

RTI Routing Service

User's Manual

Version 4.5



Your systems. Working as one.



© 2009-2012 Real-Time Innovations, Inc.
All rights reserved.
Printed in U.S.A. First printing.
March 2012.

Trademarks

Real-Time Innovations, RTI, and Connex are trademarks or registered trademarks of Real-Time Innovations, Inc. All other trademarks used in this document are the property of their respective owners.

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. The software described in this document is furnished under and subject to the RTI software license agreement. The software may be used or copied only under the terms of the license agreement.

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/>

Contents

1 Welcome to RTI Routing Service

1.1 Available Documentation	1-3
-----------------------------------	-----

2 Configuring Routing Service

2.1 Terms to Know	2-1
2.2 How to Load the XML Configuration	2-2
2.3 XML Syntax and Validation	2-3
2.4 XML Tags for Configuring Routing Service.....	2-6
2.4.1 Routing Service	2-7
2.4.2 Domain Route	2-10
2.4.3 Administration.....	2-16
2.4.4 Monitoring.....	2-18
2.4.5 Session.....	2-24
2.4.6 Routes.....	2-28
2.4.7 Auto Routes.....	2-41
2.4.8 Adapters.....	2-49
2.5 Enabling and Disabling Routing Service Entities	2-49

3 Running Routing Service

3.1 Starting Routing Service	3-1
3.2 Stopping Routing Service	3-2

4 Transforming Data with Routing Service

4.1 Transformation Usage and Configuration	4-1
4.2 Transformations Distributed with Routing Service.....	4-4
4.3 Creating New Transformations	4-5
4.3.1 Transformation Plugin API.....	4-6

5 Administering Routing Service from a Remote Location

5.1 Enabling Remote Administration.....	5-1
---	-----

5.2 Remote Commands	5-2
5.2.1 add_peer.....	5-4
5.2.2 create.....	5-4
5.2.3 delete.....	5-5
5.2.4 disable.....	5-5
5.2.5 enable.....	5-6
5.2.6 load	5-6
5.2.7 pause.....	5-6
5.2.8 resume	5-7
5.2.9 save	5-7
5.2.10 unload.....	5-7
5.2.11 update.....	5-7
5.3 Accessing Routing Service from a Connex Application	5-10

6 Monitoring Routing Service from a Remote Location

6.1 Enabling Remote Monitoring.....	6-2
6.2 Monitoring Configuration Data.....	6-2
6.2.1 Configuration Data for the Routing Service	6-3
6.2.2 Configuration Data for a Domain Route.....	6-4
6.2.3 Configuration Data for a Session.....	6-7
6.2.4 Configuration Data for a Route	6-8
6.2.5 Configuration Data for an Auto Topic Route	6-14
6.3 Monitoring Status	6-19
6.3.1 How the Statistics are Generated	6-20
6.3.2 Status Information for the Routing Service.....	6-22
6.3.3 Domain Route Status.....	6-23
6.3.4 Status Information for a Session	6-24
6.3.5 Status Information for a Route.....	6-26
6.3.6 Status Information for an Auto Route	6-27

7 Traversing Wide Area Networks

7.1 TCP Communication Scenarios	7-2
7.1.1 Communication Within a Single LAN.....	7-2
7.1.2 Symmetric Communication Across NATs	7-3
7.1.3 Asymmetric Communication Across NATs.....	7-4
7.1.4 Secure Communication.....	7-4

7.2	Configuring the TCP Transport	7-6
7.2.1	TCP Transport Initial Peers	7-6
7.2.2	Setting Up the TCP Transport Properties with the PropertyQoSPolicy	7-7
7.2.3	TCP/TLS Transport Properties.....	7-9


8 Extending Routing Service with Adapters

8.1	Adapter Usage and Configuration.....	8-2
8.2	Adapter API And Entity Model.....	8-6
8.2.1	Entity Creation	8-12
8.2.2	Stream Discovery	8-13
8.2.3	Reading Data	8-14
8.3	Creating New Adapters	8-15
8.3.1	Adapter SDK Components	8-15
8.3.2	C Adapter API.....	8-16
8.3.3	My First C Adapter	8-20
8.3.4	Debugging C Adapters	8-49
8.3.5	Java Adapter API.....	8-52
8.3.6	My First Java Adapter.....	8-54
8.3.7	Debugging Java Adapters	8-75
8.3.8	Testing an Adapter	8-80

Chapter 1 Welcome to RTI Routing Service

Welcome to *RTI® Routing Service*, an out-of-the-box solution for integrating disparate and geographically dispersed systems. It scales *RTI Connexx™* (formerly *RTI Data Distribution Service*) applications across domains, LANs and WANs, including firewall and NAT traversal. *Routing Service* also supports *Connexx-to-Connexx* bridging by allowing you to make transformations in the data along the way. This allows unmodified *Connexx* applications to communicate even if they were developed using incompatible interface definitions. This is often the case when integrating new and legacy applications or independently developed systems. Using *RTI Routing Service Adapter SDK*, you can extend *Routing Service* to interface with non-*Connexx* systems using off-the-shelf or custom developed adapters, including to third-party JMS implementations and legacy code written to the network socket API.

Traditionally, *Connexx* applications can only communicate with applications in the same domain. With *Routing Service*, you can send and receive data across domains. You can even transform and filter the data along the way! Not only can you change the actual data values, you can change the data's type. So the sending and receiving applications don't even need to use the same data structure. You can also control which data is sent by using allow and deny lists.



Simply set up *Routing Service* to pass data from one domain to another and specify any desired data filtering and transformations. No changes are required in the *Connexx* applications.

Key benefits of *Routing Service*:

- ❑ **It can significantly reduce the time and effort spent integrating and scaling *Connex*t applications across Wide Area Networks and Systems-of-Systems.**

Many systems today already rely on *Connex*t to distribute their information across a Local Area Network (LAN). However, more and more of these systems are being integrated in Wide Area Networks (WANs). With *Routing Service*, you can scale *Connex*t real-time publish/subscribe data-distribution beyond the current local networks and make it available throughout a WAN—without making any changes to existing *Connex*t applications. You can take an existing, even deployed system and integrate it with new applications or other existing systems without changing those existing systems.

- ❑ **With *Routing Service*, you can build modular systems out of existing systems.**

Data can be contained in private domains within subsystems and you can designate that only certain “global topics” can be seen across domains. The same mechanism controls the scope of discovery. Both application-level and discovery traffic can be scoped, facilitating scalable designs.

- ❑ ***Routing Service* provides secure deployment across multiple sites.** You can partition networks and protect them with firewalls and NATS and precisely control the flow of data between the network segments.

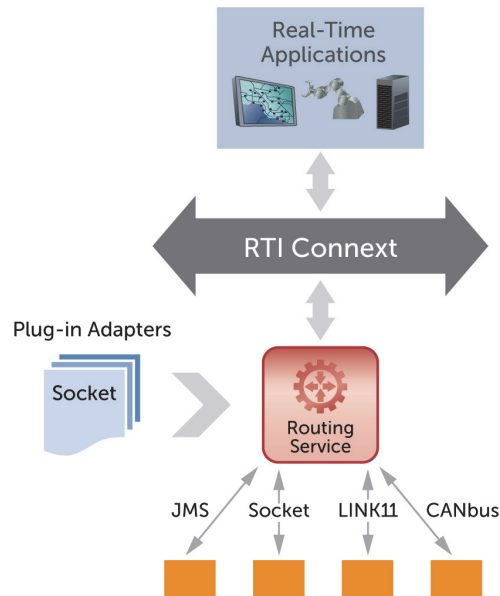
- ❑ **It allows you to manage the evolution of your data model at the subsystem level.** You can use *Routing Service* to transform data on the fly, changing topic names, type definitions, QoS, etc., seamlessly bridging different generations of *Connex*t topic definitions.

- ❑ ***Routing Service* provides features for development, integration and testing.** Multiple sites can each locally test and integrate their core application, expose selected topics of data, and accept data from remote sites to test integration connectivity, topic compatibility and specific use-cases.

- ❑ **It connects remotely to live, deployed systems** so you can perform live data analytics, fault condition analysis, and data verification.

- ❑ ***RTI Routing Service Adapter SDK* allows you to quickly build and deploy bridges to integrate *Connex*t and non-*Connex*t systems.** This can be done in a fraction of the time required to develop completely custom solutions. Bridges automatically inherit advanced *Connex*t capabilities, including automatic discovery of applications; data transformation and filtering; data lifecycle management and support across operating systems; programming languages and network transports.

RTI Routing Service Adapter SDK offers an out-of-the-box solution for interfacing with third-party protocols and technology. It includes prebuilt adapters that can be used out-of-the-box to interface with third-party Java Message Service (JMS) providers or legacy code written to the network socket API. Adapters include source code so they can be easily modified to meet application-specific requirements or serve as a template for quick creation of new custom adapters.



Quickly build and deploy bridges between natively incompatible protocols and technologies using Connex

1.1 Available Documentation

Routing Service documentation includes:

- ❑ **Getting Started Guide** ([RTI_Routing_Service_GettingStarted.pdf](#))—Highlights the benefits of *Routing Service*. It provides installation and startup instructions, and walks you through several examples so you can quickly see the benefits of using *Routing Service*.

- ❑ **Release Notes** (RTI_Routing_Service_ReleaseNotes.pdf)—Describes system requirements and compatibility, as well as any version-specific changes and known issues.
- ❑ **User's Manual** (RTI_Routing_Service_UsersManual.pdf)—Describes how to configure *Routing Service* and use it remotely.

If the optional *RTI Routing Service Adapter SDK* is installed, you will also have the following documents:

- ❑ **RTI Routing Service Adapter SDK Installation Guide** (RTI_Routing_Service_AdapterSDK_InstallationGuide.pdf)—Describes installation instructions for *RTI Routing Service Adapter SDK*.
- ❑ **RTI Routing Service Adapter SDK Release Notes** (RTI_Routing_Service_AdapterSDK_ReleaseNotes.pdf)—Describes system requirements and compatibility, as well as any version-specific changes and known issues for *RTI Routing Service Adapter SDK*.

Chapter 2 Configuring Routing Service

This document describes how to configure *Routing Service*. To see installation instructions, or to walk through some simple examples, please see the [Getting Started Guide](#).

When you start *Routing Service*, you can specify a configuration file in XML format (it is not required). In that file, you can set properties that control the behavior of the service. This chapter describes how to write a configuration file.

This chapter describes:

- ❑ [Terms to Know \(Section 2.1\)](#)
- ❑ [How to Load the XML Configuration \(Section 2.2\)](#)
- ❑ [XML Syntax and Validation \(Section 2.3\)](#)
- ❑ [XML Tags for Configuring Routing Service \(Section 2.4\)](#)

2.1 Terms to Know

Before learning how to configure *Routing Service*, you should become familiar with a few key terms and concepts.

- ❑ A *routing service* entity refers to an execution of *Routing Service*.
- ❑ A *domain route* defines a two-way mapping between two data domains. For example, a domain route could define a mapping between two different domains or between a *Connex*t domain and a JMS provider's network.
- ❑ A *session* defines a single-threaded context for routes. Data cannot be read and written from two routes in the session concurrently.

- ❑ A *route* defines a one-way mapping between an “input” stream in one domain and an “output” stream in the other domain. For example, in a route between *Connex*t and JMS, the input stream will be a *Connex*t topic and the output stream will be a JMS topic or queue.
- ❑ An *auto route* defines a set of potential routes that can be instantiated based on deny/allow filters on the stream name and registered type name.
- ❑ A *transformation* is a pluggable component that changes data from the “input” stream A to data in the “output” stream B.
- ❑ An *adapter* is a pluggable component that allows *Routing Service* to consume and produce data for different data domains. By default, *Routing Service* is distributed with a built-in *Connex*t adapter.

2.2 How to Load the XML Configuration

Routing 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*t (see Chapter 15 in the *RTI Core Libraries and Utilities User’s Manual*¹).

- ❑ **`$NDDSHOME/resource/qos_profiles_4.5x2/xml/NDDS_QOS_PROFILES.xml`**

This file contains the *Connex*t 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.

1. See `<Connex`t installation directory`>/ndds.<version>/doc/pdf/RTI_Connex`t_UsersManual.pdf.
2. *x* stands for the version letter of the current release.

The next locations are specific to *Routing Service*.

- ❑ `<Routing Service executable location>/../resource/xml/RTI_ROUTING_SERVICE.xml`

This file contains the default *Routing Service* configuration; it is loaded if it exists. **RTI_ROUTING_SERVICE.xml** defines a service that automatically routes all types and topics between domains 0 and 1.

- ❑ `<working directory>/USER_ROUTING_SERVICE.xml`

This file is loaded automatically if it exists.

- ❑ **File specified using the command line parameter `-cfgFile`**

The command-line option `-cfgFile` (see [Table 3.1 on page 3-2](#)) can be used to specify a configuration file.

- ❑ **File specified using the remote command 'load'**

The **load** command (see [Section 5.2.6](#)) allows loading an XML file remotely. The file loaded using this command replaces the file loaded using the `-cfgFile` command-line option. (*Last to be loaded.*)

You may use a combination of the above approaches.

[Figure 2.1](#) shows an example configuration file. You will learn the meaning of each line as you read the rest of this chapter.

2.3 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. *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>`.

Figure 2.1 Example XML Configuration File

```

<?xml version="1.0"?>
<dds>
  <routing_service name="TopicBridgeExample" group_name="MyGroup">
    <domain_route name="DomainRoute">
      <participant_1>
        <domain_id>0</domain_id>
      </participant_1>

      <participant_2>
        <domain_id>1</domain_id>
      </participant_2>

      <session name="Session">
        <topic_route name="SquaresToCircles">

          <input participant="1">
            <registered_type_name>ShapeType</registered_type_name>
            <topic_name>Square</topic_name>
          </input>

          <output>
            <registered_type_name>ShapeType</registered_type_name>
            <topic_name>Circle</topic_name>
          </output>

        </topic_route>
      </session>
    </domain_route>
  </routing_service>
</dds>

```

This file configures a simple bridge from Connex domain 0 to Connex domain 1 and changes the data's topic from Square to Circle. Both topics use the same data type (ShapeType). You will find this example in <Routing Service installation directory>/example/shapes/topic_bridge.xml. Additional examples are in the same directory.

Routing 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 routine service's configuration—this provides helpful features in code

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.

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 *<Routing Service installation directory>/resource/schema/rti_routing_service.dtd* and *<Routing Service installation directory>/resource/schema/rti_routing_service.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="<installation directory for RTI Routing
Service>/resource/schema/rti_routing_service.xsd">
  ...
</dds>
```

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

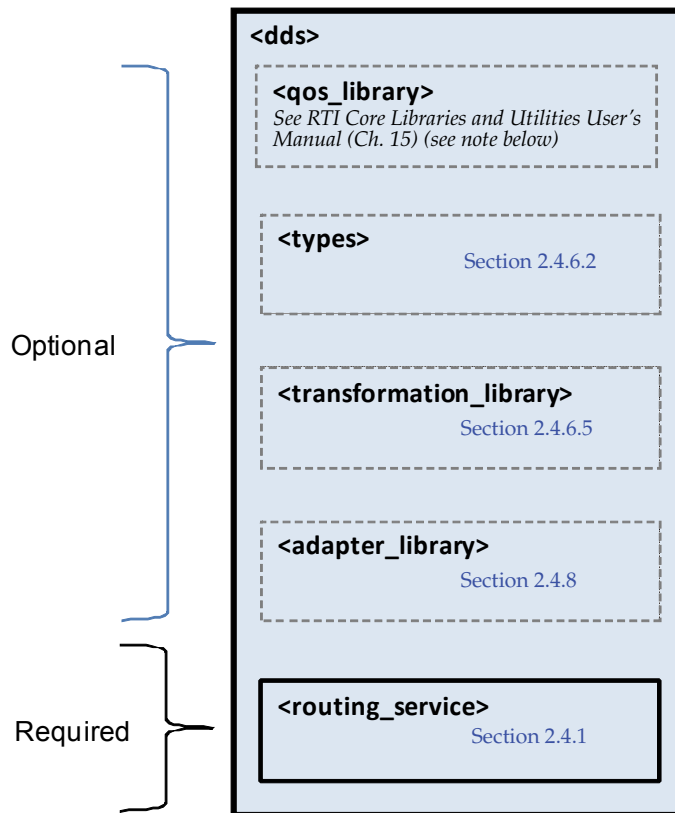
For example:

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE dds SYSTEM "<installation directory for RTI Routing Service>
/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.

2.4 XML Tags for Configuring Routing Service

This section describes the XML tags you can use in a *Routing Service* configuration file. The following diagram and [Table 2.1](#) describe the top-level tags allowed within the root `<dds>` tag.



Note: The *RTI Core Libraries and Utilities User's Manual* is located in `<Connex installation directory>/ndds.<version>/doc/pdf/RTI_Connext_UsersManual.pdf`.

Table 2.1 Top-level Tags in the Configuration File

Tags within <dds>	Description	Number of Tags Allowed
<routing_service>	Specifies a <i>Routing Service</i> configuration. See Routing Service (Section 2.4.1) .	1 or more (required)
<qos_library>	Specifies a QoS library and profiles. The contents of this tag are specified in the same manner as for a <i>Connex</i> QoS profile file—see Chapter 15 in the <i>RTI Core Libraries and Utilities User's Manual</i> . ¹	0 or more
<transformation_library>	Specifies a library of transformation plugins. See Data Transformation (Section 2.4.6.5) and Chapter 4: Transforming Data with Routing Service .	0 or more
<adapter_library>	Specifies a library of adapter plugins. See Adapters (Section 2.4.8) and Chapter 8: Extending Routing Service with Adapters .	0 or more
<types>	Defines types that can be used by the routing service. See Defining Types in the Configuration File (Section 2.4.6.2) .	0 or 1

1. See <Connex installation directory>/nlds.<version>/doc/pdf/RTI_Connex_UsersManual.pdf.

2.4.1 Routing Service

A configuration file must have at least one <routing_service> tag; this tag is used to configure an execution of *Routing Service*. A configuration file may contain multiple <routing_service> tags.

When you start *Routing Service*, you can specify which <routing_service> tag to use to configure the service using the **-cfgName** command-line parameter.

For example:

```
<dds>
  <routing_service name="Router1" group_name="Group1">
    ...
  </routing_service>

  <routing_service name="Router2" group_name="Group1">
    ...
  </routing_service>
</dds>
```

Starting *Routing Service* with the following command will use the `<routing_service>` tag with the name **Router1**:

```
rtiroutingservice -cfgFile example.xml -cfgName Router1
```

Because a configuration file may contain multiple `<routing_service>` tags, one file can be used to configure multiple *Routing Service* executions.

A *routing service* may belong to a group of several routing services identified by a common **group_name**. This common name can be used to implement a specific policy when the communication happens between routing services of the same group. For example, in the built-in *Connex* adapter, a participant will ignore other participants in the same group, as a way to avoid circular communication.

If the `<routing_service>` tag does not have a **group_name** attribute, *Routing Service* will use the following name: **RTI_RoutingService_<Host Name>_<Process ID>**, such as **RTI_RoutingService_myhost_20024**.




Table 2.2 describes the tags allowed within a `<routing_service>` tag.

Table 2.2 Routing Service Tags

Tags within <code><routing_service></code>	Description	Number of Tags Allowed
<code><annotation></code>	Contains a <code><documentation></code> tag that can be used to provide a routing service description. This description will show up when you run <i>Routing Service</i> <u>without</u> the <code>-cfgName</code> command-line option.	0 or 1
<code><domain_route></code>	Defines a mapping between two data domains. See Section 2.4.2.	1 or more (required)
<code><administration></code>	Enables and configures remote administration. See Administration (Section 2.4.3) and Chapter 5: Administering Routing Service from a Remote Location.	0 or 1

Table 2.2 Routing Service Tags

Tags within <routing_service>	Description	Number of Tags Allowed
<monitoring>	Enables and configures general remote monitoring. General monitoring settings are applicable to all the <i>Routing Service</i> entities unless they are explicitly overridden. See Monitoring (Section 2.4.4) .	0 or 1
<entity_monitoring>	Enables and configures remote monitoring for the routing_service entity.	0 or 1
<jvm>	<p>Configures the Java JVM used to load and run Java adapters such as the JMS Adapter. For example:</p> <pre><jvm> <class_path> <element> SocketAdapter.jar </element> </class_path> <options> <element>-Xms32m</element> <element>-Xmx128m</element> </options> </jvm></pre> <p>The class path for the Java adapters can be set using either the <class_path> tag or by setting the CLASSPATH environment variable. <i>Routing Service</i> will always add <Routing Service executable location>/../class/rtiroutingservicesdk.jar and <Routing Service executable location>/../class/dds.jar at the end of the user defined class path.</p> <p>You can use the <options> tag to specify options for the JVM, such as the initial and maximum Java heap sizes.</p>	0 or 1

2.4.2 Domain Route

A *domain route* defines a mapping between two data domains. Data available in either of these data domains can be routed to the other one. For example, a domain route could define a mapping between two different *Connex* domains or between a *Connex* domain and a JMS provider's network. How this data is actually read and written is defined in specific *routes*.

A domain route creates two connections, known as **connection_1** and **connection_2**. Each connection belongs to one of the two data domains.

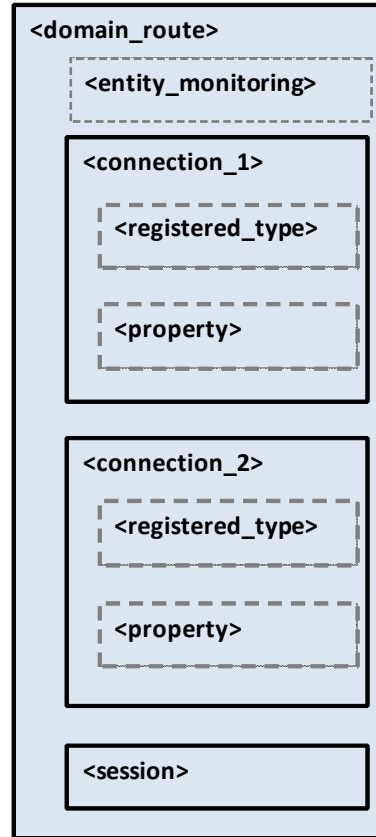
For example:

```
<dds>
  <routing_service name="Router1"
    group_name="Group1">
    <domain_route name="DomainRoute1">
      <connection_1 plugin_name="...">
        . . .
      </connection_1>

      <connection_2 plugin_name="...">
        . . .
      </connection_2>

      <session name="Session">
        . . .
      </session>
    </domain_route>
    . . .
  </routing_service>
</dds>
```

The connection tags require the specification of the attribute **plugin_name**, which will be used to associate a connection with an adapter plugin defined within `<adapter_library>` (see [Section 2.1](#)).



For *Connex* domains, the connections are specified using the tags **participant_1** and **participant_2**. Each tag has one associated DomainParticipant.

The following example routes information between two *Connex* domains.

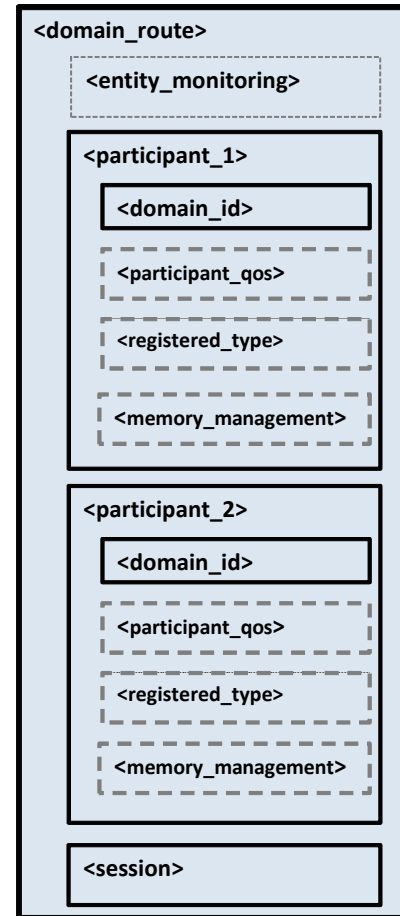
```
<dds>
  <routing_service name="Router1"
    group_name="Group1">
    <domain_route name="DomainRoute1">
      <participant_1>
        <domain_id>54</domain_id>
        ...
      </participant_1>

      <participant_2>
        <domain_id>55</domain_id>
        ...
      </participant_2>

      <session name="Session">
        ...
      </session>
    </domain_route>
    ...
  </routing_service>
</dds>
```

Configurations mixing connections and participants are allowed to provide communication between *Connex* domains and other data domains.

The following example routes information between a JMS provider network and a *Connex* domain.



```

<dds>
  <routing_service name="Router1"
    group_name="Group1">
    <domain_route name="DomainRoute1">
      <connection_1 plugin_name="adapter_library::jms">
        ...
      </connection_1>

      <participant_2>
        <domain_id>55</domain_id>
        ...
      </participant_2>

      <session name="Session">
        ...
      </session>
    </domain_route>
    ...
  </routing_service>
</dds>

```

Table 2.3 lists the tags allowed within a `<domain_route>` tag.

Table 2.3 Domain Route Tags

Tags within <code><domain_route></code>	Description	Number of Tags Allowed
<code><connection_1></code>	Applicable to non- <i>Connex</i> t domains. Configures the first connection. See Table 2.4.	1 (required)
<code><participant_1></code>	Only applicable to <i>Connex</i> t domains. Configures the first participant. See Table 2.5.	
<code><connection_2></code>	Applicable to non- <i>Connex</i> t domains. Configures the second connection. See Table 2.4.	1 (required)
<code><participant_2></code>	Only applicable to <i>Connex</i> t domains. Configures the second participant. See Table 2.5.	
<code><session></code>	Defines a single-threaded context in which data is routed according to specified routes. See Session (Section 2.4.5) .	1 or more (required)
<code><entity_monitoring></code>	Enables and configures remote monitoring for the domain route. See Monitoring (Section 2.4.4) .	0 or 1

Table 2.4 lists the tags allowed within `<connection_1>` and `<connection_2>` tags.

Table 2.4 Connection Tags

Tags within <code><connection_1/2></code>	Description	Number of Tags Allowed
<code><registered_type></code>	Registers a type name and associates it with a type representation. When you define a type in the configuration file (with the <code><types></code> tag), you have to register the type in order to use it in routes. See Route Types (Section 2.4.6.1) .	0 or more
<code><property></code>	Sequence of name/value(string) pairs that can be used to configure the parameters of the connection. For example: <pre> <property> <value> <element> <name> jms.connection.username </name> <value>myusername</value> </element> </value> </property> </pre>	0 or 1

Table 2.5 lists the tags allowed within `<participant_1>` and `<participant_2>` tags.

Table 2.5 Participant Tags

Tags within <participant_1/2>	Description	Number of Tags Allowed
<domain_id>	Sets the domain ID associated with the participant.	1 (required)
<memory_ management>	<p>Configures certain aspects of how <i>Connex</i> allocates internal memory. The configuration is per domain_route's participant and therefore affects all the contained <i>Connex</i> readers and <i>Connex</i> writers. For example:</p> <pre data-bbox="459 533 1003 934"> <domain_route name="test"> <participant_1> <domain_id>0</domain_id> ... <memory_management> <sample_buffer_min_size> X </sample_buffer_min_size> <sample_buffer_trim_to_size> true </sample_buffer_trim_to_size> </memory_management> </participant_1> ... </pre> <p>The <memory_management> tag can include the following tags:</p> <ul style="list-style-type: none"> <input type="checkbox"/> sample_buffer_min_size: For all <i>Connex</i> readers/writers, the way <i>Connex</i> allocates memory for samples is as follows: <i>Connex</i> pre-allocates space for samples up to size X in the reader and writer queues. If a sample has an actual size greater than X, the memory is allocated dynamically for that sample. The default size is DDS_LENGTH_UNLIMITED (meaning no dynamic memory is used; the maximum sample size is pre-allocated). <input type="checkbox"/> sample_buffer_trim_to_size: If set to true, after allocating dynamic memory for very large samples, that memory will be released when possible. If false, that memory will not be released but kept for future samples if needed. The default is false. <p>This feature is useful when a data type has a very high maximum size (e.g., megabytes) but most of the samples sent are much smaller than the maximum possible size (e.g., kilobytes). In this case, the memory footprint is reduced dramatically, while still correctly handling the rare cases in which very large samples are published.</p>	0 or more

Table 2.5 Participant Tags

Tags within <participant_1/2>	Description	Number of Tags Allowed
<registered_type>	Registers a type name and associates it with a type code. When you define a type in the configuration file (with the <types> tag), you have to register the type in order to use it in topic routes. See Route Types (Section 2.4.6.1) .	0 or more
<participant_qos>	<p>Sets the participant QoS.</p> <p>The contents of this tag are specified in the same manner as a <i>Connex</i> QoS profile file—see Chapter 15 in the <i>RTI Core Libraries and Utilities User's Manual</i>.¹</p> <p>If not specified, the default is used.</p> <p>You can use a <participant_qos> tag inside a <qos_library>/<qos_profile> previously defined in your configuration file by referring to it like this:</p> <pre><participant_qos base_name="MyLibrary::MyProfile" /></pre> <p>To use that profile but override just some values:</p> <pre><participant_qos base_name="MyLibrary::MyProfile"> <discovery> <initial_peers> <element>udpv4://192.168.1.12</element> <element>shmem://</element> </initial_peers> </discovery> </participant_qos></pre> <p>(This applies to all QoS tags: <publisher_qos>, <subscriber_qos> in sessions; <datareader_qos>, <datawriter_qos> in topic routes and auto topic routes.)</p>	0 or 1

1. See <Connex installation directory>/ndds.<version>/doc/pdf/RTI_Connex_UsersManual.pdf.

2.4.3 Administration

You can create a *Connex*t application that can remotely control *Routing Service*. The `<administration>` tag is used to enable remote administration and configure its behavior.

By default, remote administration is turned off in *Routing Service* for security reasons. A remote administration section is not required in the configuration file.

For example:

```
<dds>
  <routing_service>
    <administration>
      <domain_id>55</domain_id>
      <save_path>/home/david/mysaved_config.xml</save_path>
    </administration>
    ...
  </routing_service>
</dds>
```

When remote administration is enabled, *Routing Service* will create a DomainParticipant, Publisher, Subscriber, DataWriter, and DataReader. These entities are used to receive commands and send responses. You can configure these entities with QoS tags within the `<administration>` tag.

Table 2.6 lists the tags allowed within `<administration>` tag.

For more details, please see [Chapter 5: Administering Routing Service from a Remote Location](#).

Note: The command-line options used to configure remote administration take precedence over the XML configuration (see [Table 3.1 on page 3-2](#)).

Table 2.6 Remote Administration Tags

Tags within <code><administration></code>	Description	Number of Tags Allowed
<code><domain_id></code>	Specifies which domain ID <i>Routing Service</i> will use to enable remote administration.	1 (required)
<code><participant_qos></code>	Configures the DomainParticipant QoS for remote administration. If the tag is not defined, <i>Routing Service</i> will use the <i>Connex</i> t defaults.	0 or 1
<code><publisher_qos></code>	Configures the Publisher QoS for remote administration. If the tag is not defined, <i>Routing Service</i> will use the <i>Connex</i> t defaults.	0 or 1

Table 2.6 Remote Administration Tags

Tags within <administration>	Description	Number of Tags Allowed
<subscriber_qos>	Configures the Subscriber QoS for remote administration. If the tag is not defined, <i>Routing Service</i> will use the <i>Connex</i> defaults.	0 or 1
<datareader_qos>	Configures the DataReader QoS for remote administration. If the tag is not defined, <i>Routing Service</i> will use the <i>Connex</i> defaults with the following changes: reliability.kind = DDS_RELIABLE_RELIABILITY_QOS (this value cannot be changed) history.kind = DDS_KEEP_ALL_HISTORY_QOS resource_limits.max_samples = 32	0 or 1
<datawriter_qos>	Configures the DataWriter QoS for remote administration. If the tag is not defined, <i>Routing Service</i> will use the <i>Connex</i> defaults with the following changes: history.kind = DDS_KEEP_ALL_HISTORY_QOS resource_limits.max_samples = 32	0 or 1
<save_path>	Specifies the file that will contain the saved configuration. It is empty by default. A <save_path> must be specified if you want to use the save (Section 5.2.9) command. If the file specified by <save_path> already exists, the file will be overwritten when save is executed. This value is mutable when an update (Section 5.2.11) command targets a routing service. This value is sent as part of the monitoring configuration data for the routing service (see Configuration Data for the Routing Service (Section 6.2.1)).	0 or 1
<autosave_on_update>	A boolean that, if true, automatically triggers a save command when configuration updates are received. It is false by default. This value is mutable when an update (Section 5.2.11) command targets a routing service. This value is sent as part of the monitoring configuration data for the routing service (see Configuration Data for the Routing Service (Section 6.2.1)).	0 or 1

2.4.4 Monitoring

You can create a *Connex*t application that can remotely monitor the status of *Routing Service*. To enable remote monitoring and configure its behavior, use the `<monitoring>` and `<entity_monitoring>` tags.

By default, remote monitoring is turned off in *Routing Service* for security and performance reasons. A remote monitoring section is not required in the configuration file.

For example:

```
<dds>
  <routing_service>
    <enabled>true</enabled>
    <monitoring>
      <domain_id>55</domain_id>
      <status_publication_period>
        <sec>1</sec>
      </status_publication_period>
    </monitoring>
    ...
  </routing_service>
</dds>
```

Routing Service allows monitoring of the following kinds of entities:

- `<routing_service>` (see [Section 2.4.1](#))
- `<domain_route>` (see [Section 2.4.2](#))
- `<session>` (see [Section 2.4.5](#))
- `<route>` (see [Section 2.4.6](#))
- `<topic_route>` (see [Section 2.4.6](#))
- `<auto_route>` (see [Section 2.4.7](#))
- `<auto_topic_route>` (see [Section 2.4.7](#))

For each entity, *Routing Service* can publish two kinds of information:

- Entity data
- Entity status

Entity data provides information about the configuration of the entity. For example, the route data contains information such as the stream name and the type name. Entity data information is republished every time the entity is enabled, disabled or has configuration changes.

Entity status provides information about the operational status of an entity. This kind of information changes continuously and is computed and published periodically. For example, the route status contains information such as the route's latency and throughput.

For more information about entity data and status, see [Chapter 6: Monitoring Routing Service from a Remote Location](#).

When remote monitoring is enabled, *Routing Service* will create one `DomainParticipant`, one `Publisher`, five `DataWriters` for data publication (one for each kind of entity), and five `DataWriters` for status publication (one for each kind of entity). You can configure the QoS of these entities with the `<monitoring>` tag defined under `<routing_service>`.

The general remote monitoring parameters specified using the `<monitoring>` tag in `<routing_service>` (except `domain_id`, `participant_qos`, `publisher_qos`, and `datawriter_qos`) can be overwritten on a per entity basis using the `<entity_monitoring>` tag.

For example:

```
<dds>
  <routing_service>
    <monitoring>
      <domain_id>55</domain_id>
      <status_publication_period>
        <sec>1</sec>
      </status_publication_period>
    </monitoring>
    ...
    <domain_route>
      <entity_monitoring>
        <status_publication_period>
          <sec>4</sec>
        </status_publication_period>
      </entity_monitoring>
      ...
    </domain_route>
  </routing_service>
</dds>
```

Table 2.7 lists the tags allowed within `<monitoring>` tag.

Table 2.7 **Monitoring tags**

Tags within <code><monitoring></code>	Description	Number of Tags Allowed
<code><enabled></code>	<p>Enables/disables general remote monitoring.</p> <p>Setting this value to true (default value) in the <code><monitoring></code> tag under <code><routing_service></code> enables monitoring in all the entities unless they explicitly disable it by setting this tag to false in their local <code><entity_monitoring></code> tags.</p> <p>Setting this tag to false in the <code><monitoring></code> tag under <code><routing_service></code> disables monitoring in all the <i>Routing Service</i> entities. In this case, any monitoring configuration settings in the entities are ignored.</p>	0 or 1
<code><domain_id></code>	Specifies which domain ID <i>Routing Service</i> will use to enable remote monitoring.	1 (required)
<code><participant_qos></code>	<p>Configures the DomainParticipant QoS for remote monitoring.</p> <p>If the tag is not defined, <i>Routing Service</i> will use the <i>Connex</i> defaults with the following change: resource_limits.type_code_max_serialized_length = 4096</p>	0 or 1
<code><publisher_qos></code>	<p>Configures the Publisher QoS for remote monitoring.</p> <p>If the tag is not defined, <i>Routing Service</i> will use the <i>Connex</i> defaults.</p>	0 or 1
<code><datawriter_qos></code>	<p>Configures the DataWriter QoS for remote monitoring.</p> <p>If the tag is not defined, <i>Routing Service</i> will use the <i>Connex</i> defaults with the following change: durability.kind = DDS_TRANSIENT_LOCAL_DURABILITY_QOS</p>	0 or 1
<code><status_publication_period></code>	<p>Specifies the frequency at which the status of an entity is published.</p> <p>For example: <pre><status_publication_period> <sec>3</sec> <nanosec>0</nanosec> </status_publication_period></pre> </p> <p>If the tag is not defined, the period is 5 seconds.</p> <p>The status publication period defined in <code><routing_service></code> is inherited by all the entities inside <code><routing_service></code>.</p> <p>An entity can overwrite the period.</p>	0 or 1

Table 2.7 Monitoring tags

Tags within <monitoring>	Description	Number of Tags Allowed
<statistics_sampling_period>	<p>Specifies the frequency at which status statistics are gathered. Statistical variables such as latency, are part of the entity status. For example:</p> <pre><statistics_sampling_period> <sec>1</sec> <nanosec>0</nanosec> </statistics_sampling_period></pre> <p>The statistics period for a given entity should be smaller than the publication period.</p> <p>If the tag is not defined, the period is 1 second.</p> <p>The statistics sampling period defined in <routing_service> is inherited by all the entities inside <routing_service>.</p> <p>An entity can overwrite the period.</p>	0 or 1
<historical_statistics>	<p>Enables or disables the publication of statistics calculated within fixed time windows.</p> <p>By default, <i>Routing Service</i> only publishes the statistics corresponding to the window between two status publications.</p> <p>By using this tag, you can get the following additional windows:</p> <ul style="list-style-type: none"> <input type="checkbox"/> 5 seconds <input type="checkbox"/> 1 minute <input type="checkbox"/> 5 minutes <input type="checkbox"/> 1 hour <input type="checkbox"/> Up time (since the entity was enabled) <p>For example:</p> <pre><historical_statistics> <five_second>true</five_second> <one_minute>true</one_minute> <five_minute>>false</five_minute> <one_hour>true</one_hour> <up_time>>false</up_time> </historical_statistics></pre> <p>If a window is not present (inside the tag <historical_statistics>), it is considered disabled.</p> <p>Historical statistics can be overwritten on a per entity basis.</p>	0 or 1

2.4.4.1 Monitoring Configuration Inheritance

The monitoring configuration defined in `<routing_service>` is inherited by all the entities defined inside the tag.

An entity can overwrite three elements of the monitoring configuration:

- The status publication period
- The statistics sampling period
- The historical statistics windows

Each one of this three elements is inherited and can be overwritten independently using the `<entity_monitoring>` tag.

For example:

```
<dds>
  <routing_service name="MonitoringExample">
    <monitoring>
      <domain_id>55</domain_id>
      <status_publication_period>
        <sec>1</sec>
      </status_publication_period>
      <statistics_sampling_period>
        <sec>1</sec>
        <nanosec>0</nanosec>
      </statistics_sampling_period>
    </monitoring>
    ...
    <domain_route>
      <entity_monitoring>
        <status_publication_period>
          <sec>4</sec>
        </status_publication_period>
      </entity_monitoring>
      ...
    </domain_route>
  </routing_service>
</dds>
```


In the previous example, the domain route overwrites the status publication period to 4 seconds and inherits the statistics sampling period.

Table 2.8 Entity Monitoring Tags

Tags within <entity_monitoring>	Description	Number of Tags Allowed
<enabled>	Enables/disables remote monitoring for a given entity. If general monitoring is disabled this value is ignored. Default value: true	0 or 1
<status_publication_period>	Specifies the frequency at which the status of an entity is published. For example: <pre><status_publication_period> <sec>3</sec> <nanosec>0</nanosec> </status_publication_period></pre> If the tag is not defined, its value is inherited from the general monitoring settings.	0 or 1
<statistics_sampling_period>	Specifies the frequency at which status statistics are gathered. Statistical variables such as latency, are part of the entity status. For example: <pre><statistics_sampling_period> <sec>1</sec> <nanosec>0</nanosec> </statistics_sampling_period></pre> The statistics period for a given entity should be smaller than the publication period. If the tag is not defined, the period is inherited from the general monitoring settings. This tag is only present in the <entity_monitoring> tag of <route>, <topic_route>, <auto_route>, <auto_topic_route> and <routing_service>.	0 or 1

Table 2.8 Entity Monitoring Tags

Tags within <entity_monitoring>	Description	Number of Tags Allowed
<historical_statistics>	<p>Enables or disables the publication of statistics calculated within fixed time windows.</p> <p>By default, <i>Routing Service</i> only publishes the statistics corresponding to the window between two status publications.</p> <p>By using this tag, you can get the following additional windows:</p> <ul style="list-style-type: none"> <input type="checkbox"/> 5 seconds <input type="checkbox"/> 1 minute <input type="checkbox"/> 5 minutes <input type="checkbox"/> 1 hour <input type="checkbox"/> Up time (since the entity was enabled) <p>For example:</p> <pre><historical_statistics> <five_second>true</five_second> <one_minute>true</one_minute> <five_minute>false</five_minute> <one_hour>true</one_hour> <up_time>false</up_time> </historical_statistics></pre> <p>If a window is not present (inside the tag <historical_statistics>), it is considered disabled.</p> <p>If this tag is not defined, historical statistics are inherited from the general monitoring settings.</p>	0 or 1

2.4.5 Session

A <session> tag defines a single-threaded context for data routing; The data is routed according to specified routes ([Section 2.4.6](#)) and auto routes ([Section 2.4.7](#)).

Each session will have an associated session thread that will serialize access to the routes in the session.

For example:

```
<dds>
  ...
  <routing_service name="MyRoutingService">
    ...
    <domain_route>
      ...
      <session name="Session1">
        ...
        <route name="Route1" >
          ...
        </route>
        ...
      </session>
      ...
    </domain_route>
    ...
  </routing_service>
  ...
</dds>
```

Sessions that bridge *Connex* domains will create a Publisher and a Subscriber in the participants (**participant_1** or **participant_2**) associated with the *Connex* domains.

Table 2.9 lists the tags allowed within a `<session>` tag.

Table 2.9 Session Tags

Tags within <code><session></code>	Description	Number of Tags Allowed
<code><auto_route></code>	Defines a general route based on type and stream filters. See Auto Routes (Section 2.4.7) .	0 or more
<code><auto_topic_route></code>	Defines a general <i>Connex</i> topic route based on type and topic filters. See Auto Routes (Section 2.4.7) .	0 or more
<code><route></code>	Defines a data mapping between two streams. See Routes (Section 2.4.6)	0 or more
<code><topic_route></code>	Defines a data mapping between two <i>Connex</i> topics. See Routes (Section 2.4.6) .	0 or more

Table 2.9 Session Tags

Tags within <session>	Description	Number of Tags Allowed
<publisher_qos>	<p>Only applicable to <i>Connex</i>.</p> <p>Sets the QoS associated with the session Publishers. There is one Publisher per participant.</p> <p>The contents of this tag are specified in the same manner as a <i>Connex</i> QoS profile file—see Chapter 15 in the <i>RTI Core Libraries and Utilities User's Manual</i>.¹</p> <p>If the tag is not defined, <i>Routing Service</i> will use the <i>Connex</i> defaults.</p>	0 or 1
<subscriber_qos>	<p>Only applicable to <i>Connex</i>.</p> <p>Sets the QoS associated with the session Subscribers. There is one Subscriber per participant.</p> <p>The contents of this tag are specified in the same manner as a <i>Connex</i> QoS profile file—see Chapter 15 in the <i>RTI Core Libraries and Utilities User's Manual</i>.¹</p> <p>If the tag is not defined, <i>Routing Service</i> will use the <i>Connex</i> defaults.</p>	0 or 1
<thread>	<p>Sets the mask, priority and stack size of the thread associated with this session.</p> <p>Example:</p> <pre data-bbox="459 968 992 1310"> <session> <thread> <mask>MASK_DEFAULT</mask> <priority> THREAD_PRIORITY_DEFAULT </priority> <stack_size> THREAD_STACK_SIZE_DEFAULT </stack_size> </thread> ... </session> </pre> <p>Default values:</p> <pre data-bbox="496 1362 1003 1461"> mask = MASK_DEFAULT priority = THREAD_PRIORITY_DEFAULT stack_size = THREAD_STACK_SIZE_DEFAULT </pre>	0 or 1

Table 2.9 Session Tags

Tags within <session>	Description	Number of Tags Allowed
<property>	<p>Sequence of name/value(string) pairs that can be used to configure certain parameters of the session. For example:</p> <pre><property> <value> <element> <name>com.rti.socket.timeout</name> <value>1</value> </element> </value> </property></pre> <p>These properties are only used in non-<i>Connex</i>t domains.</p>	0 or 1
<wait_set>	<p>Configures the WaitSet used to sleep and notify the session thread when data is available.</p> <p>Example:</p> <pre><session> <wait_set> <max_event_count>5</max_event_count> <max_event_delay> <sec>1</sec> <nanosec>0 </nanosec> </max_event_delay> </wait_set> ... </session></pre> <p>In the previous example, the session thread wakes up and tries to read data after a 1 second timeout expires (<code>max_event_delay</code>) or after it has been notified five times across routes that new data is available (<code>max_event_count</code>).</p> <p>Default values:</p> <pre>max_event_count = 1 max_event_delay.sec = DURATION_INFINITE_SEC max_event_delay.nanosec = DURATION_INFINITE_NSEC</pre>	0 or 1
<monitoring>	<p>Enables and configures remote monitoring for the session. See Monitoring (Section 2.4.4) and Chapter 6: Monitoring Routing Service from a Remote Location.</p>	0 or 1

1. See `<Connex installation directory>/n.dds.<version>/doc/pdf/RTI_Connext_UsersManual.pdf`.

2.4.6 Routes

A route explicitly defines a mapping between an “input” data stream on one domain and an “output” data stream on the other domain.

For example, the following route defines a mapping between a *Connex* topic called Square and a JMS queue called Square.

```
<dds>
  ...
  <routing_service>
    ...
    <domain_route>
      <participant_1>
        <domain_id>54</domain_id>
      </participant_1>
      <connection_2 plugin_name="my_adapter_library::jms">
      </connection_2>
      ...
      <session name="Session1">
        ...
        <route name="DDSSquaresToJMSSquares">
          <dds_input participant="1">
            <topic_name>Square</topic_name>
            <registered_type_name>
              ShapeType
            </registered_type_name>
            ...
          </dds_input>
          <output>
            <stream_name>Square</topic_name>
            <registered_type_name>
              ShapeType
            </registered_type_name>
            ...
          </output>
          ...
        </route>
      </session>
      ...
    </domain_route>
    ...
  </routing_service>
  ...
</dds>
```

```

    </routing_service>
    ...
</dds>

```

Connex inputs and outputs within a route are defined using the XML tags **<dds_input>** and **<dds_output>**. Input and outputs from other data domains are defined using the tags **<input>** and **<output>**. A topic route is a special kind of route that defines a mapping between an “input” topic on one *Connex* domain and an “output” topic on other *Connex* domain. For example, the following topic route will subscribe to topic Square on domain 54 and will republish those samples on domain 55 as samples of topic Circle.

```

<dds>
  ...
  <routing_service>
    ...
    <domain_route>
      <participant_1>
        <domain_id>54</domain_id>
      </participant_1>
      <participant_2>
        <domain_id>55</domain_id>
      </participant_2>
      ...
      <session name="Session1">
        ...
        <topic_route name="SquaresToCircles">
          <input participant="1">
            <topic_name>Square</topic_name>
            <registered_type_name>
              ShapeType
            </registered_type_name>
            ...
          </input>
          <output>
            <topic_name>Circle</topic_name>
            <registered_type_name>
              ShapeType
            </registered_type_name>
            ...
          </output>
          ...
        </topic_route>
      </session>
    </domain_route>
    ...
  </routing_service>
  ...
</dds>

```

```

    </routing_service>
    ...
</dds>

```

In the previous example, the direction of the mapping is defined by the attribute **participant** of the tag **<input>**. Therefore, to change the above example to read Squares from domain 55 and write Circles on domain 54, we would use **<input participant="2">**. There is an equivalent attribute for non-*Connex*t inputs called **connection**.

Inputs and outputs in a route or topic route have an associated *StreamReader* and *StreamWriter*, respectively. For *Connex*t domains, the *StreamReader* will contain a *DataReader* and the *StreamWriter* will contain a *DataWriter*. The *Connex*t *DataReaders* and *DataWriters* belong to the corresponding session's *Subscriber* and *Publisher*.

The read and write operations in a route will be performed in the context of the thread associated with the session.

Routes vs. Auto Routes: A *route* is an explicit route of data for two specific streams. An *auto route* (defined with a different tag, **<auto_route>**) is a way to automatically create routes based on filters—see [Auto Routes \(Section 2.4.7\)](#).

[Table 2.10](#) lists the tags allowed within a **<route>**.

Table 2.10 **Route Tags**

Tags within <route>	Description	Number of Tags Allowed
<dds_input>	Only applicable to <i>Connex</i> t inputs. Defines the route's input topic. See Table 2.13 .	1 (required)
<input>	Applicable to non- <i>Connex</i> t inputs. Defines the route's input stream. See Table 2.13 .	
<dds_output>	Only applicable to <i>Connex</i> t outputs. Defines the route's output topic. See Table 2.13 .	1 (required)
<output>	Applicable to non- <i>Connex</i> t outputs. Defines the route's output stream. See Table 2.13 .	
<route_types>	Defines if the input connection will use types discovered in the output connection and vice versa for the creation of <i>StreamWriters</i> and <i>StreamReaders</i> in the route. See Discovering Types (Section 2.4.6.3) . Default: false	0 or 1

Table 2.10 Route Tags

Tags within <route>	Description	Number of Tags Allowed
<publish_with_original_timestamp>	When this tag is true, the data samples read from the input stream are written into the output stream with the same timestamp that was associated with them when they were made available in the input domain. This option may not be applicable in some adapter implementations in which the concept of timestamp is unsupported. Default: false	0 or 1
<transformation>	Sets a data transformation to be applied for every data sample (see Data Transformation (Section 2.4.6.5)).	0 or 1
<entity_monitoring>	Configures remote monitoring for the route. See Monitoring (Section 2.4.4) and Chapter 6: Monitoring Routing Service from a Remote Location .	0 or 1

Table 2.11 lists the tags allowed within a <topic_route>.

Table 2.11 Topic Route Tags

Tags within <topic_route>	Description	Number of Tags Allowed
<input>	Defines the topic route's input topic. See Table 2.13 .	1 (required)
<output>	Defines the topic route's output topic. See Table 2.13 .	1 (required)
<route_types>	Defines if the input domain will use types discovered in the output domain and vice versa for the creation of DataWriters and DataReaders in the topic route. See Discovering Types (Section 2.4.6.3) . Default: false	0 or 1
<publish_with_original_info>	Writes the data sample as if they came from its original writer. Setting this option to true allows having redundant routing services and prevents the applications from receiving duplicate samples. Default: false	0 or 1
<publish_with_original_timestamp>	When this tag is set to true, the data samples are written with their original source timestamp. Default: false	0 or 1

Table 2.11 Topic Route Tags

Tags within <topic_route>	Description	Number of Tags Allowed
<propagate_dispose>	Indicates whether or not disposed samples (NOT_ALIVE_DISPOSE) must be propagated by the topic route. This action maybe be overwritten by the execution of a transformation. Default: true	0 or 1
<propagate_unregister>	Indicates whether or not NOT_ALIVE_NO_WRITERS samples must be propagated by the topic route by using the <code>unregister_instance()</code> operation This action maybe be overwritten by the execution of a transformation. Default: true	0 or 1
<transformation>	Sets a data transformation to be applied for every data sample (see Data Transformation (Section 2.4.6.5)).	0 or 1
<entity_monitoring>	Configures remote monitoring for the topic route. See Monitoring (Section 2.4.4) and Chapter 6: Monitoring Routing Service from a Remote Location .	0 or 1

Table 2.12 lists the tags allowed within the input and output tags in a <route> tag.

Table 2.12 Input and Output Tags for a Route

Tags within <input>	Tags within <output>	Description	Number of Tags Allowed
<registered_type_name>		Sets the registered type name of this stream. See Route Types (Section 2.4.6.1) .	1 (required)
<stream_name>		Sets the stream name.	1 (required)
	<creation_mode>	Specifies when to create the StreamReader/StreamWriter. Default: IMMEDIATE See Creation Modes—Controlling when Stream-Readers and StreamWriters are Created (Section 2.4.6.4) .	0 or 1

Table 2.12 Input and Output Tags for a Route

Tags within <input>	Tags within <output>	Description	Number of Tags Allowed
<property>		Sequence of name/value(string) pairs that can be used to configure certain parameters of the StreamReaders/StreamWriters. For example: <pre><property> <value> <element> <name> com.rti.socket.port </name> <value>16556</value> </element> </value> </property></pre>	0 or 1

Table 2.13 lists the tags allowed within the *Connex* input and output tags, in a <route> or <topic_route> tag.

Table 2.13 Connex Input and Output Tags for a Route or Topic Route

Tags within <topic_route>/<input> and <route>/<dds_input>	Tags within <topic_route>/<output> and <route>/<dds_output>	Description	Number of Tags Allowed
<registered_type_name>		Sets the registered type name of this topic. See Route Types (Section 2.4.6.1) .	1 (required)
<topic_name>		Sets the topic name.	1 (required)
<creation_mode>		Specifies when to create the DataReader/DataWriter. Default: IMMEDIATE See Creation Modes—Controlling when StreamReaders and StreamWriters are Created (Section 2.4.6.4) .	0 or 1

Table 2.13 Connex Input and Output Tags for a Route or Topic Route

Tags within <topic_route>/<input> and <route>/<dds_input>	Tags within <topic_route>/<output> and <route>/<dds_output>	Description	Number of Tags Allowed
<content_filter>	N/A	Defines a SQL content filter for the DataReader. Example: <pre> <topic_route> ... <input> ... <content_filter> <expression> x > 100 </expression> </content_filter> ... </input> ... </topic_route> </pre>	0 or 1
<datareader_qos>	<datawriter_qos>	Sets the DataReader or DataWriter QoS. The contents of this tag are specified in the same manner as a <i>Connex</i> QoS profile file—see Chapter 15 in the <i>RTI Core Libraries and Utilities User's Manual</i> . ¹ If the tag is not defined, <i>Routing Service</i> will use the <i>Connex</i> defaults.	0 or 1

1. See <Connex installation directory>/n.dds.<version>/doc/pdf/RTI_Connex_UsersManual.pdf.

2.4.6.1 Route Types

The tag <registered_type_name> within the <input> and <output> tags contains the registered type name of the stream. The actual definition of that type can be set in the configuration file (see Section 2.4.6.2) or it can be discovered by the connections (see Section 2.4.6.3).

2.4.6.2 Defining Types in the Configuration File

To define and use a type in your XML configuration file:

1. Define your type within the **<types>** tag. (This is one of the top-level tags, see [Table 2.1](#).)
2. Register it in the connection(s)/participant(s) where you will use it.
3. Refer to it in the domain route(s) that will use it.

For example:

```
<dds>
...
<types>
  <struct name="PointType">
    ...
  </struct>
</types>
...
<routing_service name="MyRoutingService">
  ...
  <domain_route>
    <connection_1>
      ...
      <registered_type name="Position" type_name="PointType"/>
    </connection_1>
    <participant_2>
      ...
      <registered_type name="Position" type_name="PointType"/>
    </participant_2>
  ...
  <session>
    <topic_route>
      <input participant="2">
        <registered_type_name>Position
        </registered_type_name>
      </input>
      <output>
        ...
      </output>
    </topic_route>
  </session>
  ...
</domain_route>
...
</routing_service>
...
```

```
<dds>
```

The type description is done using the *Connex* XML format for type definitions. For more information, see Section 3.4 in the *RTI Core Libraries and Utilities User's Manual*.¹

2.4.6.3 Discovering Types

If a route refers to types that are not defined in the configuration file, *Routing Service* has to discover their type representation (e.g. typecode). A route cannot be created without the type representation information.

By default, the *StreamReader* creation will be tied to the discovery of types (e.g. typecodes) in the input domain and the *StreamWriter* creation will be tied to the discovery of types (e.g. typecodes) in the output domain. If you want to use types discovered in either one of the domains for the creation of both the *StreamReader* and *StreamWriter*, you must set the `<route_types>` tag to true.

In the following example, both the *StreamWriter* and *StreamReader* will be created as soon as the type **ShapeType** is discovered in either domain.

```
<topic_route>
  <route_types>true</route_types>
  <input participant="1">
    <creation_mode>IMMEDIATE</creation_mode>
    <registered_type_name>ShapeType</registered_type_name>
    ...
  </input>
  <output>
    <creation_mode>IMMEDIATE</creation_mode>
    <registered_type_name>ShapeType</registered_type_name>
    ...
  </output>
  ...
</topic_route>
```

In this next example, the *StreamReader* will be created only when the type **ShapeType** is discovered in the input domain; the *StreamWriter* will be created only when the type **ShapeType** is discovered in the output domain.

```
<topic_route>
  <route_types>>false</route_types>
  <input participant="1">
    <creation_mode>IMMEDIATE</creation_mode>
    <registered_type_name>ShapeType</registered_type_name>
```

1. See `<Connex installation directory>/ndds.<version>/doc/pdf/RTI_Connex_UsersManual.pdf`.

```

    ...
  </input>
  <output>
    <creation_mode>IMMEDIATE</creation_mode>
    <registered_type_name>ShapeType</registered_type_name>
    ...
  </output>
  ...
</topic_route>

```

2.4.6.4 Creation Modes—Controlling when StreamReaders and StreamWriters are Created

The way a route creates its StreamReader and StreamWriter and starts reading and writing data can be configured.

The `<creation_mode>` tag in a route's `<input>` and `<output>` tags controls *when* the routing service StreamReader/StreamWriter is created. [Table 2.14](#) lists the possible values for the `<creation_mode>` tag.

Table 2.14 **Creation Modes**

<code><creation_mode></code> Values	Description
IMMEDIATE (<i>default</i>)	The route StreamReader/StreamWriter is created as soon as possible; that is, as soon as the types are available. Note that if the type is defined in the configuration file, the creation will occur when the routing service starts. If the type is not defined in the configuration file, it has to be discovered; see Discovering Types (Section 2.4.6.3) .
ON_DOMAIN_MATCH	The route StreamReader is not created until the associated connection discovers a data Producer on the same stream. For example, for a <i>Connex</i> domain, <i>Routing Service</i> will not create the route DataReader until a DataWriter for the same topic is discovered on the same domain. The routing service StreamWriter is not created until the associated connection discovers a data Consumer on the same stream. For example, for a <i>Connex</i> domain, <i>Routing Service</i> will not create the route DataWriter until a DataReader for the same topic is discovered on the same domain.
ON_ROUTE_MATCH	The routing service StreamReader/StreamWriter is not created until its counterpart in the route is created.

Table 2.14 **Creation Modes**

<creation_mode> Values	Description
ON_DOMAIN_AND_ROUTE_MATCH	Both conditions must be true.
ON_DOMAIN_OR_ROUTE_MATCH	At least one of the conditions must be true.

Route Destruction:

The same rules that are applied to create the route StreamWriter and StreamReader also apply to their destruction. When the condition that triggered the creation of that entity becomes false, the entity is destroyed. (Note that IMMEDIATE will never become false.)

For example, if the creation mode of a topic route's **<input>** tag is ON_DOMAIN_MATCH, when all the matching user DataWriters in the input domain are deleted, the topic route's DataReader is deleted.

Example 1

In this example, data is routed as soon as a user DataWriter is publishing it on the first domain.

```
<topic_route>
  <input participant="1">
    <creation_mode>ON_DOMAIN_MATCH</creation_mode>
    ...
  </input>
  <output>
    <creation_mode>ON_ROUTE_MATCH</creation_mode>
    ...
  </output>
</topic_route>
```

Example 2

In this example, data is not routed until a user DataWriter is publishing and a user DataReader is already expecting it.

```
<topic_route>
  <input participant="1">
    <creation_mode>ON_DOMAIN_AND_ROUTE_MATCH</creation_mode>
    ...
  </input>
  <output>
    <creation_mode>ON_DOMAIN_AND_ROUTE_MATCH</creation_mode>
    ...
  </output>
</topic_route>
```


Example 3

In this example, all the data is received by the topic route's DataReader, because it is created as soon as a user DataWriter is discovered on the first domain. However, the data is not resent until a user DataReader on the other domain subscribes to it.

```
<topic_route>
  <input participant="1">
    <creation_mode>ON_DOMAIN_MATCH</creation_mode>
    ...
  </input>
  <output>
    <creation_mode>ON_DOMAIN_AND_ROUTE_MATCH</creation_mode>
    ...
  </output>
</topic_route>
```

2.4.6.5 Data Transformation

A route can transform the incoming data using a *transformation*, an object created by a transformation plugin.

For example, the following transformation switches the coordinates of the input sample: *x* becomes *y*, and *y* becomes *x*.

```
<topic_route name="SquareSwitchCoord">
  <input participant="1">
    <topic_name>Square</topic_name>
    <registered_type_name>ShapeType</registered_type_name>
  </input>
  <output>
    <topic_name>Square</topic_name>
    <registered_type_name>ShapeType</registered_type_name>
  </output>

  <transformation
    plugin_name="transformationLib::assign">
    <property>
      <value>
        <element>
          <name>X</name>
          <value>Y</value>
        </element>
        <element>
          <name>Y</name>
          <value>X</value>
        </element>
      </value>
    </property>
  </transformation>
```

```

        </element>
      </value>
    </property>
  </transformation>
</topic_route>

```

To include a transformation in a route:

1. Implement the transformation plugin API and generate a shared library. See [Chapter 4: Transforming Data with Routing Service](#) for more information.
2. Register that library in the configuration file by creating a `<transformation_plugin>` tag inside a `<transformation_library>` tag. (As noted in [Table 2.1](#), `<transformation_library>` is a top-level tag.)
3. Instantiate a transformation by creating a `<transformation>` tag inside a `<route>` or a `<topic_route>` tag.

[Table 2.15](#) lists the tags allowed within a `<transformation>` tag.

For additional information about transformations see [Chapter 4: Transforming Data with Routing Service](#).

Table 2.15 Transformation Tags

Tags within <code><transformation></code>	Description	Number of Tags Allowed
<code><input_type_name></code>	Type name of the data samples this transformation receives	0 or 1
<code><output_type_name></code>	Type name of the data samples this transformation creates	0 or 1
<code><property></code>	Sequence of name/value(string) pairs that can be used to configure certain parameters of the transformation. For example: <pre> <property> <value> <element> <name>scaling_factor</name> <value>2</value> </element> </value> </property> </pre>	0 or 1

2.4.7 Auto Routes

The tag `<auto_route>` defines a set of potential routes, with the same input and output type and same input and output stream name. A route can eventually be instantiated when a new stream is discovered with a type name and a stream name that match the filters in the auto route. When this happens, a route is created (but not necessarily started; see [Section 2.4.6.4](#)) with the configuration defined in the auto route tag.

For example:

```
<dds>
  ...
  <routing_service>
    ...
    <domain_route>
      <participant_1>
        <domain_id>54</domain_id>
      </participant_1>
      <connection_2 plugin_name="my_adapter_library::jms">
      </connection_2>
      ...
      <session>
        ...
        <auto_route name="AutoRoute1">
          ...
          <dds_input participant="1">
            <allow_topic_name_filter>*</allow_topic_name_filter>
            <allow_registered_type_name_filter>
              ShapeType
            </allow_registered_type_name_filter>
            ...
          </dds_input>
          <output>
            <allow_stream_name_filter>A*
            </allow_stream_name_filter>
            <allow_registered_type_name_filter>
              B*
            </allow_registered_type_name_filter>
            ...
          </output>
        </auto_route>
      </session>
      ...
    </domain_route>
    ...
  </routing_service>
```

...
</dds>

The above auto route will lead to the creation of a route every time any topic of type `ShapeType` is discovered on the *Connex*t domain or a JMS queue/topic starting with A with a type starting with B is discovered on the output JMS connection.

For example, discovering the topic “Triangle” of “ShapeType” will trigger the creation of a topic route that routes triangles from the *Connex*t domain to the JMS domain. Discovering a topic “Atopic” of type “Btype” on the JMS domain will trigger the creation of a topic route that routes “Atopic” from the *Connex*t domain to the JMS domain.

*Connex*t inputs and outputs within an auto route are defined using the XML tags `<dds_input>` and `<dds_output>`. Input and outputs from other data domains are defined using the tags `<input>` and `<output>`.

An auto topic route is a special kind of route that defines a mapping between two *Connex*t domains.

Please see the following tables for more information on allowable tags:

- ❑ Table 2.16 on page 2-42 lists the tags allowed within an `<auto_route>` tag.
- ❑ Table 2.17 on page 2-43 lists the tags allowed within an `<auto_topic_route>` tag.
- ❑ Table 2.18 on page 2-44 lists the tags allowed within `<input>` and `<output>` tags nested within an `<auto_route>` tag.
- ❑ Table 2.19 on page 2-47 lists the tags allowed within the `<dds_input>` and `<dds_output>` tags nested within an `<auto_topic_route>` or a `<topic_route>` tag.

Table 2.16 Auto Route Tags

Tag within <code><auto_route></code>	Description	Number of Tags Allowed
<code><dds_input></code>	Only applicable to <i>Connex</i> t inputs. Defines the auto route’s input stream (topic). See Auto Routes (Section 2.4.7)	1 (required)
<code><input></code>	Applicable to non- <i>Connex</i> t inputs. Defines the auto route’s input stream. See Auto Routes (Section 2.4.7) .	

Table 2.16 Auto Route Tags

Tag within <auto_route>	Description	Number of Tags Allowed
<dds_output>	Only applicable to <i>Connex</i> outputs. Defines the auto route's output stream (topic). See Auto Routes (Section 2.4.7) .	1 (required)
<output>	Applicable to non- <i>Connex</i> outputs. Defines the auto route's output stream. See Auto Routes (Section 2.4.7) .	
<publish_with_original_timestamp>	When this tag is true, the data samples read from the input streams are written into the output streams with the same timestamp that was associated with them when they were made available in the input domain. This option may not be applicable in some adapter's implementations where the concept of timestamp is not supported. Default: false	0 or 1
<entity_monitoring>	Enables and configures remote monitoring for the auto route. See Monitoring (Section 2.4.4) and Chapter 6: Monitoring Routing Service from a Remote Location	0 or 1

Table 2.17 Auto-topic Route Tags

Tag within <auto_topic_route>	Description	Number of Tags Allowed
<input>	Defines the auto topic route's input topic. See Auto Routes (Section 2.4.7) .	1 (required)
<output>	Defines the auto topic route's output topic. See Auto Routes (Section 2.4.7) .	1 (required)
<publish_with_original_info>	The topic routes are created with this configuration. When this flag is set to true, if you have N topic routes for the same topic (in different routers or in the same one), each sample that was written from a DataWriter in the input domain will be routed N times, but DataReaders on the output domain will only see one. Default: false	0 or 1

Table 2.17 Auto-topic Route Tags

Tag within <auto_topic_route>	Description	Number of Tags Allowed
<publish_with_ original_timestamp>	The topic routes are created with this configuration. When this tag is set to true, the data samples are written with their original source timestamp. Default: false	0 or 1
<propagate_dispose>	The topic routes are created with this configuration. When this flag is set to true, the NOT_ALIVE_DISPOSE samples received by the topic routes' DataReaders are not published with the topic routes' DataWriters. Default: true	0 or 1
<propagate_unregister >	The topic routes are created with this configuration. When this flag is set to true, the NOT_ALIVE_NO_WRITERS samples received by the topic routes' DataReaders are not published with the topic routes' DataWriters as unregister samples. Default: true	0 or 1
<entity_monitoring>	Enables and configures remote monitoring for the auto topic route. See Monitoring (Section 2.4.4) and Chapter 6: Monitoring Routing Service from a Remote Location	0 or 1

Table 2.18 Input and Output Tags for the <auto_route> Tag

Tags within <input>	Tags within <output>	Description	Number of Tags Allowed
<creation_mode>		The routes are created with this configuration. See Creation Modes—Controlling when StreamReaders and StreamWriters are Created (Section 2.4.6.4) .	0 or 1
<datareader_qos>	<datawriter_qos>	The topic routes are created with this configuration. The contents of this tag are specified in the same manner as for a <i>Connex</i> QoS profile file—see Chapter 15 in the <i>RTI Core Libraries and Utilities User's Manual</i> . ¹ If the tag is not defined, <i>Routing Service</i> will use the <i>Connex</i> defaults.	0 or 1

Table 2.18 Input and Output Tags for the <auto_route> Tag

Tags within <input>	Tags within <output>	Description	Number of Tags Allowed
<allow_stream_name_filter>		A stream name filter. ² You may use a comma-separated list to specify more than one filter. Default:* (allow all)	0 or 1
<allow_registered_type_name_filter>		A registered type name filter. ² You may use a comma-separated list to specify more than one filter. Default:* (allow all)	0 or 1
<deny_registered_type_name_filter>		A registered type name filter ² that should be denied (excluded). This is applied after the <allow_registered_type_name_filter>. You may use a comma-separated list to specify more than one filter. Default: Not applied	0 or 1
<deny_stream_name_filter>		A stream name filter ² that should be denied (excluded). This is applied after the <allow_stream_name_filter>. You may use a comma-separated list to specify more than one filter. Default:* (allow all)	0 or 1
<deny_topic_name_filter>		A topic name filter ² that should be denied (excluded). This is applied after the <allow_topic_name_filter>. You may use a comma-separated list to specify more than one filter. Default: Not applied	0 or 1

Table 2.18 Input and Output Tags for the <auto_route> Tag

Tags within <input>	Tags within <output>	Description	Number of Tags Allowed
<property>		<p>The topic routes are created with this configuration. Sequence of name/value(string) pairs that can be used to configure certain parameters of the StreamReaders/StreamWriters associated with the routes created from the auto route. For example:</p> <pre data-bbox="646 541 1076 827"> <property> <value> <element> <name> com.rti.socket.port </name> <value>16556</value> </element> </value> </property> </pre>	0 or 1

1. See <Connex installation directory>/ndds.<version>/doc/pdf/RTI_Connext_UsersManual.pdf.
2. As defined by the POSIX fnmatch API (1003.2-1992 section B.6).

Table 2.19 Connex Input and Output Tags for the <auto_topic_route> and <auto_route> Tags

Tags within <auto_topic_route>/<input> and <auto_route>/<dds_input>	Tags within <auto_topic_route>/<output> and <auto_route>/<dds_output>	Description	Number of Tags Allowed
		<p>The topic routes are created with this configuration. See Creation Modes—Controlling when StreamReaders and StreamWriters are Created (Section 2.4.6.4).</p>	0 or 1
		<p>The topic routes are created with a SQL content filter topic with this expression.</p> <pre> <auto_topic_route> ... <input> ... <content_filter> <expression> x > 100 </expression> </content_filter> ... </input> ... </auto_topic_route> </pre>	0 or 1
<content_filter>	N/A		
<datareader_qos>	<datawriter_qos>	<p>The topic routes are created with this configuration.</p> <p>The contents of this tag are specified in the same manner as for a <i>Connex</i> QoS profile file—see Chapter 15 in the <i>RTI Core Libraries and Utilities User’s Manual</i>.¹</p> <p>If the tag is not defined, <i>Routing Service</i> will use the <i>Connex</i> defaults.</p>	0 or 1
		<p>A topic name filter.²</p> <p>You may use a comma-separated list to specify more than one filter.</p> <p>Default: * (allow all)</p>	0 or 1
		<allow_topic_name_filter>	

Table 2.19 Connex Input and Output Tags for the <auto_topic_route> and <auto_route> Tags

Tags within <auto_topic_route>/<input> and <auto_route>/<dds_input>	Tags within <auto_topic_route>/<output> and <auto_route>/<dds_output>	Description	Number of Tags Allowed
<allow_registered_type_name_filter>		A registered type name filter. ² You may use a comma-separated list to specify more than one filter. Default:* (allow all)	0 or 1
<deny_topic_name_filter>		A topic name filter ² that should be denied (excluded). This is applied after the <allow_topic_name_filter>. You may use a comma-separated list to specify more than one filter. Default: Not applied	0 or 1
<deny_registered_type_name_filter>		A registered type name filter ² that should be denied (excluded). This is applied after the <allow_registered_type_name_filter>. You may use a comma-separated list to specify more than one filter. Default: Not applied	0 or 1

1. See <Connex installation directory>/ndds.<version>/doc/pdf/RTI_Connex_UsersManual.pdf.

2. As defined by the POSIX fnmatch API (1003.2-1992 section B.6).

2.4.8 Adapters

Adapters are pluggable components that allow *Routing Service* to consume and produce data for different data domains (e.g., *Connex*, JMS, Socket, etc.). By default, *Routing Service* is distributed with a built-in *Connex* adapter. Any other adapters must be registered within the `<adapter_library>` tag.

To support new data domains:

1. Implement the adapter plugin API in Java or C. See [Chapter 8: Extending Routing Service with Adapters](#) for more information.
2. Register the plugin in the configuration file by creating an `<adapter_plugin>` tag or a `<java_adapter_plugin>` inside an `<adapter_library>` tag. (As noted in [Table 2.1](#), `<adapter_library>` is a top-level tag.)
3. Instantiate an adapter connection by creating a `<connection>` tag inside a `<domain_route>` tag that refers to the adapter plugin.

For additional information about adapters see [Chapter 8: Extending Routing Service with Adapters](#).

2.5 Enabling and Disabling Routing Service Entities

The *Routing Service* entities associated with the tags `<routing_service>`, `<domain_route>`, `<route>`, `<topic_route>`, `<auto_route>`, and `<auto_topic_route>` can be created enabled or disabled using the attribute `enabled`.

By default, the value of the `enabled` attribute is true.

For example:

```
<dds>
  <routing_service name="TopicBridgeExample"
    group_name="rti.router.default" enabled="true">
    <domain_route name="DomainRoute" enabled="false">
      <participant_1>
        <domain_id>0</domain_id>
      </participant_1>
      <participant_2>
        <domain_id>1</domain_id>
      </participant_2>
    </domain_route>
  </routing_service>
</dds>
```

```
<session name="Session">
  <topic_route name="SquaresToCircles" enabled="false">
    <input participant="1">
      <registered_type_name>
        ShapeType
      </registered_type_name>
      <topic_name>Square</topic_name>
    </input>
    <output>
      <registered_type_name>
        ShapeType
      </registered_type_name>
      <topic_name>Circle</topic_name>
    </output>
  </topic_route>
</session>
</domain_route>
</routing_service>
</dds>
```

When an entity is created disabled, it can be enabled remotely using the commands [enable](#) (Section 5.2.5) and [disable](#) (Section 5.2.4). A **routing_service** can be created disabled by setting the attribute **enabled** to false or by using the **-noAutoEnable** command-line option. The command-line parameter takes precedence over the XML attribute value.

Chapter 3 Running Routing Service

This chapter describes:

- ❑ Starting Routing Service (Section 3.1)
- ❑ Stopping Routing Service (Section 3.2)

3.1 Starting Routing Service

Routing Service runs as a separate application. The script to run the executable is located in *<Routing Service installation directory>/scripts*.

Routing Service supports loading Java adapters. If your configuration is set up to load a Java adapter, follow these steps:

1. **On Windows Systems:** To use a Java adapter, you must have the Visual Studio 2005 service pack 1 redistributable libraries. You can obtain this package from Microsoft or RTI (see the *RTI Core Libraries and Utilities Release Notes*¹ for details).
2. Make sure Java 1.5 or higher is available.
3. Make sure you add the directory of the Java Virtual Machine dynamic library to your environment variable: LD_LIBRARY_PATH (on UNIX-based systems) or Path (on Windows systems). For example:

```
setenv LD_LIBRARY_PATH \  
    ${LD_LIBRARY_PATH}:/local/java/jdk1.5.0_07/jre/lib/i386/client
```

1. See *<Connex installation directory>/nlds.<version>/doc/pdf/RTI_Connext_ReleaseNotes.pdf*.

To start Routing Service, enter:

```
cd <installation directory for Routing Service>
scripts/rtiroutingservice [options]
```

Example:

```
cd <installation directory for Routing Service>
scripts/rtiroutingservice \
    -cfgFile example/shapes/topic_bridge.xml -cfgName example
```

Table 3.1 describes the command-line options.

3.2 Stopping Routing Service

To stop *Routing Service*, press **Ctrl-c**. *Routing Service* will perform a clean shutdown.

Table 3.1 RTI Routing Service Command-line Options

Option	Description
-appName <name>	Assigns a name to the execution of the <i>Routing Service</i> . Remote commands and status information will refer to the routing service using this name. See the <i>Routing Service User's Manual</i> for more information. In addition, the name of <i>DomainParticipants</i> created by <i>Routing Service</i> will be based on this name. Default: The name given with -cfgName if present, otherwise it is "RTI_Routing_Service".
-domainIdBase <ID>	Sets the base domain ID. This value is added to the domain IDs in the configuration file. For example, if you set -domainIdBase to 50 and use domain IDs 0 and 1 in the configuration file, then the <i>Routing Service</i> will use domains 50 and 51. Default: 0
-help	Displays help information.
-licenseFile <file>	Specifies the license file (path and filename). Only applicable to licensed versions of <i>Routing Service</i> . If not specified, <i>Routing Service</i> looks for the license as described in Installing the License File (Section 2.3) in the Getting Started Guide .

Table 3.1 RTI Routing Service Command-line Options

Option	Description
<code>-identifyExecution</code>	<p>Appends the host name and process ID to the service name provided with the <code>-appName</code> option. This helps ensure unique names for remote administration and monitoring.</p> <p>For example: <code>MyRoutingService_myhost_20024</code></p>
<code>-maxObjectsPerThread <int></code>	Parameter for the <code>DomainParticipantFactory</code> .
<code>-noAutoEnable</code>	<p>Starts <i>Routing Service</i> in a disabled state.</p> <p>Use this option if you plan to enable <i>Routing Service</i> remotely, as described in the <i>User's Manual</i>.</p> <p>This option overwrites the value of the <code>enable</code> attribute in the <code><routing_service></code> tag.</p>
<code>-remoteAdministrationDomainId <ID></code>	<p>Enables remote administration and sets the domain ID for remote communication.</p> <p>When remote administration is enabled, <i>Routing Service</i> will create a <i>DomainParticipant</i>, <i>Publisher</i>, <i>Subscriber</i>, <i>DataWriter</i>, and <i>DataReader</i> in the designated domain. The QoS values for these entities are described in the <i>Routing Service User's Manual</i>.</p> <p>This option overwrites the value of the tag <code><domain_id></code> within a <code><administration></code> tag. (See the <i>Routing Service User's Manual</i> for information on configuring remote access).</p> <p>Default: remote administration is not enabled unless it is enabled from the XML file.</p>
<code>-remoteMonitoringDomainId <ID></code>	<p>Enables remote monitoring and sets the domain ID for status publication.</p> <p>When remote monitoring is enabled, <i>Routing Service</i> will create one <i>DomainParticipant</i>, one <i>Publisher</i>, five <i>DataWriters</i> for data publication (one for each kind of entity), and five <i>DataWriters</i> for status publication (one for each kind of entity). The QoS values for these entities are described in the <i>Routing Service User's Manual</i>.</p> <p>This option overwrites the value of the tag <code><domain_id></code> within a <code><monitoring></code> tag. (See the <i>Routing Service User's Manual</i> for information on configuring remote monitoring).</p> <p>Default: remote monitoring is not enabled unless it is enabled from the XML file.</p>
<code>-sleep <seconds></code>	Sleeps for x seconds before executing the next command. This is useful when issuing commands from a command file (see <code>-cmdName</code>).
<code>-stopAfter <sec></code>	Stops the service after the specified number of seconds.

Table 3.1 RTI Routing Service Command-line Options

Option	Description
-use42eAlignment	<p>Enables compatibility with <i>RTI Data Distribution Service 4.2e</i>.</p> <p>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.</p> <p>Default: disabled</p>
-verbosity <n>	<p>Controls what type of messages are logged:</p> <ul style="list-style-type: none"> 0 - Silent 1 - Exceptions (<i>Connex</i> and <i>Routing Service</i>) (default) 2 - Warnings(<i>Routing Service</i>) 3 - Information (<i>Routing Service</i>) 4 - Warnings (<i>Connex</i> and <i>Routing Service</i>) 5 - Tracing (<i>Routing Service</i>) 6 - Tracing (<i>Connex</i> and <i>Routing Service</i>) <p>Each verbosity level, <i>n</i>, includes all the verbosity levels smaller than <i>n</i>.</p>
-version	<p>Prints the <i>Routing Service</i> version number.</p>

Chapter 4 Transforming Data with Routing Service

As described in [Data Transformation \(Section 2.4.6.5\)](#), a route can transform the incoming data using a *transformation*, which is an object created by a transformation plugin.

Transformation plugins implement the transformation API and must be provided as shared libraries that *Routing Service* will load dynamically.

Currently, the transformation plugin API is only supported in C.

This chapter describes:

- ❑ [Transformation Usage and Configuration \(Section 4.1\)](#)
- ❑ [Transformations Distributed with Routing Service \(Section 4.2\)](#)
- ❑ [Creating New Transformations \(Section 4.3\)](#)

4.1 Transformation Usage and Configuration

In the XML configuration file, transformation plugins must be defined within a transformation library.

For example:

```
<dds>
  <transformation_library name="MyTransfLib">
    <transformation_plugin name="MyTransfPlugin">
      <dll>mytransformation</dll>
      <create_function>
```

```
        MyTransfPlugin_create
    </create_function>
</transformation_plugin>
...
</transformation_library>
...
<routing_service>
    ...
</routing_service>
...
</dds>
```

Table 4.1 on page 4-3 lists the tags allowed within `<transformation_plugin>`.

Table 4.2 on page 4-3 lists the tags allowed within a `<transformation>` tag.

Once a transformation plugin is registered, a route can use it to create a data transformation. For example, the following route uses a transformation to switch the coordinates of the input sample: x becomes y , and y becomes x .

```
<topic_route name="SquareSwitchCoord">
  <input participant="1">
    <topic_name>Square</topic_name>
    <registered_type_name>ShapeType</registered_type_name>
  </input>
  <output>
    <topic_name>Square</topic_name>
    <registered_type_name>ShapeType</registered_type_name>
  </output>
  <transformation plugin_name="MyTransfLib:MyTransPlugin">
    <property>
      <value>
        <element>
          <name>X</name>
          <value>Y</value>
        </element>
        <element>
          <name>Y</name>
          <value>X</value>
        </element>
      </value>
    </property>
  </transformation>
</topic_route>
```

Table 4.1 Transformation Plugin Tags

Tags within <transformation_plugin>	Description	Number of Tags Allowed
<dll>	<p>Shared library containing the implementation of the transformation plugin. The <dll> tag may specify the exact name of the file (for example, lib/lib-mytransformation.so) or a general name (no file extension) which will be completed as follows:</p> <p><dll> value: dir/mytransformation</p> <p>Final Path (UNIX-based systems): dir/libmytransformation.so</p> <p>Final Path (Windows systems): dir/mytransformation.dll</p> <p>If the library specified with the <dll> tag cannot be opened (because the library path is not in the Path environment variable on Windows or the LD_LIBRARY_PATH environment variables on UNIX-based systems), <i>Routing Service</i> will look for the library in <Routing Service installation directory>/bin/<architecture>.</p>	1 (required)
<create_function>	<p>This tag will contain the name of the function used to create the transformation plugin (see Section 4.3.1).</p> <p>The function must be implemented in the shared library.</p>	1 (required)

Table 4.2 Transformation Tags

Tags within <transformation>	Description	Number of Tags Allowed
<property>	<p>Sequence of name/value(string) pairs that can be used to configure the parameters of the transformation. For example:</p> <pre><property> <value> <element> <name>X</name> <value>Y</value> </element> <element> <name>Y</name> <value>X</value> </element> </value> </property></pre> <p>In this example, the properties are used to define the field assignments. The semantics associated with the transformation property value depends on the plugin implementation.</p>	0 or 1

4.2 Transformations Distributed with Routing Service

Routing Service provides a transformation that is able to map fields of the input type to fields of the output type using the property tag inside the transformation to provide this mapping. For example:

```
<dds>
  ...
  <transformation_library name="TransformationLib">
    <transformation_plugin name="Assignment">
      <dll>rtirsassigntransf</dll>
      <create_function>
        RTI_RoutingServiceAssignTransformationPlugin_create
      </create_function>
    </transformation_plugin>
  </transformation_library>
  ...
  <routing_service name="MyService">
    <domain_route name="MyDomainRoute">
      <session name="MySession">
        <route name="MyRoute">
          ...
          <transformation plugin_name="TransformationLib::Assignment">
            <property>
              <value>
                <element>
                  <name>X</name>
                  <value>Y</value>
                </element>
                <element>
                  <name>Y</name>
                  <value>X</value>
                </element>
              </value>
            </property>
          </transformation>
        </route>
        ...
      </session>
      ...
    </domain_route>
    ...
  </routing_service>
  ...
</dds>
```

```

    </routing_service>
    ...
</dds>

```

This transformation plugin is implemented in the shared library, *<Routing Service installation directory>/bin/<architecture>/librtirsassigntransf.so* (or *rtirsassigntransf.dll* for Windows systems).

Important:

The assign transformation only supports the assignment of primitive fields (including strings) that are not part of arrays or sequences. For example:

```

<transformation plugin_name="TransformationLib::Assignment">
  <property>
    <value>
      <element>
        <name>position.x</name>
        <value>position.y</value>
      </element>
      <element> <!-- not supported -->
        <name>x[0]</name>
        <value>y[0]</value>
      </element>
      <element> <!-- supported -->
        <name>position</name>
        <value>position</value>
      </element>
    </value>
  </property>
</transformation>

```

4.3 Creating New Transformations

Routing Service provides a transformation SDK in C to support the creation of custom transformation plugins.

The SDK contains two main components:

- API header file: *<Routing Service installation directory>/include/routingservice/routingservice_transformation.h*.

The transformation plugin will include this header.

- ❑ Infrastructure library: <Routing Service installation directory>/lib/<architecture>/librtirsinfrastructure.so (for UNIX-based systems) and <Routing Service installation directory>/lib/<architecture>/rtirsinfrastructure.dll (for Windows systems).

The transformation plugin will link with this library.

Transformation plugins working with TypeCode and DynamicData must also link with the *Connex* libraries.

Important: Because RTI only distributes the release version of *Routing Service*, your transformation should be linked against the release version of the *Connex* shared libraries when needed.

4.3.1 Transformation Plugin API

Every transformation plugin will implement a plugin constructor (entry point to the shared library) that will be used by *Routing Service* to create a plugin instance.

```
typedef struct RTI_RoutingServiceTransformationPlugin *
  (*RTI_RoutingServiceTransformationPlugin_create) (
    RTI_RoutingServiceEnvironment * env);
```

The structure `RTI_RoutingServiceTransformationPlugin` will contain the plugin implementation as a set of function pointers.

```
struct RTI_RoutingServiceTransformationPlugin {
  RTI_RoutingServiceTransformationPlugin_DeleteFcn
  transformation_plugin_delete;
  RTI_RoutingServiceTransformationPlugin_CreateTransformationFcn
  transformation_plugin_create_transformation;
  RTI_RoutingServiceTransformationPlugin_DeleteTransformationFcn
  transformation_plugin_delete_transformation;
  RTI_RoutingServiceTransformation_TransformFcn
  transformation_transform;
  RTI_RoutingServiceTransformation_ReturnLoanFcn
  transformation_return_loan;
  RTI_RoutingServiceConfigurableEntity_UpdateFcn
  transformation_update;
  void * user_object;
};
```

The rest of this section introduces the different transformation functions. For detailed information about the API, please see the online (HTML) *Routing Service* documentation.

❑ delete

Deletes the transformation plugin instance.

```
typedef void (*RTI_RoutingServiceTransformationPlugin_delete)(
    struct RTI_RoutingServiceTransformationPlugin * plugin,
    RTI_RoutingServiceEnvironment * env);
```

❑ create_transformation

Creates a new transformation. The function is called when the route containing the transformation is ready to forward data.

```
typedef RTI_RoutingServiceTransformation
(*RTI_RoutingServiceTransformationPlugin_create_transformation)(
    struct RTI_RoutingServiceTransformationPlugin * plugin,
    const struct RTI_RoutingServiceTypeInfo * input_type_info,
    const struct RTI_RoutingServiceTypeInfo * output_type_info,
    const struct RTI_RoutingServiceProperties * properties,
    RTI_RoutingServiceEnvironment * env);
```

The behavior of the transformation can be configured using the **properties** parameter.

❑ delete_transformation

Deletes a transformation. The function is called when the route containing the transformation is disabled.

```
typedef void
(*RTI_RoutingServiceTransformationPlugin_delete_transformation)(
    struct RTI_RoutingServiceTransformationPlugin * plugin,
    RTI_RoutingServiceTransformation transformation,
    RTI_RoutingServiceEnvironment * env);
```

The **transformation** parameter corresponds to the value returned by the function **create_transformation()**.

❑ transform

This function is called in a route to transform a sequence of input data samples into a sequence of output data samples.

```
typedef void (*RTI_RoutingServiceTransformation_transform)(
    RTI_RoutingServiceTransformation transformation,
    RTI_RoutingServiceSample ** out_sample_lst,
    RTI_RoutingServiceSampleInfo ** out_info_lst,
    unsigned int * out_count,
    RTI_RoutingServiceSample * in_sample_lst,
    RTI_RoutingServiceSampleInfo * in_info_lst,
```

```
    unsigned int in_count,  
    RTI_RoutingServiceEnvironment * env);
```

When the routing service is done using the output samples, it will 'return the loan' to the transformation by calling the **return_loan()** operation.

The **transformation** parameter corresponds to the value returned by the function **create_transformation()**.

❑ **return_loan**

Indicates to the transformation that the routing service is done accessing the sequence of data samples obtained by an earlier invocation of **transform()**.

```
typedef void (*RTI_RoutingServiceTransformation_return_loan)(  
    RTI_RoutingServiceTransformation transformation,  
    RTI_RoutingServiceSample * sample_lst,  
    RTI_RoutingServiceSampleInfo * info_lst,  
    unsigned int count,  
    RTI_RoutingServiceEnvironment * env);
```

The **transformation** parameter corresponds to the value returned by the function **create_transformation()**.

❑ **update**

This function is called when the configuration of a transformation changes as a result of a remote **update** command.

```
typedef void (*RTI_RoutingServiceTransformation_UpdateFcn)(  
    RTI_RoutingServiceTransformation transformation,  
    const struct RTI_RoutingServiceProperties * properties,  
    RTI_RoutingServiceEnvironment * env);
```


Chapter 5 Administering Routing Service from a Remote Location

Routing Service can be controlled remotely by sending commands through a special *Connex* Topic. Any *Connex* application can be implemented to send these commands and receive the corresponding responses. A shell application that sends/receives these commands is provided with *Routing Service*.

The script for the shell application is in `<Routing Service installation directory>/scripts/rtirssh`.

Entering `rtirssh -help` will show you the command-line options:

```
RTI Routing Service Shell
Usage: rtirssh [options]...
Options:
  -domainId <integer>   Domain id for the remote configuration
  -timeout <seconds>    Max time to wait a remote response
  -cmdFile <file>       Run commands in this file
  -help                  Displays this information
```

5.1 Enabling Remote Administration

By default, remote administration is disabled in *Routing Service* for security reasons.

To enable remote administration you can use the `<administration>` tag (see [Section 2.4.3](#)) or the `-remoteAdministrationDomainId <ID>` command-line parameter, which enables remote administration and sets the domain ID for remote communication. For more information about the command-line options, see [Section 3.1 in the Getting Started Guide](#).

When remote administration is enabled, *Routing Service* will create a DomainParticipant, Publisher, Subscriber, DataWriter, and DataReader in the designated domain. (The QoS values for these entities are described in [Section 2.4.3](#).)

5.2 Remote Commands

This section describes the remote commands using the shell interface; [Section 5.3](#) explains how to use remote administration from a *Connex* application.

Remote commands:

- **add_peer** <target_routing_service> <domain_route_name> p1 | p2 <peer_list>
- **create** <target_routing_service>
domain_route | session | topic_route | auto_route [<parent_entity_name>]
<xml_url> [remote | local]
- **delete** <target_routing_service> [<entity_name>]
- **disable** <target_routing_service> [<entity_name>]
- **enable** <target_routing_service> [<entity_name>]
- **load** <target_routing_service> <cfg_name><xml_url> [remote | local]
- **pause** <target_routing_service> [<entity_name>]
- **resume** <target_routing_service> [<entity_name>]
- **save** <target_routing_service>
- **enable** <target_routing_service>
- **update** <target_routing_service> [<entity_name>] [<xml_url> | <assignment_expr>]
[remote | local]

Parameters:

- <assignment_expr> can be used instead of <xml_url> to modify single values in an entity configuration.

The assignment expression has the form:

<fully qualified value name> = <value>

For example:

```
update ShapeRouter DomainRoute1::Session1::SquareToCircles
  topic_route.input.datareader_qos.deadline.period.sec = 3
```

```
update ShapeRouter DomainRoute1::Session1::SquareToCircles
  topic_route.input.content_filter.expression = "x < 30"
```

- ❑ `<domain_route_name>` is the fully qualified name of a domain route entity
- ❑ `<entity_name>` is a fully qualified name. For example, consider the following XML configuration:

```
<routing_service name="ShapeRoutingService">
  ...
  <domain_route name="DomainRoute1">
    ...
    <session name="Session">
      <topic_route name="SquaresToCircles">
        ...
```

The above XML configuration would allow you to use commands such as:

- `enable ShapeRoutingService DomainRoute1::Session::SquaresToCircles`
- `enable ShapeRoutingService DomainRoute1`

Note that the fully qualified name does not include the name of the routing service.

- ❑ `<peer_list>` is a comma-separated list of peers
- ❑ `<target_routing_service>` can be:
 - The application name of a routing service, such as **"MyRoutingService1"**, as specified at start-up with the command-line option **-appName**
 - A regular expression¹ for a routing service name, such as **"MyRoutingService*"**
- ❑ `<xml_url>` can be:
 - A file URL, such as **file:///home/user/myconfig.xml**
 - A string URL, such as:

```
str://"<topic_route><input><datareader_qos>...
</datareader_qos></input></topic_route>"
```

If you omit the URL schema name, *Routing Service* will assume a file name; for example, **/home/user/myconfig.xml** is equivalent to **file:///home/user/myconfig.xml**.

1. As defined by the POSIX fnmatch API (1003.2-1992 section B.6)

In either case, the XML code can represent either a whole DTD-valid configuration file (such as the one you specify through the command-line when you start *Routing Service*) or a snippet of XML that only refers to a specific entity (this is further explained in [Section 5.2.11](#)).

The [**remote** | **local**] parameter is used with file URLs to indicate if the file is local to the shell (local) or local to the routing service (remote). If the file is local to the shell (local), the shell application will read it and will send it as a string URL. If the file is local to the routing service (remote), the shell will send it as a file URL that will be read by the routing service. The default value is **remote**.

If a relative path is specified, the path will be relative to the working directory in which the routing service (if **remote** is specified) or shell (if **local** is specified) is running.

5.2.1 add_peer

```
add_peer <target_routing_service> <domain_route_name> p1|p2 <peer_list>
```

The **add_peer** command passes the *peer_list* to the underlying DomainParticipant's **add_peer()** function. It is only valid for DomainParticipants in a domain route.

<domain_route_name> is like *<entity_name>*, but must be a domain route entity.

p1|p2 specifies if the DomainParticipant associated with *<participant_1>* or *<participant_2>* configuration is selected.

<peer_list> is a comma-separated list of peers.

5.2.2 create

```
create <target_routing_service>  
    domain_route|session|topic_route|auto_route [<parent_entity_name>]  
    <xml_url> [remote|local]
```

The **create** command is similar to [update \(Section 5.2.11\)](#), but the configuration is applied to a newly created entity instead of an existing one.

The second parameter (*domain_route* | *session* | *topic_route* | *auto_route*) is the kind of entity to be created. If the kind is a *domain_route*, there will be no parent. For the other kinds (*session*, *topic_route*, or *auto_route*), a *<parent_entity_name>* must be specified.

<xml_url> and [**remote** | **local**] are the same as used in [update \(Section 5.2.11\)](#), except that only XML snippets matching the entity kind are allowed. A full file (starting with *<dds>...*) is not valid.

For example (this would be entered as a single command, with no line-breaks):

```
create example topic_route DomainRoute::Session
str://"<topic_route name="TrianglesToTriangles">
<input participant="1"><registered_type_name>ShapeType
</registered_type_name><topic_name>Triangle</topic_name></input>
<output><registered_type_name>ShapeType</registered_type_name>
<topic_name>Triangle</topic_name></output></topic_route>"
```

5.2.3 delete

```
delete <target_routing_service> [<entity_name>]
```

You can invoke the **delete** command on domain routes, routes and auto routes. It acts like the **disable** (Section 5.2.4) command, but also purges the configuration data for the target entity.

For example:

```
delete example DomainRoute::Session::CirclesToCircles
```

A deleted entity cannot be re-enabled, but a new one can be created.

5.2.4 disable

```
disable <target_routing_service> [<entity_name>]
```

The **disable** command disables a routing service entity by destroying its sub-entities and corresponding *Connex* objects:

- ❑ Routing service—When a routing service is disabled, all of its domain routes are destroyed. You do not need to specify the *entity_name* to disable a routing service.
- ❑ Domain route—When a domain route is disabled, all its routes, topic routes, auto routes, and auto topic routes are destroyed, as well as both Connections (DomainParticipants for *Connex*). All the session threads are stopped and their corresponding adapter sessions (Publisher and Subscriber for *Connex*) are also deleted.
- ❑ Route, topic route, auto route and auto topic route—When a route, topic route, auto route, or auto topic route is disabled, its StreamReaders and StreamWriters are destroyed, so data will no longer be routed.

5.2.5 enable

```
enable <target_routing_service> [<entity_name>]
```

The **enable** command enables an entity that has been disabled or marked as **'enabled=false'** in the configuration file.

This command can be used to enable the following entities:

- ❑ Routing service—When a routing service is enabled, it uses the currently loaded configuration and starts. You don't need to specify the *entity_name* to enable a routing service.
- ❑ Domain route—When a domain route is enabled, it creates the participants, routes, topic routes, auto routes, and auto topic routes that it contains. The routes, topic routes, auto routes, and auto topic routes will be created enabled or disabled depending on their current configuration. Enabling a domain route is required to start routing data from the input domain to the output domain.
- ❑ Route, topic route, auto route, and auto topic route—Enabling a route, topic route, auto route or auto topic route is a necessary condition to start routing data between input and output streams. However, data routing will not start until the StreamWriter and StreamReader associated with a route are created (see [Section 2.4.6.4](#) for additional information).

5.2.6 load

```
load <target_routing_service> <cfg_name> <xml_url> [remote/local]
```

The **load** command loads specific XML configuration code. The *target_routing_service* must be disabled. For more information, see [How to Load the XML Configuration \(Section 2.2\)](#).

The XML code received must represent a valid routing service configuration file. The name of the `<routing_service>` tag to load is identified with `<cfg_name>`.

5.2.7 pause

```
pause <target_routing_service> <entity_name>
```

When the pause command is called in a route, the session thread containing this route will stop reading data from the route's StreamReader.

For routing service, domain routes, auto routes, and auto topic routes, the execution of this command will pause the contained topic routes and routes.

5.2.8 resume

```
resume <target_routing_service> <entity_name>
```

When the resume command is called in a route, the session thread containing this route will continue reading data from the route's StreamReader.

For routing service, domain routes, auto routes and auto topic routes, the execution of this command will resume the contained topic routes and routes.

5.2.9 save

```
save <target_routing_service>
```

This command writes the current configuration to a file. The file itself is specified with <save_path> (see [page 2-17](#)). If <save_path> has not been specified, the **save** command will fail. If the file specified by <save_path> already exists, the file will be overwritten.

The saved configuration is *functionally* equivalent to the loaded XML file plus any updates (either from an **update** command or other remote commands that change the configuration, such as **add_peer**). However it may not be *textually* equivalent. For example, the saved XML configuration may explicitly contain default values that were not in the initial XML.

5.2.10 unload

```
unload <target_routing_service>
```

The **unload** command unloads the current configuration that the *target_routing_service* is using, so you can change it with a subsequent **load** ([Section 5.2.6](#)) command.

The *target_routing_service* must be disabled for this command to succeed.

5.2.11 update

```
update <target_routing_service> [<entity_name>]
      [<xml_url>|<assignment_expr>] [remote|local]
```

The **update** command changes the configuration of a specific entity. [Table 5.1](#) shows the parameters that can be changed for each entity.

If you try to change an immutable parameter in an entity that is enabled, you will receive an error message. To change an immutable parameter, you must disable the routing service entity, change the parameter, and then enable the routing service entity again.

Table 5.1 Changeable Parameters

Entity	Mutable (changeable any time)	Immutable (only changeable when entity is <i>disabled</i>) ¹
Routing Service	<monitoring>/<enabled> <monitoring>/<status_publication_period> <entity_monitoring>/<enabled> <entity_monitoring>/<status_publication_period> <administration>/<save_path> <administration>/<autosave_on_update>	<monitoring>/<statistics_sampling_period> <monitoring>/<historical_statistics> <monitoring>/<domain_id> <entity_monitoring>/<statistics_sampling_period> <entity_monitoring>/<historical_statistics> <administration>/ <i><all except save_path and autosave_on_update></i>
Domain route	<connection_x>: Mutable properties in <property> (adapter specific) <participant_x>: Mutable QoS policies in <participant_qos> <entity_monitoring>/<enabled> <entity_monitoring>/<status_publication_period>	<connection_x>: Immutable properties in <property> (adapter specific). <participant_qos>: Immutable QoS policies in <participant_qos> <entity_monitoring>/<statistics_sampling_period> <entity_monitoring>/<historical_statistics>
Session	For non- <i>Connex</i> t adapter: Mutable properties in <property> (adapter specific) For <i>Connex</i> t adapter: Mutable QoS policies in <publisher_qos> and <subscriber_qos> <entity_monitoring>/<enabled> <entity_monitoring>/<status_publication_period>	For non- <i>Connex</i> t adapter: Mutable properties in <property> (adapter specific) For <i>Connex</i> t adapter: Immutable QoS policies in <publisher_qos> and <subscriber_qos> <entity_monitoring>/<statistics_sampling_period> <entity_monitoring>/<historical_statistics>
Route	Mutable properties in <property> (adapter specific) Mutable properties in <transformation>/<property> (transformation specific)	Immutable properties in <property> (adapter specific) Immutable properties in <transformation>/<property> (transformation specific)
Topic Route	Mutable QoS policies in <datawriter_qos> and <datareader_qos> Mutable properties in <transformation>/<property> (transformation specific) <route_types> <propagate_dispose> <propagate_unregister> <publish_with_original_info> <content_filter>/<parameter> <entity_monitoring>/<enabled> <entity_monitoring>/<status_publication_period>	Immutable QoS policies in <datawriter_qos> and <datareader_qos> <creation_mode> <content_filter>/<expression> <entity_monitoring>/<statistics_sampling_period> <entity_monitoring>/<historical_statistics>

Table 5.1 Changeable Parameters

Entity	Mutable (changeable any time)	Immutable (only changeable when entity is <i>disabled</i>) ¹
Auto Route	Mutable properties in <property> (adapter specific)	Immutable properties in <property> (adapter specific)
Auto Topic Route	Mutable QoS policies in <datawriter_qos> and <datareader_qos> <propagate_dispose> <propagate_unregister> <publish_with_original_info> <content_filter>/<parameter> <entity_monitoring>/<enabled> <entity_monitoring>/<status_publication_period>	Immutable QoS policies in <datawriter_qos> and <datareader_qos> <creation_mode> <allow_topic_name_filter> <allow_registered_type_name_filter> <deny_topic_name_filter> <deny_registered_type_name_filter> <content_filter>/<expression> <entity_monitoring>/<statistics_sampling_period> <entity_monitoring>/<historical_statistics>

1. Monitoring parameters can also be changed when monitoring is disabled (even when the entity is enabled).

You can send an XML snippet (or an assignment expression) that only contains the values you want to change for that entity, or you can send a whole well-formed configuration file.

- ❑ If you send an XML snippet (or an assignment expression), only the changes you specify will take effect.

For example, suppose you send this command:

```
update ShapeRouter DomainRoute1::Session1::SquareToCircles
  str://"<topic_route><input><datareader_qos><deadline><period>
  <sec>1</sec></period></deadline></datareader_qos></input>
  </topic_route>"
```

or

```
update ShapeRouter DomainRoute1::Session1::SquareToCircles
  topic_route.input.datareader_qos.deadline.period.sec = 1
```

The topic route **DomainRoute1::Session1::SquareToCircles** will only change the **period** value in the Deadline QoS for that particular DataReader.

Now suppose that later on you send this command:

```
update ShapeRouter DomainRoute1::Session1::SquareToCircles
  str://"<topic_route><input><datareader_qos><property>
  <value><element><name>MyProp</name><value>MyValueRemote</value>
```

```
</element></value></property><datareader_qos></input>
</topic_route>"
```

This would only change the Property QoS; the Deadline QoS would keep the setting from the prior command.

- ❑ If you send a well-formed configuration file (starting with `<dds><routing_service>`), the properties in the route (QoS values in the topic route) will be completely replaced with the properties (QoS values) defined in the XML code. If a QoS value for a topic route is not defined in the XML code, *Routing Service* will use the *Connex* default.

5.3 Accessing Routing Service from a Connex Application

You can send commands to control a *Routing Service* instance from your own *Connex* application. You will need to create a *DataWriter* for a specific topic and type. Then, you can send a sample that contains a command and its parameters. Optionally, you can create a *DataReader* for a specific topic to receive the results of the execution of your commands.

The topics are:

- ❑ `rti/routing_service/administration/command_request`
- ❑ `rti/routing_service/administration/command_response`

The types are:

- ❑ `RTI::RoutingService::Administration::CommandRequest`
- ❑ `RTI::RoutingService::Administration::CommandResponse`

You can find the IDL definitions for these types in `<Routing Service installation directory>/resource/idl/RoutingServiceAdministration.idl`.

The QoS configuration of your *DataWriter* and *DataReader* must be compatible with the one used by the routing service (see how this QoS is configured in [Section 2.4.3](#)).

When you send an XML string URL (`str://"<xml_code>"`) with the **load** and **update** commands, if the string is longer than `XML_URL_MAX_LENGTH` (in the IDL file), you will have to split the string and send several samples, setting the **"final"** field to false in all but the last sample.

The following example shows how to send a command to update the Deadline QoS policy for a topic route's *DataReader*:

```

/* Create entities: participant, publisher, topic, datawriter...*/
/* ... */
RTI_RoutingService_CommandRequest * cmdRequest =
    RTI_RoutingService_CommandRequestTypeSupport::create_data();
/* By specifying an unique ID for this command, you will be able
   to identify its response later on */
cmdRequest->id.host = /* host ID */;
cmdRequest->id.app = /* process ID */;
cmdRequest->id.invocation = ++invocationCounter;
/* Send this command to a routing service called MyRouter */
strcpy(cmdRequest->target_router, "MyRouter");
/* The command type is update */
cmdRequest->command._d = RTI_ROUTING_SERVICE_COMMAND_UPDATE;
/* Specify entity name to update and the XML code to define
   the new configuration */
strcpy(cmdRequest->command._u.entity_desc.name,
    "DomainRoute1::Session1::TopicRoute1");
/* When we use an XML snippet, the first tag we specify is
   that of the entity, <topic_route> in this case */
strcpy(cmdRequest->command._u.entity_desc.xml_url.content,
    "str://\"<topic_route>\
        <input>\
            <datareader_qos>\
                <deadline>\
                    <period>\
                        <sec>10</sec>\
                    </period>\
                </deadline>\
            </datareader_qos>\
        </input>\
    </topic_route>\"");
/* The content above is small enough to send it in one sample.
   Otherwise (if the length were > XML_URL_MAX_LENGTH) we would have
   to split it in multiple partial strings, each < XML_URL_MAX_LENGTH,
   and set final = 0 for all the samples but the last one */
cmdRequest->command._u.entity_desc.xml_url.final = 1;
RTI_RoutingService_CommandRequest_writer->write(cmdRequest, ...);

```


Chapter 6 Monitoring Routing Service from a Remote Location

You can monitor *Routing Service* remotely by subscribing to special *Connex*t topics. By subscribing to these topics, any *Connex*t application can receive information about the configuration and operational status of *Routing Service*.

Being able to monitor the state of a *Routing Service* instance is an important tool that allows you to detect problems. For example, looking at the latency statistics for a route might show you that the performance of a transformation in the route is not as expected. Looking at the input samples per second in the different sessions, you might see that one session is receiving most of the traffic. In that case, you could reassign some of the routes to other sessions to improve load balancing.

Routing Service can publish status for the following kinds of entities:

1. *Routing Service* itself (<routing_service>)
2. Domain Route <domain_route>)
3. Session (<session>)
4. Route (<route> and <topic_route>)
5. Topic Route (<auto_route> and <auto_topic_route>)

For each of the above kinds of entities, *Routing Service* creates two topics:

- ❑ **rti/routing_service/monitoring/<tag>_data** describes the entity's *configuration*
- ❑ **rti/routing_service/monitoring/<tag>_status_set** describes the entity's *operational status*

With the corresponding types:

- ❑ **RTI::RoutingService::Monitoring::<tag>Data**

❑ `RTI::RoutingService::Monitoring::<tag>StatusSet`

Where `<tag>` is one of the following entity kind tags: **RoutingService**, **DomainRoute**, **Session**, **Route**, or **AutoRoute**.

6.1 Enabling Remote Monitoring

By default, remote monitoring is disabled in *Routing Service* for security and performance reasons.

To enable remote monitoring, you can use the `<monitoring>` tag (see [Section 2.4.4](#)) or the `-remoteMonitoringDomainId` command-line parameter, which enables remote monitoring and sets the domain ID for data publication. For more information about the command-line options, see [Section 3.1 in the Getting Started Guide](#).

When remote monitoring is enabled, *Routing Service* creates:

- ❑ 1 `DomainParticipant`
- ❑ 1 `Publisher`
- ❑ 5 `DataWriters` for publishing *configuration data* (one for each kind of entity)
- ❑ 5 `DataWriters` for publishing *status* (one for each kind of entity).

The QoS values for these entities are described in [Section 2.4.4](#).

6.2 Monitoring Configuration Data

Configuration data for *Routing Service* entities is published in entity *data* topics. These topics are similar to the *Connext* builtin topics (`DCPSParticipant`, `DCPSPublication`, and `DCPSSubscription`) that provide information about the configuration of remote *Connext* entities.

This configuration data is published when:

- ❑ An entity is created or enabled.
- ❑ An entity is disabled or destroyed (a dispose message is published).
- ❑ The entity's configuration is modified using the remote command **"update"** (see [Section 5.2.11](#)).

- ❑ The entity's configuration is modified due to certain events in *Routing Service*. For example, discovery events may trigger the creation of StreamWriters and StreamReaders in a route.

The following sections describe the data available for each kind of *Routing Service* entity.

- ❑ Configuration Data for the Routing Service (Section 6.2.1)
- ❑ Configuration Data for a Domain Route (Section 6.2.2)
- ❑ Configuration Data for a Session (Section 6.2.3)
- ❑ Configuration Data for a Route (Section 6.2.4)
- ❑ Configuration Data for an Auto Topic Route (Section 6.2.5)

Each section describes the IDL for the topics' underlying data types. The IDL is also in the file `<Routing Service installation directory>/resource/idl/RoutingServiceMonitoring.idl`.

6.2.1 Configuration Data for the Routing Service

The topic that publishes configuration data is called `rti/routing_service/monitoring/routing_service_data`. This topic describes the configuration of the routing service but not its contained entities.

The IDL definition of the data type is:

```
struct RoutingServiceAdministrationData {
    string<EXPRESSION_MAX_LENGTH> save_path;
    boolean autosave_on_update;
};
struct RoutingServiceData {
    string<ENTITY_NAME_MAX_LENGTH> name; //@key
    string<ENTITY_NAME_MAX_LENGTH> group_name;
    string<ENTITY_NAME_MAX_LENGTH> host_name;
    long host_id;
    long app_id;
    RoutingServiceAdministrationData administration;
};
```

Table 6.1 on page 6-4 describes the members of the `RoutingServiceData` data type.

Table 6.1 RoutingServiceData

Field Name	Key	Description
name	✓	Name of the routing service instance. The name associated with the routing service instance can be assigned explicitly using the command-line parameter -appName . If -appName is not used, the <routing_service> tag name provided with -cfgName is used. If you use -identifyExecution , the host name and process ID are appended to the name. For example: RTI_RoutingService_myhost_1234
group_name		Name of the group to which the routing service belongs. Routing services in the same group will not communicate with each other. The group name is assigned using the attribute group_name in the <routing_service> tag. If the attribute is not defined, the group name is automatically set to RTI_RoutingService_<HostName>_<Process ID>
host_name		Name of the host where the routing service is running.
host_id		Identifies the host where the routing service instance is running.
app_id		Process (task) ID of the routing service instance.
administration. save_path		Specifies the file that will contain the saved configuration.
administration. auto_save_on_update		A boolean that, if true, automatically triggers a save command when configuration updates are received.

Routing service data samples are published when:

- The routing service instance is enabled.
- The routing service instance is disabled (dispose sample).
- Monitoring is enabled via remote administration.

6.2.2 Configuration Data for a Domain Route

The topic that publishes domain route configuration data is called **rti/routing_service/monitoring/domain_route_data**. The domain route data describes the configuration of the domain route and its connections but not its contained entities. Each connection can be defined with two different types, depending on if it is a *Connex* connection

(<participant_1> or <participant_2>) or a generic connection using an adapter (<connection_1 or <connection_2>).

The IDL definition of the data type `RTI::RoutingService::Monitoring::DomainRouteData` is:

```
struct DomainRouteParticipantData {
    long domain_id;
    BuiltinTopicKey_t participant_key;
};

struct DomainRouteAdapterConnectionData {
    string<ENTITY_NAME_MAX_LENGTH> plugin_name;
    sequence<Property, MAX_PROPERTIES> property;
};

union DomainRouteConnectionData switch(AdapterKind) {
    case RTI_ROUTING_SERVICE_DDS_ADAPTER_KIND:
        DomainRouteParticipantData dds;
    case RTI_ROUTING_SERVICE_GENERIC_ADAPTER_KIND:
        DomainRouteAdapterConnectionData generic;
};

struct DomainRouteData {
    string<ENTITY_NAME_MAX_LENGTH> routing_service_name; //@key
    string<ENTITY_NAME_MAX_LENGTH> name; //@key

    DomainRouteConnectionData connection_1;
    DomainRouteConnectionData connection_2;
};
```

Table 6.2 describes the members of the `DomainRouteData` data type.

Table 6.2 **DomainRouteData**

Field Name	Key	Description
routing_service_name	✓	The routing service name (assigned using <code>-appName</code>).
name	✓	The domain route name. This is configured using the <code>name</code> attribute in the <code><domain_route></code> tag.

Table 6.2 DomainRouteData

Field Name	Key	Description
connection_1		The configuration of the <connection_1> or <participant_1>. If it is a <connection_1>, the union discriminator is RTI_ROUTING_SERVICE_GENERIC_ADAPTER_KIND; if it is a <participant_1>, the union discriminator is RTI_ROUTING_SERVICE_DDS_ADAPTER_KIND
connection_2		The configuration of the <connection_2> or <participant_2>. If it is a <connection_2>, the union discriminator is RTI_ROUTING_SERVICE_GENERIC_ADAPTER_KIND; if it is a <participant_2>, the union discriminator is RTI_ROUTING_SERVICE_DDS_ADAPTER_KIND
connection_1.generic.plugin_name		The name of the plugin used by the first connection (<connection_1>)
connection_2.generic.plugin_name		The name of the plugin used by the second connection (<connection_2>)
connection_1.generic.property		The sequence of properties defined in the tag <property> inside <connection_1>
connection_2.generic.property		The sequence of properties defined in the tag <property> inside <connection_2>
cconnection_1.dds.domain_id		Domain ID of the first domain route participant. This domain ID is configured using the XML tag <domain_id> inside <participant_1>.
cconnection_2.dds.domain_id		Domain ID of the second domain route participant. This domain ID is configured using the XML tag <domain_id> inside <participant_1>.
connection_1.dds.participant_key		Unique identifier for the first participant.
connection_2.dds.participant_key		Unique identifier for the second participant.

A domain route using *Connex*t can be correlated with its corresponding participants using the fields `connection_1.dds.participant_key` and/or `connection_2.dds.participant_key`.

For example, let's assume that we want to get the value of the `PropertyQosPolicy` associated with the first `DomainParticipant` of a domain route. To do that, we would subscribe

to the participant *Connex* builtin topic and look for a sample where the key member is equal to `participant1_key`. From this sample, we can get the `PropertyQosPolicy` by accessing the member called **property**.

For additional information on how to subscribe to *Connex* builtin topics, see the *RTI Core Libraries and Utilities User's Manual*¹.

Domain-route data samples are published when:

- The domain route is enabled.
- The domain route is disabled (dispose sample).
- Monitoring is enabled via remote administration.

6.2.3 Configuration Data for a Session

The topic that publishes session configuration data is called `rti/routing_service/monitoring/session_data`. The session data describes the configuration of the session but not its contained entities.

The IDL definition of the data type is:

```
struct SessionData {
    string<ENTITY_NAME_MAX_LENGTH> routing_service_name; //@key
    string<ENTITY_NAME_MAX_LENGTH> domain_route_name; //@key
    string<ENTITY_NAME_MAX_LENGTH> name; //@key

    long enabled_route_count;

    sequence<Property, MAX_PROPERTIES> property;
};
```

Table 6.3 describes the fields in the `SessionData` data type.

Table 6.3 **SessionData**

Field Name	Key	Description
<code>routing_service_name</code>	✓	The routing service name (assigned using <code>-appName</code>).
<code>domain_route_name</code>	✓	The domain route name.
<code>name</code>	✓	The session name, which is configured with the name attribute in the <code><session></code> tag.

1. See `<Connex installation directory>/ndds.<version>/doc/pdf/RTI_Connex_UsersManual.pdf`.

Table 6.3 **SessionData**

Field Name	Key	Description
enabled_route_count		The number of enabled routes.
property		The sequence of properties defined in the tag <property> inside <session>

Session data samples are published when:

- The session is enabled.
- The session is disabled (dispose sample).
- An auto route/route inside the session is enabled.
- An auto route/route inside the session is disabled.
- Monitoring is enabled via remote administration.

6.2.4 Configuration Data for a Route

The topic that publishes route configuration data is called **rti/routing_service/monitoring/route_data**.

The IDL definition of the data type `RTI::RoutingService::Monitoring::RouteData` is:

```

struct TransformationData {
    string<ENTITY_NAME_MAX_LENGTH> plugin_name;
    sequence<Property, MAX_PROPERTIES> property;
};

struct RouteAdapterData {
    sequence<Property, MAX_PROPERTIES> property;
};

struct RouteDdsInputData {
    long domain_id;
    BuiltinTopicKey_t datareader_key;
    string<EXPRESSION_MAX_LENGTH> content_filter_expression;
};

```

```
union RouteInputAdapterData switch(AdapterKind) {
    case RTI_ROUTING_SERVICE_DDS_ADAPTER_KIND:
        RouteDdsInputData dds;
    case RTI_ROUTING_SERVICE_GENERIC_ADAPTER_KIND:
        RouteAdapterData generic;
};

struct RouteDdsOutputData {
    long domain_id;
    BuiltinTopicKey_t datawriter_key;
};

union RouteOutputAdapterData switch(AdapterKind) {
    case RTI_ROUTING_SERVICE_DDS_ADAPTER_KIND:
        RouteDdsOutputData dds;
    case RTI_ROUTING_SERVICE_GENERIC_ADAPTER_KIND:
        RouteAdapterData generic;
};

struct RouteInputData {
    string<TOPIC_NAME_MAX_LENGTH> stream_name;
    string<TYPE_NAME_MAX_LENGTH> registered_type_name;
    long connection;
    RouteCreationMode creation_mode;
    DDSEntityState state;
    RouteInputAdapterData adapter_data;
};

struct RouteOutputData {
    string<TOPIC_NAME_MAX_LENGTH> stream_name;
    string<TYPE_NAME_MAX_LENGTH> registered_type_name;
    RouteCreationMode creation_mode;
    DDSEntityState state;
    RouteOutputAdapterData adapter_data;
};
```

```

struct RouteData {
    string<ENTITY_NAME_MAX_LENGTH> routing_service_name; // @key
    string<ENTITY_NAME_MAX_LENGTH> domain_route_name; // @key
    string<ENTITY_NAME_MAX_LENGTH> session_name; // @key
    string<ENTITY_NAME_MAX_LENGTH> name; // @key

    string<ENTITY_NAME_MAX_LENGTH> auto_route_name;

    boolean propagate_dispose;
    boolean propagate_unregister;
    boolean publish_with_original_info;
    boolean publish_with_original_timestamp;

    boolean route_types;

    RouteInputData input;
    RouteOutputData output;

    sequence<TransformationData, MAX_TRANSFORMATIONS> transformations;
};

```

Table 6.4 describes the fields in the RouteData topic data type.

Table 6.4 **RouteData**

Field Name	Key	Description
routing_service_name	✓	The routing service name (assigned with -appName).
domain_route_name	✓	The domain route name.
session_name	✓	The session name.
name	✓	The route name, which is configured using the name attribute in the <route> or <topic_route> tag.
auto_route_name		If the route is contained in an auto-route, this field contains the auto-route name. Otherwise, the field is initialized with the empty string.
input		The configuration of the route's input, as contained in the tag <input> or <dds_input> inside <route> or <topic_route>
output		The configuration of the route's output, as contained in the tag <output> or <dds_output> inside <route> or <topic_route>

Table 6.4 RouteData

Field Name	Key	Description
input.connection		Index of the input connection or participant (1 or 2). The value of this field is used to determine whether the input of this route is the domain route's connection 1/participant 1 or the connection 2/participant 2.
input.stream_name		Input stream name. The input stream name is configured using the tag <topic_name> inside <topic_route>/<input> or inside <route>/<dds_input> or the tag <stream_name> inside <route>/<input>.
input.registered_type_name		Input registered name. The input registered name is configured using the tag <registered_type_name> inside <topic_route>/<input>, <route>/<dds_input> or <route>/<input>.
output.stream_name		Output stream name. The output stream name is configured using the tag <topic_name> inside <topic_route>/<output> or inside <route>/<dds_output> or the tag <stream_name> inside <route>/<output>.
output.registered_type_name		Output registered name. The output registered name is configured using the tag <registered_type_name> inside <topic_route>/<output>, <route>/<dds_output> or <route>/<output>.
input.creation_mode		Indicates when the StreamReader is created in the input. The input creation mode is configured using the tag <creation_mode>.
output.creation_mode		Indicates when the StreamWriter is created in the output. The output creation mode is configured using the tag <creation_mode>.
input.adapter_data		Contains the configuration of the route's input that is specific to either the <i>Connex</i> adapter or a generic adapter. When a generic input is defined (<route>/<input>) then the union discriminator is RTI_ROUTING_SERVICE_GENERIC_ADAPTER_KIND; if it is a <i>Connex</i> input (<topic_route>/<input> or <route>/<dds_input>), then the union discriminator is RTI_ROUTING_SERVICE_DDS_ADAPTER_KIND

Table 6.4 RouteData

Field Name	Key	Description
output.adapter_data		Contains the configuration of the route's output that is specific to either the <i>Connex</i> t adapter or a generic adapter. When a generic output is defined (<route>/<output>) then the union discriminator is RTI_ROUTING_SERVICE_GENERIC_ADAPTER_KIND; if it is a <i>Connex</i> t output (<topic_route>/<output> or <route>/<dds_output>), then the union discriminator is RTI_ROUTING_SERVICE_DDS_ADAPTER_KIND
input.adapter_data.dds.domain_id		(<i>Connex</i> t input only) Domain ID of the input participant
output.adapter_data.dds.domain_id		(<i>Connex</i> t output only) Domain ID of the output participant
input.adapter_data.dds.datareader_key		(<i>Connex</i> t input only) Unique identifier for the DataReader. The value of this field is meaningful only when the state is RTI_ROUTING_SERVICE_CREATED_AND_ENABLED.
output.adapter_data.dds.datawriter_key		(<i>Connex</i> t output only) Unique identifier for the DataWriter. The value of this field is only meaningful when datawriter_state is RTI_ROUTING_SERVICE_CREATED_AND_ENABLED.
input.adapter_data.dds.content_filter_expression		(<i>Connex</i> t input only) Content filter expression associated with the content filter for the topic route DataReader. The expression is configured using the tag <content_filter>/<expression> inside <topic_route>/<input> or <route>/<dds_input>
input.adapter_data.generic.property		(Not applicable for <i>Connex</i> t input) The properties used to configure this route's StreamReader, specified with the tag <property> inside <route>/<input>
output.adapter_data.generic.property		(Not applicable for <i>Connex</i> t output) The properties used to configure this route's StreamWriter, specified with the tag <property> inside <route>/<output>
propagate_dispose		(<i>Connex</i> t topic routes only) Indicates if the topic route propagates NOT_ALIVE_DISPOSE samples. The propagation of NOT_ALIVE_DISPOSE samples is configured using the tag <propagate_dispose> in <topic_route>.

Table 6.4 RouteData

Field Name	Key	Description
propagate_unregister		(<i>Connex</i> t topic routes only) Indicates if the topic route propagates NOT_ALIVE_NO_WRITERS samples. The propagation of NOT_ALIVE_NO_WRITERS samples is configured using the tag <propagate_unregister> in <topic_route>.
publish_with_original_info		(<i>Connex</i> t topic routes only) Indicates if the topic route publishes the samples with original writer info. Setting this option to true allows redundant topic routes and prevents the applications from receiving duplicate samples. The publication with original writer info is configured using the tag <publish_with_original_info> inside <topic_route>.
publish_with_original_timestamp		Indicates if the route is configured to publish the output samples with the same timestamp as that of the input sample.
route_types		Indicates if the input connection will use types discovered in the output connection and viceversa for the creation of StreamWriters and StreamReaders. The route types flag is configured using the tag <route_types> inside <route> or <topic_route>.
transformations		List of transformations associated with a route. For each transformation you will be able to retrieve the transformation plugin name, and the properties. Transformations are defined using the <transformation> tag inside <route> or <topic_route>. Note: in this version, only one transformation per route is supported.
input.state		Indicates whether or not the StreamReader associated with a route is created.
output.state		Indicates whether or not the StreamWriter associated with a route is created.

The correlation between a route using *Connex*t and its DataReader and DataWriter can be done using the fields **datareader_key** and **datawriter_key**.

For example, let's assume that we want to retrieve the value of the DurabilityQosPolicy associated with the route's DataWriter. To do that, we would subscribe to the publication *Connex*t builtin topic and we would look for a sample where the key member is equal to **datawriter_key**. From this sample, we can get the DurabilityQosPolicy value accessing the member durability.

For additional information on how to subscribe to the *Connex* builtin topics, see the *RTI Core Libraries and Utilities User's Manual*¹.

Route data samples are published when:

- The route is enabled.
- The route is disabled (dispose sample).
- The route configuration is modified using the remote command update.
- The route's `StreamReader` is created.
- The route's `StreamReader` is destroyed.
- The route's `StreamWriter` is created.
- The route's `StreamWriter` is destroyed.
- Monitoring is enabled via remote administration.

6.2.5 Configuration Data for an Auto Topic Route

The topic that publishes auto route configuration data is called **`rti/routing_service/monitoring/auto_route_data`**.

The IDL definition of the data type `RTI::RoutingService::Monitoring::AutoRouteData` is:

```
struct AutoRouteAdapterData {
    sequence<Property, MAX_PROPERTIES> property;
};

struct AutoRouteDdsInputData {
    long domain_id;
    string<EXPRESSION_MAX_LENGTH> content_filter_expression;
};

union AutoRouteInputAdapterData switch(AdapterKind) {
    case RTI_ROUTING_SERVICE_DDS_ADAPTER_KIND:
        AutoRouteDdsInputData dds;
    case RTI_ROUTING_SERVICE_GENERIC_ADAPTER_KIND:
        AutoRouteAdapterData generic;
};

struct AutoRouteDdsOutputData {
    long domain_id;
```

1. See `<Connex installation directory>/ndds.<version>/doc/pdf/RTI_Connext_UsersManual.pdf`.

```

};

union AutoRouteOutputAdapterData switch(AdapterKind) {
    case RTI_ROUTING_SERVICE_DDS_ADAPTER_KIND:
        AutoRouteDdsOutputData dds;
    case RTI_ROUTING_SERVICE_GENERIC_ADAPTER_KIND:
        AutoRouteAdapterData generic;
};

struct AutoRouteInputData {
    string<TOPIC_NAME_MAX_LENGTH> allow_stream_name_filter;
    string<TYPE_NAME_MAX_LENGTH> allow_registered_type_name_filter;
    string<TOPIC_NAME_MAX_LENGTH> deny_stream_name_filter;
    string<TYPE_NAME_MAX_LENGTH> deny_registered_type_name_filter;
    long connection;
    RouteCreationMode creation_mode;
    AutoRouteInputAdapterData adapter_data;
};

struct AutoRouteOutputData {
    string<TOPIC_NAME_MAX_LENGTH> allow_stream_name_filter;
    string<TYPE_NAME_MAX_LENGTH> allow_registered_type_name_filter;
    string<TOPIC_NAME_MAX_LENGTH> deny_stream_name_filter;
    string<TYPE_NAME_MAX_LENGTH> deny_registered_type_name_filter;

    RouteCreationMode creation_mode;
    AutoRouteOutputAdapterData adapter_data;
};

struct AutoRouteData {
    string<ENTITY_NAME_MAX_LENGTH> routing_service_name; //@key
    string<ENTITY_NAME_MAX_LENGTH> domain_route_name; //@key
    string<ENTITY_NAME_MAX_LENGTH> session_name; //@key
    string<ENTITY_NAME_MAX_LENGTH> name; //@key

    boolean propagate_dispose;
    boolean propagate_unregister;
    boolean publish_with_original_info;
    boolean publish_with_original_timestamp;
    long enabled_route_count;
    AutoRouteInputData input;
    AutoRouteOutputData output;
};

```

Table 6.5 describes the fields in the AutoRouteData data type.

Auto-topic-route data samples are published when:

- The auto route is enabled.
- The auto route is disabled (dispose sample).
- The auto route configuration is modified using the remote command update.
- A new route is created from the auto route
- Monitoring is enabled via remote administration.

Table 6.5 **AutoRouteData**

Field Name	Key	Description
routing_service_name	✓	The routing service name (assigned using <code>-appName</code>).
domain_route_name	✓	The domain route name.
session_name	✓	The session name.
name	✓	The auto route name, which is configured using the name attribute in the <code><auto_route></code> or <code><auto_topic_route></code> tags.
input		The configuration of the auto route input, as contained in the tag <code><input></code> or <code><dds_input></code> inside <code><auto_route></code> or <code><auto_topic_route></code>
output		The configuration of the auto route output, as contained in the tag <code><output></code> or <code><dds_output></code> inside <code><auto_route></code> or <code><auto_topic_route></code>
input.connection		Index of the input connection or participant (1 or 2). The value of this field is used to determine whether the input of this auto route is the domain route's connection 1/participant 1 or the connection 2/participant 2.
input.allow_stream_name_filter		Topics that do not pass this filter in the input participant will not trigger the creation of routes. This filter is configured using the tag <code><allow_topic_name_filter></code> inside <code><auto_topic_route>/<input></code> or inside <code><auto_route>/<dds_input></code> or the tag <code><allow_stream_name_filter></code> inside <code><auto_route>/<input></code>
input.allow_registered_type_name_filter		Topic with types that do not pass this filter in the input participant will not trigger the creation of routes. This filter is configured using the tag <code><allow_registered_type_name_filter></code> inside <code><auto_topic_route>/<input></code> , <code><auto_route>/<dds_input></code> , or <code><auto_route>/<input></code> .

Table 6.5 AutoRouteData

Field Name	Key	Description
output.allow_stream_name_filter		Topics that do not pass this filter in the output participant will not trigger the creation of routes. This filter is configured using the tag <allow_topic_name_filter> inside <auto_topic_route>/<output> or inside <auto_route>/<dds_output> or the tag <allow_stream_name_filter> inside <auto_route>/<output>.
output.allow_registered_type_name_filter		Topics with types that do not pass this filter in the output participant will not trigger the creation of routes. This filter is configured using the tag <allow_registered_type_name_filter> inside <auto_topic_route>/<output>, <auto_route>/<dds_ioutput>, or <auto_route>/<output>.
input.deny_stream_name_filter		Topics that pass this filter in the input participant will not trigger the creation of routes. This filter is configured using the tag <deny_topic_name_filter> inside <auto_topic_route>/<input>. or inside <auto_route>/<dds_input> or the tag <deny_stream_name_filter> inside <auto_route>/<input>.
input.deny_registered_type_name_filter		Topics with types that pass this filter in the input participant will not trigger the creation of routes. The input deny registered type name filter is configured using the tag <deny_registered_type_name_filter> inside <auto_topic_route>/<input>, <auto_route>/<dds_input>, or <auto_route>/<input>.
output.deny_stream_name_filter		Topics that pass this filter in the output participant will not trigger the creation of routes. The output deny topic name filter is configured using the tag <deny_topic_name_filter> inside <auto_topic_route>/<output>. or inside <auto_route>/<dds_output> or the tag <deny_stream_name_filter> inside <auto_route>/<output>.
output.deny_registered_type_name_filter		Topics with types that pass this filter in the output participant will not trigger the creation of routes. The output deny registered type name filter is configured using the tag <deny_registered_type_name_filter> inside <auto_topic_route>/<output>, <auto_route>/<dds_output> , or <auto_route>/<output>.

Table 6.5 **AutoRouteData**

Field Name	Key	Description
input.creation_mode		Indicates when the StreamReader is created in the input. The input creation mode is configured using the tag <creation_mode>.
output.creation_mode		Indicates when the StreamWriter is created in the output. The output creation mode is configured using the tag <creation_mode>..
input.adapter_data		Contains the configuration of the auto route's input that is specific to either the <i>Connex</i> t adapter or a generic adapter. When a generic input is defined (<auto_route>/<input>), the union discriminator is RTI_ROUTING_SERVICE_GENERIC_ADAPTER_KIND; if it is a <i>Connex</i> t input (<auto_topic_route>/<input> or <auto_route>/<dds_input>), the union discriminator is RTI_ROUTING_SERVICE_DDS_ADAPTER_KIND.
output.adapter_data		Contains the configuration of the auto_route's output that is specific to either the <i>Connex</i> t adapter or a generic adapter. When a generic output is defined (<auto_route>/<output>), the union discriminator is RTI_ROUTING_SERVICE_GENERIC_ADAPTER_KIND; if it is a <i>Connex</i> t output (<auto_topic_route>/<output> or <auto_route>/<dds_output>), the union discriminator is RTI_ROUTING_SERVICE_DDS_ADAPTER_KIND.
input.adapter_data.dds.domain_id		(<i>Connex</i> t input only) Domain ID of the input participant
output.adapter_data.dds.domain_id		(<i>Connex</i> t output only) Domain ID of the output participant
input.adapter_data.dds.content_filter_expression		(<i>Connex</i> t input only) Content filter expression associated with the content filter for the topic route DataReader. The expression is configured using the tag <content_filter>/<expression> inside <topic_route>/<input> or <route>/<dds_input>
input.adapter_data.generic.property		(Not applicable for <i>Connex</i> t input) The properties used to configure this route's StreamReader, specified with the tag <property> inside <route>/<input>

Table 6.5 **AutoRouteData**

Field Name	Key	Description
output.adapter_data. generic.property		(Not applicable for <i>Connex</i> output) The properties used to configure this route's StreamWriter, specified with the tag <property> inside <route>/<output>
propagate_dispose		(<i>Connex</i> auto_topic routes only) Indicates if the topic route propagates NOT_ALIVE_DISPOSE samples. The propagation of NOT_ALIVE_DISPOSE samples is configured using the tag <propagate_dispose> in <topic_route>.
propagate_unregister		(<i>Connex</i> auto_topic routes only) Indicates if the topic routes propagate NOT_ALIVE_NO_WRITERS samples. The propagation of NOT_ALIVE_NO_WRITERS samples is configured using the tag <propagate_unregister> in <auto_topic_route>.
publish_with_original_info		(<i>Connex</i> auto_topic routes only) Indicates if the topic routes publish the samples with original writer information. Setting this option to true allows redundant topic routes and prevents applications from receiving duplicate samples. The publication with original writer info is configured using the tag <publish_with_original_info> inside <auto_topic_route>.
publish_with_original_timestamp		Indicates if the routes are configured to publish the output samples with the same timestamp as that of the input sample.
enabled_route_count		The number of enabled routes associated with the auto route.

6.3 Monitoring Status

Operational status for *Routing Service* entities is published in entity **status_set** topics. This information changes continuously and is computed and published periodically.

The status information for the different entities is composed primarily of statistics. [Section 6.3.1](#) explains how these statistics are calculated and published. These sections describe the status information associated with each kind of entity:

- ❑ [Status Information for the Routing Service \(Section 6.3.2\)](#)
- ❑ [Domain Route Status \(Section 6.3.3\)](#)
- ❑ [Status Information for a Session \(Section 6.3.4\)](#)

- ❑ [Status Information for a Route \(Section 6.3.5\)](#)
- ❑ [Status Information for an Auto Route \(Section 6.3.6\)](#)

Each section describes the IDL for the topics' underlying data types. The IDL is also in the file `<Routing Service installation directory>/resource/idl/RoutingServiceMonitoring.idl`.

6.3.1 How the Statistics are Generated

6.3.1.1 Statistics Publication

Routing Service reports multiple statistics as part of the different status sets. For example, for a route the status contains statistical metrics about the input and output samples per second (throughput).

```
struct RouteStatusSet {  
    ...  
    StatisticVariable input_samples_per_s;  
    StatisticVariable output_samples_per_s;  
    ...  
};
```

The statistical information is published periodically in the form **StatisticVariables**.

The period at which statistics are published is configurable using the tag `<status_publication_period>` (see [Section 2.4.4](#)).

For a given variable, *Routing Service* computes the metrics in **StatisticMetrics** during specific time frames.

```
struct StatisticMetrics {  
    unsigned long long period_ms;  
    long long count;  
    float mean;  
    float minimum;  
    float maximum;  
    float std_dev;  
};  
  
struct StatisticVariable {  
    StatisticMetric publication_period_metrics;  
    sequence<StatisticMetrics, MAX_HISTORICAL_METRICS> historical_metrics;  
};
```


The **count** is the sum of all the values received during the time frame. For example, in the case of **input_sample_per_s** and **output_sample_p_s**, **count** is the number of samples received during the time frame. For latency, **count** is the sum of all the latency times for the samples received during the time frame.

If status publication is enabled (see [Section 2.4.4](#)), *Routing Service* always publishes the statistics corresponding to the time between two status publications (**publication_period_metrics**). You can also select additional windows on a per entity basis using the tag `<historical_statistics>` (see [Section 2.4.4](#)). The sequence **historical_metrics** in **StatisticVariable** contains values corresponding to the windows that have been enabled:

- 5-sec. metrics correspond to activity in the last five seconds.
- 1-min. metrics correspond to activity in the last minute.
- 5-min. metrics correspond to activity in the last five minutes.
- 1-hour metrics correspond to activity in the last hour.
- Up-time metrics correspond to activity since the entity was enabled.

Each window has a field called **period_ms** that identifies its size in milliseconds. For the **publication_period_metrics**, this field contains the publication period. For the up-time metrics, this field contains the time since the entity was enabled. For the other windows, this field contains a fixed value that identifies the window size (5000 for the 5-second window, 60000 for the one-minute window, etc).

6.3.1.2 Statistics Calculation

The accuracy of the statistics calculation process is determined by the value of the statistics sampling period. This period specifies how often statistics are gathered and is configured on a per entity basis using the tag `<statistics_sampling_period>` (see [Section 2.4.4](#)).

As a general rule, the **statistics_sampling_period** of an entity must be smaller than its **status_publication_period**. A small **statistics_sampling_period** provides more accurate statistics at expense of increasing the memory consumption and decreasing performance.

6.3.2 Status Information for the Routing Service

The topic that publishes routing service status is called **rti/routing_service/monitoring/routing_service_status_set**.

The IDL definition of the data type is:

```
struct RoutingServiceStatusSet {
    string<ENTITY_NAME_MAX_LENGTH> name; //@key

    StatisticVariable cpu_usage_percentage;
    StatisticVariable physical_memory_kb;
    StatisticVariable total_memory_kb;
};
```

Table 6.6 describes the fields in the **RoutingServiceStatusSet** data type.

Table 6.6 **RoutingServiceStatusSet**

Field Name	Key	Description
name	✓	Name of the routing service instance. The name associated with the routing service instance can be assigned explicitly by using the -appName command-line parameter. If -appName is not used, the <routing_service> tag name provided with -cfgName is used. If you use the -identifyExecution command-line parameter, the host name and the process ID are appended to the name. For example: RTI_RoutingService_myhost_1234
cpu_usage_percentage		Statistic variable that provides the percentage of CPU usage of the routing service process over different time windows. This variable is only supported on Windows and Linux systems.
physical_memory_kb		Statistic variable that provides the physical memory utilization of the routing service process. This variable is only supported on Windows and Linux systems.
total_memory_kb		Statistic variable that provides the virtual memory utilization of the routing service process. This variable is only supported on Windows and Linux systems.

6.3.3 Domain Route Status

The topic that publishes domain route status is called `rti/routing_service/monitoring/domain_route_status_set`.

The domain route status aggregates the statistics of the routes contained in it: the mean of the means in the routes, the absolute maximum and minimum across routes, the mean of the standard deviation and the total count.

The IDL definition of the data type is:

```
struct DomainRouteStatusSet {
    string<ENTITY_NAME_MAX_LENGTH> routing_service_name; //@key
    string<ENTITY_NAME_MAX_LENGTH> name; //@key

    StatisticVariable input_samples_per_s;
    StatisticVariable input_bytes_per_s;
    StatisticVariable output_samples_per_s;
    StatisticVariable output_bytes_per_s;
    StatisticVariable latency_s;
};
```

Table 6.7 describes the fields in the `DomainRouteStatusSet` data type.

Table 6.7 `DomainRouteStatusSet`

Field Name	Key	Description
<code>routing_service_name</code>	✓	The routing service name (assigned with <code>-appName</code>).
<code>name</code>	✓	The domain route name. The domain route name is configured using the <code>name</code> attribute in the <code><domain_route></code> tag.
<code>input_samples_per_s</code>		Statistic variable that provides information about the input samples per second across routes. Input samples refer to the samples that are taken by the sessions from the routes's <code>StreamReaders</code> .
<code>input_bytes_per_s¹</code>		Statistic variable that provides information about the input bytes per second across routes. Input bytes refer to the bytes that are taken by the sessions from the routes's <code>StreamReaders</code> . These bytes only refer to the serialized samples. The protocol headers (UDP, RTPS) are not included.

Table 6.7 **DomainRouteStatusSet**

Field Name	Key	Description
output_samples_per_s		Statistic variable that provides information about the output samples per second across routes. Output samples refer to the samples that are published out by the session threads using the route's StreamWriters.
output_bytes_per_s		Statistic variable that provides information about the output bytes per second across routes. Output bytes refer to the bytes that are published out by the session threads using the route's StreamWriters. The variable only considers the bytes of the serialized samples. Protocol headers (UDP, RTPS) are not included.
latency_s		Statistic variable that provides information about the latency in seconds across routes. The latency in a route refers to the time elapsed between the sample read and write. This is a good metric to monitor the health and performance of transformations.

1. The throughput measured in bytes can only be computed if the samples are DynamicData samples. If not, only the throughput measured in samples per second is available. This statement applies to all the statistic variables described in this chapter that measure throughput in bytes per second.

6.3.4 Status Information for a Session

The topic that publishes session status is called **rti/routing_service/monitoring/session_status_set**.

The session status aggregates the statistics of the routes contained in it: the mean of the means in the routes, the absolute maximum and minimum across routes, the mean of the standard deviation and the total count.

The IDL definition of the data type is:

```
struct SessionStatusSet {
    string<ENTITY_NAME_MAX_LENGTH> routing_service_name; //@key
    string<ENTITY_NAME_MAX_LENGTH> domain_route_name; //@key
    string<ENTITY_NAME_MAX_LENGTH> name; //@key

    StatisticVariable input_samples_per_s;
    StatisticVariable input_bytes_per_s;
    StatisticVariable output_samples_per_s;
    StatisticVariable output_bytes_per_s;
```

```
StatisticVariable latency_s;
};
```

Table 6.8 describes the fields in the **SessionStatusSet** data type.

Table 6.8 **SessionStatusSet**

Field Name	Key	Description
routing_service_name	✓	The routing service name (assigned with <code>-appName</code>).
domain_route_name	✓	The domain route name
name	✓	The session name. The domain route name is configured using the name attribute in the <code><session></code> tag.
input_samples_per_s		Statistic variable that provides information about the input samples per second across routes. Input samples refer to the samples that are taken by the session from the routes's StreamReaders.
input_bytes_per_s		Statistic variable that provides information about the input bytes per second across routes. Input bytes refer to the bytes that are taken by the sessions from the routes's StreamReaders. These bytes only refer to the serialized samples. The protocol headers (UDP, RTPS) are not included.
output_samples_per_s		Statistic variable that provides information about the output samples per second across routes. Output samples refer to the samples that are published out by the session thread using the route's StreamWriters.
output_bytes_per_s		Statistic variable that provides information about the output bytes per second across routes. Output bytes refer to the bytes that are published out by the session thread using the route's StreamWriters. The variable only considers the bytes of the serialized samples. Protocol headers (UDP, RTPS) are not included.
latency_s		Statistic variable that provides information about the latency in seconds across routes. The latency in a route refers to the time elapsed between the sample read and write. This is a good metric to monitor the health and performance of transformations.

6.3.5 Status Information for a Route

The topic that publishes route status is called `rti/routing_service/monitoring/route_status_set`.

The IDL definition of the data type is:

```
struct RouteStatusSet {
    string<ENTITY_NAME_MAX_LENGTH> routing_service_name; //@key
    string<ENTITY_NAME_MAX_LENGTH> domain_route_name; //@key
    string<ENTITY_NAME_MAX_LENGTH> session_name; //@key
    string<ENTITY_NAME_MAX_LENGTH> name; //@key

    StatisticVariable input_samples_per_s;
    StatisticVariable input_bytes_per_s;
    StatisticVariable output_samples_per_s;
    StatisticVariable output_bytes_per_s;
    StatisticVariable latency_s;
};
```

Table 6.9 describes the fields in the `RouteStatusSet` data type.

Table 6.9 `RouteStatusSet`

Member Name	Key	Description
<code>routing_service_name</code>	✓	The routing service name (assigned with <code>-appName</code>).
<code>domain_route_name</code>	✓	The domain route name
<code>session_name</code>	✓	The session name.
<code>name</code>	✓	The route name. The route name is configured using the <code>name</code> attribute in the <code><topic_route></code> or <code><route></code> tags.
<code>input_samples_per_s</code>		Statistic variable that provides information about the input samples per second in the route. Input samples refer to the samples that are taken by the session from the route's <code>StreamReader</code> .
<code>input_bytes_per_s</code>		Statistic variable that provides information about the input bytes per second in the route. Input bytes refer to the bytes that are taken by the session from the route's <code>StreamReader</code> . These bytes only refer to the serialized samples. The protocol headers (UDP, RTPS) are not included.

Table 6.9 RouteStatusSet

Member Name	Key	Description
output_samples_per_s		Statistic variable that provides information about the output samples per second in the routes. Output samples refer to the samples that are published out by the session thread using the route's StreamWriters.
output_bytes_per_s		Statistic variable that provides information about the output bytes per second in routes. Output bytes refer to the bytes that are published out by the session thread using the route's StreamWriter. The variable only considers the bytes of the serialized samples. Protocol headers (UDP, RTPS) are not included.
latency_s		Statistic variable that provides information about the latency in seconds in the routes. The latency in a route refers to the time elapsed between the sample read and write. This is a good metric to monitor the health and performance of transformations.

6.3.6 Status Information for an Auto Route

The topic that publishes auto route status is called **rti/routing_service/monitoring/route_status_set**.

The auto route status aggregates the statistics of the routes created from it: the mean of the means in the routes, the absolute maximum and minimum across routes, the mean of the standard deviation and the total count.

The IDL definition of the data type is:

```
struct AutoRouteStatusSet {
    string<ENTITY_NAME_MAX_LENGTH> routing_service_name; // @key
    string<ENTITY_NAME_MAX_LENGTH> domain_route_name; // @key
    string<ENTITY_NAME_MAX_LENGTH> session_name; // @key
    string<ENTITY_NAME_MAX_LENGTH> name; // @key

    StatisticVariable input_samples_per_s;
    StatisticVariable input_bytes_per_s;
    StatisticVariable output_samples_per_s;
    StatisticVariable output_bytes_per_s;
    StatisticVariable latency_s;
};
```

Table 6.10 describes the fields in the **AutoRouteStatusSet** data type.

Table 6.10 **AutoRouteStatusSet**

Member Name	Key	Description
routing_service_name	✓	The routing service name (assigned with -appname).
domain_route_name	✓	The domain route name.
session_name	✓	The session name.
name	✓	The auto route name. The auto route name is configured using the name attribute in the <auto_topic_route> or <auto_route> tags.
input_samples_per_s		Statistic variable that provides information about the input samples per second across routes. Input samples refer to the samples that are taken by the session from the auto routes's StreamReaders.
input_bytes_per_s		Statistic variable that provides information about the input bytes per second across routes. Input bytes refer to the bytes that are taken by the session from the auto routes's StreamReaders. These bytes only refer to the serialized samples. The protocol headers (UDP, RTPS) are not included.
output_samples_per_s		Statistic variable that provides information about the output samples per second across routes. Output samples refer to the samples that are published out by the session thread using the auto route's StreamWriters.
output_bytes_per_s		Statistic variable that provides information about the output bytes per second across routes. Output bytes refer to the bytes that are published out by the session thread using the auto route's StreamWriters. The variable only considers the bytes of the serialized samples. Protocol headers (UDP, RTPS) are not included.
latency_s		Statistic variable that provides information about the latency in seconds across routes. The latency in a route refers to the time elapsed between the sample read and write. This is a good metric to monitor the health and performance of transformations.

Chapter 7 Traversing Wide Area Networks

Many systems today already rely on *Connex*t to distribute their information across a Local Area Network (LAN). However, more and more of these systems are being integrated in Wide Area Networks (WANs). With *Routing Service*, you can scale *Connex*t real-time publish/subscribe data-distribution beyond the current local networks and make it available throughout a WAN.

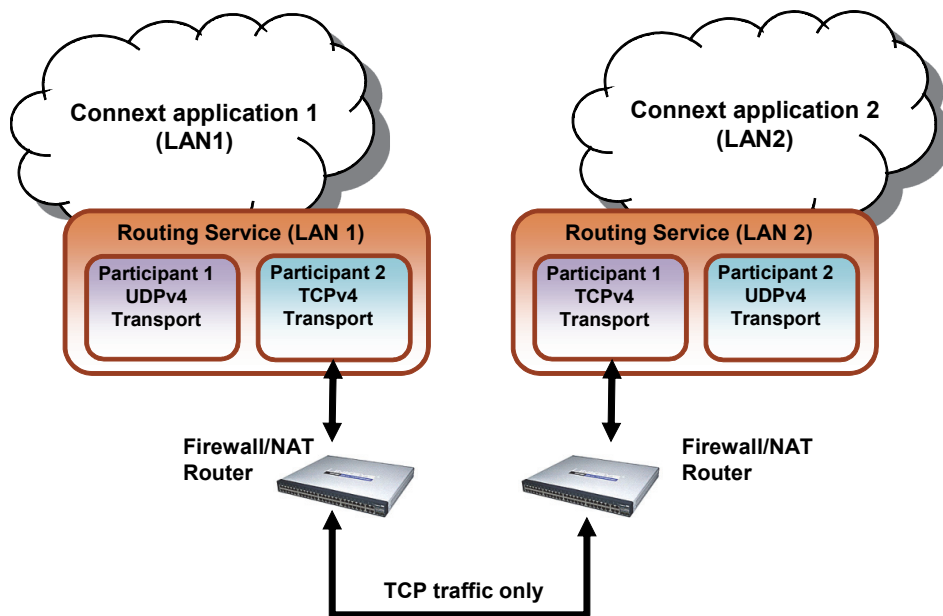
Out of the box, *Routing Service* only uses UDPv4 and Shared Memory transports to communicate with other *Routing Services* and *Connex*t applications. This configuration is appropriate for systems running within a single LAN. However, using UDPv4 introduces several problems when trying to communicate with *Connex*t applications running in different LANs:

- ❑ UDPv4 traffic is usually filtered out by the LAN firewalls for security reasons.
- ❑ Forwarded ports are usually TCP ports.
- ❑ Each LAN may run in its own private IP address space and use NAT (Network Address Translation) to communicate with other networks.

To overcome these issues, *Routing Service* is distributed with a TCP transport that is NAT friendly. The transport can be configured via XML using the PropertyQosPolicy of the *Routing Service's* participants. [Figure 7.1](#) shows a typical scenario where two *Routing Services* are used to bridge two *Connex*t applications running in two different LANs.

The next sections explain how to use and configure the TCP transport with *Routing Service*.

Figure 7.1 WAN Communication Using TCP Transport



7.1 TCP Communication Scenarios

The TCP transport distributed with *Routing Service* can be used to address multiple communication scenarios that go from simple communication within a single LAN to complex communication scenarios across LANs where NATs and firewalls may be involved.

7.1.1 Communication Within a Single LAN

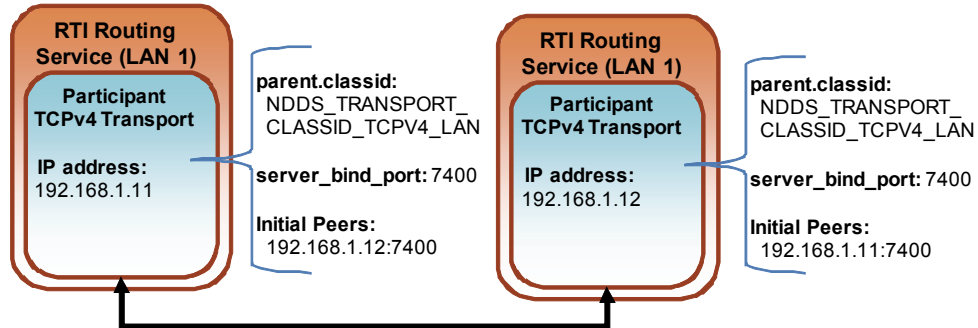
TCP transport can be used as an alternative to UDPv4 to communicate *Connex* applications running inside the same LAN.

Figure 7.2 shows how to configure the TCP transport in this scenario.

`parent.classid`, `transport_mode` and `server_bind_port` are transport properties configured using the `PropertyQosPolicy` of the participant.

Initial Peers represents the peers to which the participant will be announced to. Usually, these peers are configured using the DiscoveryQosPolicy of the participant or the environment variable `NDDS_DISCOVERY_PEERS`. For information on the format of initial peers, see [Section 7.2.1](#).

Figure 7.2 Communication within a Single LAN



7.1.2 Symmetric Communication Across NATs

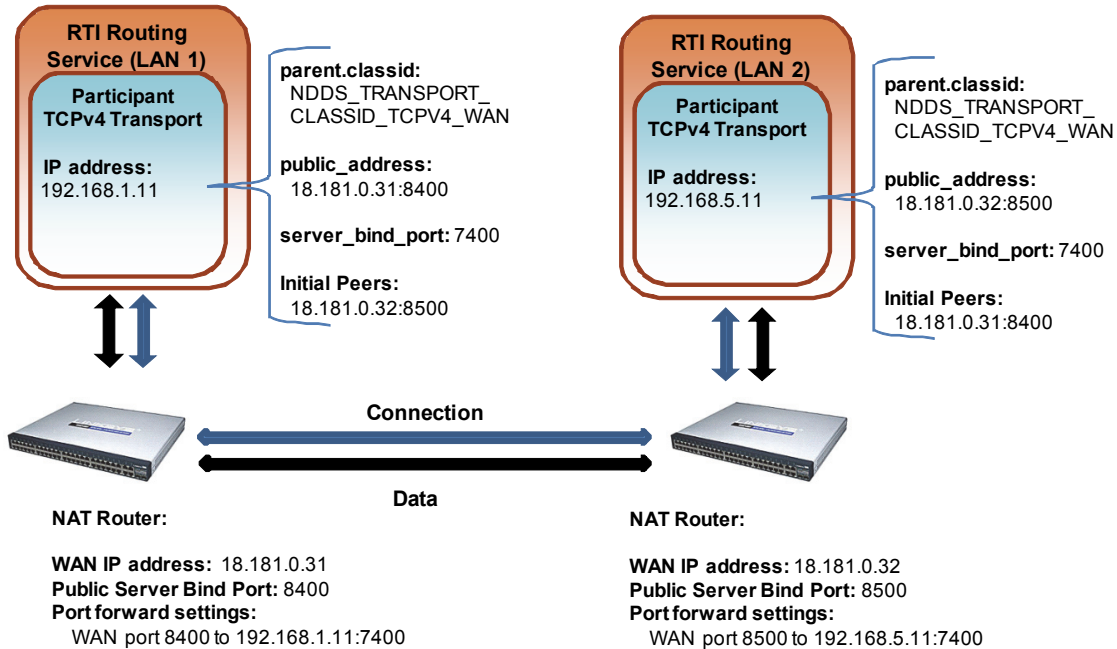
In NAT communication scenarios, each one of the LANs has a private IP address space. The communication with other LANs is done through NAT routers that translate private IP addresses and ports into public IP addresses and ports.

In symmetric communication scenarios, any instance of *Routing Service* can initiate TCP connections with other routing services. [Figure 7.3](#) shows how to configure the TCP transport in this scenario.

Notice that initial peers refer to the public address of the *Routing Service* instances and not the LAN address. In addition, the transport associated with a *Routing Service* instance will have to be configured with its **public_address** so that this information can be propagated as part of the discovery process.

Because the public address and port of the *Routing Service* instances must be known before the communication is established, the NAT Routers will have to be configured statically to translate (forward) the private **server_bind_port** into a public **port**. This process is known as “static NAT” or “port forwarding” and it allows traffic originating in outer networks to reach designated peers in the LAN behind the NAT router.

Figure 7.3 Symmetric Communication across NATs



7.1.3 Asymmetric Communication Across NATs

This scenario is similar to the previous one, except in this case the TCP connections can be initiated only by the *Routing Service* instance in LAN1. For security reasons, incoming connections to LAN1 are not allowed. Figure 7.4 shows how to configure the TCP transport in this scenario.

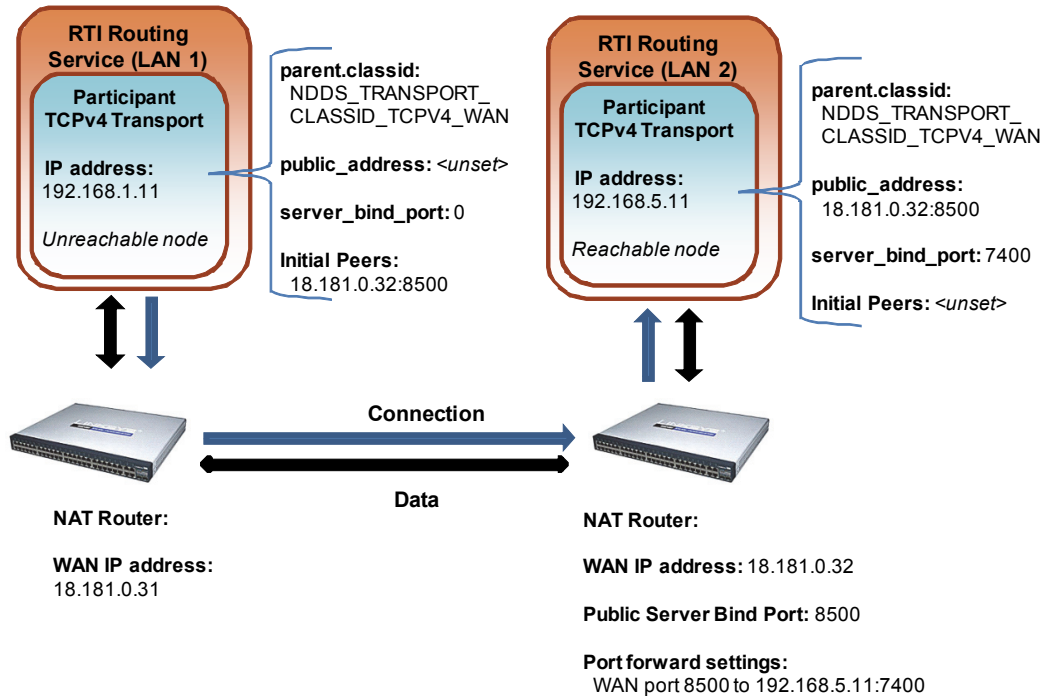
7.1.4 Secure Communication

Security can be added on top of any of the above scenarios. You can have secure communication within the same LAN or across NATs.

To enable secure communication, modify the previous configurations as follows:

- ❑ Change the transport class ID property (**parent.classid**) to be one of the following values:
 - NDDS_TRANSPORT_CLASSID_TLsv4_LAN
 - NDDS_TRANSPORT_CLASSID_TLsv4_WAN

Figure 7.4 Asymmetric Communication Across NATs



Notice that the Routing Service on LAN 1 now does not have a `public_address` set (and its `server_bind_port` is set to zero), meaning that it cannot be reached from the outside network.

- Set at least a certificate of authority (through either the `tls.verify.ca_file` or `tls.verify.ca_path` properties), and the certificate identity (through either the `tls.identity.certificate_chain`, or `tls.identity.certificate_chain_file` properties)
- Make sure to use `'tls_v4_lan'` or `'tls_v4_wan'` in the initial peers list as the prefix for all destination addresses.

To see the differences between a WAN scenario and the same scenario with TLS enabled, you can compare the two example configuration files:

- `shapes/tcp_transport.xml`
- `shapes/tcp_transport_tls.xml`

7.2 Configuring the TCP Transport

The TCP transport is distributed as a shared library in *<Routing Service installation directory>/bin/<architecture>*. The library is called **nddstransporttcp.dll** on Windows and **libnddstransporttcp.so** on UNIX-based systems.

For an example on how to use and configure the TCP transport with *Routing Service* see [Example 8 - Using the TCP Transport with Routing Service \(Section 4.8\)](#) in the *Getting Started Guide*.

As seen in the example, you can configure the properties of the transport in the XML configuration file using the appropriate name/value pairs in the DomainParticipant's PropertyQoSProperty. This will cause *Routing Service* to dynamically load the TCP transport library at run time and then implicitly create and register the transport plugin with *Connex*.

7.2.1 TCP Transport Initial Peers

With the TCP transport, the addresses of the initial peers (NDDS_DISCOVERY_PEERS) that will be contacted during the discovery process have the following format:

```
For WAN communication: tcpv4_wan://<IP address or hostname>:<port>
For LAN communication: tcpv4_lan://<IP address or hostname>:<port>
For WAN+TLS communication: tlsv4_wan://<IP address or hostname>:port
For LAN+TLS communication: tlsv4_lan://<IP address or hostname>:port
```

For example:

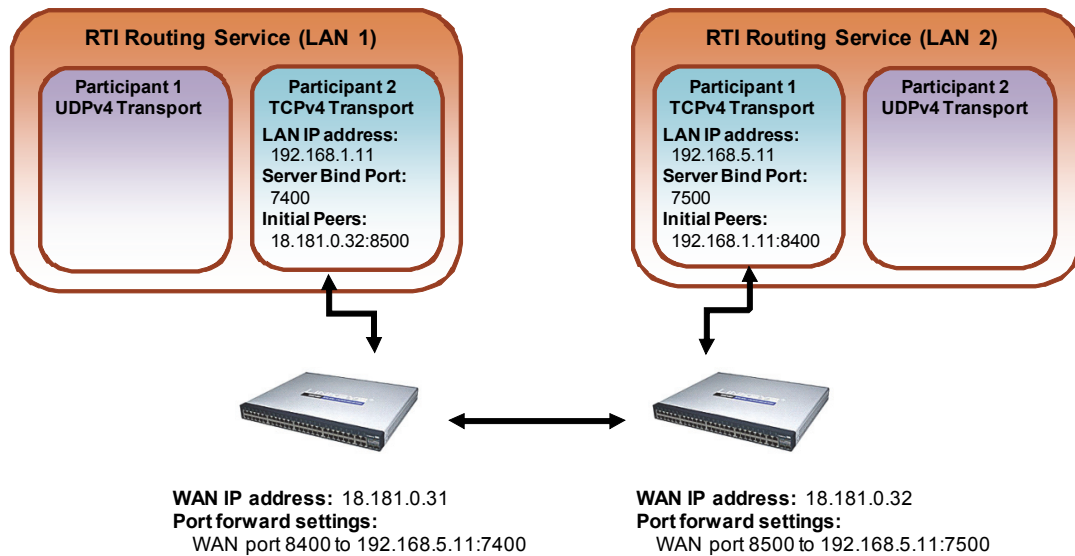
```
setenv NDDS_DISCOVERY_PEERS tcpv4_wan://10.10.1.165:7400,tcpv4_wan:/
/10.10.1.111:7400,tcpv4_lan://192.168.1.1:7500
```

When the TCP transport is configured for LAN communication (with the **parent.classid** property), the IP address is the LAN address of the peer and the port is the server port used by the transport (the **server_bind_port** property).

When the TCP transport is configured for WAN communication (with the **parent.classid** property), the IP address is the WAN or public address of the peer and the port is the public port that is used to forward traffic to the server port in the TCP transport.

When TLS is enabled, the transport settings are similar to WAN and LAN over TCP.

Figure 7.5 Initial Peers in WAN Communication



7.2.2 Setting Up the TCP Transport Properties with the PropertyQoSPolicy

The PropertyQoSPolicy allows you to set up name/value pairs of data and attach them to an entity, such as a DomainParticipant. The configuration of the TCP transport with *Routing Service* is done using the PropertyQoSPolicy of the Domain Participants that are going to use the transport.

For a list of the properties that you can set for the TCP transport, see the *RTI Core Libraries and Utilities User's Manual*¹.

In the following example, participant_1 will communicate with other participants on the same LAN using UDP and Shared Memory transports; participant_2 will communicate with other participants in different LANs using the TCP transport.

```
<dds>
  <routing_service name="MyRoutingService">
    <domain_route name="MyDomainRoute">
      <participant_1>
        <domain_id>56</domain_id>
      </participant_1>
      <participant_2>
        <domain_id>57</domain_id>
      </participant_2>
    </domain_route>
  </routing_service>
</dds>
```

1. See <Connex installation directory>/n.dds.<version>/doc/pdf/RTI_Connext_UsersManual.pdf.

```
<participant_qos>
  <transport_builtin>
    <mask>MASK_NONE</mask>
  </transport_builtin>
  <property>
    <value>
      <element>
        <name>dds.transport.load_plugins</name>
        <value>dds.transport.TCPv4.tcpl</value>
      </element>
      <element>
        <name>
          dds.transport.TCPv4.tcpl.library
        </name>
        <value>libnddstransporttcp.so</value>
      </element>
      <element>
        <name>
          dds.transport.TCPv4.tcpl.create_function
        </name>
        <value>NDDS_Transport_TCPv4_create</value>
      </element>
      <element>
        <name>
          dds.transport.TCPv4.tcpl.parent.classid
        </name>
        <value>
          NDDS_TRANSPORT_CLASSID_TCPV4_WAN
        </value>
      </element>
      <element>
        <name>
          dds.transport.TCPv4.tcpl.public_address
        </name>
        <value>18.181.0.31:8400</value>
      </element>
      <element>
        <name>
          dds.transport.TCPv4.tcpl.server_bind_port
        </name>
        <value>7400</value>
      </element>
    </value>
  </property>
</participant_qos>
</participant_2>
```



```

        </domain_route>
    </routing_service>
</dds>

```

7.2.3 TCP/TLS Transport Properties

Table 7.1 describes the TCP and TLS transport properties.

Table 7.1 TCP/TLS Transport Properties (over LAN or WAN) — `NDDS_Transport_TCPv4_Property_t`

Property Name (prefix with 'dds.transport.TCPv4. tcp1.')	Description
<code>dds.transport.load_plugins</code> (Note: this does not take a prefix)	Required Comma-separated strings indicating the prefix names of all plugins that will be loaded by <i>Connex</i> . For example: " <code>dds.transport.TCPv4.tcp1</code> ". You will use this string as the prefix to the property names. See Footnote 1 on page 7-19 . Note: you can load up to 8 plugins.
<code>library</code>	Required Must be " <code>nddstransporttcp</code> ". This library needs to be in the path during run time (in the <code>LD_LIBRARY_PATH</code> environment variable on UNIX systems, in <code>PATH</code> for Windows systems).
<code>create_function</code>	Required Must be " <code>NDDS_Transport_TCPv4_create</code> ".
<code>aliases</code>	Used to register the transport plugin returned by <code>NDDS_Transport_TCPv4_create()</code> (as specified by <code><TCP_prefix>.create_function</code>) to the <i>DomainParticipant</i> . Aliases should be specified as a comma-separated string, with each comma delimiting an alias. Default: the transport prefix (see Footnote 1)
<code>parent.classid</code>	Must be set to one of the following values: <code>NDDS_TRANSPORT_CLASSID_TCPV4_LAN</code> for TCP communication within a LAN <code>NDDS_TRANSPORT_CLASSID_TLsv4_LAN</code> for TLS communication within a LAN <code>NDDS_TRANSPORT_CLASSID_TCPV4_WAN</code> for TCP communication across LANs and firewalls <code>NDDS_TRANSPORT_CLASSID_TLsv4_WAN</code> for TLS communication across LAN and firewalls Default: <code>NDDS_TRANSPORT_CLASSID_TCPV4_LAN</code>

Table 7.1 TCP/TLS Transport Properties (over LAN or WAN) — NDDS_Transport_TCPv4_Property_t

Property Name (prefix with 'dds.transport.TCPv4. tcp1.')	Description
parent.gather_send_buffer_count_max	<p>Specifies the maximum number of buffers that <i>Connex</i>t can pass to the send() function of the transport plugin.</p> <p>The transport plugin send() operation supports a gather-send concept, where the send() call can take several discontinuous buffers, assemble and send them in a single message. This enables <i>Connex</i>t to send a message from parts obtained from different sources without first having to copy the parts into a single contiguous buffer.</p> <p>However, most transports that support a gather-send concept have an upper limit on the number of buffers that can be gathered and sent. Setting this value will prevent <i>Connex</i>t from trying to gather too many buffers into a send call for the transport plugin.</p> <p><i>Connex</i>t requires all transport-plugin implementations to support a gather-send of least a minimum number of buffers. This minimum number is defined as <code>NDDS_TRANSPORT_PROPERTY_GATHER_SEND_BUFFER_COUNT_MIN</code>.</p> <p>Default: 128</p>
parent.message_size_max	<p>The maximum size of a message, in bytes, that can be sent or received by the transport plugin.</p> <p>If you set this higher than the default, the <i>DomainParticipant's</i> buffer_size (in the <i>RECEIVER_POOL</i> QoSPolicy, see the <i>RTI Core Libraries and Utilities User's Manual</i>²) should also be changed.</p> <p>Default: 9216</p>
parent.allow_interfaces_list	<p>A list of strings, each identifying a range of interface addresses that can be used by the transport.</p> <p>Interfaces must be specified as comma-separated strings, with each comma delimiting an interface.</p> <p>For example: 10.10.*, 10.15.*</p> <p>If the list is non-empty, this "white" list is applied before parent.deny_interfaces_list.</p> <p>Default: All available interfaces are used.</p>

Table 7.1 TCP/TLS Transport Properties (over LAN or WAN) — NDDS_Transport_TCPv4_Property_t

Property Name (prefix with 'dds.transport.TCPv4. tcp1.')	Description
parent. deny_interfaces_list	<p>A list of strings, each identifying a range of interface addresses that will not be used by the transport.</p> <p>If the list is non-empty, deny the use of these interfaces.</p> <p>Interfaces must be specified as comma-separated strings, with each comma delimiting an interface.</p> <p>For example: 10.10.*</p> <p>This "black" list is applied after parent.allow_interfaces_list and filters out the interfaces that should not be used.</p> <p>Default: No interfaces are denied</p>
send_socket_ buffer_size	<p>Size, in bytes, of the send buffer of a socket used for sending. On most operating systems, setsockopt() will be called to set the SENDBUF to the value of this parameter.</p> <p>This value must be greater than or equal to parent.message_size_max or -1.</p> <p>The maximum value is operating system-dependent.</p> <p>Default: -1 (setsockopt() (or equivalent) will not be called to size the send buffer of the socket)</p>
recv_socket_ buffer_size	<p>Size, in bytes, of the receive buffer of a socket used for receiving.</p> <p>On most operating systems, setsockopt() will be called to set the RECVBUF to the value of this parameter.</p> <p>This value must be greater than or equal to parent.message_size_max or -1. The maximum value is operating-system dependent.</p> <p>Default: -1 (setsockopt() (or equivalent) will not be called to size the receive buffer of the socket)</p>

Table 7.1 TCP/TLS Transport Properties (over LAN or WAN) — NDDS_Transport_TCPv4_Property_t

Property Name (prefix with 'dds.transport.TCPv4. tcp1.')	Description
ignore_loopback_ interface	<p>Prevents the transport plugin from using the IP loopback interface.</p> <p>This property is ignored when <code>parent.classid</code> is <code>NDDS_TRANSPORT_CLASSID_TCPV4_WAN</code> or <code>NDDS_TRANSPORT_CLASSID_TLsv4_WAN</code>.</p> <p>Two values are allowed:</p> <p>0: Enable local traffic via this plugin. The plugin will use and report the IP loopback interface only if there are no other network interfaces (NICs) up on the system.</p> <p>1: Disable local traffic via this plugin. This means “do not use the IP loopback interface, even if no NICs are discovered.” This setting is useful when you want applications running on the same node to use a more efficient plugin like shared memory instead of the IP loopback.</p> <p>Default: 1</p>
ignore_nonrunning_ interfaces	<p>Prevents the transport plugin from using a network interface that is not reported as RUNNING by the operating system.</p> <p>The transport checks the flags reported by the operating system for each network interface upon initialization. An interface which is not reported as UP will not be used. This property allows the same check to be extended to the IFF_RUNNING flag implemented by some operating systems. The RUNNING flag means that "all resources are allocated" and may be off if no link is detected (e.g., the network cable is unplugged).</p> <p>Two values are allowed:</p> <p>0: Do not check the RUNNING flag when enumerating interfaces, just make sure the interface is UP.</p> <p>1: Check the flag when enumerating interfaces, and ignore those that are not reported as RUNNING. This can be used on some operating systems to cause the transport to ignore interfaces that are enabled but not connected to the network.</p> <p>Default: 1</p>

Table 7.1 TCP/TLS Transport Properties (over LAN or WAN) — NDDS_Transport_TCPv4_Property_t

Property Name (prefix with 'dds.transport.TCPv4. tcp1.')	Description
transport_priority_mask	<p>Mask for the transport priority field. This is used in conjunction with transport_priority_mapping_low/transport_priority_mapping_high to define the mapping from <i>Connex</i> transport priority to the IPv4 TOS field. Defines a contiguous region of bits in the 32-bit transport priority value that is used to generate values for the IPv4 TOS field on an outgoing socket.</p> <p>For example, the value 0x0000ff00 causes bits 9-16 (8 bits) to be used in the mapping. The value will be scaled from the mask range (0x0000 -0xff00 in this case) to the range specified by low and high.</p> <p>If the mask is set to zero, then the transport will not set IPv4 TOS for send sockets.</p> <p>Default: 0</p>
transport_priority_mapping_low	<p>Sets the low and high values of the output range to IPv4 TOS.</p> <p>These values are used in conjunction with transport_priority_mask to define the mapping from <i>Connex</i> transport priority to the IPv4 TOS field. Defines the low and high values of the output range for scaling.</p>
transport_priority_mapping_high	<p>Note that IPv4 TOS is generally an 8-bit value.</p> <p>Default <code>transport_priority_mapping_low</code>: 0</p> <p>Default <code>transport_priority_mapping_high</code>: 0xFF</p>
server_socket_backlog	<p>Determines the maximum length of the queue of pending connections.</p> <p>Default: 5</p>
public_address	<p>Required for WAN communication</p> <p>Public IP address and port (WAN address and port) associated with the transport instantiation. The address and port must be separated with ':'. For example: 10.10.9.10:4567</p> <p>This field is only used when <code>parent.classid</code> is <code>NDDS_TRANSPORT_CLASSID_TCPV4_WAN</code> or <code>NDDS_TRANSPORT_CLASSID_TLsv4_WAN</code>.</p> <p>The public address and port are necessary to support communication over a WAN that involves Network Address Translators (NATs). Typically, the address is the public address of the IP router that provides access to the WAN. The port is the IP router port that is used to reach the private <code>server_bind_port</code> inside the LAN from the outside. This value is expressed as a string in the form: <code>ip[:port]</code>, where <code>ip</code> represents the IPv4 address and <code>port</code> is the external port number of the router.</p> <p>Note that host names are not allowed in the <code>public_address</code> because they may resolve to an internet address that is not what you want (i.e., 'localhost' may map to your local IP or to 127.0.0.1).</p>

Table 7.1 TCP/TLS Transport Properties (over LAN or WAN) — `NDSS_Transport_TCPv4_Property_t`

Property Name (prefix with 'dds.transport.TCPv4. tcp1.')	Description
<code>server_bind_port</code>	<p>Private IP port (inside the LAN) used by the transport to accept TCP connections.</p> <p>If this property is set to zero, the transport will disable the internal server socket, making it impossible for external peers to connect to this node. In this case, the node is considered unreachable and will communicate only using the asynchronous mode with other (reachable) peers.</p> <p>For WAN communication, this port must be forwarded to a public port in the NAT-enabled router that connects to the outer network.</p> <p>Default: 7400</p>
<code>read_buffer_allocation</code>	<p>Allocation settings applied to read buffers.</p> <p>These settings configure the initial number of buffers, the maximum number of buffers and the buffers to be allocated when more buffers are needed.</p> <p>Default:</p> <ul style="list-style-type: none"> <input type="checkbox"/> <code>read_buffer_allocation.initial_count = 2</code> <input type="checkbox"/> <code>read_buffer_allocation.max_count = -1</code> (unlimited) <input type="checkbox"/> <code>read_buffer_allocation.incremental_count = -1</code> (number of buffers will keep doubling on each allocation until it reaches <code>max_count</code>)
<code>write_buffer_allocation</code>	<p>Allocation settings applied to buffers used for an asynchronous (non-blocking) write.</p> <p>These settings configure the initial number of buffers, the maximum number of buffers, and the buffers to be allocated when more buffers are needed.</p> <p>Default:</p> <ul style="list-style-type: none"> <input type="checkbox"/> <code>write_buffer_allocation.initial_count = 4</code> <input type="checkbox"/> <code>write_buffer_allocation.max_count = 1000</code> <input type="checkbox"/> <code>write_buffer_allocation.incremental_count = 10</code> <p>Note that for the write buffer pool, the <code>max_count</code> is not set to unlimited. This is to avoid having a fast writer quickly exhaust all the available system memory, in case of a temporary network slowdown. When this write buffer pool reaches the maximum, the low-level send command of the transport will fail; at that point <i>Connex</i>t will take the appropriate action (retry to send or drop it), according to the application's QoS (if the transport is used for reliable communication, the data will still be sent eventually).</p>

Table 7.1 TCP/TLS Transport Properties (over LAN or WAN) — NDDS_Transport_TCPv4_Property_t

Property Name (prefix with 'dds.transport.TCPv4. tcp1.')	Description
control_buffer_allocation	<p>Allocation settings applied to buffers used to serialize and send control messages. These settings configure the initial number of buffers, the maximum number of buffers, and the buffers to be allocated when more buffers are needed.</p> <p>Default:</p> <ul style="list-style-type: none"> <input type="checkbox"/> control_buffer_allocation.initial_count = 2 <input type="checkbox"/> control_buffer_allocation.max_count = -1 (unlimited) <input type="checkbox"/> control_buffer_allocation.incremental_count = -1 (number of buffers will keep doubling on each allocation until it reaches max_count)
control_message_allocation	<p>Allocation settings applied to control messages. These settings configure the initial number of messages, the maximum number of messages, and the messages to be allocated when more messages are needed.</p> <p>Default:</p> <ul style="list-style-type: none"> <input type="checkbox"/> control_message_allocation.initial_count = 2 <input type="checkbox"/> control_message_allocation.max_count = -1 (unlimited) <input type="checkbox"/> control_message_allocation.incremental_count = -1 (number of messages will keep doubling on each allocation until it reaches max_count)
control_attribute_allocation	<p>Allocation settings applied to control messages attributes. These settings configure the initial number of attributes, the maximum number of attributes, and the attributes to be allocated when more attributes are needed.</p> <p>Default:</p> <ul style="list-style-type: none"> <input type="checkbox"/> control_attribute_allocation.initial_count = 2 <input type="checkbox"/> control_attribute_allocation.max_count = -1 (unlimited) <input type="checkbox"/> control_attribute_allocation.incremental_count = -1 (number of attributes will keep doubling on each allocation until it reaches max_count)

Table 7.1 TCP/TLS Transport Properties (over LAN or WAN) — NDDS_Transport_TCPv4_Property_t

Property Name (prefix with 'dds.transport.TCPv4. tcp1.')	Description
force_asynchronous_send	<p>Forces an asynchronous send. When this parameter is set to 0, the TCP transport will attempt to send data as soon as the internal send() function is called. When it is set to 1, the transport will make a copy of the data to send and enqueue it in an internal send buffer. Data will be sent as soon as the low-level socket buffer has space.</p> <p>Normally setting it to 1 delivers better throughput in a fast network, but will result in a longer time to recover from various TCP error conditions. Setting it to 0 may cause the low-level send() function to block until the data is physically delivered to the lower socket buffer. For an application writing data at a very fast rate, it may cause the caller thread to block if the send socket buffer is full. This could produce lower throughput in those conditions (the caller thread could prepare the next packet while waiting for the send socket buffer to become available).</p> <p>Default: 0</p>
max_packet_size	<p>The maximum size of a TCP segment.</p> <p>This parameter is only supported on Linux architectures.</p> <p>By default, the maximum size of a TCP segment is based on the network MTU for destinations on a local network, or on a default 576 for destinations on non-local networks. This behavior can be changed by setting this parameter to a value between 1 and 65535.</p> <p>Default: -1 (default behavior)</p>
enable_keep_alive	<p>Configures the sending of KEEP_ALIVE messages in TCP.</p> <p>Setting this value to 1 causes a KEEP_ALIVE packet to be sent to the remote peer if a long time passes with no other data sent or received.</p> <p>This feature is implemented only on architectures that provide a low-level implementation of the TCP keep-alive feature.</p> <p>On Windows systems, the TCP keep-alive feature can be globally enabled through the system's registry: <code>\HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Tcpip\Parameters</code>. Refer to MSDN documentation for more details.</p> <p>On Solaris systems, most of the TCP keep-alive parameters can be changed through the kernel properties.</p> <p>Default: 0</p>
keep_alive_time	<p>Specifies the interval of inactivity, in seconds, that causes TCP to generate a KEEP_ALIVE message.</p> <p>This parameter is only supported on Linux architectures.</p> <p>Default: -1 (OS default value)</p>

Table 7.1 TCP/TLS Transport Properties (over LAN or WAN) — NDDS_Transport_TCPv4_Property_t

Property Name (prefix with 'dds.transport.TCPv4. tcp1.')	Description
keep_alive_interval	Specifies the interval, in seconds, between KEEP_ALIVE retries. This parameter is only supported on Linux architectures. Default: -1 (OS default value)
keep_alive_retry_count	The maximum number of KEEP_ALIVE retries before dropping the connection. This parameter is only supported on Linux architectures. Default: -1 (OS default value)
disable_nagle	Disables the TCP nagle algorithm. When this property is set to 1, TCP segments are always sent as soon as possible, which may result in poor network utilization. Default: 0
logging_verbosity_bitmap	Bitmap that specifies the verbosity of log messages from the transport. Logging values: <input type="checkbox"/> -1 (0xffffffff): do not change the current verbosity <input type="checkbox"/> 0x00: silence <input type="checkbox"/> 0x01: errors <input type="checkbox"/> 0x02: warnings <input type="checkbox"/> 0x04: local <input type="checkbox"/> 0x08: remote <input type="checkbox"/> 0x10: period <input type="checkbox"/> 0x80: other (used for control protocol tracing) Default: -1 Note: the logging verbosity is a global property shared across multiple instances of the TCP transport. If you create a new TCP Transport instance with logging_verbosity_bitmap different than -1, the change will affect all the other instances as well. The default TCP transport verbosity is errors and warnings. Note: The option of 0x80 (other) is used only for tracing the internal control protocol. Since the output is very verbose, this feature is enabled only in the debug version of the <i>TCP Transport</i> library (<i>libnddstransporttcpd.so / LIBNDDSTRANSPORTD.LIB</i>).

Table 7.1 TCP/TLS Transport Properties (over LAN or WAN) — NDDS_Transport_TCPv4_Property_t

Property Name (prefix with 'dds.transport.TCPv4. tcp1.')	Description
outstanding_connection_cookies	<p>Maximum number of outstanding connection cookies allowed by the transport when acting as server.</p> <p>A connection cookie is a token provided by a server to a client; it is used to establish a data connection. Until the data connection is established, the cookie cannot be reused by the server.</p> <p>To avoid wasting memory, it is good practice to set a cap on the maximum number of connection cookies (pending connections).</p> <p>When the maximum value is reached, a client will not be able to connect to the server until new cookies become available.</p> <p>Range: 1 or higher, or -1 (which means an unlimited number).</p> <p>Default: 100</p>
outstanding_connection_cookies_life_span	<p>Maximum lifespan (in seconds) of the cookies associated with pending connections.</p> <p>If a client does not connect to the server before the lifespan of its cookie expires, it will have to request a new cookie.</p> <p>Range: 1 second or higher, or -1</p> <p>Default : -1, which means an unlimited amount of time (effectively disabling the feature).</p>
tls.verify.ca_file	<p>A string that specifies the name of a file containing Certificate Authority certificates. The file should be in PEM format. See the OpenSSL manual page for SSL_load_verify_locations for more information.</p> <p>To enable TLS, ca_file or ca_path is required; both may be specified (at least one is required).</p>
tls.verify.ca_path	<p>A string that specifies paths to directories containing Certificate Authority certificates. Files should be in PEM format and follow the OpenSSL-required naming conventions. See the OpenSSL manual page for SSL_CTX_load_verify_locations for more information.</p> <p>To enable TLS, ca_file or ca_path is required; both may be specified (at least one is required).</p>
tls.verify.verify_depth	<p>Maximum certificate chain length for verification.</p>
tls.verify.crl_file	<p>Name of the file containing the Certificate Revocation List.</p> <p>File should be in PEM format.</p>
tls.cipher.cipher_list	<p>List of available TLS ciphers. See the OpenSSL manual page for SSL_set_cipher_list for more information on the format of this string.</p>

Table 7.1 TCP/TLS Transport Properties (over LAN or WAN) — NDDS_Transport_TCPv4_Property_t

Property Name (prefix with 'dds.transport.TCPv4. tcp1.')	Description
tls.cipher. dh_param_files	List of available Diffie-Hellman (DH) key files. For example: "foo.pem:512,bar.pem:256" means: dh_param_files[0].file = foo.pem, dh_param_files[0].bits = 512, dh_param_files[1].file = bar.pem, dh_param_files[1].bits = 256
tls.cipher.engine_id	String ID of OpenSSL cipher engine to request.
tls.identity. certificate_chain	A string containing an identifying certificate chain (in PEM format). An identifying certificate is required for secure communication. The string must be sorted starting with the certificate to the highest level (root CA). Either certificate_chain or certificate_chain_file is required. You must set exactly one of these. Do not set both of them (this would produce a configuration error).
tls.identity. certificate_chain_file	A string that specifies the name of a file containing an identifying certificate chain (in PEM format). An identifying certificate is required for secure communication. The file must be sorted starting with the certificate to the highest level (root CA). Optionally, a private key may be appended to this file. If a private key is not appended to this file, then either private_key or private_key_file is required. Either certificate_chain or certificate_chain_file is required. You must set exactly ONE of these. Do not set both of them (this would produce a configuration error).
tls.identity. private_key_password	A string that specifies the password for private key.
tls.identity. private_key	A string containing a private key (in PEM format). Either private_key or private_key_file may be specified. Do not set both of them (this would produce a configuration error). If both are unspecified (NULL), the private key must be appended to the certificate chain file.
tls.identity. private_key_file	A string that specifies the name of a file containing a private key (in PEM format). Either private_key or private_key_file may be specified. Do not set both of them (this would produce a configuration error). If both are unspecified (NULL), the private key must be appended to the certificate chain file.

1. Assuming you used 'dds.transport.TCPv4.tcp1' as the alias to load the plugin. If not, change the prefix to match the string used with dds.transport.load_plugins.

2. See <Connext installation directory>/nlds.<version>/doc/pdf/RTI_Connext_UsersManual.pdf.


Chapter 8 Extending Routing Service with Adapters

As described in [Section 2.4.8](#), *adapters* are pluggable components that allow *Routing Service* to consume and produce data for different data domains (e.g., *Connex*, JMS, Socket, etc.).

By default, *Routing Service* is distributed with a built-in *Connex* adapter. Any other adapter plugins must be provided as shared libraries or Java classes and registered within the `<adapter_library>` tag.

The figure to the right describes the *Routing Service* adapter architecture.

Input adapters are used to collect data samples from different data domains, such as *Connex* or JMS. The input samples are processed by the *Routing Service* engine and are passed along to custom transformations if they are defined. Finally, the transformed data is provided to the output adapters.



The adapter plugin API is used to create new adapters; it is supported in C and Java.

The rest of this chapter describes:

- ❑ [Adapter Usage and Configuration \(Section 8.1\)](#)
- ❑ [Adapter API And Entity Model \(Section 8.2\)](#)
- ❑ [Creating New Adapters \(Section 8.3\)](#)

8.1 Adapter Usage and Configuration

Adapter plugins must be registered within an adapter library in the XML configuration file.

For example:

```
<?xml version="1.0"?>
<dds>
  <adapter_library name="MyAdapterLibrary">
    <adapter_plugin name="MyCAdapterPlugin">
      <dll>mycadapter</dll>
      <create_function>
        MyCAdapterPlugin_create
      </create_function>
    </adapter_plugin>
    <java_adapter_plugin name="MyJavaAdapter">
      <class_name>com.rti.adapters.MyJavaAdapter</class_name>
    </java_adapter_plugin>
  </adapter_library>
  ...
</dds>
```

C adapters are registered using the tag `<adapter_plugin>`; Java adapters use the tag `<java_adapter_plugin>`.

[Table 8.1](#) lists the tags allowed within `<adapter_plugin>`.

[Table 8.2](#) lists the tags allowed within `<java_adapter_plugin>`.

Once the adapter plugins are registered, they can be used to create connections within a domain route (see [Section 2.4.2](#)).

For example:

```
<dds>
  <routing_service name="Router1"
    group_name="Group1">
    <domain_route name="DomainRoute1">
      <connection_1 plugin_name="MyAdapterLibrary::MyCAdapterPlugin">
        . . .
      </connection_1>

      <connection_2 plugin_name="MyAdapterLibrary::MyJavaAdapter">
        . . .
      </connection_2>

      <session name="Session">
        . . .
      </session>
    </domain_route>
    . . .
  </routing_service>
</dds>
```

Table 8.1 C Adapter Plugin Tags

Tags within <adapter_plugin>	Description	Number of Tags Allowed
<dll>	<p>Required</p> <p>Shared library containing the implementation of the adapter plugin.</p> <p>The <dll> tag may specify the exact name of the file (for example, lib/libmyadapter.so) or a general name (no file extension) which will be completed as follows:</p> <ul style="list-style-type: none"> • <dll> value: dir/myadapter • Final Path (UNIX-based systems): dir/libmyadapter.so • Final Path (Windows systems): dir/myadapter.dll <p>If the library specified with the <dll> tag cannot be opened (because the library path is not in the Path environment variable on a Windows system or the LD_LIBRARY_PATH environment variables on a UNIX-based system), <i>Routing Service</i> will look for the library in <i><Routing Service installation directory>/bin/<architecture></i>.</p>	1
<create_function>	<p>Required</p> <p>This tag must contain the name of the function used to create the adapter plugin.</p> <p>The function must be implemented in the adapter shared library.</p>	1
<property>	<p>Sequence of name/value(string) pairs that can be used to configure the parameters of the adapter. For example:</p> <pre data-bbox="451 1035 959 1253"><property> <value> <element> <name>username</name> <value>myusername</value> </element> </value> </property></pre>	0 or 1

Table 8.2 Java Adapter Plugin Tags

Tags within <java_adapter_plugin>	Description	Number of Tags Allowed
<class_name>	<p>Required</p> <p>Name of the class that implements the adapter plugin. For example: com.rti.adapters.JMSAdapter</p> <p>The classpath required to run the Java adapter must be part of the <i>Routing Service</i> JVM configuration. See Routing Service Tags (Table 2.2) for additional information on JVM creation and configuration with the routing service.</p>	1
<property>	<p>Sequence of name/value(string) pairs that can be used to configure the parameters of the adapter. For example:</p> <pre data-bbox="548 674 1057 894"><property> <value> <element> <name>username</name> <value>myusername</value> </element> </value> </property></pre>	0 or 1

8.2 Adapter API And Entity Model

There are five main classes in the adapter class model:

1. **Adapter:** An *Adapter* is a factory for *Connections*.
See Table 8.3, “Adapter Operations,” on page 8-8.
2. **Connection:** A *Connection* provides access to a data domain (such as a *Connex*t domain or JMS provider network) and is a factory for *Sessions*, *StreamReaders* and *StreamWriters*.

In the built-in *Connex*t adapter, a *Connection* is mapped to a *DomainParticipant*.

In an XML configuration file, connections are associated with the tags `<connection_1>` and `<connection_2>` within a domain route (see Section 2.4.2).

See Table 8.4, “Connection Operations,” on page 8-8.

3. **Session:** A *Session* is a concurrency unit within a connection that has an associated set of *StreamReaders* and *StreamWriters*. Access to the *StreamReaders* and *StreamWriters* in the same *Session* is serialized by *Routing Service* (two *StreamReaders/StreamWriters* cannot be accessed concurrently).

In the built-in *Connex*t adapter, a *Session* is mapped to a *Publisher/Subscriber* pair.

In an XML file, *Sessions* are associated with the tag `<session>` (see Section 2.4.5).

See Table 8.5, “Session Operations,” on page 8-8.

4. **StreamReader:** A *StreamReader* provides a way to read samples of a specific type from a data domain.

In the built-in *Connex*t adapter, a *StreamReader* is mapped to a *DataReader*.

In an XML file, *StreamReaders* are associated with the tag `<input>` within `<route>` or `<auto_route>` (see Section 2.4.6).

See Table 8.6, “StreamReader Operations,” on page 8-11.

5. **StreamWriter:** A *StreamWriter* provides a way to write samples of a specific type in a data domain.

In the built-in *Connex*t adapter, a *StreamWriter* is mapped to a *DataWriter*.

In an XML file, *StreamWriters* are associated with the tag `<output>` within `<route>` or `<auto_route>` (see Section 2.4.6).

See Table 8.7, “StreamWriter Operations,” on page 8-11.

Figure 8.1 describes the adapter class model.

Figure 8.1 **Adapter Class Model**

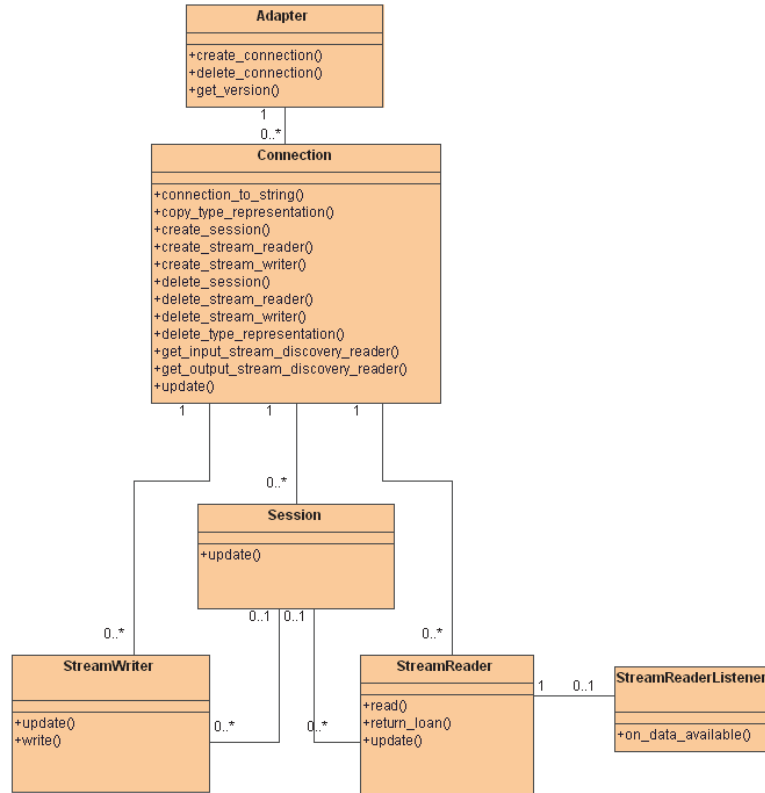


Table 8.5 Session Operations

Operation	Description
update	Updates the configuration of a session. This method is called when the update command is received by the routing service session (<session> tag) containing the adapter session (see Section 5.2.11). Implementation of this API is optional.

Table 8.3 Adapter Operations

Operation	Description
create_connection	Creates a new connection. Connection objects are created when the domain routes that contain them are enabled. Implementation of this API is required.
delete_connection	Deletes a previously created connection. Connection objects are deleted when the domain routes that contain them are disabled. Implementation of this API is required.
getVersion	Returns the Adapter's version. This method is only available in Java. In C, the version of the adapter is set on a member called plugin_version in the plugin structure RTI_RoutingServiceAdapterPlugin (see Section 8.3.2). The version of the adapter is only used for logging purposes. Implementation of this API is required.

Table 8.4 Connection Operations

Operation	Description
connection_to_string	Returns the string representation of a connection for logging purposes. Implementation of this API is optional. If the API is not implemented, <i>Routing Service</i> will use the fully qualified name of the adapter plugin.
create_session	Creates a new session. Connection session objects are created when the associated routing service sessions are enabled. Implementation of this API is optional.

Table 8.4 Connection Operations

Operation	Description
<code>delete_session</code>	<p>Deletes a previously created session.</p> <p>Connection session objects are deleted when the routing service sessions that contain them are disabled.</p> <p>Implementation of this API is optional.</p>
<code>create_stream_reader</code>	<p>Creates a new <code>StreamReader</code> within a routing service route.</p> <p>This method is called when the route is enabled and the 'creation mode' condition associated with the <input> tag becomes true (see Section 2.4.6.4).</p> <p>One of the parameters received by the <code>create_stream_reader()</code> operation is the <code>StreamReaderListener</code>. The <code>StreamReaderListener</code> interface provides a callback which will be used by the adapter to notify <i>Routing Service</i> of the existence of new data.</p> <p>Implementation of this API is required only when there are routes using the adapter to receive data.</p>
<code>delete_stream_reader</code>	<p>Deletes a previously created <code>StreamReader</code>.</p> <p>This method is called when the route is disabled or when the 'creation mode' condition associated with the <input> tag becomes false (see Section 2.4.6.4).</p> <p>Implementation of this API is required only when there are routes using the adapter to receive data.</p>
<code>create_stream_writer</code>	<p>Creates a new <code>StreamWriter</code> within a routing service route.</p> <p>This method is called when the route is enabled and the 'creation mode' condition associated with the <output> tag becomes true (see Section 2.4.6.4).</p> <p>Implementation of this API is required only when there are routes using the adapter to produce data.</p>
<code>delete_stream_writer</code>	<p>Deletes a previously created <code>StreamWriter</code>.</p> <p>This method is called when the route is disabled or when the 'creation mode' condition associated with the <output> tag becomes false (see Section 2.4.6.4).</p> <p>Implementation of this API is required only when there are routes using the adapter to produce data.</p>

Table 8.4 Connection Operations

Operation	Description
get_output_stream_discovery_reader	<p>Returns a <code>StreamReader</code> that is used by <i>Routing Service</i> to discover output streams. An output stream is a stream to which <code>StreamWriters</code> can write data. Disposed scenarios, where an output stream disappears, are also notified using the discovery <code>StreamReader</code>.</p> <p>For additional information, see Stream Discovery (Section 8.2.2).</p> <p>Implementation of this API is optional. However, if none of the adapters in a domain route implement the discovery API, the routes' types must be declared in the configuration file.</p>
get_input_stream_discovery_reader	<p>Returns a <code>StreamReader</code> that is used by <i>Routing Service</i> to discover input streams. An input stream is a stream from which a <code>StreamReader</code> can read data. Disposed scenarios, where an input stream disappears, are also notified using the discovery <code>StreamReader</code>.</p> <p>For additional information, see Stream Discovery (Section 8.2.2).</p> <p>Implementation of this API is optional. However, if none of the adapters in a domain route implement the discovery API, the routes' types must be declared in the configuration file.</p>
copy_type_representation	<p>Copies a type representation object (<code>RoutingServiceTypeRepresentation</code>).</p> <p>The format of the type representation is given by the representation kind. For example, if the representation kind is <code>RTI_ROUTING_SERVICE_TYPE_REPRESENTATION_DYNAMIC_TYPE</code>, the type representation will be a <i>Connex</i> <code>TypeCode</code>.</p> <p>This method is part of the adapter discovery API and is used by <i>Routing Service</i> to copy the type representation of discovered streams (see Stream Discovery (Section 8.2.2)).</p> <p>Implementation of this API is optional and tied to the implementation of <code>get_input_stream_discovery_reader()</code> and <code>get_output_stream_discovery_reader()</code>.</p>
delete_type_representation	<p>Deletes a previously created type-representation object.</p> <p>This method is part of the adapter discovery API.</p> <p>Implementation of this API is optional and tied to the implementation of <code>get_input_stream_discovery_reader()</code> and <code>get_output_stream_discovery_reader()</code>.</p>
update	<p>Updates the connection's configuration.</p> <p>This method is called when the update command is received by the domain route containing the connection (see Section 5.2.11).</p> <p>Implementation of this API is optional.</p>

Table 8.6 **StreamReader Operations**

Operation	<p style="text-align: center;">Description</p> <p style="text-align: center;"><i>The StreamReader API is required only when the adapter is used to receive data. Otherwise, it is optional.</i></p>
update	<p>Updates the configuration of a StreamReader providing a new set of properties. This method is called after the update command is received by the routing service route containing the StreamReader (see Section 5.2.11).</p> <p>Implementation of this API is optional.</p>
read	<p>Reads a collection of data samples and sample infos from the StreamReader. When <i>Routing Service</i> is done using the samples, it will 'return the loan' to the StreamReader by calling return_loan().</p> <p>Implementation of this API is required if the adapter is used to receive data.</p>
return_loan	<p>Returns the loan on the read samples and infos.</p> <p><i>Routing Service</i> calls this method to indicate that it is done accessing the collection of data samples and sample infos obtained by an earlier invocation to read.</p> <p>Implementation of this API is required if the adapter is used to receive data.</p>

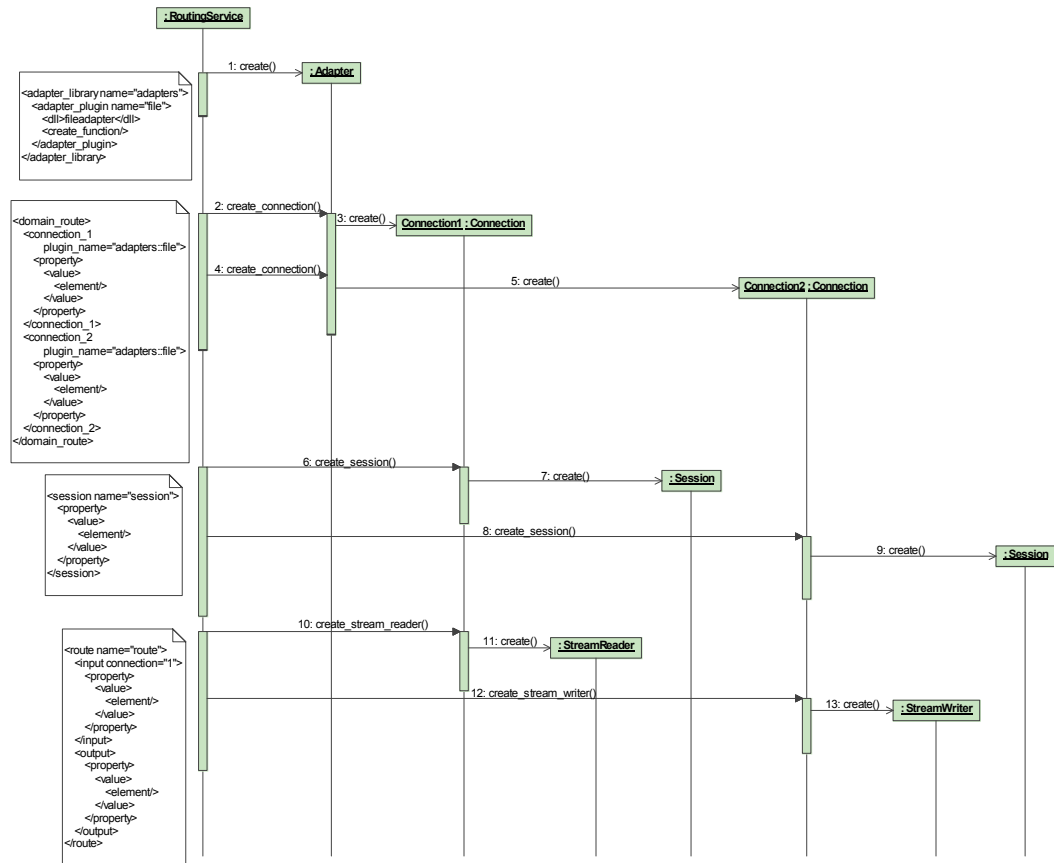
Table 8.7 **StreamWriter Operations**

Operation	<p style="text-align: center;">Description</p> <p style="text-align: center;"><i>The StreamWriter API is only required when the adapter is used to produce data. Otherwise it is optional.</i></p>
update	<p>Updates the configuration of a StreamWriter providing a new set of properties. This method is called after the update command is received by the routing service route containing the StreamWriter (see Section 5.2.11).</p> <p>Implementation of this API is optional.</p>
write	<p>Writes a collection of data samples and sample infos in the data domain associated with the StreamWriter.</p> <p>Implementation of this API is required if the adapter is used to produce data.</p>

8.2.1 Entity Creation

The sequence diagram in [Figure 8.2](#) shows how the different *Routing Service* entities are created.

Figure 8.2 Entity Creation Sequence Diagram



- An Adapter object is created when the first domain route that refers to it is enabled.
- A Connection object is created when the domain route (`<domain_route>`) that contain it is enabled.
- A Session object is created when the associated routing service session (`<session>`) is enabled.

- ❑ A route's `StreamReader` is created when the route is enabled and the 'creation mode' condition associated with the `<input>` tag becomes true (see [Section 2.4.6.4](#)).
- ❑ A route's `StreamWriter` is created when the route is enabled and the 'creation mode' condition associated with the `<output>` tag becomes true (see [Section 2.4.6.4](#)).

8.2.2 Stream Discovery

A route cannot forward data until the type representations (e.g., `TypeCode`) associated with the input and output streams are available.

If a route refers to types that are not defined in the configuration file, *Routing Service* has to discover their type representation (e.g., `TypeCode`) before creating `StreamReaders` and `StreamWriters`. The adapter discovery API is used to provide stream and type information in a data domain to *Routing Service*.

The discovery API consists of four methods:

- ❑ `Connection::get_input_stream_discovery_reader()`
- ❑ `Connection::get_output_stream_discovery_reader()`
- ❑ `Connection::copy_type_representation()`
- ❑ `Connection::delete_type_representation()`

The first two methods provide access to `StreamReaders` used to discover streams in the data domain associated with a connection.

The input `StreamReader` (`get_input_stream_discovery_reader()`) provides information about input streams. An input stream is a stream from which a `StreamReader` read data. Disposed scenarios, where an input stream disappears, are also notified using the input `StreamReader`.

In the built-in *Connex* adapter, the input `StreamReader` is associated with the publication built-in `DataReader` of the `DomainParticipant`.

The output `StreamReader` (`get_output_stream_discovery_reader()`) provides information about output streams. An output stream is a stream to which `StreamWriters` can write data. Disposed scenarios, where an output stream disappears, are also notified using the output `StreamReader`.

In the built-in *Connex* adapter, the output `StreamReader` is associated with the subscription built-in `DataReader` of the `DomainParticipant`.

The samples provided by the discovery `StreamReaders` have the type `RoutingServiceStreamInfo`.

```

struct RTI_RoutingServiceStreamInfo {
    int disposed;
    char * stream_name;
    struct RTI_RoutingServiceTypeInfo type_info;
};

```

The `disposed` member is used to indicate whether the stream is a new discovered stream or a disposed stream.

The `type_info` member provides information about the type associated with the stream.

```

struct RTI_RoutingServiceTypeInfo {
    char * type_name;
    RTI_RoutingServiceTypeRepresentationKind type_representation_kind;
    RTI_RoutingServiceTypeRepresentation type_representation;
};

```

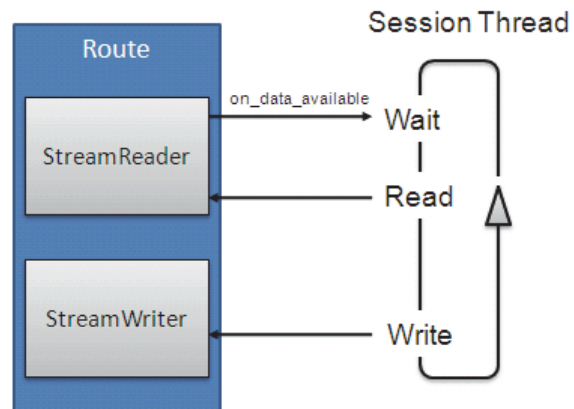
The content associated with the `type_representation` depends on the `type_representation_kind`. For example, if the representation kind is `RTI_ROUTING_SERVICE_TYPE_REPRESENTATION_DYNAMIC_TYPE`, the `type_representation` member will contain a *Connex* `TypeCode`. The method `copy_type_representation()` is used by *Routing Service* to copy the type representation associated with a discovered stream.

8.2.3 Reading Data

Routing Service uses the session threads (there is one per `<session>` tag) to read data from `StreamReaders`.

Each session thread will block waiting for new data using a `WaitSet`. When a `StreamReader` receives new data, it will use the `StreamReaderListener`'s `on_data_available()` callback operation to wake up the session thread associated with it. After that, the session thread will invoke the `StreamReader`'s `read()` operation to get the new data.

The figure to the right describes how the session thread reads samples from a `StreamReader`.



8.3 Creating New Adapters

Routing Service provides an adapter SDK in C and Java to support the creation of new adapter plugins.

The *Routing Service Adapter SDK* is distributed as a separate component that must be installed over an existing installation of *Routing Service*. For more information, see the *Routing Service Adapter SDK SDK Installation Guide*.

8.3.1 Adapter SDK Components

After installing *Routing Service Adapter SDK*, the components in [Table 8.8](#) will be available in the *Routing Service* root folder.

Table 8.8 **Adapter SDK Components**

Component	Description
<i>Release Notes and Installation Guide</i>	Adapter SDK release notes and installation guide. <Routing Service home>/doc/pdf
Adapter SDK Programming Guide	Chapter 8 in the <i>Routing Service User's Manual</i> (this chapter).
API Specification	C and Java API specification in HTML and PDF format. The C API specification describes the Adapter and Transformation API (see Chapter 4). The Java API specification describes the Adapter API. <Routing Service home>/ReadMe.html <Routing Service home>/doc/pdf
Adapter Sample Code	The SDK provides three buildable adapter implementations, two in C (file and socket) and one in Java (JMS). For instructions on compiling and using the sample adapters, see Section 4.9 , Section 4.10 , and Section 4.11 in the <i>Getting Started Guide</i> . Sample Code: <Routing Service home>/adapters Sample Configuration Files: <Routing Service home>/example/shapes
SDK .jar file (rtirsadapter.jar)	The SDK .jar file provides the necessary interfaces and support classes to implement Java adapters (see Section 8.3.5). In addition, the JAR file also includes an implementation of a test adapter (Test-Adapter) that can be used to test new input adapters implementations. JAR Location: <Routing Service home>/class/rtirsadapter.jar

Table 8.8 Adapter SDK Components

Component	Description
SDK infrastructure shared library (<code>[lib]rtirsinfrastructure.dll,.so</code>)	The infrastructure library provides environment (see Section 8.3.2.1) and properties management functions for C adapters. The C adapters will have to link with this library. Library Location: <code><Routing Service home>/bin/<architecture>/[lib]rtirsinfrastructure.dll,.so</code>
SDK header files	The C adapters will have to include two SDK header files: routing_service_adapter.h: This header file defines the adapter API. routing_service_infrastructure.h: This header file defines the public interface of the infrastructure library. Header Location: <code><Routing Service home>/include/routing_service_infrastructure.h</code> <code><Routing Service home>/include/routing_service_adapter.h</code>

8.3.2 C Adapter API

This section does not intend to give complete information on all the C API functions, but rather to describe the aspects of the API that are specific to the C language.

For detailed information about the C API, please see the online (HTML) *Routing Service* documentation.

Every adapter plugin will implement a plugin constructor (entry point to the shared library) that will be used by *Routing Service* to create a plugin instance.

```
typedef
struct RTI_RoutingServiceAdapterPlugin * (
    *RTI_RoutingServiceAdapterPlugin_CreateFcn) (
        const struct RTI_RoutingServiceProperties * properties,
        RTI_RoutingServiceEnvironment * env);
```

The entry point function is specified in the configuration file using the tag `<create_function>` within `<adapter_plugin>` (see [Section 8.1](#)).

The structure `RTI_RoutingServiceAdapterPlugin` will contain the plugin implementation as a set of function pointers. This structure also encapsulates the plugin version information that will be used by *Routing Service* for logging purposes.

```
struct RTI_RoutingServiceAdapterPlugin {
    int _init;
    struct RTI_RoutingServiceVersion _rs_version;
```

```

/* The version of the adapter */
struct RTI_RoutingServiceVersion plugin_version;

RTI_RoutingServiceAdapterPlugin_DeleteFcn
    adapter_plugin_delete;

/* Adapter API */
RTI_RoutingServiceAdapterPlugin_CreateConnectionFcn
    adapter_plugin_create_connection;
RTI_RoutingServiceAdapterPlugin_DeleteConnectionFcn
    adapter_plugin_delete_connection;

/* Connection API */
RTI_RoutingServiceConnection_CreateSessionFcn
    connection_create_session;
RTI_RoutingServiceConnection_DeleteSessionFcn
    connection_delete_session;
RTI_RoutingServiceConnection_CreateStreamReaderFcn
    connection_create_stream_reader;
RTI_RoutingServiceConnection_DeleteStreamReaderFcn
    connection_delete_stream_reader;
RTI_RoutingServiceConnection_CreateStreamWriterFcn
    connection_create_stream_writer;
RTI_RoutingServiceConnection_DeleteStreamWriterFcn
    connection_delete_stream_writer;
RTI_RoutingServiceConnection_GetDiscoveryReaderFcn
    connection_get_input_stream_discovery_reader;
RTI_RoutingServiceConnection_GetDiscoveryReaderFcn
    connection_get_output_stream_discovery_reader;
RTI_RoutingServiceConnection_CopyTypeRepresentationFcn
    connection_copy_type_representation;
RTI_RoutingServiceConnection_DeleteTypeRepresentationFcn
    connection_delete_type_representation;
RTI_RoutingServiceConnection_GetAttributesFcn
    connection_get_attributes;
RTI_RoutingServiceConnection_ToStringFcn
    connection_to_string;
RTI_RoutingServiceAdapterEntity_UpdateFcn
    connection_update;

/* Session API*/
RTI_RoutingServiceAdapterEntity_UpdateFcn
    session_update;

```

```
/* Stream Reader API */
RTI_RoutingServiceStreamReader_ReadFcn
    stream_reader_read;
RTI_RoutingServiceStreamReader_ReturnLoanFcn
    stream_reader_return_loan;
RTI_RoutingServiceAdapterEntity_UpdateFcn
    stream_reader_update;

/* Stream Writer API */
RTI_RoutingServiceStreamWriter_WriteFcn
    stream_writer_write;
RTI_RoutingServiceAdapterEntity_UpdateFcn
    stream_writer_update;

void * user_object;
};
```

The adapter plugin instance created by the entry point function must be initialized using the macro **RTI_RoutingServiceAdapterPlugin_initialize** that is part of the adapter API. For example:

```
struct RTI_RoutingServiceAdapterPlugin *
MyAdapterPlugin_create(
    const struct RTI_RoutingServiceProperties * properties,
    RTI_RoutingServiceEnvironment * env)
{
    struct RTI_RoutingServiceAdapterPlugin * adapter = NULL;
    struct RTI_RoutingServiceVersion version = {1,0,0,0};

    adapter = calloc(1, sizeof(struct RTI_RoutingServiceAdapterPlugin));

    if (adapter == NULL) {
        RTI_RoutingServiceEnvironment_set_error(
            env, "Memory allocation error");
        return NULL;
    }

    RTI_RoutingServiceAdapterPlugin_initialize(adapter);

    adapter->plugin_version = version;

    /*
     * Assign the function pointers
     */
}
```

8.3.2.1 Environment

The last parameter of each adapter API is the environment (`RTI_RoutingServiceEnvironment`). This parameter is used to get information about the *Routing Service* execution such as the version or the verbosity. The environment is also used by the adapter implementations to provide error notification.

8.3.2.2 Adapter verbosity

The C adapter implementations can access the verbosity level used to run *Routing Service* by using the following environment function:

```
RTI_RoutingServiceVerbosity
RTI_RoutingServiceEnvironment_get_verbosity(
    const RTI_RoutingServiceEnvironment * self);
```

The mapping between the command-line option `-verbosity` and the `RTI_RoutingServiceVerbosity` enumeration is as follows:

Table 8.9 Mapping between `-verbosity` and `RTI_RoutingServiceVerbosity`

<code>-verbosity</code>	<code>RTI_RoutingServiceVerbosity</code>
0	<code>RTI_ROUTING_SERVICE_VERBOSITY_NONE</code>
1	<code>RTI_ROUTING_SERVICE_VERBOSITY_EXCEPTION</code>
2	<code>RTI_ROUTING_SERVICE_VERBOSITY_WARN</code>
3 and 4	<code>RTI_ROUTING_SERVICE_VERBOSITY_INFO</code>
5 and 6	<code>RTI_ROUTING_SERVICE_VERBOSITY_DEBUG</code>

8.3.2.3 Version Information

Routing Service and the different adapter implementations are identified by a version number.

The adapter version is provided to *Routing Service* using the member `plugin_version` in the `RTI_RoutingServiceAdapterPlugin` structure. This member must be initialized in the adapter entry point function; it is used by *Routing Service* for logging purposes.

The *Routing Service* version is provided to the C adapters through the environment. The adapters can access this information with the following function:

```
void RTI_RoutingServiceEnvironment_get_version(
    const RTI_RoutingServiceEnvironment * self,
    struct RTI_RoutingServiceVersion * version);
```

8.3.3 My First C Adapter

This section shows how to create a simple C adapter on Windows and UNIX-based systems. It is not intended to give complete coverage of the entire adapter API, but rather to introduce the adapter technology and provide the basic process for developing a C adapter.

The new Adapter will be a simple file adapter where the input adapter reads lines from a text file and the output adapter saves the provided lines to an output text file.

A more flexible and complex file adapter that is able to work with structured information is provided under `<Routing Service home>/adapters/file`.

The source code and projects that you will create in the next sections are provided in `<Routing Service home>/adapters/tutorial/C`.

8.3.3.1 Setting the Environment on the Development Machine

There are a few things to take care of before you start developing the simple file adapter.

1. Install *Connex*; see the *Routing Service Release Notes* for the compatible version of *Connex*.

The C adapter will use `TypeCode` as the type representation format and `DynamicData` as the data representation format. This will require linking against the *Connex* libraries and including the header files defining the `TypeCode` and `DynamicData` APIs.

For information on how to install *Connex*, see the *RTI Core Libraries and Utilities Getting Started Guide*¹.

As part of the installation process, make sure that `NDDSHOME` points to the *Connex* installation directory.

2. Set the environment variable `ROUTINGSERVICEHOME`

Set `ROUTINGSERVICEHOME` to the *Routing Service* installation directory. (*Routing Service* itself does not require that you set the environment variable. It is used to build and compile new adapters).

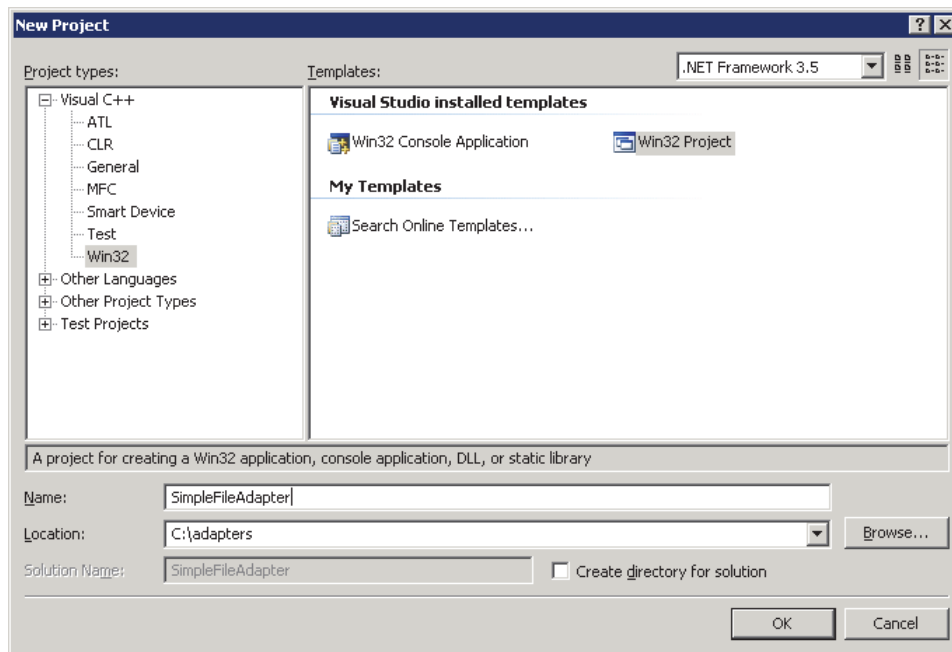
8.3.3.2 Creating a Visual Studio Project (Only for Windows systems)

In this section you will create a Visual Studio project for the adapter dynamic library. We will use Microsoft® Visual Studio® 2008.


1. Start Microsoft Visual Studio 2008.

1. See `<Connex installation directory>/ndds.<version>/doc/pdf/RTI_Connex_GettingStarted.pdf`.

2. Select **File, New, Project, Visual C++, Win32, Win32 Project**. Name the project **SimpleFileAdapter** and select a location.



3. Select **Application Settings** and choose **DLL**. Click **Finish**.



4. Create a new file called **SimpleFileAdapter.c** with the following content. This file will contain the adapter implementation.

```

/*****
/*                               Simple File Adapter                               */
/*****

#include <stdio.h>
#include <string.h>
#ifdef RTI_WIN32
    #include <process.h>
#else
    #include <pthread.h>
#endif

#include "ndds/ndds_c.h"
#include "routingservice/routingservice_adapter.h"

```

```

#ifdef RTI_WIN32
    /* Disable strtok, fopen warnings */
    #pragma warning( disable : 4996 )
    #define DllExport __declspec( dllexport )
#else
    #define DllExport
#endif

/*-----*/
/*          Simple File Adapter: Connection          */
/*-----*/

/*-----*/
/*          Simple File Adapter: StreamReader        */
/*-----*/

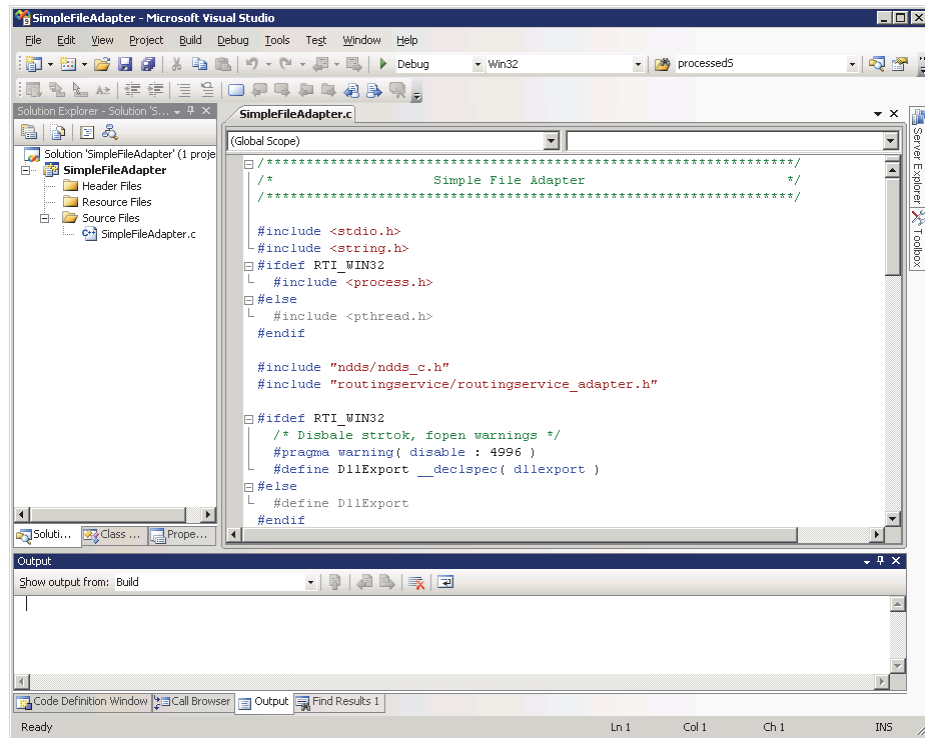
/*-----*/
/*          Simple File Adapter: StreamWriter        */
/*-----*/

/*-----*/
/*          Simple File Adapter: Adapter            */
/*-----*/

/*
 * Entry point to the adapter plugin
 */
DllExport
struct RTI_RoutingServiceAdapterPlugin *
SimpleFileAdapter_create(
    const struct RTI_RoutingServiceProperties * properties,
    RTI_RoutingServiceEnvironment * env)
{
    return NULL;
}

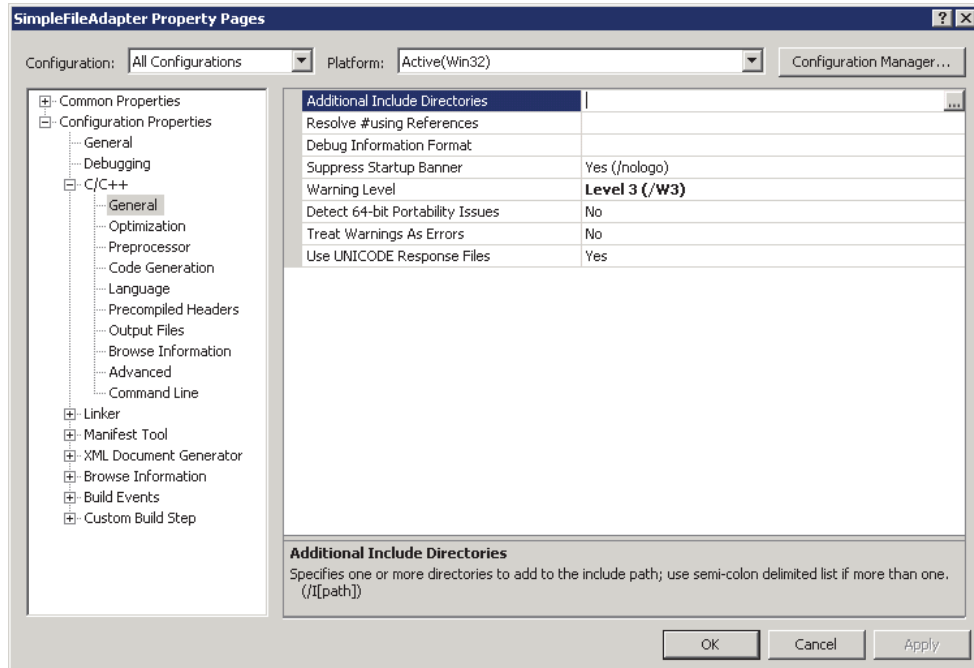
```

5. Add the new file to the project SimpleFileAdapter.



6. Right-click on **SimpleFileAdapter, Properties.**

- In the **Configuration** combo box, select **All Configurations.**
- Select **Configuration Properties, C/C++, General.**



- Add the following to **Additional Include Directories:**

```
$(NDDSHOME)\include;$(NDDSHOME)\include\ndds;  
$(ROUTINGSERVICEHOME)\include
```

- Select **Configuration Properties, Linker, General;** add the following to **Additional Library Directories:**

```
$(NDDSHOME)\lib\i86Win32VS2008;  
$(ROUTINGSERVICEHOME)\bin\i86Win32VS2008
```

- Select **Configuration Properties, Linker, Input;** add the following to **Additional Dependencies:**

```
rtirsinfrastructure.lib nddsc.lib nddscore.lib  
netapi32.lib advapi32.lib user32.lib WS2_32.lib
```

- In the **Configuration** combo box, select **Debug.**

- Select **Configuration Properties, C/C++, Preprocessor**; replace the contents of **Preprocessor Definitions** with:

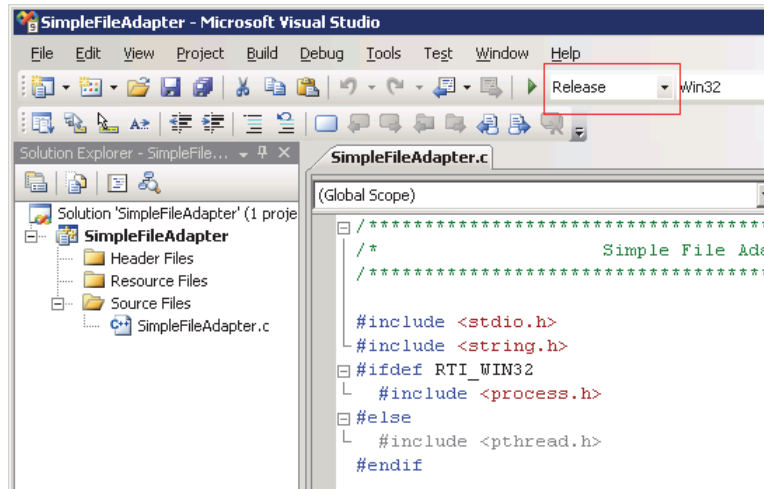
```
WIN32;WIN32_LEAN_AND_MEAN;NDDS_DLL_VARIABLE;RTI_WIN32;_DEBUG
```

- In the Configuration combo box, select **Release**.
- Select **Configuration Properties, C/C++, Preprocessor**; replace the contents of **Preprocessor Definitions** with:

```
WIN32;WIN32_LEAN_AND_MEAN;NDDS_DLL_VARIABLE;RTI_WIN32;NDEBUG
```

- Click **OK**.

7. In the Solution Configuration combo box, select **Release**.



8. Build the SimpleFileAdapter project and verify that there are no errors.

8.3.3.3 Creating an Adapter makefile (Only for UNIX-based systems)

In this section you will create a makefile to generate and compile the adapter shared library.

1. The makefile that you will generate is intended to be used with the GNU distribution of the **make** utility. On modern Linux systems, the **make** binary typically is GNU **make**. On other systems, GNU make is called **gmake**. The instructions below use **gmake**. Make sure that the GNU make binary is on your path before continuing.

2. Create a directory that will contain the adapter makefile and implementation. The rest of this section assumes that `/opt/adapters/simplefile` is the adapter directory.
3. In `/opt/adapters/simplefile`, create a file called `makefile` with the following content.

```
#####
# Makefile to build libsimplfileadapter.so
#####

ARCH = i86Linux2.6gcc4.1.1

c_cc = gcc
c_ld = gcc

ifeq ($(DEBUG),1)
c_cc_flags = -m32 -g
else
c_cc_flags = -m32
endif

c_ld_flags = -m32 -static-libgcc
syslibs = -ldl -lnsl -lm -lpthread

DEFINES_ARCH_SPECIFIC = -DRTI_UNIX -DRTI_LINUX
DEFINES = $(DEFINES_ARCH_SPECIFIC)

INCLUDES = -I. -I$(NDDSHOME)/include -I$(NDDSHOME)/include/ndds \
           -I$(ROUTINGSERVICEHOME)/include

LIBS = -L$(NDDSHOME)/lib/$(ARCH) -L$(ROUTINGSERVICEHOME)/bin/$(ARCH) \
       -lnddsc -lnddscore -lrtirsinfrastructure $(syslibs) $(extralibs)

COMMONSOURCES = SimpleFileAdapter.c
SHAREDLIB      = lib/$(ARCH)/libsimplfileadapter.so
DIRECTORIES   = lib.dir lib/$(ARCH).dir objs.dir objs/$(ARCH).dir
COMMONOBSJS   = $(COMMONSOURCES:%.c=objs/$(ARCH)/%.o)

$(ARCH) : $(DIRECTORIES) $(COMMONOBSJS) $(SHAREDLIB)

$(SHAREDLIB) : $(COMMONOBSJS)
               $(c_cc) $(c_ld_flags) -shared -o $@ $^ $(LIBS)

objs/$(ARCH)/%.o : %.c
```

```

$(c_cc) $(c_cc_flags) -o $@ $(DEFINES) $(INCLUDES) -c $<

# Here is how we create those subdirectories automatically.
%.dir :
    @echo "Checking directory $*"
    @if [ ! -d $* ]; then \
        echo "Making directory $*"; \
        mkdir -p $* ; \
    fi;
clean:
    @rm -rf ./objs
    @rm -rf ./lib

```

The above makefile assumes that the architecture is i86Linux2.6gcc4.1.1. If you are building for a different architecture, you can use the above makefile as an example.

4. Create a new file called **SimpleFileAdapter.c** with the following content. This file will contain the adapter implementation.

```

/*****
/*                               Simple File Adapter                               */
*****/

#include <stdio.h>
#include <string.h>
#ifdef RTI_WIN32
    #include <process.h>
#else
    #include <pthread.h>
#endif

#include "ndds/ndds_c.h"
#include "routingservice/routingservice_adapter.h"

#ifdef RTI_WIN32
    /* Disable strtok, fopen warnings */
    #pragma warning( disable : 4996 )
    #define DllExport __declspec( dllexport )
#else
    #define DllExport
#endif

/*-----*/
/*                               Simple File Adapter: Connection                               */
/*-----*/

```



```

/*-----*/
/*          Simple File Adapter: StreamReader          */
/*-----*/

/*-----*/
/*          Simple File Adapter: StreamWriter          */
/*-----*/

/*-----*/
/*          Simple File Adapter: Adapter              */
/*-----*/

/*
 * Entry point to the adapter plugin
 */
DllExport
struct RTI_RoutingServiceAdapterPlugin *
SimpleFileAdapter_create(
    const struct RTI_RoutingServiceProperties * properties,
    RTI_RoutingServiceEnvironment * env)
{
    return NULL;
}

```

5. Compile the SimpleFileAdapter skeleton by executing **gmake** from the adapter directory.

```
> gmake
```

After compilation, you will find the adapter library in **/opt/adapters/simplefile/lib/<architecture>**. The next few sections will show you how to complete the adapter implementation.

8.3.3.4 Initializing the Adapter Entry Point Function

Every adapter plugin must implement a plugin constructor (entry point to the dynamic library) that will be used by *Routing Service* to create a plugin instance (see [Section 8.3.2](#)). In this example, the entry point is the function **SimpleFileAdapter_create** in the file **SimpleFileAdapter.c**. You have to initialize this function to create a new plugin.

```

/*
 * Plugin destructor
 */
void SimpleFileAdapter_delete(
    struct RTI_RoutingServiceAdapterPlugin * adapter,
    RTI_RoutingServiceEnvironment * env)

```

```
{
    free(adapter);
}

/*
 * Entry point to the adapter plugin
 */
DllExport
struct RTI_RoutingServiceAdapterPlugin *
SimpleFileAdapter_create(
    const struct RTI_RoutingServiceProperties * properties,
    RTI_RoutingServiceEnvironment * env)
{
    struct RTI_RoutingServiceAdapterPlugin * adapter = NULL;
    struct RTI_RoutingServiceVersion version = {1,0,0,0};
    int verbosity;

    verbosity = RTI_RoutingServiceEnvironment_get_verbosity(env);
    if (verbosity == RTI_ROUTING_SERVICE_VERBOSITY_DEBUG) {
        printf("CALL SimpleFileAdapter_create\n");
    }

    adapter = calloc(1, sizeof(struct RTI_RoutingServiceAdapterPlugin));
    if (adapter == NULL) {
        RTI_RoutingServiceEnvironment_set_error(
            env, "Memory allocation error");
        return NULL;
    }

    RTI_RoutingServiceAdapterPlugin_initialize(adapter);

    adapter->plugin_version = version;

    /*
     * Assign the function pointers
     */
    adapter->adapter_plugin_delete = SimpleFileAdapter_delete;

    return (struct RTI_RoutingServiceAdapterPlugin *) adapter;
}
```

The structure `RTI_RoutingServiceAdapterPlugin` contains the plugin implementation as a set of function pointers. For now, you only need to implement `adapter_plugin_delete` that deletes the plugin instances created by `SimpleFileAdapter_create()`. You will ini-

tialize the other pointers in the plugin structure as you implement the adapter functionality.

The entry point function receives two parameters: the adapter **properties** and the environment, **env**.

The **properties** parameter (not used by the SimpleFileAdapter) is used to configure the adapter instance. The values contained in this parameter are provided as (name,value) pairs using the tag <property> within <adapter_plugin> (see [Adapter Usage and Configuration \(Section 8.1\)](#)).

The environment parameter, **env**, is part of every function in the adapter API. This parameter is used to get information about the *Routing Service* execution such as the version or the verbosity. In addition, the environment is also used to notify *Routing Service* of any error in the adapter execution.

8.3.3.5 Implementing the Adapter Connection

The adapter plugin instances are connection factories. Connection objects provide access to data domains such as *Connex*t domains or JMS network providers and they are configured using the XML tags <connection_1> and <connection_2> in a <domain_route> (see [Section 2.4.2](#)). In the SimpleFileAdapter example, the connection objects will provide access to a directory on your computer's file system.

The next step consist on implementing the functions that create and delete a connection. Insert the following code in the "Simple File Adapter: Connection" section of **SimpleFileAdapter.c**.

```
/*
 * Connection.
 */
struct SimpleFileAdapterConnection {
    char * directory;
};

/*
 * Deletes a connection.
 */
void SimpleFileAdapter_delete_connection(
    struct RTI_RoutingServiceAdapterPlugin * adapter,
    RTI_RoutingServiceConnection connection,
    RTI_RoutingServiceEnvironment * env)
{
    struct SimpleFileAdapterConnection * cx =
        (struct SimpleFileAdapterConnection *)connection;
    int verbosity;
```

```
    verbosity = RTI_RoutingServiceEnvironment_get_verbosity(env);

    if (verbosity == RTI_ROUTING_SERVICE_VERBOSITY_DEBUG) {
        printf("CALL SimpleFileAdapter_delete_connection\n");
    }

    if (cx->directory != NULL) {
        free(cx->directory);
    }
    free(cx);
}

/*
 * Creates a connection.
 */
RTI_RoutingServiceConnection
SimpleFileAdapter_create_connection(
    struct RTI_RoutingServiceAdapterPlugin * adapter,
    const char * routing_service_name,
    const char * routing_service_group_name,
    const struct RTI_RoutingServiceStreamReaderListener *
input_disc_listener,
    const struct RTI_RoutingServiceStreamReaderListener *
output_disc_listener,
    const struct RTI_RoutingServiceTypeInfo ** registeredTypes,
    int registeredTypeCount,
    const struct RTI_RoutingServiceProperties * properties,
    RTI_RoutingServiceEnvironment * env)
{
    const char * directory;
    struct SimpleFileAdapterConnection * cx;
    int verbosity;

    verbosity = RTI_RoutingServiceEnvironment_get_verbosity(env);

    if (verbosity == RTI_ROUTING_SERVICE_VERBOSITY_DEBUG) {
        printf("CALL SimpleFileAdapter_create_connection\n");
    }

    cx = calloc(1, sizeof(struct SimpleFileAdapterConnection));

    if (cx == NULL) {
        RTI_RoutingServiceEnvironment_set_error(
            env, "Memory allocation error");
        return NULL;
    }
}
```

```

    }

    directory = RTI_RoutingServiceProperties_lookup_property(
        properties,
        "directory");

    if (directory == NULL) {
        RTI_RoutingServiceEnvironment_set_error(
            env, "directory property is required");
        free(cx);
        return NULL;
    }
    cx->directory = (char *)malloc(strlen(directory)+1);

    if (cx->directory == NULL) {
        RTI_RoutingServiceEnvironment_set_error(
            env, "Memory allocation error");
        free(cx);
        return NULL;
    }

    strcpy(cx->directory, directory);
    return cx;
}

```

From the implementation, you can see that the connection object encapsulates the name of the directory from which the StreamReaders and StreamWriters will read and write files.

The value of the `RTI_RoutingServiceAdapterPlugin` structure created in `SimpleFileAdapter_create()` must be updated to contain the two new functions.

```

adapter->adapter_plugin_create_connection =
    SimpleFileAdapter_create_connection;
adapter->adapter_plugin_delete_connection =
    SimpleFileAdapter_delete_connection;

```

8.3.3.6 Implementing the StreamReader

The connection objects are factories of StreamReaders. A *StreamReader* provides a way to read data samples of a specific type from a data domain.

In the configuration file, *StreamReaders* are associated with the tag `<input>` within `<route>` or `<auto_route>` (see [Section 2.4.6](#)).

The StreamReaders created by the SimpleFileAdapter connections read text files from the connection directory.

The data samples provided to *Routing Service* (using the read operation) are Dynamic-Data with the following IDL type:

```
struct TextLine {
    string<1024> value;
};
```

When a `SimpleFileAdapter StreamReader` is created, the name of the file is the input stream name with a `.txt` extension. You can use the `read_period` property to control how often the `StreamReader` notifies *Routing Service* about new lines. For example:

```
<route name="route">
  <input connection="1">
    <stream_name>HelloWorld</stream_name>
    <registered_type_name>TextLine</registered_type_name>
    <property>
      <value>
        <element>
          <name>read_period</name>
          <value>1000</value>
        </element>
      </value>
    </property>
  </input>
  ...
</route>
```

In the above example, the input `StreamReader` will read lines from a file called **HelloWorld.txt** and provide one line per second to *Routing Service*.

The next step is to implement the `StreamReader` functionality. You will implement five new functions:

- ❑ **SimpleFileAdapterStreamReader_read()**: This function will be called by *Routing Service* after being notified that there are new lines available. Although the signature of the function allows returning more than one sample (line), for the sake of simplicity, the implementation only returns one line each time the function is called.
- ❑ **SimpleFileAdapterStreamReader_return_loan()**: The loan on the samples provided by `SimpleFileAdapterStreamReader_read()` is returned to the adapter using this function. The `SimpleFileAdapter` implementation of `return_loan()` is empty because:
 - The read operation does not create new samples and always returns a single sample stored in the `StreamReader`.

- Two calls to `SimpleFileAdapterStreamReader_read()` cannot occur in parallel.
- ❑ **SimpleFileAdapterStreamReader_run:** *Routing Service* will not call the read operation until it is notified of the presence of new data (see [Section 8.2.3](#)). To provide data notification, the `StreamReader` implementation creates a thread that wakes up after `read_period` and notifies *Routing Service* of new data if the end of the file has not been reached yet. `SimpleFileAdapterStreamReader_run` is the function executed by the notification thread.
- ❑ **SimpleFileAdapterConnection_delete_stream_reader:** This function is called to destroy a `StreamReader`. The implementation will finalize the notification thread and close the file handle.
- ❑ **SimpleFileAdapterConnection_create_stream_reader:** This function is called when a new `StreamReader` is created. Among other things, the implementation will open the file that will be read and create the notification thread.

Insert the following code in the “Simple File Adapter: StreamReader” section of **SimpleFileAdapter.c**.

```

/*
 * StreamReader.
 */
struct SimpleFileAdapterStreamReader {
    int run;
#ifdef RTI_WIN32
    HANDLE thread;
#else
    pthread_t thread;
#endif
    DDS_DynamicData * sample[1];
    struct DDS_Duration_t readPeriod;
    struct RTI_RoutingServiceStreamReaderListener listener;
    FILE * fHandle;
};

/*
 * Returns sample loan
 */
void SimpleFileAdapterStreamReader_return_loan(
    RTI_RoutingServiceStreamReader stream_reader,
    RTI_RoutingServiceSample * sample_list,
    RTI_RoutingServiceSampleInfo * info_list,
    int count,
    RTI_RoutingServiceEnvironment * env)
{

```

```
int verbosity;

verbosity = RTI_RoutingServiceEnvironment_get_verbosity(env);

if (verbosity == RTI_ROUTING_SERVICE_VERBOSITY_DEBUG) {
    printf("CALL SimpleFileAdapterStreamReader_return_loan\n");
}

return;
}

/*
 * Reads one line from the StreamReader file
 */
void SimpleFileAdapterStreamReader_read(
    RTI_RoutingServiceStreamReader stream_reader,
    RTI_RoutingServiceSample ** sample_list,
    RTI_RoutingServiceSampleInfo ** info_list,
    int * count,
    RTI_RoutingServiceEnvironment * env)
{
    DDS_ReturnCode_t retCode;
    char line[2048];
    char * str;
    struct SimpleFileAdapterStreamReader * self =
        (struct SimpleFileAdapterStreamReader *) stream_reader;
    int verbosity;
    int length;

    verbosity = RTI_RoutingServiceEnvironment_get_verbosity(env);

    if (verbosity == RTI_ROUTING_SERVICE_VERBOSITY_DEBUG) {
        printf("CALL SimpleFileAdapterStreamReader_read\n");
    }

    *count = 0;
    *sample_list = NULL;

    /*
     * We don't provide sample info in this adapter, which
     * is an optional feature
     */
    *info_list = NULL;

    DDS_DynamicData_clear_all_members(self->sample[0]);
}
```



```

str = fgets(line, sizeof(line), self->fHandle);

if (!str) {
    return;
}

length = strlen(str);
if (length > 0 && str[length-1] == '\n') {
    str[length-1] = '\0';
    if (length > 1 && str[length-2] == '\r') {
        str[length-2] = '\0';
    }
}

retCode = DDS_DynamicData_set_string(
    self->sample[0], "value",
    DDS_DYNAMIC_DATA_MEMBER_ID_UNSPECIFIED,
    line);

if (retCode != DDS_RETCODE_OK) {
    RTI_RoutingServiceEnvironment_set_error(
        env, "Error assigning value=%s", line);
    return;
}

*sample_list = (RTI_RoutingServiceSample *)self->sample;
*count = 1;
return;
}

/*
 * Notification thread.
 *
 * This thread will notify of data availability in the file.
 */
void * SimpleFileAdapterStreamReader_run(void * threadParam)
{
    struct SimpleFileAdapterStreamReader * self =
        (struct SimpleFileAdapterStreamReader *) threadParam;

    while (self->run) {
        NDDS_Utility_sleep(&self->readPeriod);

        if (!feof(self->fHandle)) {
            self->listener.on_data_available(
                self, self->listener.listener_data);
        }
    }
}

```

```
        }
    }

    return NULL;
}

/*
 * Deletes a StreamReader.
 */
void SimpleFileAdapterConnection_delete_stream_reader(
    RTI_RoutingServiceConnection connection,
    RTI_RoutingServiceStreamReader stream_reader,
    RTI_RoutingServiceEnvironment * env)
{
    struct SimpleFileAdapterStreamReader * reader =
        (struct SimpleFileAdapterStreamReader *) stream_reader;
#ifdef RTI_WIN32
    void * value = NULL;
#endif

    int verbosity;

    verbosity = RTI_RoutingServiceEnvironment_get_verbosity(env);

    if (verbosity == RTI_ROUTING_SERVICE_VERBOSITY_DEBUG) {
        printf("CALL SimpleFileAdapterConnection_delete_stream_reader\n");
    }

    reader->run = 0;

#ifdef RTI_WIN32
    WaitForSingleObject(reader->thread, INFINITE);
#else
    pthread_join(reader->thread, &value);
#endif

    if (reader->fHandle) {
        fclose(reader->fHandle);
    }

    if (reader->sample[0]) {
        DDS_DynamicData_delete(reader->sample[0]);
    }

    free(reader);
}
```

```

/*
 * Creates a StreamReader
 */
RTI_RoutingServiceStreamReader
SimpleFileAdapterConnection_create_stream_reader(
    RTI_RoutingServiceConnection connection,
    RTI_RoutingServiceSession session,
    const struct RTI_RoutingServiceStreamInfo * stream_info,
    const struct RTI_RoutingServiceProperties * properties,
    const struct RTI_RoutingServiceStreamReaderListener * listener,
    RTI_RoutingServiceEnvironment * env)
{
    const char * readPeriodStr;
    unsigned int readPeriod;
    char * file;
    struct SimpleFileAdapterConnection * self =
        (struct SimpleFileAdapterConnection *)connection;
    struct SimpleFileAdapterStreamReader * reader = NULL;
    struct DDS_DynamicDataProperty_t dynamicDataProps =
        DDS_DynamicDataProperty_t_INITIALIZER;
    int error = 0;
#ifdef RTI_WIN32
    pthread_attr_t threadAttr;
#endif
    int verbosity;

    verbosity = RTI_RoutingServiceEnvironment_get_verbosity(env);

    if (verbosity == RTI_ROUTING_SERVICE_VERBOSITY_DEBUG) {
        printf("CALL SimpleFileAdapterConnection_create_stream_reader\n");
    }

    /* Create StreamReader */
    reader = calloc(1, sizeof(struct SimpleFileAdapterStreamReader));

    if (reader == NULL) {
        RTI_RoutingServiceEnvironment_set_error(
            env, "Memory allocation error");
        return NULL;
    }

    reader->sample[0] = DDS_DynamicData_new(
        (struct DDS_TypeCode *) stream_info->type_info.type_representation,
        &dynamicDataProps);
}

```

```
if (reader->sample[0] == NULL) {
    RTI_RoutingServiceEnvironment_set_error(
        env, "Failure creating dynamic data sample");
    free(reader);
    return NULL;
}

/* Open input file */
file = malloc(strlen(self->directory) + strlen("/") +
              strlen(stream_info->stream_name) +
              strlen(".txt") + 1);

if (file == NULL) {
    RTI_RoutingServiceEnvironment_set_error(
        env, "Memory allocation error");
    DDS_DynamicData_delete(reader->sample[0]);
    free(reader);
    return NULL;
}

sprintf(file, "%s/%s.txt", self->directory, stream_info->stream_name);

reader->fHandle = fopen(file, "r");

if (reader->fHandle == NULL) {
    RTI_RoutingServiceEnvironment_set_error(
        env, "Error opening %s", file);
    free(file);
    DDS_DynamicData_delete(reader->sample[0]);
    free(reader);
    return NULL;
}

free(file);

/* Creates notification thread */
readPeriodStr =
    RTI_RoutingServiceProperties_lookup_property(
        properties,
        "read_period");

if (readPeriodStr != NULL) {
    readPeriod = atoi(readPeriodStr);
} else {
    readPeriod = 1000; /* 1 Sec */
}
```

```

reader->readPeriod.sec = readPeriod / 1000;
reader->readPeriod.nanosec = (readPeriod % 1000) * 1000000;
reader->run = 1;
reader->listener = *listener;

#ifdef RTI_WIN32
reader->thread = (HANDLE) _beginthread(
    (void(__cdecl*)(void*))SimpleFileAdapterStreamReader_run,
    0, (void*)reader);

if (!reader->thread) {
    error = 1;
}
#else
pthread_attr_t threadAttr;
pthread_attr_init(&threadAttr);
pthread_attr_setdetachstate(&threadAttr, PTHREAD_CREATE_JOINABLE);
error = pthread_create(
    &reader->thread,
    &threadAttr,
    SimpleFileAdapterStreamReader_run,
    (void *)reader);
pthread_attr_destroy(&threadAttr);
#endif

if (error) {
    RTI_RoutingServiceEnvironment_set_error(
        env, "Error creating notification thread");
    DDS_DynamicData_delete(reader->sample[0]);
    free(reader);
    fclose(reader->fHandle);
    return NULL;
}

return reader;
}

```

The value of the `RTI_RoutingServiceAdapterPlugin` structure created in `SimpleFileAdapter_create()` must be updated to contain the `StreamReader` functions.

```

adapter->connection_create_stream_reader =
    SimpleFileAdapterConnection_create_stream_reader;
adapter->connection_delete_stream_reader =
    SimpleFileAdapterConnection_delete_stream_reader;
adapter->stream_reader_read =
    SimpleFileAdapterStreamReader_read;

```

```

adapter->stream_reader_return_loan =
    SimpleFileAdapterStreamReader_return_loan;

```

8.3.3.7 Implementing the StreamWriter

The connection objects are factories of *StreamWriters*. A *StreamWriter* provides a way to write samples of a specific type into a data domain.

In the configuration file, *StreamWriters* are associated with the tag `<output>` within `<route>` or `<auto_route>` (see [Section 2.4.6](#)).

The SimpleFileAdapter *StreamWriters* create new files into the connection directory and store the lines read from the routes' inputs.

The data samples provided to the write operation of the *StreamWriters* are Dynamic-Data with the following IDL type:

```

struct TextLine {
    string<1024> value;
};

```

When a SimpleFileAdapter *StreamWriter* is created, the name of the file is the output stream name with ".txt" extension. For debugging purposes, the *StreamWriter* can be configured to print the written samples on the console:

```

<route name="route">
    ...
    <output>
        <stream_name>HelloWorld</stream_name>
        <registered_type_name>TextLine</registered_type_name>
        <property>
            <value>
                <element>
                    <name>print_to_stdout</name>
                    <value>1</value>
                </element>
            </value>
        </property>
    </output>
</route>

```

In the previous example, the output *StreamWriter* will store the lines provided by *Routing Service* on a file called **HelloWorld.txt**. It will also print the lines on the screen.

Insert the following code in the "Simple File Adapter: StreamWriter" section of **SimpleFileAdapter.c**.

```

/*
 * StreamWriter

```

```

*/
struct SimpleFileAdapterStreamWriter {
    int printToStdout;
    FILE * fHandle;
};

int SimpleFileAdapterStreamWriter_write(
    RTI_RoutingServiceStreamWriter stream_writer,
    const RTI_RoutingServiceSample * sample_list,
    const RTI_RoutingServiceSampleInfo * info_list,
    int count,
    RTI_RoutingServiceEnvironment * env)
{
    int i, samplesWritten;
    DDS_DynamicData * sample;
    DDS_ReturnCode_t retCode;
    char * line;
    struct SimpleFileAdapterStreamWriter * self =
        (struct SimpleFileAdapterStreamWriter *) stream_writer;
    int verbosity;

    verbosity = RTI_RoutingServiceEnvironment_get_verbosity(env);

    if (verbosity == RTI_ROUTING_SERVICE_VERBOSITY_DEBUG) {
        printf("CALL SimpleFileAdapterStreamWriter_write\n");
    }

    if (!self->fHandle) {
        return 0;
    }

    samplesWritten = 0;

    for (i=0; i<count; i++) {
        sample = (DDS_DynamicData *)sample_list[i];

        line = NULL;

        retCode = DDS_DynamicData_get_string(
            sample, &line, NULL, "value",
            DDS_DYNAMIC_DATA_MEMBER_ID_UNSPECIFIED);

        if (retCode != DDS_RETCODE_OK) {
            RTI_RoutingServiceEnvironment_set_error(
                env, "Error assigning value");
        } else {

```

```
        samplesWritten++;
    }

    fputs(line, self->fHandle);
    fputs("\n", self->fHandle);
    fflush(self->fHandle);

    if (self->printToStdout) {
        printf("%s\n", line);
        fflush(stdout);
    }

    DDS_String_free(line);
}

return samplesWritten;
}

/*
 * Deletes a StreamWriter
 */
void SimpleFileAdapterConnection_delete_stream_writer(
    RTI_RoutingServiceConnection connection,
    RTI_RoutingServiceStreamWriter stream_writer,
    RTI_RoutingServiceEnvironment * env)
{
    struct SimpleFileAdapterStreamWriter * writer =
        (struct SimpleFileAdapterStreamWriter *) stream_writer;
    int verbosity;

    verbosity = RTI_RoutingServiceEnvironment_get_verbosity(env);

    if (verbosity == RTI_ROUTING_SERVICE_VERBOSITY_DEBUG) {
        printf("CALL SimpleFileAdapterConnection_delete_stream_writer\n");
    }

    if (writer->fHandle) {
        fclose(writer->fHandle);
    }

    free(writer);
}

/*
 * Creates a StreamWriter
 */
```



```

RTI_RoutingServiceStreamWriter
SimpleFileAdapterConnection_create_stream_writer(
    RTI_RoutingServiceConnection connection,
    RTI_RoutingServiceSession session,
    const struct RTI_RoutingServiceStreamInfo * stream_info,
    const struct RTI_RoutingServiceProperties * properties,
    RTI_RoutingServiceEnvironment * env)
{
    const char * printToStdoutStr;
    char * file;
    struct SimpleFileAdapterConnection * self =
        (struct SimpleFileAdapterConnection *)connection;
    struct SimpleFileAdapterStreamWriter * writer = NULL;
    int verbosity;

    verbosity = RTI_RoutingServiceEnvironment_get_verbosity(env);

    if (verbosity == RTI_ROUTING_SERVICE_VERBOSITY_DEBUG) {
        printf("CALL SimpleFileAdapterConnection_create_stream_writer\n");
    }

    /* Create StreamWriter */
    writer = calloc(1, sizeof(struct SimpleFileAdapterStreamWriter));

    if (writer == NULL) {
        RTI_RoutingServiceEnvironment_set_error(
            env, "Memory allocation error");
        return NULL;
    }

    /* Open output file */
    file = malloc(strlen(self->directory) + strlen("/") +
        strlen(stream_info->stream_name) +
        strlen(".txt") + 1);

    if (file == NULL) {
        RTI_RoutingServiceEnvironment_set_error(
            env, "Memory allocation error");
        free(writer);
        return NULL;
    }

    sprintf(file, "%s/%s.txt", self->directory, stream_info->stream_name);

    writer->fHandle = fopen(file, "w+");
}

```

```
if (writer->fHandle == NULL) {
    RTI_RoutingServiceEnvironment_set_error(
        env, "Error opening %s", file);
    free(file);
    free(writer);
    return NULL;
}

free(file);

/* Creates notification thread */
printToStdoutStr =
    RTI_RoutingServiceProperties_lookup_property(
        properties,
        "print_to_stdout");

if (printToStdoutStr != NULL) {
    writer->printToStdout = atoi(printToStdoutStr);
} else {
    writer->printToStdout = 0;
}

return writer;
}
```

The value of the `RTI_RoutingServiceAdapterPlugin` structure created in `SimpleFileAdapter_create()` must be updated to contain the `StreamWriter` functions.

```
adapter->connection_create_stream_writer =
    SimpleFileAdapterConnection_create_stream_writer;
adapter->connection_delete_stream_writer =
    SimpleFileAdapterConnection_delete_stream_writer;
adapter->stream_writer_write =
    SimpleFileAdapterStreamWriter_write;
```

8.3.3.8 Running the SimpleFileAdapter

This section describes the steps required to use and run the `SimpleFileAdapter` with *Routing Service*. You will create a configuration file with a single route that reads a Hello-World text file from an input directory and saves it into an output directory.

1. If you have not done it yet, compile and build the `SimpleFileAdapter`.
2. Under the adapter project directory (`c:\adapters\SimpleFileAdapter1` on Windows systems; `/opt/adapters/simplefile1` on UNIX-based systems) create two directories called `input` and `output`.

3. In the input directory create a file called **HelloWorld.txt** with the following content.

```
Hello World 1!
Hello World 2!
Hello World 3!
Hello World 4!
Hello World 5!
Hello World 6!
Hello World 7!
Hello World 8!
Hello World 9!
Hello World 10!
```

4. In the adapter project directory create a *Routing Service* XML configuration file called **simple_file_adapter.xml** with the following content.

Replace the value of the “directory” property under both connections with the location of the input and output directories.

Replace the content of the dll tag under adapter_plugin with the location of the release version of the SimpleFileAdapter shared library.

```
<?xml version="1.0"?>
<dds>
  <adapter_library name="adapters">
    <adapter_plugin name="simple_file">
      <dll>
        c:\adapters\SimpleFileAdapter\Release\SimpleFileAdapter.dll
      </dll>
      <create_function>SimpleFileAdapter_create</create_function>
    </adapter_plugin>
  </adapter_library>

  <types>
    <struct name="TextLine">
      <member name="value" type="string" stringMaxLength="2048"/>
    </struct>
  </types>

  <routing_service name="file_to_file">
    <domain_route name="domain_route">
      <connection_1 plugin_name="adapters::simple_file">
        <registered_type name="TextLine" type_name="TextLine"/>
        <property>
          <value>
```

1. Your directory may be different if you did not use the default locations.

```
        <element>
          <name>directory</name>
          <value>c:\adapters\SimpleFileAdapter\input</value>
        </element>
      </value>
    </property>
  </connection_1>

  <connection_2 plugin_name="adapters::simple_file">
    <registered_type name="TextLine" type_name="TextLine"/>
    <property>
      <value>
        <element>
          <name>directory</name>
          <value>c:\adapters\SimpleFileAdapter\output</value>
        </element>
      </value>
    </property>
  </connection_2>

  <session name="session">
    <route name="route">
      <input connection="1">
        <stream_name>HelloWorld</stream_name>
        <registered_type_name>TextLine</registered_type_name>
      </input>
      <output>
        <stream_name>HelloWorld</stream_name>
        <registered_type_name>TextLine</registered_type_name>
        <property>
          <value>
            <element>
              <name>print_to_stdout</name>
              <value>1</value>
            </element>
          </value>
        </property>
      </output>
    </route>
  </session>
</domain_route>
</routing_service>
</dds>
```

5. Start *Routing Service* by entering the following in a command shell.

On UNIX-based systems:

```
> cd <SimpleFileAdapter project directory>
> $ROUTINGSERVICEHOME/scripts/rtiroutingservice
  -cfgFile simple_file_adapter.xml -cfgName file_to_file
```

On Windows systems:

```
> cd <SimpleFileAdapter project directory>
> %ROUTINGSERVICEHOME%\scripts\rtiroutingservice
  -cfgFile simple_file_adapter.xml -cfgName file_to_file
```

6. On the screen you will see:

```
RTI Routing Service <version> started (with name file_to_file)
Hello World 1!
Hello World 2!
Hello World 3!
Hello World 4!
Hello World 5!
Hello World 6!
Hello World 7!
Hello World 8!
Hello World 9!
Hello World 10!
```

7. Verify that a file called **HelloWorld.txt has been generated into the output directory. The content of this file should be identical to the content of the same file in the input directory.****8.3.4 Debugging C Adapters**

When you develop a custom adapter you will need to debug it and test it. This section talks about the tools and APIs that you have available to debug and detect problems in *Routing Service* adapters written in C.

The first debugging capability is provided by the *Routing Service* SDK. The adapter SDK provides a way to access the verbosity level of *Routing Service* through the usage of the environment function **RTI_RoutingServiceEnvironment_get_verbosity**. It is highly recommendable that as part of the adapter implementation you instrument the code by adding status messages that will be printed with the INFO and DEBUG verbosity levels. This level of instrumentation will help you to capture run-time information for troubleshooting.

The second debugging capability is provided by third party tools. On a Windows system, you can debug the adapter shared libraries using Visual Studio. On a UNIX-based system, you can use GDB, the GNU Project debugger.


8.3.4.1 Debugging the Adapter with Visual Studio—only for Windows Systems

Let's see how to debug the adapter library with Visual Studio 2008 using the SimpleFileAdapter implemented in [Section 8.3.3](#).

1. Start Microsoft Visual Studio 2008 and open the solution **SimpleFileAdapter**.
2. In the Solution Configuration combo box select **Debug** configuration and recompile the SimpleFileAdapter project.
3. Edit **simple_file_adapter.xml**, the configuration file generated in [Section 8.3.3.8](#). Replace the library in the <dll> tag with the debug version of the adapter. For example:

```
<dll>c:\adapters\SimpleFileAdapter\Debug\SimpleFileAdapter.dll</dll>
```

4. Right-click on **SimpleFileAdapter**, **Properties**



- In the configuration combo box select **Debug**.
- Under **Configuration Properties, Debugging**; go to “Command” and add the following:
`${ROUTINGSERVICEHOME}\bin\i86Win32VS2008\rtiroutingservice`

- Under **Configuration Properties, Debugging**; go to “Command Arguments” and add the following:

```
-cfgFile c:\adapters\SimpleFileAdapter\simple_file_adapter.xml1
-cfgName file_to_file
```

- Click **OK**.

5. Open the file **SimpleFileAdapter.c** and insert breakpoints in the functions that you want to debug. Then press **F5** to run *Routing Service* and debug the adapter.

If you get an information window that says there is no debugging information in *rtiroutingservice*, press **YES**. Although *rtiroutingservice* does not have debugging symbols, your adapter was built with debug information and you should not have any problems debugging it.

8.3.4.2 Debugging the Adapter with GDB—only for UNIX-based systems

Let’s see how to debug the adapter library with **gdb** using the **SimpleFileAdapter** implemented in [Section 8.3.3](#).

1. Go to the directory containing the **SimpleFileAdapter** makefile and build the debug version of the shared library as follows:

```
> gmake clean
> gmake DEBUG=1
```

The debug version of the adapter replaces the release version because is generated in the same location.

2. Edit the configuration file `simple_file_adapter.xml` generated in [Section 8.3.3.8](#) and replace the library in the `<dll>` tag with the debug version of the adapter. For example:

```
<dll>/opt/adapters/simplefile/lib/i86Linux2.6gcc4.1.1/libsimplefileadapter.so</dll>
```

3. Run **gdb**:

```
> gdb $ROUTINGSERVICEHOME/bin/i86Linux2.6gcc4.1.1/rtiroutingservice
```

4. Insert breakpoints in the functions that you want to debug. For example:

```
(gdb) b SimpleFileAdapter_create_connection
Function "SimpleFileAdapter_create_connection" not defined.
Make breakpoint pending on future shared library load? (y or [n]) y
Breakpoint 1 (SimpleFileAdapter_create_connection) pending.
```

1. The location of your configuration file may be different. Replace the value with the right location.

- Execute *Routing Service* and debug your adapter.

```
(gdb) r -cfgFile simple_file_adapter.xml -cfgName file_to_file
```

8.3.5 Java Adapter API

This section does not intend to give complete information on the entire Java API, but rather to describe the aspects of the Java API that are specific to the Java language.

For detailed information about the Java API, please see the online (HTML) *Routing Service* documentation.

The Java Adapter API defines the interfaces in [Table 8.10](#).

Table 8.10 **Java Adapter API Interfaces**

Interface	Description
<code>com.rti.routing.service.adapter.Adapter</code>	Required The Adapter interface defines methods to: <ul style="list-style-type: none"> get the adapter version create/destroy connections
<code>com.rti.routing.service.adapter.Connection</code>	Required The Connection interface defines methods to: <ul style="list-style-type: none"> create/destroy Sessions create/destroy StreamReaders create/destroy StreamWriters update the Connection configuration
<code>com.rti.routing.service.adapter.DiscoveryConnection</code>	The DiscoveryConnection interface defines methods to: <ul style="list-style-type: none"> get the discovery StreamReaders (see Section 8.2.2) copy/delete TypeRepresentations
<code>com.rti.routing.service.adapter.Session</code>	Required The Session interface defines methods to update the Session configuration

Table 8.10 Java Adapter API Interfaces

Interface	Description
<code>com.rti.routing.service.adapter.StreamReader</code>	<p>Required for input adapters</p> <p>The <code>StreamReader</code> interface defines methods to:</p> <ul style="list-style-type: none"> • read samples • return the loan on the read samples • update the <code>StreamReader</code> configuration
<code>com.rti.routing.service.adapter.StreamWriter</code>	<p>Required for output adapters</p> <p>The <code>StreamWriter</code> interface defines methods to:</p> <ul style="list-style-type: none"> • write samples • update the <code>StreamWriter</code> configuration

8.3.5.1 Adapter Entry Point

Every Java adapter must create an Adapter class that implements the `com.rti.routing.service.adapter.Adapter` interface.

Adapter classes are registered with *Routing Service* using the tag `<class_name>` within `<java_adapter_plugin>` (see [Section 8.1](#)).

When *Routing Service* creates a new adapter object it will look for the following constructor:

```
MyAdapter(java.util.Properties properties)
```

If the constructor does not exist, *Routing Service* will use the default constructor without arguments.

```
MyAdapter()
```

8.3.5.2 Error Notification

Routing Service must be notified about errors in the adapter's logic. To do so, use the following exception: `com.rti.routing.service.adapter.infrastructure.AdapterException`

8.3.5.3 Adapter Verbosity

The property `rti.routing.service.verbosity` provided to the Adapter constructor can be used to get the verbosity level used to run *Routing Service*.

[Table 8.11](#) describes the mapping between the command-line option `-verbosity` and the values of the property `"rti.routing.service.verbosity"`.

Table 8.11 Mapping between `-verbosity` and `rti.routing.service.verbosity`

<code>-verbosity</code>	<code>rti.routing.service.verbosity</code>
0	none
1	exception
2	warn
3 and 4	info
5 and 6	debug

8.3.6 My First Java Adapter

This section shows how to create a simple Java adapter on Windows and UNIX-based architectures. It is not intended to give complete coverage of the entire adapter API, but rather to introduce the adapter technology and provide the basics of the development process of a Java adapter.

The new Adapter will be a simple file adapter where the input adapter reads lines from a text file and the output adapter saves the provided lines to an output text file.

The source code and scripts that you will create in the next sections are provided in `<Routing Service home>/adapters/tutorial/Java`.

8.3.6.1 Setting the Environment on the Development Machine

There are a few things to take care of before you start developing the simple file adapter.

1. Set the environment variable **ROUTINGSERVICEHOME**

Set the environment variable `ROUTINGSERVICEHOME` to the *Routing Service* installation directory. (*Routing Service* itself does not require that you set the environment variable. It is used to build, compile and run the example adapter).

2. **On Windows Systems:** To use a Java adapter, you must have the Visual Studio 2005 service pack 1 redistributable libraries. You can obtain this package from Microsoft or RTI (see the *RTI Core Libraries and Utilities Release Notes*¹ for details).

3. Make sure Java 1.5 or higher is available.

Ensure that appropriate `javac`, `jar` and `jdb` (for debugging) executables are on your path. They can be found in the `bin` directory of your JDK installation.

1. See `<Connex installation directory>/ndds.<version>/doc/pdf/RTI_Connext_ReleaseNotes.pdf`.

4. Make sure you add the directory of the Java Virtual Machine dynamic library to your environment variable: `LD_LIBRARY_PATH` (on UNIX-based systems) or `Path` (on Windows systems). For example:

```
setenv LD_LIBRARY_PATH
${LD_LIBRARY_PATH}:/local/java/jdk1.5.0_07/jre/lib/i386/client
```

8.3.6.2 Creating a Build Script for UNIX-based Systems

In this section, you will create a shell script to compile the Java adapter.

1. Create a directory that will contain the build script and the adapter implementation. The rest of this section assumes that you will use `/opt/adapters/simplefile` as the adapter directory.
2. In `/opt/adapters/simplefile`, create a file called `build.sh` with the following content.

```
#!/bin/sh
#####
## RTI Routing Service File Simple Adapter ##
#####

#####

# Java compiler
JAVAC=javac
JAR=jar
# Path to RTI Routing Service Adapter API
ADAPTER_CLASSPATH="$ROUTINGSERVICEHOME/class/rtirsadapter.jar"
# Path to RTI Connex Java API
DDS_CLASSPATH="$ROUTINGSERVICEHOME/class/nddsjava.jar"

ALL_SRC=`find routingservice/adapter/simplefile -name \*.java`

mkdir -p class

# Builds all files from 'routingservice' to 'class'
echo "Building all the sources in 'rtiroutingservice' into 'class'
directory..."
$JAVAC -d class -sourcepath . -classpath
$ADAPTER_CLASSPATH:$DDS_CLASSPATH $ALL_SRC
$JAR cf class/simplefileadapter.jar -C class routingservice
rm -rf class/routingservice
```

8.3.6.3 Creating a Build Script for Windows Systems

In this section you will create a script to compile the Java adapter.

1. Create a directory that will contain the build script and the adapter implementation. The rest of this section assumes that you will use `c:\adapters\SimpleFileAdapter` as the adapter directory.
2. In `c:\adapters\SimpleFileAdapter`, create a file called `build.cmd` with the following content.

```
@ECHO OFF
REM #####
REM # RTI Routing Service Simple File Adapter #
REM #####

SETLOCAL enabledelayedexpansion

REM Get rid of quotes
SET ROUTINGSERVICEHOME_NQ=%ROUTINGSERVICEHOME:="=%

REM Path to Java
SET JAVAC=javac.exe
SET JAR=jar.exe

REM Path to RTI Routing Service Adapter API
SET ADAPTER_CLASSPATH="%ROUTINGSERVICEHOME_NQ%\class\rtirs-
adapter.jar"

REM Path to RTI Connex Java API
SET DDS_CLASSPATH="%ROUTINGSERVICEHOME_NQ%\class\nddsjava.jar"

REM Ensure the 'objs' directory exists
IF NOT EXIST class (
    MD class
)

ECHO Building all the sources in 'routingervice' into 'class'
directory...
FOR /R routingervice %%F IN (*.java) DO %JAVAC% -d class -source-
path . -classpath "%ADAPTER_CLASSPATH%;%DDS_CLASSPATH%" "%%F"
%JAR% cf class/simplefileadapter.jar -C class routingervice
RD /S /Q class\routingervice
```

8.3.6.4 Implementing the Adapter Class

In this section you will create the adapter class for the simple file adapter.

Every Java adapter has to create an Adapter class that implements the `com.rti.routing-service.adapter.Adapter` interface.

Adapter classes are registered with *Routing Service* using the tag `<class_name>` within `<java_adapter_plugin>` (Section 8.1).

Using your favorite Java editor, create a file called **SimpleFileAdapter.java** under **<Adapter directory>¹/routing-service/adapter/simplefile**.

Insert the following content:

```

/*****
/*                               Simple File Adapter                               */
/*****

package routing-service.adapter.simplefile;

import com.rti.routing-service.adapter.Adapter;
import com.rti.routing-service.adapter.Connection;
import com.rti.routing-service.adapter.StreamReaderListener;
import com.rti.routing-service.adapter.infrastructure.AdapterException;
import com.rti.routing-service.adapter.infrastructure.Version;
import java.util.Properties;

/**
 * Simple file adapter.
 */
public class SimpleFileAdapter implements Adapter {
    String verbosity;

    /**
     * Entry point to the adapter.
     */
    public SimpleFileAdapter(Properties props) {
        verbosity = props.getProperty("rti.routing-service.verbosity");

        if (verbosity.equals("debug")) {
            System.out.println("CREATE " + getClass().getName());
        }
    }
}

```

1. `c:\adapters\SimpleFileAdapter` for Windows systems, or `/opt/adapters/simplefile` for UNIX-based systems.

```
/**
 */
public Connection createConnection(
    String routingServiceName,
    String routingServiceGroupName,
    StreamReaderListener inputStreamDiscoveryListener,
    StreamReaderListener outputStreamDiscoveryListener,
    Properties properties) throws AdapterException
{
    if (verbosity.equals("debug")) {
        System.out.println("CALL " + getClass().getName() +
            ".createConnection");
    }

    return new SimpleFileAdapterConnection(properties, verbosity);
}

/**
 */
public void deleteConnection(Connection connection)
    throws AdapterException
{
    if (verbosity.equals("debug")) {
        System.out.println("CALL " + getClass().getName() +
            ".deleteConnection");
    }
}

/**
 * Returns the adapter version.
 */
public Version getVersion() {
    return new Version(1,0,0,0);
}
}
```

To create a `SimpleFileAdapter` object, *Routing Service* will use the constructor **SimpleFileAdapter(Properties props)**.

The `props` parameter is used to configure the adapter object. Some of the values can be set from the XML configuration file using the tag `<property>` within `<java_adapter_plugin>` and other values are set by *Routing Service*. One of the predefined values is "rti.routingservice.verbosity". This property provides information about the verbosity level used to run *Routing Service* (see [Section 8.3.5.3](#)).

Adapter objects are factories for Connection objects.

8.3.6.5 Implementing the Connection Class

Connection objects provide access to data domains such as *Connex* domains or JMS network providers and they are configured using the XML tags <connection_1> and <connection_2> in a <domain_route> (see [Section 2.4.2](#)). In the SimpleFileAdapter example, the connection objects will provide access to a directory in your computer's file system.

The next step consists of implementing the Connection Java class.

Create a file called **SimpleFileAdapterConnection.java** under <Adapter directory>¹/routingservice/adapter/simplefile.

Insert the following content:

```

/*****
/*          Simple File Adapter Connection          */
/*****

package routingservice.adapter.simplefile;

import java.util.Properties;
import com.rti.routingservice.adapter.Connection;
import com.rti.routingservice.adapter.Session;
import com.rti.routingservice.adapter.StreamReader;
import com.rti.routingservice.adapter.StreamReaderListener;
import com.rti.routingservice.adapter.StreamWriter;
import com.rti.routingservice.adapter.infrastructure.AdapterException;
import com.rti.routingservice.adapter.infrastructure.StreamInfo;

/**
 * Simple file connection.
 */
public class SimpleFileAdapterConnection implements Connection {
    private String verbosity;
    private String directory = null;

    /**
     */
    SimpleFileAdapterConnection(Properties properties, String verbosity)
        throws AdapterException
    {
        this.verbosity = verbosity;

```

1. c:\adapters\SimpleFileAdapter for Windows systems, or /opt/adapters/simplefile for UNIX-based systems

```
        directory = properties.getProperty("directory");

        if (directory == null) {
            throw new AdapterException(0,
                "directory property is required");
        }
    }

    /**
     */
    public Session createSession(Properties properties)
        throws AdapterException {

        /* We dont need a session for the simple file adapter but
        we cannot return null */
        return new Session() {
            public void update(Properties properties)
                throws AdapterException {
            }
        };
    }

    /**
     */
    public void deleteSession(Session session)
        throws AdapterException {
    }

    /**
     */
    public StreamReader createStreamReader(
        Session session,
        StreamInfo streamInfo,
        Properties properties,
        StreamReaderListener listener) throws AdapterException
    {
        if (verbosity.equals("debug")) {
            System.out.println("CALL " + getClass().getName() +
                ".createStreamReader");
        }

        return new SimpleFileAdapterStreamReader(
            listener, streamInfo,
            properties, directory,
            verbosity);
    }
}
```



```
/**
 */
public void deleteStreamReader(StreamReader streamReader)
    throws AdapterException
{
    if (verbosity.equals("debug")) {
        System.out.println("CALL " + getClass().getName() +
            ".deleteStreamReader");
    }

    ((SimpleFileAdapterStreamReader) streamReader).close();
}

/**
 */
public StreamWriter createStreamWriter(
    Session session,
    StreamInfo streamInfo,
    Properties properties) throws AdapterException
{
    if (verbosity.equals("debug")) {
        System.out.println("CALL " + getClass().getName() +
            ".createStreamWriter");
    }

    return new SimpleFileAdapterStreamWriter(
        streamInfo,
        properties,
        directory,
        verbosity);
}

/**
 */
public void deleteStreamWriter(StreamWriter streamWriter)
    throws AdapterException
{
    if (verbosity.equals("debug")) {
        System.out.println("CALL " + getClass().getName() +
            ".deleteStreamWriter");
    }

    ((SimpleFileAdapterStreamWriter) streamWriter).close();
}
```

```
    /**
    */
    public Properties getAttributes() throws AdapterException {
        throw new AdapterException(0, "operation not supported");
    }

    /**
    */
    public void update(Properties properties) throws AdapterException {
    }
}
```

Connection objects are configurable using properties (name/value pairs). The properties are set using the `<property>` tag within `<connection_x>`. For the `SimpleFileAdapter` example, there is one property called **directory** that is used to specify the directory containing the files to read/write.

For example:

```
<connection_1 plugin_name="adapters::simple_file">
  <registered_type name="TextLine" type_name="TextLine"/>
  <property>
    <value>
      <element>
        <name>directory</name>
        <value>/tmp/SimpleFileAdapter/input</value>
      </element>
    </value>
  </property>
</connection_1>
```

Connection objects are factories for `Session`, `StreamReader` and `StreamWriter` objects. In the next sections you will implement `StreamReader` and `StreamWriters`. `Session` objects are not used in this example.

8.3.6.6 Implementing the `StreamReader` Class

A *StreamReader* provides a way to read data samples of a specific type from a data domain.

In the configuration file, *StreamReaders* are associated with the tag `<input>` within `<route>` or `<auto_route>` (see [Section 2.4.6](#)).

The `StreamReaders` created by the `SimpleFileAdapter` connections read text files from the connection directory.

The data samples provided to *Routing Service* (using the read operation) are `DynamicData` with the following IDL type:

```
struct TextLine {
    string<1024> value;
};
```

When a `SimpleFileAdapter StreamReader` is created, the name of the file is the input stream name with a `.txt` extension. The frequency at which the `StreamReader` notifies *Routing Service* of new lines is configurable using the `read_period` property. For example:

```
<route name="route">
  <input connection="1">
    <stream_name>HelloWorld</stream_name>
    <registered_type_name>TextLine</registered_type_name>
    <property>
      <value>
        <element>
          <name>read_period</name>
          <value>1000</value>
        </element>
      </value>
    </property>
  </input>
  ...
</route>
```

In the previous example, the input `StreamReader` will read the lines of a file called `HelloWorld.txt` and it will provide one line per second to *Routing Service*.

The next step consist on the implementation of the `StreamReader` class. There are three main methods:

❑ `read()`

This method will be called by *Routing Service* after being notified that there are new lines available. Although the signature of the method allows returning more than one sample (line), for the sake of simplicity, the implementation only returns one line every time the method is called.

Routing Service will not call the `read` operation until it is notified of the presence of new data (see [Section 8.2.3](#)). To provide data notification, the `StreamReader` implementation creates a thread (`NotificationThread`) that wakes up after `read_period` and notifies *Routing Service* of new data if the end of the file has not been reached yet.

❑ `return_loan()`

The loan on the samples provided by `read()` is returned to the `StreamReader` using this method. The `SimpleFileAdapter` implementation of `return_loan` is empty because of these reasons:

- The read operation does not create new samples and it always return a single sample stored in the `StreamReader`.
- Two calls to `read()` cannot occur in parallel.

❑ `update()`

The update methods will be called when the `read_period` is changed using remote administration.

Create a file called `SimpleFileAdapterStreamReader.java` under `<Adapter directory>`¹/`routingsservice/adapter/simplefile`.

Insert the following content:

```

/*****
/*          Simple File Adapter Stream Reader          */
*****/

package routingsservice.adapter.simplefile;

import java.io.File;
import java.io.BufferedReader;
import java.io.FileReader;
import java.io.IOException;
import java.util.List;
import java.util.Properties;

import com.rti.dds.dynamicdata.DynamicData;

import com.rti.dds.typecode.TypeCode;
import com.rti.routingsservice.adapter.StreamReader;
import com.rti.routingsservice.adapter.StreamReaderListener;
import com.rti.routingsservice.adapter.infrastructure.AdapterException;
import com.rti.routingsservice.adapter.infrastructure.StreamInfo;

public class SimpleFileAdapterStreamReader implements StreamReader {
    private String verbosity;
    private int readPeriod;
    private String fileName = null;

```

1. `c:\adapters\SimpleFileAdapter` for Windows systems, or `/opt/adapters/simplefile` for UNIX-based systems.

```

private BufferedReader fileReader = null;
private NotificationThread notificationThread = null;
private DynamicData dynamicData = null;

/**
 */
private void parseProperties(Properties properties) {
    String readPeriodStr;

    readPeriodStr = properties.getProperty("read_period");

    if (readPeriodStr == null) {
        readPeriod = 1000;
    } else {
        readPeriod = new Integer(readPeriodStr).intValue();
    }
}

/**
 */
SimpleFileAdapterStreamReader(
    StreamReaderListener listener,
    StreamInfo streamInfo,
    Properties properties,
    String directory,
    String verbosity) throws AdapterException
{
    this.verbosity = verbosity;
    parseProperties(properties);
    fileName = streamInfo.getStreamName() + ".txt";

    try {
        fileReader = new BufferedReader(new FileReader(
            new File(directory, fileName)));
    } catch (IOException e) {
        throw new AdapterException(0, "error opening " + fileName);
    }

    dynamicData = new DynamicData(
        (TypeCode)streamInfo.getTypeInfo().getTypeRepresentation(),
        DynamicData.PROPERTY_DEFAULT);

    notificationThread = new NotificationThread(
        this, listener, fileReader, readPeriod);
    notificationThread.start();
}

```

```
/**
 */
void close() throws AdapterException {
    try {
        notificationThread.terminate();
        notificationThread.join();

        if (fileReader != null) {
            fileReader.close();
        }
    } catch (InterruptedException e) {
        throw new AdapterException(0,
            "error finishing notification thread");
    } catch (IOException e) {
        throw new AdapterException(0, "error closing " + fileName);
    }
}

/**
 */
public void read(List<Object> sampleList, List<Object> infoList)
    throws AdapterException
{
    String line;

    if (verbosity.equals("debug")) {
        System.out.println("CALL " + getClass().getName() + ".read");
    }

    try {
        sampleList.clear();
        infoList.clear();
        dynamicData.clear_all_members();

        if (fileReader.ready()) {
            line = fileReader.readLine();
            dynamicData.set_string("value",
                DynamicData.MEMBER_ID_UNSPECIFIED,
                line);
            sampleList.add(dynamicData);
        }
    } catch (IOException e) {
        throw new AdapterException(0, "error reading from file " +
            fileName, e);
    } catch (Exception e) {
```

```

        throw new AdapterException(0, "error reading", e);
    }
}

/**
 */
public void returnLoan(List<Object> sampleList, List<Object> infoList)
    throws AdapterException
{
    if (verbosity.equals("debug")) {
        System.out.println("CALL " + getClass().getName() +
            ".returnLoan");
    }
}

/**
 */
public void update(Properties properties) throws AdapterException {
    parseProperties(properties);
    notificationThread.setReadPeriod(readPeriod);
}

/**
 * Notification thread
 *
 * This thread will notify of data availability in the file.
 */
class NotificationThread extends Thread {
    private BufferedReader fileReader = null;
    private int notificationPeriod;
    private boolean _terminate;
    private StreamReaderListener listener = null;
    private StreamReader streamReader = null;

    /**
     */
    NotificationThread(StreamReader streamReader,
        StreamReaderListener listener,
        BufferedReader fileReader,
        int notificationPeriod) {
        this.listener = listener;
        this.fileReader = fileReader;
        this.notificationPeriod = notificationPeriod;
        this.streamReader = streamReader;
        _terminate = false;
    }
}

```

```
/**
 */
public void run() {
    while (!_terminate) {
        try {
            Thread.sleep(notificationPeriod);

            if (fileReader.ready()) {
                listener.onDataAvailable(streamReader);
            }
        } catch (Exception e) {}
    }
}

/**
 */
public void terminate() {
    _terminate = true;
}

/**
 */
public void setReadPeriod(int readPeriod) {
    notificationPeriod = readPeriod;
}
}
}
```

8.3.6.7 Implementing the StreamWriter Class

A *StreamWriter* provides a way to write samples of a specific type into a data domain.

In the configuration file, *StreamWriters* are associated with the tag `<output>` within `<route>` or `<auto_route>` (see [Section 2.4.6](#)).

The SimpleFileAdapter *StreamWriters* create new files in the connection directory and store the lines read from the routes' inputs.

The data samples provided to the *StreamWriters*' write operation are *DynamicData* with the following IDL type:

```
struct TextLine {
    string<1024> value;
};
```


When a `SimpleFileAdapter StreamWriter` is created, the name of the file is the output stream name with a `.txt` extension. For debugging purposes, the `StreamWriter` can be configured to print the written samples on the console:

```
<route name="route">
  ...
  <output>
    <stream_name>HelloWorld</stream_name>
    <registered_type_name>TextLine</registered_type_name>
    <property>
      <value>
        <element>
          <name>print_to_stdout</name>
          <value>1</value>
        </element>
      </value>
    </property>
  </output>
</route>
```

In the above example, the output `StreamWriter` will store the lines provided by *Routing Service* on a file called **HelloWorld.txt**. It will also print the lines on the screen.

Insert the following code in the “Simple File Adapter: StreamWriter” section of **SimpleFileAdapter.c**.

Create a file called **SimpleFileAdapterStreamWriter.java** under `<Adapter directory>1/routing-service/adapter/simplefile`.

Insert the following content:

```

/*****
/*                               Simple File Adapter Stream Writer                               */
*****/

package routing-service.adapter.simplefile;

import java.io.File;
import java.io.FileWriter;
import java.io.BufferedWriter;
import java.io.IOException;
import java.util.List;
import java.util.ListIterator;
import java.util.Properties;
```

1. `c:\adapters\SimpleFileAdapter` for Windows systems, or `/opt/adapters/simplefile` for UNIX-based systems.

```
import com.rti.dds.dynamicdata.DynamicData;
import com.rti.routingservice.adapter.StreamWriter;
import com.rti.routingservice.adapter.infrastructure.AdapterException;
import com.rti.routingservice.adapter.infrastructure.StreamInfo;

public class SimpleFileAdapterStreamWriter implements StreamWriter {
    private String verbosity = null;
    private String fileName = null;
    private boolean printToStdout;
    private BufferedWriter fileWriter = null;

    /**
     */
    private void parseProperties(Properties properties) {
        int printToStdoutInt;
        String printToStdoutStr;

        printToStdoutStr = properties.getProperty("print_to_stdout");

        if (printToStdoutStr == null) {
            printToStdout = false;
        } else {
            printToStdoutInt = new Integer(printToStdoutStr).intValue();

            if (printToStdoutInt != 0) {
                printToStdout = true;
            } else {
                printToStdout = false;
            }
        }
    }

    /**
     */
    SimpleFileAdapterStreamWriter(
        StreamInfo streamInfo,
        Properties properties,
        String directory,
        String verbosity) throws AdapterException
    {
        this.verbosity = verbosity;
        parseProperties(properties);

        fileName = streamInfo.getStreamName() + ".txt";
    }
}
```

```

    try {
        fileWriter = new BufferedWriter(new FileWriter(new File(
            directory, fileName));
    } catch (IOException e) {
        throw new AdapterException(0, "error opening " + fileName);
    }
}

/**
 */
void close() throws AdapterException {
    try {
        if (fileWriter != null) {
            fileWriter.close();
        }
    } catch (IOException e) {
        throw new AdapterException(0, "error closing " + fileName);
    }
}

/**
 */
public int write(List<Object> sampleList, List<Object> infoList)
    throws AdapterException
{
    String line;
    ListIterator iterator = sampleList.listIterator();
    DynamicData dynamicData = null;

    if (verbosity.equals("debug")) {
        System.out.println("CALL " + getClass().getName() + ".write");
    }

    try {
        while (iterator.hasNext()) {
            dynamicData = (DynamicData)iterator.next();
            line = dynamicData.get_string("value",
                DynamicData.MEMBER_ID_UNSPECIFIED);
            fileWriter.write(line);
            fileWriter.newLine();

            if (printToStdout) {
                System.out.println(line);
            }
        }
    }
}

```

```
        } catch (IOException e) {
            throw new AdapterException(0, "error writing to file " +
                fileName, e);
        } catch (Exception e) {
            throw new AdapterException(0, "error writing", e);
        }
    }

    return 0;
}

/**
 * /
public void update(Properties properties) throws AdapterException {
    parseProperties(properties);
}
}
```

8.3.6.8 Running the SimpleFileAdapter

This section describes the steps required to use and run the SimpleFileAdapter with *Routing Service*. You will create a configuration file with a single route that reads a Hello-World text file from an input directory and saves it into an output directory.

1. Compile and build the SimpleFileAdapter.

UNIX-based systems:

```
> cd /opt/adapters/simplefile
> ./build.sh
```

Windows systems:

```
> cd c:\adapters\SimpleFileAdapter
> build.cmd
```

2. In the adapter project directory (**c:\adapters\SimpleFileAdapter¹** on Windows systems; **/opt/adapters/simplefile¹** on UNIX-based systems), create two directories called **input** and **output**.
3. In the input directory create a file called **HelloWorld.txt** with the following content.

```
Hello World 1!
Hello World 2!
Hello World 3!
Hello World 4!
Hello World 5!
Hello World 6!
```

1. Your directory may be different if you did not use the default locations.

```

Hello World 7!
Hello World 8!
Hello World 9!
Hello World 10!

```

4. In the adapter project directory, create a *Routing Service* XML configuration file called **simple_file_adapter.xml** with the following content. Replace the value of the **directory** property under both connections with the location of the input and output directories.

```

<?xml version="1.0"?>
<dds>
  <adapter_library name="adapters">
    <java_adapter_plugin name="simple_file">
      <class_name>routing.service.adapter.simplefile.SimpleFileAdapter</class_name>
    </java_adapter_plugin>
  </adapter_library>

  <types>
    <struct name="TextLine">
      <member name="value" type="string" stringMaxLength="2048"/>
    </struct>
  </types>

  <routing_service name="file_to_file">
    <jvm>
      <class_path>
        <element>./class/simplefileadapter.jar</element>
      </class_path>
    </jvm>

    <domain_route name="domain_route">
      <connection_1 plugin_name="adapters::simple_file">
        <registered_type name="TextLine" type_name="TextLine"/>
        <property>
          <value>
            <element>
              <name>directory</name>
              <value>/opt/adapters/simplefile/input</value>
            </element>
          </value>
        </property>
      </connection_1>
    </domain_route>
  </routing_service>

```

```
<connection_2 plugin_name="adapters::simple_file">
  <registered_type name="TextLine" type_name="TextLine"/>
  <property>
    <value>
      <element>
        <name>directory</name>
        <value>/opt/adapters/simplefile/output</value>
      </element>
    </value>
  </property>
</connection_2>

<session name="session">
  <route name="route">
    <input connection="1">
      <stream_name>HelloWorld</stream_name>
      <registered_type_name>
        TextLine
      </registered_type_name>
    </input>
    <output>
      <stream_name>HelloWorld</stream_name>
      <registered_type_name>
        TextLine
      </registered_type_name>
      <property>
        <value>
          <element>
            <name>print_to_stdout</name>
            <value>1</value>
          </element>
        </value>
      </property>
    </output>
  </route>
</session>
</domain_route>
</routing_service>
</dds>
```

5. Start *Routing Service* by entering the following in a command shell.

On UNIX-based systems:

```
> cd <SimpleFileAdapter project directory>
> $ROUTINGSERVICEHOME/scripts/rtiroutingservice
   -cfgFile simple_file_adapter.xml -cfgName file_to_file
```

On Windows systems:

```
> cd <SimpleFileAdapter project directory>
> %ROUTINGSERVICEHOME%\scripts\rtiroutingservice
   -cfgFile simple_file_adapter.xml -cfgName file_to_file
```

6. On the screen you will see:

```
RTI Routing Service <version> started (with name file_to_file)
Hello World 1!
Hello World 2!
Hello World 3!
Hello World 4!
Hello World 5!
Hello World 6!
Hello World 7!
Hello World 8!
Hello World 9!
Hello World 10!
```

7. Verify that a file called **HelloWorld.txt** has been generated into the output directory. The content of this file should be identical to the content of the same file in the input directory.

8.3.7 Debugging Java Adapters

When you develop a custom adapter, you will need to debug it and test it. This section describes the tools and APIs that you have available to debug and detect problems in *Routing Service* adapters written in Java.

The first debugging capability is provided by the *Routing Service Adapter SDK*. The adapter SDK provides a way to access the verbosity level of *Routing Service* as a property called `rti.routing.service.verbosity`, which can be obtained from the properties passed to the adapter constructor. It is highly recommended that, as part of the adapter implementation, you instrument the code by adding status messages that will be printed with the INFO and DEBUG verbosity levels. This level of instrumentation will help you to capture run-time information for troubleshooting.

The second debugging capability is provided by third party tools. The rest of this section shows how to debug a Java adapter using **jdb** (the command-line debugger of Java) and **NetBeans™ IDE 6.9**.

8.3.7.1 Enabling Debugging in the Routing Service JVM

Before you start debugging with **jdb** or NetBeans, you have to enable debugging in the JVM-created *Routing Service*.

1. If you have not done so already, stop the existing *Routing Service* execution by pressing CTRL-C.
2. Edit `java_simple_adapter.xml` and replace the content of the JVM tag with:

```
<jvm>
  <class_path>
    <element>./class/simplefileadapter.jar</element>
  </class_path>
  <options>
    <element>-Xdebug</element>
    <element>
      -Xrunjwdp:transport=dt_socket,address=8192,server=y,suspend=y
    </element>
  </options>
</jvm>
```

The JVM option `-Xdebug` is used to enable debugging.

The JVM option `-Xrunjwdp` loads the JDBA reference implementation of JDWP (Java Debug Wire Protocol) and starts listening on port 8192 to communicate with a separate debugger application such as `jdb` and NetBeans.

For additional details on Java debugging see:

<http://java.sun.com/javase/technologies/core/toolsapis/jpda>

3. Save the changes.
4. Run *Routing Service*.

On UNIX-based systems:

```
> cd <SimpleFileAdapter project directory>
> $ROUTINGSERVICEHOME/scripts/rtiroutingservice
   -cfgFile simple_file_adapter.xml -cfgName file_to_file
```

On Windows systems:

```
> cd <SimpleFileAdapter project directory>
> %ROUTINGSERVICEHOME%\scripts\rtiroutingservice
   -cfgFile simple_file_adapter.xml -cfgName file_to_file
```

You should see output like this:

```
Listening for transport dt_socket at address: 1024
```

At this point, the execution of *Routing Service* is suspended and waiting for a debugger to attach.

8.3.7.2 Debugging with JDB

jdb is the command-line debugger of Java. This section is not intended to give complete coverage on all the **jdb** functionality and commands, but rather to provide basic information on how to attach to the *Routing Service* JVM and start debugging.

For more information about JDB see the following web page:

<http://java.sun.com/j2se/1.5.0/docs/tooldocs/windows/jdb.html>

1. In a separate shell window, start **jdb**:

UNIX-based systems:

```
> cd /opt/adapters/simplefile
> jdb -sourcepath . -attach 1024
```

Windows systems:

```
> cd c:\adapters\SimpleFileAdapter
> jdb -sourcepath . -attach 1024
```

2. Set breakpoints in the methods or classes that you would like to debug.

For example, to set a breakpoint in the `SimpleFileAdapter` constructor enter the following:

```
main[1] stop in routingservice.adapter.simplefile.SimpleFileAdapter.<init>
```

3. Resume the execution of *Routing Service* by entering:

```
main[1] cont
```

You will see output similar to:

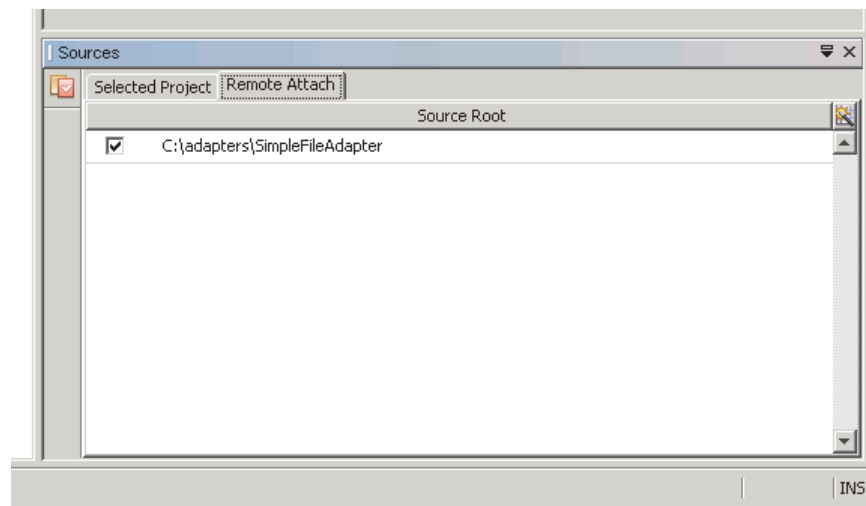
```
> Set deferred breakpoint routingservice.adapter.simplefile.SimpleFileAdapter.<init>
Breakpoint hit: "thread=main", routingservice.adapter.simplefile.SimpleFileAdapter.<init>(), line=23 bci=0
23     public SimpleFileAdapter(Properties props) {
```

Use the command `help` to get a list of the command that will allow you to continue the debugging process.

8.3.7.3 Debugging with NetBeans

NetBeans is an IDE for developing and debugging Java applications. This section is not intended to give complete coverage of all the NetBeans debugger functionality, but rather to provide basic information on how to attach the NetBeans debugger to the *Routing Service* JVM and start debugging.

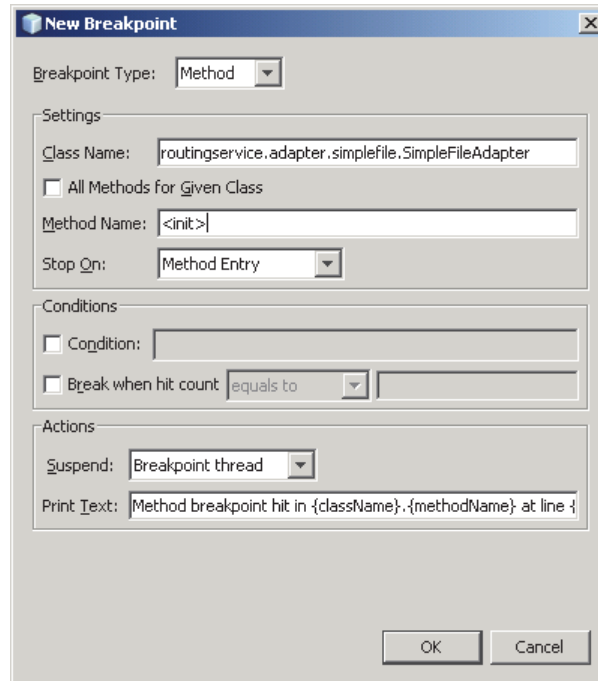
1. Verify that NetBeans IDE 6.9 is installed on your system. The installation of NetBeans is beyond the scope of this document; please refer to NetBeans documentation.
2. Start NetBeans.
3. Make the adapter source code available to the debugger.
 - a. Select **Window, Debugging, Sources**.
 - b. Right-click the **Remote Attach** window and select **Add Source Root**.
 - c. Enter the adapter directory.



- Set breakpoints in the methods or classes that you would like to debug.

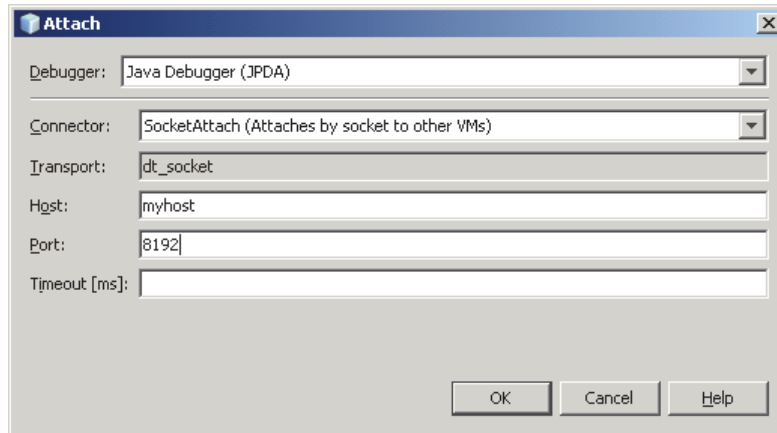
For example, to set a breakpoint in the SimpleFileAdapter constructor, follow the following steps:

- Select **Debug, New Breakpoint**.
- In the **New Breakpoint** window, select **Method** as the breakpoint type and provide **routing.service.adapter.simplefile.SimpleFileAdapter** as the class name and **<init>** as the method name.



- Press **OK**.

5. Attach the debugger to *Routing Service* JVM.
 - a. Select **Debug, Attach Debugger**.
 - b. For the **Host**, enter the name of the host where *Routing Service* is running.
 - c. For the **Port**, enter **8192**.



- d. Press **OK** to start debugging the adapter.

8.3.8 Testing an Adapter

A simple Java test adapter is provided with *Routing Service Adapter SDK*. You will find the class, `com.rti.routing.service.adapter.test.TestAdapter`, in `rtirsadapter.jar`.

This is a convenient way to test your own adapters. The `TestAdapter` is used as an output adapter that counts the number of samples that meet certain conditions defined in the configuration file.

Your adapter will act as the input and its samples will be passed to the `TestAdapter`. If the number of samples received by the `TestAdapter` is not between a defined range when you stop *Routing Service*, you will see a failure message. (Success or failure is determined when you stop *Routing Service* and it destroys the adapter.)

To use the `TestAdapter` to test your input adapter:

1. Write a configuration file in which your adapter is the input for one or more routes and the `TestAdapter` is the output.

Configure the `TestAdapter` with the expected number of samples within a range specified using the properties `MinExpectedSamples` and `MaxExpectedSamples` in the `<output>` tag.

2. Run *Routing Service* using that configuration file.
3. Wait the amount of time your adapter may require.
4. Stop *Routing Service*. The TestAdapter will print a failure or success message.

You can avoid steps 3-4 by starting *Routing Service* with the **-stopAfter <seconds>** command-line option.

If you run *Routing Service* with **-verbosity 3** (or higher), the TestAdapter will also print the DynamicData samples as they arrive.

For an example of how to use and configure the TestAdapter, see *<Routing Service home>/example/testing/test_adapter.xml*. This example tests the simple C file adapter introduced in previous sections.

You can also write your own adapter to extend the TestAdapter class. The source code is in **rtirsadapter.jar**.

