

RTI Database Integration Service

User's Manual

Version 5.2.3



© 2016 Real-Time Innovations, Inc.
All rights reserved.
Printed in U.S.A. First printing.
April 2016.

Trademarks

Real-Time Innovations, RTI, NDDS, RTI Data Distribution Service, DataBus, Connex, Micro DDS, the RTI logo, 1RTI and the phrase, “Your Systems. Working as one,” are registered trademarks, trademarks or service marks of Real-Time Innovations, Inc. All other trademarks belong to their respective owners.

Third Party Copyright Notices

The Oracle® TimesTen® In-Memory Database and the Oracle® Database are products of Oracle.

Copy and Use Restrictions

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form (including electronic, mechanical, photocopy, and facsimile) without the prior written permission of Real-Time Innovations, Inc. All software and documentation (whether in hard copy or electronic form) enclosed are subject to the license agreement. The software and documentation may be used or copied only under the terms of the license agreement.

The programs in this book have been included for their instructional value. RTI does not offer any warranties or representations in respect of their fitness for a particular purpose, nor does RTI accept any liability for any loss or damage arising from their use.

Technical Support

Real-Time Innovations, Inc.
232 E. Java Drive
Sunnyvale, CA 94089
Phone: (408) 990-7444
Email: support@rti.com
Website: <https://support.rti.com/>

Available Documentation

The following documentation is available for *RTI[®] Database Integration Service*:

- ❑ The [Release Notes](#), **RTI_Database_Integration_Service_ReleaseNotes.pdf**. This document provides an overview of the current release's features and lists changes since the previous release, system requirements, supported architectures, and compatibility with other products.
- ❑ The [Getting Started Guide](#), **RTI_Database_Integration_Service_GettingStarted.pdf**. This document provides installation instructions, a short 'Hello World' tutorial, and troubleshooting tips.
- ❑ This User's Manual starts with an overview of *RTI Database Integration Service's* basic concepts, terminology, and unique features. It then describes how to develop and implement applications that use *RTI Database Integration Service*.

Additional Resources:

- ❑ The *ODBC API Reference* from Microsoft is available from [http://msdn.microsoft.com/en-us/library/ms714562\(VS.85\).aspx](http://msdn.microsoft.com/en-us/library/ms714562(VS.85).aspx).
- ❑ The documentation for the Oracle TimesTen In-Memory Database can be found in the **doc/** directory of the Oracle TimesTen installation.
- ❑ The documentation for Oracle databases can be found here: <http://www.oracle.com/technology/documentation/index.html>.
- ❑ The documentation for MySQL databases can be found here: <http://dev.mysql.com/doc/refman/5.1/en/index.html>.

Contents

1	Welcome to RTI Database Integration Service.....	1-1
1.1	Intended Readers.....	1-1
1.2	Paths Mentioned in Documentation.....	1-1
1.3	Background Reading.....	1-3
2	Introduction.....	2-1
2.1	The Edge to Enterprise Integration Solution.....	2-1
2.2	Database Integration Service's Unique Features.....	2-2
2.2.1	Interconnecting Standards.....	2-3
2.2.2	Connectivity To Edge Devices.....	2-3
2.2.3	Flexibility and Scalability.....	2-4
2.2.4	Matching Real-Time Performance.....	2-4
2.2.5	High Availability.....	2-4
2.2.6	Additional Benefits of Database Integration Service.....	2-4
3	Architecture.....	3-1
3.1	Database Integration Service Architecture.....	3-1
3.1.1	Database Integration Service Daemon.....	3-1
3.1.2	Database Integration Service's Unique Features.....	3-2
3.2	Capturing Real-Time Data in a DBMS.....	3-3
3.3	Remote Real-Time Notification of Table Changes.....	3-4
3.4	Bidirectional Integration.....	3-4
3.5	Bridging between Domains.....	3-5
3.6	High-Rate Data Streams Cached before Storage.....	3-5
3.7	Real-Time Database Replication.....	3-6
4	Using Database Integration Service.....	4-1
4.1	Introduction to the Database Integration Service Daemon.....	4-2
4.1.1	How to Run the Database Integration Service Daemon with Oracle.....	4-2
4.1.2	How to Run the Database Integration Service Daemon with MySQL.....	4-4
4.1.3	How to Run the Database Integration Service Daemons as Windows Services ..	4-6
4.1.4	Typecodes.....	4-6
4.2	Command-Line Parameters.....	4-7
4.3	Environment Variables.....	4-10
4.4	Configuration File.....	4-11
4.4.1	How to Load the XML Configuration.....	4-11
4.4.2	XML Syntax and Validation.....	4-12
4.4.3	Top-Level XML Tags.....	4-13
4.4.4	Database Configuration Using the Database Integration Service XML Tag.....	4-15

4.5	Meta-Tables.....	4-26
4.5.1	Publications Table.....	4-26
4.5.2	Subscriptions Table.....	4-38
4.5.3	Table Info	4-54
4.5.4	Log Table.....	4-55
4.6	User-Table Creation.....	4-57
4.7	Support for Extensible Types	4-59
4.8	Enabling RTI Distributed Logger in Database Integration Service.....	4-60
4.9	Enabling RTI Monitoring Library in Database Integration Service	4-60
5	IDL/SQL Semantic and Data Mapping	5-1
5.1	Semantic Mapping.....	5-1
5.2	Data Representation Mapping.....	5-3
5.2.1	IDL to SQL Mapping.....	5-3
5.2.2	Primitive Types Mapping.....	5-6
5.2.3	Oracle In-Memory Database Cache Mapping	5-9
5.2.4	Bit Field Mapping.....	5-9
5.2.5	Enum Types Mapping.....	5-11
5.2.6	Simple IDL Structures.....	5-11
5.2.7	Complex IDL Structures	5-11
5.2.8	Array Fields.....	5-12
5.2.9	Sequence Fields.....	5-13
5.2.10	NULL Values	5-14
5.2.11	Sparse Data Types.....	5-14
A	Error Codes.....	A-1
B	Database Limits	B-1
B.1	Maximum Columns for Oracle 11g.....	B-2
B.2	Maximum Columns for MySQL.....	B-2

Chapter 1 Welcome to RTI Database Integration Service

Welcome to *RTI® Database Integration Service*—a high-performance solution for integrating applications and data across real-time and enterprise systems from RTI.

Database Integration Service is the integration of two complementary technologies: data-centric publish-subscribe middleware and relational database management systems (RDBMS). This powerful integration allows your applications to uniformly access data from real-time/embedded and enterprise data sources via *RTI Connex*[™] *DDS*, or via database interfaces. Since both these technologies are data-centric and complementary, they can be combined to enable a new class of applications. In particular, *Connex* can be used to produce a truly decentralized, distributed RDBMS, while RDBMS technology can be used to provide persistence for real-time data.

1.1 Intended Readers

This document presents the general concepts behind *Database Integration Service's* architecture and provides basic information on how to develop applications using *Database Integration Service*.

The chapters assume general knowledge of relational databases and SQL, familiarity with the ODBC API, IDL and the *Connex* API, and a working knowledge of the C/C++ programming languages.

1.2 Paths Mentioned in Documentation

The documentation refers to:

□ <NDDSHOME>

This refers to the installation directory for *Connex DDS*.

The default installation paths are:

- Mac OS X systems:
/Applications/rti_connex_dds-version
- UNIX-based systems, non-*root* user:
/home/your user name/rti_connex_dds-version

- UNIX-based systems, *root* user:
/opt/rti_connex_dds-version
- Windows systems, user without Administrator privileges:
<your home directory>\rti_connex_dds-version
- Windows systems, user with Administrator privileges:
C:\Program Files\rti_connex_dds-version (for 64-bits machines) or
C:\Program Files (x86)\rti_connex_dds-version (for 32-bit machines)

You may also see \$NDDSHOME or %NDDSHOME%, which refers to an environment variable set to the installation path.

Wherever you see <NDDSHOME> used in a path, replace it with your installation path.

Note for Windows Users: When using a command prompt to enter a command that includes the path **C:\Program Files** (or any directory name that has a space), enclose the path in quotation marks. For example:

```
"C:\Program Files\rti_connex_dds-version\bin\rtiddsgen"
```

or if you have defined the NDDSHOME environment variable:

```
"%NDDSHOME%\bin\rtiddsgen"
```

❑ RTI Workspace directory, **rti_workspace**

The RTI Workspace is where all configuration files for the applications and example files are located. All configuration files and examples are copied here the first time you run *RTI Launcher* or any script in <NDDSHOME>/bin. The default path to the RTI Workspace directory is:

- Mac OS X systems:
/Users/your user name/rti_workspace
- UNIX-based systems:
/home/your user name/rti_workspace
- Windows systems:
your Windows documents folder\rti_workspace

Note: '*your Windows documents folder*' depends on your version of Windows.

For example, on Windows 7, the folder is **C:\Users\your user name\Documents**; on Windows Server 2003, the folder is **C:\Documents and Settings\your user name\Documents**.

You can specify a different location for the **rti_workspace** directory. See the *RTI Connex DDS Core Libraries Getting Started Guide* for instructions.

❑ <path to examples>

Examples are copied into your home directory the first time you run *RTI Launcher* or any script in <NDDSHOME>/bin. This document refers to the location of these examples as <path to examples>. Wherever you see <path to examples>, replace it with the appropriate path.

By default, the examples are copied to **rti_workspace/version/examples**.

So the paths are:

- Mac OS X systems:
/Users/your user name/rti_workspace/version/examples
- UNIX-based systems:
/home/your user name/rti_workspace/version/examples
- Windows systems:
your Windows documents folder\rti_workspace\version\examples

Note: 'your Windows documents folder' is described above.

You can specify that you do not want the examples copied to the workspace. See the *RTI Connexx DDS Core Libraries Getting Started Guide* for instructions.

1.3 Background Reading

For information on distributed systems and databases:

- ❑ George Coulouris, Jean Dollimore, Tim Kindberg. *Distributed Systems: Concepts and Design (3rd edition)*. Addison-Wesley, 2000
- ❑ M. Tamer Ozsu, Patrick Valduriez. *Principles of Distributed Database Systems (2nd Edition)*. Prentice Hall, 1999
- ❑ Andrew S. Tanenbaum, Maarten van Steen. *Distributed Systems: Principles and Paradigms (1st edition)*. Prentice Hall, 2002

For information on real-time systems:

- ❑ Qing Li, Caroline Yao. *Real-Time Concepts for Embedded Systems*. CMP Books, 2003
- ❑ Doug Abbott. *Linux for Embedded and Real-Time Applications*. Butterworth-Heinemann, 2002
- ❑ David E. Simon. *An Embedded Software Primer*. Addison-Wesley, 1999

For information on SQL:

- ❑ Joe Celko. *Joe Celko's SQL for Smarties: Advanced SQL Programming (expanded 2nd Edition)*. Morgan Kaufmann, 1999
- ❑ Rick van der Lans. *Introduction to SQL: Mastering the Relational Database Language (3rd edition)*. Addison-Wesley, 1999

For information on ODBC:

- ❑ Microsoft Corporation. *Microsoft ODBC 2.0 Software Development Kit and Programmer's Reference*. Microsoft Press, 1997

For information on the C programming language:

- ❑ Brian W. Kernighan, Dennis M. Ritchie. *The C programming Language (2nd edition)*. Prentice Hall Software Series, 1988

Chapter 2 Introduction

This chapter presents a conceptual view of *Database Integration Service's* architecture and highlights its unique features. It includes the following sections:

- ❑ [The Edge to Enterprise Integration Solution \(Section 2.1\)](#)
- ❑ [Database Integration Service's Unique Features \(Section 2.2\)](#)

2.1 The Edge to Enterprise Integration Solution

Database Integration Service is a solution for integrating existing applications, including service-oriented architectures (SOAs), with high performance real-time applications, data, and edge devices. *Database Integration Service* provides a run-time bridge between RTI's high-throughput, embeddable messaging infrastructure, *Connex DDS*, and integration and data management standards such as SQL, XML, Web services and JMS. This allows developers to benefit from the performance, scalability, Quality of Service (QoS) control and broad platform support provided by RTI's messaging technology while retaining interoperability with existing enterprise applications, see [Figure 2.1](#).


Database Integration Service is a fully standards-based solution for the integration of enterprise applications and high performance real-time applications. Enterprise applications typically use Structured Query Language (SQL) and Extensible Markup Language (XML) for data access and Java Message Service (JMS), Service-Oriented Architectures (SOAs), and Web services for integration.

The most commonly adopted standard in high performance real-time systems, for both integration and data access, is Object Management Group's (OMG) Data Distribution Service for Real-Time Systems (DDS). *Database Integration Service* enables interoperability by bridging between enterprise and embedded standards at the data and communications levels, allowing existing applications to be integrated with few or no changes.

Database Integration Service also mitigates the performance mismatch between enterprise and embedded applications. Real-time applications using *Connex DDS* can deliver messages and data at rates between 10 and 100 times faster than can be accepted by enterprise applications. *Database Integration Service* decouples applications at the data level, preventing high throughput real-time messages and data from overwhelming applications managing business processes by buffering and storing all data in an in-memory DBMS, the Oracle TimesTen In-Memory database which is bundled with *Database Integration Service*.

Oracle TimesTen is a memory-optimized database delivering exceptionally high performance, enabling it to easily support the high data throughput of real-time and embedded applications.

Figure 2.1 Database Integration Service Bridges Embedded and Enterprise Applications



Using the Oracle TimesTen Cache Connect to Oracle, the data and messages buffered by the in-memory database can be synchronized and stored in a disk-based Oracle database for access by enterprise applications directly or through enterprise-level integration options such as a BPEL (Business Process Execution Language) process manager.


Multiple copies of *Database Integration Service* can also run in parallel on separate systems. Different instances can be responsible for a discrete set of messages and data, maximizing throughput and providing load balancing. Or each instance can bridge the same messages and data providing high availability and fault tolerance.

2.2 Database Integration Service's Unique Features

In this section, a few of the unique qualities and features of *Database Integration Service* are discussed in greater detail. [Figure 2.2](#) shows an example system where *Database Integration Service* serves as the central integration technology to interconnect the real-time, embedded world with the analysis and high-level decision-making processes of the enterprise world.

In [Figure 2.2](#), sensors of physical processes produce data that must be filtered, fused, and stored for use in business processes. In addition, multiple user consoles must have concurrent access to both raw and fused data. *Database Integration Service* is used to store the raw data at high rates into an Oracle TimesTen In-Memory database where Oracle Cache Connect is then used to propagate the data into an Oracle 11g database for permanent storage. Enterprise applications used for analysis or applying other business logic can access the data stored in Oracle using SQL or other standards such as JMS, XML, and HTTP via a BPEL process manager.

Figure 2.2 Example System Using Database Integration Service



Database Integration Service is the bridge that connects real-time/high performance to complex analysis, edge devices to business systems, and embedded to enterprise.

2.2.1 Interconnecting Standards

Until recently, distributed real-time systems were built using custom-developed data structures and algorithms to store and manipulate data in combination with a commercial, or even proprietary, data-distribution middleware layer. This was necessary to meet real-time performance requirements. However, in recent years, DDS, a standard for data distribution, has emerged as the premier method to integrate and build distributed real-time systems.

For decades in enterprise systems, standards for communications, data representation and data storage has enabled the tremendous growth of software applications for business processes worldwide. The standards such as SQL, ODBC, JMS, HTML, XML, and WDSL have greatly increased the interoperability of those business systems.

Database Integration Service is the first commercial product that interconnects the DDS standard newly established in the embedded world to the common standards of the enterprise world. With *Database Integration Service*, enterprise applications have direct access to real-time data, and real-time applications have access to the plethora of processes and logic that has been developed to configure and direct actions based on business decisions.

2.2.2 Connectivity To Edge Devices

For edge devices, such as sensors and hand-helds, *Database Integration Service* integrates *Connex DDS* applications with databases. Applications can publish data into relational databases and subscribe to changes in relational databases using the standard *Connex DDS* application programming interface. Integration between *Connex DDS* and relational database applications is

supported by an IDL-to-SQL mapping that allows both types of applications to access a uniform data model.

2.2.3 Flexibility and Scalability

By leveraging *Connex DDS* Quality-of-Service (QoS) settings, *Database Integration Service* supports an unprecedented variety of deployment configurations to accommodate a wide range of scenarios, from reliable point-to-point delivery to best-effort multicasting that enables real-time transaction streaming to large numbers of subscribers. By setting QoS policies, system throughput, response time, reliability, footprint, and network bandwidth consumption can be tuned to meet application requirements. Previously, a system was hard-coded with parameters set for a specific operation profile during integration. In contrast, *Database Integration Service* provides run-time configurable policy settings, which greatly enhances system deployment flexibility.

2.2.4 Matching Real-Time Performance

Database Integration Service integration with the Oracle TimesTen In-Memory Database allows the user to capture data into standard relational databases at rates far exceeding most application requirements. Data in Oracle TimesTen can be synchronized in the background to a permanent Oracle database using the TimesTen option Cache Connect to Oracle.

Thus the in-memory database can act as a fast data cache in front of terabytes of data storage and tremendous analysis capability of a disk-based database. *Database Integration Service* solves the “impedance mismatch” between real-time and enterprise applications.

2.2.5 High Availability

Availability is an essential requirement for most distributed real-time applications. Systems built in the Defense and Aerospace industries are typically safety critical and are required to operate in crisis situations. In telecommunications, a minute of system downtime may mean many millions of dollars in lost revenue. With *Database Integration Service*, automatic data caching and replication can serve as the foundation technology for high-availability. Applications can use *Database Integration Service* to maintain copies of SQL database tables on two or more hosts in the network. In the event of a host failure, copies of the tables are available from other hosts to continue operation.

Database Integration Service's automated replication management and no-single-point-of-failure guarantees the availability of critical information. With *Database Integration Service*, tables can be stored on multiple hosts, allowing applications and services to concurrently read and write in multiple tables. Conflict resolution can be based on application-defined timestamps.

2.2.6 Additional Benefits of Database Integration Service

- ❑ Achieve quick time-to-market
 - Start application development immediately using well-known interfaces.
 - Minimize time-consuming custom programming.
 - Easily integrate into existing solutions using industry- standard interfaces.
- ❑ Reduce development costs
 - Use widely available modeling and database tools.
 - Eliminate expensive complex coding for real-time data management and communication.
 - Integrate edge devices, distributed real-time data management, and enterprise databases using a single set of standard Application Programming Interfaces.

- ❑ Deliver cutting-edge solutions
 - Process massive amounts of information across networks in real-time.
 - Turn near-instantaneous responses to (remote) critical events into a business advantage.
 - Seamlessly integrate networked applications, services, and devices.
- ❑ Minimize operational costs
 - Maintain complex networked applications with near-zero administration.
 - Dynamically add or change system components.
 - Run on common hardware platforms and networks.
- ❑ Reduce risks
 - Guarantee continuous system availability through dynamic replication management.
 - Rely on continuous high-quality technical support.
 - Build on years of experience in the world's most demanding real-time application domains.

Chapter 3 Architecture

This chapter presents a more detailed view of *RTI Database Integration Service*'s architecture and highlights the different ways that *RTI Database Integration Service* can be used to integrate systems. It includes the following sections:

- ❑ [Database Integration Service Architecture \(Section 3.1\)](#)
- ❑ [Capturing Real-Time Data in a DBMS \(Section 3.2\)](#)
- ❑ [Remote Real-Time Notification of Table Changes \(Section 3.3\)](#)
- ❑ [Bidirectional Integration \(Section 3.4\)](#)
- ❑ [Bridging between Domains \(Section 3.5\)](#)
- ❑ [High-Rate Data Streams Cached before Storage \(Section 3.6\)](#)
- ❑ [Real-Time Database Replication \(Section 3.7\)](#)

3.1 Database Integration Service Architecture

The *Database Integration Service* architecture is designed to integrate existing systems that use the *Connex* DDS API or relational databases with minimal modification to working applications. In many situations, existing applications do not have to change at all.


As seen in [Figure 3.1](#), *Database Integration Service* consists of a daemon that acts like bridge between two software development domains. One uses the OMG Data Distribution Service API to publish and subscribe to data that may be generated at high rates with real-time constraints. The other applies algorithms and logic representing business processes to megabytes, gigabytes or terabytes of data stored in relational databases.

3.1.1 Database Integration Service Daemon

The *Database Integration Service* Daemon oversees the incoming (subscribed) data and outgoing (published) data. It enables automatic storage of data published by *Connex* DDS applications in a database by mapping a Topic to a table in the database and storing an instance of a Topic as a row in that table. Also, the daemon can automatically publish changes in a database table as a Topic. Users have total control of the Quality of Service that the daemon uses for publishing and subscribing to *Connex* DDS data.

The *Database Integration Service* Daemon uses the *Connex* DDS API, as well as SQL through the ODBC API. In addition, there is a custom interface for each supported database management system (DBMS). The three currently supported DBMSs are Oracle 11g, Oracle TimesTen In-

Figure 3.1 Database Integration Service Architecture



Memory Database 11.2.1, and MySQL 5.1. There is a separate daemon executable for each of the specific DBMSs.

As illustrated in [Figure 3.1](#), Oracle TimesTen In-Memory Database and Oracle Database 11g can be used in combination. In this case, TimesTen acts as a front-end cache to Oracle Database, providing high-performance access to real-time data.

3.1.2 Database Integration Service's Unique Features

Database Integration Service offers a unique set of features that enable seamless integration of real-time/embedded *Connext DDS* applications and enterprise services:

❑ Storage of Connext DDS Data in a DBMS

Database Integration Service automatically stores received values of specified Topics in a database. Once the data is propagated to the database, it can be accessed by a user application via regular SQL queries.

❑ Publication of DBMS Data via Connext DDS

Database Integration Service automatically publishes changes in specified database tables. Changes made via the SQL API (with the INSERT, UPDATE and DELETE statements) will be published into the network via *Connext DDS*, so real-time/embedded applications and devices can respond to time-critical changes with near-zero latency.

❑ Mapping Between IDL to SQL Data Types

Database Integration Service provides automatic mapping between an IDL data type representation and a SQL table schema representation. This mapping is used to directly translate a table record to a *Connext DDS* data structure and vice-versa. Previously, this translation had to be done by custom-developed code.

❑ History

Database Integration Service can store a history of received values of a data instance. Normally, an instance of a topic is mapped to a single row in the associated database with the IDL key used as the primary key for the table. But when *Database Integration Service's* data history feature is enabled, multiple samples of a topic instance can be stored across multiple rows in the same table of the database, supporting both real-time and off-line analysis based on historical data.


❑ Configurable QoS

Database Integration Service exposes many of the QoS attributes defined by the DDS standard. This gives the user full control over the quality of service when capturing real-time data or subscribing to changes in the database. Supported QoS attributes include reliability, durability, multicasting, delivery ordering, and many others.

3.2 Capturing Real-Time Data in a DBMS

Figure 3.2 shows how *Database Integration Service* can be used to capture real-time data streams generated by embedded *Connex* DDS applications into one or more tables in a [in-memory] DBMS. In this scenario, the *Database Integration Service* Daemon has been configured with user-customizable QoSs to subscribe to Topics. When new values arrive, the daemon stores the data in the appropriate table in the database. Mapping the Topic described by IDL to the equivalent SQL table schema is done automatically by the daemon with no user configuration necessary.


Figure 3.2 Storing Published Connex DDS Data in a SQL Database



3.3 Remote Real-Time Notification of Table Changes


Figure 3.3 shows how *Database Integration Service* can be used to notify remote *Connex* DDS applications running in embedded devices of time-critical changes in the database. In this scenario, the *Database Integration Service* Daemon has been configured with user-customizable QoSs to publish Topics whenever the specified table changes in the database. Mapping the SQL table schema to the equivalent Topic described by IDL is done automatically by the daemon—no user configuration necessary.

Figure 3.3 Notifying Remote Applications of Changes in a Database



3.4 Bidirectional Integration

Figure 3.4 shows a system that integrates the capabilities described in the last two sections. *Database Integration Service* provides bidirectional dataflow between embedded *Connex* DDS applications and enterprise database systems. This approach can typically be used to create a closed-loop system, where sensory data is collected, processed, and analyzed in an in-memory database, and the resulting analysis creates state changes that are fed back to remote sensors and devices to control their behavior and mode of operation.


Figure 3.4 **Bidirectional Integration**

3.5 Bridging between Domains

Figure 3.5 shows how *Database Integration Service* can be used as a bridge between two domains by configuring the *Database Integration Service Daemon* to subscribe to data in one domain and publishing the same data in a different domain. Data sent by *Connex DDS* applications in the first domain are stored by the daemon in a local in-memory table. Since changes in the table are sent by the daemon into a second domain, the data is ultimately received by *Connex DDS* applications in the second domain. There is no feedback cancellation needed since the data is being bridged across domains. Usually domain bridges have to be written by users and modified whenever data types or Topics change. Using *Database Integration Service*, no programming is required to create a high performance bridge for any topic of any data type between any domains.

3.6 High-Rate Data Streams Cached before Storage

While disk-based persistent databases can store terabytes of data, the performance of such DBMSs is usually too low to capture data of real-time applications streaming at ultra-high rates of tens of thousands to over a million samples per second. Figure 3.6 shows how *Database Integration Service* can use the Oracle TimesTen In-Memory database as a front-end cache to the persistent Oracle database, with Cache Connect to Oracle transferring table data from memory to disk in the background.


Figure 3.5 **Bridging Domains**

This solution enables the archival of high throughput *Connex DDS* data streams that would otherwise be uncaptureable by standard database technologies.

3.7 Real-Time Database Replication


By running multiple *Database Integration Service* daemons on different nodes connected to different databases, and configuring all of the daemons to publish and subscribe to the same table, changes made by applications to a table on one node can be automatically replicated to tables on all of the other nodes. Figure 3.7 shows how *Database Integration Service* can be used to perform lazy table replication between distributed databases.

Figure 3.6 High-rate Data Streams are Cached Before Storing onto Disk



With lazy replication, an update is sent to the subscribers after the transaction is committed into the local database. The advantages of lazy replication are short response time and high concurrency, since locks in the data cache are immediately released after a transaction commits and before it is sent on the network.

Figure 3.7 Replicating Tables Across Databases



If you need remote table initialization and application timestamp-based conflict resolution, you can enable this either by using the *Database Integration Service* configuration properties or by setting individual QoS values. This is described further in [Configuration File \(Section 4.4\)](#).

Even different DBMSs can be synchronized by the *Database Integration Service* Daemon with table changes in Oracle TimesTen In-Memory databases propagated to corresponding tables in Oracle 11g or MySQL 5.1 databases, and vice versa.

Chapter 4 Using Database Integration Service

This chapter provides detailed information on using the *Database Integration Service* Daemon to subscribe to and store data received as Topics into relational databases, as well as to publish as Topic changes in relational database tables.

The contents of this chapter assume you have a working knowledge of *Connexit DDS* and relational databases, especially the MySQL, Oracle and/or Oracle TimesTen In-Memory databases. The chapter also assumes familiarity with IDL (Interface Definition Language), the DDS and SQL specifications and APIs. Finally, you should be able to create and run applications using *Connexit DDS* to publish and subscribe to data, as well as applications that can access Oracle databases using SQL through either ODBC or JDBC interfaces.

Users can configure the *Database Integration Service* Daemon to subscribe to Topics and store received values in a table, or to publish database changes as Topics using a combination of methods:

- ❑ Command-line parameters
- ❑ Environment variables
- ❑ Configuration file
- ❑ Configuration tables in the database

This chapter includes the following sections:

- ❑ [Introduction to the Database Integration Service Daemon \(Section 4.1\)](#)
- ❑ [Command-Line Parameters \(Section 4.2\)](#)
- ❑ [Environment Variables \(Section 4.3\)](#)
- ❑ [Configuration File \(Section 4.4\)](#)
- ❑ [Meta-Tables \(Section 4.5\)](#)
- ❑ [User-Table Creation \(Section 4.6\)](#)
- ❑ [Support for Extensible Types \(Section 4.7\)](#)
- ❑ [Enabling RTI Distributed Logger in Database Integration Service \(Section 4.8\)](#)
- ❑ [Enabling RTI Monitoring Library in Database Integration Service \(Section 4.9\)](#)

4.1 Introduction to the Database Integration Service Daemon

Database Integration Service bridges the world of *Connext DDS* and the world of relational databases. The main element of the bridge is an executable that must run on the same host as the database management system (DBMS). This executable is called the *Database Integration Service Daemon*.

Database Integration Service uses *Connext DDS* and supports three databases: Oracle Database 11g, Oracle TimesTen In-Memory Database 11.2.1, and MySQL 5.1. There is a separate executable that you must run depending on which database you are using.

- ❑ Oracle Database: `rtirtc_oracle.exe`
- ❑ Oracle TimesTen In-Memory Database: `rtirtc_timesten.exe`
- ❑ MySQL: `rtirtc_mysql.exe`

These executables can be executed as foreground processes during development or as background processes or as a service on Windows systems. You can configure the general behavior of the *Database Integration Service Daemon* by using command-line parameters, environment variables and configuration files. Meta-tables in the database are used to configure the specific topics and tables that are bridged by the daemon.

Besides using compatible versions of *Connext DDS* and Oracle/Oracle TimesTen/MySQL databases (see the [Release Notes](#) for a list of compatible versions), the *Database Integration Service Daemon* expects that *typecodes* for the IDL types used by *Connext DDS* applications have been generated and are being propagated. If typecodes for IDL types were not generated, you must create the tables (used by the daemon for storing or publishing data) yourself or declare the types in the configuration files.

4.1.1 How to Run the Database Integration Service Daemon with Oracle

To run *Database Integration Service* correctly with an Oracle database, you must complete the procedures described in this section.

These procedures are not needed when using *Database Integration Service* with the Oracle TimesTen In-Memory database.

To work with an Oracle database, there is a shared library distributed with *Database Integration Service* that must be installed correctly on the host of the Oracle database server. Communication by the *Database Integration Service Daemon* with the Oracle server is accomplished through external procedures executed by the server when triggers installed by the daemon are fired. These external procedures are provided in the shared library (on UNIX-based systems) or DLL (on Windows systems) called `[lib]rtirtc_oracleq[.so,.dll]`.

This library is distributed with *Database Integration Service* and can be found in the `lib/<architecture>` directory of the installation directory. The correct version of the library to use depends on the platform on which Oracle server is running. For example, `<architecture>` can be

- ❑ `x64Linux2.6cc4.1.1` for Red Hat Enterprise Linux 5 systems on 64-bit x86 processors
- ❑ `i86Linux2.6cc4.1.1` for Red Hat Enterprise Linux 5 systems on 32-bit x86 processors
- ❑ `i86Win32` for Windows systems on 32-bit x86 processors

Since the library, `[lib]rtirtc_oracleq[.so,.dll]`, internally uses *Connext DDS*, the corresponding shared libraries, `[lib]nddsc[.so,.dll]` and `[lib]nddscore[.so,.dll]`, distributed with *Connext DDS* must also be installed on the Oracle server host.

Notes:

- ❑ The exact platforms that are supported on Windows and Linux systems may be different for Oracle versus Oracle TimesTen In-Memory databases. Please consult the [Release Notes](#) for specific details of the supported platforms for your release of *Database Integration Service*.
- ❑ Not only do the libraries have to be present on the Oracle server host, the Oracle server must also be configured to find the libraries. There are separate procedures for the `librtirtc_oracleq` and `libnndsxxx` libraries. These procedures are detailed below.

4.1.1.1 Installing and Configuring the Oracle Server to Access (lib)rtirti_oracleq(.so,.dll)

There are two options for installing `[[lib]rtirti_oracleq[.so,.dll]]`:

1. Copy the appropriate version of the library into either `$ORACLE_HOME/bin` or `$ORACLE_HOME/lib` on the server host. `$ORACLE_HOME` is the installation directory of the Oracle DBMS.
- or
2. Copy the appropriate version of the library into any directory on the server host. It can even be used directly from the *Database Integration Service* installation directory if that directory can be accessed by the Oracle server.

With the second option, the location of the `[[lib]rtirti_oracleq[.so,.dll]]` library must be defined in the file `extproc.ora` (located at `$ORACLE_HOME/hs/admin` on UNIX operating systems and at `ORACLE_HOME\hs\admin` on Windows) or in the file `listener.ora` (located at `$ORACLE_HOME/network/admin`) using the `ENVS` parameter.

Additional information on how to load external procedures can be found in the Oracle manual by following this URL: http://download.oracle.com/docs/cd/E11882_01/appdev.112/e10471/adfns_externproc.htm

Important: With either option, if `[[lib]rtirti_oracleq[.so,.dll]]` is not located in `$ORACLE_HOME/bin`, the `rtirtc_oracle` daemon executable must be started with the additional command-line option `"-queuelibpath <directory containing [[lib]rtirti_oracleq[.so,.dll]]>"`— see [Command-Line Parameters \(Section 4.2\)](#).

4.1.1.2 Installing (lib)nndsc(.so,.dll) and (lib)nndscore(.so,.dll) on the Oracle Server

The shared library, `[[lib]rtirti_oracleq[.so,.dll]]`, installed in the previous section will need access to additional shared libraries provided by *Connex* DDS.

The libraries `[[lib]nndsc[.so,.dll]]` and `[[lib]nndscore[.so,.dll]]` should be copied to the server host from the appropriate `lib/<platform>` directory in the installation of *Connex* DDS. Then follow the procedure below to add these files to the library search path for the Oracle server.

UNIX-based Systems

The directory containing the *Connex* DDS libraries should be added to the environment variable `LD_LIBRARY_PATH`. This environment variable must be set in the environment of the user who started the Oracle server.

A better method for setting this environment variable is in the `extproc.ora` file or with the `ENVS` parameter in the file `listener.ora`.

Refer to the Oracle Net manual and this link for more information on the `listener.ora` file: http://download.oracle.com/docs/cd/E11882_01/network.112/e10835/listener.htm.

Refer to this link for more information on `extproc.ora`: http://download.oracle.com/docs/cd/E11882_01/appdev.112/e10471/adfns_externproc.htm.

Windows Systems

Using the dialog opened with **Start, Settings, Control Panel, System, Advanced tab, Environment Variables** button, add the directory (with backslash ‘\’ and semicolon separators ‘;’) containing the *Connext DDS* libraries to the **System** variable “**Path**”. You will need to reboot the computer for this change to take effect.

A better method for setting this environment variable (it only requires restarting the Oracle database service and the Oracle listener service) is using the **extproc.ora** file or the **ENVS** parameter in the file **listener.ora**.

Refer to the Oracle Net manual and this link for more information on the **listener.ora** file: http://download.oracle.com/docs/cd/E11882_01/network.112/e10835/listener.htm.

Refer to this link for more information on **extproc.ora**: http://download.oracle.com/docs/cd/E11882_01/appdev.112/e10471/adfnsexternproc.htm.

4.1.2 How to Run the Database Integration Service Daemon with MySQL

Before *Database Integration Service* will run correctly with a MySQL database, the procedures described in this section must be completed.

4.1.2.1 Installing MySQL ODBC 5.1.6 driver

The *Database Integration Service* Daemon requires the installation of the *MySQL ODBC 5.1.6* driver (or higher). The driver is not bundled with the MySQL server and must be installed separately.

The ODBC connector can be downloaded from <http://dev.mysql.com/downloads/connector/odbc/5.1.html>.

The installation guide can be found at <http://dev.mysql.com/doc/refman/5.1/en/connector-odbc-installation.html>.

The MySQL ODBC driver requires an ODBC driver manager. In Windows, the ODBC driver manager is automatically installed with the OS. For Linux systems we recommend the use of *UnixODBC 2.2.12* (or higher), a complete, free/open ODBC solution for Unix and Linux systems. The driver manager can be downloaded from <http://www.unixodbc.org>.

4.1.2.2 Installing and Configuring the MySQL Server to Access ([lib]rtirc_mysqlq[.so,.dll])

To work with a MySQL database, there is a shared library distributed with *Database Integration Service* that must be installed correctly on the host of the MySQL database server. Communication by the *Database Integration Service* Daemon with the MySQL server is accomplished through user-defined functions (UDF) executed by the MySQL server when triggers installed by the *Database Integration Service* Daemon are fired. These functions are provided in a shared library (on UNIX-based systems) or DLL (on Windows systems) called **[lib]rtirc_mysqlq[.so,.dll]**.

This library is distributed with *Database Integration Service* and can be found in the **lib/<platform>** directory of the installation directory. The correct version of the library to use depends on the platform on which MySQL server is running. For example, **<platform>** can be:

- ❑ **x64Linux2.6cc4.1.1** for Red Hat Enterprise Linux 5 systems on 64-bit x86 processors
- ❑ **i86Linux2.6cc4.1.1** for Red Hat Enterprise Linux 5 systems on 32-bit x86 processors
- ❑ **i86Win32** for Windows systems on 32-bit x86 processors

To install **[lib]rtirc_mysqlq[.so,.dll]** copy the appropriate version of **[lib]rtirc_mysqlq[.so,.dll]** into the MySQL server’s plugin directory (the directory named by the **plugin_dir** system vari-

able). The plugin directory can be changed by setting the value of **plugin_dir** when the MySQL server is started. For example, you can set its value in the **my.cnf** configuration file:

```
[mysqld]
plugin_dir=/path/to/plugin/directory
```

For additional information about the plugin directory see the following link:
<http://dev.mysql.com/doc/refman/5.1/en/install-plugin.html>

4.1.2.3 Installing libnddsc(.so,.dll) and libnddscore(.so,.dll) on the MySQL Server

Since the library **librtirtc_mysqlq[.so,.dll]** internally uses *Connex* DDS, the corresponding shared libraries [**lib**]nddsc[.so,.dll] and [**lib**]nddscore[.so,.dll] distributed with *Connex* DDS also need to be installed on the MySQL server host.

Note: Please consult the [Release Notes](#) for specific details of the supported platforms for your release of *Database Integration Service* to MySQL.

The libraries **nddsc** and **nddscore** must be located in a directory that is searched by the system dynamic linker.

UNIX-based Systems:

- ❑ Copy the shared libraries to a directory such as **/usr/lib**.
- ❑ Alternatively, add the libraries to the environment variable **LD_LIBRARY_PATH** that must be set for the user who starts the MySQL server. This method requires restarting the MySQL server.

Windows Systems:

- ❑ Copy the .dll files to the system directory (**WINDOWS\System32** or **WINDOWS\System**).
- ❑ Alternatively, you can add the directories containing the libraries to the **System** variable **Path** as follows:

Using the dialog opened with **Start, Settings, Control Panel, System, Advanced tab, Environment Variables** button, add the directories (with backslash '\ ' and semicolon separators ';') containing the libraries to the **System** variable "Path". If the MySQL server is running as a service, you will need to reboot the computer for this change to take effect.

4.1.2.4 Starting the MySQL Server in ANSI_QUOTES mode

The MySQL server can operate in different sql modes. The *Database Integration Service* Daemon requires the MySQL server to be configured in ANSI_QUOTES mode. Under that configuration, the MySQL server treats "" as an identifier quote character instead of a string quote character.

To verify if the MySQL server is already configured in ANSI_QUOTES mode, run the following SQL statement:

```
SELECT @@global.sql_mode;
```

If the string 'ANSI_QUOTES' is not part of the result, the MySQL server needs to be configured in ANSI_QUOTES mode using the option **--sql_mode='ANSI_QUOTES'** to start the server

That same effect can be achieved at runtime by executing the following SQL statement:

```
SET GLOBAL sql_mode = 'ANSI_QUOTES'
```

Note: The specific configuration of the MySQL server may require the use of additional SQL mode strings when starting the server.

4.1.3 How to Run the Database Integration Service Daemons as Windows Services

On Windows, the *Database Integration Service* Daemons, `rtirtc_oracle.exe`, `rtirtc_timesten.exe` and `rtirtc_mysql.exe`, can be run as system services. During the installation process, you may choose to install these daemons as Windows services, which can then be controlled through the **Start, Programs, Administrative Tools, Services** application. The *Database Integration Service* services will be installed in manual mode. Use the Services application to change this automatic to have the services start when the Windows machine boots up.

The configuration file used by the Windows services is the default file, `<Database Integration Service installation directory>/resource/xml/RTI_REAL_TIME_CONNECT.xml`.

You can change the location of the configuration file by running the Windows service with the command line option, `-cfgFile` (see [Section 4.2](#)).

4.1.4 Typecodes

Typecodes are runtime parsible descriptions of data, generated for user data types from an IDL file by the *Connex* DDS utility `rtiddsgen`. Typecodes are automatically propagated during the discovery process of *Connex* DDS applications. Unless the user has specifically disabled `rtiddsgen` from generating typecodes, applications built with types generated by `rtiddsgen` should be propagating typecodes for all of the Topics that they use, and thus are compatible with *Database Integration Service*. Please consult *Connex* DDS documentation for more information about typecodes and their generation.

An important note is that typecodes can become quite large as the corresponding IDL type becomes more complex. By default, *Connex* DDS applications allocate 2048 bytes to store a typecode. The default size for the *Database Integration Service* Daemon is 2048 bytes as well. The typecode size is controlled by the QoS parameter, `DomainParticipantQos::resource_limits.type_code_max_serialized_length`, in the *Connex* DDS API. In *Database Integration Service*, you can change the typecode limit using XML QoS Profiles (see [Table 4.2](#)).

If the *Database Integration Service* Daemon discovers Topics that have typecodes that (a) are larger than what it has been configured to handle or (b) have no associated typecodes at all, the daemon will not be able to subscribe to or publish those topics unless the user manually creates the corresponding tables in the database or defines the topic types in the configuration file (see [Table 4.2](#)). The only way to determine whether or not this situation exists is to examine the log messages printed by the daemon.

A status message will indicate when there is no typecode found for a Topic. This message may have been generated because the typecode associated with the topic is too large for the daemon. By increasing the `DomainParticipantQos::resource_limits.type_code_max_serialized_length` QoS policy, the daemon can be configured to handle larger typecodes for complex IDL types.

The *Database Integration Service* Daemon will store all the typecodes that it receives with discovered Topics. These typecodes may be used by the daemon to create user-accessible tables in the database from which changes are published or data received via *Connex* DDS is stored. See [Publications Table \(Section 4.5.1\)](#) and [Subscriptions Table \(Section 4.5.2\)](#) for more information of how typecodes are used by the daemon.

4.2 Command-Line Parameters

Any user can start a *Database Integration Service* Daemon. The user name/ID and password with which it connects to a database is specified in the configuration file, see [Table 4.9](#).

When starting a *Database Integration Service* Daemon, the following command-line parameters are supported; the **-cfgName** parameter is required.

```
Usage: rtirtc_mysql [options]
Options:
  -cfgFile      <file> Configuration file. This parameter is optional
                    since the configuration can be loaded from
                    other locations
  -cfgName      <name> Configuration name. This parameter is required
                    and it is used to find a <real_time_connect>
                    matching tag in the configuration files
  -appName      <name> Application name
                    Used to name the domain participants
                    Default: -cfgName
  -logFile      <file> Log file
  -noDaemon     Run as a regular process. Messages are sent to
                    stdout and stderr
  -use42eAlignment Enables compatibility with
                    RTI Data Distribution Service 4.2e
                    This option should be used when compatibility
                    with 4.2e is required and the topic data types
                    contain double, long long, unsigned long long,
                    or long double members
  -queueLibPath <path> Location of the library rtirtc_oracleq
                    Default: $ORACLE_HOME/bin
  -queueDomainId <int> Domain ID of the channel connecting the MySQL
                    server with RTI Database Integration Service
                    Default: 1
  -dbTransport <1|2> Transport used to communicate the MySQL server
                    with RTI Database Integration Service
                    * 1: UDPv4
                    * 2: Shared Memory
                    Default: 2 (Shared memory)
  -typeMode     <0|1> Type mode
                    Specifies whether the names and semantics of
                    the data types follow Oracle or TimesTen type
                    rules
                    Default: 0 (Oracle type mode)
  -verbosity    [0-6] RTI Database Integration Service verbosity
                    * 0 - silent
                    * 1 - exceptions (Core Libraries and Service)
                    * 2 - warnings (Service)
                    * 3 - information (Service)
                    * 4 - warnings (Core Libraries and Service)
                    * 5 - tracing (Service)
                    * 6 - tracing (Core Libraries and Service)
                    Default: 1 (exceptions)
  -version      Prints the RTI Database Integration Service version
  -help         Displays this information
```

Table 4.1 Command-line Options

Option	Description
-appName <application name>	Assigns a name to the <i>Database Integration Service</i> execution. The application name is used to set the EntityNameQosPolicy of the <i>DomainParticipants</i> created by <i>Database Integration Service</i> .
-cfgFile <configuration file>	Specifies an XML configuration file for <i>Database Integration Service</i> . The parameter is optional since the <i>Database Integration Service</i> configuration can be loaded from other locations. See Section 4.4.1 for further details.
-cfgName <configuration name>	Required. Specifies the name of the configuration to load. The <i>Database Integration Service</i> Daemon will look for the first tag <real_time_connect> with that name. (See Configuration File (Section 4.4) .)
-dbTransport <1 2>	This parameter is only available for the <code>rtirtc_oracle-[-.exe]</code> for Oracle Database 11g and the <code>rtirtc_mysql-[-.exe]</code> for MySQL. By default, <i>Database Integration Service</i> uses shared memory to communicate with the MySQL and Oracle database servers. The -dbTransport parameter can be used to change the communication transport. There are two possible values: 1: UDPv4 2: Shared memory (default) Note: Shared memory communication between the <i>Database Integration Service</i> Daemon and the database servers does not work on Windows 2003, Windows Vista or Windows 7 systems when the <i>Database Integration Service</i> daemon runs with the option -nodaemon and the database server runs as a service. For this use case, communication can be enabled by using UDPv4 as the transport.
-help	Prints out a usage message listing the command-line parameters.
-logFile <log file>	Pathname of the file to be used for log messages. If specified, log messages will automatically be stored in the file.
-noDaemon	Start as a normal process. Without this option, running the <i>Database Integration Service</i> Daemon executable will start a daemon process on Linux systems, or start a service on Windows systems. As a daemon, no log messages of any kind are printed to stdout or stderr. However, by specifying this option, the daemon will start as a regular process, which can be run as a background process using the standard OS with the command-line option (" & "), and log messages will be printed to stdout and stderr.

Table 4.1 **Command-line Options**

Option	Description
-queueDomainId <domain ID>	<p>This parameter is only available for the <code>rtirtc_oracle[-.exe]</code> for Oracle Database 11g and the <code>rtirtc_mysql[-.exe]</code> for MySQL.</p> <p>The <i>Database Integration Service</i> Daemon uses <i>Connext DDS</i> to communicate with the MySQL and Oracle 11g servers. This command-line option sets the domain ID used for the connection between the daemon and the servers.</p> <p>Default: 1</p>
-queueLibPath <directory containing [lib]rtirtc_oracleq[.so,.dll]>	<p>This parameter is only available for the <code>rtirtc_oracle[-.exe]</code> for Oracle Database 11g:</p> <p>The Oracle server must find and load <code>[lib]rtirti_oracle[.so,.dll]</code> in order to connect to the <i>Database Integration Service</i> Daemon. See How to Run the Database Integration Service Daemon with Oracle (Section 4.1.1) for more information. By default, the daemon will ask the Oracle server to look in the directory <code>\$ORACLE_HOME/bin</code>.</p> <p>If <code>[lib]rtirti_oracle[.so,.dll]</code> is installed on the server in a different directory, then the <i>Database Integration Service</i> Daemon must be started with this option set to that directory.</p>
-typeMode	<p>This parameter is only available for <code>rtirtc_timesten[.exe]</code></p> <p>Specifies whether the names and semantics of the data types in the TimesTen database follow Oracle or TimesTen type rules.</p> <p>There are two possible values:</p> <ul style="list-style-type: none"> 0: Oracle type mode (default) 1: TimesTen type mode <p>For additional information about TypeMode refer to Oracle® TimesTen In-Memory Database Reference</p> <p>Note: To work with Oracle In-Memory Database Cache this option must be set to 0.</p>

Table 4.1 Command-line Options

Option	Description
-verbosity <verbosity level>	<p><i>Database Integration Service</i> verbosity level:</p> <p>0 - No verbosity</p> <p>1 - Exceptions (<i>Connex</i>t DDS and <i>Database Integration Service</i>) (default)</p> <p>2 - Warnings (<i>Database Integration Service</i>)</p> <p>3 - Information (<i>Database Integration Service</i>)</p> <p>4 - Warnings (<i>Connex</i>t DDS and <i>Database Integration Service</i>)</p> <p>5 - Tracing (<i>Database Integration Service</i>)</p> <p>6 - Tracing (<i>Connex</i>t DDS and <i>Database Integration Service</i>)</p> <p>Each verbosity level, <i>n</i>, includes all the verbosity levels smaller than <i>n</i>.</p> <p>As the <i>Database Integration Service</i> Daemon runs, it may generate log messages reflecting error conditions, warning messages or general execution status. The messages may be produced by the daemon or by <i>Connex</i>t DDS.</p> <p>The messages produced by the daemon can be redirected to three possible destinations: stdout/stderr, a file, and log tables in the databases to which it is connected. Each of these destinations may be enabled independently of each other. The first two, stdout/stderr and file, are controlled by command line parameters discussed above, and the last, log table, is controlled in the configuration of a connection, as discussed in Database Connection Options (Section 4.4.4.3).</p> <p>In this <i>Database Integration Service</i> version, the messages produced by <i>Connex</i>t DDS can be redirected only to stdout/stderr.</p>
-version	Prints the <i>Database Integration Service</i> version.

4.3 Environment Variables

Since the *Database Integration Service* Daemon will be making connections to databases using ODBC, on UNIX-based systems, the following environment variables may be used to find DSNs (data source names) via ODBCINI files.

- ❑ **ODBCINI**: location of INI file for database connections. If not set, ODBCINI will be set to “**\$HOME/.odbc.ini**”, where **\$HOME** is the home directory of the user who started the daemon.
- ❑ **SYSODBCINI**: location of system INI file, used if the DSN is not found in the file specified by **ODBCINI**.

If the *Database Integration Service* Daemon cannot find a valid DSN in any ODBC.INI file, then no connections to any databases can be made.

On a Windows system, the equivalent functionality of the ODBCINI file is found in the Windows registry. You create and modify DSNs using the application found in **Start, Programs, Administrative Tools, Data Sources (ODBC)**.

4.4 Configuration File

When you start *Database Integration Service*, you can provide a configuration file in XML format (it is not required). Among other things, this file can be used to specify the set of databases to which the daemon will connect and the properties of the database connections.

This section describes:

- ❑ [How to Load the XML Configuration \(Section 4.4.1\)](#)
- ❑ [XML Syntax and Validation \(Section 4.4.2\)](#)
- ❑ [Top-Level XML Tags \(Section 4.4.3\)](#)
- ❑ [Database Configuration Using the Database Integration Service XML Tag \(Section 4.4.4\)](#)

4.4.1 How to Load the XML Configuration

Database Integration Service loads its XML configuration from multiple locations. This section presents the various approaches, listed in load order.

The first three locations only contain QoS Profiles and are inherited from *Connex DDS* (see Chapter 15 in the *RTI Connex DDS Core Libraries User's Manual*).

<NDDSHOME> is described in [Paths Mentioned in Documentation \(Section 1.2\)](#).

- ❑ <NDDSHOME>/resource/xml/NDDS_QOS_PROFILES.xml
This file contains the *Connex DDS* default QoS values; it is loaded automatically if it exists. (*First to be loaded.*)
- ❑ File in NDDS_QOS_PROFILES
The files (or XML strings) separated by semicolons referenced in this environment variable are loaded automatically.
- ❑ <working directory>/USER_QOS_PROFILES.xml
This file is loaded automatically if it exists.

The next locations are specific to *Database Integration Service*.

- ❑ <NDDSHOME>/resource/xml/RTI_REAL_TIME_CONNECT.xml
This file contains the default *Database Integration Service* configuration and QoS Profiles; it is loaded if it exists. The default configuration does not work out-of-the-box because it requires setting the parameters that configure the database connections such as dsn, username and password (see [Section 4.4.4](#)).
- ❑ <working directory>/USER_REAL_TIME_CONNECT.xml
This file is loaded automatically if it exists.
- ❑ File specified using the command line parameter -cfgFile
The command-line option -**cfgFile** (see [Section 4.2](#)) can be used to specify a configuration file.

You may use a combination of the above approaches.

4.4.2 XML Syntax and Validation

The XML configuration file must follow these syntax rules:

- ❑ The syntax is XML; the character encoding is UTF-8.
- ❑ Opening tags are enclosed in `<>`; closing tags are enclosed in `</>`.
- ❑ A tag value is a UTF-8 encoded string. Legal values are alphanumeric characters. The routing service's parser will remove all leading and trailing spaces¹ from the string before it is processed.

For example, "`<tag> value </tag>`" is the same as "`<tag>value</tag>`".

- ❑ All values are case-sensitive unless otherwise stated.
- ❑ Comments are enclosed as follows: `<!-- comment -->`.
- ❑ The root tag of the configuration file must be `<dds>` and end with `</dds>`.

Database Integration Service provides DTD and XSD files that describe the format of the XML content. We recommend including a reference to one of these documents in the XML file that contains the *Database Integration Service*'s configuration—this provides helpful features in code editors such as Visual Studio and Eclipse, including validation and auto-completion while you are editing the XML file.

The DTD and XSD definitions of the XML elements are in `<Database Integration Service installation directory>/resource/schema/rti_real_time_connect.dtd` and `<Database Integration Service installation directory>/resource/schema/rti_real_time_connect.xsd`, respectively.

To include a reference to the XSD document in your XML file, use the attribute `xsi:noNamespaceSchemaLocation` in the `<dds>` tag. For example:

```
<?xml version="1.0" encoding="UTF-8"?>
<dds xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:noNamespaceSchemaLocation= "<NDDSHOME>/resource/schema/
  rti_real_time_connect.xsd">
  ...
</dds>
```

`<NDDSHOME>` is described in [Paths Mentioned in Documentation \(Section 1.2\)](#).

To include a reference to the DTD document in your XML file, use the `<!DOCTYPE>` tag.

For example:

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE dds SYSTEM "<NDDSHOME>/resource/schema/rti_routing_service.dtd">
<dds>
  ...
</dds>
```

We recommend including a reference to the XSD file in the XML documents; this provides stricter validation and better auto-completion than the corresponding DTD file.

1. Leading and trailing spaces in enumeration fields will not be considered valid if you use the distributed XSD document to do validation at run-time with a code editor.

4.4.3 Top-Level XML Tags

Let's look at an example configuration file. You will learn the meaning of each line as you read the rest of the sections.

```
<?xml version="1.0"?>
<dds>
  <real_time_connect name="Example">
    <database_mapping_options>
      <identifier_separator_char>$
    </identifier_separator_char>
    </database_mapping_options>

    <mysql_connection>
      <dsn>Example_MySQL</dsn>
      <user_name>Student</user_name>
      <password>mypsswr</password>
    </mysql_connection>

    <oracle_connection>
      <dsn>Example_Oracle</dsn>
      <user_name>Student</user_name>
      <password>mypsswr</password>
    </oracle_connection>

    <timesten_connection>
      <dsn>Example_TT</dsn>
    </timesten_connection>
  </real_time_connect>
</dds>
```



Table 4.2 describe the top-level tags allowed within the root <dds> tag. Notice that the <real_time_connect> tag is required.

Table 4.2 Top-Level Tags

Tags within <dds>	Description	Number of Tags Allowed
<qos_library>	<p>Specifies a QoS library and profiles.</p> <p>The contents of this tag are specified in the same manner as for an <i>Connex DDS</i> QoS profile file—see Chapter 15 in the <i>RTI Connex DDS Core Libraries User's Manual</i>.</p> <p>The profiles you specify here can be used in three ways.</p> <ul style="list-style-type: none"> <input type="checkbox"/> By setting the attribute is_default_qos in the tag <qos_profile> to true. In this case, that profile is the default configuration for all the Entities created by the <i>Database Integration Service</i> daemon. <input type="checkbox"/> By referring to a profile using the XML tag <profile_name> within <publication> and <subscription> (see Section 4.4.4.4). <input type="checkbox"/> By referring to a profile in the profile_name column of the tables RTIDDS_PUBLICATIONS or RTIDDS_SUBSCRIPTIONS (see Section 4.5.1 and Section 4.5.2). 	0 or more

Table 4.2 Top-Level Tags

Tags within <dds>	Description	Number of Tags Allowed
<real_time_connect>	<p>Specifies a <i>Database Integration Service</i> configuration.</p> <p>This tag is used to specify the set of databases to which the daemon will connect.</p> <p>Note: There is no way to dynamically configure the <i>Database Integration Service</i> Daemon Daemon to connect to a database after it has started. All database connections must be specified within this tag before the daemon starts.</p> <p>See Table 4.3 for a description of the elements contained inside <real_time_connect>.</p>	1 or more <i>(required)</i>
<types>	<p>Defines types that can be used to create database tables.</p> <p>The type description is done using the Connex DDS XML format for type definitions. For more information, see Section 3.4 in the RTI Connex DDS Core Libraries User's Manual.</p> <p>For example:</p> <pre><types> <struct name="Point"> <member name="x" type="long"/> <member name="y" type="long"/> </struct> </types></pre> <p><i>Database Integration Service</i> supports automatic table creation by using the types defined within this tag or the typecode sent by <i>Connex DDS</i> applications.</p> <p>See Section 4.6 for additional information on user table creation.</p>	0 or 1

Because a configuration file may contain multiple <real_time_connect> tags, one file can be used to configure multiple daemon executions. When you start *RTI Real-Time Connect*, you have to use the **-cfgName** option to specify which <real_time_connect> tag to use.

For example:

```
<dds>
  ...
  <real_time_connect name="rtcA">
    ...
  </ real_time_connect >

  <real_time_connect name="rtcB">
    ...
  </real_time_connect>
  ...
</dds>
```

Starting Database Integration Service with the following command will use the <real_time_connect> tag with the name rtcA:

```
rtirtc_mysql -cfgFile example.xml -cfgName rtcA
```

If there is no <real_time_connect> tag matching the name provided with the command line option **-cfgName**, the daemon will report an error and it will list the available configurations.

4.4.4 Database Configuration Using the Database Integration Service XML Tag

Table 4.3 describes the tags allowed with the <real_time_connect> section of the XML file.

Table 4.3 Database Integration Service Tags

Tags within <real_time_connect>	Description	Number of Tags Allowed
<database_mapping_options>	Configures how the IDL identifier names are mapped to the database column names. See Section 4.4.4.2.	0 or 1
<general_options>	Contains attributes that are independent of any particular database connection made by the <i>Database Integration Service</i> Daemon. See Section 4.4.4.1.	0 or 1
<mysql_connection>	Configures a connection to a MySQL database. See Section 4.4.4.3.	1 or more (<i>required</i>) if running rtirtc_mysql ; 0 or more (ignored) if running other DBMS version of the daemon
<oracle_connection>	Configures a connection to an Oracle database. See Section 4.4.4.3.	1 or more (<i>required</i>) if running rtirtc_oracle ; 0 or more (ignored) if running other DBMS version of the daemon
<timesten_connection>	Configures a connection to an Oracle TimesTen database. See Section 4.4.4.3.	1 or more (<i>required</i>) if running rtirtc_timesten ; 0 or more (ignored) if running other DBMS version of the daemon

4.4.4.1 General Options

Table 4.4 describes the general options; these attributes are independent of any particular database connection made by the *Database Integration Service* Daemon.

Table 4.4 General Options Tags

Tags within <general_options >	Description	Number of Tags Allowed
<administration>	See Table 4.5, "Administration Tags"	0 or 1
<enable_table_replication>	<i>Database Integration Service</i> can be configured to perform real-time, lazy database replication (see Section 3.7) by setting this attribute to true. Default: false	0 or 1
<max_objects_per_thread>	This parameter controls the maximum number of objects per thread that <i>Connext DDS</i> can store. If you run into problems related to the creation of Entities, increasing this number may be necessary. See the <i>RTI Connext DDS Core Libraries User's Manual</i> for more information. Default: <i>Connext DDS</i> default (1024)	0 or 1

Table 4.4 General Options Tags

Tags within <general_options >	Description	Number of Tags Allowed
<typecode_from_table_schema>	<p>This tag can be used to enable typecode generation from table schemas.</p> <p>If this parameter is set to true and a publication or subscription is created for a database table without an associated typecode, <i>Database Integration Service</i> will create the typecode from the table schema.</p> <p>The new typecode will be made available to other <i>Connex DDS</i> applications or RTC daemons via discovery traffic.</p> <p>When <i>Database Integration Service</i> is used for table replication, the default value for this parameter is true allowing automatic table creation in the replicas.</p> <p>Default: false (except when <code>enable_table_replication</code> is set to true).</p>	0 or 1

Table 4.5 Administration Tags

Tags within <administration>	Description	Number of Tags Allowed
<distributed_logger>	<p>Configures <i>RTI Distributed Logger</i>.</p> <p>See Enabling RTI Distributed Logger in Database Integration Service (Section 4.8)</p>	0 or 1
<domain_id>	Specifies which domain ID <i>Database Integration Service</i> will use to send log messages when <i>Distributed Logger</i> is enabled.	1 (required)
<profile_name>	<p>Determines which QoS profile to use when creating the DomainParticipant that will be used for <i>Distributed Logger</i> (when enabled). The value is the fully qualified name of the QoS Profile, represented as a string with this form:</p> <p><QoS profile library name>::<QoS profile name></p>	0 or 1

4.4.4.1.1 Enabling Table Replication

Enabling database replication will automatically configure the QoS values of publications and subscriptions to provide conflict resolution and table initialization (see [Table 4.6](#) and [Table 4.7](#)). The attribute also enables automatic table creation (see `<typecode_from_table_schema>` in [Table 4.4](#)) and propagation of NULL values.

Table 4.6 DataWriter QoS Changes when `<enable_table_replication>` is True

QoS Change	Purpose
<code>reliability.kind = RELIABLE_RELIABILITY_QOS</code>	Enables reliability
<code>destination_order.kind =</code> <code>BY_SOURCE_TIMESTAMP_DESTINATIONORDER_QOS</code> <code>destination_order.source_timestamp_tolerance.sec = 0</code> <code>destination_order.source_timestamp_tolerance.nanosec = 0</code> <code>ownership.kind = SHARED_OWNERSHIP_QOS</code>	Performs conflict resolution
<code>protocol.serialize_key_with_dispose = true</code> <code>writer_data_lifecycle.autodispose_unregistered_instances = false</code>	Propagates delete operations
<code>durability.kind = TRANSIENT_LOCAL_DURABILITY_QOS</code>	Sends table contents to late joiners (table initialization)

Table 4.6 **DataWriter QoS Changes when <enable_table_replication> is True**

QoS Change	Purpose
history.depth = 1 history.kind = KEEP_LAST_HISTORY_QOS	Keeps one record per primary key value
writer_resource_limits.instance_replacement = DDS_DISPOSED_INSTANCE_REPLACEMENT writer_resource_limits.replace_empty_instances = DDS_BOOLEAN_FALSE	Enables replacement of deleted rows

Table 4.7 **DataReader QoS Changes when <enable_table_replication> is True**

QoS Change	Purpose
reliability.kind = RELIABLE_RELIABILITY_QOS	Enables reliability
destination_order.source_timestamp_tolerance.sec = DURATION_INFINITE_SEC destination_order.source_timestamp_tolerance.nanosec = DURATION_INFINITE_NSEC ownership.kind = SHARED_OWNERSHIP_QOS;	Performs conflict resolution Note: <enable_table_replication> sets some QoS related to conflict resolution, but it does not enable the feature. See Section 4.4.4.1.2 for additional details
protocol.propagate_dispose_of_unregistered_instances = true	Enables propagation of delete operations
durability.kind = TRANSIENT_LOCAL_DURABILITY_QOS;	Sends table contents to late joiners (table initialization)
history.kind = KEEP_LAST_HISTORY_QOS;	Keeps one sample per primary key value

These QoS changes have priority over the values set using QoS Profiles. However, they can be overwritten per publication and per subscription by setting the corresponding fields in the RTIDDS_PUBLICATIONS and RTIDDS_SUBSCRIPTIONS meta tables (see [Section 4.5.1](#) and [Section 4.5.2](#)).

4.4.4.1.2 Conflict Resolution

Because there are no global (network-wide) locks on records when a transaction is being executed, conflicts can occur. The best way to avoid conflicts is to have only one host modify a specific row (instance) or table (topic), but that is not always possible. The second best way is to design the application in such a way that conflicts can never occur, for instance due to data flow dependencies. But that also is often hard to achieve.

By default, conflict resolution is not enabled when you set <enable_table_replication> to true. You can enable conflict resolution by setting the column `dr.destination_order.kind` in **RTIDDS_SUBSCRIPTIONS** to **BY_SOURCE_TIMESTAMP_DESTINATIONORDER_QOS** (see [Section 4.5.2.1.40](#)). With this setting, eventual consistency can be guaranteed. Conflicts can cause a temporary inconsistency between the databases, but eventually these are resolved by the *Database Integration Service* conflict-resolution mechanism. By default, conflicts are resolved using a timestamp corresponding to the system time when the update occurred. You can overwrite this behavior by providing your own timestamp in a separate database column (see [Section 4.5.1.1.9](#)).

If you do not need conflict resolution, you can disable it by setting the column `dr.destination_order.kind` in **RTIDDS_SUBSCRIPTIONS** to **BY_RECEPTION_TIMESTAMP_DESTINATIONORDER_QOS**.

4.4.4.1.3 Table Initialization

When a host starts using a table in the distributed shared database, it is essential that the local table is up-to-date. *Database Integration Service* supports two approaches to filling the local table's contents:

1. If all the rows in the table are updated frequently, it is sufficient to apply these updates to the data cache.
2. The table can be synchronized by explicitly requesting the table's contents from the other hosts. This is called table initialization.

If table initialization is not needed, you can disable it by setting the columns **dw.durability.kind** in `RTIDDS_PUBLICATIONS` and **dw.durability.kind** in `RTIDDS_SUBSCRIPTIONS` to `VOLATILE_DURABILITY_QOS`.

4.4.4.2 Database Mapping Options

[Table 4.8](#) describes the options that are allowed with the `<database_mapping_options>` tag.

Table 4.8 Database Mapping Options

Tags within <code><database_mapping_options></code>	Description	Number of Tags Allowed
<code><close_bracket_char></code>	<p>Sets the closing bracket character that is used in the index component of the arrays and sequences members names.</p> <p>See Array Fields (Section 5.2.8) and Sequence Fields (Section 5.2.9) for more information about the mapping of IDL arrays and sequences into SQL columns.</p> <p>The default value of <code>'</code> will generate columns names that must be referenced using double quotes.</p> <p>Default: <code>'</code></p>	0 or 1
<code><identifier_separator_char></code>	<p>Sets the character that is used as a separator in the hierarchical names generated when mapping IDL fields into SQL table columns.</p> <p>The attribute is also used to configure the separator character for the columns in the meta tables.</p> <p>The default value of <code>'</code> in Oracle and Oracle TimesTen will generate columns names that must be referenced using double quotes. See IDL to SQL Mapping (Section 5.2.1) for more information about double quoted identifiers.</p> <p>Default: <code>'</code> For Oracle/TimesTen; <code>\$</code> for MySQL</p>	0 or 1

Table 4.8 Database Mapping Options

Tags within <database_mapping_ options>	Description	Number of Tags Allowed
<idl_member_prefix_max_length>	<p>Controls the prefix length of the IDL member identifiers that will be used to truncate column names when a table is automatically created.</p> <p>If the default value (-1) is used, <i>Database Integration Service</i> will not truncate IDL member identifiers when these are used to create column names.</p> <p>If a positive value, <i>n</i>, is provided, <i>Database Integration Service</i> will use the first <i>n</i> characters from the IDL member identifier to compose the associated column name.</p> <p>A value of 0 tells <i>Database Integration Service</i> to compose the column name using only the last characters of the identifiers, as defined by <idl_member_suffix_max_length>.</p> <p>This value can be overridden per table by assigning a value to the <code>idl_member_prefix_max_length</code> column in the meta-tables.</p> <p>Default: -1 (unlimited)</p>	0 or 1
<idl_member_suffix_max_length>	<p>Controls the suffix length of the IDL member identifiers that will be used to truncate column names when a table is automatically created.</p> <p>If a positive value, <i>n</i>, is provided, <i>Database Integration Service</i> will use the last <i>n</i> characters from the IDL member identifier to compose the associated column name.</p> <p>A value of 0 tells <i>Database Integration Service</i> to compose the column name using only the first characters of the identifiers, as defined by <idl_member_prefix_max_length>.</p> <p>This value can be overridden per table by assigning a value to the <code>idl_member_suffix_max_length</code> column in the meta-tables.</p> <p>Note that although <idl_member_prefix_max_length> and <idl_member_suffix_max_length> can be individually set to zero, they cannot be both zero at the same time.</p> <p>Default: -1 (unlimited)</p>	0 or 1
<open_bracket_char>	<p>Sets the opening bracket character that is used in the index component of the arrays and sequences members names.</p> <p>See Array Fields (Section 5.2.8) and Sequence Fields (Section 5.2.9) for more information about the mapping of IDL arrays and sequences into SQL columns.</p> <p>The default value of '[' will generate columns names that must be referenced using double quotes.</p> <p>Default: '['</p>	0 or 1

4.4.4.3 Database Connection Options

The database connection tags in the XML file direct the *Database Integration Service* Daemon to connect to a particular database as specified by a DSN (data source name) and configure the connection.

The database connection tags are DBMS-specific:

- ❑ `<mysql_connection>`
- ❑ `<oracle_connection>`
- ❑ `<timesten_connection>`

A `<real_time_connect>` tag may have multiple database connection tags. The DBMS-specific *Database Integration Service* Daemon will only parse the tags that apply to it. As explained earlier, the *Database Integration Service* Daemon will make a connection to a database using the DSN attribute for every connection tag that it parses. This is the only way to direct the daemon to connect to a database. No other connections will be made after startup.

Example:

```
<real_time_connect name="MyRtc">
  <mysql_connection>
    <dsn>Example_MySQL</dsn>
    <user_name>Student</user_name>
    <password>mypsswr</password>
    <send_period>100</send_period>
    <database_logging>
      <enabled>true</enabled>
      <history_depth>100</history_depth>
    </database_logging>
    <publications>
      <publication>...</publication>
    </publications>
    <subscriptions>
      <subscription>...</subscription>
    </subscriptions>
  </mysql_connection>
</real_time_connect>
```




Table 4.9 describes tags allowed within all three types of `<database_connection>` tags. Tables 4.10 through 4.12 describe additional tags for each connection type.

Table 4.9 Common Tags for all Database Connections

Common Tags for <mysql_connection>, <oracle_connection>, <timesten_connection>	Description	Number of Tags Allowed
<database_logging> <enabled> <history_depth>	<p>If enabled, the <i>Database Integration Service</i> Daemon's log messages will be stored in a table named "RTIRTC_LOG" in the database specified by the DSN.</p> <p>Optionally, you can specify the history depth of the log. This value limits the size of the table, RTIRTC_LOG, in the database that the daemon uses for logging messages. The default is 1000 rows, and a value of -1 implies no limit. When the table is filled, new log messages will replace the oldest messages, effectively using the table as a circular buffer.</p> <p>Default: disabled</p>	0 or 1
<dsn>	<p>You must specify a valid DSN that is found in a ODBCINI file or the Windows registry (see Environment Variables (Section 4.3)). The <i>Database Integration Service</i> Daemon will make a connection to this DSN.</p> <p><i>Database Integration Service</i> detects changes in an Oracle TimesTen data store by reading from the transaction log. Consequently, <i>Database Integration Service</i> will not work for DSN configurations where logging is turned off (Logging=0). This limitation applies only to Oracle TimesTen In-Memory Database, not to Oracle Database 11g or MySQL.</p>	1 (required)
<password>	Specifies the password to connect to the database.	<p>1 (required) in <mysql_connection> and <oracle_connection></p> <p>0 or 1 (optional) in <timesten_connection></p>
<publications>	This tags allows inserting publications in the table RTIRTC_PUBLICATIONS when the daemon starts up. See Initial Subscriptions and Publications (Section 4.4.4.4) .	0 or 1
<send_period>	<p>The send_period value specifies the milliseconds interval at which the <i>Database Integration Service</i> Daemon publishes database changes. The value must be greater than or equal to 0. With a value of 0 the daemon publishes database changes as soon as they are available. A shorter time interval reduces latency.</p> <p>Default: 100 ms</p>	0 or 1
<subscriptions>	This tags allows inserting subscriptions in the table RTIRTC_SUBSCRIPTIONS when the daemon starts up. See Initial Subscriptions and Publications (Section 4.4.4.4) .	0 or 1

Table 4.9 Common Tags for all Database Connections

Common Tags for <mysql_connection>, <oracle_connection>, <timesten_connection>	Description	Number of Tags Allowed
<user_name>	<p>Specifies the user name to connect to the database.</p> <p>This attribute is mandatory for Oracle and MySQL databases, since there is no default user for database connections.</p> <p>It is optional for Oracle TimesTen: if this attribute is missing, then the connection to the database will use the username of the UID who owns the <i>Database Integration Service</i> Daemon process.</p>	<p>1 <i>(required)</i> in <mysql_connection> and <oracle_connection></p> <p>0 or 1 <i>(optional)</i> in <timesten_connection></p>

Table 4.10 Tags for MySQL Connections

Additional Tags Allowed within <mysql_connection>	Description	Number of Tags Allowed
<transaction_max_duration>	<p>Provides an estimation of the maximum duration of a database transaction. If a table change is not committed in the interval specified by this attribute, it will not be published to <i>Connext DDS</i>.</p> <p>Uncommitted table changes are stored in a per-table queue. The maximum size of that queue can be configured setting the value of the changes_queue_maximum_size column in the RTIDDS_PUBLICATIONS table (see Section 4.5.1.1.21).</p> <p>If a change in the uncommitted changes queue has not been committed after transaction_max_duration milliseconds, it will be discarded by the <i>Database Integration Service</i> Daemon.</p> <p>With a value of -1, the <i>Database Integration Service</i> Daemon will not discard changes into the uncommitted queue until they are committed.</p> <p>Default: 5000</p>	0 or 1

Table 4.11 Tags for Oracle Connections

Additional Tags Allowed within <oracle_connection>	Description	Number of Tags Allowed
<blob_default_size>	<p>Limits the size of BLOB columns for tables that do not have an associated TypeCode (see Section 4.1.4) in the RTIRTC_TBL_INFO meta table (see Section 4.5.3).</p> <p>If a user table has an associated TypeCode, the maximum size of the BLOB columns is determined on a per-column basis using the TypeCode information.</p> <p>Default: 65536</p>	0 or 1
<clob_default_size>	<p>Limits the size of CLOB and NCLOB columns for tables that do not have an associated TypeCode (see Section 4.1.4) in the RTIRTC_TBL_INFO meta table (see Section 4.5.3).</p> <p>If a user table has an associated TypeCode, the maximum size of the CLOB and NCLOB columns is determined on a per-column basis using the TypeCode information.</p> <p>Default: 65536</p>	0 or 1

Table 4.12 Tags for TimesTen Connections

Additional Tags Allowed within <timesTen_connection>	Description	Number of Tags Allowed
<log_acknowledgement_period >	Sets the transaction log acknowledgement period in seconds. The transaction log will be acknowledged at the period specified by the tag value. This allows TimesTen to purge old transaction log files when no longer needed. Default: 30 seconds	0 or 1
<xla_buffer_size>	This attribute is only used if the Oracle TimesTen DSN specifies a diskless connection. See Setting the XLA Staging Buffer Size for Diskless Connections (Section 4.4.4.5) Default: 409600	0 or 1

4.4.4.4 Initial Subscriptions and Publications

As explained in [Meta-Tables \(Section 4.5\)](#), the daemon is configured to publish and subscribe to data in the database by inserting entries in two meta-tables, RTIDDS_PUBLICATIONS and RTIDDS_SUBSCRIPTIONS. In your XML configuration you can specify initial values to insert in these tables.

For example:

```

<mysql_connection>
...
  <subscriptions delete="true">
    <subscription>
      <table_owner>user</table_owner>
      <table_name>mytable1</table_name>
      <domain_id>54</domain_id>
      <topic_name>mytopic1</topic_name>
      <type_name>mytype1</type_name>
    </subscription>
    ...
    <subscription>
      ...
    </subscription>
    ...
  </subscriptions>

  <publications>
    <publication overwrite="true">
      <table_owner>user</table_owner>
      <table_name>mytable2</table_name>
      <domain_id>54</domain_id>
      <topic_name>mytopic2</topic_name>
      <type_name>mytype2</type_name>
    </publication>

    <publication>
      ...
    </publication>
    ...
  </publications>
</mysql_connection>

```

Within `<subscriptions>` and `<publications>` tags, you can specify as many `<subscription>` and `<publication>` tags as you want. The content of each tag inside `<subscription>/<publication>` represents the value for a column with the same name in the table `RTIDDS_SUBSCRIPTIONS/RTIDDS_PUBLICATIONS`. Each of these `<subscription>/<publication>` tags may result in the insertion or update of a row in the corresponding meta-table.

All the rows in the tables can be deleted before inserting new rows if the attribute “delete” in `<publications>/<subscriptions>` is set to true.

If a `<publication>` or `<subscription>` already exists in its table (the primary key is the same), then the insertion won’t succeed. However you can set the attribute “overwrite” to true. In that case, if the insertion fails, an update is performed on that row.

Table 4.13 Subscriptions Tags

Tags Allowed within <code><subscriptions></code>	Description	Number of Tags Allowed
<code><subscription></code>	Configures a subscription by inserting or updating a row in the table <code>RTIDDS_SUBSCRIPTIONS</code> . See Table 4.14 on page 4-24 .	1 or more

Table 4.14 Subscription Tags

Tags Allowed within <code><subscription></code>	Description	Number of Tags Allowed
<code><domain_id></code>	Inserts the tag value into the column with the same name in the table <code>RTIDDS_SUBSCRIPTIONS</code> See Section 4.5.2 .	1 <i>(required)</i>
<code><table_name></code>		
<code><table_owner></code>		
<code><topic_name></code>		
<code><cache_initial_size></code>	Inserts the tag value into the column with the same name in the table <code>RTIDDS_SUBSCRIPTIONS</code> If the value is not specified, NULL is inserted. See Section 4.5.2 .	0 or 1
<code><cache_maximum_size></code>		
<code><commit_type></code>		
<code><delete_on_dispose></code>		
<code><filter_duplicates></code>		
<code><idl_member_prefix_max_length></code>		
<code><idl_member_suffix_max_length></code>		
<code><ordered_store></code>		
<code><persist_state></code>		
<code><process_batch></code>		
<code><process_period></code>		
<code><profile_name></code>		
<code><table_history_depth></code>		
<code><type_name></code>		

Note that there are columns in the tables `RTIDDS_PUBLICATIONS` and `RTIDDS_SUBSCRIPTIONS` that don’t have a corresponding tag inside `<publications>` and `<subscriptions>`. Those columns represent configuration of QoS. However, you can configure the quality of service by using `<profile_name>`, where you can refer to a QoS profile in your own XML configuration file or in any of the other QoS profile files loaded by the daemon (see [How to Load the XML Configuration \(Section 4.4.1\)](#)).

Table 4.15 **Publications Tags**

Tags Allowed within <publications>	Description	Number of Tags Allowed
<publication>	Configures a publication by inserting or updating a row in the table RTIDDS_PUBLICATIONS. See Table 4.16 on page 4-25 .	1 or more

Table 4.16 **Publication Tags**

Tags Allowed within <publication>	Description	Number of Tags Allowed
<domain_id>	Inserts the tag value into the column with the same name in the table RTIDDS_PUBLICATIONS. See Section 4.5.1	1 <i>(required)</i>
<table_name>		
<table_owner>		
<topic_name>		
<idl_member_prefix_max_length>	Inserts the tag value into the column with the same name in the table RTIDDS_PUBLICATIONS. If the value is not specified, NULL is inserted. See Section 4.5.1	0 or 1
<idl_member_suffix_max_length>		
<resolution_column >		
<profile_name>		
<table_history_depth>		
<type_name>		

4.4.4.5 Setting the XLA Staging Buffer Size for Diskless Connections

For diskless connections, the minimum size of the Oracle TimesTen XLA staging buffer must be at least as big as the largest single SQL transaction executed through the database connection. The size of a transaction is related to the aggregated size of data values that have changed as a result of the transaction. For example, if a transaction modifies every column in a row of a table, then the transaction is at least as large as the size (in bytes) of a row.

Note that a single transaction may modify multiple rows of a table. It is up to the user to determine what the largest transaction size in a database may be and set the XLABUFFERSIZE attribute for the *Database Integration Service* Daemon appropriately. However, the maximum transaction size is only the minimum value that should be set for XLABUFFERSIZE. You may need to set XLABUFFERSIZE to be much larger, since the XLA staging buffer in Oracle TimesTen must hold more than a single non-committed transaction simultaneously if there are multiple threads or processes accessing the same database at the same time. Recall that this discussion only pertains to Oracle TimesTen DSNs that specify a diskless connection.

If the XLABUFFERSIZE is too small, then SQLExecute or SQLExecDirect statements that were working will return an error indicating that a buffer is full. You should see these errors in your own applications as well as in the *Database Integration Service* Daemon if the daemon is subscribing to Topics and trying to store the received data in the database. These ODBC errors will appear in the daemon log. If you encounter this situation, you should increase the value of XLABUFFERSIZE appropriately. The Oracle TimesTen native error codes that you may see when the XLABUFFERSIZE is too small are:

- ❑ TT8009: Transaction Log API Buffer size too small or too large
- ❑ TT986: Log buffer overflow; transaction must rollback
- ❑ TT987: Log record larger than log buffer; transaction must rollback

More details about Oracle TimesTen error codes can be found in the *Oracle TimesTen In-Memory Database Error Messages and SNMP Traps Release 11.2.1*.

4.5 Meta-Tables

After the *Database Integration Service* Daemon has started and successfully made a connection to a database, the user will still need to configure the daemon to publish table changes as Topics as well as subscribe to Topics for storing received data into a table. This configuration is done by inserting entries into two tables **RTIDDS_PUBLICATIONS** and **RTIDDS_SUBSCRIPTIONS**. These tables will be created by the *Database Integration Service* Daemon if they do not already exist in the database.

The two tables are referred to as *meta-tables* since their data is not user data but information used by the daemon to create DataWriters and DataReaders, as well as corresponding user tables in the database. The tables are just ordinary tables that users can create themselves before starting the *Database Integration Service* Daemon if so desired. However, if the user chooses to do so, it is important that the tables be created with the exact tables schemas presented below, otherwise the daemon may not work correctly. If the daemon finds existing meta-tables upon startup, it will process every row in the tables as if they were newly inserted. The meta-tables **RTIDDS_PUBLICATIONS** and **RTIDDS_SUBSCRIPTIONS** can be populated using the <publication> and <subscription> tags in the configuration file (see [Section 4.4.4.4](#)) or running Insert/Update SQL statements.

There are two more meta-tables created by the *Database Integration Service* Daemon:

- ❑ The meta-table **RTIRTC_TBL_INFO** will store the typecode associated with the user tables created automatically by the RTC daemon (see [Section 4.5.3](#)).
- ❑ The meta-table **RTIRTC_LOG** will be created to store log messages generated by the *Database Integration Service* Daemon. Use of this table is controlled by command-line parameters and the connection discussed in [Section 4.2](#) and [Section 4.4.4.3](#).

The following sections discuss the usage of these tables and describe the actions taken by the daemon when these tables are modified:

- ❑ [Publications Table \(Section 4.5.1\)](#)
- ❑ [Subscriptions Table \(Section 4.5.2\)](#)
- ❑ [Table Info \(Section 4.5.3\)](#)
- ❑ [Log Table \(Section 4.5.4\)](#)

4.5.1 Publications Table

When entries (rows) are added to the meta-table **RTIDDS_PUBLICATIONS**, the *Database Integration Service* Daemon will try to create a DataWriter (and Publisher along with a DomainParticipant if required) and use it to send changes to the designated user table via the *Connex* DDS.

If the **RTIDDS_PUBLICATIONS** table does not exist at startup, the *Database Integration Service* Daemon will create it with the table owner set to the user name of the database connection as specified in the daemon's configuration file, see [Section 4.4](#). The schema and meaning of the columns of this table are described in the next section.

Users may insert new rows or modify the column values of existing rows in this table at any-time. For a new row, the daemon will first check to see if the designated user table exists. If so, it

will immediately create the DataWriter with the QoS values as specified by the entry. The name of the Topic to publish may be specified by the **topic_name** column or be automatically constructed as **<table_owner>.<table_name>** if the **topic_name** entry is NULL.


If the user table does not exist, the *Database Integration Service* Daemon will look for the typecode associated with the topic defined in the **topic_name** column. If it finds the typecode, the daemon will create the user table with a SQL table schema derived from the typecode following the IDL type to SQL type mapping described in [Chapter 5: IDL/SQL Semantic and Data Mapping](#). Then the daemon will proceed to create the associated DataWriter. More about the creation of user tables by the daemon can be found in [User-Table Creation \(Section 4.6\)](#).

How the daemon discovers and stores typecodes is described in [Typecodes \(Section 4.1.4\)](#). If the *Database Integration Service* Daemon has not yet have a typecode associated with the **topic_name**, it will defer the creation of the DataWriter until the typecode is discovered. When a new typecode is discovered, the daemon will scan all rows in the **RTIDDS_PUBLICATIONS** meta-table and create the user tables and DataWriters for entries that were pending on the discovery of the typecode.

The daemon will also create the DataWriter if there is an entry in the **RTIDDS_PUBLICATIONS** table without an associated typecode, but the user subsequently creates the corresponding table.

If user applications modify an existing row in the **RTIDDS_PUBLICATIONS** table, the *Database Integration Service* Daemon will first delete the DataWriter that was created for that entry (if it exists) and then go through the same process of trying to create the user table and DataWriter as if the row was newly inserted. If user applications delete an existing row in the **RTIDDS_PUBLICATIONS** table, the *Database Integration Service* Daemon will delete the associated DataWriter (if it exists).

A flow chart describing this logic is provided below.



4.5.1.1 Publications Table Schema

The `RTIDDS_PUBLICATIONS` table is created with the following SQL statement.

Oracle TimesTen (with command-line option `-typeMode 1`):

```

Create Table RTIDDS_PUBLICATIONS (
  table_owner VARCHAR(128) NOT NULL,
  table_name VARCHAR(128) NOT NULL,
  domain_id INTEGER NOT NULL,
  topic_name VARCHAR(200),
  type_name VARCHAR(200),
  table_history_depth INTEGER,
  resolution_column VARCHAR(255),
  idl_member_prefix_max_length INTEGER,
  idl_member_suffix_max_length INTEGER,
  profile_name VARCHAR(255),
  "pub.present.access_scope" VARCHAR(25),
  "pub.present.ordered_access" TINYINT,
  "pub.partition.name" VARCHAR(256),
  "dw.durability.kind" VARCHAR(30),
  "dw.liveliness.lease_dur.sec" INTEGER,
  "dw.liveliness.lease_dur.nsec" INTEGER,
  "dw.deadline.period.sec" INTEGER,
  "dw.deadline.period.nsec" INTEGER,

```

```

"dw.history.kind" VARCHAR(21),
"dw.history.depth" INTEGER,
"dw.ownership.kind" VARCHAR(23),
"dw.ownership_strength.value" INTEGER,
"dw.publish_mode.kind" VARCHAR(29),
"dw.res_limits.max_samples" INTEGER,
"dw.res_limits.max_instances" INTEGER,
PRIMARY KEY(table_owner,table_name,domain_id,topic_name)
)

```

Oracle TimesTen (with command-line option -typeMode 0):

```

Create Table RTIDDS_PUBLICATIONS (
  table_owner TT_VARCHAR(128) NOT NULL,
  table_name TT_VARCHAR(128) NOT NULL,
  domain_id TT_INTEGER NOT NULL,
  topic_name TT_VARCHAR(200),
  type_name TT_VARCHAR(200),
  table_history_depth TT_INTEGER,
  resolution_column TT_VARCHAR(255),
  idl_member_prefix_max_length TT_INTEGER,
  idl_member_suffix_max_length TT_INTEGER,
  profile_name TT_VARCHAR(255),
  "pub.present.access_scope" TT_VARCHAR(25),
  "pub.present.ordered_access" TT_TINYINT,
  "pub.partition.name" TT_VARCHAR(256),
  "dw.durability.kind" TT_VARCHAR(30),
  "dw.liveliness.lease_dur.sec" TT_INTEGER,
  "dw.liveliness.lease_dur.nsec" TT_INTEGER,
  "dw.deadline.period.sec" TT_INTEGER,
  "dw.deadline.period.nsec" TT_INTEGER,
  "dw.history.kind" TT_VARCHAR(21),
  "dw.history.depth" TT_INTEGER,
  "dw.ownership.kind" TT_VARCHAR(23),
  "dw.ownership_strength.value" TT_INTEGER,
  "dw.publish_mode.kind" TT_VARCHAR(29),
  "dw.res_limits.max_samples" TT_INTEGER,
  "dw.res_limits.max_instances" TT_INTEGER,
  PRIMARY KEY(table_owner,table_name,domain_id,topic_name)
)

```

Oracle Database 11g:

```

Create Table RTIDDS_PUBLICATIONS (
  table_owner VARCHAR(128) NOT NULL,
  table_name VARCHAR(128) NOT NULL,
  domain_id NUMBER(10) NOT NULL,
  topic_name VARCHAR(200),
  type_name VARCHAR(200),
  table_history_depth NUMBER(10),
  resolution_column VARCHAR(255),
  idl_member_prefix_max_length NUMBER(10),
  idl_member_suffix_max_length NUMBER(10),
  profile_name VARCHAR(255),
  "pub.present.access_scope" VARCHAR(25),
  "pub.present.ordered_access" NUMBER(3),
  "pub.partition.name" VARCHAR(256),
  "dw.durability.kind" VARCHAR(30),
  "dw.liveliness.lease_dur.sec" NUMBER(10),
  "dw.liveliness.lease_dur.nsec" NUMBER(10),

```

```

"dw.deadline.period.sec" NUMBER(10),
"dw.deadline.period.nsec" NUMBER(10),
"dw.history.kind" VARCHAR(21),
"dw.history.depth" NUMBER(10),
"dw.ownership.kind" VARCHAR(23),
"dw.ownership_strength.value" NUMBER(10),
"dw.publish_mode.kind" VARCHAR(29),
"dw.res_limits.max_samples" NUMBER(10),
"dw.res_limits.max_instances" NUMBER(10),
PRIMARY KEY(table_owner,table_name,domain_id,topic_name)
)

```

MySQL¹:

```

Create Table RTIDDS_PUBLICATIONS (
  table_owner VARCHAR(128) NOT NULL,
  table_name VARCHAR(128) NOT NULL,
  domain_id INTEGER NOT NULL,
  topic_name VARCHAR(200),
  type_name VARCHAR(200),
  table_history_depth INTEGER,
  resolution_column VARCHAR(255),
  idl_member_prefix_max_length INTEGER,
  idl_member_suffix_max_length INTEGER,
  profile_name VARCHAR(255),
  "pub.present.access_scope" VARCHAR(25),
  "pub.present.ordered_access" TINYINT,
  "pub.partition.name" VARCHAR(256),
  "dw.durability.kind" VARCHAR(30),
  "dw.liveliness.lease_dur.sec" INTEGER,
  "dw.liveliness.lease_dur.nsec" INTEGER,
  "dw.deadline.period.sec" INTEGER,
  "dw.deadline.period.nsec" INTEGER,
  "dw.history.kind" VARCHAR(21),
  "dw.history.depth" INTEGER,
  "dw.ownership.kind" VARCHAR(23),
  "dw.ownership_strength.value" INTEGER,
  "dw.publish_mode.kind" VARCHAR(29),
  "dw.res_limits.max_samples" INTEGER,
  "dw.res_limits.max_instances" INTEGER,
  changes_queue_maximum_size INTEGER,
  RTIRTC_SCN BIGINT DEFAULT 0,
  PRIMARY KEY(table_owner,table_name,domain_id,topic_name)
)

```

Users should use the same SQL statement in their own applications if they want to create and populate this table before the *Database Integration Service* Daemon is started. [Table 4.17](#) describes how each column is used by the daemon in creating and using DataWriters that publish table changes.

Table 4.17 RTIDDS_PUBLICATIONS Table Schema

Column Name	SQL Type	Null-able	Default if NULL	Described in...
table_owner ^a	VARCHAR(128)	No	N/A	Section 4.5.1.1.4
table_name ^a	VARCHAR(128)	No	N/A	Section 4.5.1.1.4

1. See [Starting the MySQL Server in ANSI_QUOTES mode \(Section 4.1.2.4\)](#).

Table 4.17 RTIDDS_PUBLICATIONS Table Schema

Column Name	SQL Type	Null-able	Default if NULL	Described in...
domain_id ^a	INTEGER	No	N/A	Section 4.5.1.1.5
topic_name ^a	VARCHAR(200)	Yes	<table_owner>.<table_name>	Section 4.5.1.1.6
type_name	VARCHAR(200)	Yes	<topic_name>	Section 4.5.1.1.7
table_history_depth	INTEGER	Yes	0	Section 4.5.1.1.8
resolution_column	VARCHAR(255)	Yes	None	Section 4.5.1.1.9
idl_member_prefix_max_length	INTEGER	Yes	Value specified in the configuration file	Section 4.5.1.1.10
idl_member_suffix_max_length	INTEGER	Yes	Value specified in the configuration file	Section 4.5.1.1.10
profile_name	VARCHAR(255)	Yes	Database Integration Service will not use a profile to create the publication	Section 4.5.1.1.11
pub.present.access_scope	VARCHAR(25)	Yes	INSTANCE_PRESENTATION_QOS	Section 4.5.1.1.12
pub.present.ordered_access	TINYINT	Yes	0 (false)	Section 4.5.1.1.12
pub.partition.name	VARCHAR(256)	Yes	Empty string partition	Section 4.5.1.1.13
dw.durability.kind	VARCHAR(30)	Yes	VOLATILE_DURABILITY_QOS	Section 4.5.1.1.14
dw.liveliness.lease_dur.sec	INTEGER	Yes	Infinite	Section 4.5.1.1.15
dw.liveliness.lease_dur.nsec	INTEGER	Yes	Infinite	Section 4.5.1.1.15
dw.deadline.period.sec	INTEGER	Yes	Infinite	Section 4.5.1.1.16
dw.deadline.period.nsec	INTEGER	Yes	Infinite	Section 4.5.1.1.16
dw.history.kind	VARCHAR(21)	Yes	KEEP_LAST_HISTORY_QOS	Section 4.5.1.1.17
dw.history.depth	INTEGER	Yes	1	Section 4.5.1.1.17
dw.ownership.kind	VARCHAR(23)	Yes	SHARED_OWNERSHIP_QOS	Section 4.5.1.1.18
dw.ownership_strength.value	INTEGER	Yes	0	Section 4.5.1.1.18
dw.publish_mode.kind	VARCHAR(29)	Yes	SYNCHRONOUS_PUBLISH_MODE_QOS	Section 4.5.1.1.19
dw.res_limits.max_samples	INTEGER	Yes	Infinite	Section 4.5.1.1.20
dw.res_limits.max_instances	INTEGER	Yes	Infinite	Section 4.5.1.1.20
changes_queue_maximum_size	INTEGER	Yes	Infinite	Section 4.5.1.1.21
RTIRTC_SCN	BIGINT	Yes	Next SCN number	Section 4.5.1.1.22

a. Primary key column

4.5.1.1.4 table_owner, table_name

These columns specify the user table for which changes will be published using a DataWriter. Because a DBMS uses a combination of <table_owner>.<table_name> to identify a table, both of these columns must have valid values should the user want these entries to refer to an existing table.

If no table exists in the database with the identifier “<table_owner>.<table_name>” at the time that the daemon sees this entry in the RTIDDS_PUBLICATIONS meta-table, it will create a user table with this name automatically, see [User-Table Creation \(Section 4.6\)](#).

Note: In MySQL, the value of the table_owner column corresponds to the table schema or database name.

4.5.1.1.5 domain_id

This column specifies the domain ID that will be used to publish changes in the table. Before creating a DataWriter, if no DomainParticipant has previously been created with the domain ID, the *Database Integration Service* Daemon will create a DomainParticipant with the specified ID.

If the publications entry has associated a QoS profile, *Database Integration Service* will use the values in this profile to create the participant. The participant will also be configured using the QoS values of a profile when the attribute, `is_default_qos`, is set to 1 in that profile (see the *RTI Connext DDS Core Libraries User's Manual* for additional details).

4.5.1.1.6 topic_name

This column defines the Topic that will be used to publish the changes in the associated table. The `<topic_name>` need to match the Topic used by subscriptions in user applications that expect to receive data changes from the table. If the *Database Integration Service* Daemon has discovered the typecode associated with the `<topic_name>` and the user table does not exist in the database, the daemon will use the typecode to create the table using entries in the `<table_owner>` and `<table_name>` column. See [User-Table Creation \(Section 4.6\)](#) for more details.

4.5.1.1.7 type_name

This column defines the registered name of the type associated with the Topic defined using the column `<topic_name>`. If the user table does not exist in the database, the daemon will use the type name to find a typecode in the XML configuration file. See [User-Table Creation \(Section 4.6\)](#) for more details.

4.5.1.1.8 table_history_depth

The `<table_history_depth>` column in the `RTIDDS_PUBLICATIONS` determines whether or not the *Database Integration Service* Daemon will create the user table with additional meta-columns that support the storing of historic, or past, values of instances of Topics by DataReaders created with the `RTIDDS_SUBSCRIPTIONS` table. It is only used if the daemon creates the table because it does not exist.

More about the ability to store historic data in the table as well as the added meta-columns can be found in [table_history_depth \(Section 4.5.2.1.27\)](#) and in [User-Table Creation \(Section 4.6\)](#).

This column is useful in the case that the user wants the *Database Integration Service* Daemon to both publish and subscribe to a Topic for the same user table. The value set in `<table_history_depth>` will enable the daemon to create the user table correctly if the user wants to store more than a single value for an instance of the Topic in the table.

The possible values for `<table_history_depth>` column are:

- NULL or 0

These values should be used if the user does not want to store more than a single value for an instance of a Topic in the table.

If the *Database Integration Service* Daemon creates the table, it will not add any meta-columns for table history to the table schema.

- Any other value

For any non-zero value in this column, the *Database Integration Service* Daemon will add meta-columns for table history to the table schema when it creates the user table automatically.

A table's schema or definition cannot be changed to accommodate the table-history meta-columns after a table has been created. So a non-zero value for this column is useful if the user wants the table to be created with the ability to store historic values in support of entries in the `RTIDDS_SUBSCRIPTIONS` table that may be made later.

4.5.1.1.9 resolution_column

This column is used to designate one of the columns of the user table for use as the timestamp when data changes are published with the DataWriter. Instead of using the system time, when a row in the user table changes, the *Database Integration Service* Daemon will take the current value of the designated column and use it in the `DataWriter::write_w_timestamp()` method when publishing the value of the row.

The possible values for the `<resolution_column>` column are:

NULL

If this column is NULL, then the *Database Integration Service* Daemon will just call `DDS-DataWriter::write()` to publish the table changes. This implies that the source timestamp used by *Connex DDS* will be the system time when the write occurred.

"column_name"

The column name of any column in the user table that has a valid type. The column must be one of the following SQL types: INTEGER, SMALLINT, BIGINT, or TIMESTAMP.

If the user directs the daemon to use a column from the user table as the timestamp, then it is *imperative* to the proper operation of the publication that the value in the timestamp column is *monotonically increasing* with every table change. So when a change is made to a row of the table, the value in the column `<resolution_column>` must be larger than the last value of this column that was published.

The `<resolution_column>` can be used with the `<dr.destination_order.kind>` column of the `RTIDDS_SUBSCRIPTIONS` table to implement a conflict resolution policy in a system where *Database Integration Service* is used to implement database table replication across a network. See `dr.destination_order.kind` (Section 4.5.2.1.40) and `TableReplicationMode` on page 4-16 for more information.

4.5.1.1.10 idl_member_prefix_max_length, idl_member_suffix_max_length

These columns define how *Database Integration Service* maps IDL member identifiers into column names. In particular, they control how the column names are formed by using as a prefix n characters from the identifier's prefix and m characters from the identifier's suffix.

They can assume any value greater than or equal to -1. They cannot both be set to zero.

If a positive value n is provided for `idl_member_prefix_max_length`, *Database Integration Service* will use the first n characters from the IDL member identifier to compose the associated column name. A value of 0 tells *Database Integration Service* to compose the column name using only the last characters of the identifiers, as defined by the `'idl_member_suffix_max_length'` column. A value of -1, instructs *Database Integration Service* to use all the available characters.

If a positive value n is provided for `idl_member_suffix_max_length`, *Database Integration Service* will use the last n characters from the IDL member identifier to compose the associated column name. A value of 0 tells *Database Integration Service* to compose the column name using only the first characters of the identifiers, as defined by the `'idl_member_prefix_max_length'` column. A value of -1, instructs *Database Integration Service* to use all the available characters.

4.5.1.1.11 **profile_name**

This column specifies the name of the QoS Profile that *Database Integration Service* will use to create the publication.

The name must have the following format:

<QoS profile library name>::<QoS profile name>

See the *Connex DDS* documentation for a complete description of QoS Profiles.

The QoS values specified in the publication table (if they are not NULL) take precedence over the same values specified in the QoS profile.

4.5.1.1.12 **pub.present.access_scope, pub.present.ordered_access**

These two columns map directly to the `DDS_PresentationQosPolicy` of the `DDS_PublisherQos` used by the Publisher that is created with the `DataWriter` for publishing changes to the table. The `DDS_PresentationQosPolicy` specifies how the samples representing changes to data instances are presented to a subscribing application.

The specific columns affect the relative order of changes seen by subscribers to the table. The values of these columns must be coordinated with the values of the `DDS_PresentationQosPolicy` used by the Subscriber in the receiving application or else published data may not be received by the subscriber.

The possible values for the <**pub.present.access_scope**> column are:

- "INSTANCE_PRESENTATION_QOS" (default value if the column is NULL)
- "TOPIC_PRESENTATION_QOS"
- "GROUP_PRESENTATION_QOS"

The possible values for the <**pub.present.ordered_access**> column are:

- 0 (default value if the column is NULL)
- 1

For the best performance of the *Database Integration Service* Daemon, you should set <**pub.present.access_scope**> to "TOPIC_PRESENTATION_QOS" and <**pub.present.ordered_access**> to 1. This will require that the corresponding values in the `DDS_PresentationQosPolicy` of the Subscriber in the receiving applications to be changed to those values as well.

See the *Connex DDS* documentation for more details on how this QoS policy may be used. See also [sub.present.access_scope](#), [sub.present.ordered_access](#) (Section 4.5.2.1.36).

4.5.1.1.13 **pub.partition.name**

For publishing table changes, *Database Integration Service* creates a `DataWriter` per table. The `pub.partition.name` column maps directly to the `DDS_PartitionQosPolicy` of the `DDS_PublisherQos` used by the Publisher that is created with the `DataWriter`. The `DDS_PartitionQosPolicy` introduces a logical partition concept inside the 'physical' partition concept introduced by the domain ID. A Publisher can communicate with a Subscriber only if they have some partition in common. The value of the `pub.partition.name` column specifies a list of partitions separated by commas to which the Publisher belongs.

See the *Connex DDS* documentation for more details on how this QoS policy may be used. See also [sub.partition.name](#) (Section 4.5.2.1.37)

4.5.1.1.14 `dw.durability.kind`

This column maps directly to the `DDS_DurabilityQosPolicy` of the DataWriter created to publish table changes. By changing this policy, the *Database Integration Service* Daemon can be configured to resend past changes to the database table to remote applications as soon as their subscriptions are discovered.

Only changes made by local applications to the table will be sent. That is, if the daemon is configured to subscribe to and store changes into the table from remote DataWriters, those changes are not sent. In addition, the changes will only be sent if the DataReader is created with a reliable setting for its `DDS_ReliabilityQosPolicy`.

The number of past changes that will be sent is limited by the values of the `<dw.history.kind>`, `<dw.history.depth>` and `<dw.res_limits.max_samples>` columns.

The possible values for the `<dw.durability.kind>` column are:

`"VOLATILE_DURABILITY_QOS"`

This value prevents the daemon from sending past changes to the table to newly discovered DataReaders.

This also the default value if the column is NULL.

`"TRANSIENT_LOCAL_DURABILITY_QOS"`

This value will enable the daemon to send past changes to the table to newly discovered DataReaders. The DataReaders must be created with reliable `DDS_ReliabilityQosPolicy`.

NOTE: If a table exists when the *Database Integration Service* Daemon creates a DataWriter, the daemon will initialize the DataWriter with the current contents of the table such that those values will be sent to new DataReaders with their `DDS_DurabilityQosPolicy` set to `"TRANSIENT_LOCAL_DURABILITY_QOS"`.

See the *Connex* DDS documentation for more details on how this QoS policy may be used. See also `dr.durability.kind` (Section 4.5.2.1.38) and `dr.reliability.kind` (Section 4.5.2.1.39).

4.5.1.1.15 `dw.liveliness.lease_dur`

These columns specify the lease duration for the `DDS_LivelinessQosPolicy` for the DataWriter created to publish table changes. The user may need to change the lease duration if remote applications have modified their DataReaders' corresponding `DDS_LivelinessQosPolicy` to non-default values.

The possible values of the `<dw.liveliness.lease_dur.sec>` (seconds) and `<dw.liveliness.lease_dur.nsec>` (nanoseconds) columns are:

An infinite lease duration is specified if both columns are NULL or contain the value 2147483647 ($2^{31} - 1$). This is the DDS default value.

A non-zero value representing the number of seconds and nanoseconds for the lease duration.

Note: `DDS_LivelinessQosPolicy.kind` is always set to `DDS_AUTOMATIC_LIVELINESS_QOS`.

See the *Connex* DDS documentation for more details on how this QoS policy may be used. See also `dr.liveliness.lease_dur` (Section 4.5.2.1.41).

4.5.1.1.16 `dw.deadline.period`

These columns specify the deadline period for the `DDS_DeadlineQosPolicy` for the DataWriter created to publish table changes. The user may need to change the deadline period if remote

applications have modified their DataReaders' corresponding DDS_DeadlineQoSPolicy to non-default values.

The possible values of the `<dw.deadline.period.sec>` (seconds) and `<dw.deadline.period.nsec>` (nanoseconds) columns are:

- ❑ An infinite deadline period is specified if both columns are NULL or contain the value 2147483647 ($2^{31} - 1$). This is the DDS default value.
- ❑ A non-zero value representing the number of seconds and nanoseconds for the deadline period.

The DDS_DeadlineQoSPolicy sets a commitment by the DataWriter to publish a value for every data instance to DataReaders every deadline period. If this value is set to a non-infinite value, user applications must update the value of every instance of the Topic stored in the table within each deadline period or the contract with DataReaders that subscribe to the changes to the table will be violated.

See the *Connexit DDS* documentation for more details on how this QoS policy may be used. See also [dr.deadline.period](#) (Section 4.5.2.1.42).

4.5.1.1.17 dw.history.kind, dw.history.depth

These columns directly map to the DDS_HistoryQoSPolicy for the DataWriter created to publish table changes. The values set for this QoS policy affect the DDS_ReliabilityQoSPolicy and the DDS_DurabilityQoSPolicy.

Using a "KEEP_ALL_HISTORY_QOS" will ensure that reliable DataReaders will receive every change to the table reliably. With a "KEEP_LAST_HISTORY_QOS," the *Database Integration Service* Daemon will only guarantee that the last `<dw.history.depth>` changes for each data instance are sent reliably.

If the `<dw.durability.kind>` column of the row is set to "TRANSIENT_LOCAL_DURABILITY_QOS", then these columns determine how many past data changes are sent to new subscribers to table changes.

The possible values of the `<dw.history.kind>` and `<dw.history.depth>` columns are:

- ❑ "KEEP_LAST_HISTORY_QOS"
 - For this setting, the column `<dw.history.depth>` determines how many published changes for each data instance in the table are stored in the DataWriter to support reliability or durability.
 - `<dw.history.depth>` should be set to an integer greater than 0. The default value for history depth is 1 if this column is NULL.
- ❑ "KEEP_ALL_HISTORY_QOS" (default value if the column is NULL)
 - This setting implies that the DataWriter created to publish table changes will store all of the changes to the table that it has sent. The total number of changes that can be stored is limited by the value in the `<dw.res_limits-max_samples>` column.
 - For this setting, the value in `<dw.history.depth>` is ignored.

See the *Connexit DDS* documentation for more details on how this QoS policy may be used. See also [dw.durability.kind](#) (Section 4.5.1.1.14) and [dw.res_limits.max_samples](#), [dw.res_limits.max_instances](#) (Section 4.5.1.1.20).

4.5.1.1.18 **dw.ownership.kind, dw.ownership_strength.value**

These columns directly map to the DDS_OwnershipQoSPolicy and DDS_Ownership-Strength-QoSPolicy for the DataWriter created to publish table changes. These policies control whether or not DataReaders are allowed to receive changes to an instance of a Topic from multiple DataWriters simultaneously.

The possible values of the <dw.ownership.kind> and <dw.ownership_strength.value> columns are:

- "SHARED_OWNERSHIP_QOS" (default value if the column is NULL)

This setting allows DataReaders to receive updates for an instance of a Topic from multiple DataWriters at the same time.

- "EXCLUSIVE_OWNERSHIP_QOS"

This setting prevents a DataReader from receiving changes from more than a single DataWriter for an instance of a Topic at the same time.

The DataReader will receive changes for a topic instance from the DataWriter with the greatest value of ownership strength. If the liveliness of the DataWriter fails or if the DataWriter fails to write within a deadline period, then the DataReader will receive published changes to the topic instance from the DataWriter with the next highest ownership strength.

The ownership strength set is set in the <dw.ownership_strength.value> column. The default value is 0 if the column is NULL.

See the *Connex* DDS documentation for more details on how this QoS policy may be used. See also [dw.liveliness.lease_dur](#) (Section 4.5.1.1.15), [dw.deadline.period](#) (Section 4.5.1.1.16), and [dr.ownership.kind](#) (Section 4.5.2.1.44).

4.5.1.1.19 **dw.publish_mode.kind**

This column controls the type of DataWriter that *Database Integration Service* will create for publishing data. This column can take the following values:

- ASYNCHRONOUS_PUBLISH_MODE_QOS
- SYNCHRONOUS_PUBLISH_MODE_QOS (default value)

Asynchronous DataWriters were introduced in *Connex* DDS to support large data packets (greater than 64Kb). If IDL data types exceed the 64Kb limit and reliable communication is used, **dw.publish_mode.kind** must be set to 'ASYNCHRONOUS_PUBLISH_MODE_QOS'.

See the *Connex* DDS documentation for more details on the differences between synchronous and asynchronous DataWriters.

4.5.1.1.20 **dw.res_limits.max_samples, dw.res_limits.max_instances**

These columns set some parameters for the DDS_ResourceLimitsQoSPolicy for the DataWriter created to publish table changes. In particular, they control the amount of memory that the system is allowed to allocate for storing published data values as well as the total number of data instances (different primary keys) that can be handled by the DataWriter.

A value of -1 for either of these columns means infinite. An infinite setting means that the DataWriter is allowed to allocate memory as needed to store published table changes and manage new keys.

- The default value for dw.res_limits.max_samples (if set to NULL) is 32.
- The default value for dw.res_limits.max_instances (if set to NULL) is -1.

The number of keys that the DataWriter is allowed to manage places an upper limit on the number of rows that the related table in the database can have.

See the *Connexit DDS* documentation for more details on how this QoS policy may be used.

4.5.1.1.21 `changes_queue_maximum_size`

This column is available only for connections to a MySQL database. The value of the column configures the maximum size of the queue that maintains the list of uncommitted changes. Note that there is a separate queue per table.

A value of -1 is used to indicate unlimited size.

4.5.1.1.22 `RTIRTC_SCN`

The System Change Number (SCN) column is available only for connections to a MySQL database. The value of this column is automatically maintained by *Database Integration Service* and is usually of no interest to the application. For more information about the `RTIRTC_SCN` column see [Section 4.6](#).

4.5.2 Subscriptions Table

When entries (rows) are added to the meta-table `RTIDDS_SUBSCRIPTIONS`, the *Database Integration Service* Daemon will try to create a DataReader (and Subscriber along with a DomainParticipant if required) and use it to receive data via the *Connexit DDS* for a Topic and store values into the designated user table.

If the `RTIDDS_SUBSCRIPTIONS` table does not exist at startup, the *Database Integration Service* Daemon will create it with the table owner set to the user name of the database connection as specified in the daemon's configuration file, see [Section 4.4](#). The schema and meaning of the columns of this table are described in the next section.

You may insert new rows or modify the column values of existing rows in this table at any time. For a new row, the daemon will first check to see if the designated user table exists. If so, it will immediately create the DataReader with the QoS values specified by the entry. The name of the Topic to subscribe to may be specified by the `topic_name` column or automatically constructed as `<table_owner>.<table_name>` if the `topic_name` entry is NULL.

If the user table does not exist, the *Database Integration Service* Daemon will look for the typecode associated with the type defined in the `topic_name` column. If it finds the typecode, the daemon will create the user table with a SQL table schema derived from the typecode following the IDL type to SQL type mapping described in [Chapter 5: IDL/SQL Semantic and Data Mapping](#). Then the daemon will proceed to create the associated DataReader. More about the creation of user tables by the daemon can be found in [User-Table Creation \(Section 4.6\)](#).

How the daemon discovers and stores typecodes is described in [Typecodes \(Section 4.1.4\)](#).

If the *Database Integration Service* Daemon does not yet have a typecode associated with the `topic_name`, it will defer the creation of the DataReader until the typecode is discovered. When a new typecode is discovered, the daemon will scan all rows in the `RTIDDS_SUBSCRIPTIONS` meta-table and create the user tables and DataReaders for entries that were pending on the discovery of the typecode.


The daemon will also create the DataReader if there is an entry in the `RTIDDS_SUBSCRIPTIONS` table without an associated typecode, but the user subsequently creates the corresponding table.

If user applications modify an existing row in the `RTIDDS_SUBSCRIPTIONS` table, the *Database Integration Service* Daemon will first delete the DataReader that was created for that entry (if

it exists) and then go through the same process of trying to create the user table and DataReader as if the row was newly inserted.

If user applications delete an existing row in the **RTIDDS_SUBSCRIPTIONS** table, the *Database Integration Service* Daemon will delete the associated DataReader (if it exists).

A flow chart describing this logic is provided below.



4.5.2.1 Subscriptions Table Schema

The **RTIDDS_SUBSCRIPTIONS** table is created with the following SQL statement.

Oracle TimesTen (with command-line option `-typeMode 1`):

```

Create Table RTIDDS_SUBSCRIPTIONS (
  table_owner VARCHAR(128) NOT NULL,
  table_name VARCHAR(128) NOT NULL,
  domain_id INTEGER NOT NULL,
  topic_name VARCHAR(200),
  type_name VARCHAR(200),
  table_history_depth INTEGER,
  process_batch INTEGER,
  "process_period.sec" INTEGER,
  "process_period.nsec" INTEGER,
  commit_type VARCHAR(17),
  cache_maximum_size INTEGER,
  cache_initial_size INTEGER,
  delete_on_dispose INTEGER,
  idl_member_prefix_max_length INTEGER,

```

```

    idl_member_suffix_max_length INTEGER,
    profile_name VARCHAR(255),
    filter_duplicates TINYINT,
    ordered_store TINYINT,
    persist_state TINYINT,
    "sub.present.access_scope" VARCHAR(25),
    "sub.present.ordered_access" TINYINT,
    "sub.partition.name" VARCHAR(256),
    "dr.durability.kind" VARCHAR(30),
    "dr.reliability.kind" VARCHAR(27),
    "dr.destination_order.kind" VARCHAR(43),
    "dr.liveliness.lease_dur.sec" INTEGER,
    "dr.liveliness.lease_dur.nsec" INTEGER,
    "dr.deadline.period.sec" INTEGER,
    "dr.deadline.period.nsec" INTEGER,
    "dr.history.kind" VARCHAR(21),
    "dr.history.depth" INTEGER,
    "dr.ownership.kind" VARCHAR(23),
    "dr.time_filter.min_sep.sec" INTEGER,
    "dr.time_filter.min_sep.nsec" INTEGER,
    "dr.res_limits.max_samples" INTEGER,
    "dr.res_limits.max_instances" INTEGER,
    "dr.unicast.receive_port" INTEGER,
    "dr.multicast.receive_address" VARCHAR(39),
    "dr.multicast.receive_port" INTEGER,
    PRIMARY KEY(table_owner,table_name,domain_id,topic_name)
)

```

Oracle TimesTen (with command-line option -typeMode 0):

```

Create Table RTIDDS_SUBSCRIPTIONS (
    table_owner TT_VARCHAR(128) NOT NULL,
    table_name TT_VARCHAR(128) NOT NULL,
    domain_id TT_INTEGER NOT NULL,
    topic_name TT_VARCHAR(200),
    type_name TT_VARCHAR(200),
    table_history_depth TT_INTEGER,
    process_batch TT_INTEGER,
    "process_period.sec" TT_INTEGER,
    "process_period.nsec" TT_INTEGER,
    commit_type TT_VARCHAR(17),
    cache_maximum_size TT_INTEGER,
    cache_initial_size TT_INTEGER,
    delete_on_dispose TT_INTEGER,
    idl_member_prefix_max_length TT_INTEGER,
    idl_member_suffix_max_length TT_INTEGER,
    profile_name TT_VARCHAR(255),
    filter_duplicates TT_TINYINT,
    ordered_store TT_TINYINT,
    persist_state TT_TINYINT,
    "sub.present.access_scope" TT_VARCHAR(25),
    "sub.present.ordered_access" TT_TINYINT,
    "sub.partition.name" TT_VARCHAR(256),
    "dr.durability.kind" TT_VARCHAR(30),
    "dr.reliability.kind" TT_VARCHAR(27),
    "dr.destination_order.kind" TT_VARCHAR(43),
    "dr.liveliness.lease_dur.sec" TT_INTEGER,
    "dr.liveliness.lease_dur.nsec" TT_INTEGER,
    "dr.deadline.period.sec" TT_INTEGER,
    "dr.deadline.period.nsec" TT_INTEGER,

```

```

"dr.history.kind" TT_VARCHAR(21),
"dr.history.depth" TT_INTEGER,
"dr.ownership.kind" TT_VARCHAR(23),
"dr.time_filter.min_sep.sec" TT_INTEGER,
"dr.time_filter.min_sep.nsec" TT_INTEGER,
"dr.res_limits.max_samples" TT_INTEGER,
"dr.res_limits.max_instances" TT_INTEGER,
"dr.unicast.receive_port" TT_INTEGER,
"dr.multicast.receive_address" TT_VARCHAR(39),
"dr.multicast.receive_port" TT_INTEGER,
PRIMARY KEY(table_owner,table_name,domain_id,topic_name)
)

```

Oracle Database 11g:

```

Create Table RTIDDS_SUBSCRIPTIONS (
  table_owner VARCHAR(128) NOT NULL,
  table_name VARCHAR(128) NOT NULL,
  domain_id NUMBER(10) NOT NULL,
  topic_name VARCHAR(200),
  type_name VARCHAR(200),
  table_history_depth NUMBER(10),
  process_batch NUMBER(10),
  "process_period.sec" NUMBER(10),
  "process_period.nsec" NUMBER(10),
  commit_type VARCHAR(17),
  cache_maximum_size NUMBER(10),
  cache_initial_size NUMBER(10),
  delete_on_dispose NUMBER(10),
  idl_member_prefix_max_length NUMBER(10),
  idl_member_suffix_max_length NUMBER(10),
  profile_name VARCHAR(255),
  filter_duplicates NUMBER(3),
  ordered_store NUMBER(3),
  persist_state NUMBER(3),
  "sub.present.access_scope" VARCHAR(25),
  "sub.present.ordered_access" NUMBER(3),
  "sub.partition.name" VARCHAR(256),
  "dr.durability.kind" VARCHAR(30),
  "dr.reliability.kind" VARCHAR(27),
  "dr.destination_order.kind" VARCHAR(43),
  "dr.liveliness.lease_dur.sec" NUMBER(10),
  "dr.liveliness.lease_dur.nsec" NUMBER(10),
  "dr.deadline.period.sec" NUMBER(10),
  "dr.deadline.period.nsec" NUMBER(10),
  "dr.history.kind" VARCHAR(21),
  "dr.history.depth" NUMBER(10),
  "dr.ownership.kind" VARCHAR(23),
  "dr.time_filter.min_sep.sec" NUMBER(10),
  "dr.time_filter.min_sep.nsec" NUMBER(10),
  "dr.res_limits.max_samples" NUMBER(10),
  "dr.res_limits.max_instances" NUMBER(10),
  "dr.unicast.receive_port" NUMBER(10),
  "dr.multicast.receive_address" VARCHAR(39),
  "dr.multicast.receive_port" NUMBER(10),
  PRIMARY KEY(table_owner,table_name,domain_id,topic_name)
)

```

MySQL¹:

```

Create Table RTIDDS_SUBSCRIPTIONS (

```

```

table_owner VARCHAR(128) NOT NULL,
table_name VARCHAR(128) NOT NULL,
domain_id INTEGER NOT NULL,
topic_name VARCHAR(200),
type_name VARCHAR(200),
table_history_depth INTEGER,
process_batch INTEGER,
"process_period.sec" INTEGER,
"process_period.nsec" INTEGER,
commit_type VARCHAR(17),
cache_maximum_size INTEGER,
cache_initial_size INTEGER,
delete_on_dispose INTEGER,
idl_member_prefix_max_length INTEGER,
idl_member_suffix_max_length INTEGER,
profile_name VARCHAR(255),
filter_duplicates TINYINT,
ordered_store TINYINT,
persist_state TINYINT,
"sub.present.access_scope" VARCHAR(25),
"sub.present.ordered_access" TINYINT,
"sub.partition.name" VARCHAR(256),
"dr.durability.kind" VARCHAR(30),
"dr.reliability.kind" VARCHAR(27),
"dr.destination_order.kind" VARCHAR(43),
"dr.liveliness.lease_dur.sec" INTEGER,
"dr.liveliness.lease_dur.nsec" INTEGER,
"dr.deadline.period.sec" INTEGER,
"dr.deadline.period.nsec" INTEGER,
"dr.history.kind" VARCHAR(21),
"dr.history.depth" INTEGER,
"dr.ownership.kind" VARCHAR(23),
"dr.time_filter.min_sep.sec" INTEGER,
"dr.time_filter.min_sep.nsec" INTEGER,
"dr.res_limits.max_samples" INTEGER,
"dr.res_limits.max_instances" INTEGER,
"dr.unicast.receive_port" INTEGER,
"dr.multicast.receive_address" VARCHAR(39),
"dr.multicast.receive_port" INTEGER,
RTIRTC_SCN BIGINT DEFAULT 0,
PRIMARY KEY(table_owner,table_name,domain_id,topic_name)
)

```

Users should use the same SQL statement in their own applications if they want to create and populate this table before the *Database Integration Service* Daemon is started. [Table 4.18](#) describes how each column is used by the daemon in creating DataReaders and storing received data into tables.

Table 4.18 **RTIDDS_SUBSCRIPTIONS Table Schema**

Column Name	SQL Type	Null-able	Default if NULL	Described in...
table_owner ^a	VAR-CHAR(128)	No	N/A	Section 4.5.2.1.23
table_name ^a	VAR-CHAR(128)	No	N/A	Section 4.5.2.1.23

1. See [Starting the MySQL Server in ANSI_QUOTES mode \(Section 4.1.2.4\)](#).

Table 4.18 RTIDDS_SUBSCRIPTIONS Table Schema

Column Name	SQL Type	Null-able	Default if NULL	Described in...
domain_id ^a	INTEGER	No	N/A	Section 4.5.2.1.24
topic_name ^a	VAR-CHAR(200)	YES	<table_owner>.<table_name>	Section 4.5.2.1.25
type_name	VAR-CHAR(200)	YES	<topic_name>	Section 4.5.2.1.26
table_history_depth	INTEGER	YES	0	Section 4.5.2.1.27
process_batch	INTEGER	YES	10	Section 4.5.2.1.28
process_period.sec	INTEGER	YES	0	Section 4.5.2.1.28
process_period.nsec	INTEGER	YES	100000000	Section 4.5.2.1.28
commit_type	VARCHAR(17)	YES	COMMIT_ON_PROCESS	Section 4.5.2.1.28
cache_maximum_size	INTEGER	YES	0	Section 4.5.2.1.29
cache_initial_size	INTEGER	YES	0	Section 4.5.2.1.29
delete_on_dipose	INTEGER	YES	0	Section 4.5.2.1.30
idl_member_prefix_max_length	INTEGER	YES	Value specified in the configuration file	Section 4.5.2.1.31
idl_member_suffix_max_length	INTEGER	YES	Value specified in the configuration file	Section 4.5.2.1.31
profile_name	VAR-CHAR(255)	YES	Database Integration Service will not use a profile to create the publication	Section 4.5.2.1.32
filter_duplicates	TINYINT	YES	0	Section 4.5.2.1.33
ordered_store	TINYINT	YES	1	Section 4.5.2.1.34
persist_state	TINYINT	YES	0	Section 4.5.2.1.35
sub.present.access_scope	VARCHAR(25)	YES	INSTANCE_PRESENTATION_QOS	Section 4.5.2.1.36
sub.present.ordered_access	TINYINT	YES	0 (false)	Section 4.5.2.1.36
sub.partition.name	VAR-CHAR(256)	YES	Empty partition string	Section 4.5.2.1.37
dr.durability.kind	VARCHAR(30)	YES	VOLATILE_DURABILITY_QOS	Section 4.5.2.1.38
dr.reliability.kind	VARCHAR(27)	YES	BEST_EFFORT_RELIABILITY_QOS	Section 4.5.2.1.39
dr.destination_order.kind	VARCHAR(43)	YES	BY_RECEPTION_TIMESTAMP_DESTINATIONORDER_QOS	Section 4.5.2.1.40
dr.liveliness.lease_dur.sec	INTEGER	YES	Infinite	Section 4.5.2.1.41
dr.liveliness.lease_dur.nsec	INTEGER	YES	Infinite	Section 4.5.2.1.41
dr.deadline.period.sec	INTEGER	YES	Infinite	Section 4.5.2.1.42
dr.deadline.period.nsec	INTEGER	YES	Infinite	Section 4.5.2.1.42
dr.history.kind	VARCHAR(21)	YES	KEEP_LAST_HISTORY_QOS	Section 4.5.2.1.43
dr.history.depth	INTEGER	YES	1	Section 4.5.2.1.43
dr.ownership.kind	VARCHAR(23)	YES	SHARED_OWNERSHIP_QOS	Section 4.5.2.1.44
dr.time_filter.min_sep.sec	INTEGER	YES	0	Section 4.5.2.1.45
dr.time_filter.min_sep.nsec	INTEGER	YES	0	Section 4.5.2.1.45
dr.res_limits.max_samples	INTEGER	YES	Infinite	Section 4.5.2.1.46
dr.res_limits.max_instances	INTEGER	YES	Infinite	Section 4.5.2.1.46
dr.unicast.receive_port	INTEGER	YES	0	Section 4.5.2.1.47
dr.multicast_receive_address	VARCHAR(15)	YES	None	Section 4.5.2.1.48

Table 4.18 RTIDDS_SUBSCRIPTIONS Table Schema

Column Name	SQL Type	Null-able	Default if NULL	Described in...
dr.multicast.receive_port	INTEGER	YES	0	Section 4.5.2.1.49
RTIRTC_SCN	BIGINT	YES	Next SCN number	Section 4.5.2.1.50

a. Primary key column.

4.5.2.1.23 table_owner, table_name

These columns specify the user table into which data received by a DataReader will be stored. Because a DBMS uses a combination of <table_owner>.<table_name> to identify a table, both of these columns must have valid values should the user want these entries to refer to an existing table.

If no table exists in the database with the identifier “<table_owner>.<table_name>” at the time that the daemon sees this entry in the **RTIDDS_SUBSCRIPTIONS** meta-table, it will create a user table with this name automatically, see [User-Table Creation \(Section 4.6\)](#).

Note: In MySQL, the value of the table_owner column corresponds to the table schema or database name.

4.5.2.1.24 domain_id

This column specifies the domain ID that will be used to subscribe to Topics whose values will be stored in the table. Before creating a DataReader, if no DomainParticipant has previously been created with the domain ID, the *Database Integration Service* Daemon will create a DomainParticipant with the specified ID.

If the subscriptions entry has an associated QoS profile, *Database Integration Service* will use the values in this profile to create the participant. The participant will also be configured using the QoS values of a profile when the attribute, **is_default_qos**, is set to 1 in that profile (see the *RTI Connex DDS Core Libraries User’s Manual* for additional details).

4.5.2.1.25 topic_name

This column defines the Topic that will be subscribed to and whose received values will be stored in the associated table. The “<topic_name>” entry needs to match the Topic used by the DataWriters that are sending data changes.

If the *Database Integration Service* Daemon has discovered the typecode associated with the <topic_name> and the user table does not exist in the database, the daemon will use the typecode to create the table using entries in the <table_owner> and <table_name> column. See [User-Table Creation \(Section 4.6\)](#) for more details.

4.5.2.1.26 type_name

This column defines the registered name of the type associated with the Topic defined using the column <topic_name>. If the user table does not exist in the database, the daemon will use the type name to find a typecode in the XML configuration file. See [User-Table Creation \(Section 4.6\)](#) for more details.

4.5.2.1.27 table_history_depth

This column determines the number of values of each instance received by the DataReader that can be stored in the table by the *Database Integration Service* Daemon. For non-keyed Topics, there is only a single instance, thus the <table_history_depth> would correspond to the maximum size of the table (in rows).

For keyed Topics, the *Database Integration Service* Daemon may store up to <table_history_depth> values of each instance of the Topic that the DataReader receives. When the history depth is reached, the rows are reused as a circular buffer with the newest values replacing the oldest.

To support this capability, the associated user table may be created with additional columns, *meta-columns*, to help the *Database Integration Service* Daemon manage history for a table. Whether or not meta-columns need to be added to support table history is based on the value of the entry in <table_history_depth>.

The two meta-columns for supporting table history are:

RTIRTC_HISTORY_SLOT: INTEGER

This column is also added to the Primary Key of the table. There is usually no need for users to access this column, it is only used by the daemon. It is only needed since many DBMS systems do not allow you to alter the value of a Primary Key column.

RTIRTC_HISTORY_ORDER: INTEGER

This value of this column is maintained by the *Database Integration Service* Daemon when it stores data received via *Connex DDS* into the table. The column stores a strictly incrementing counter that represents the received sequence number (starting at 0) of the data that is stored in that row.

User should use a combination of the instance key and the value of **RTIRTC_HISTORY_ORDER** to find the latest data received for an instance in the table.

The possible values for the <table_history_depth> column are:

NULL or 0

Only the current value of an instance of the Topic is stored. For non-keyed topics, this implies the table will only have a single row. For keyed topics, each instance will correspond to a single row in the table. This is the most common value for tables that are published with *Connex DDS*.

No meta-columns are added to help manage history.

1

Exactly the same behavior as NULL or 0, a single value is stored in the table per instance of the Topic. However, table-history meta-columns *are* added to the table schema if the *Database Integration Service* Daemon creates the user table automatically.

This value is useful for preparing the table to store more than a single value per instance after the table is created. Because table schema cannot be changed to accommodate the table-history meta-columns after a table has been created, using a value of 1 for this column is useful if the user wants to store historic values of instances, but does not know how many instances to store at the time the entry is made.

n > 1

Meta-columns will be added to accommodate the storing of historic values for instances. The last n values received for an instance will be stored by the table.

-1

Meta-columns will be added to accommodate the storing of historic values for instances. *All* values received by the DataReader will be stored by the table.

See [User-Table Creation \(Section 4.6\)](#) for more information on meta-columns.

4.5.2.1.28 `process_batch`, `process_period`, `commit_type`

These columns allow users to tune the *Database Integration Service* Daemon for optimal throughput performance. When the daemon receives data from a `DataReader`, it may be configured to delay storing the data into a table and/or committing the transaction until more data arrives. For a data streams with high throughput, thousands of samples per seconds, the ability for the daemon to process incoming data in batches greatly improves the efficiency and ultimately the maximum sustainable throughput rate for a given Topic.

The trade-off is latency. The more data that is processed in a single batch, the more efficiently the processing can occur. However, A greater delay between the receiving of the data by the daemon and the time that it can be accessed by user applications in the database.

The column `<process_batch>` controls how many data samples are processed at a time by the *Database Integration Service* Daemon. Instead of executing SQL UPDATE or INSERT every time data is received, the daemon only stores the data after it receives a certain number of samples set by `<process_batch>`. If the value `<process_batch>` is greater than 1, then it is essential that the `<process_period.[sec,nsec]>` is set to be non-zero. Thus, the daemon will process stored data periodically, even if the total number of data samples received is less than `<process_batch>`.

`<process_period.[sec,nsec]>` is an upper limit on the amount of delay that will be incurred before received data is stored in the database. The period can be set to 0 only if `<process_batch>` is set to 1. This means that the daemon will store each data sample as it is received so there is no need for periodic processing of the received samples. Use these values to have the daemon store the data with minimal latency (at the cost of lower overall throughput).

Finally, using the column `<commit_type>`, you can choose whether or not the SQL UPDATE/INSERT statements are committed when each data sample is stored or after all of the data being processed have been stored. There is significant performance enhancement if the storing of multiple data samples is committed as a single transaction.

However, if there is a problem during an SQL commit, for example, the transaction log of the database is full, then the entire transaction is rolled back which means that none of the received data in that batch will be stored in the table. If the storing of each data sample is committed separately, then an error committing any one sample will only result in the loss of that sample.

The possible values of the `<process_batch>` column are:

`n > 0`

The daemon will process data samples in batches of `n`. A value less than or equal to 0 will result in an error that is logged by the daemon.

A value of `n = 1` means that the daemon will store each data sample as it arrives.

The default value is 10 if this column is NULL.

The possible values of the `<process_period.sec>` (seconds) and `<process_period.nsec>` (nanoseconds) columns are:

0

If both columns are 0, then the daemon will not commit received samples periodically.

`n > 0`

A background thread will process received but un-stored data at the period specified by these columns. It is essential that a non-zero period be used if `<process_batch>` is greater than 1 to insure that all received data is eventually stored.

The default value for process period is 0.1 seconds (0 sec, 100000000 nanosec) if both columns are NULL.

The possible values of the `<commit_type>` column are:

- ❑ “COMMIT_ON_PROCESS” (default value if the columns are NULL)

This value will direct the *Database Integration Service* Daemon to commit the storage of a batch of data as a single transaction. This will result in higher performance at the risk of losing more data than necessary when the transaction is rolled-back because an error with the database.

- ❑ “COMMIT_ON_SAMPLE”

This value will direct the daemon to commit the storage of each data sample as a separate transaction. Although the daemon will use more resources, if an error occurs when a transaction is committed, only that data sample is lost.

4.5.2.1.29 `cache_maximum_size`, `cache_initial_size`

These columns control the size of a cache, used to store keys that exist in the table, that the *Database Integration Service* Daemon maintains for each `DataReader`. When a data instance is received, the daemon first checks the cache to see if a row corresponding to the data already exists in the table. If the key is in the cache, then the daemon executes an SQL UPDATE to store the data in the table.

If the key does not exist in the cache, then the *Database Integration Service* Daemon will INSERT a row with the key instead. The key cache can greatly enhance the performance of the daemon in storing data into the database by saving an SQL operation each time data is received. Without a cache, the daemon would need to execute 2 SQL statements. to store data; with the cache, only 1.

The trade off is the memory used to store keys versus the performance gain.

The default values of `<cache_maximum_size>` and `<cache_initial_size>` are 0 if the columns are NULL. The sizes are specified as the number of keys.

For small tables, the cache could be sized to hold all of the keys. Thus the size of the cache would be the maximum number of rows in the table. However, this is not practical for large tables and thus the cache will be smaller.

4.5.2.1.30 `delete_on_dispose`

This column configures the behavior of the *Database Integration Service* Daemon when a `DataWriter` disposes an instance stored into the database. When `delete_on_dispose` is initialized to 0 (the default value), the rows corresponding to the instance will not be deleted from the database. If `delete_on_dispose` is initialized to 1, all the rows associated with the instance will be deleted from the database.

4.5.2.1.31 `idl_member_prefix_max_length`, `idl_member_suffix_max_length`

These columns define how *Database Integration Service* maps IDL member identifiers into column names. In particular, they control how the column names are formed by using as a prefix n characters from the identifier’s prefix and m characters from the identifier’s suffix.

They can assume any value greater than or equal to -1. They cannot both be set to zero.

If a positive value n is provided for `idl_member_prefix_max_length`, *Database Integration Service* will use the first n characters from the IDL member identifier to compose the associated column name. A value of 0 tells *Database Integration Service* to compose the column name using only the last characters of the identifiers, as defined by the ‘`idl_member_suffix_max_length`’ column. A value of -1, instructs *Database Integration Service* to use all the available characters.

If a positive value n is provided for `idl_member_suffix_max_length`, *Database Integration Service* will use the last n characters from the IDL member identifier to compose the associated column

name. A value of 0 tells *Database Integration Service* to compose the column name using only the first characters of the identifiers, as defined by the 'idl_member_prefix_max_length' column. A value of -1, instructs *Database Integration Service* to use all the available characters.

4.5.2.1.32 profile_name

This column specifies the name of the QoS Profile that *Database Integration Service* will use to create the subscription.

The name must have the following format:

<QoS profile library name>::<QoS profile name>

See the *RTI Connext DDS Core Libraries User's Manual* for a complete description of QoS Profiles.

QoS values specified in the subscription table (if they are not NULL) take precedence over the same values specified in the QoS profile.

4.5.2.1.33 filter_duplicates

There are multiple scenarios in which *Database Integration Service* may receive duplicate samples (see Chapter 11, *Mechanisms for Achieving Information Durability and Persistence*, in the *RTI Connext DDS Core Libraries User's Manual*). For example, if *RTI Persistence Service* is used in the system, *Database Integration Service* could receive the same sample from the original writer and from *RTI Persistence Service*.

The **filter_duplicates** column specifies whether or not duplicates should be filtered by the *Database Integration Service* Daemon. If duplicates are not filtered, the subscription data table may end up containing duplicate rows.

Note: Durable Reader State configurations (see Section 11.4 in the *RTI Connext DDS Core Libraries User's Manual*) are ignored by *Database Integration Service*. Duplicate filtering and subscription state persistence are implemented by the *Database Integration Service* Daemon.

4.5.2.1.34 ordered_store

This column specifies whether or not the samples associated with a DataWriter identified by a virtual GUID 'x' (see Chapter 11, *Mechanisms for Achieving Information Durability and Persistence*, in the *RTI Connext DDS Core Libraries User's Manual*) must be stored in the database in order. The field only applies when **filter_duplicates** (Section 4.5.2.1.33) is set to 1.

Ordered storage means that given a DataWriter with virtual GUID 'x', a sample with virtual sequence number 'sn+1' will be stored after a virtual sample with virtual sequence number 'sn'. If there is only one DataWriter with virtual GUID 'x' in the system (for example, if there are no *RTI Persistence Services*) the value of this column does not affect behavior.

Note: *Database Integration Service* stores samples in the database as soon as they are received by the *Database Integration Service* subscriptions (*Connext DDS DataReaders*). If **ordered_store** is set to 1 and there are multiple DataWriters with the same virtual GUID in the system, old samples will not be stored in the database. A sample with sequence number 'sn' will be ignored if a sample with sequence number 'sn+1' for the same virtual writer has been already stored in the database.

4.5.2.1.35 persist_state

This column specifies whether or not the state of a *Database Integration Service* subscription must be persisted into the database. The field only applies when **filter_duplicates** (Section 4.5.2.1.33) is set to 1. The subscription state is used on restore by *Database Integration Service* in order to avoid receiving duplicate samples.

4.5.2.1.36 **sub.present.access_scope, sub.present.ordered_access**

These two columns map directly to the DDS_PresentationQosPolicy of the DDS_SubscriberQos used by the Subscriber that is created with the DataReader for storing received data in the table. The DDS_PresentationQosPolicy specifies how the data instances sent by a publishing application are ordered before they are received.

The values of these columns must be coordinated with the values of the DDS_PresentationQosPolicy used by the Publisher in the sending application or else published data may not be received by the DataReader.

The possible values for the **<sub.present.access_scope>** column are:

- "INSTANCE_PRESENTATION_QOS"
This also the default value if the column is NULL.
- "TOPIC_PRESENTATION_QOS"
- "GROUP_PRESENTATION_QOS"

The possible values for the **<sub.present.ordered_access>** column are:

- 0 (default value if the column is NULL)
- 1

For the best performance of the *Database Integration Service* Daemon, you should set **<sub.present.access_scope>** to "TOPIC_PRESENTATION_QOS" and **<sub.present.ordered_access>** to 1. This will require that the corresponding values in the DDS_PresentationQosPolicy of the Publisher in the sending applications to be changed to those values as well.

See the *Connexit DDS* documentation for more details on how this QoS policy may be used. See also [pub.present.access_scope](#), [pub.present.ordered_access](#) (Section 4.5.1.1.12).

4.5.2.1.37 **sub.partition.name**

For capturing data in a table, *Database Integration Service* creates a DataReader per Topic. The *sub.partition.name* column maps directly to the DDS_PartitionQosPolicy of the DDS_SubscriberQos used by the Subscriber that is created with the DataReader. The DDS_PartitionQosPolicy introduces a logical partition concept inside the 'physical' partition concept introduced by the domain ID. A Subscriber can communicate with a Publisher only if they have some partition in common. The value of the **sub.partition.name** column specifies a list of partitions separated by commas to which the Subscriber belongs.

See the *Connexit DDS* documentation for more details on how this QoS policy may be used. See also [pub.partition.name](#) (Section 4.5.1.1.13)

4.5.2.1.38 **dr.durability.kind**

This column maps directly to the DDS_DurabilityQosPolicy of the DataReader created to subscribe to Topic data that is stored in the table. By changing this policy, the DataReader can be configured to request for past values published for the Topic to be sent by existing applications soon as their matching DataWriters are discovered.

The DataWriter's DDS_DurabilityQosPolicy must also be set appropriately to permit the sending of historic, or past, published data. In addition, the column **<dr.reliability.kind>** for the entry must be set to "RELIABLE_RELIABILITY_QOS" for historic data to be received.

The possible values for the **<dr.durability.kind>** column are:

- "VOLATILE_DURABILITY_QOS" (default value if the column is NULL)

This value means that the DataReader does not request past data to be sent.

“TRANSIENT_LOCAL_DURABILITY_QOS”

This value requests that existing DataWriters of the Topic send past data that they are storing to the DataReader.

See the *Connexit DDS* documentation for more details on how this QoS policy may be used. See also [dw.durability.kind](#) (Section 4.5.1.1.14) and [dr.reliability.kind](#) (Section 4.5.2.1.39).

4.5.2.1.39 **dr.reliability.kind**

This column sets the DDSReliabilityQosPolicy for the DataReader created to subscribe to Topic data that is stored in the table. The value in this column determines whether or not DataWriters will send their data reliably to the DataReader.

If the value for **<dr.durability.kind>** is “TRANSIENT_LOCAL_DURABILITY_QOS”, then the value for this column must be set to “RELIABLE_RELIABILITY_QOS”.

The possible values for the **<dr.reliability.kind>** column are:

“BEST_EFFORT_RELIABILITY_QOS” (default value if the column is NULL)

This value means that the DataWriters will send their data to the DataReader using best efforts. Data may be lost if the system is too busy.

“RELIABLE_RELIABILITY_QOS”

This value means that the DataWriters will send their data to the DataReader using a reliable protocol. The exact semantics of the reliable connection is controlled by the DDS_HistoryQosPolicy of both the DataWriter and DataReader.

See the *Connexit DDS* documentation for more details on how this QoS policy may be used. See also [dw.durability.kind](#) (Section 4.5.1.1.14), [dw.history.kind](#), [dw.history.depth](#) (Section 4.5.1.1.17), and [dr.history.kind](#), [dr.history.depth](#) (Section 4.5.2.1.43).

4.5.2.1.40 **dr.destination_order.kind**

This column sets the DestinationOrderQosPolicy for the DataReader created to subscribe to Topic data that is stored in the table. The value in this column determines how the DataReader treats data received for the same instance of the Topic from different DataWriters.

When a data instance is received, a timestamp associated with the data is compared to the timestamp of the last value of the data instance. If the time of the new data is older than the time of the last data received (for that instance), then the new data is dropped.

What this column does is to set which timestamp (the one associated with the source of the data when it was sent or the one associated with the data when it was received) the DataReader will use.

This column has no practical effect unless the value of the **<dr.ownership.kind>** column is “SHARED_OWNERSHIP_QOS”.

The possible values for the **<dr.destination_order.kind>** column are:

“BY_RECEPTION_TIMESTAMP_DESTINATIONORDER_QOS”
(default value if the column is NULL)

This configures the DataReader to use the timestamp of when the data was received to determine whether or not to drop the data. In practice, this setting means all data received from all DataWriters will be accepted since the timestamp will always be newer for the new data.

“BY_SOURCE_TIMESTAMP_DESTINATIONORDER_QOS”

This value means that the DataReader will use the timestamp that was sent with the data in determining whether or not to accept the data. This timestamp was added by the DataWriter when the data was published. Because different DataWriters may run in applications on different machines, it is likely that the clocks on the different machines are only synchronized to a certain resolution or not synchronized at all.

Thus the DataReader may receive data with timestamps older than the last data that received and thus drop those data. However if all DataReaders of the same Topic used the source timestamp to filter the data, then all DataReaders will end up with the same final value for a data instance.

If DataReaders used the reception timestamp, the DataReaders may end up with different final values because data from different DataWriters may be received in a different order by different DataReaders.

See the *Connexst DDS* documentation for more details on how this QoS policy may be used. See also [resolution_column](#) (Section 4.5.1.1.9).

4.5.2.1.41 `dr.liveliness.lease_dur`

These columns specify the lease duration for the `DDS_LivelinessQoSPolicy` for the DataReader created to subscribe to Topic data that is stored in the table. This value is useful when there are redundant DataWriters that publish values for the same data instance for the Topic and the value set for the `<dr.ownership.kind>` column is `“EXCLUSIVE_OWNERSHIP_QOS”`.

The liveliness of a DataWriter is monitored by the DataReader. These columns control how quickly the DataReader can determine that the DataWriter with the highest ownership strength has lost liveliness because heartbeat packets or data were not received within the liveliness lease duration. When liveliness is lost, the DataReader will then receive the data instance from the DataWriter with the next highest ownership strength that is still alive.

The possible values of the `<dr.liveliness.lease_dur.sec>` (seconds) and `<dr.liveliness.lease_dur.nsec>` (nanoseconds) columns are:

- ❑ An infinite lease duration is specified if both columns are NULL or contain the value 2147483647 ($2^{31} - 1$). This is the DDS default value.
- ❑ A non-zero value representing the number of seconds and nanoseconds for the lease duration.

Note: The `DDS_LivelinessQoSPolicy.kind` is always set to `DDS_AUTOMATIC_LIVELINESS_QOS`.

See the *Connexst DDS* documentation for more details on how this QoS policy may be used. See also [dw.liveliness.lease_dur](#) (Section 4.5.1.1.15), [dr.ownership.kind](#) (Section 4.5.2.1.44), and [dw.ownership.kind, dw.ownership_strength.value](#) (Section 4.5.1.1.18).

4.5.2.1.42 `dr.deadline.period`

These columns specify the deadline period for the `DDS_DeadlineQoSPolicy` for the DataReader created to subscribe to Topic data that is stored in the table. By setting the values in this column, the user is setting an expectation that DataWriters will publish new values for data instances at least as fast as the deadline period.

The possible values of the `<dr.deadline.period.sec>` (seconds) and `<dr.deadline.period.nsec>` (nanoseconds) columns are:

- ❑ An infinite deadline period is specified if both columns are NULL or contain the value 2147483647 ($2^{31} - 1$). This is the DDS default value.

- ❑ A non-zero value representing the number of seconds and nanoseconds for the deadline period.

See the *Connexst DDS* documentation for more details on how this QoS policy may be used. See also [dw.deadline.period](#) (Section 4.5.1.1.16).

4.5.2.1.43 **dr.history.kind, dr.history.depth**

These columns directly map to the DDS_HistoryQoSPolicy for the DataReader created to subscribe to Topic data that is stored in the table. The values set for this QoSPolicy affect the DDS_ReliabilityQoSPolicy.

Using a “KEEP_ALL_HISTORY_QOS” will ensure that reliable DataReaders will receive every change to the table reliably. With a “KEEP_LAST_HISTORY_QOS”, the *Database Integration Service* Daemon will only guarantee that the last <dr.history.depth> changes for each data instance are received reliably.

The possible values of the <dr.history.kind> and <dr.history.depth> columns are:

- ❑ “KEEP_LAST_HISTORY_QOS”

For this setting, the column <dr.history.depth> determines the maximum number of values for each data instance that be buffered in the DataReader before the *Database Integration Service* Daemon stores the received values into the table.

<dr.history.depth> should be set to an integer greater than 0. The default value for history depth is 1 if this column is NULL.

- ❑ “KEEP_ALL_HISTORY_QOS” (default value if the column is NULL)

This setting implies that the DataReader created to subscribe to Topic data has an unlimited queue in which to save received data before the data is stored in the table. The actual size of the queue is limited by the value in <dr.res_limits.max_samples> column.

For this setting, the value in <dr.history.depth> is ignored.

See the *Connexst DDS* documentation for more details on how this QoS policy may be used. See also [dr.res_limits.max_samples](#), [dr.res_limits.max_instances](#) (Section 4.5.2.1.46).

4.5.2.1.44 **dr.ownership.kind**

These columns directly map to the DDS_OwnershipQoSPolicy and DDS_Ownership-Strength-QoSPolicy for the DataReader created to subscribe to Topic data that is stored in the table. These policies control whether or not the DataReader is allowed to receive changes to an instance of a Topic from multiple DataWriters simultaneously.

The possible values of the <dr.ownership.kind> column are:

- ❑ “SHARED_OWNERSHIP_QOS” (default value if the column is NULL)

This setting allows the DataReader to receive updates for an instance of a Topic from multiple DataWriters at the same time.

- ❑ “EXCLUSIVE_OWNERSHIP_QOS”

This setting prevents the DataReader from receiving changes from more than a single DataWriter for an instance of a Topic at the same time.

The DataReader will receive changes for a topic instance from the DataWriter with the greatest value of ownership strength. If the liveliness of the DataWriter fails or if the DataWriter fails to write within a deadline period, then the DataReader will receive published changes to the topic instance from the DataWriter with the next highest ownership strength.

See the *Connexst DDS* documentation for more details on how this QoS policy may be used. See also [dr.liveliness.lease_dur](#) (Section 4.5.2.1.41), [dr.deadline.period](#) (Section 4.5.2.1.42), and [dw.ownership.kind](#), [dw.ownership_strength.value](#) (Section 4.5.1.1.18).

4.5.2.1.45 dr.time_filter.min_sep

This column specifies the minimum separation duration between subsequent samples for the DDS_TimeBasedFilterQosPolicy for the DataReader created to subscribe to Topic data that is stored in the table. By setting the values in these columns, the user configures the DataReader to see at most one change every the minimum_separation period.

The possible values of the `<dr.time_filter.min_sep.sec>` (seconds) and `<dr.time_filter.min_sep.nsec>` (nanoseconds) columns are:

- A 0 minimum separation duration is specified if both columns are NULL or contain the value 0. This is the DDS default value. With this value, the DataReader is potentially interested in all the samples.
- A non-zero value representing the number of seconds and nanoseconds for the minimum separation duration. That value must be smaller than the deadline period and contained in the interval [0, 1 year].

See the *Connexst DDS* documentation for more details on how this QoS policy may be used.

4.5.2.1.46 dr.res_limits.max_samples, dr.res_limits.max_instances

These columns set some parameters for the DDS_ResourceLimits QosPolicy for the DataReader created to subscribe to Topic data that is stored in the table. In particular, they control the amount of memory that the system is allowed to allocate for storing published data values as well as the total number of data instances (different primary keys) that can be handled by the DataReader.

A value of -1 for either of these columns means infinite. This is also the default value for these columns if they are NULL. An infinite setting means that the DataReader is allowed to allocate memory as needed to store received table changes and manage new keys.

The number of keys that the DataReader is allowed to manage places an upper limit on the number of rows that the related table in the database can have.

See the *Connexst DDS* documentation for more details on how this QoS policy may be used.

4.5.2.1.47 dr.unicast.receive_port

This column is used to configure the unicast port on which the DataReader will receive data. When the default value (NULL or 0) is used, the actual port number is determined by a formula as a function of the domain ID.

4.5.2.1.48 dr.multicast.receive_address

This column is used to set a multicast address for the DataReader to receive values for the Topic. The column maps to the DDS_TransportMulticastQosPolicy of the DataReader.

The possible values for the `<dr.multicast.receive_address>` column are:

- NULL
 - A NULL column means that the DataReader will receive Topic data using unicast.
- A string that contains a valid multicast address in the form "xxx.xxx.xxx.xxx".
 - The DataReader for the table will subscribe to the Topic on the multicast address provided.

See the *Connext DDS* documentation for more details on how this QoS policy may be used.

4.5.2.1.49 `dr.multicast.receive_port`

This column configures the multicast port on which the DataReader will receive data. When the default value (NULL or 0) is used, the actual port number is determined by a formula as a function of the domain ID.

Note that the value of this field is ignored when `dr.multicast.receive_address` is NULL

4.5.2.1.50 `RTIRTC_SCN`

The System Change Number (SCN) column is available only for connections to a MySQL database. The value of this column is automatically maintained by *Database Integration Service* and is usually of no interest to the application. For more information about the `RTIRTC_SCN` column see [Section 4.6](#).

4.5.3 Table Info

The meta-table `RTIRTC_TBL_INFO` stores meta information associated with the user tables.

When a table is automatically created by the *Database Integration Service* Daemon (see [Section 4.6](#)), its TypeCode is stored in `RTIRTC_TBL_INFO` as a sequence of octets. When the *Database Integration Service* Daemon is restarted, the persisted TypeCodes corresponding to existing publications and subscriptions will be made available to other *Connext DDS* applications.

4.5.3.1 Table Info Schema

The `RTIRTC_TBL_INFO` table is created with the following SQL statement:

Oracle TimesTen (with command-line option `-typeMode 1`):

```
Create Table RTIRTC_TBL_INFO (
    table_owner VARCHAR(128) NOT NULL,
    table_name VARCHAR(128) NOT NULL,
    type_code VARBINARY(65000),
    PRIMARY KEY(table_owner,table_name)
)
```

Oracle TimesTen (with command-line option `-typeMode 0`):

```
Create Table RTIRTC_TBL_INFO (
    table_owner TT_VARCHAR(128) NOT NULL,
    table_name TT_VARCHAR(128) NOT NULL,
    type_code VARBINARY(65000),
    PRIMARY KEY(table_owner,table_name)
)
```

Oracle:

```
Create Table RTIRTC_TBL_INFO (
    table_owner VARCHAR(128) NOT NULL,
    table_name VARCHAR(128) NOT NULL,
    type_code BLOB,
    PRIMARY KEY(table_owner,table_name)
)
```

MySQL:

```
Create Table RTIRTC_TBL_INFO (
    table_owner VARCHAR(128) NOT NULL,
    table_name VARCHAR(128) NOT NULL,
    type_code VARBINARY(65000),
```

```

RTIRTC_SCN BIGINT DEFAULT 0,
PRIMARY KEY(table_owner,table_name)
)

```

Table 4.19 describes the meta-table columns.

Table 4.19 **RTIRTC_TBL_INFO Table Schema**

Column Name	SQL Type	Nullable	Default if NULL	Described In
table_owner	VARCHAR(128)	No	N/A	Section 4.5.3.1.51
table_name	VARCHAR(128)	No	N/A	
type_code	VARCHAR(65000)	Yes	NULL	Section 4.5.3.1.52

4.5.3.1.51 table_owner, table_name

These columns specify the user table associated with the meta information described in the other columns.

Because a DBMS uses a combination of `<table_owner>.<table_name>` to identify a table, both of these columns must have valid values.

Note: In MySQL, the value of the `table_owner` column corresponds to the table schema or database name.

4.5.3.1.52 type_code

This column contains the TypeCode information associated to the user table identified by `<table_owner>.<table_name>`.

The TypeCode information stored in this table is used when publications and subscriptions are created after the *Database Integration Service* Daemon is restarted.

4.5.4 Log Table

A meta-table named `RTIRTC_LOG` is used to store log messages generated by the daemon. Whether or not this table is created and used depends on the `-loglevel` option (see [Section 4.2](#)) and the `LOGTODB` and `LOGHISTORY` *Database Integration Service* Daemon connection attributes (see [Section 4.4.4.2](#)).

Users should treat the contents of this table as *read-only*. There is no reason for users to modify this table. The number of rows in the Log table is controlled by the `LOGHISTORY` connection attribute. If set to `-1`, the table will hold as many log messages as generated by the *Database Integration Service* Daemon. Otherwise, the daemon will only store the last `n` log messages as specified by `LOGHISTORY`, using the table as a circular buffer.

Users may use the “id” column to determine the last log message that was generated by the daemon (see [Section 4.5.4.1.53](#)).

4.5.4.1 Log Table Schema

The `RTIRTC_LOG` table is created with the following SQL statement.

Oracle TimesTen (with command-line option `-typeMode 1`):

```

Create Table RTIRTC_LOG (
  id INTEGER NOT NULL,
  ts TIMESTAMP NOT NULL,
  type VARCHAR(7) NOT NULL,
  function VARCHAR(64) NOT NULL,
  line INTEGER,

```

```

code INTEGER,
native_code INTEGER,
message VARCHAR(2048) NOT NULL)

```

Oracle TimesTen (with command-line option -typeMode 0):

```

Create Table RTIRTC_LOG (
id TT_INTEGER NOT NULL,
ts TT_TIMESTAMP NOT NULL,
type TT_VARCHAR(7) NOT NULL,
function TT_VARCHAR(64) NOT NULL,
line TT_INTEGER,
code TT_INTEGER,
native_code TT_INTEGER,
message TT_VARCHAR(2048) NOT NULL)

```

Oracle:

```

Create Table RTIRTC_LOG (
id NUMBER(10) NOT NULL,
ts TIMESTAMP NOT NULL,
type VARCHAR(7) NOT NULL,
function VARCHAR(64) NOT NULL,
line NUMBER(10),
code NUMBER(10),
native_code NUMBER(10),
message VARCHAR(2048) NOT NULL)

```

MySQL:

```

Create Table RTIRTC_LOG (
id INTEGER NOT NULL,
ts TIMESTAMP NOT NULL,
type VARCHAR(7) NOT NULL,
function VARCHAR(64) NOT NULL,
line INTEGER,
code INTEGER,
native_code INTEGER,
message VARCHAR(2048) NOT NULL)

```

Each column of the Log meta-table stores a different portion of a log message generated by the *Database Integration Service Daemon*. [Table 4.20](#) describes these columns.

Table 4.20 **RTIRTC_LOG Table Schema**

Column Name	SQL Type	Nullable	Default if NULL	Described in...
id	INTEGER	NO	N/A	Section 4.5.4.1.53
ts	TIMESTAMP	NO	N/A	Section 4.5.4.1.54
type	VARCHAR(7)	NO	N/A	Section 4.5.4.1.55
function	VARCHAR(64)	NO	N/A	Section 4.5.4.1.56
line	INTEGER	NO	None	Section 4.5.4.1.56
code	INTEGER	YES	None	Section 4.5.4.1.57
native_code	INTEGER	YES	None	Section 4.5.4.1.57
message	VARCHAR(1024)	NO	N/A	Section 4.5.4.1.57

4.5.4.1.53 id

This column stores a strictly incrementing integer for each log message that is generated by the daemon. The largest value in the **id** column is the last message that was produced.

4.5.4.1.54 ts

This column stores the system timestamp of when the log message was generated.

4.5.4.1.55 type

This column stores the kind of log message. Possible values are: "ERROR", "WARNING", "STATUS", and "SPECIAL". "SPECIAL" messages are ones that are always printed independently of the log level.

4.5.4.1.56 function, line

These two columns contain the function name and line number of the *Database Integration Service* Daemon code where the message was generated. It is useful only to support engineers at RTI.

4.5.4.1.57 code, native_code, message

The **code** column contains the *Database Integration Service* error code that correspond to the message. This column will have NULL entries for messages of type "STATUS".

The **native_code** column will contain the error code of any external APIs, e.g., ODBC, OS, *Connect DDS*, that the daemon has called and returned an error. This column may have NULL entries.

Finally, the **message** column will contain a statement with details on why the message was generated.

For a complete list of possible error codes and messages that can be generated by the *Database Integration Service* Daemon, please see [Appendix A](#).

4.6 User-Table Creation

The *Database Integration Service* Daemon may create tables automatically for user applications in the database when entries are made in the **RTIDDS_PUBLICATIONS** or **RTIDDS_SUBSCRIPTIONS** meta-tables (see Sections 4.5.1 and 4.5.2). The daemon will create the table with the table owner and table name specified in an entry in one of those tables if:

1. There is no existing table in the database with the same **<table_owner>.<table_name>** identifier.
and
 2. A type corresponding to the **<type_name>** column for the entry has been defined in the XML configuration file (see [Section 4.4.3](#)).
- or
- A typecode corresponding to the **<topic_name>** column for the entry has been discovered.

If either condition above is not satisfied, then the daemon will not create the user table. If the user table already exists, then the daemon will attempt to use that table when publishing or subscribing to Topics. It is up to the user to create the table with a schema that maps to the Topic IDL type. See [Data Representation Mapping \(Section 5.2\)](#) for details on how SQL table schemas and IDL types are mapped to each other.

If the table does not exist and there is no XML definition for the type and the typecode for the IDL type specified by the entry is unknown, the *Database Integration Service* Daemon will defer

creation of the table until the typecode has been discovered from other applications on the network that are using *Connex DDS*. See [Typecodes \(Section 4.1.4\)](#) for more details on how the daemon uses typecodes.

If the table is created by the *Database Integration Service Daemon*, the daemon may add up to 5 additional columns (6 in MySQL) that store meta-data used by the daemon when storing data received via *Connex DDS* or sending table changes via *Connex DDS*. Although optional, there are specific operating scenarios where these meta-columns are required for the proper operation of the daemon. We suggest that the user understands the purpose of the meta-columns, and if the user applications create the tables used by the *Database Integration Service Daemon*, the user code itself should add the meta-columns to the table schema when appropriate. (There is no specific order required for the new columns.)

The meta-columns that may be created are:

❑ **RTIDDS_DOMAIN_ID and RTIRTC_REMOTE**

These two SQL INTEGER columns are always added to the tables created by the daemon. These additional columns are used by the daemon when user has created entries in both the **RTIDDS_PUBLICATIONS** and **RTIDDS_SUBSCRIPTIONS** meta-tables for the same user table. In that situation, changes to the table made by local user applications will be published via *Connex DDS* at the same time that the daemon itself may store data into the table received via *Connex DDS*.

Database Integration Service Daemon uses these meta-columns in order to prevent the republishing of tables values that were changed because they were received via *Connex DDS*. User applications that create the table do not need to add these columns if the daemon is configured only to publish data from the table or to store data into the table.

However, it is essential that these columns do exist for the situation where both publications and subscriptions are tied to the same table. If the meta-columns are omitted, then when *Database Integration Service Daemon* receives data via *Connex DDS*, it will be echoed (republished) as a change to the table.

❑ **RTIRTC_KEY**

This SQL INTEGER column is added by the daemon if the IDL type that is used to create the table does not contain any fields marked as a topic key (i.e., non-keyed IDL types). In such cases, the `<RTIRTC_KEY>` column will be added to the table as the primary key column. The value in that column will always be 0. Thus, there is only a single instance of the Topic which means the table will only ever have a single row (subject to whether or not the user wants the table to store historical value of data instances, see the details for the `<RTIRTC_-HISTORY_SLOT>` and `<RTIRTC_HISTORY_ORDER>` meta-columns below).

If the IDL type does have key fields, then the fields will be mapped into columns that are marked as primary keys. This meta-column is not added, and the table can contain as many rows as there are different instance keys (primary keys).

❑ **RTIRTC_HISTORY_SLOT and RTIRTC_HISTORY_ORDER**

These SQL INTEGER columns are used to implement the ability of the *Database Integration Service Daemon* to store multiple values (historical) of the same data instance into a table. Usually, a single data instance maps to a single row of a table. As new values for the instance is received by the daemon, the value of the same row is changed.

However, users may use the `<table_history_depth>` columns (see Sections [4.5.1.1.8](#) and [4.5.2.1.27](#)) of the **RTIDDS_PUBLICATIONS** and **RTIDDS_SUBSCRIPTIONS** meta-tables to direct the daemon to store multiple past values of a data instance. These values

are be stored in multiple rows of a table. To support table history, the daemon must add the meta-columns <RTIRTC_HISTORY_SLOT> and <RTIRTC_HISTORY_ORDER> to a table. They will only be added if the <table_history_depth> column for an entry is non-NULL and has a non-0 value.

The <RTIRTC_HISTORY_SLOT> is an auto-increment column that will also be added as a primary key column of the table.

The <RTIRTC_HISTORY_ORDER> is a column that will contain a number that is incremented as data is stored into the table. The oldest row of an instance will have the lowest value for this column whereas the most recent row of an instance will have the highest value.

❑ RTIRTC_SCN

The System Change Number (SCN) meta-column (SQL_BIGINT) is only required for connections to a MySQL database. The SCN meta-column is used to detect committed changes in a table. Its value is automatically assigned by the MySQL server.

Each time there is a change in a table row or a new row is inserted, the MySQL server assigns a new SCN value to the column RTIRTC_SCN. The assignment is done during the execution of the BEFORE UPDATE/INSERT trigger installed by the *Database Integration Service* Daemon.

4.7 Support for Extensible Types

Database Integration Service includes partial support for the "Extensible and Dynamic Topic Types for DDS" specification from the Object Management Group (OMG)¹. This section assumes that you are familiar with Extensible Types and you have read the *RTI Connex DDS Core Libraries Getting Started Guide Addendum for Extensible Types*.

- ❑ The *Database Integration Service* Daemon can publish and subscribe to topics associated with final and extensible types.
- ❑ The *Database Integration Service* Daemon can publish and/or subscribe to only one version of the types associated with a topic within a domain.
- ❑ The DataReaders created when a new entry is inserted in the RTIDDS_SUBSCRIPTIONS table are configured with the TypeConsistencyEnforcementQosPolicy's **kind** set to DISALLOW_TYPE_COERCION. You cannot overwrite this setting.
- ❑ You can select the type version for a given topic within a domain in two different ways:
 - By providing the type description via XML and then referring to that description using the column **type_name** in the RTIDDS_SUBSCRIPTIONS and/or RTIDDS_PUBLICATIONS tables.
 - By pre-creating the database table used to store/publish the data with the right schema.

Since the DataReaders created by the *Database Integration Service* Daemon have their TypeConsistencyEnforcementQosPolicy's **kind** set to DISALLOW_TYPE_COERCION, they will not match with *DataWriters* whose types are not equal to the *DataReader*'s type. You can learn more in the *RTI Connex DDS Core Libraries Getting Started Guide Addendum for Extensible Types*² (see the section on Rules for Type-Consistency Enforcement).

1. The OMG specification is available here: <http://www.omg.org/spec/DDS-XTypes/>.

4.8 Enabling RTI Distributed Logger in Database Integration Service

Database Integration Service provides integrated support for *RTI Distributed Logger*.

When you enable *Distributed Logger*, *Database Integration Service* will publish its log messages to *Connex DDS*. Then you can use *RTI Monitor*¹ to visualize the log message data. Since the data is provided in a *Connex DDS* topic, you can also use *rtiddsspy* or even write your own visualization tool.

To enable *Distributed Logger*, modify the *Database Integration Service* XML configuration file. In the `<general_options><administration>` section, add the `<distributed_logger>` tag as shown in the example below.

```
<real_time_connect name="default">
  <general_options>
    <administration>
      <domain_id>0</domain_id>
      <distributed_logger>
        <enabled>true</enabled>
      </distributed_logger>
    </administration>
  </general_options>
  ...
</real_time_connect>
```

Replace the value of `<domain_id>` with the domain ID that *Database Integration Service* will use to send log messages when *Distributed Logger* is enabled.

There are more configuration tags that you can use to control *Distributed Logger's* behavior. For example, you can specify a filter so that only certain types of log messages are published. For details, see the *Distributed Logger* section of the *RTI Connex DDS Core Libraries User's Manual*.

4.9 Enabling RTI Monitoring Library in Database Integration Service

To enable monitoring of the Entities that are created by *Database Integration Service*, you must specify the property `rti.monitor.library` in the QoS of the participants that you want to monitor. For example:

```
<participant_qos>
  <property>
    <value>
      <element>
        <name>rti.monitor.library</name>
        <value>rtimonitoring</value>
        <propagate>>false</propagate>
      </element>
    </value>
  </property>
</participant_qos>
```

2. `<Connex DDS installation directory>/n.dds.<version>/doc/pdf/RTI_ConnexDDS_CoreLibraries_GettingStarted_ExtensibleTypesAddendum.pdf`

1. *RTI Monitor* is a separate GUI application that can run on the same host as your application or on a different host.

The QoS associated with the DomainParticipants that are created by *Database Integration Service* can be configured in three different ways:

- ❑ By setting the attribute **is_default_qos** in the tag <qos_profile> containing the <participant_qos> to true. In this case, that profile is the default configuration for all the Entities created by the *Database Integration Service* Daemon.

For a list of XML files where you can declare the QoS Profile, see [Section 4.4.1](#)

- ❑ By referring to a profile using the XML tag <profile_name> within <publication> and <subscription> (see [Section 4.4.4.4](#)).
- ❑ By referring to a profile in the **profile_name** column of the tables RTIDDS_PUBLICATIONS or RTIDDS_SUBSCRIPTIONS (see [Section 4.5.1](#) and [Section 4.5.2](#)).

Notice that since *Database Integration Service* is statically linked with *RTI Monitoring Library*, you do not need to have it in your Path (on Windows systems) or LD_LIBRARY_PATH (on UNIX-based systems).

For details on how to configure the monitoring process, see the *Monitoring Library* section of the *RTI Connext DDS Core Libraries User's Manual*.

Chapter 5 IDL/SQL Semantic and Data Mapping

This chapter describes the semantic and data representation mapping that *RTI Database Integration Service* uses to connect DDS-based applications such as *Connex DDS* to MySQL, Microsoft SQL Server, Oracle, and Oracle TimesTen In-Memory databases.

Connex DDS provides an API to send and receive data between networked applications following a publish/subscribe paradigm. Oracle provides both file-based and in-memory products that allow applications to store and retrieve data following a relational database paradigm. The corresponding standard API in the database world is SQL. Both the *Connex DDS* and SQL APIs have various language bindings in C/C++ and Java.

How *Database Integration Service* maps actions (semantics) and data types (data representation) from *Connex DDS* to SQL and vice versa is described in the following sections.

- ❑ [Semantic Mapping \(Section 5.1\)](#)
- ❑ [Data Representation Mapping \(Section 5.2\)](#)

5.1 Semantic Mapping

Connex DDS applications publish and subscribe to topics which are named data structures using functions like `DDSDataWriter::write()` and `DDSDataReader::read()`. Relational databases contain tables that applications access data using SQL operations such as **INSERT**, **UPDATE**, **DELETE** and **SELECT**. [Table 5.1](#) describes the mapping between *Connex DDS* and relational database semantic models.

Table 5.1 **Connex DDS-DBMS Semantic Models**

Connex DDS	Relational Database	Details
Accessed via <i>Connex DDS</i> API Various language bindings (e.g. C/C++, Java)	Accessed via SQL Various language bindings (e.g. C/C++ and ODBC, Java and JDBC)	
Data structures Defined by IDL (Interface Description Language).	Tables Defined by table schema.	Fields in data structures are mapped to columns of a table. Each row of a table represents a different value for a data structure. The exact mapping of IDL data structures to table schemas is described in Data Representation Mapping (Section 5.2) .

Table 5.1 Connex DDS-DBMS Semantic Models

Connex DDS	Relational Database	Details
<p>Topic</p> <p>Identified by a name string.</p> <p>DataWriter can publish values for Topics and DataReaders can subscribe</p>	<p>Table</p> <p>Identified by a name string.</p> <p>Applications can write values or read values from tables using SQL.</p>	<p>Topic names and table names do not have to be the same when making a correspondence between a Topic and a database table.</p>
<p>Data values</p>	<p>Rows in table</p> <p>No history: A single row in a table.</p> <p>History: Multiple rows in a table.</p>	<p>When the <i>Database Integration Service</i> Daemon table history option is turned OFF (see Sections 4.5.1.1.8 and 4.5.2.1.27), only the last value of a topic instance is stored in the table. So a non-keyed topic will be stored in a single row whereas for keyed topics, there will be as many rows as there are topic instances.</p> <p>When the <i>Database Integration Service</i> Daemon table history option is turned ON, each instance will occupy up to a user-settable maximum number of rows so that the last N values received for the Topic are stored in the table. When N values have been stored, the N rows are used as a circular buffer so that new values received will overwrite the oldest values stored.</p>
<p>Key</p> <p>IDL data types may contain one or more fields that are used to distinguish different instances of the Topic.</p>	<p>Primary key</p> <p>Most relational databases require table schemes to identify one or more columns to act as the primary key for the table.</p>	<p>Keys are mapped to the primary keys of a table. When a table is created by the <i>Database Integration Service</i> Daemon, the columns corresponding to the IDL key fields will be created as primary key columns.</p> <p>For tables created by user code, the correspondence of IDL key fields to table primary key columns must be set correctly.</p>
<p>DDSDataWriter::write()</p>	<p>SQL INSERT or UPDATE</p>	<p>Values published for Topics will be stored into a database table by the <i>Database Integration Service</i> Daemon.</p> <p>Table rows modified by SQL INSERT or UPDATE commands will be published by the <i>Database Integration Service</i> Daemon as values of Topics.</p>
<p>DDSDataReader::take()</p> <p>DDSDataReader::read()</p>	<p>SQL SELECT</p>	
<p>DDSDataWriter::dispose()</p>	<p>SQL DELETE</p>	<p>When SQL DELETE is used to delete a row from a table, the <i>Database Integration Service</i> Daemon will call DDSDataWriter::dispose() to dispose the instance corresponding to the row.</p> <p>If a user application calls DDSDataWriter::dispose() to dispose an instance, the <i>Database Integration Service</i> Daemon may be configured to delete or keep the corresponding rows.</p>

5.2 Data Representation Mapping

In *Connext DDS*, data is stored in data structures or classes defined using the Interface Definition Language (IDL). In relational databases, data is stored in tables defined using SQL table schemas. While there is a good correspondence of IDL primitive data types to SQL data types, this mapping is not one-to-one. Both IDL and SQL have data types that the other does not define nor has an unambiguous mapping. In addition, many complex data structures in IDL such as unions and data structures that contain embedded data structures do not have equivalents in SQL.

This section describes the mapping used by the *Database Integration Service* Daemon when taking data received with DDS and storing it in tables, or taking data from tables and publishing it with *Connext DDS*.

- ❑ [IDL to SQL Mapping \(Section 5.2.1\) on Page 5-3](#)
- ❑ [Primitive Types Mapping \(Section 5.2.2\) on Page 5-6](#)
- ❑ [Oracle In-Memory Database Cache Mapping \(Section 5.2.3\) on Page 5-9](#)
- ❑ [Bit Field Mapping \(Section 5.2.4\) on Page 5-9](#)
- ❑ [Enum Types Mapping \(Section 5.2.5\) on Page 5-11](#)
- ❑ [Simple IDL Structures \(Section 5.2.6\) on Page 5-11](#)
- ❑ [Complex IDL Structures \(Section 5.2.7\) on Page 5-11](#)
- ❑ [Array Fields \(Section 5.2.8\) on Page 5-12](#)
- ❑ [Sequence Fields \(Section 5.2.9\) on Page 5-13](#)
- ❑ [NULL Values \(Section 5.2.10\) on Page 5-14](#)
- ❑ [Sparse Data Types \(Section 5.2.11\) on Page 5-14](#)

5.2.1 IDL to SQL Mapping

Identifiers are used for the names of table columns in SQL and names of fields within an IDL structure. SQL identifiers are a superset of IDL identifiers. Because of that, an IDL identifier can always be used as a SQL identifier. However, there are some SQL identifiers that cannot be used as IDL identifiers. For example, SQL allows special characters like '#' to be part of an identifier, whereas IDL does not.

There are two kind of SQL identifiers: *quoted identifiers* and *basic identifiers*. The quoted identifiers can use any combination of characters. Those identifiers need to be surrounded by double quotes when referenced. The definition of a basic identifier changes depending on the database vendor.

In *Oracle TimesTen*, a basic identifier can consist of any of letters (A to Z), decimal digits (0 to 9), \$, #, @, or underscore (_). The first character must be a letter.

In *Oracle*, a basic identifier can consist of any of letters (A to Z), decimal digits (0 to 9), \$, #, or underscore (_). The first character must be a letter.

In *MySQL*, a basic identifier can consist of any letters (A to Z), decimal digits (0 to 9), \$, or underscore (_).

In *Microsoft SQL Server*, a basic identifier can consist of any letters (A to Z), decimal digits (0 to 9), \$, #, or @. The first character may be a letter, @, or #.

If the daemon creates the user table, it will use quoted strings for the identifiers of the table and column names only if they cannot be considered as basic identifiers according to the previous

definitions. Thus, user applications should also use quoted strings when referring to those column and table names in their SQL statements.

Ordinarily, the name of a field in an IDL data structure can just be used as the name of a column in a table. In fact, for those data types with clear and obvious mappings, the column name can be independent of the field name used in the IDL type. However, because there is no one-to-one mapping of all IDL data types to all SQL data types, for certain types, the column names used in SQL table schemas *must* follow certain conventions that tie them to the names of the fields of IDL types from which they are mapped. This is true for only the small subset of primitive IDL data types and for the complex IDL data types that would otherwise have ambiguous mappings, i.e., multiple ways to map IDL to SQL or vice versa.

The *Database Integration Service* Daemon scans for, identifies and uses special mappings of column names when serializing and deserializing IDL data to and from a table in a database. There are two types of special mappings, hierarchical naming and suffixes.

Hierarchical Naming

Complex IDL types may have fields that are actually embedded structures, so a field may actually contain multiple values. In SQL, columns usually contain a single value for each column element, although there are a few types like `BINARY(x)`, `CHAR(x)`, `VARBINARY(x)` and `VARCHAR(x)` that can store multiple values of the same type in a single column element. To map complex IDL types to SQL table schemas, embedded data structures are unfolded so that elements of an embedded structure are stored individually in separate columns.

When the *Database Integration Service* Daemon creates a table schema from a Topic, it will automatically flatten hierarchical data structures into tables. In doing so, the names of columns that store the fields of embedded structures will have hierarchical names. For example, given this IDL definition:

```
struct bar {          struct foo {
    long one;         bar element;
    long two;        };
};
```

The table constructed from a Topic which uses the `foo` type would have the following schema by default:

```
CREATE table foo ( INTEGER element.one, INTEGER element.two )
```

The *Database Integration Service* Daemon allows the configuration of the separator character (‘.’) using the attribute *IdentifierSeparatorChar* defined in the general options of the configuration file (described in [Section 4.4.4.1](#)).

While for most embedded structures, the hierarchical naming of columns is not needed for the *Database Integration Service* Daemon to handle type translation correctly, the proper hierarchical naming of columns is essential for the daemon to serialize and deserialize IDL unions and sequences. These types are variable in length, however the table must have enough columns to hold the maximum size of the IDL data type. Hierarchical naming allows the *Database Integration Service* Daemon to identify columns that form an embedded, complex element.

For variable-length types (other than sequences of “char”, “wchar” or “octet”), an extra column with the suffix “#length” is also added to the table to hold the current length of the type. Also, each column that represents a field in an element of the variable-length type must have a suffix “[x]” in its name that identifies the index of the element, where

$$x = 0 \text{ to } (\text{max_length} - 1)$$

During the serialization and deserialization process, the daemon will usually be working with less than the maximum length of data, and thus, will need to use the hierarchical naming along with the suffix to determine which columns belong to unused elements that should be skipped.

This hierarchical flattening operation of member names may lead to very long column names in the generated table and can easily exceed the maximum number of characters supported by the database (some databases limit the column names to 30 characters).

To reduce the length of the generated names, you can instruct *Database Integration Service* to consider only the first n and the last m characters of the flattened name, and eventually resolve any conflict by using a progressive number between the prefix and the suffix. The two tags `<idl_member_prefix_max_length>` and `<idl_member_suffix_max_length>` (see [page 4-19](#)), defined in the configuration file (described in [Section 4.4](#)) and the columns `idl_member_prefix_max_length` and `idl_member_suffix_max_length` in the meta-tables (described in [Section 4.5.1.1.10](#)) tell the daemon the values to use. (The values defined in the meta-table have precedence over the values defined in the configuration file.)

Suffixes

Suffixes are also needed for column names when multiple IDL primitive types map into the same SQL type. Because there are more IDL primitive types than SQL primitive types, a full mapping will result in the use of the same SQL type to hold more than one IDL type. For example, an IDL “long double” has no equivalent in SQL. Thus, a SQL BINARY(16) does double duty and is used to store both an IDL “long double” as well as an IDL “octet[16]”.

If a “long double” could be treated the same as an “octet[16]” by the *Database Integration Service* Daemon, then there would be no issue and no special name mapping would be needed. However, because the representation of a “long double” is Endianess-dependent while an “octet[16]” is not, the *Database Integration Service* Daemon *must* use the column name to decide whether or not a SQL BINARY(16) value needs to be byte swapped or not when converting to an IDL data type. Since “long double” has no equivalent SQL type, a “.ld” must be appended to the name of a SQL BINARY(16) column that is used to store one.

Similarly, a suffix of “.str” is used to indicate that a SQL VARCHAR(x) stores IDL “string”, which is a NULL-terminated sequence of the primitive type “char”. Without the suffix in the column name, a SQL VARCHAR(x) naturally stores a sequence of chars—the IDL type “sequence <char,x>”.

For the Oracle database, but not the Oracle TimesTen In-Memory database, the IDL “octet”, “octet[x]” and “sequence<octet,x>” are all stored in the Oracle type RAW. A suffix of “.bin” is used to distinguish between using RAW to store “octet” and “octet[x]” which can be treated the same, and “sequence<octet,x>” which must be treated differently by the *Database Integration Service* Daemon.

NOTE: Because of the use of suffixes in the mapping of identifiers of certain IDL datatypes, the identifiers “**str**”, “**ld**”, and “**bin**” are reserved keywords that should not be used as the name of fields in IDL structures. For example, the following IDL definitions have the same SQL mapping which would result in the incorrect treatment of the type “**Foo2**” by the daemon. Each would result in a table schema that would have the ten columns named “**my_field[0].str**”, “**my_field[1].str**”, ..., “**my_field[2].str**”.

```

struct Foo1 {
    string<10> my_field;
};

struct Bar {
    sequence<char,10> str;
};

struct Foo2 {
    struct Bar my_field;
}

```

5.2.2 Primitive Types Mapping

The following tables show the mapping between basic types in IDL and SQL:

- ❑ [Table 5.2, “Basic Types in IDL and SQL \(TimesTen\),” on page 5-6](#)
- ❑ [Table 5.3, “Basic Types in IDL and SQL \(Oracle\),” on page 5-7](#)
- ❑ [Table 5.4, “Basic Types in IDL and SQL \(MySQL\),” on page 5-8](#)
- ❑ [Table 5.5, “Basic Types in IDL and SQL \(SQL Server\),” on page 5-8](#)

Table 5.2 **Basic Types in IDL and SQL (TimesTen)**

IDL Type	IDL Field Name	SQL Type (TypeMode 0)	SQL Type (TypeMode 1)	Table Column Name
char ^a	my_field	TT_CHAR(1) or CHAR(1)	CHAR(1)	“my_field”
char[x] ^a	my_field	TT_CHAR(x) or CHAR(x)	CHAR(x)	“my_field”
sequence<char,x>	my_field	TT_VARCHAR(x) or VARCHAR2(x)	VARCHAR(x)	“my_field”
wchar ^b	my_field	TT_NCHAR(x) or NCHAR(1)	NCHAR(1)	“my_field”
wchar[x] ^b	my_field	TT_NCHAR(x) or NCHAR(x)	NCHAR(x)	“my_field”
sequence<wchar,x>	my_field	TT_NVARCHAR(x) or NVARCHAR2(x)	NVARCHAR(x)	“my_field”
octet ^c	my_field	BINARY(1)	BINARY(1)	“my_field”
octet[x] ^c	my_field	BINARY(x)	BINARY(x)	“my_field”
sequence<octet,x> ^c	my_field	VARBINARY(x)	VARBINARY(x)	“my_field”
boolean	my_field	TT_TINYINT	TINYINT	“my_field”
short	my_field	TT_SMALLINT	SMALLINT	“my_field”
unsigned short	my_field	TT_SMALLINT	SMALLINT	“my_field”
unsigned	my_field	TT_INTEGER	INTEGER	“my_field”
unsigned long	my_field	TT_INTEGER	INTEGER	“my_field”
double	my_field	BINARY_DOUBLE	DOUBLE	“my_field”
float	my_field	BINARY_FLOAT	REAL	“my_field”
string<x>	my_field	TT_VARCHAR(x) or VARCHAR2(x)	VARCHAR(x)	“my_field.str” ^d
wstring<x>	my_field	TT_NVARCHAR(x) or NVARCHAR2(x)	NVARCHAR(x)	“my_field.str” ^d
long long	my_field	TT_BIGINT	BIGINT	“my_field”

Table 5.2 Basic Types in IDL and SQL (TimesTen) (Continued)

IDL Type	IDL Field Name	SQL Type (TypeMode 0)	SQL Type (TypeMode 1)	Table Column Name
unsigned long long	my_field	TT_BIGINT	BIGINT	"my_field"
long double	my_field	BINARY(16)	BINARY(16)	"my_field.Id" ^e
unsigned long long	my_field	TT_TIMESTAMP	TIMESTAMP	"my_field"

- The format on the wire of "char" and "char[x]" is the same.
- The format on the wire of "wchar" and "wchar[x]" is the same.
- The format on the wire of "octet" and "octet[x]" is the same.
- The ".str" suffix is used to distinguish between "(w)string<x>" and "sequence<(w)char,x>".
- The ".Id" suffix is used to distinguish between "octet[x]" and "long double".

Table 5.3 Basic Types in IDL and SQL (Oracle)

IDL Type	IDL Field Name	SQL Type	Table Column Name
char ^a	my_field	CHAR(1)	"my_field"
char[x] ^a	my_field	CHAR(x)	"my_field"
sequence<char,x>	my_field	VARCHAR2(x) if x <=4000; otherwise CLOB	"my_field"
wchar ^b	my_field	NCHAR(1)	"my_field"
wchar[x] ^b	my_field	NCHAR(x)	"my_field"
sequence<wchar,x>	my_field	NVARCHAR2(x) if x <=4000; otherwise NCLOB	"my_field"
octet ^c	my_field	RAW(1)	"my_field.bin" ^d
octet[x] ^c	my_field	RAW(x)	"my_field.bin" ^d
sequence<octet,x> ^c	my_field	RAW(x) if x <= 2000; otherwise BLOB	"my_field"
boolean	my_field	NUMBER(3)	"my_field"
short	my_field	NUMBER(5)	"my_field"
unsigned short	my_field	NUMBER(5)	"my_field"
unsigned	my_field	NUMBER(10)	"my_field"
unsigned long	my_field	NUMBER(10)	"my_field"
double	my_field	BINARY_DOUBLE	"my_field"
float	my_field	BINARY_FLOAT	"my_field"
string<x>	my_field	VARCHAR2(x) if x <= 4000; otherwise CLOB	"my_field.str" ^e
wstring<x>	my_field	NVARCHAR2(x) if x <=4000; otherwise NCLOB	"my_field.str" ^d
long long	my_field	NUMBER(20)	"my_field"
unsigned long long	my_field	NUMBER(20)	"my_field"
long double	my_field	RAW(16)	"my_field.Id" ^f
unsigned long long	my_field	TIMESTAMP	"my_field"

- The format on the wire of "char" and "char[x]" is the same.
- The format on the wire of "wchar" and "wchar[x]" is the same.
- The format on the wire of "octet" and "octet[x]" is the same.
- The ".bin" suffix is only used for Oracle databases (not Oracle TimesTen). The reason is that Oracle doesn't support SQL BINARY(x) or VARBINARY(x). Binary data must be always stored as variable-length raw data using the Oracle type RAW. Thus, Oracle RAW can be used to store IDL "octet", "octet[x]" and "sequence<octet,x>". The ".bin" suffix allows the daemon to determine which is actually stored for Oracle databases only.
- The ".str" suffix is used to distinguish between "(w)string<x>" and "sequence<(w)char,x>".

f. The “.ld” suffix is used to distinguish between “octet[x]” and “long double”.

Table 5.4 Basic Types in IDL and SQL (MySQL)

IDL Type	IDL Field Name	SQL Type	Table Column Name
char ^a	my_field	CHAR(1)	“my_field”
char[x] ^a	my_field	CHAR(x)	“my_field”
sequence<char,x>	my_field	VARCHAR(x)	“my_field”
wchar ^b	my_field	NCHAR(1)	“my_field”
wchar[x] ^b	my_field	NCHAR(x)	“my_field”
sequence<wchar,x>	my_field	NVARCHAR(x)	“my_field”
octet ^c	my_field	BINARY(1)	“my_field”
octet[x] ^c	my_field	BINARY(x)	“my_field”
sequence<octet,x> ^c	my_field	VARBINARY(x)	“my_field”
boolean	my_field	TINYINT	“my_field”
short	my_field	SMALLINT	“my_field”
unsigned short	my_field	SMALLINT	“my_field”
unsigned	my_field	INTEGER	“my_field”
unsigned long	my_field	INTEGER	“my_field”
double	my_field	DOUBLE	“my_field”
float	my_field	FLOAT	“my_field”
string<x>	my_field	VARCHAR(x)	“my_field.str” ^d
wstring<x>	my_field	NVARCHAR(x)	“my_field.str” ^d
long long	my_field	BIGINT	“my_field”
unsigned long long	my_field	BIGINT	“my_field”
long double	my_field	BINARY(16)	“my_field.ld” ^e
unsigned long long	my_field	DATETIME	“my_field”

a. The format on the wire of “char” and “char[x]” is the same.

b. The format on the wire of “wchar” and “wchar[x]” is the same.

c. The format on the wire of “octet” and “octet[x]” is the same.

d. The “.str” suffix is used to distinguish between “(w)string<x>” and “sequence<(w)char,x>”.

e. The “.ld” suffix is used to distinguish between “octet[x]” and “long double”.

Table 5.5 Basic Types in IDL and SQL (SQL Server)

IDL Type	IDL Field Name	SQL Type	Table Column Name
char	my_field	[char](1)	[my_field]
char[x]	my_field	[char](x)	[my_field]
sequence<char,x>	my_field	[varchar](x)	[my_field]
wchar	my_field	[nchar](1)	[my_field]
wchar[x]	my_field	[nchar](x)	[my_field]
sequence<wchar,x>	my_field	[nvarchar](x)	[my_field]
octet	my_field	[binary](1)	[my_field]
octet[x]	my_field	[binary](x)	[my_field]
sequence<octet,x>	my_field	[varbinary](x)	[my_field]

Table 5.5 Basic Types in IDL and SQL (SQL Server)

IDL Type	IDL Field Name	SQL Type	Table Column Name
boolean	my_field	[tinyint]	[my_field]
short	my_field	[smallint]	[my_field]
unsigned short	my_field	[smallint]	[my_field]
long	my_field	[int]	[my_field]
unsigned long	my_field	[int]	[my_field]
double	my_field	[float](53)	[my_field]
float	my_field	[real]	[my_field]
string<x>	my_field	[varchar](x)	[my_field\$str] ^a
wstring<x>	my_field	[nvarchar](x)	[my_field\$str] ^a
long long	my_field	[bigint]	[my_field]
unsigned long long	my_field	[bigint]	[my_field]
long double	my_field	[binary](16)	[my_field\$id]
unsigned long long ^b	my_field	[datetime]	[my_field]

a. The ".str" suffix is used to distinguish between "(w)string<x>" and "sequence<(w)char,x>".

b. The [datetime] type is converted to/from a 64-bit value expressed as microseconds in the Gregorian epoch.

5.2.3 Oracle In-Memory Database Cache Mapping

When *Database Integration Service* to TimesTen is used with Oracle In-Memory Database Cache, the definition of cache groups requires mapping the Oracle SQL types to TimesTen SQL types. [Table 5.6](#) describes the mapping supported by *Database Integration Service*.

Table 5.6 Mapping Between Oracle and TimesTen types

Oracle Type	TimesTen Type (TypeMode 0)
CHAR(x)	CHAR(x)
VARCHAR2(x)	VARCHAR2(x)
NCHAR(x)	NCHAR(x)
NVARCHAR2(x)	NVARCHAR2(x)
RAW(x)	VARBINARY(x)
NUMBER(3)	TT_TINYINT
NUMBER(5)	TT_SMALLINT
NUMBER(10)	TT_INTEGER
BINARY_DOUBLE	BINARY_DOUBLE
BINARY_FLOAT	BINARY_FLOAT
NUMBER(20)	TT_BIGINT
TIMESTAMP	This mapping is not supported

5.2.4 Bit Field Mapping

IDL bit-field type is an RTI extension that maps directly to C/C++ bit fields but are stored in primitive types in Java with only the specified number of bits being significant. When mapped to SQL, a full primitive SQL type is used to store the value, but only a subset of the bits are sig-

nificant. A suffix must be added to the column name to indicate to the *Database Integration Service* Daemon which bits to serialize when translating the table data into an IDL structure.

The following tables show the mapping of bit fields between IDL and SQL:

- ❑ [Table 5.7, “Bit Fields in IDL and SQL \(TimesTen\)”](#)
- ❑ [Table 5.8, “Bit Fields in IDL and SQL \(Oracle\)”](#)
- ❑ [Table 5.9, “Bit Fields in IDL and SQL \(MySQL\)”](#)
- ❑ [Table 5.10, “Bit Fields in IDL and SQL \(SQL Server\)”](#)

Table 5.7 **Bit Fields in IDL and SQL (TimesTen)**

IDL Type	IDL Field Name	SQL Type (TypeMode 0)	SQL Type (TypeMode 1)	Table Column Name ^a
char	my_field:x	TT_CHAR(1) or CHAR(1)	CHAR(1)	“my_field:x”
wchar	my_field:x	TT_NCHAR(1) or NCHAR(1)	NCHAR(1)	“my_field:x”
octet	my_field:x	BINARY(1)	BINARY(1)	“my_field:x”
short	my_field:x	TT_SMALLINT	SMALLINT	“my_field:x”
unsigned short	my_field:x	TT_SMALLINT	SMALLINT	“my_field:x”
long	my_field:x	TT_INTEGER	INTEGER	“my_field:x”
unsigned long	my_field:x	TT_INTEGER	INTEGER	“my_field:x”

a. The column storing the last bit field in a set of bits will use name of “my_field!:x”.

Table 5.8 **Bit Fields in IDL and SQL (Oracle)**

IDL Type	IDL Field Name	SQL Type	Table Column Name ^a
char	my_field:x	CHAR(1)	“my_field:x”
wchar	my_field:x	NCHAR(1)	“my_field:x”
octet	my_field:x	BINARY(1)	“my_field:x”
short	my_field:x	NUMBER(5)	“my_field:x”
unsigned short	my_field:x	NUMBER(5)	“my_field:x”
long	my_field:x	NUMBER(10)	“my_field:x”
unsigned long	my_field:x	NUMBER(10)	“my_field:x”

a. The column storing the last bit field in a set of bits will use name of “my_field!:x”.

Table 5.9 **Bit Fields in IDL and SQL (MySQL)**

IDL Type	IDL Field Name	SQL Type	Table Column Name ^a
char	my_field:x	CHAR(1)	“my_field:x”
wchar	my_field:x	NCHAR(1)	“my_field:x”
octet	my_field:x	BINARY(1)	“my_field:x”
short	my_field:x	SMALLINT	“my_field:x”
unsigned short	my_field:x	SMALLINT	“my_field:x”

Table 5.9 Bit Fields in IDL and SQL (MySQL) (Continued)

IDL Type	IDL Field Name	SQL Type	Table Column Name ^a
long	my_field:x	INTEGER	"my_field:x"
unsigned long	my_field:x	INTEGER	"my_field:x"

a. The column storing the last bit field in a set of bits will use name of "my_field!x".

Table 5.10 Bit Fields in IDL and SQL (SQL Server)

IDL Type	IDL Field Name	SQL Type	Table Column Name ^a
char	my_field:x	[char](1)	"my_field:x"
wchar	my_field:x	[nchar](1)	"my_field:x"
octet	my_field:x	[binary](1)	"my_field:x"
short	my_field:x	[smallint]	"my_field:x"
unsigned short	my_field:x	[smallint]	"my_field:x"
long	my_field:x	[int]	"my_field:x"
unsigned long	my_field:x	[int]	"my_field:x"

a. The column storing the last bit field in a set of bits will use name of "my_field!x".

5.2.5 Enum Types Mapping

IDL enumeration fields are mapped to columns of type SQL INTEGER. No special naming is required.

5.2.6 Simple IDL Structures

Simple IDL structures containing only basic or primitive types directly map to SQL schemas with fields in the structure becoming columns in the table. Table 5.11 shows the mapping of a simple structure between IDL and SQL

Table 5.11 Simple Structures in IDL and SQL

IDL Types	SQL Table Schema
<pre> struct MyStruct { long my_key_field; //@key^a short my_short_field; }; </pre>	<pre> CREATE TABLE "MyStructContainer" ("my_key_field" INTEGER NOT NULL, "my_short_field" SMALLINT NOT NULL, PRIMARY KEY(my_key_field)); </pre>

a. IDL fields marked as keys are mapped to the primary keys of SQL tables.

5.2.7 Complex IDL Structures

IDL structures that contain more complex fields, fields that are structures, unions, or sequences and arrays of types other than "octet", "char" or "wchar" are mapped to SQL tables by flattening the embedded structures so that their fields are all at the top (and only) level.

Structure fields

Elements of embedded structures map into individual table columns with names that are hierarchically composed from the name of the field in the embedded structure and the name of the embedded structure field itself. This naming convention is not required for serialization to work

properly. Just as long as the column types map to the types of the embedded structure, then the *Database Integration Service* Daemon will properly handle the data irrelevant of the actual column name.

Table 5.12 shows the mapping of a complex structure between IDL and SQL.

Table 5.12 Nested Structures in IDL and SQL

IDL Type	SQL Table Schema
<pre> struct MyStruct { short my_short_field; long my_long_field; }; struct MyStructContainer { long my_key_field; //@key MyStruct my_struct_field; }; </pre>	<pre> CREATE TABLE "MyStructContainer" ("my_key_field" INTEGER NOT NULL, "my_struct_field.my_short_field" SMALLINT NOT NULL, "my_struct_field.my_long_field" INTEGER NOT NULL, PRIMARY KEY(my_key_field)); </pre>

Union fields

IDL unions are mapped by adding an extra column with the name “_d” to represent the discriminator that is used to indicate which type is actually stored by the union. These unions are also known as “switched unions”. All of the individual union fields are mapped to corresponding columns in a table. However, only one of these columns will contain valid data as indicated by the discriminator column, “_d”.

If the *Database Integration Service* Daemon creates the table from an IDL containing an union, it will generate the data columns with hierarchical names from the name of the union field and the name of the union itself. In addition, the values of the switch/case statement in the IDL union are encoded into the names of the data columns as well, e.g., “.c(0,1).”, “.c(2).”, “.def.”.

This naming convention **is required** for the proper serialization and deserialization of unions. The *Database Integration Service* Daemon uses the name of the fields when processing an IDL union to know which column(s) correspond to the value of the discriminator.

Table 5.13 shows the mapping of an union between IDL and SQL.

Table 5.13 Union Fields in IDL and SQL

IDL Type	SQL Table Schema
<pre> union MyUnion switch(long) { case 0: case 1: long my_long_field; case 2: double my_double_field; default: short my_short_field; }; struct MyUnionContainer { long my_key_field; //@key MyUnion my_union_field; }; </pre>	<pre> CREATE TABLE "MyUnionContainer" ("my_key_field" INTEGER NOT NULL, "my_union_field._d" INTEGER NOT NULL, "my_union_field.c(0,1).my_long_field" INTEGER, "my_union_field.c(2).my_double_field" DOUBLE, "my_union_field.c(def).my_short_field" SMALLINT, PRIMARY KEY(my_key_field)); </pre>

5.2.8 Array Fields

For array fields where the array type is different from “octet”, “char” and “wchar”, an IDL array type is stored as consecutive columns of the same type in a SQL table. If the *Database Integration*

Service Daemon creates a table from an IDL type that contains an array, it will create the column names using a naming convention that prevents name collisions. By default, the daemon simply adds the suffix “[i]”, where “i” is the array index of that element (beginning at 0 for the first index). The open bracket and close bracket characters can be configured using the tags in the configuration file `<open_bracket_char>` and `<close_bracket_char>` (see [page 4-19](#)). Note, this naming convention is not required for the *Database Integration Service Daemon* to serialize/deserialize IDL array fields.

Note that array fields of type “octet”, “char” and “wchar” are mapped into a single column element of the corresponding SQL types BINARY(x), CHAR(x) and WCHAR(x), respectively. [Table 5.14](#) shows a mapping of an array field between IDL and SQL.

Table 5.14 **Array Fields in IDL and SQL**

IDL Type	SQL Table Schema
<pre>struct MyArrayContainer { long my_key_field; //@key short my_arr_field[2]; };</pre>	<pre>CREATE TABLE "MyArrayContainer" ("my_key_field" INTEGER NOT NULL, "my_arr_field[0]" SMALLINT NOT NULL, "my_arr_field[1]" SMALLINT NOT NULL, PRIMARY KEY("my_key_field"));</pre>

5.2.9 Sequence Fields

Sequences are basically variable-sized arrays that have a maximum length and carry an additional integer that indicates the current size. The mapping of IDL sequences to a table schema is similar to the array mapping, with the following differences:

- ❑ An extra column is added with the suffix “#length”, used to store the current length of the sequence.
- ❑ The total number of columns created is equal to the maximum number of elements that the sequence can hold, although the number of columns containing valid data at a given time is stored in the “#length” column.
- ❑ The naming convention of adding the suffix “[i]” to each column is **required** for the *Database Integration Service Daemon* to handle the mapping between IDL and SQL correctly. The open bracket and close bracket characters can be configured using the tags `<open_bracket_char>` and `<close_bracket_char>` (see [page 4-19](#)).
- ❑ Sequence elements can contain the NULL value since not all elements may be used at a given time.

Note: Sequences of the IDL types “char”, “wchar” or “octet” map directly into the variable-length SQL types VARCHAR, VARWCHAR, and VARBINARY, respectively. [Table 5.15](#) shows the mapping of a sequence field between IDL and SQL.

Table 5.15 **Sequence Fields in IDL and SQL**

IDL Type	SQL Table Schema
<pre>struct MySequenceContainer { long my_key_field; //@key sequence<short,4> my_seq_field; };</pre>	<pre>CREATE TABLE "MySequenceContainer" ("my_key_field" INTEGER NOT NULL, "my_seq_field#length" INTEGER NOT NULL, "my_seq_field[0]" SMALLINT, "my_seq_field[1]" SMALLINT, "my_seq_field[2]" SMALLINT, "my_seq_field[3]" SMALLINT, PRIMARY KEY("my_key_field"));</pre>

5.2.10 NULL Values

Null values exist in SQL databases but do not have an equivalent in IDL. The *Database Integration Service* daemon converts NULL values into '0'-values when publishing from a SQL table, in the following way:

- ❑ numerical types: 0
- ❑ fixed-length string types (CHAR, NCHAR): ""
- ❑ variable-length types (VARCHAR, NVARCHAR, VARBINARY): length 0
- ❑ binary: every byte is set to 0
- ❑ timestamp: 0

5.2.11 Sparse Data Types

Sparse Data Types follow the same mapping as structures (see [Section 5.2.6](#) and [Section 5.2.7](#)). The fields that are not required or primary keys are created with the nullable attribute.

Table 5.16 Simple Sparse Type

Type in Pseudo Language	SQL Table Schema
<pre>sparse MySparse ^a { long my_key_field; //@key short my_short_field; long my_long_field; //@required };</pre>	<pre>Create Table "MySparseContainer" ("my_key_field" INTEGER NOT NULL, "my_short_field" SMALLINT, "my_long_field" INTEGER NOT NULL, PRIMARY_KEY("my_key_field"));</pre>

a. Sparse types must be built dynamically. There is no IDL construct sparse.

Appendix A Error Codes

Table A.1 lists the native error and warning messages that may be logged by the *Database Integration Service* Daemon. While some of these messages may actually provide enough information by themselves to help users fix the problem, many have to be used along with other data to help with debugging the issue.

Often several of these messages will be logged for a single problem. A failure at a lower layer will cause log messages to be printed at various levels of the *Database Integration Service* Daemon logic. These messages will be valuable to you and to RTI support engineers in debugging issues with *Database Integration Service*.

Table A.1 Database Integration Service Errors and Warnings

Code	Message	Details
0 - 1023 Database Integration Service Daemon errors These messages are produced by the logic of the <i>Database Integration Service</i> Daemon itself.		
0	Unexpected error	Should never occur. Contact support@rti.com if seen.
1	<message>	General error.
2	Error storing RTI DDS sample in table '<table>'	There was an error when storing value received with <i>Connex</i> DDS into the database.
3	Error creating <entity>	
4	Error creating <entity> associated to the table '<table>'	
5	Error getting <entity>	
6	<meta-table> entry not valid	There was an entry in a meta-table (RTIDDS_PUBLICATIONS or RTIRTCSUBSCRIPTIONS) that was not valid.
7	Error creating '<type>' SQL statement	There was an error when creating or preparing a SQL statement.
8	Error creating table '<table>'	
9	Error opening RTI DDS connection	There was a problem initializing <i>Connex</i> DDS.

Table A.1 Database Integration Service Errors and Warnings

Code	Message	Details
10	The type of the column '<column>' is not valid	<p>The meta-columns RTIDDS_DOMAIN_ID and RTIRTC_REMOTE must be added to tables that the user creates and wished to connect to via <i>Database Integration Service</i>. They will be automatically if the <i>Database Integration Service</i> Daemon creates the table.</p> <p>If the user creates the table and adds the two columns, they must be of type INTEGER.</p> <p>This message is produced if these columns exist and are of the wrong type.</p>
11	Error publishing record/instance	The <i>Database Integration Service</i> Daemon had a problem publishing a table change as a Topic.
12	Error disposing record/instance	The <i>Database Integration Service</i> Daemon had a problem disposing of an instance of Topic when the user deleted a row in a table.
13	Error initializing <module> module	The <i>Database Integration Service</i> Daemon had problems initializing an internal code module.
14	The definition of environment variable '%s' is required	
15	Error opening the database connection associated to the DSN '<DSN>'	
16	Error enabling database log	
17	<string> too long (maximum length: <length>)	
18	Error creating connection to database log	
19	The value of the column 'column' in the table '<meta-table>' is not valid	
20	Error generating '<type>' SQL statement string Error generating SQL statement string	The <i>Database Integration Service</i> Daemon had a problem in generating the string for preparing or executing a SQL statement.
21	Error skipping parameter for the field '<column>'	
22	Error binding parameter for the column '<column>' Error binding parameters	

Table A.1 Database Integration Service Errors and Warnings

Code	Message	Details
23	The column '<column>' has an unexpected SQL Type	Supported SQL types are: SQL_CHAR SQL_WCHAR SQL_VARCHAR SQL_WVARCHAR SQL_BINARY SQL_VARBINARY SQL_INTEGER SQL_SMALLINT SQL_TINYINT SQL_BIGINT SQL_REAL SQL_FLOAT SQL_DOUBLE SQL_TIMESTAMP
24	Error opening configuration file 'filename'	
25	Error reading configuration file 'filename'	
26	Error parsing configuration file 'filename'	
27	The maximum length for a <type> field is <length>	
28	Error serializing record	A problem occurred when serializing a table row for publishing as a Topic.
29	Error deserializing RTI DDS sample	A problem occurred when deserializing data received via <i>Connex</i> DDS for storing into a table.
30	Error creating key cache	Could not create cache of known instance keys see cache_maximum_size , cache_initial_size (Section 4.5.2.1.29).
31	Error inserting key in cache	Problem occurred while storing instance key in cache see cache_maximum_size , cache_initial_size (Section 4.5.2.1.29).
32	Null pointer argument	A precondition failed in which a NULL pointer was passed to an internal daemon function.
33	Error reading table '<table>'	
34	The resolution column '<column>' does not exist in the table '<table>'	The RTIDDS_PUBLICATIONS table contained an entry in the 'resolution_column' which does not match the name of an existing column in the corresponding table. See resolution_column (Section 4.5.1.1.9).
35	The type of the resolution column '<column>' in the table '<table>' is not valid. The type of the resolution column can be: SQL_INTEGER, SQL_SMALLINT, SQL_BIGINT and SQL_TIMESTAMP	A column specified in the column 'resolution_column' in the RTIDDS_PUBLICATIONS table is not of an acceptable type. See resolution_column (Section 4.5.1.1.9).
36	Error gathering instance information	Error gathering <i>Connex</i> DDS instance information through the execution of the associated SELECT statement
37	Invalid metatable schema	The schema of the metatables is not valid. It is possible that those tables were created with a previous version of RTI RTC.

Table A.1 Database Integration Service Errors and Warnings

Code	Message	Details
38	Error deleting key from cache	
39	Error deleting a row from '<table>'	There was a problem deleting a row from a user data table.
41	Error creating publication/subscription for the '<table>' without primary key.	The user tried to create a publication/subscription for a table without a primary key. <i>Database Integration Service to Oracle</i> does not support tables without primary keys.
43	Key not supported	Non-primitive IDL keys are not supported. When <i>Database Integration Service</i> tries to create a table with complex keys, it will report this error message. Example with supported keys: <pre>struct SupportedKeysSt { string id_str; //@key long id_long; //@key short id_short; //@key };</pre> Example with unsupported keys: <pre>struct KeySt { long id_long; } struct NonSupportedKeysSt { KeySt id_st; //@key }</pre>
44	Error creating subscriber state queue	
45	Error updating subscription state	
46	Error creating '<object>'	
47	Path too long	The path to the configuration file is too long.
48	Error creating database publication cache	The <i>Database Integration Service</i> Daemon had a problem creating the publication database cache.
49	Error adding record to publication cache	The <i>Database Integration Service</i> Daemon had a problem adding a new record to the publication database cache.
50	Column name length exceeds maximum length	The maximum length of a column name in <i>Database Integration Service</i> is 30 characters. To control the length of a column name, use the configuration tags <code><idl_member_prefix_max_length></code> and <code><idl_member_suffix_max_length></code> under <code><database_mapping_options></code> . See Section 4.4.4.2 for additional details.

Table A.1 Database Integration Service Errors and Warnings

Code	Message	Details
1024 - 2047 Connex DDS-related errors		
These messages are produced through the interaction of the <i>Database Integration Service</i> Daemon with <i>Connex DDS</i> . More information on each error can be found by examining the native <i>Connex DDS</i> errors codes that will be logged with these messages.		
1024	<message>	General <i>Connex DDS</i> error message.
1025	Error getting <entity> default QoS	
1026	Error getting <entity> QoS	
1027	Error setting <entity> QoS	
1028	Error creating <entity>	
1029	Error getting <entity>	
1030	Error enabling <entity>	
1031	Error cloning type code	
1032	Error reading RTI DDS samples	
1033	Error setting <entity> user data	
1034	Error disposing RTI DDS instance	
1035	Error unregistering RTI DDS instance	
1036	Error writing RTI DDS sample	
1037	Error ignoring <entity>	
1038	Error creating <waitset type> waitset	
1039	Error waiting in <waitset type> waitset	
1040	Error getting builtin transport property	
1041	Error setting builtin transport property	
1042	Error creating <waitset type> guard condition	
1043	Error attaching condition	
1044	Error registering type '<type>'	
1045	Error taking REDA buffer	
1046	Error creating mutex	
1047	Error creating <thread> thread	
1048	Error creating REDA fast buffer	
1049	Error taking semaphore	
1050	Error giving semaphore	
1051	Error creating worker factory	
1052	Error creating worker	
1053	Error creating clock	
1054	Error creating event manager	
1055	Error creating timer	
1056	Error posting event	
1057	Error getting time	
1058	Error creating semaphore	
1059	Error loading DDS XML Profile	
1060	Error getting TypeCode	

Table A.1 Database Integration Service Errors and Warnings

Code	Message	Details
1061	Error cloning TypeCode	
1062	Error parsing TypeCode	
1063	Error creating TypeCode	
2048 - 4095 ODBC-related errors		
These message are produced through the interaction of the <i>Database Integration Service</i> Daemon with the database through the ODBC driver. More information on each error can be found by examining the native ODBC errors codes that will be logged with these messages.		
2048	<message> <message>: <ODBC driver error message>	General ODBC error message.
4096 - 8191 DBMS Log Connection-related errors		
These messages are produced through the interaction of the <i>Database Integration Service</i> Daemon with the <i>Connext DDS</i> . More information on each error can be found by examining the native <i>Connext DDS</i> errors codes that will be logged with these messages.		
4096	<message>	General DBMS log connection error message.
8192 - 16383 OS-related errors		
These messages are produced through the interaction of the <i>Database Integration Service</i> Daemon with the operating system. More information on each error can be found by examining the native OS errors codes that will be logged with these messages.		
8192	<message>	General OS error message.
8193	Error handling OS signals	
8194	Unable to set signal handler for <signal>	
8195	Error getting the host name	
8196	Error allocating memory for <object>	
From 16384 Warning messages		
These are warning messages that may be logged by the <i>Database Integration Service</i> Daemon.		
16384	Timestamps prior to '1970-01-01 00:00:00.00' cannot be used for conflict resolution. The RTC daemon will always use '1970-01-01 00:00:00,00' as the timestamp for those cases	
16385	The Timestamp value of the resolution column is NULL. The RTC daemon will use the value '1970-01-01 00:00:00,00'.	
16386	Diskless log buffer overflow	This warning appears only with Oracle TimesTen databases.
16387	IDL member identifier collision	The prefix/suffix-based name associated with member <i>A</i> in IDL type <i>T</i> collides with the name of another member inside the same type. <i>Database Integration Service</i> will resolve the conflict using an index.
16388	Invalid configuration parameter	
16389	Ignored QoS value	A QoS value has been ignored by <i>Database Integration Service</i> .

Table A.1 Database Integration Service Errors and Warnings

Code	Message	Details
16392	Dynamic loading of monitoring library is not supported	<i>RTI Monitoring Library</i> is statically linked. In <i>Database Integration Service</i> there is no need to load this library dynamically.
32769	Requested incompatible QoS	The QoS of a <i>Database Integration Service</i> subscription is incompatible with the QoS of a <i>Connex DDS</i> publication. The name of the policy that is incompatible is shown in this warning message.
32770	Offered incompatible QoS	The QoS of a <i>Database Integration Service</i> publication is incompatible with the QoS of a <i>Connex DDS</i> subscription. The name of the policy that is incompatible is shown in this warning message.
32771	Sample lost message	A <i>Database Integration Service</i> subscription lost a sample.
32772	Sample rejected message	A <i>Database Integration Service</i> subscription rejected a sample.

Appendix B Database Limits

The maximum number of columns is limited by the underlying database product. The maximum length of a column is independent of the database and it is limited to 30 characters. [Table B.1](#) notes the database limits of *Database Integration Service*.

Table B.1 **Real-Time Connect Database Limits**

	Oracle 11g	Oracle TimesTen 11.2.1	MySQL 5.1	Microsoft SQL Server 2012
Maximum number of columns	For subscriptions: 1000 For publications: 625 - 650 See Section B.1	1,000	4096, although the effective maximum may be considerably smaller, see Section B.2 .	1,024
Maximum column-name length (characters)	30	30	30 ^a	30 ^b
CLOB/BLOB support	YES	Not supported by database	Not supported by RTC (maximum record size is 65535 bytes)	Maximum record size is 8,000 bytes ^c
CHAR maximum size (bytes)	2000	8300	255 ^d	8,000
VARCHAR maximum size (bytes)	4000	4194304	65535 ^d	8,000 ^c
BINARY maximum size (bytes)	2000	8300	255 ^d	8,000
VARBINARY maximum size (bytes)	2000	4194304	65535 ^d	8,000 ^c

- a. Limit is imposed by *Real-Time Connect* (not MySQL, which allows column names of up to 64 characters).
- b. Limit is imposed by *Real-Time Connect* (not SQL Server, which allows column names of up to 128 characters).
- c. [varbinary](max) and [varchar](max) are not supported by *Real-Time Connect*.
- d. The maximum size of a row in MySQL 5.1 is limited to 65535. For example, you cannot have two fields of type VARCHAR(40000) because the total width of the columns would exceed 65535 bytes. For additional information on this restriction, see <http://dev.mysql.com/doc/refman/5.1/en/column-count-limit.html>.

B.1 Maximum Columns for Oracle 11g

For *Database Integration Service* subscriptions, the maximum number of columns is 1000. This limit is imposed by the maximum number of columns in an Oracle 11g table.

For *Database Integration Service* publications, the maximum number of columns is limited by the maximum size of the PL/SQL programs that *Database Integration Service* installs in the Oracle 11g server. The size of the PL/SQL programs depends on the column data types. For tables in which all the columns are numbers, RTI has been able to publish with approximately 625 columns. For tables in which all the columns are of type VARCHAR2, RTI has been able to publish with approximately 650 columns.

If the maximum number of columns is exceeded for *Database Integration Service* publications, you will see an error message similar to:

```
[DDSQLDaemonCore_onUpdateMetaTableEntry,line 3126:ERROR:4096:5009]
[CNA:CNAMonitorTable] ODBC call failed: ORA-24344: success with
compilation error
```

B.2 Maximum Columns for MySQL

The exact limit for the number of columns is driven by two factors:

- ❑ The maximum row size for a table, not counting BLOBs, is 65535.
- ❑ The maximum size of the meta information (schema) associated with a table is 64 KB. The meta information includes the column names. The longer the column names, the smaller the maximum number of columns.

[Table B.2](#) provides information about the maximum number of columns that RTI was able to use for different column-name lengths and column types.

For more information about MySQL restrictions on the number of columns, see <http://dev.mysql.com/doc/refman/5.1/en/column-count-limit.html>.

Table B.2 Max. Number of Columns Used by RTI for Different Column Types and Name Lengths

Column Type	Column-Name Length (characters)	Maximum Number of Columns
INTEGER	30	1283 ^a
DOUBLE		
CHAR (50)		
VARCHAR (50)		
INTEGER	15	1823 ^a
DOUBLE		
CHAR (50)	15	1308 ^b
VARCHAR (50)	15	1283 ^b

a. Limited by schema size

b. Limited by row size.