

However, if the user later adds an SP by hand using the name;1, this results in a duplicate name exception being thrown while the file is loaded (i.e., spName=spName;number).

[0059] Relationships are a superset of the foreign key constraints that are stored in the database and can also represent relationships that are enforced by triggers, the business logic, or any other means. Relationships can be specified across databases within the RSD and have a unique name within the scope of the database that contains them.

[0060] Two more validation rules specific to SQL Server are AliasTypes and PrimaryKeys. AliasTypes has a unique 1-part name within the scope of the database in which they are defined, and are not referenced across databases. PrimaryKeys has a unique 1-part name within the scope of the database in which they are defined.

[0061] Columns has a unique name within the scope of the Table to which they belong.

[0062] Parameters has a unique name within the scope of the Routine (Stored Procedure or UDF) to which they belong.

[0063] If any of the syntactic or semantic constraints are violated, the system throws an exception that at a minimum, lists the filename/url, structure type (Table, Relationship, etc.) and the nature and/or root cause of the error. In more robust implementations, a listing of the line numbers in the persistent RSD format that contain RSD language constructs that are causing the errors can be provided. Furthermore, a listing of all the syntactic and semantic errors in one exception (or allowing the user to iterate through a list of exceptions) can also be provided to help speed the debugging process.

[0064] In some cases, multi-part names are used to reference structures, relationships, and types either within the RSD or from the MSD (Mapping Schema Definition) architecture, in which case escaping may be required if the name does not conform to the RSD naming guidelines.

[0065] Within the RSD architecture, file names can be referenced by other structures. The references that are possible are the following. Relationships can reference Tables or CustomTables using 1-, 2-, or 3-part names. For a 1-part name, the table name is resolved using the database that contains the Relationship and the Default Schema. Since relationships are not stored in the context of a schema, they do not have a "current" schema. For a 2-part name, the table name is resolved using the database that contains the Relationship. For a 3-part name, the table name is specified explicitly.

[0066] Alias types can be referenced from Columns or Parameters using a 0-part name. Names are resolved using the current (parent) database for the structure that contains the Column or Parameter to resolve the name. AliasTypes cannot be referenced across databases.

[0067] CustomTables has command instances that can reference other RSD structures using 1-, 2- or 3-part names. For a 1-part name, the structure name is resolved using the current database and the current schema. For a 2-part name, the structure name is resolved using the current database. For a 3-part name, the structure name is specified explicitly.

[0068] With respect to RSD reference validation, name references resolve to valid structures in the logical RSD, although they may be stored in separate physical segments. The one exception is Relationships that does not attempt to validate the remote end of the Relationship, if the structure is not available.

[0069] When referencing elements in the RSD domain from a mapping file of the mapping component, default attributes are used on the database and schema to resolve 1-part and 2-part names to fully-qualified names. References from the mapping file to the RSD file take place within the scope of a DataSource.

[0070] All references from the MSD to the RSD structures/relationships need to be valid in the logical RSD that is in scope for a particular logical mapping.

[0071] Case-sensitivity of name references depends on the "CaseSensitive" property in the RSD. Name references within the RSD and, from the MSD to the RSD are treated as case-sensitive. The implication is that authors should avoid name mismatches, but it allows the representation of all structures from all case-sensitive as well as case-insensitive database instances.

[0072] Using namespaces avoids name clashes (i.e., situations where the same tag name is used in different contexts). For instance, a namespace can identify whether an address is a postal address, an e-mail address, or an IP address. Tag names within a namespace must be unique.

[0073] There are some SQL Server specific elements in this namespace, so that SQL is treated as a first class citizen in the mapping framework, but none of the SQL specific elements are required so as not to preclude the RSD from describing other DBMS (Database Management System) platforms, as well by including other namespaces.

[0074] In the event that a loading component encounters validation errors it caches the information and continues processing in order to return a comprehensive list to the user of the problems with the file, rather than performing 1-off debugging. XML syntax and schema exceptions halt processing.

[0075] RSD Format

[0076] The XSD for RSD is the implementation of the following description of the RSD format.

[0077] 1. Databases

[0078] This is the root tag of the RSD file, and represents zero or more of the databases that can be accessed with a single connection. The DatabaseCollection element is referenced by the DataSource in the mapping file by pointing to the file or files that define the Databases element.

Tag	Card	Comment
<u>Attributes</u>		
DefaultDatabase	Opt	Must be the name of a Database element within the Databases. If no default is specified by the user, it will be the first database in the Databases. Used to enable 2-part naming.