The value has been adjusted 1 hour ahead to compensate for the difference between GMT-7 and GMT-8.	
MODIFYTIMEZONE (timestamp, endTimeZone)	Return a timestamp in the same manner as modifytimezone(timestamp, startTimeZone, endTimeZone), but will assume that the startTimeZone is the same as the server process.	Timestamp is a timestamp; endTimeZone is a string, returns a timestamp
FROM_UNIXTIME (unix_timestamp)	Return the Unix timestamp (in seconds) as a Timestamp value	Unix timestamp (in seconds)

4.3.1. Parsing Date Datatypes from Strings

Teiid does not implicitly convert strings that contain dates presented in different formats, such as '19970101' and '31/1/1996' to date-related datatypes. You can, however, use the `parseDate`, `parseTime`, and `parseTimestamp` functions, described in the next section, to explicitly convert strings with a different format to the appropriate datatype. These functions use the convention established within the `java.text.SimpleDateFormat` class to define the formats you can use with these functions. You can learn more about how this class defines date and time string formats by visiting the [Sun Java Web site](http://java.sun.com/j2se/1.4.2/docs/api/java/text/SimpleDateFormat.html) [http://java.sun.com/j2se/1.4.2/docs/api/java/text/SimpleDateFormat.html] .

For example, you could use these function calls, with the formatting string that adheres to the `java.text.SimpleDateFormat` convention, to parse strings and return the datatype you need:

String	Function Call To Parse String
'19970101'	<code>parseDate(myDateString, 'yyyyMMdd')</code>
'31/1/1996'	<code>parseDate(myDateString, 'dd"/"MM"/"yyyy')</code>
'22:08:56 CST'	<code>parseTime (myTime, 'HH:mm:ss z')</code>
'03.24.2003 at 06:14:32'	<code>parseTimestamp(myTimestamp, 'MM.dd.yyyy "at" hh:mm:ss')</code>

4.3.2. Specifying Time Zones

Time zones can be specified in several formats. Common abbreviations such as EST for "Eastern Standard Time" are allowed but discouraged, as they can be ambiguous. Unambiguous time zones are defined in the form continent or ocean/largest city. For example, America/New_York, America/Buenos_Aires, or Europe/London. Additionally, you can specify a custom time zone by GMT offset: GMT[+/-]HH:MM.

For example: GMT-05:00

4.4. Type Conversion Functions

Within your queries, you can convert between datatypes using the CONVERT or CAST keyword. See also [Data Type Conversions](#) .

Function	Definition
CONVERT(x, type)	Convert x to type, where type is a Teiid Base Type
CAST(x AS type)	Convert x to type, where type is a Teiid Base Type

These functions are identical other than syntax; CAST is the standard SQL syntax, CONVERT is the standard JDBC/ODBC syntax.

4.5. Choice Functions

Choice functions provide a way to select from two values based on some characteristic of one of the values.

Function	Definition	Datatype Constraint
IFNULL(x,y)	If x is null, return y; else return x	x, y, and the return type must be the same type but can be any type
NVL(x,y)	If x is null, return y; else return x	x, y, and the return type must be the same type but can be any type
NULLIF(param1, param2)	Equivalent to case when (param1 = param2) then null else param1	param1 and param2 must be compatible comparable types
COALESCE(x,y+)	Returns the first non-null parameter	x and all y's can be any compatible types

IFNULL and NVL are aliases of each other. They are the same function.

4.6. Decode Functions

Decode functions allow you to have the Teiid Server examine the contents of a column in a result set and alter, or decode, the value so that your application can better use the results.

Function	Definition	Datatype Constraint
DECODESTRING(x, y)	Decode column x using string of value pairs y and return the decoded column as a string	all string
DECODESTRING(x, y, z)	Decode column x using string of value pairs y with delimiter z and return the decoded column as a string	all string
DECODEINTEGER(x, y)	Decode column x using string of value pairs y and return the decoded column as an integer	all string parameters, return integer
DECODEINTEGER(x,y,z)	Decode column x using string of value pairs y with delimiter z and return the decoded column as an integer	all string parameters, return integer

Within each function call, you include the following arguments:

1. x is the input value for the decode operation. This will generally be a column name.
2. y is the literal string that contains a delimited set of input values and output values.
3. z is an optional parameter on these methods that allows you to specify what delimiter the string specified in y uses.

For example, your application might query a table called PARTS that contains a column called IS_IN_STOCK which contains a Boolean value that you need to change into an integer for your application to process. In this case, you can use the DECODEINTEGER function to change the Boolean values to integers:

```
SELECT DECODEINTEGER(IS_IN_STOCK, 'false, 0, true, 1') FROM PartsSupplier.PARTS;
```

When the Teiid System encounters the value false in the result set, it replaces the value with 0.

If, instead of using integers, your application requires string values, you can use the DECODESTRING function to return the string values you need:

```
SELECT DECODESTRING(IS_IN_STOCK, 'false, no, true, yes, null') FROM
PartsSupplier.PARTS;
```

In addition to two input/output value pairs, this sample query provides a value to use if the column does not contain any of the preceding input values. If the row in the IS_IN_STOCK column does not contain true or false, the Teiid Server inserts a null into the result set.

When you use these DECODE functions, you can provide as many input/output value pairs if you want within the string. By default, the Teiid System expects a comma delimiter, but you can add a third parameter to the function call to specify a different delimiter:

```
SELECT DECODESTRING(IS_IN_STOCK, 'false:no:true:yes:null',':') FROM
PartsSupplier.PARTS;
```

You can use keyword null in the DECODE string as either an input value or an output value to represent a null value. However, if you need to use the literal string null as an input or output value (which means the word null appears in the column and not a null value) you can put the word in quotes: "null".

```
SELECT DECODESTRING( IS_IN_STOCK, 'null,no,"null",no,nil,no,false,no,true,yes' ) FROM
PartsSupplier.PARTS;
```

If the DECODE function does not find a matching output value in the column and you have not specified a default value, the DECODE function will return the original value the Teiid Server found in that column.

4.7. Lookup Function

The Lookup function allows you to cache a group's data in memory and access it through a scalar function. This caching accelerates response time to queries that use the lookup groups, known in business terminology as lookup tables or code groups.

A StatePostalCodes group used to translate postal codes to complete state names might represent an example of this type of lookup group. One element, PostalCode, represents a key element. Other groups within the model refer to this two-letter code. A second element, StateDisplayName, would represent the complete name of the state. Hence, a query to this lookup group would typically provide the PostalCode and expect the StateDisplayName in response.

When you call this function for any combination of codeGroup, returnElement, and keyElement for the first time, the Teiid System caches the result. The Teiid System uses this cached map for all queries, in all sessions, that later access this lookup group. The codeGroup requires use of the

fully-qualified name, and the returnElement and keyElement parameters should use shortened column names.

Because the Teiid System caches and indexes this information in memory, this function provides quick access after the Teiid System initially caches the lookup group. The Teiid System unloads these cached lookup groups when you stop and restart the Teiid System. Thus, you should not use this function for data that is subject to updates. Instead, you can use it against static data that does not change over time.



Note

- The keyElement column is expected to contain unique key values. If the column contains duplicate values, only the last loaded value will be used for lookup purposes. In some cases, this may cause unexpected results, so it is strongly recommended that only columns without duplicate values be used as the keyElement. The lookup caches can be flushed via the svcmgr.
- Cached lookup groups might consume significant memory. You can limit the number and maximum size of these code groups by setting properties of the QueryService through the Teiid Console.

Function	Definition	Datatype Constraint
LOOKUP(codeGroup, returnElement, keyElement, keyValue)	In the lookup group codeGroup, find the row where keyElement has the value keyValue and return the associated returnElement	codeGroup must be a fully-qualified string literal containing metadata identifiers, keyValue datatype must match datatype of the keyElement, return datatype matches that of returnElement. returnElement and keyElement parameters should use their shortened names.

4.7.1. Clearing the Cache

You can force a cache clearing by using the expert mode of the svcmgr command, found under the \bin directory of your Teiid server installation.

Launch the appropriate command:

1. `svcmgr.cmd` (Windows)
2. `svcmgr.sh` (Solaris or Linux)

From the command line enter `ClearCodeTableCaches`.

4.8. System Functions

System functions provide access to information in the Teiid system from within a query.

Function	Definition	Datatype Constraint
<code>USER()</code>	Retrieve the name of the user executing the query	return is string
<code>ENV(key)</code>	Retrieve an environment property. The only key currently allowed is 'sessionid', although this will expand in the future.	key in {string}, return is string
<code>COMMANDPAYLOAD()</code>	Retrieve the string form of the command payload or null if no command payload was specified. The command payload is set by a method on the Teiid JDBC API extensions on a per-query basis.	Returns a string
<code>COMMANDPAYLOAD(key)</code>	Cast the command payload object to a <code>java.util.Properties</code> object and look up the specified key in the object	key in {string}, return is string

4.9. XML Functions

XML functions provide functionality for working with XML data.

Function	Definition	Datatype Constraint
<code>XPATHVALUE(doc, xpath)</code>	Takes a document and an XPATH query and returns a string value for the result. An attempt is made to provide a meaningful result for non-text nodes.	Doc in {string, xml} and xpath in {string}. Return value is a string.

4.10. Security Functions

Security functions provide the ability to interact with the security system.

Function	Definition	Datatype Constraint
<code>hasRole(roleType, roleName)</code>	Whether the current caller has the role <code>roleName</code> of <code>roleType</code> .	<code>roleType</code> must be one of ('data','admin' , 'repository')

Function	Definition	Datatype Constraint
		and roleName must be a string, the return type is Boolean.

4.11. User Defined Functions

If you need to extend Teiid's scalar function library, then Teiid provides a means to define custom scalar functions or User Defined Functions(UDF). The following steps need to be taken in creating a UDF.

4.11.1. UDF Definition

The FunctionDefinition.xmi file provides metadata to the query engine on User Defined Functions. See our product document on "Creating User-defined Functions" for a more extensive reference on creating that file through the Designer Tool.

The following are used to define a UDF.

- *Function Name* When you create the function name, keep these requirements in mind:
 - You cannot use a reserved word, which includes existing Teiid System function names. You cannot overload existing Teiid System functions.
 - The function name must be unique among user-defined functions for the number of arguments. You can use the same function name for different numbers of types of arguments. Hence, you can overload your user-defined functions.
 - The function name can only contain letters, numbers, and the underscore (_). Your function name must start with a letter.
 - The function name cannot exceed 128 characters.
- *Input Parameters* - defines a type specific signature list. All arguments are considered required.
- *Return Type* - the expected type of the returned scalar value.
- *Pushdown* - can be one of REQUIRED, NEVER, ALLOWED. Indicates the expected pushdown behavior. If NEVER or ALLOWED are specified then a Java implementation of the function should be supplied.
- *invocationClass/invocationMethod* - optional properties indicating the static method to invoke when the UDF is not pushed down.
- *Deterministic* - if the method will always return the same result for the same input parameters.

4.11.2. Source Supported UDF

While Teiid provides an extensive scalar function library, it contains only those functions that can be evaluated within the query engine. In many circumstances, especially for performance, a user defined function allows for calling a source specific function.

For example, suppose you want to use the Oracle-specific functions `score` and `contains`:

```
SELECT score(1), ID, FREEDATA FROM Docs WHERE contains(freedata, 'nick', 1) > 0
```

The `score` and `contains` functions are not part of built-in scalar function library. While you could write your own custom scalar function to mimic their behavior, it's more likely that you would want to use the actual Oracle functions that are provided by Oracle when using the Oracle Free Text functionality.

In addition to the normal steps outlined in the section to create and install a function model (FunctionDefinitions.xmi), you will need to extend the appropriate connector(s).

For example, to extend the Oracle Connector

- *Required* - extend OracleCapabilities and set up SCORE and CONTAINS as supported functions (this lets Teiid know that the connector can accept these functions).
- Optionally extend the OracleSQLTranslator to insert new FunctionModifiers to handle translation of these functions. Given that the syntax of these functions is same as other typical functions, this probably isn't needed - the default translation should work.
- Create a new connector type - the easiest way is to export the Oracle ANSI connector type from the Console and just modify the properties such as the connector name (to differentiate it from base Oracle connector) and the capabilities class (to use the extended version) and possibly the translation class (if that was extended for b. Also, connector classpath needs to be extended to include a new jar of your changes above.
- Install the code as an extension module and add your new connector type in the Console.

4.11.3. Non-pushdown Support for User-Defined Functions

Non-pushdown support requires a Java function that matches the metadata supplied in the FunctionDefinitions.xmi file. You must create a Java method that contains the function's logic. This Java method should accept the necessary arguments, which the Teiid System will pass to it at runtime, and function should return the calculated or altered value.

4.11.3.1. Java Code

Code Requirements

- The java class containing the function method must be defined public.
- The function method must be public and static.
- Number of input arguments must match the function metadata defined in section [Install user-defined functions](#)
- Any exception can be thrown, but Teiid will rethrow the exception as a `FunctionExecutionException`.

Example 4.1. Sample code

```
package userdefinedfunctions;

public class TempConv {

    /**
     * Converts the given Celsius temperature to Fahrenheit, and returns the
     * value.
     * @param doubleCelsiusTemp
     * @return Fahrenheit
     */
    public static Double celsiusToFahrenheit(Double doubleCelsiusTemp){
        if (doubleCelsiusTemp == null) {
            return null;
        }
        return (doubleCelsiusTemp)*9/5 + 32;
    }
}
```

4.11.3.2. Post Code Activities

1. After coding the functions you should compile the Java code into a Java Archive (JAR) file, so that you can add it to the Teiid System as an Extension Module.
2. After adding the jar file as an extension module, the name of jar file need to be added to user defined functions classpath using Console tool.

4.11.4. Installing user-defined functions

Once a user-defined function model (FunctionDefinitions.xmi) has been created in in the Designer Tool, it should be installed by replacing the existing version under the Extension Modules (for the Enterprise product this will be done through the Console). That will allow the query engine to know about and use functions

4.12. Nondeterministic Function Handling

Teiid categorizes functions by varying degrees of determinism. When a function is evaluated and to what extent the result can be cached are based upon its determinism level.

1. **Deterministic** - the function will always return the same result for the given inputs. Deterministic functions are evaluated by the engine as soon as all input values are known, which may occur as soon as the rewrite phase. Some functions, such as the lookup function, are not truly deterministic, but is treated as such for performance. All functions not categorized below are considered deterministic.
2. **Session Deterministic** - the function will return the same result for the given inputs under the same user session. This category includes the `hasRole`, `env`, and `user` functions. Session deterministic functions are evaluated by the engine as soon as all input values are known, which may occur as soon as the rewrite phase. If a session deterministic function is evaluated during the creation of a prepared processing plan, then the resulting plan will be cached only for the user's session.
3. **Command Deterministic** - the result of function evaluation is only deterministic within the scope of the user command. This category include the `curdate`, `curtime`, `now`, and `commandpayload` functions. Command deterministic functions are delayed in evaluation until processing to ensure that even prepared plans utilizing these functions will be executed with relevant values. Command deterministic function evaluation will occur prior to pushdown - however multiple occurrences of the same command deterministic time function are not guaranteed to evaluate to the same value.
4. **Nondeterministic** - the result of function evaluation is fully nondeterministic. This category includes the `rand` function and UDFs marked as nondeterministic. Nondeterministic functions are delayed in evaluation until processing with a preference for pushdown. If the function is not pushed down, then it may be evaluated for every row in it's execution context (for example if the function is used in the select clause).

Procedures

5.1. Procedure Language

Teiid supports a procedural language for defining *virtual procedures*. These are similar to stored procedures in relational database management systems. You can use this language to define the transformation logic for decomposing INSERT, UPDATE, and DELETE commands against virtual tables; these are known as *update procedures*.

5.1.1. Command Statement

A command statement executes a *SQL command*, such as SELECT, INSERT, UPDATE, DELETE, or EXECUTE, against one or more other models (and their underlying data sources).

Example 5.1. Example Command Statements

```
SELECT * FROM MyModel.MyTable WHERE ColA > 100;
INSERT INTO MyModel.MyTable (ColA,ColB) VALUES (50, 'hi');
```

5.1.2. Dynamic SQL Command

Dynamic SQL allows for the execution of an arbitrary SQL command in a virtual procedure. Dynamic SQL is useful in situations where the exact command form is not known prior to execution.

Usage:

```
EXECUTE STRING <expression> [AS <variable> <type> [, <variable> <type>]*
    [INTO <variable>]]
[USING <variable>=<expression> [,<variable>=<expression>]*] [UPDATE
    <literal>]
```

Syntax Rules:

- The "AS" clause is used to define the projected symbols names and types returned by the executed SQL string. The "AS" clause symbols will be matched positionally with the symbols returned by the executed SQL string. Non-convertible types or too few columns returned by the executed SQL string will result in an error.
- The "INTO" clause will project the dynamic SQL into the specified temp table. With the "INTO" clause specified, the dynamic command will actually execute a statement that behaves like an

INSERT with a QUERY EXPRESSION. If the dynamic SQL command creates a temporary table with the "INTO" clause, then the "AS" clause is required to define the table's metadata.

- The "USING" clause allows the dynamic SQL string to contain special element symbols that are bound at runtime to specified values. This allows for some independence of the SQL string from the surrounding procedure variable names and input names. In the dynamic command "USING" clause, each variable is specified by short name only. However in the dynamic SQL the "USING" variable must be fully qualified to "USING.". The "USING" clause is only for values that will be used in the dynamic SQL as legal expressions. It is not possible to use the "USING" clause to replace table names, keywords, etc. This makes using symbols equivalent in power to normal bind (?) expressions in prepared statements. The "USING" clause helps reduce the amount of string manipulation needed. If a reference is made to a USING symbol in the SQL string that is not bound to a value in the "USING" clause, an exception will occur.
- The "UPDATE" clause is used to specify the *updating model count*. Accepted values are (0,1,*). 0 is the default value if the clause is not specified.

Example 5.2. Example Dynamic SQL

```
...
/* Typically complex criteria would be formed based upon inputs to the procedure.
In this simple example the criteria is references the using clause to isolate
the SQL string from referencing a value from the procedure directly */
DECLARE string criteria = 'Customer.Accounts.Last = USING.LastName';
/* Now we create the desired SQL string */
DECLARE string sql_string = 'SELECT ID, First || " " || Last AS Name, Birthdate FROM
Customer.Accounts WHERE ' || criteria;
/* The execution of the SQL string will create the #temp table with the columns (ID, Name,
Birthdate).
Note that we also have the USING clause to bind a value to LastName, which is referenced in
the criteria. */
EXECUTE STRING sql_string AS ID integer, Name string, Birthdate date INTO #temp USING
LastName='some name';
/* The temp table can now be used with the values from the Dynamic SQL */
loop on (SELECT ID from #temp) as myCursor
...
```

Here is an example showing a more complex approach to building criteria for the dynamic SQL string. In short, the virtual procedure AccountAccess.GetAccounts has inputs ID, LastName, and bday. If a value is specified for ID it will be the only value used in the dynamic SQL criteria. Otherwise if a value is specified for LastName the procedure will detect if the value is a search string. If bday is specified in addition to LastName, it will be used to form compound criteria with LastName.

Example 5.3. Example Dynamic SQL with USING clause and dynamically built criteria string

```

...
DECLARE string crit = null;
IF (AccountAccess.GetAccounts.ID IS NOT NULL)
  crit = '(Customer.Accounts.ID = using.ID)';
ELSE IF (AccountAccess.GetAccounts.LastName IS NOT NULL)
BEGIN
  IF (AccountAccess.GetAccounts.LastName == '%')
    ERROR "Last name cannot be %";
  ELSE IF (LOCATE('%', AccountAccess.GetAccounts.LastName) < 0)
    crit = '(Customer.Accounts.Last = using.LastName)';
  ELSE
    crit = '(Customer.Accounts.Last LIKE using.LastName)';
  IF (AccountAccess.GetAccounts.bday IS NOT NULL)
    crit = '(' || crit || ' and (Customer.Accounts.Birthdate = using.BirthDay)';
END
ELSE
  ERROR "ID or LastName must be specified.";
EXECUTE STRING 'SELECT ID, First || " " || Last AS
  Name, Birthdate FROM Customer.Accounts WHERE ' || crit USING
  ID=AccountAccess.GetAccounts.ID, LastName=AccountAccess.GetAccounts.LastName,
  BirthDay=AccountAccess.GetAccounts.Bday;
...

```

Known Limitations and Work-Arounds

- The use of dynamic SQL command results in an assignment statement requires the use of a temp table.

Example 5.4. Example Assignment

```

EXECUTE STRING <expression> AS x string INTO #temp;
DECLARE string VARIABLES.RESULT = SEELCT x FROM #temp;

```

- The construction of appropriate criteria will be cumbersome if parts of the criteria are not present. For example if "criteria" were already NULL, then the following example results in "criteria" remaining NULL.

Example 5.5. Example Dangerous NULL handling

```
...
criteria = '(' || criteria || ' and (Customer.Accounts.Birthdate = using.BirthDay)');
```

The preferred approach is for the user to ensure the criteria is not NULL prior its usage. If this is not possible, a good approach is to specify a default as shown in the following example.

Example 5.6. Example NULL handling

```
...
criteria = '(' || nvl(criteria, '(1 = 1)') || ' and (Customer.Accounts.Birthdate = using.BirthDay)');
```

- If the dynamic SQL is an UPDATE, DELETE, or INSERT command, and the user needs to specify the "AS" clause (which would be the case if the number of rows effected needs to be retrieved). The user will still need to provide a name and type for the return column if the into clause is specified.

Example 5.7. Example with AS and INTO clauses

```
/* This name does not need to match the expected update command symbol "count". */
EXECUTE STRING <expression> AS x integer INTO #temp;
```

- Unless used in other parts of the procedure, tables in the dynamic command will not be seen as sources in the Designer.
- When using the "AS" clause only the type information will be available to the Designer. ResultSet columns generated from the "AS" clause then will have a default set of properties for length, precision, etc.

5.1.3. Declaration Statement

A declaration statement declares a variable and its type. After you declare a variable, you can use it in that block within the procedure and any sub-blocks. A variable is initialized to null by default, but can also be assigned the value of an expression as part of the declaration statement.

Usage:

```
DECLARE <type> [VARIABLES.]<name> [= <expression>];
```

Example Syntax

- declare integer x;
- declare string VARIABLES.myvar = 'value';

Syntax Rules:

- You cannot redeclare a variable with a duplicate name in a sub-block
- The VARIABLES group is always implied even if it is not specified.

5.1.4. Assignment Statement

An assignment statement assigns a value to a variable by either evaluating an expression or executing a SELECT command that returns a column value from a single row.

Usage:

```
<variable reference> = <expression>;
```

Example Syntax

- myString = 'Thank you';
- VARIABLES.x = SELECT Column1 FROM MyModel.MyTable;

5.1.5. If Statement

An IF statement evaluates a condition and executes either one of two blocks depending on the result. You can nest IF statements to create complex branching logic. A dependent ELSE statement will execute its block of code only if the IF statement evaluates to false.

Example 5.8. Example If Statement

```
IF ( var1 = 'North America')  
BEGIN  
  ...statement...  
END ELSE  
BEGIN  
  ...statement...  
END
```



Note

NULL values should be considered in the criteria of an IF statement. IS NULL criteria can be used to detect the presence of a NULL value.

5.1.6. Loop Statement

A LOOP statement is an iterative control construct that is used to cursor through a result set.

Usage:

```
LOOP ON <select statement> AS <cursorname>
BEGIN
  ...
END
```

5.1.7. While Statement

A WHILE statement is an iterative control construct that is used to execute a set of statements repeatedly whenever a specified condition is met.

Usage:

```
WHILE <criteria>
BEGIN
  ...
END
```

5.1.8. Continue Statement

A CONTINUE statement is used inside a LOOP or WHILE construct to continue with the next loop by skipping over the rest of the statements in the loop. It must be used inside a LOOP or WHILE statement.

5.1.9. Break Statement

A BREAK statement is used inside a LOOP or WHILE construct to break from the loop. It must be used inside a LOOP or WHILE statement.

5.1.10. Error Statement

An ERROR statement declares that the procedure has entered an error state and should abort. This statement will also roll back the current transaction, if one exists. Any valid expression can be specified after the ERROR keyword.

Example 5.9. Example Error Statement

```
ERROR 'Invalid input value: ' || nvl(Acct.GetBalance.AcctID, 'null');
```

5.2. Virtual Procedures

Virtual procedures are defined using the Teiid procedural language. A virtual procedure has zero or more input parameters, and a result set return type. Virtual procedures support the ability to execute queries and other SQL commands, define temporary tables, add data to temporary tables, walk through result sets, use loops, and use conditional logic.

5.2.1. Virtual Procedure Definition

Usage:

```
CREATE VIRTUAL PROCEDURE
BEGIN
  . . .
END
```

The CREATE VIRTUAL PROCEDURE line demarcates the beginning of the procedure. The BEGIN and END keywords are used to denote block boundaries. Within the body of the procedure, any valid *statement* may be used.

The last command statement executed in the procedure will be return as the result. The output of that statement must match the expected result set and parameters of the procedure.

5.2.2. Procedure Input Parameters

Virtual procedures can take zero or more input parameters. Each input has the following information that is used during runtime processing:

- Name - The name of the input parameter
- Datatype - The design-time type of the input parameter
- Default value - The default value if the input parameter is not specified
- Nullable - NO_NULLS, NULLABLE, NULLABLE_UNKNOWN; parameter is optional if nullable, and is not required to be listed when using named parameter syntax

You reference an input to a virtual procedure by using the fully-qualified name of the param (or less if unambiguous). For example, MyModel.MyProc.Param1.

Example 5.10. Example of Referencing an Input Parameter for 'GetBalance' Procedure

```
CREATE VIRTUAL PROCEDURE
BEGIN
```

```
SELECT Balance FROM MyModel.Accts WHERE MyModel.Accts.AccountID =  
MyModel.GetBalance.AcctID;  
END
```

5.2.3. Example Virtual Procedures

This example is a LOOP that walks through a cursored table and uses CONTINUE and BREAK.

Example 5.11. Virtual Procedure Using LOOP, CONTINUE, BREAK

```
CREATE VIRTUAL PROCEDURE  
BEGIN  
  DECLARE double total;  
  DECLARE integer transactions;  
  LOOP ON (SELECT amt, type FROM CashTxnTable) AS txncursor  
  BEGIN  
    IF(txncursor.type <> 'Sale')  
    BEGIN  
      CONTINUE;  
    END ELSE  
    BEGIN  
      total = (total + txncursor.amt);  
      transactions = (transactions + 1);  
      IF(transactions = 100)  
      BEGIN  
        BREAK;  
      END  
    END  
  END  
  SELECT total, (total / transactions) AS avg_transaction;  
END
```

This example is uses conditional logic to determine which of two SELECT statements to execute.

Example 5.12. Virtual Procedure with Conditional SELECT

```
CREATE VIRTUAL PROCEDURE  
BEGIN  
  DECLARE string VARIABLES.SORTDIRECTION;
```

```

VARIABLES.SORTDIRECTION = PartsVirtual.OrderedQtyProc.SORTMODE;
IF ( ucase(VARIABLES.SORTDIRECTION) = 'ASC' )
BEGIN
    SELECT * FROM PartsVirtual.SupplierInfo WHERE QUANTITY >
PartsVirtual.OrderedQtyProc.QTYIN ORDER BY PartsVirtual.SupplierInfo.PART_ID;
END ELSE
BEGIN
    SELECT * FROM PartsVirtual.SupplierInfo WHERE QUANTITY >
PartsVirtual.OrderedQtyProc.QTYIN ORDER BY PartsVirtual.SupplierInfo.PART_ID DESC;
END
END

```

5.2.4. Executing Virtual Procedures

You execute procedures using the SQL `EXECUTE` command. If the procedure has defined inputs, you specify those in a sequential list, or using "name=value" syntax. You must use the name of the input parameter, scoped by the full procedure name if the parameter name is ambiguous in the context of other elements or variables in the procedure.

A virtual procedure call will return a result set just like any `SELECT`, so you can use this in many places you can use a `SELECT`. However, within a virtual procedure itself you cannot always use an `EXEC` directly. Instead, you use the following syntax:

```
SELECT * FROM (EXEC ...) AS x
```

The following are some examples of how you can use the results of a virtual procedure call within a virtual procedure definition:

- `LOOP` instruction - you can walk through the results and do work on a row-by-row basis
- Assignment instruction - you can run a command and set the first column / first row value returned to a variable
- `SELECT * INTO #temp FROM (EXEC ...) AS x` - you can select the results from a virtual procedure into a temp table, which you can then query against as if it were a physical table.

5.3. Update Procedures

Virtual tables are abstractions above physical sources. They typically union or join information from multiple tables, often from multiple data sources or other views. Teiid can perform update operations against virtual tables. Update commands - `INSERT`, `UPDATE`, or `DELETE` - against a virtual table require logic to define how the tables and views integrated by the virtual table

are affected by each type of command. This transformation logic is invoked when an update command is issued against a virtual table. Update procedures define the logic for how a user's update command against a virtual table should be decomposed into the individual commands to be executed against the underlying physical sources. Similar to [virtual procedures](#), update procedures have the ability to execute queries or other commands, define temporary tables, add data to temporary tables, walk through result sets, use loops, and use conditional logic.

5.3.1. Update Procedure Definition

Usage:

```
CREATE PROCEDURE
BEGIN
  . . .
END
```

The CREATE VIRTUAL PROCEDURE line demarcates the beginning of the procedure. The BEGIN and END keywords are used to denote block boundaries. Within the body of the procedure, any valid [statement](#) may be used.

5.3.2. Special Variables

You can use a number of special variables when defining your update procedure.

5.3.2.1. INPUT Variables

Every attribute in the virtual table whose UPDATE and INSERT transformations you are defining has an equivalent variable named INPUT.<column_name>

When an INSERT or an UPDATE command is executed against the virtual table, these variables are initialized to the values in the INSERT VALUES clause or the UPDATE SET clause respectively.

In an UPDATE procedure, the default value of these variables, if they are not set by the command, is null. In an INSERT procedure, the default value of these variables is the default value of the virtual table attributes, based on their defined types. See [CHANGING Variables](#) for distinguishing defaults from passed values.

5.3.2.2. CHANGING Variables

Similar to INPUT Variables, every attribute in the virtual table whose UPDATE and INSERT transformations you are defining has an equivalent variable named CHANGING.<column_name>

When an INSERT or an UPDATE command is executed against the virtual table, these variables are initialized to `true` or `false` depending on whether the INPUT variable was set by the command.

For example, for a virtual table with columns A, B, C:

If User Executes...	Then...
INSERT INTO VT (A, B) VALUES (0, 1)	CHANGING.A = true, CHANGING.B = true, CHANGING.C = false
UPDATE VT SET C = 2	CHANGING.A = false, CHANGING.B = false, CHANGING.C = true

5.3.2.3. ROWS_UPDATED Variable

Teiid returns the value of the VARIABLES.ROWS_UPDATED variable as a response to an update command executed against the virtual table. Your procedure must set the value that returns when an application executes an update command against the virtual table, which triggers invocation of the update procedure. For example, if an UPDATE command is issued that affects 5 records, the ROWS_UPDATED should be set appropriately so that the user will receive '5' for the count of records affected.

5.3.3. Update Procedure Command Criteria

You can use a number of special SQL clauses when defining UPDATE or DELETE procedures. These make it easier to do variable substitutions in WHERE clauses or to check on the change state of variables without using a lot of conditional logic.

5.3.3.1. HAS CRITERIA

You can use the HAS CRITERIA clause to check whether the user's command has a particular kind of criteria on a particular set of attributes. This clause evaluates to either true or false. You can use it anywhere you can use a criteria within a procedure.

Usage:

```
HAS [criteria operator] CRITERIA [ON (element list)]
```

Syntax Rules

- The criteria operator, can be one of =, <, >, <=, >=, <>, LIKE, IS NULL, or IN.
- If the ON clause is present, HAS CRITERIA will return true only if criteria was present on all of the specified elements.
- The elements in a HAS CRITERIA ON clause always refer to virtual elements.

Some samples of the HAS CRITERIA clause:

SQL	Result
HAS CRITERIA	Checks simply whether there was any criteria at all.

SQL	Result
HAS CRITERIA ON (element1, element2)	Checks whether the criteria uses element1 and element2.
HAS = CRITERIA ON (element1)	Checks whether the criteria has a comparison criteria with = operator.
HAS LIKE CRITERIA	Checks whether the criteria has a match criteria using LIKE.

The HAS CRITERIA predicate is most commonly used in an IF clause, to determine if the user issued a particular form of command and to respond appropriately.

5.3.3.2. TRANSLATE CRITERIA

You can use the TRANSLATE CRITERIA clause to convert the criteria from the user application's SQL command into the form required to interact with the target source or view tables. The TRANSLATE CRITERIA statement uses the SELECT transformation to infer the element mapping. This clause evaluates to a translated criteria that is evaluated in the context of a command.

Usage:

```
TRANSLATE [criteria operator] CRITERIA [ON (element list)] [WITH (mapping list)]
```

Syntax Rules

- The criteria operator, can be one of =, <, >, <=, >=, <>, LIKE, IS NULL, or IN.
- If the ON clause is present, TRANSLATE CRITERIA will only form criteria using the specified elements.
- The elements in a TRANSLATE CRITERIA ON clause always refer to virtual elements.

You can use these mappings either to replace the default mappings generated from the SELECT transformation or to specify a reverse expression when a virtual element is defined by an expression.

Some samples of the HAS TRANSLATE clause:

SQL	Result
TRANSLATE CRITERIA	Translates any user criteria using the default mappings.
TRANSLATE CRITERIA WITH (element1 = 'A', element2 = INPUT.element2 + 2)	Translates any criteria with some additional mappings: element1 is always mapped to 'A' and element2 is mapped to the incoming element2 value + 2.

SQL	Result
TRANSLATE = CRITERIA ON (element1)	Translates only criteria that have = comparison operator and involve element1.

The TRANSLATE CRITERIA, ON clause always refers to virtual elements. The WITH clause always has items with form <elem> = <expression>, where the <elem> is a virtual element and the <expression> specifies what that virtual element should be replaced with when TRANSLATE CRITERIA translates the virtual criteria (from UPDATE or DELETE) into a physical criteria in the command. By default, a mapping is created based on the SELECT clause of the SELECT transformation (virtual column gets mapped to expression in SELECT clause at same position).

5.3.4. Update Procedure Processing

1. The user application submits the SQL command through one of SOAP, JDBC, or ODBC.
2. The virtual table that this SQL command is executed against is detected.
3. The correct procedure is chosen depending upon whether the command is an INSERT, UPDATE, or DELETE.
4. The procedure is executed. The procedure itself can contain SQL commands of its own which can be of different types than the command submitted by the user application that invoked the procedure.
5. Commands, as described in the procedure, as issued to the individual physical data sources or other views.
6. A value representing the number of rows changed is returned to the calling application.

Transaction Support

Teiid utilizes XA transactions for both participating in global transactions and for demarcating its own local and command scoped transactions. *JBoss Transactions* [<http://www.jboss.org/jbosstm/>] is used by Teiid as its internal transaction manager. See [this documentation](http://www.jboss.org/jbosstm/docs/index.html) [<http://www.jboss.org/jbosstm/docs/index.html>] for the advanced features provided by JBoss Transactions.

Table 6.1. Teiid Transaction Scopes

Scope	Description
Command	Treats the user command as if all source commands are executed within the scope of the same transaction. The <i>AutoWrap</i> execution property controls the behavior of command level transactions.
Local	The transaction boundary is local defined by a single client session.
Global	Teiid participates in a global transaction as an XA Resource.

6.1. AutoWrap Execution Property

Since user level commands may execute multiple source commands, users can specify the *AutoWrap* execution property to control the transactional behavior of a user command when not in a local or global transaction.

Table 6.2. AutoWrap Settings

Setting	Description
OFF	Do not wrap each command in a transaction. Individual source commands may commit or rollback regardless of the success or failure of the overall command.
ON	Wrap each command in a transaction. This mode is the safest, but may be burdensome on performance.
OPTIMISTIC	<i>This is the default setting.</i> Will not automatically wrap a command in a transaction, instead throw an exception if the command executed is <i>transactionally unsafe</i> . to execute outside of a transaction.
PESSIMISTIC	Will automatically wrap commands in a transaction, but only if the command seems to be <i>transactionally unsafe</i> .

The concept of command safety with respect to a transaction is determined by Teiid based upon command type and available metadata. Whenever any INSERT, UPDATE, DELETE, or

EXECUTE (with update count greater than 0) command is detected and the success or failure of that command is not the same as the user level command, then the command is deemed unsafe without a transaction.

The update count may be set on dynamic SQL as part of the command and on all other procedures as part of the procedure metadata in the model.

6.2. Updating Model Count

The term "updating model count" refers to the number of times any model is updated during the execution of a command. It is used to determine whether a transaction, of any scope, is required to safely execute the command.

Table 6.3. Updating Model Count Settings

Count	Description	Default For
0	No updates are performed by this command.	Dynamic SQL instructions
1	Indicates that only one model is updated by this command (and its subcommands). Also the success or failure of that update corresponds to the success or failure of the command. It should not be possible for the update to succeed while the command fails. Execution is not considered transactionally unsafe.	Physical procedures, inserts, updates, and deletes
*	Any number greater than 1 indicates that execution is transactionally unsafe and an XA transaction will be required.	XQuery commands

By default Teiid will calculate the updating model count for a user query (which may be composed of many subcommands) given the above table. Thus a command with a single update as a subcommand may still require a transaction if the update is not tied to the success of the command.

Explicit values for updating model count may be set using Teiid Designer and directly in the [dynamic SQL command](#).

6.3. JDBC and Transactions

6.3.1. JDBC API Functionality

The transaction scopes above map to these JDBC modes:

- Command - Connection autoCommit property set to true.
- Local - Connection autoCommit property set to false. The transaction is committed by setting autoCommit to true or calling `java.sql.Connection.commit`. The transaction can be rolled back by a call to `java.sql.Connection.rollback`

- Global - the XAResource interface provided by an XAConnection is used to control the transaction. Note that XAConnections are available only if Teiid is consumed through its XADatasource, `com.metamatrix.jdbc.MMDatasource` . JEE containers or data access APIs typically control XA transactions on behalf of application code.

6.3.2. J2EE Usage Models

J2EE provides three ways to manage transactions for beans:

- Client-controlled – the client of a bean begins and ends a transaction explicitly.
- Bean-managed – the bean itself begins and ends a transaction explicitly.
- Container-managed – the app server container begins and ends a transaction automatically.

In any of these cases, transactions may be either local or XA transactions, depending on how the code and descriptors are written. Some kinds of beans (stateful session beans and entity beans) are not required by the spec to support non-transactional sources, although the spec does allow an app server to optionally support this with the caution that this is not portable or predictable. Generally speaking, to support most typical EJB activities in a portable fashion requires some kind of transaction support.

6.4. Limitations and Workarounds

- The client setting of transaction isolation level is not used. The transaction isolation level can be set on each XA connector, however this isolation level is fixed and cannot be changed at runtime for specific connections/commands.
- Since the client transaction isolation level is not used, Teiid internally assumes a level of `READ_COMMITTED`. This implies that explicit transactions are not required for user level commands performing multiple reads.
- Temporary tables are not transactional. For example, a global temporary table will retain all inserts performed during a local transaction that was rolled back.
- Connectors may be set to immutable to prevent their participation in transactions. This is useful in situations where update commands are being issued against a source that lacks XA transaction capabilities.

System Tables

7.1. VDB and Model Metadata

7.1.1. System.VirtualDatabases

This table supplies information about the currently connected virtual database, of which there is always exactly one (in the context of a connection).

Column Name	Type	Description
Name	string	The name of the VDB
Version	string	The version of the VDB

7.1.2. System.Models

This table supplies information about all the models in the virtual database, including the system model itself (System).

Column Name	Type	Description
Name	string	Model name
Version	string	Model version
IsPhysical	boolean	True if source model, false for view
SupportsWhereAll	boolean	Model supports queries with no criteria
SupportsOrderBy	boolean	Model supports ORDER BY queries
SupportsJoin	boolean	Model supports queries with joins
SupportsDistinct	boolean	Model supports SELECT DISTINCT queries
SupportsOuterJoin	boolean	Model supports queries with outer joins
MaxSetSize	integer	Max number of values to pass in an IN value set for a dependent join
UID	string	Unique ID
Description	string	Description
PrimaryMetamodelURI	string	URI for the primary metamodel describing this model

7.1.3. System.ModelProperties

This table supplies user-defined properties on models based on metamodel extensions. Normally, this table is empty if no metamodel extensions are being used.

Column Name	Type	Description
ModelName	string	Model name
Name	string	Property name
Value	string	Property value
UID	string	Model unique ID

7.2. Table Metadata

7.2.1. System.Groups

This table supplies information about all the groups (tables, views, documents, etc) in the virtual database.

Column Name	Type	Description
ModelName	string	Model name
FullName	string	Full group name
Name	string	Short group name
Type	string	Table type (Table, View, Document, ...)
NameInSource	string	Name of this group in the source
IsPhysical	boolean	True if this is a source model
UpperName	string	Upper-case full group name for easier matching
SupportsUpdates	boolean	True if group can be updated
UID	string	Group unique ID
Cardinality	integer	Approximate number of rows in the group
Description	string	Description
IsSystem	boolean	True if in system model

7.2.2. System.GroupProperties

This table supplies user-defined properties on groups based on metamodel extensions. Normally, this table is empty if no metamodel extensions are being used.

Column Name	Type	Description
ModelName	string	Model name
GroupFullName	string	Full group name
Name	string	Property name
Value	string	Property value
GroupName	string	Short group name
GroupUpperName	string	Full upper-case group name
ID	string	Group unique ID

7.2.3. System.Elements

This table supplies information about all the elements (columns, tags, attributes, etc) in the virtual database.

Column Name	Type	Description
ModelName	string	Model name
GroupName	string	Short group name
GroupFullName	string	Full group name
Name	string	Element name (not qualified)
Position	integer	Position in group (1-based)
NameInSource	string	Name of element in source
Data Type	string	Teiid runtime data type name
Scale	integer	Number of digits after the decimal point
ElementLength	integer	Element length (mostly used for strings)
sLengthFixed	boolean	Whether the length is fixed or variable
SupportsSelect	boolean	Element can be used in SELECT
SupportsUpdates	boolean	Values can be inserted or updated in the element
IsCaseSensitive	boolean	Element is case-sensitive
IsSigned	boolean	Element is signed numeric value
IsCurrency	boolean	Element represents monetary value
IsAutoIncremented	boolean	Element is auto-incremented in the source

Column Name	Type	Description
NullType	string	Nullability: "Nullable", "No Nulls", "Unknown"
MinRange	string	Minimum numeric value
MaxRange	string	Maximum numeric value
SearchType	string	Searchability: "Searchable", "All Except Like", "Like Only", "Unsearchable"
Format	string	Format of string value
DefaultValue	string	Default value
JavaClass	string	Java class that will be returned
Precision	integer	Number of digits in numeric value
CharOctetLength	integer	Measure of return value size
Radix	integer	Radix for numeric values
GroupUpperName	string	Upper-case full group name
UpperName	string	Upper-case element name
UID	string	Element unique ID
Description	string	Description

7.2.4. System.ElementProperties

This table supplies user-defined properties on groups based on metamodel extensions. Normally, this table is empty if no metamodel extensions are being used.

Column Name	Type	Description
ModelName	string	Model name
GroupFullName	string	Full group name
ElementName	string	Element name
Name	string	Property name
Value	string	Property value
GroupName	string	Short group name
ElementUpperName	string	Upper-case element name
GroupUpperName	string	Upper-case group name
UID	string	Element unique ID

7.2.5. System.Keys

This table supplies information about primary, foreign, and unique keys.

Column Name	Type	Description
ModelName	string	Model name
GroupFullName	string	Full group name
Name	string	Key name
Description	string	Description
NameInSource	string	Name of key in source system
Type	string	Type of key: "Primary", "Foreign", "Unique", etc
IsIndexed	boolean	True if key is indexed
GroupName	string	Short group name
GroupUpperName	string	Upper-case full group name
RefKeyUID	string	Referenced key UID (if foreign key)
UID	string	Key unique ID

7.2.6. System.KeyProperties

This table supplies user-defined properties on keys based on metamodel extensions. Normally, this table is empty if no metamodel extensions are being used.

Column Name	Type	Description
Column Name	Type	Description
ModelName	string	Model name
GroupFullName	string	Full group name
KeyName	string	Key name
Name	string	Extension property name
Value	string	Extension property value
GroupName	string	Short group name
GroupUpperName	string	Upper-case full group name
UID	string	Key unique ID

7.2.7. System.KeyElements

This table supplies information about the elements referenced by a key.

Column Name	Type	Description
ModelName	string	Model name
GroupFullName	string	Full group name

Column Name	Type	Description
Name	string	Element name
KeyName	string	Key name
KeyType	string	Key type: "Primary", "Foreign", "Unique", etc
GroupName	string	Short group name
GroupUpperName	string	Upper case full group name
RefKeyUID	string	Referenced key UID
UID	string	Key UID
Position	integer	Position in key

7.3. Procedure Metadata

7.3.1. System.Procedures

This table supplies information about the procedures in the virtual database.

Column Name	Type	Description
ModelName	string	Model name
Name	string	Procedure name
NameInSource	string	Procedure name in source system
ReturnsResults	boolean	Returns a result set
ModelUID	string	Model UID
UID	string	Procedure UID
Description	string	Description
FullName	string	Full procedure name

7.3.2. System.ProcedureProperties

This table supplies user-defined properties on procedures based on metamodel extensions. Normally, this table is empty if no metamodel extensions are being used.

Column Name	Type	Description
ModelName	string	Model name
ProcedureName	string	Procedure name
Name	string	Property name
Value	string	Property value
UID	string	Procedure UID

7.3.3. System.ProcedureParams

This supplies information on procedure parameters.

Column Name	Type	Description
ModelName	string	Model name
ProcedureName	string	Procedure name
Name	string	Parameter name
DataType	string	Teiid runtime data type name
Position	integer	Position in procedure args
Type	string	Parameter direction: "In", "Out", "InOut", "ResultSet", "ReturnValue"
Optional	boolean	Parameter is optional
Precision	integer	Precision of parameter
TypeLength	integer	Length of parameter value
Scale	integer	Scale of parameter
Radix	integer	Radix of parameter
NullType	string	Nullability: "Nullable", "No Nulls", "Unknown"

7.4. Datatype Metadata

7.4.1. System.DataTypes

This table supplies information on *datatypes*.

Column Name	Type	Description
Name	string	Teiid design-time type name
IsStandard	boolean	Always false
IsPhysical	boolean	Always false
TypeName	string	Design-time type name (same as Name)
JavaClass	string	Java class returned for this type
Scale	integer	Max scale of this type
TypeLength	integer	Max length of this type
NullType	string	Nullability: "Nullable", "No Nulls", "Unknown"

Column Name	Type	Description
IsSigned	boolean	Is signed numeric?
IsAutoIncremented	boolean	Is auto-incremented?
IsCaseSensitive	boolean	Is case-sensitive?
Precision	integer	Max precision of this type
Radix	integer	Radix of this type
SearchType	string	Searchability: "Searchable", "All Except Like", "Like Only", "Unsearchable"
UID	string	Data type unique ID
RuntimeType	string	Teiid runtime data type name
BaseType	string	Base type
Description	string	Description of type

7.4.2. System.DataTypeProperties

This table supplies user-defined properties on data types based on metamodel extensions. Normally, this table is empty if no metamodel extensions are being used.

Column Name	Type	Description
DataType	string	Data type name
Name	string	Property name
Value	string	Property value
UID	string	Data type UID

Connectors

8.1. Common Connector Properties

These properties are common to all connectors.

Property	Description	Type	Default
ConnectorBindingName	The name of the connector binding. Must be unique across all connector bindings. This property is not editable, and is not seen in the Designer in the connector properties table.	string	
ConnectorClass	The class name of the custom connector class that connects to the data source. Required. Not editable.	string	
ConnectorClassPath	The class path of the custom connector class that connects to the data source. Required.	string	
ConnectorMaxThreads	The maximum number of connector worker threads. Required.	integer	
ConnectorThreadTTL	The maximum length of time a connector thread may live in idle state. Required.	integer - milliseconds	
ExceptionOnMaxRows	Specifies whether or not an exception is thrown when the number of rows for a query exceeds the value of MaxResultRows. If this flag is set to false, then no more rows than MaxResultRows will be returned, but no exception will be thrown.	boolean	
MaxResultRows	The maximum number of rows to be processed by the connector from a source. A custom connector should stop adding rows to the ResultsCollector when the	integer	

Property	Description	Type	Default
	number of rows reaches this value. Optional – if not specified, all rows will be processed.		

8.2. Source Security

Teiid can use the security mechanisms of individual data sources during execution.

The use of source security is driven off the credentials JDBC property.

There are 3 basic use cases:

1. Specify the credentials for all sources
2. Pass the session logon credentials to all sources
3. Some combination of the previous two

8.2.1. Specific Source Credentials

This use case is for when you want to specify the logon credentials for all the source systems in the JDBC URL. The form of the credentials property is:

```
credentials=(system=<sys1>,user=<uname>, password=<pwd>/system=<sys2>,...)
```



Note

Any property can be specified for a system, using the name=value syntax, as long as the connector associated with the source understands the property or can pass it to the source.

The following properties are "well-known" and required:

- system - The name of the system; this must correspond to the connector binding name
- user - The username that can be used to connect to the system; must be valid user of the system being accessed
- password - The password for the user in the target system; the user and password are used to authenticate to the target system.

Example 8.1. Example URL With Two Source Credentials

```
jdbc:metamatrix:MyVDB@mm://  
HostA:5001;credentials=(system=s1,user=u1,password=p1/  
system=s2,user=u2,password=p2)
```


8.2.2. Using Session Credentials

This use case uses the session credentials for all source systems.

Example 8.2. Example URL With Session Credentials

```
jdbc:metamatrix:MyVDB@mm://HostA:5001;credentials=defaultToLogon
```

8.2.3. Session and Specific Source Credentials

The above techniques can be used in combination with one another as needed.

Example 8.3. Combination of Source Credential Approaches

```
jdbc:metamatrix:MyVDB@mm://HostA:5001;  
credentials=defaultToLogon,(system=s1,user=u1,password=p1)
```

8.2.4. Configuring Connectors For Source Security

When source security is used, the connector bindings for the sources must be configured appropriately.

- Connector binding name - this is the name used in the system part of the credentials property
- For JDBC connectors, the extension connection factory class name must be set for user instead of single identity. For example the Oracle JDBC connector would use `com.metamatrix.connector.jdbc.oracle.OracleUserIdentityConnectionFactory`

Federated Planning

Teiid at its core is a federated relational query engine. This query engine allows you to treat all of your data sources as one virtual database and access them in a single SQL query. This allows you to focus on building your application, not on hand-coding joins, and other relational operations, between data sources.

9.1. Overview

When the query engine receives an incoming SQL query it performs the following operations:

1. Parsing - validate syntax and convert to internal form
2. Resolving - link all identifiers to metadata and functions to the function library
3. Validating - validate SQL semantics based on metadata references and type signatures
4. Rewriting - rewrite SQL to simplify expressions and criteria
5. Logical plan optimization - the rewritten canonical SQL is converted into a logical plan for in-depth optimization. The Teiid optimizer is predominantly rule-based. Based upon the query structure and hints a certain rule set will be applied. These rules may trigger in turn trigger the execution of more rules. Within several rules, Teiid also takes advantage of costing information. The logical plan optimization steps can be seen by using the `OPTION DEBUG` clause and are described in the [query planner](#) section.
6. Processing plan conversion - the logic plan is converted into an executable form where the nodes are representative of basic processing operations. The final processing plan is displayed as the [query plan](#).

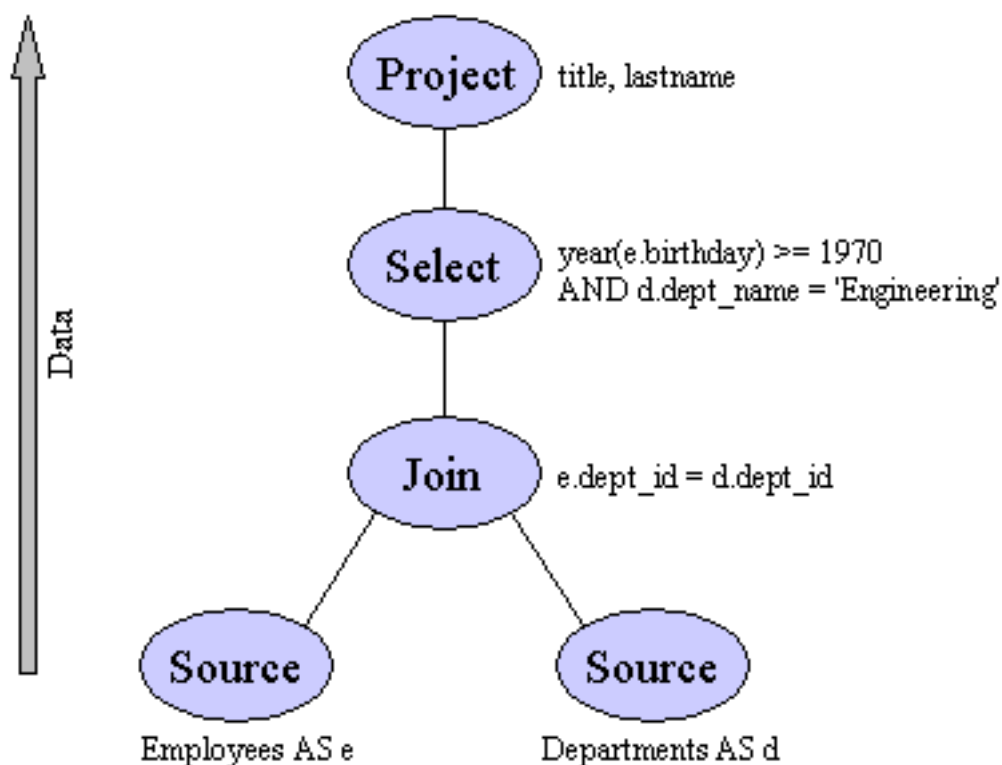
The logical query plan is a tree of operations used to transform data in source tables to the expected result set. In the tree, data flows from the bottom (tables) to the top (output). The primary logical operations are *select* (select or filter rows based on a criteria), *project* (project or compute column values), *join*, *source* (retrieve data from a table), *sort* (ORDER BY), *duplicate removal* (SELECT DISTINCT), *group* (GROUP BY), and *union* (UNION).

For example, consider the following query that retrieves all engineering employees born since 1970.

Example 9.1. Example query

```
SELECT e.title, e.lastname FROM Employees AS e JOIN  
Departments AS d ON e.dept_id = d.dept_id WHERE year(e.birthdate) >= 1970 AND d.dept_name  
= 'Engineering'
```

Logically, the data from the Employees and Departments tables are retrieved, then joined, then filtered as specified, and finally the output columns are projected. The canonical query plan thus looks like this:



Data flows from the tables at the bottom upwards through the join, through the select, and finally through the project to produce the final results. The data passed between each node is logically a result set with columns and rows.

Of course, this is what happens *logically*, not how the plan is actually executed. Starting from this initial plan, the query planner performs transformations on the query plan tree to produce an equivalent plan that retrieves the same results faster. Both a federated query planner and a relational database planner deal with the same concepts and many of the same plan transformations. In this example, the criteria on the Departments and Employees tables will be pushed down the tree to filter the results as early as possible.

In both cases, the goal is to retrieve the query results in the fastest possible time. However, the relational database planner does this primarily by optimizing the access paths in pulling data from storage.

In contrast, a federated query planner is less concerned about storage access because it is typically pushing that burden to the data source. The most important consideration for a federated query planner is minimizing data transfer.

9.2. Federated Optimizations

9.2.1. Access Patterns

Access patterns are used on both physical and virtual sources to specify the need for criteria against a set of elements. Failure to supply the criteria will result in a planning error, rather than a run-away source query. Access patterns can be applied in a set such that only one of the access patterns is required to be satisfied.

Currently any form of criteria may satisfy an access pattern as long as it contains references to affect elements.

9.2.2. Pushdown

In federated database systems pushdown refers to decomposing the user level query into source queries that perform as much work as possible on their respective source system. Pushdown analysis requires knowledge of source system capabilities, which is provided to Teiid through the Connector API. Any work not performed at the source is then processed in Federate's relational engine.

Based upon capabilities, Teiid will manipulate the query plan to ensure that each source performs as much joining, filtering, grouping, etc. as possible. In many cases, such as with join ordering, planning is a combination of *standard relational techniques* and, cost based and heuristics for pushdown optimization.

Criteria and join push down are typically the most important aspects of the query to push down when performance is a concern. See *Query Plans* on how to read a plan to ensure that source queries are as efficient as possible.

9.2.3. Dependent Joins

A special optimization called a dependent join is used to reduce the rows returned from one of the two relations involved in a multi-source join. In a dependent join, queries are issued to each source sequentially rather than in parallel, with the results obtained from the first source used to restrict the records returned from the second. Dependent joins can perform some joins much faster by drastically reducing the amount of data retrieved from the second source and the number of join comparisons that must be performed.

The conditions when a dependent join is used are determined by the query planner based on *access patterns*, hints, and costing information.

Teiid supports the MAKEDEP and MAKENOTDEP hints. These can be placed in either the *OPTION clause* or directly in the *FROM clause*. As long as all can be met, the MAKEDEP and MAKENOTDEP hints override any use of costing information.



Tip

The MAKEDEP hint should only be used if the proper query plan is not chosen by default. You should ensure that your costing information is representative of the actual source cardinality. An inappropriate MAKEDEP hint can force an inefficient join structure and may result in many source queries.

9.2.4. Copy Criteria

Copy criteria is an optimization that creates additional predicates based upon combining join and where clause criteria. For example, equi-join predicates (`source1.table.column = source2.table.column`) are used to create new predicates by substituting `source1.table.column` for `source2.table.column` and vice versa. In a cross source scenario, this allows for where criteria applied to a single side of the join to be applied to both source queries

9.2.5. Projection Minimization

Teiid ensures that each pushdown query only projects the symbols required for processing the user query. This is especially helpful when querying through large intermediate view layers.

9.2.6. Partial Aggregate Pushdown

Partial aggregate pushdown allows for grouping operations above multi-source joins and unions to be decomposed so that some of the grouping and aggregate functions may be pushed down to the sources.

9.2.7. Optional Join

The optional join hint indicates to the optimizer that a join clause should be omitted if none of its columns are used in either user criteria or output columns in the result. This hint is typically only used in view layers containing multi-source joins.

The optional join hint is applied as a comment on a join clause.

Example 9.2. Example Optional Join Hint

```
select a.column1, b.column2 from a inner join /* optional */ b on a.key = b.key
```

Suppose that the preceding example defined a view layer X. If X is queried in such a way as to not need `b.column2`, then the optional join hint will cause `b` to be omitted from the query plan. The result would be the same as if X were defined as:

```
select a.column1 from a
```



Tip

When a join clause is omitted via the optional join hint, the relevant join criteria is not applied. Thus it is possible that the query results may not have the same cardinality or even the same row values as when the join is fully applied.

Left/right outer joins where the inner side values are not used and whose rows under go a distinct operation will automatically be treated as an optional join and does not require a hint.

9.2.8. Standard Relational Techniques

Teiid also incorporates many standard relational techniques to ensure efficient query plans.

- Rewrite analysis for function simplification and evaluation.
- Boolean optimizations for basic criteria simplification.
- Removal of unnecessary view layers.
- Removal of unnecessary sort operations.
- Advanced search techniques through the left-linear space of join trees.
- Parallelizing of source access during execution.

9.3. Federated Failure Modes

9.3.1. Partial Results

Teiid provides the capability to obtain "partial results" in the event of data source unavailability. This is especially useful when unioning information from multiple sources, or when doing a left outer join, where you are 'appending' columns to a master record but still want the record if the extra info is not available.

If one or more data sources are unavailable to return results, then the result set obtained from the remaining available sources will be returned. In the case of joins, an unavailable data source essentially contributes zero tuples to the result set.

9.3.1.1. Setting Partial Results Mode

Partial results mode is off by default but can be turned on by default for all queries in a Connection with either `setPartialResultsMode("true")` on a DataSource or `partialResultsMode=true` on a JDBC URL. In either case, partial results mode may be overridden on a per-query basis by setting the execution property on the Statement. To set this property, cast to the Teiid Statement JDBC API extension interface `com.metamatrix.jdbc.api.Statement`

Example 9.3. Example - Setting Partial Results Mode

```
Statement statement = ...obtain statement from Connection...

com.metamatrix.jdbc.api.Statement mmStatement =
    (com.metamatrix.jdbc.api.Statement) statement;

mmStatement.setExecutionProperty(
    ExecutionProperties.PROP_PARTIAL_RESULTS_MODE, "true");
```

This property can be set before each execution (via an execute method) on a Statement (or PreparedStatement or CallableStatement).

9.3.1.2. Source Unavailability

A source is considered to be 'unavailable' if the connector binding associated with the source issues an exception in response to a query. The exception will be propagated to the query processor, where it will become a warning in the result set.



Warning

Since Teiid supports multi-source cursoring, it is possible that the unavailability of a data source will not be determined until after the first batch of results have been returned to the client. This can happen in the case of unions, but not joins. In this situation, there will be no warnings in the result set when the client is processing the first batch of results. The client will be responsible for periodically checking the status of warnings in the results object as results are being processed, to see if a new warning has been added due to the detection of an unavailable source. [Note that client applications have no notion of 'batches', which are purely a server-side entity. Client apps deal only with records.]

For each source that is excluded from a query, a warning will be generated describing the source and the failure. These warnings can be obtained from the Statement.getWarnings() method. This method returns a SQLWarning object but in the case of "partial results" warnings, this will be an object of type com.metamatrix.jdbc.api.PartialResultsWarning. This class can be used to obtain a list of all the failed connectors by name and to obtain the specific exception thrown by each connector.

Example 9.4. Example - Printing List of Failed Sources

```
statement.setExecutionProperty( PROP_PARTIAL_RESULTS_MODE, "true");
ResultSet results = statement.executeQuery("SELECT Name FROM Accounts");
```



```

SQLWarning warning = statement.getWarnings();
if(warning instanceof PartialResultsWarning) {
    PartialResultsWarning partialWarning = (PartialResultsWarning)warning;
    Collection failedConnectors = partialWarning.getFailedConnectors();
    Iterator iter = failedConnectors.iterator();
    while(iter.hasNext()) {
        String connectorName = (String) iter.next();
        SQLException connectorException = partialWarning.getConnectorException(connectorName);
        System.out.println(connectorName + ": " + ConnectorException.getMessage());
    }
}

```

9.4. Query Plans

When integrating information using a federated query planner, it is useful to be able to view the query plans that are created, to better understand how information is being accessed and processed, and to troubleshoot problems.

A query plan is a set of instructions created by a query engine for executing a command submitted by a user or application. The purpose of the query plan is to execute the user's query in as efficient a way as possible.

9.4.1. Getting a Query Plan

You can get a query plan any time you execute a command. The SQL options available are as follows:

- `OPTION SHOWPLAN` - Returns the plan in addition to any results
- `OPTION PLANONLY` - Returns the plan, does not execute the command though

With the above options, the query plan is available from the Statement object by casting to the `com.metamatrix.jdbc.api.Statement` interface.

Example 9.5. Retrieving a Query Plan

```

ResultSet rs = statement.executeQuery("select ...");
com.metamatrix.jdbc.api.Statement mmstatement =
    (com.metamatrix.jdbc.api.Statement)statement;
PlanNode queryPlan = mmstatement.getPlanDescription();
System.out.println(XMLOutputVisitor.convertToXML(queryPlan);

```

The query plan is made available automatically in several of Teiid's tools.

9.4.2. Analyzing a Query Plan

Once a query plan has been obtained you will most commonly be looking for:

- Source pushdown -- what parts of the query that got pushed to each source
- Join ordering
- Join algorithm used - merge or nested loop.
- Presence of federated optimizations, such as dependent joins.
- Join criteria type mismatches.

All of these issues presented above will be present subsections of the plan that are specific to relational queries. If you are executing a procedure or generating an XML document, the overall query plan will contain additional information related the surrounding procedural execution.

A query plan consists of a set of nodes organized in a tree structure. As with the above example, you will typically be interested in analyzing the textual form of the plan.

In a procedural context the ordering of child nodes implies the order of execution. In most other situation, child nodes may be executed in any order even in parallel. Only in specific optimizations, such as dependent join, will the children of a join execute serially.

9.4.3. Relational Plans

Relational plans represent the actually processing plan that is composed of nodes that are the basic building blocks of logical relational operations. Physical relational plans differ from logical relational plans in that they will contain additional operations and execution specifics that were chosen by the optimizer.

The nodes for a relational query plan are:

- Access - Access a source. A source query is sent to the connector binding associated with the source. [For a dependent join, this node is called Dependent Select.]
- Project - Defines the columns returned from the node. This does not alter the number of records returned. [When there is a subquery in the Select clause, this node is called Dependent Project.]
- Project Into - Like a normal project, but outputs rows into a target table.
- Select - Select is a criteria evaluation filter node (WHERE / HAVING). [When there is a subquery in the criteria, this node is called Dependent Select.]

- Join - Defines the join type, join criteria, and join strategy (merge or nested loop).
- Union - There are no properties for this node, it just passes rows through from its children
- Sort - Defines the columns to sort on, the sort direction for each column, and whether to remove duplicates or not.
- Dup Removal - Same properties as for Sort, but the removeDups property is set to true
- Group - Groups sets of rows into groups and evaluates aggregate functions.
- Null - A node that produces no rows. Usually replaces a Select node where the criteria is always false (and whatever tree is underneath). There are no properties for this node.
- Plan Execution - Executes another sub plan.
- Limit - Returns a specified number of rows, then stops processing. Also processes an offset if present.
- Dependent Feeder - This node accepts its input stream and forwards to its parent unchanged but also feeds all dependent sources that need the stream of data. Thus, this node actually performs no work within the tree, just diverts a copy of the tuple stream to listening nodes.
- Dependent Wait - This node waits until a criteria requiring dependent values below this node has the necessary data to continue. At that point, it continues processing on its subplan and merely forwards data from the child to the parent.

9.4.3.1. Node Statistics

Every node has a set of statistics that are output. These can be used to determine the amount of data flowing through the node.

Statistic	Description	Units
Node Output Rows	Number of records output from the node	count
Node Process Time	Time processing in this node only	millisec
Node Cumulative Process Time	Elapsed time from beginning of processing to end	millisec
Node Cumulative Next Batch Process Time	Time processing in this node + child nodes	millisec

Statistic	Description	Units
Node Next Batch Calls	Number of times a node was called for processing	count
Node Blocks	Number of times a blocked exception was thrown by this node or a child	count

In addition to node statistics, some nodes display cost estimates computed at the node.

Cost Estimates	Description	Units
Estimated Node Cardinality	Estimated number of records that will be output from the node; -1 if unknown	count

9.5. Query Planner

For each sub-command in the user command an appropriate kind of sub-planner is used (relational, XML, XQuery, procedure, etc).

Each planner has three primary phases:

1. Generate canonical plan
2. Optimization
3. Plan to process converter - converts plan data structure into a processing form

9.5.1. Relational Planner

The `GenerateCanonical` class generates the initial (or “canonical” plan). This plan is based on the typical logical order that a SQL query gets executed. A SQL select query has the following possible clauses (all but `SELECT` are optional): `SELECT`, `FROM`, `WHERE`, `GROUP BY`, `HAVING`, `ORDER BY`, `LIMIT`. These clauses are logically executed in the following order:

1. `FROM` (read and join all data from tables)
2. `WHERE` (filter rows)
3. `GROUP BY` (group rows into collapsed rows)
4. `HAVING` (filter grouped rows)
5. `SELECT` (evaluate expressions and return only requested columns)
6. `INTO`

7. ORDER BY (sort rows)

8. LIMIT (limit result set to a certain range of results)

These clause translate into the following types of planning nodes:

- FROM: Source node for each from clause item, Join node (if >1 table)
- WHERE: Select node
- GROUP BY: Group node
- GROUP BY: Group node
- SELECT: Project node and DupRemoval node (for SELECT DISTINCT)
- INTO: Project node with a SOURCE Node
- INTO: Project node with a SOURCE Node
- LIMIT: Limit node
- UNION, EXCEPT, INTERSECT: SetOp Node

There is also a Null Node that can be created as the result of rewrite or planning optimizations. It represents a node that produces no rows

Relational optimization is based upon rule execution that evolves the initial plan into the execution plan. There are a set of pre-defined rules that are dynamically assembled into a rule stack for every query. The rule stack is assembled based on the contents of the user's query and its transformations. For example, if there are no virtual layers, then RuleMergeVirtual, which merges virtual layers together, is not needed and will not be added to the stack. This allows the rule stack to reflect the complexity of the query.

Logically the plan node data structure represents a tree of nodes where the source data comes up from the leaf nodes (typically Access nodes in the final plan), flows up through the tree and produces the user's results out the top. The nodes in the plan structure can have bidirectional links, dynamic properties, and allow any number of child nodes. Processing *plan nodes* in contrast typical have fixed properties, and only allow for binary operations - due to algorithmic limitations.

Below are some of the rules included in the planner:

- RuleRemoveSorts - removes sort nodes that do not have an effect on the result. This most common when a view has an non-limited ORDER BY.
- RulePlaceAccess - insert an Access node above every physical Source node. The source node represents a table typically. An access node represents the point at which everything below the access node gets pushed to the source. Later rules focus on either pushing stuff under the

access or pulling the access node up the tree to move more work down to the data sources. This rule is also responsible for placing .

- RulePushSelectCriteria - pushes select criteria down through unions, joins, and views into the source below the access node. In most cases movement down the tree is good as this will filter rows earlier in the plan. We currently do not undo the decisions made by PushSelectCriteria. However in situations where criteria cannot be evaluated by the source, this can lead to sub optimal plans.

One of the most important optimization related to pushing criteria, is how the criteria will be pushed trough join. Consider the following plan tree that represents a subtree of the plan for the query "select ... from A inner join b on (A.x = B.x) where A.y = 3"

```
SELECT (B.y = 3)
 |
JOIN - Inner Join on (A.x = B.x)
 /  \
SRC (A) SRC (B)
```

Note: SELECT nodes represent criteria, and SRC stands for SOURCE.

It is always valid for inner join and cross joins to push (single source) criteria that are above the join, below the join. This allows for criteria originating in the user query to eventually be present in source queries below the joins. This result can be represented visually as:

```
JOIN - Inner Join on (A.x = B.x)
 /  \
 /  SELECT (B.y = 3)
 |    |
SRC (A) SRC (B)
```

The same optimization is valid for criteria specified against the outer side of an outer join. For example:

```
SELECT (B.y = 3)
 |
JOIN - Right Outer Join on (A.x = B.x)
 /  \
SRC (A) SRC (B)
```

Becomes

```

      JOIN - Right Outer Join on (A.x = B.x)
    /  \
  /      \
 / SELECT (B.y = 3)
|        |
|        |
SRC (A) SRC (B)

```

However criteria specified against the inner side of an outer join needs special consideration. The above scenario with a left or full outer join is not the same. For example:

```

      SELECT (B.y = 3)
      |
      JOIN - Left Outer Join on (A.x = B.x)
    /  \
  /      \
 / SRC (A) SRC (B)

```

Can become (available only after 5.0.2):

```

      JOIN - Inner Join on (A.x = B.x)
    /  \
  /      \
 / SELECT (B.y = 3)
|        |
|        |
SRC (A) SRC (B)

```

Since the criterion is not dependent upon the null values that may be populated from the inner side of the join, the criterion is eligible to be pushed below the join – but only if the join type is also changed to an inner join.

On the other hand, criteria that are dependent upon the presence of null values CANNOT be moved. For example:

```

      SELECT (B.y is null)
      |
      JOIN - Left Outer Join on (A.x = B.x)
    /  \
  /      \
 / SRC (A) SRC (B)

```

This plan tree must have the criteria remain above the join, since the outer join may be introducing null values itself. This will be true regardless of which version of Teiid is used.

- RulePushNonJoinCriteria – this rule will push criteria out of an on clause if it is not necessary for the correctness of the join.
- RuleRaiseNull – this rule will raise null nodes to their highest possible point. Raising a null node removes the need to consider any part of the old plan that was below the null node.
- RuleMergeVirtual - merges virtual layers together. Virtual layers are connected by nesting canonical plans under source leaf nodes of the parent plan. Each canonical plan is also sometimes referred to as a “query frame”. Merge virtual attempts to merge child frames into the parent frame. The merge involves renaming any symbols in the lower frame that overlap with symbols in the upper frame. It also involves merging the join information together.
- RuleRemoveOptionalJoins – removes optional join nodes from the plan tree as soon as possible so that planning will be more optimal.
- RulePlanJoins – this rule attempts to find an optimal ordering of the joins performed in the plan, while ensuring that dependencies are met. This rule has three main steps. First it must determine an ordering of joins that satisfy the present. Second it will heuristically create joins that can be pushed to the source (if a set of joins are pushed to the source, we will not attempt to create an optimal ordering within that set. More than likely it will be sent to the source in the non-ANSI multi-join syntax and will be optimized by the database). Third it will use costing information to determine the best left-linear ordering of joins performed in the processing engine. This third step will do an exhaustive search for 6 or less join sources and is heuristically driven by join selectivity for 7 or more sources.
- RuleCopyCriteria - this rule copies criteria over an equality criteria that is present in the criteria of a join. Since the equality defines an equivalence, this is a valid way to create a new criteria that may limit results on the other side of the join (especially in the case of a multi-source join).
- RuleCleanCriteria - this rule cleans up criteria after all the other rules.
- RuleMergeCriteria - looks for adjacent criteria nodes and merges them together. It looks for adjacent identical conjuncts and removes duplicates.
- RuleRaiseAccess - this rule attempts to raise the Access nodes as far up the plan as possible. This is mostly done by looking at the source’s capabilities and determining whether the operations can be achieved in the source or not.
- RuleChooseDependent - this rule looks at each join node and determines whether the join should be made dependent and in which direction. Cardinality, the number of distinct values, and primary key information are used in several formulas to determine whether a dependent join is likely to be worthwhile. The dependent join differs in performance ideally because a fewer number of values will be returned from the dependent side. Also, we must consider the number of values passed from independent to dependent side. If that set is larger than the max number of values in an IN criteria on the dependent side, then we must break the query into a set of queries and combine their results. Executing each query in the connector has some overhead and that is taken into account. Without costing information a lot of common cases

where the only criteria specified is on a non-unique (but strongly limiting) field are missed. A join is eligible to be dependent if:

1. there is at least one equi-join criterion, i.e. $tablea.col = tableb.col$
2. the join is not a full outer join and the dependent side of the join is on the inner side of the join

The join will be made dependent if one of the following conditions, listed in precedence order, holds:

1. There is an unsatisfied access pattern that can be satisfied with the dependent join criteria
2. The potential dependent side of the join is marked with an option `makedep`
3. (4.3.2) if costing was enabled, the estimated cost for the dependent join (5.0+ possibly in each direction in the case of inner joins) is computed and compared to not performing the dependent join. If the costs were all determined (which requires all relevant table cardinality, column ndv, and possibly nrv values to be populated) the lowest is chosen.
4. If key metadata information indicates that the potential dependent side is not “small” and the other side is “not small” or (5.0.1) the potential dependent side is the inner side of a left outer join.

Dependent join is the key optimization we use to efficiently process multi-source joins.

Instead of reading all of source A and all of source B and joining them on $A.x = B.x$, we read all of A then build a set of A.x that are passed as a criteria when querying B. In cases where A is small and B is large, this can drastically reduce the data retrieved from B, thus greatly speeding the overall query.

- RuleChooseJoinStrategy – Determines the base join strategy. Currently this is a decision as to whether to use a merge join rather than the default strategy, which is a nested loop join. Ideally the choice of a hash join would also be evaluated here. Also costing should be used to determine the strategy cost.
- - RuleCollapseSource - this rule removes all nodes below an Access node and collapses them into an equivalent query that is placed in the Access node.
- RuleAssignOutputElements - this rule walks top down through every node and calculates the output columns for each node. Columns that are not needed are dropped at every node. This is done by keeping track of both the columns needed to feed the parent node and also keeping track of columns that are “created” at a certain node.
- RuleValidateWhereAll - this rule validates a rarely used model option.
- RuleAccessPatternValidation – validates that all access patterns have been satisfied.

9.5.2. Procedure Planner

The procedure planner is fairly simple. It converts the statements in the procedure into instructions in a program that will be run during processing. This is mostly a 1-to-1 mapping and very little optimization is performed.

9.5.3. XML Planner

The XML Planner creates an XML plan that is relatively close to the end result of the Procedure Planner – a program with instructions. Many of the instructions are even similar (while loop, execute SQL, etc). Additional instructions deal with producing the output result document (adding elements and attributes).

The XML planner does several types of planning (not necessarily in this order):

- Document selection - determine which tags of the virtual document should be excluded from the output document. This is done based on a combination of the model (which marks parts of the document excluded) and the query (which may specify a subset of columns to include in the SELECT clause).
- Criteria evaluation - breaks apart the user's criteria, determine which result set the criteria should be applied to, and add that criteria to that result set query.
- Result set ordering - the query's ORDER BY clause is broken up and the ORDER BY is applied to each result set as necessary
- Result set planning - ultimately, each result set is planned using the relational planner and taking into account all the impacts from the user's query
- Program generation - a set of instructions to produce the desired output document is produced, taking into account the final result set queries and the excluded parts of the document. Generally, this involves walking through the virtual document in document order, executing queries as necessary and emitting elements and attributes.

XML programs can also be recursive, which involves using the same document fragment for both the initial fragment and a set of repeated fragments (each a new query) until some termination criteria or limit is met.

9.5.4. XQuery Planner

XQuery planning uses the embedded Saxon XQuery engine to compile the query. There is no direct optimization by Teiid.

Architecture

10.1. Terminology

- VM or Process – a processing node and Java VM. Now typically called a Process or “node” depending on context.
- Host – a machine that is “hosting” one or more VMs.
- Host controller – an app that runs on a host and can stop/start/control a VM
- VM controller – the component of a VM that starts up the VM, connects to the distributed message bus and registry, and triggers the service controller to start all the services for this VM.
- Service – a subsystem running in a VM (often in many VMs) and providing a related set of functionality
- Registry – a distributed service registry (one instance in each VM and host controller) that shares the status of current system (existence of VMs and services), and provides remote access to other services via RMI
- Service Controller – a component within the VM that the registry uses to start/stop/control services
- Repository database – the repository database stores the system configuration and other service-specific information (sessions, users, entitlements, vdfs, etc).

In addition to these main components, the service platform provides a core set of services available to applications built on top of the service platform. These services are:

- Session – the Session service manages active session information. Active sessions are stored in a distributed cache and shared between Session services in each VM. Sessions are also persisted in the server repository database.
- Membership – the Membership service manages authentication, users, and groups. This was redesigned in the 5.5 release to provide primary support for LDAP authentication and authorization. Custom membership modules can allow be developed as needed.
- Authorization – the Authorization service manages user entitlements. This service persists entitlements information in the repository database. Entitlements use is optional (as specified in the configuration) and off by default.

10.2. Data Management

10.2.1. Cursors and Batching

Teiid cursors all results, regardless of whether they are from one source or many sources, and regardless of what type of processing (joins, unions, etc.) have been performed on the results.

Teiid processes results in batches. A batch is simply a set of records. The number of rows in a batch is determined by the buffer system properties Processor Batch Size (within query engine) and Connector Batch Size (created at connectors).

Client applications have no direct knowledge of batches or batch sizes, but rather specify fetch size. However the first batch, regardless of fetch size is always proactively returned to synchronous clients. Subsequent batches are returned based on client demand for the data. Pre-fetching is utilized at both the client and connector levels.

10.2.2. Buffer Management

The buffer manager manages memory for all result sets used in the query engine. That includes result sets read from a connector binding, result sets used temporarily during processing, and result sets prepared for a user. Each result set is referred to in the buffer manager as a tuple source.

When retrieving batches from the buffer manager, the size of a batch in bytes is estimated and then allocated against the max and session session limits. If a limit is exceeded and memory space cannot be cleared for the batch, then processing can optionally give up its timeslice and try again.

10.2.2.1. Memory Management

The buffer manager has two storage managers - a memory manager and a disk manager. The buffer manager maintains the state of all the batches, and determines when batches must be moved from memory to disk.

10.2.2.2. Disk Management

Each tuple source has a dedicated file (named by the ID) on disk. This file will be created only if at least one batch for the tuple source had to be swapped to disk. The file is random access. The connector batch size and processor batch size properties define how many rows can exist in a batch and thus define how granular the batches are when stored into the storage manager. Batches are NOT removed from the file when they are swapped back into memory because that would require removing data out of the middle of the file and updating all the indexes which would be very expensive. Thus the disk storage manager never removes a particular batch. Batches are always read and written from the storage manager whole.

The disk storage manager has a cap on the maximum number of open files to prevent running out of file handles. In cases with heavy buffering, this can cause wait times while waiting for a file handle to become available - customers may want to increase the number of open files allowed (a configuration parameter defaulted to 10).

10.2.3. Cleanup

When a tuple source is no longer needed, it is removed from the buffer manager. The buffer manager will remove it from both the memory storage manager and the disk storage manager. The disk storage manager will delete the file. In addition, every tuple source is tagged with a "group

name" which is typically the session ID of the client. When the client's session is terminated (by closing the connection, server detecting client shutdown, or administrative termination), a call is sent to the buffer manager to remove all tuple sources for the session. This is a final cleanup mechanism that removes all state associated with a session.

In addition, when the query engine is shutdown, the buffer manager is shut down, which will remove all state from the disk storage manager and cause all files to be closed. In general, these mechanisms mean that the engine should always shut down with 0 open files. When the query engine is stopped, it is safe to delete any files in the buffer directory as they are not used across query engine restarts and must be due to a system crash where buffer files were not cleaned up.

10.3. Query Termination

10.3.1. Canceling Queries

If the client issues a 'cancel' command, then no results from the batch currently being processed in the server will be returned to the client.

When a query is canceled, processing will be stopped in the query engine and in all connectors involved in the query. The semantics of what a connector does in response to a cancellation command is dependent on the connector implementation. For example, JDBC connectors will asynchronously call cancel on the underlying JDBC driver, which may or may not actually support this method.

10.3.2. Timeouts

Timeouts in Teiid are managed on the client-side, in the JDBC API (which underlies both SOAP and ODBC access). Timeouts are only relevant for the first record returned. If the first record has not been received by the client within the specified timeout period, a 'cancel' command is issued to the server for the request and no results are returned to the client. The cancel command is issued by the JDBC API without the client's intervention.

10.4. Processing

10.4.1. Join Algorithms

Nested loop does the most obvious processing – for every row in the outer source, it compares with every row in the inner source. Nested loop is only used when the join criteria has no equi-join predicates.

Merge join first sorts the input sources on the joined columns. You can then walk through each side in parallel (effectively one pass through each sorted source) and when you have a match, emit a row. Because the inputs are sorted, you can skip through large portions of the input without comparing if one side is less than the other. In general, merge join is on the order of $n+m$ rather than $n*m$ in nested loop. When n and m are large, this makes a huge difference. Merge join is the default algorithm. It cannot support full outer join or non-equality criteria, but other than that handles almost all common cases well.

Any of the Join Algorithms above can be made into a dependent join (however hash joins would need new logic). The decision to implement a dependent join is considered after the join algorithm is chosen, and does not currently influence the algorithm selection.

10.4.2. Sort Based Algorithms

Sorting is used as the basis of the Sort (ORDER BY), Grouping (GROUP BY), and DupRemoval (SELECT DISTINCT) operations. The sort algorithm is a multi-pass merge-sort that does not require all of the result set to ever be in memory yet uses the maximal amount of memory allowed by the buffer manager.

It consists of two phases. The first phase (“sort”) will take an unsorted input stream and produce one or more sorted input streams. Each pass reads as much of the unsorted stream as possible, sorts it, and writes it back out as a new stream. Since the stream may be more than can fit in memory, this may result in many sorted streams.

The second phase (“merge”) consists of a set of phases that grab the next batch from as many sorted input streams as will fit in memory. It then repeatedly grabs the next tuple in sorted order from each stream and outputs merged sorted batches to a new sorted stream. At completion of the pass, all input streams are dropped. In this way, each pass reduces the number of sorted streams. When only one stream remains, it is the final output.

Appendix A. BNF Grammar

A.1. Terminals

<DEFAULT> SKIP : { " " "\t" "\n" "\r" }
<DEFAULT> MORE : { "/" * : IN_MULTI_LINE_COMMENT }
<IN_MULTI_LINE_COMMENT> SPECIAL : { <MULTI_LINE_COMMENT: "*" /> : DEFAULT }
<IN_MULTI_LINE_COMMENT> MORE : { <~[]> }
<DEFAULT> TOKEN : { <STRING: "string"> <BOOLEAN: "boolean"> <BYTE: "byte"> <SHORT: "short"> <CHAR: "char"> <INTEGER: "integer"> <LONG: "long"> <BIGINTEGER: "biginteger"> <FLOAT: "float"> <DOUBLE: "double"> <BIGDECIMAL: "bigdecimal"> <DATE: "date"> <TIME: "time"> <TIMESTAMP: "timestamp"> <OBJECT: "object"> <BLOB: "blob"> <CLOB: "clob"> <XML: "xml"> }
<DEFAULT> TOKEN : { <CAST: "cast"> <CONVERT: "convert"> <TIMESTAMPADD: "timestampadd"> <TIMESTAMPDIFF: "timestampdiff"> <COUNT: "count"> <SUM: "sum"> <AVG: "avg"> <MIN: "min"> <MAX: "max"> }
<DEFAULT> TOKEN : { <ALL: "all"> <AND: "and"> <ANY: "any"> <AS: "as"> <ASC: "asc"> <BEGIN: "begin"> <BETWEEN: "between"> <BREAK: "break"> <BY: "by"> <CASE: "case"> <CONTINUE: "continue"> <CREATE: "create"> <CRITERIA: "criteria"> <CROSS: "cross"> <DEBUG: "debug"> <DECLARE: "declare"> <DELETE: "delete"> <DESC: "desc"> <DISTINCT: "distinct"> <DROP: "drop"> <ELSE: "else"> <END: "end"> <ERROR: "error"> <ESCAPE: "escape"> <EXCEPT: "except"> <EXEC: "exec"> <EXECUTE: "execute"> <EXISTS: "exists"> <FALSE: "false"> <FN: "fn"> <FOR: "for"> <FROM: "from"> <FULL: "full"> <GROUP: "group"> <HAS: "has"> <HAVING: "having"> <IF: "if"> <IN: "in"> <INNER: "inner"> <INSERT: "insert"> <INTERSECT: "intersect"> <INTO: "into"> <IS: "is"> <JOIN: "join"> <LEFT: "left"> <LIKE: "like"> <LIMIT: "limit"> <LOCAL: "local"> <LOOP: "loop"> <MAKEDEP: "makedep"> <MAKENOTDEP: "makenotdep"> <NOCACHE: "nocache"> <NOT: "not"> <NULL: "null"> <ON: "on"> <OJ: "oj"> <OPTION: "option"> <OR: "or"> <ORDER: "order"> <OUTER: "outer"> <PLANONLY: "planonly"> <PROCEDURE: "procedure"> <RIGHT: "right"> <SELECT: "select"> <SET: "set"> <SHOWPLAN: "showplan"> <SOME: "some"> <TABLE: "table"> <TEMPORARY: "temporary"> <THEN: "then"> <TRANSLATE: "translate"> <TRUE: "true"> <UNION: "union"> <UNKNOWN: "unknown"> <UPDATE: "update"> <USING: "using"> <VALUES: "values"> <VIRTUAL: "virtual"> <WHEN: "when"> <WHERE: "where"> <WITH: "with"> <WHILE: "while"> }
<DEFAULT> TOKEN : { <SQL_TSI_FRAC_SECOND: "SQL_TSI_FRAC_SECOND"> <SQL_TSI_SECOND: "SQL_TSI_SECOND"> <SQL_TSI_MINUTE: "SQL_TSI_MINUTE"> <SQL_TSI_HOUR: "SQL_TSI_HOUR"> <SQL_TSI_DAY: "SQL_TSI_DAY"> <SQL_TSI_WEEK: "SQL_TSI_WEEK"> <SQL_TSI_MONTH: "SQL_TSI_MONTH"> <SQL_TSI_QUARTER: "SQL_TSI_QUARTER"> <SQL_TSI_YEAR: "SQL_TSI_YEAR"> }
<DEFAULT> TOKEN : { <ALL_IN_GROUP: (<GROUP_PART> <MMUID_PART>) <PERIOD> <STAR>> <VARIABLE: <ID> <MMUID>> <#ID: <GROUP_PART> ((<PERIOD>

```
| <SLASH>) (<QUOTED_ID> | <MMUUID_PART>))?)>| <#ELEMENT: <GROUP_PART>
(<PERIOD> | <SLASH>) <QUOTED_ID>>| <#GROUP_PART: ("#")? (<QUOTED_ID>
(<PERIOD> | <SLASH>))? <QUOTED_ID>>| <#QUOTED_ID: <DOTTED_ID> | "\"
<DOTTED_ID> "\">| <#DOTTED_ID: <ID_PART> ((<PERIOD> | <SLASH>) <ID_PART>)*>|
<#ID_PART: ("@"?)? <LETTER> (<ID_CHAR>)*>| <#ID_CHAR: <LETTER> | "_" | <DIGIT>>|
<#MMUUID: <MMUUID_PART> (<PERIOD> <MMUUID_PART>)?>| <#MMUUID_PART:
"mmuuid:" (<MMUUID_CHAR>)*>| <#MMUUID_CHAR: ["a"-“f”] | ["0"-“9”] | “-”>| <DATETYPE:
“{” “d”>| <TIMETYPE: “{” “t”>| <TIMESTAMPSTYPE: “{” “ts”>| <BOOLEANTYPE: “{”
“b”>| <INTEGERTYPE: (<MINUS>)? (<DIGIT>)+>| <FLOATVAL: (<MINUS>)? (<DIGIT>)*
<PERIOD> (<DIGIT>)+ (“e”,“E”) (“+”,“-”)? (<DIGIT>)+>?>| <STRINGVAL: (“N”)? (<STRINGA>
| <STRINGB>)>| <#STRINGA: "\" (~[“\”])* (“\” (~[“\”])* “\”>| <#STRINGB: "\" (~[“\”])*
(“\” (~[“\”])* “\”>| <#LETTER: ["a"-“z”,“A"-“Z”] | [“\u0153”-“\ufffd”]>| <#DIGIT: ["0"-“9”]>|
<#COLON: “:”>}
```

```
<DEFAULT> TOKEN : {<COMMA: “,”>| <PERIOD: “.”>| <LPAREN: “(”>| <RPAREN: “)”>|
<LBRACE: “{”>| <RBRACE: “}”>| <EQ: “=”>| <NE: “<>”>| <NE2: “!=”>| <LT: “<”>| <LE: “<=”>|
<GT: “>”>| <GE: “>=”>| <STAR: “*”>| <SLASH: “/”>| <PLUS: “+”>| <MINUS: “-”>| <QMARK: “?”>|
<DOLLAR: “$”>| <SEMICOLON: “;”>| <CONCAT_OP: “||”>}
```

A.2. Non-Terminals

command ::=	((<i>createUpdateProcedure</i>) <i>userCommand</i> <i>callableStatement</i>) (<SEMICOLON>)? <EOF>
userCommand ::=	(<i>queryExpression</i> <i>storedProcedure</i> <i>insert</i> <i>update</i> <i>delete</i> <i>dropTable</i> <i>createTempTable</i>)
dropTable ::=	<DROP> <TABLE> <VARIABLE>
createTempTable ::=	<CREATE> <LOCAL> <TEMPORARY> <TABLE> <VARIABLE> <LPAREN> <i>createElementsWithTypes</i> <RPAREN>
errorStatement ::=	<ERROR> <i>expression</i>
statement ::=	(<i>ifStatement</i> <i>loopStatement</i> <i>whileStatement</i> <i>delimitedStatement</i>)
delimitedStatement ::=	(<i>sqlStatement</i> <i>errorStatement</i> <i>assignStatement</i> <i>declareStatement</i> <i>continueStatement</i> <i>breakStatement</i>) <SEMICOLON>
block ::=	(<i>statement</i> (<BEGIN> (<i>statement</i>)* <END>))
breakStatement ::=	<BREAK>
continueStatement ::=	<CONTINUE>
whileStatement ::=	<WHILE> <LPAREN> <i>criteria</i> <RPAREN> <i>block</i>
loopStatement ::=	<LOOP> <ON> <LPAREN> <i>queryExpression</i> <RPAREN> <AS> <VARIABLE> <i>block</i>
ifStatement ::=	<IF> <LPAREN> <i>criteria</i> <RPAREN> <i>block</i> (<ELSE> <i>block</i>)?

criteriaSelector ::=	(((<EQ> <NE> <NE2> <LE> <GE> <LT> <GT> <IN> <LIKE> (<IS> <NULL>) <BETWEEN>))? <CRITERIA> (<ON> <LPAREN> <VARIABLE> (<COMMA> <VARIABLE>)* <RPAREN>)?)
hasCriteria ::=	<HAS> <i>criteriaSelector</i>
declareStatement ::=	<DECLARE> <i>dataType</i> <VARIABLE> (<EQ> <i>assignStatementOperand</i>)?
assignStatement ::=	<VARIABLE> <EQ> <i>assignStatementOperand</i>
assignStatementOperand ::=	((<i>insert</i>) <i>update</i> <i>delete</i> <i>storedProcedure</i> (<i>expression</i>) <i>queryExpression</i>)
sqlStatement ::=	((<i>dynamicCommand</i>) <i>userCommand</i>)
translateCriteria ::=	<TRANSLATE> <i>criteriaSelector</i> (<WITH> <LPAREN> <VARIABLE> <EQ> <i>expression</i> (<COMMA> <VARIABLE> <EQ> <i>expression</i>)* <RPAREN>)?
createUpdateProcedure ::=	<CREATE> (<VIRTUAL>)? (<UPDATE>)? <PROCEDURE> <i>block</i>
dynamicCommand ::=	(<EXECUTE> <EXEC>) <STRING> <i>expression</i> (<AS> <i>createElementsWithTypes</i> (<INTO> <VARIABLE>))? (<USING> <i>setClauseList</i>)? (<UPDATE> ((<INTERVAL>) (<STAR>)))?)
setClauseList ::=	<VARIABLE> <EQ> (<COMMA> <VARIABLE> <EQ>)*
createElementsWithTypes ::=	<VARIABLE> <i>dataType</i> (<COMMA> <VARIABLE> <i>dataType</i>)*
callableStatement ::=	<LBRACE> (<QMARK> <EQ>)? <VARIABLE> <VARIABLE> (<LPAREN> (<i>executeUnnamedParams</i>) <RPAREN>)? <RBRACE> (<i>option</i>)?
storedProcedure ::=	(((<EXEC> <EXECUTE>) <VARIABLE> <LPAREN> (<i>executeNamedParams</i> <i>executeUnnamedParams</i>) <RPAREN>) (<i>option</i>)?)
executeUnnamedParams ::=	(<i>expression</i> (<COMMA> <i>expression</i>)*)?
executeNamedParams ::=	(<i>paramName</i> <EQ> <i>expression</i> (<COMMA> <i>paramName</i> <EQ> <i>expression</i>)*)
paramName ::=	<VARIABLE>
insert ::=	<INSERT> <INTO> <VARIABLE> (<LPAREN> <VARIABLE> (<COMMA> <VARIABLE>)* <RPAREN>)? ((<VALUES> <i>rowValues</i>) (<i>queryExpression</i>)) (<i>option</i>)?
rowValues ::=	<LPAREN> <i>expression</i> (<COMMA> <i>expression</i>)* <RPAREN>

Appendix A. BNF Grammar

update ::=	<UPDATE> <VARIABLE> <SET> <i>setClauseList</i> (<i>where</i>)? (<i>option</i>)?
delete ::=	<DELETE> <FROM> <VARIABLE> (<i>where</i>)? (<i>option</i>)?
queryExpression ::=	<i>queryExpressionBody</i>
queryExpressionBody ::=	<i>queryTerm</i> ((<UNION> <EXCEPT>) (<ALL> <DISTINCT>)? <i>queryTerm</i>)* (<i>orderBy</i>)? (<i>limit</i>)? (<i>option</i>)?
queryTerm ::=	<i>queryPrimary</i> (<INTERSECT> (<ALL> <DISTINCT>)? <i>queryPrimary</i>)*
queryPrimary ::=	(<i>query</i> (<LPAREN> <i>queryExpressionBody</i> <RPAREN>))
query ::=	<i>select</i> (<i>into</i>)? (<i>from</i> (<i>where</i>)? (<i>groupBy</i>)? (<i>having</i>)?)?
into ::=	<INTO> (<VARIABLE>)
select ::=	<SELECT> (<ALL> (<DISTINCT>))? (<STAR> (<i>selectSymbol</i> (<COMMA> <i>selectSymbol</i>)*))
selectSymbol ::=	((<ALL_IN_GROUP>) (<i>expression</i>) ((<AS>)? (<VARIABLE> <STRINGVAL>))?)
aggregateSymbol ::=	((<COUNT> <LPAREN> <STAR> <RPAREN>) ((<COUNT> <SUM> <AVG> <MIN> <MAX>) <LPAREN> (<DISTINCT>)? <i>expression</i> <RPAREN>))
from ::=	<FROM> (<i>tableReference</i> (<COMMA> <i>tableReference</i>)*)
tableReference ::=	((<LBRACE> <OJ> <i>tableReferenceUnescaped</i> <RBRACE>) <i>tableReferenceUnescaped</i>)
tableReferenceUnescaped ::=	(<i>joinedTable</i> <i>tablePrimary</i>)
joinedTable ::=	<i>tablePrimary</i> ((<i>crossJoin</i> <i>qualifiedJoin</i>))+
crossJoin ::=	((<CROSS> <UNION>) <JOIN> <i>tablePrimary</i>)
qualifiedJoin ::=	(((<RIGHT> (<OUTER>)?) (<LEFT> (<OUTER>)?)) (<FULL> (<OUTER>)?) <INNER>)? <JOIN> <i>tableReference</i> <ON> <i>criteria</i>)
tablePrimary ::=	(<i>unaryFromClause</i> <i>subqueryFromClause</i> (<LPAREN> <i>joinedTable</i> <RPAREN>)) ((<MAKEDEP>) (<MAKENOTDEP>))?
subqueryFromClause ::=	<LPAREN> (<i>queryExpression</i> <i>storedProcedure</i>) <RPAREN> (<AS>)? <VARIABLE>
unaryFromClause ::=	(<VARIABLE> ((<AS>)? <VARIABLE>)?)
where ::=	<WHERE> <i>criteria</i>
criteria ::=	<i>compoundCritOr</i>
compoundCritOr ::=	<i>compoundCritAnd</i> (<OR> <i>compoundCritAnd</i>)*
compoundCritAnd ::=	<i>notCrit</i> (<AND> <i>notCrit</i>)*

notCrit ::=	(<NOT>)? <i>primary</i>
primary ::=	(<i>predicate</i> (<LPAREN> <i>criteria</i> <RPAREN>))
predicate ::=	(<i>subqueryCompareCriteria</i> <i>compareCrit</i> <i>matchCrit</i> <i>betweenCrit</i> <i>setCrit</i> <i>existsCriteria</i> <i>hasCriteria</i> <i>translateCriteria</i> <i>isNullCrit</i>)
compareCrit ::=	<i>expression</i> (<EQ> <NE> <NE2> <LT> <LE> <GT> <GE>) <i>expression</i>
subquery ::=	<LPAREN> (<i>queryExpression</i> <i>storedProcedure</i>) <RPAREN>
subqueryCompareCriteria ::=	<i>expression</i> (<EQ> <NE> <NE2> <LT> <LE> <GT> <GE>) (<ANY> <SOME> <ALL>) <i>subquery</i>
matchCrit ::=	(<i>expression</i> (<NOT>)? <LIKE> <i>expression</i> (<i>escapeChar</i> (<LBRACE> <i>escapeChar</i> <RBRACE>))?)
escapeChar ::=	<ESCAPE> <STRINGVAL>
betweenCrit ::=	<i>expression</i> (<NOT>)? <BETWEEN> <i>expression</i> <AND> <i>expression</i>
isNullCrit ::=	<i>expression</i> <IS> (<NOT>)? <NULL>
setCrit ::=	<i>expression</i> (<NOT>)? <IN> ((<i>subquery</i>) (<LPAREN> <i>expression</i> (<COMMA> <i>expression</i>)* <RPAREN>))
existsCriteria ::=	<EXISTS> <i>subquery</i>
groupBy ::=	<GROUP> <BY> (<i>groupByItem</i> (<COMMA> <i>groupByItem</i>)*)
groupByItem ::=	<i>expression</i>
having ::=	<HAVING> <i>criteria</i>
orderby ::=	<ORDER> <BY> (<VARIABLE> <STRINGVAL> <INTERVAL>) (<ASC> <DESC>)? (<COMMA> (<VARIABLE> <STRINGVAL> <INTERVAL>) (<ASC> <DESC>)?)*
limit ::=	<LIMIT> (<INTERVAL> <QMARK>) (<COMMA> (<INTERVAL> <QMARK>))?
option ::=	<OPTION> (<SHOWPLAN> <PLANONLY> <DEBUG> <MAKEDEP> <VARIABLE> (<COMMA> <VARIABLE>)* <MAKENOTDEP> <VARIABLE> (<COMMA> <VARIABLE>)* <NOCACHE> (<VARIABLE> (<COMMA> <VARIABLE>)*)?)*
expression ::=	<i>concatExpression</i>
concatExpression ::=	(<i>plusExpression</i> (<CONCAT_OP> <i>plusExpression</i>)*)
plusExpression ::=	(<i>timesExpression</i> (<i>plusOperator</i> <i>timesExpression</i>)*)

Appendix A. BNF Grammar

plusOperator ::=	(<PLUS> <MINUS>)
timesExpression ::=	(<i>basicExpression</i> (<i>timesOperator</i> <i>basicExpression</i>) [*])
timesOperator ::=	(<STAR> <SLASH>)
basicExpression ::=	(<QMARK> <i>literal</i> (<LBRACE> <FN> <i>function</i> <RBRACE>) (<i>aggregateSymbol</i>) (<i>function</i>) (<VARIABLE>) (<LPAREN> <i>expression</i> <RPAREN>) <i>subquery</i> <i>caseExpression</i> <i>searchedCaseExpression</i>)
caseExpression ::=	<CASE> <i>expression</i> (<WHEN> <i>expression</i> <THEN> <i>expression</i>) ⁺ (<ELSE> <i>expression</i>)? <END>
searchedCaseExpression ::=	<CASE> (<WHEN> <i>criteria</i> <THEN> <i>expression</i>) ⁺ (<ELSE> <i>expression</i>)? <END>
function ::=	((<CONVERT> <LPAREN> <i>expression</i> <COMMA> <i>dataType</i> <RPAREN>) (<CAST> <LPAREN> <i>expression</i> <AS> <i>dataType</i> <RPAREN>) ((<TIMESTAMPADD> <TIMESTAMPDIFF>) <LPAREN> <i>intervalType</i> <COMMA> <i>expression</i> <COMMA> <i>expression</i> <RPAREN>) ((<LEFT> <RIGHT> <CHAR>) <LPAREN> (<i>expression</i> (<COMMA> <i>expression</i>) [*])? <RPAREN>) ((<INSERT>) <LPAREN> (<i>expression</i> (<COMMA> <i>expression</i>) [*])? <RPAREN>) ((<TRANSLATE>) <LPAREN> (<i>expression</i> (<COMMA> <i>expression</i>) [*])? <RPAREN>) (<VARIABLE> <LPAREN> (<i>expression</i> (<COMMA> <i>expression</i>) [*])? <RPAREN>))
dataType ::=	(<STRING> <BOOLEAN> <BYTE> <SHORT> <CHAR> <INTEGER> <LONG> <BIGINTEGER> <FLOAT> <DOUBLE> <BIGDECIMAL> <DATE> <TIME> <TIMESTAMP> <OBJECT> <BLOB> <CLOB> <XML>)
intervalType ::=	(<SQL_TSI_FRAC_SECOND> <SQL_TSI_SECOND> <SQL_TSI_MINUTE> <SQL_TSI_HOUR> <SQL_TSI_DAY> <SQL_TSI_WEEK> <SQL_TSI_MONTH> <SQL_TSI_QUARTER> <SQL_TSI_YEAR>)
literal ::=	(<STRINGVAL> <INTEGerval> <FLOATVAL> <FALSE> <TRUE> <UNKNOWN> <NULL> (<BOOLEANType> <STRINGVAL> <RBRACE>) (<TIMESTAMPType> <STRINGVAL> <RBRACE>) (<DATeType> <STRINGVAL> <RBRACE>) (<TIMeType> <STRINGVAL> <RBRACE>))