

RTI Code Generator

User's Manual

Version 3.1.1



© 2022 Real-Time Innovations, Inc.
All rights reserved.
Printed in U.S.A. First printing.
March 2022.

Trademarks

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

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.

This is an independent publication and is neither affiliated with, nor authorized, sponsored, or approved by, Microsoft Corporation.

The security features of this product include software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>). This product includes cryptographic software written by Eric Young (eay@cryptsoft.com). This product includes software written by Tim Hudson (tjh@cryptsoft.com).

Notice

Any deprecations noted in this document serve as notice under the Real-Time Innovations, Inc. Maintenance Policy #4220.

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

Chapter 1 Introduction	1
Chapter 2 Paths Mentioned in Documentation	2
Chapter 3 Command-Line Arguments for rtiddsgen	4
Chapter 4 Generating Example Code	
4.1 is_default_qos (true vs. false)	14
4.2 Listeners vs. WaitSets	15
Chapter 5 Generated Files	16
Chapter 6 Customizing the Generated Code	21
Chapter 7 Optimizing the Code Generation Process	
7.1 Optimization Levels	24
7.2 How the Optimizations are Applied	25
7.2.1 Inline expansion of nested types	25
7.2.2 Serialization of consecutive members with a single copy	26
7.2.3 Rules for Inline Expansion	26
7.2.4 Guidelines	29
Chapter 8 Boosting Performance	
8.1 Using Server Mode	30
8.2 Using JVM Optimization	32

Chapter 1 Introduction

RTI® *Code Generator* creates the code needed to define and register a user data type with *RTI Connex*t® *DDS*.

Using *Code Generator* is optional if:

- You are using dynamic types (see *Managing Memory for Built-in Types* (Section 3.2.7) in the [RTI Connex](#)t *DDS Core Libraries Users Manual).*
- You are using one of the built-in types (see *Built-in Data Types* (Section 3.2) in the [RTI Connex](#)t *DDS Core Libraries Users Manual).*

To use *Code Generator*, you will need to provide a description of your data type(s) in an IDL or XML file. You can define multiple data types in the same type-definition file. For details on these files, see the [RTI Connex](#)t *DDS Core Libraries User's Manual* (Sections 3.3 and 3.4).

Chapter 2 Paths Mentioned in Documentation

The documentation refers to:

- **<NDDSHOME>**

This refers to the installation directory for *RTI® Connex® DDS*. The default installation paths are:

- macOS® systems:
/Applications/rti_connex_dds-6.1.1
- Linux systems, non-*root* user:
/home/<your user name>/rti_connex_dds-6.1.1
- Linux systems, *root* user:
/opt/rti_connex_dds-6.1.1
- Windows® systems, user without Administrator privileges:
<your home directory>\rti_connex_dds-6.1.1
- Windows systems, user with Administrator privileges:
C:\Program Files\rti_connex_dds-6.1.1

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

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

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

```
"C:\Program Files\rti_connex_tdds-6.1.1\bin\rtiddsgen"
```

Or if you have defined the **NDDSHOME** environment variable:

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

- *<path to examples>*

By default, examples are copied into your home directory the first time you run *RTI Launcher* or any script in **<NDDSHOME>/bin**. This document refers to the location of the copied examples as *<path to examples>*.

Wherever you see *<path to examples>*, replace it with the appropriate path.

Default path to the examples:

- macOS systems: **/Users/<your user name>/rti_workspace/6.1.1/examples**
- Linux systems: **/home/<your user name>/rti_workspace/6.1.1/examples**
- Windows systems: **<your Windows documents folder>\rti_workspace\6.1.1\examples**

Where 'your Windows documents folder' depends on your version of Windows. For example, on Windows 10, the folder is **C:\Users\<your user name>\Documents**.

Note: You can specify a different location for **rti_workspace**. You can also specify that you do not want the examples copied to the workspace. For details, see *Controlling Location for RTI Workspace and Copying of Examples* in the *RTI Connex DDS Installation Guide*.

Chapter 3 Command-Line Arguments for `rtiddsgen`

On Windows systems: Before running `rtiddsgen`, run the "vcvars" batch file that is appropriate for your architecture in the same command prompt that you will use to run `rtiddsgen`. For example: `vcvarsall.bat x86`. The location of the "vcvars" batch file varies depending on your version of Visual Studio. Consult the Microsoft Visual Studio documentation for the exact location and parameters for your installation. Alternatively, run `rtiddsgen` from the Visual Studio Command Prompt under the Visual Studio Tools folder. Or, use the command-line argument `-ppDisable`. See `-ppDisable` in [Table 3.1 Options for `rtiddsgen` on page 6](#).

If you are generating code for *Connex DDS*, the options are:

```
rtiddsgen    [-help]
             [-allocateWithMalloc]
             [-alwaysUseStdVector]
             [-autoGenFiles <architecture>]
             [-constructor]
             [-create <typefiles| examplefiles|makefiles>]
             [-convertToIdl | -convertToXML | -convertToXsd]
             [-D <name>[=<value>]]
             [-d <outdir>]
             [-disableXSDValidation]
             [-dllExportMacroSuffix <suffix>]
             [-dotnet <modern|legacy>]
             [-enableEscapeChar]
             [-example <architecture>]
             [-exampleTemplate]
             [-I <directory>]
             [[-inputIdl <IDLInputFile.idl> | [-inputXml <XMLInputFile.xml>
             | [-inputXsd <IDLInputFile.idl>]]]
             [-language <Ada|C|C++|C++11|C++/CLI|C#|Java>]
             [-namespace]
             [-obfuscate]
             [-optimization <level>]
             [-package <packagePrefix>]
             [-platform <architecture>]
```

```

[-ppDisable]
[-ppPath <path to preprocessor>]
[-ppOption <option>]
[-qualifiedEnumerator]
[-reader]
[-replace]
[-sequenceSize <unbounded sequences size>]
[-sharedLib]
[-showTemplates]
[-strict]
[-stringSize <unbounded strings size>]
[-typeSequenceSuffix <suffix>]
[-U <name>]
[-unboundedSupport]
[-update <typefiles| examplefiles|makefiles>]
[-use52Keyhash]
[-use526Keyhash]
[-useStdString]
[-V <name< [=<value>]]]
[-verbosity [1-3]]
[-version]
[-virtualDestructor]
[-writer]

```

If you are generating code for *RTI Connex DDS Micro*, the options are:

```

rtiddsgen [-help]
[-create <typefiles| examplefiles|makefiles>]
[-convertToIdl | -convertToXML]
[-D <name>[=<value>]]
[-d <outdir>]
[-enableEscapeChar]
[-example]
[-exampleTemplate]
[-I <directory>]
[[[-inputIdl] <IDLInputFile.idl> | [-inputXml] <XMLInputFile.xml>]
[-language <C|C++>]
[-micro]
[-namespace]
[-ppDisable]
[-ppPath <path to preprocessor>]
[-ppOption <option>]
[-reader]
[-replace]
[-showTemplates]
[-sequenceSize <unbounded sequences size>]
[-strict]
[-stringSize <unbounded strings size>]
[-U <name>]
[-update <typefiles| examplefiles|makefiles>]
[-V <name< [=<value>]]]
[-verbosity [1-3]]
[-version]
[-writer]

```


Table 3.1 Options for `rtiddsgen` describes the options.

Note: Before using a makefile created by *Code Generator* to compile an application, make sure the `${NDDSHOME}` environment variable is set as described in Set Up Environment Variables (`rtisetenv`), in "Hands-On 1" of *Introduction to Publish/Subscribe*, in the [RTI Connext DDS Getting Started Guide](#).

Table 3.1 Options for `rtiddsgen`

Option	Description
<code>-allocateWithMalloc</code>	Use this flag to obtain backward-compatibility when allocating optional members with <code>DDS_Heap_malloc</code> in C++.
<code>-alwaysUseStdVector</code>	Only applies if <code>-language C++11</code> is specified. Generates code that maps all sequences to <code>std::vector</code> , even bounded sequences that would otherwise map to <code>rti::core::bounded_sequence</code> . Alternatively, the <code>@use_vector</code> annotation can be applied to each sequence member.
<code>-autoGenFiles <architecture></code>	Updates the auto-generated files, i.e, the typefiles and makefile/project files. To see the valid options for <code><architecture></code> per language, run <code>rtiddsgen</code> with the <code>-help</code> option, or use the string "universal" (<code>-autoGenFiles universal</code>) to generate compatible publisher/subscriber code for all supported platforms. The universal architecture will not generate makefiles/project files. This is a shortcut for: <code>-update typefiles -update makefiles -platform <architecture></code>
<code>-constructor</code>	Only applies if <code>-language C++</code> is also specified. Generates the types default constructor, copy constructor, copy assignment operator, and destructor. Using this option will also disable the generation of the following TypeSupport methods: <code>create_data(_ex)</code> , <code>delete_data(_ex)</code> , <code>initialize_data(_ex)</code> , <code>finalize_data(_ex)</code> , <code>copy_data</code> .
<code>-create <typefiles examplefiles makefiles></code>	Creates the files indicated (typefiles, examplefiles, or makefiles) if they do not exist. For example: <code>rtiddsgen -language C++11 -create typefiles test.idl</code> If the files already exist, the files are not modified and a warning is printed. There can be multiple <code>-create</code> options. If you specify both <code>-create</code> and <code>-update</code> for the same file type, only <code>-update</code> will be applied. If you use <code>-create makefiles</code> , the <code>-platform <arch></code> option is required. For example: <code>rtiddsgen -language c -create makefiles -platform x64Darwin17clang9.0 test.idl</code>
<code>-convertToIdl</code>	Converts the input type description file into IDL format. This option creates a new file with the same name as the input file and a <code>.idl</code> extension.
<code>-convertToXML</code>	Converts the input type description file into XML format. This option creates a new file with the same name as the input file and a <code>.xml</code> extension.
<code>-convertToXsd</code>	Converts the input type description file into XSD format. This option creates a new file with the same name as the input file and a <code>.xsd</code> extension.
<code>-D <name>[=<value>]</code>	Defines preprocessor macros. On Windows systems, enclose the argument in quotation marks: <code>-D "<name>[=<value>]"</code>

Table 3.1 Options for rtiddsgen

Option	Description
<code>-d <outdir></code>	Generates the output in the specified directory. By default, <i>Code Generator</i> will generate files in the directory where the input type-definition file is found.
<code>-disableXSDValidation</code>	Causes <i>Code Generator</i> not to check that the input XSD file is well-formed. The use of <i>this option is not recommended</i> in general, as <i>Code Generator</i> may receive invalid inputs.
<code>-dllExportMacroSuffix <suffix></code>	Defines the suffix of the macro that is used to export symbols when building Windows DLLs. The default macro is <code>NDDS_USER_DLL_EXPORT</code> . When this option is specified, the name of the macro is <code>NDDS_USER_DLL_EXPORT_<Suffix></code> .
<code>-dotnet <modern legacy></code>	<p>Defines the C# API version to use.</p> <ul style="list-style-type: none"> • -dotnet legacy (used with -language C#) generates code for the legacy C# language, which was used in 6.1.0 and earlier releases. (By default, when you specify -language C#, <i>Code Generator</i> generates code for the new C# API introduced in release 6.1.0.) • -dotnet modern (used with -language C#) uses the new C# API. You don't need to specify it, because -language C# uses the new C# API already, but you can use this option to explicitly specify the new C# API in your build system if desired. <p>Starting in release 6.1.1, the legacy C# language is deprecated, and the -dotnet parameter will be removed in a future release. Using -dotnet legacy will give a warning that the old binding will be removed in a future release; however, <i>Code Generator</i> will still continue generating code for the legacy API.</p>
<code>-enableEscapeChar</code>	Enables use of the escape character '_' in IDL identifiers.
<p>For <i>Connex DDS Professional</i>:</p> <p><code>-example <architecture></code></p> <p>For <i>Connex DDS Micro</i>:</p> <p><code>-example</code></p>	<p>Generates type files, example files, and a makefile.</p> <p>This is a shortcut for:</p> <pre><code>-create typefiles -create examplefiles -create makefiles -platform <architecture></code></pre> <p>To see the valid options for <code><architecture></code> per language, run <i>rtiddsgen</i> with the -help option, or use the string "universal" (-example universal) to generate compatible publisher/subscriber code for all supported platforms. The universal architecture will not generate makefiles/project files.</p> <p>For <i>Connex DDS Micro</i>, you do not need to specify an <code><architecture></code>. <i>rtiddsgen</i> for <i>Connex DDS Micro</i> generates CMakefiles, which can be used to compile for any architecture.</p>

Table 3.1 Options for rtiddsgen


Option	Description
-exampleTemplate	<p>For <i>Connexx DDS Professional</i>:</p> <p>Generates a custom example application using the specified custom publisher and subscriber templates, instead of the default example files. Place your custom templates in <code>\$NDDSHOME/resource/app/app_support/rtiddsgen/templates/example/<language>/<exampleTemplateDirectoryName>/</code>. Name them publisher.vm and subscriber.vm; you can also add a QoS template (<code>qosProfile.vm</code>) file to your example template directory (if you don't, <i>Code Generator</i> will use the default one, <code>\$NDDSHOME/resource/app/app_support/rtiddsgen/templates/example/qosProfile.vm</code>). You can base your templates on those that <i>Connexx DDS</i> provides in <code>\$NDDSHOME/resource/app/app_support/rtiddsgen/templates/example/<language></code>.</p>  <p>← Create your own templates if desired</p> <p>← Provided by Connexx DDS</p> <p>The -exampleTemplate option must be used in combination with one of the following command-line options:</p> <ul style="list-style-type: none"> • -create examplefiles • -update examplefiles • -example <architecture> <p>For example:</p> <pre>rtiddsgen -language C++11 -example x64Darwin17clang9.0 -exampleTemplate myCustomTemplate foo.idl</pre> <p>"myCustomTemplate" is your <exampleTemplateDirectoryName> in <code>\$NDDSHOME/resource/app/app_support/rtiddsgen/templates/example/<language>/<exampleTemplateDirectoryName>/</code>.</p> <p>To get a full list of the available custom templates in your <i>Connexx DDS</i> installation, use the -showTemplates option.</p> <p>The -exampleTemplate option is not supported in Ada.</p>

Table 3.1 Options for rtiddsgen

Option	Description
-exampleTemplate	<p>For <i>Connex DDS Micro</i>, -exampleTemplate behaves the same as it does for <i>Connex DDS Professional</i>, except for the following differences:</p> <p>Place your custom templates in <code>\$RTMEHOME/rtiddsgen/resource/rtiddsgen/templates/example/<language>/<exampleTemplateDirectoryName>/</code>. Name them publisher.vm and subscriber.vm. The following files are also in the <code>\$RTMEHOME/rtiddsgen/resource/rtiddsgen/templates/example/<language></code> directory: application.vm, applicationHeader.vm, and readme.vm.</p> <p>Template files mag.vm and mag_qos.vm are optional. If present in the directory, they can be used to generate <i>RTI Micro Application Generator</i> (MAG) applications. See <i>Connex DDS Micro Application Generation</i> in the RTI Connex DDS XML-Based Application Creation Getting Started Guide.</p> <p>For example:</p> <pre>rtiddsgen -micro -language C -example -exampleTemplate <exampleTemplateName> foo.idl</pre> <p>qosProfile.vm is not supported by <i>Connex DDS Micro</i>.</p>
-help	Prints out the command-line options for <i>rtiddsgen</i> .
-l <directory>	Adds to the list of directories to be searched for type-definition files (IDL or XML files). Note: A type-definition file in one format cannot include a file in another format.
-inputIdl	Indicates that the input file is an IDL file, regardless of the file extension.
-inputXml	Indicates that the input file is an XML file, regardless of the file extension.
-inputXsd	Indicates that the input file is an XSD file, regardless of the file extension.
IDLInputFile.idl	A file containing IDL descriptions of your data types. If -inputIdl is not used, the file must have a .idl extension.
<p>For <i>Connex DDS Professional</i>:</p> <p>-language <Ada C C++ C++11 C++/CLI C# Java></p> <p>For <i>Connex DDS Micro</i>:</p> <p>-language <C C++></p>	<p>Specifies the language to use for the generated files. The default language is C++.</p> <p>Notes:</p> <ul style="list-style-type: none"> The C++03 -language option (for <i>Connex DDS Professional</i>) has been DEPRECATED. Release 6.1.0 is the last release that supports it. For the Modern C++ API, use -language C++11; for the Traditional C++ API, use -language C++. The C# -language will generate code using the new C# API. To use the legacy version used by default in release 6.1.0 and earlier, use -language C# -dotnet legacy. (The C++/CLI -language option will give a warning that the old binding will be removed in a future release. <i>Code Generator</i> will still continue generating code for the old API, but C++/CLI is not supported in the new API.)
-micro	Generates code and support files for <i>Connex DDS Micro</i> , instead of for <i>Connex DDS Professional</i> . Use -micro -help to list the options supported by <i>Code Generator</i> when targeting <i>RTI Connex DDS Micro</i> .
-namespace	Specifies the use of C++ namespaces, for traditional C++ only. For modern C++, C++/CLI, and C#, it is implied—namespaces are always used.
-obfuscate	Generates an obfuscated IDL file from the input file. Note that even if the input type is XML, this option generates an obfuscated IDL file.

Table 3.1 Options for rtiddsngen

Option	Description
<code>-optimization <level></code>	<p>Level of optimization of the code:</p> <ul style="list-style-type: none"> • 0: No optimization. • 1: The compiler generates extra code for typedefs but optimizes its use. If a type that is used is a typedef that can be resolved to a primitive, enum, or aggregated type (struct, union, or value type), the generated code will invoke the code of the most basic type to which the typedef can be resolved. This level can be used if the generated code is not expected to be modified. This is the only optimization level supported for Java, C#, and C++/CLI languages. • 2: (Default) This optimization level applies only to C, C++, C++11 and higher, microC, microC++, and Ada languages. With this optimization level, <i>rtiddsngen</i> optimizes the serialization/deserialization of structures and valuetypes by using more aggressive techniques. These techniques include inline expansion of nested types and serialization/deserialization of a set of consecutive members with a single copy function invocation (memcpy) when the memory layout (C, C++ structure layout) is the same as the wire layout (XCDR). See Chapter 7 Optimizing the Code Generation Process on page 24 for more information. <p>2 is the default for C, C++, C++11 and higher, microC, microC++, and Ada languages (but you can change it to 0 or 1). 1 is always used for Java, C#, and C++/CLI languages, and you cannot change it.</p>
<code>-package <packagePrefix></code>	<p>Specifies the root package into which generated classes will be placed. It applies to Java only. If the type-definition file contains module declarations, those modules will be considered subpackages of the package specified here.</p>
<code>-platform <architecture></code>	<p>Required if <code>-create makefiles</code> or <code>-update makefiles</code> is used.</p> <p>To see the valid options for <code><architecture></code> per language, run <i>rtiddsngen</i> with the <code>-help</code> option, or use the string "universal" (<code>-platform universal</code>) to generate compatible publisher/subscriber code for all supported platforms. The universal architecture will not generate makefiles/project files.</p>
<code>-ppDisable</code>	<p>Disables the preprocessor.</p> <p><i>Code Generator</i> supports the standard preprocessor directives defined by the IDL specification, such as <code>#if</code>, <code>#endif</code>, <code>#include</code>, and <code>#define</code>.</p> <p>To support these directives, <i>Code Generator</i> calls an external C preprocessor before parsing the IDL file. On Windows systems, the preprocessor is 'cl.exe.' On other architectures, the preprocessor is 'cpp.' You can change the default preprocessor with the <code>-ppPath</code> option. If you do not want to run the preprocessor, use the <code>-ppDisable</code> option.</p> <p>The argument <code>-ppDisable</code> tells the code generator not to attempt to invoke the C preprocessor on the IDL file prior to generating code. In some cases, the IDL file contains no preprocessor directives, so no preprocessing is necessary. Therefore, you may want to use the <code>-ppDisable</code> option to tell <i>Code Generator</i> to skip calling the preprocessor. However, if the preprocessor executable is already on your system path (on Windows systems, it will be on your path if you've run the Visual Studio "vcvars" batch file that is appropriate for your architecture), you don't need to use <code>-ppDisable</code>.</p> <p>In summary, <code>-ppDisable</code> is a useful, quick remedy for simple examples (including HelloWorld examples) for which the IDL is simple and does not need <i>Code Generator</i> to invoke the preprocessor. The longterm recommendation, however, is to have the preprocessor in your path (or use the <code>-ppOption</code>). Linux architectures often already have the preprocessor in the path; Windows architectures may require you to run a "vcvars" batch file as described at the beginning of this chapter to put the preprocessor in your path.</p>
<code>-ppOption <option></code>	<p>Specifies a preprocessor option. This option can be used multiple times to provide the command-line options for the specified preprocessor. See <code>-ppPath <path to preprocessor></code>.</p>

Table 3.1 Options for rtiddsgen

Option	Description
-ppPath <path to preprocessor>	<p>Specifies the preprocessor. If you only specify the name of an executable (not a complete path to that executable), the executable must be found in your Path. The default value is cpp for non-Windows architectures and cl.exe for Windows architectures.</p> <p>If you use -ppPath to provide the full path and filename for cl.exe or the cpp preprocessor, you must also use -ppOption <option> to set the following preprocessor options:</p> <p>If you use a non-default path for cl.exe, you also need to set: <pre>-ppOption /nologo -ppOption /C -ppOption /E -ppOption /X</pre></p> <p>If you use a non-default path for cpp, you also need to set: <pre>-ppOption -C</pre></p>
-qualifiedEnumerator	Uses the fully qualified name for enumerator values including the enum value.
-reader	Generates support for a <i>DataReader</i> (only with -micro).
-replace	<p>Deprecated option. Instead, use -update <typefiles examplefiles makefiles> for the proper files (typefiles, examplefiles, makefiles).</p> <p>This option is maintained for backwards compatibility. It allows <i>Code Generator</i> to overwrite any existing generated files.</p> <p>If it is not present and existing files are found, <i>Code Generator</i> will print a warning but will not overwrite them.</p>
-sequenceSize <unbounded sequences size>	Sets the size assigned to unbounded sequences. The default value is 100 elements.
-sharedLib	Generates makefiles that compile with the <i>Connex DDS</i> shared libraries (by default, the makefile will link with the static libraries)
-showTemplates	Prints and generates an XML file containing a list of the available example templates in your <i>Connex DDS</i> installation, organized per language. When you use the -exampleTemplate option, you can specify one of these example templates. <i>Code Generator</i> will then generate the example you specified instead of the default one.
-strict	<p>Enforces compliance with the OMG 'Extensible and Dynamic Topic Types for DDS' specification, version 1.3 by turning informational (INFO) messages into errors. (Informational messages are described in -verbosity [1-3].)</p> <p>Currently in this release, by default, <i>Code Generator</i> will report an informational (INFO) message for keyed derived structures when running <i>rtiddsgen</i>, but still generate code:</p> <pre>INFO com.rti.ndds.nddsgen.antlr.annotation.AnnotationApplier struct/valuetype derived from a struct/valuetype should not contain @key fields. This is not compliant with the XTypes specification. You can make this message an error by using the -strict flag</pre> <p>If you want <i>Code Generator</i> to enforce the specification, use the -strict option. This option will report an error when there are keyed derived structures, and not generate code.</p> <p>See the RTI Code Generator Release Notes or the <i>Migration Guide</i> on the RTI Community Portal (https://community.rti.com/documentation) for enforcement details that might occur in future releases.</p>
-stringSize <unbounded strings size>	Sets the size assigned to unbounded strings, not counting a terminating NULL character. The default value is 255 bytes.
-typeSequenceSuffix <suffix>	Assigns a suffix to the name of the implicit sequence defined for IDL types. The option is compatible with languages C, C++, C++/CLI, C#, and Java. By default, the suffix is 'Seq'. For example, given the type 'Foo', the name of the sequence would be 'FooSeq'.

Table 3.1 Options for rtiddsgen

Option	Description
-U <name>	Cancels any previous definition of <name>.
-unboundedSupport	<p>Generates code that supports unbounded sequences and strings. This option is not supported in Ada. When the option is used, the command-line options -sequenceSize and -stringSize are ignored.</p> <p>This option also affects the way unbounded sequences are deserialized. When a sequence is being received into a sample from the <i>DataReader</i>'s cache, the old memory for the sequence will be deallocated and memory of sufficient size to hold the deserialized data will be allocated. When initially constructed, sequences will not preallocate any elements having a maximum of zero elements.</p> <p>For more information on using the -unboundedSupport option, including some required QoS settings, see these sections in the RTI Connext DDS Core Libraries User's Manual:</p> <ul style="list-style-type: none"> • <i>Sequences (Section 3.1.1)</i> • <i>Strings and Wide Strings (Section 3.1.2)</i> • <i>DDS Sample and Instance Memory Management (Chapter 22)</i>
-update <typefiles examplefiles makefiles>	<p>Creates the files indicated if they do not exist. For example:</p> <pre>rtiddsgen -language C++11 -update typefiles test.idl</pre> <p>If the files already exist, -update overwrites the files without printing a warning.</p> <p>There can be multiple -update options.</p> <p>If you specify both -create and -update for the same file type, only the -update will be applied.</p> <p>If you use -update makefiles, the -platform <arch> option is required. For example:</p> <pre>rtiddsgen -language c -update makefiles -platform x64Darwin17clang9.0 test.idl</pre>
-use52Keyhash	This option should be used when compatibility with 5.2.3 and earlier General Access Releases (GARs) is required when using keyed mutable types (related to RTI Issue IDs CODEGENII-501 and CODEGENII-693).
-use526Keyhash	This option should be used when compatibility with 5.2.6 is required when using keyed mutable types (related to RTI Issue ID CODEGENII-693).
-useStdString	<p>Use 'std::string' instead of 'char*' when generating code for IDL strings when the language option is C++.</p> <p>Using this option will automatically enable constructor generation. Therefore you can use this option with or without -constructor and achieve the same result.</p>
-V <name> [= <value>]	Defines a user variable that can be used in the templates as \$userVarList.name or \$user-VarList.name.equals(value) . The variables defined with this option are case sensitive.
-verbosity [1-3]	<p>Sets the <i>Code Generator</i> verbosity:</p> <p>1: Exceptions</p> <p>2: Exceptions and warnings</p> <p>3: Exceptions, warnings and information (Default)</p>
-version	Displays the version of <i>Code Generator</i> being used, such as 2.x.y, as well as the version of the templates being used (xxxx-xxxx-xxxx).
-virtualDestructor	<p>Only applies if -language C++ is also specified.</p> <p>Generates a virtual destructor for the generated types in C++. Using this option will automatically enable the -constructor option.</p> <p>Note that using this option will affect filtering performance when using ContentFilteredTopics or QueryConditions.</p>

Table 3.1 Options for rtiddsgen

Option	Description
-writer	Generates support for a <i>DataWriter</i> (only with -micro).
XMLInputFile.xml	A file containing XML descriptions of your data types. If -inputXml is not used, the file must have an .xml extension.

Chapter 4 Generating Example Code

Most *Connex DDS Getting Started Guides* (such as the [RTI Connex DDS Getting Started Guide](#) or the [RTI Connex DDS Core Libraries Getting Started Guide Addendum for Embedded Systems](#)) have you create a simple "Hello World" application to get started. For example, on macOS:

```
$ rtiddsgen -language c++11 -example x64Darwin17clang9.0 hello_world.idl
```

Note: Before using a makefile created by *Code Generator* to compile an application, make sure the `${NDDSHOME}` environment variable is set as described in Set Up Environment Variables (rtisetenv), in "Hands-On 1" of *Introduction to Publish/Subscribe*, in the [RTI Connex DDS Getting Started Guide](#).

After you get familiar with the simple Hello World example, you can create a more advanced example by specifying the 'advanced' template option. The advanced template is already included with *Code Generator* (see [-exampleTemplate](#) and [-showTemplates](#)). For example, on macOS:

```
rtiddsgen -language C++11 -example x64Darwin17clang9.0 -exampleTemplate advanced hello_world.idl
```

There are a couple of main differences between the simple and advanced examples that *Code Generator* generates.

Note: The advanced generated example is not supported on Android™ or INTEGRITY® platforms.

4.1 is_default_qos (true vs. false)

Connex DDS loads QoS profiles from a file named `USER_QOS_PROFILES.xml` in your current working directory. (*Connex DDS* may also look for this file in other locations; see *How to Load XML-Specified QoS Settings* in the [RTI Connex DDS Core Libraries User's Manual](#)). The simple example sets `is_default_qos=true` in the `USER_QOS_PROFILES.xml` file. It creates the DDS entities without specifying a profile, so it uses the default from `USER_QOS_PROFILES.xml`.

The advanced example also loads QoS from the `USER_QOS_PROFILES.xml` file. However, the advanced example sets `is_default_qos=false`, so there is no default provided by the XML. The example explicitly specifies the QoS profile to use from the XML file when creating DDS entities.

Setting `is_default_qos=true` is a convenient way to get you started quickly, but in production applications you should explicitly specify which QoS profile to use, instead of relying on a default. See the *Basic QoS* chapter, in the [RTI Connex DDS Getting Started Guide](#).¹

4.2 Listeners vs. WaitSets

Both the simple and advanced examples use *WaitSets* to block a thread until data is available. This is the safest way to get data, because it does not affect any middleware threads (it is blocking the application's main thread until data is available).

In addition, the advanced example installs *Listeners* on both the *DataReader* and *DataWriter* with callbacks that you can implement to accomplish a desired behavior. These *Listener* callbacks are triggered for various events, such as discovering a matched *DataWriter* or *DataReader*. *Listener* callbacks are called back from a middleware thread, which you should not block. There are benefits to using *Listeners* for non-data callbacks, because you will not miss events. However, if you block or do slow processing in a *Listener*, it can cause undesired behavior such as data loss. See the *Listeners* chapter, in the [RTI Connex DDS Core Libraries User's Manual](#).

¹The *Getting Started Guide* example is more complex than the advanced example generated by *Code Generator*; however, it does not use *Listeners* like the *Code Generator* example does.

Chapter 5 Generated Files

The following tables show the files that *Code Generator* creates for an example IDL file called **Hello.idl**.

- [Table 5.1 C, Traditional C++, C++/CLI, Legacy C# Files Created for Example “Hello.idl”](#)
- [Table 5.2 Modern C++ Files Created for Example “Hello.idl”](#)
- [Table 5.3 C# Files Created for Example “Hello.idl”](#)
- [Table 5.4 Java Files Created for Example “Hello.idl”](#)
- [Table 5.5 Ada Files Created for Example “Hello.idl”](#)

Table 5.1 C, Traditional C++, C++/CLI, Legacy C# Files Created for Example “Hello.idl”

Generated Files	Description
The following files are required for the user data type. The source files should be compiled and linked with your application. The header files are required to use the data type in source. You should not modify these files unless you intend to customize the generated code supporting your type.	
Hello.[c, cxx, cpp] HelloSupport.[c, cxx, cpp] HelloPlugin.[c, cxx, cpp]	Generated code for the data types. These files contain the implementation for your data types.
Hello.h HelloSupport.h HelloPlugin.h	Header files that contain declarations used in the implementation of your data types.
The following optional files are generated when you use the -example <architecture> command-line option. You may modify and use these files as a way to create simple applications that publish or subscribe to the user data type.	
application.h (traditional C++ only)	Helper code for the Hello_publisher.cxx and Hello_subscriber.cxx applications

Table 5.1 C, Traditional C++, C++/CLI, Legacy C# Files Created for Example “Hello.idl”

Generated Files	Description
Hello_publisher.[c, cxx, cpp, cs]	Example code for an application that publishes the user data type. This example shows the basic steps to create all of the DDS objects needed to send data. You will need to modify the code to set and change the values being sent in the data structure. Otherwise, just compile and run.
Hello_subscriber.[c, cxx, cpp, cs]	Example code for an application that subscribes to the user data type. This example shows the basic steps to create all of the DDS objects needed to receive data. No modification of this file is required. It is ready for you to compile and run.
Hello-<architecture>.sln Hello_publisher-<architecture>.v[c, cs, cx]proj Hello_subscriber-<architecture>.v[c, cs, cx]-proj	Microsoft Visual Studio solution and project files, generated only for Visual Studio-based architectures. (An example <architecture> is armv8Linux4gcc7.3.0.) To compile the generated source code, open the workspace file and build the two projects.
makefile_Hello_<architecture>	Makefile for non-Visual-Studios-based architectures. An example <architecture> is armv8Linux4gcc7.3.0.
USER_QOS_PROFILES.xml	The Quality of Service (QoS) configuration of the DDS entities in the generated example is loaded from this file.

Table 5.2 Modern C++ Files Created for Example “Hello.idl”

Generated Files	Description
The following files are required for the user data type. The source files should be compiled and linked with your application. The header files are required to use the data type in source. You should not modify these files unless you intend to customize the generated code supporting your type.	
Hello.cxx HelloPlugin.cxx	Generated code for the data types. These files contain the implementation for your data types.
Hello.hpp HelloPlugin.hpp	Header files that contain declarations used in the implementation of your data types.
The following optional files are generated when you use the -example <architecture> command-line option. You may modify and use these files as a way to create simple applications that publish or subscribe to the user data type.	
application.hpp	Helper code for the Hello_publisher.cxx and Hello_subscriber.cxx applications
Hello_publisher.cxx	Example code for an application that publishes the user data type. This example shows the basic steps to create all of the DDS objects needed to send data. You will need to modify the code to set and change the values being sent in the data structure. Otherwise, just compile and run.

Table 5.2 Modern C++ Files Created for Example “Hello.idl”

Generated Files	Description
Hello_subscriber.cxx	Example code for an application that subscribes to the user data type. This example shows the basic steps to create all of the DDS objects needed to receive data. No modification of this file is required. It is ready for you to compile and run.
Hello-<architecture>.sln Hello_publisher-<architecture>.vcxproj Hello_subscriber-<architecture>.vcxproj	Microsoft Visual Studio solution and project files, generated only for Visual Studio-based architectures. (An example <architecture> is armv8Linux4gcc7.3.0.) To compile the generated source code, open the workspace file and build the two projects.
makefile_Hello_<architecture>	Makefile for non-Visual-Studios-based architectures. An example <architecture> is armv8Linux4gcc7.3.0.
USER_QOS_PROFILES.xml	The Quality of Service (QoS) configuration of the DDS entities in the generated example is loaded from this file.

Table 5.3 C# Files Created for Example “Hello.idl”

Generated Files	Description
The following files are required for the user data type. The source files should be compiled with your .NET project. You should not modify these files.	
Hello.cs HelloPlugin.cs	Generated code for the data types. These files contain the implementation for your data types.
The following optional files are generated when you use the -example <platform> command-line option (where <platform> is the .NET platform identifier, such as net5 or netcoreapp3.1). You may modify and use these files as a way to create simple applications that publish or subscribe to the user data type.	
HelloProgram.cs	Code for running the HelloPublisher.cs and HelloSubscriber.cs applications
HelloPublisher.cs	Example code for an application that publishes the user data type. This example shows the basic steps to create all of the DDS objects needed to send data. You will need to modify the code to set and change the values being sent in the data structure. Otherwise, just compile and run.
HelloSubscriber.cs	Example code for an application that subscribes to the user data type. This example shows the basic steps to create all of the DDS objects needed to receive data. No modification of this file is required. It is ready for you to compile and run.
Hello.csproj NuGet.Config	Used to build and run your application. NuGet.Config tells the .NET compiler where to find the Rti.ConnectedDs NuGet package.

Table 5.3 C# Files Created for Example “Hello.idl”

Generated Files	Description
USER_QOS_PROFILES.xml	The Quality of Service (QoS) configuration of the DDS entities in the generated example is loaded from this file.

Table 5.4 Java Files Created for Example “Hello.idl”

Data Type	Generated Files	Description
Since the Java language requires individual files to be created for each class, <i>Code Generator</i> will generate a source file for every IDL construct that translates into a class in Java.		
Constants	Hello.java	Class associated with the constant
Enums	Hello.java	Class associated with enum type
Structures/Unions	Hello.java HelloSeq.java HelloDataReader.java HelloDataWriter.java HelloTypeSupport.java	Structure/Union class Sequence class DDS <i>DataReader</i> and <i>DataWriter</i> classes Support (serialize, deserialize, etc.) class
Typedef of sequences or arrays	Hello.java HelloSeq.java HelloTypeSupport.java	Wrapper class Sequence class Support (serialize, deserialize, etc.) class
The following optional files are generated when you use the -example <i><architecture></i> command-line option. You may modify and use these files as a way to create simple applications that publish or subscribe to the user data type.		
Last structure/union	HelloPublisher.java HelloSubscriber.java	Example code for applications that publish or subscribe to the user data type. You should modify the code in the publisher application to set and change the value of the published data. Otherwise, both files should be ready to compile and run.
	makefile_Hello_ <i><architecture></i>	Makefile for non-Windows-based architectures. An example <i><architecture></i> is <code>armv8Linux4gcc7.3.0</code> .
	USER_QOS_PROFILES.xml	The Quality of Service (QoS) configuration of the DDS entities in the generated example is loaded from this file.
Structures/Unions/ Typedefs/Enums	HelloTypeCode.java	Type code class associated with the IDL type, Hello

Table 5.5 Ada Files Created for Example “Hello.idl”

Generated Files	Description
Hello[h,.c]	Generated code for the data types, which contain the implementation for the data types, and header files that contain declarations used in the implementation of the data types.
HelloSupport[h,.c]	
HelloPlugin[h,.c]	
hello_idl_file[.adb,.ads]	
hello_idl_file-hello_datawriter.ads	DataReader and DataWriter classes and serialize/deserialize methods.
hello_idl_file-hello_datareader.ads	
hello_idl_file-hello_typesupport[.adb,.ads]	
hello_idl_file-hello_mettypesupport[.adb,.ads]	These files are generated only for types that support Zero Copy transfer over shared memory (that is, are annotated with @transfer_mode(SHMEM_REF) in the IDL file).
hello_idl_file-hello_publisher[.adb,.ads] (in the samples / directory)	Example code for an application that publishes the user data type. You will need to modify the code to set and change the values being sent in the data structure. Otherwise, just compile and run. The subscriberlistener file implements the on_data_available() callback.
hello_idl_file-hello_subscriber[.adb,.ads] (in the samples / directory)	
hello_idl_file-hello_subscriberlistener[.adb,.ads] (in the samples / directory)	
hello.gpr	Project files using Ada-like syntax. These files define the build-related characteristics of the application. These characteristics include the list of sources, the location of those sources, the location of the generated object files, the name of the main program, and the options for the various tools involved in the build process. Each of them is for a different set of files (hello-samples is for the examples, hello_c is for the c files and hello is for rest of the ada files.)
hello_c.gpr	
hello-samples.gpr (in the samples directory)	
The following optional file is generated when you use the -example <architecture> command-line option. You may modify and use this file as a way to create simple applications that publish or subscribe to the user data type.	
USER_QOS_PROFILES.xml	The Quality of Service (QoS) configuration of the DDS entities in the generated example is loaded from this file.

Chapter 6 Customizing the Generated Code

Code Generator allows you to customize the generated code for different languages by changing the provided templates. This version does not allow you to create new output files.

You can load new templates using the following command in an existing template, where `<pathToTemplate>` is the absolute path to the folder with the customized template:

```
#parse("<pathToTemplate>/template.vm")
```

If that **template.vm** file contains macros, you can use it within the original template. If **template.vm** contains just plain text without macros, that text will be included directly in the original file.

You can customize the behavior of a template by using the predefined set of variables provided with *Code Generator*. For more information, see the tables in **RTI_rtiddsgen_template_variables.xlsx**.

This file contains two different sheets: Language-Templates and Template variables. The Language-Template sheet shows the correspondence between the Velocity Templates used and the generated files for each language. If, for example, we want to add a method in C in the **Hello.c** file, we would need to modify the template **typeBody.vm** under the **templates/c** directory.

The scope of a template can be:

- **type**: If we generate a file with that template for each type in the IDL file. For example in Java, where we generate a `TypeSupport` file for each type in the IDL.
- **file**: If we generate a file with that template for each IDL file. For example in C, we generate a single plugin file containing all the types `Plugin` information.
- **lastTopLevelType**: If we generate a file with that template for the last top-level type in the IDL file. This is commonly used for the publisher/subscriber examples.

- **module:** If we generate a file with that template for each module in the IDL file. This is used in Ada, where there are files that contain all the types of a module.
- **topLevelType:** if we generate a file with that template for each type in the idl file. This is used in ADA where the publisher/subscriber files are only generated for top level types

The table also shows the `top_level` variables that can be used for that templates. These variables are explained in the sheet Template variables. For example in Java, the main unit of variables are the `constructMap` which is a `HashMap` of variables that represent a type. In C, we will have as the main unit the `constructMapList`, which is a `List` of `constructMap`. In the Template variables sheet, we can see which variables are contained in each `constructMap`, the `constructKind` or type that it is applicable to and the value that it contains depending on the language we use.

One important variable that contains the `constructMap` for a type is the `memberFieldMapList`. This list represent the members contained within the type. Each member is also represented as a `HashMap` whose variables are also described in the Template variables sheet.

Apart from that there are environmental or general variables that are not related with the types that are defined within a `HashMap` called `envMap`.

Let's see how to use these variables with an example. Suppose we want to generate a method in C that prints the members for a structure and, if it is an array or sequence, its corresponding size. For this IDL:

```
module Mymodule{
    struct MyStruct{
        int32 longMember;
        int32 arrayMember [2][100];
        sequence<char,2> sequenceMember;
        sequence <int32, 5> arrayOfSequenceMember[28];
    };
};
```

We want to generate this:

```
void MyModule_MyStruct_specialPrint(){
    printf(" longMember \n");
    printf(" arrayMember is an array [2, 100] \n ");
    printf(" sequenceMember is a sequence <2> \n");
    printf(" arrayOfSequenceMember is an array [28] is a sequence <5> ");
}
```

The code in the template would look like this:

```
## We go through all the list of types
#foreach ($node in $constructMapList)
##We only want the method for structs
#*--*##if ($node.constructKind.equals("struct"))
void ${node.nativeFQName}_specialPrint(){
##We go through all the members and call to the macros that check if they are array or
sequences
#*----*##foreach($member in $node.memberFieldMapList)
print("${member.name} #isArray($member) #isASeq($member) \n");
```

```
*----*##end
}
*--*##end
#end
```

The **isAnArray** macro checks if the member is an array (i.e, has the variable **dimensionList**) and in that case, prints it:

```
#macro (isAnArray $member)
#ife($member.dimensionList) is an array $member.dimensionList #end
#end
```

The **isASeq** macro checks if the member is a sequence (i.e, has the variable **seqSize**) and in that case, prints it:

```
#macro (isASeq $member)
#ife($member.seqSize) is a sequence <$member.seqSize> #end
#end
```

You can add new variables to the templates using the **-V <name> [=<value>]** command-line option when starting *Code Generator*. This variable will be added to the `userVarList` hashMap. You can refer to it in the template as **\$userVarList.name** or **\$userVarList.name.equals(value)**.

For more information on velocity templates, see <https://velocity.apache.org/engine/1.5/user-guide.html>.

Chapter 7 Optimizing the Code Generation Process

The cost of serialization and deserialization operations increases with type complexity and sample size. It can become a significant contributor to the latency required to send and receive a sample. *Code Generator* provides the command-line option **-optimization**, which can be used to indicate the level of optimization of the serialize/deserialize operations. This command-line option allows selecting one of three different levels.

7.1 Optimization Levels

0: No optimization

1: *rtiddsgen* generates extra code for typedefs but optimizes its use. If a type that is used is a typedef that can be resolved to a primitive, enum, or aggregated type (struct, union, or value type), the generated code will invoke the code of the most basic type to which the typedef can be resolved. This level can be used if the generated code for typedef is not expected to be modified. This is the only optimization level supported for Java, C#, and C++/CLI languages.

For example:

```
typedef int32 Latitude;
typedef int32 Longitude;

struct Position {
    Latitude x;
    Longitude y;
};
```

With optimization 0, the serialization of a sample with type `Position` will require calling the serialize methods for `Latitude` and `Longitude`. For example:

```
LatitudePlugin_serialize(...) {
    serialize_long(...)
}
```

```
LongitudePlugin_serialize(...) {
    serialize_long(...)
}

Position_serialize(...) {
    LatitudePlugin_serialize(...)
    LongitudePlugin_serialize(...)
}
```

With optimization 1, *rtiddsgen* resolves Latitude and Longitude to their most primitive types for serialization purposes, resulting in a more efficient serialization. In this case, *rtiddsgen* will save two function/method calls.

```
Position_serialize(...) {
    serialize_long(...)
    serialize_long(...)
}
```

2: This optimization level is the default if not specified. (You can also explicitly specify it.) This optimization level applies only to C, C++, C++11, microC, microC++, and Ada languages. With this optimization level, *rtiddsgen* optimizes the serialization/deserialization of structures and valuetypes by using more aggressive techniques. These techniques include inline expansion of nested types and serialization/deserialization of a set of consecutive members with a single copy function invocation (`memcpy`) when the memory layout (C, C++ structure layout) is the same as the wire layout (XCDR).

7.2 How the Optimizations are Applied

In *Code Generator*, the optimizations (*inline expansion of nested types* and *serialization of consecutive members with a single copy*) are related. Inline expansion of a nested structure is only done when the C/C++ memory layout with standard packing of the structure matches the XCDR layout. (In this case, the structure's members can be serialized with a single `memcpy`.) If the C/C++ memory layout with standard packing of the structure matches the XCDR layout, then *rtiddsgen* tries first to do the inline expansion, then the serialization of consecutive members with a single copy.

7.2.1 Inline expansion of nested types

Inline expansion is an optimization in which *Code Generator* replaces a type definition with another one in which nested types are flattened out. This is done to remove extra function calls during serialization/deserialization. For example:

```
struct Point {
    int32 x;
    int32 y;
};

struct Dimension {
    int32 height;
    int32 width;
};
```

```
struct Rectangle {
    Point leftTop;
    Dimension size;
};
```

With optimization level 2, *Code Generator* replaces the definition of `Rectangle` with the following equivalent definition:

```
struct Rectangle {
    int32 leftTop_x;
    int32 leftTop_y;
    int32 size_height;
    int32 size_width;
};
```

This optimization is only done for serialization/deserialization. The generated type in C/C++ continues using `Point` and `Dimension`.

7.2.2 Serialization of consecutive members with a single copy

In the previous `Rectangle` example, *Code Generator*, using optimization level 2, further optimizes the serialization and deserialization by serializing a `Rectangle` sample with a single copy operation (`memcpy`) instead of four.

Before optimization:

```
Rectangle_serialize(...) {
    memcpy(..., 4) // leftTop_x
    memcpy(..., 4) // leftTop_y
    memcpy(..., 4) // size_height
    memcpy(..., 4) // size_width
}
```

After optimization:

```
Rectangle_serialize(...) {
    memcpy(..., 16) // leftTop_x
}
```

This optimization is only applicable when the memory layout of the C/C++ structure is equivalent to the serialization layout, which uses the XCDR version 1 or version 2 format.

7.2.3 Rules for Inline Expansion

To be inlinable, a structure 'MyStruct' has to meet the following two requirements:

- It has to have a C/C++-friendly XCDR layout.
- No members of 'MyStruct' should be marked with the `@min`, `@max`, or `@range` annotations.

A struct/valuetype 'MyStruct' has a C/C++-friendly XCDR layout when all of the following conditions apply:

- MyStruct is marked as `@final` or `@appendable` when the data representation is XCDR version 1. Mutable structures are not inlinable.
- MyStruct does not have a base type.
- MyStruct contains only primitive members, or complex members composed only of primitive members. A primitive member is a member with any of the following types: `int8`¹, `uint8`², `int16`, `int32`, `int64`, `uint16`, `uint32`, `uint64`, `float`, `double`, `octet`, and `char`. The following primitive types are not supported for inlining purposes: `long double`, `wchar`, `boolean`, `enum`.

```
struct Dimension {
    int32 height;
    int32 width;
}; // Inlinable

struct Dimension {
    string label; // Inlinable structures cannot contain strings
    int32 height;
    int32 width;
}; // Not Inlinable
```

- With any initial alignment (1, 2, 4, 8) greater than the alignment of the first member of the struct, there is no padding between the members that are part of MyStruct. To apply this rule, consider these alignments and sizes for primitive types:

Table 7.1 Alignments and Sizes for Primitive Types

Primitive Type	Alignment (bytes)	Size (bytes)
<code>int8</code> ³	1	1
<code>uint8</code> ⁴	1	1
<code>int16</code>	2	2
<code>uint16</code>	2	2
<code>int32</code>	4	4

¹This type is supported only at the API level. It is still considered an octet for type matching purposes.

²This type is supported only at the API level. It is still considered an octet for type matching purposes.

³This type is supported only at the API level. It is still considered an octet for type matching purposes.

⁴This type is supported only at the API level. It is still considered an octet for type matching purposes.

Primitive Type	Alignment (bytes)	Size (bytes)
uint32	4	4
int64	8	8
uint64	8	8
float	4	4
double	8	8
octet	1	1
char	1	1

```
struct Dimension {
    int32 height;
    int16 width;
}; // Inlinable. Independently of the alignment of the starting memory address (4 or 8),
there is no padding between height and width
```

```
struct Dimension {
    int16 height;
    int32 width;
}; // Not Inlinable. Starting in a memory address aligned to 4 will require adding a
padding of two bytes between height and width
```

- With any initial alignment (1, 2, 4, 8) greater than the alignment of the first member of the struct, there is no padding between the elements of an array of MyStruct.

```
struct Dimension {
    int32 height;
    int16 width;
}; // Not inlinable. Let's assume an array of two dimensions Dimension[2]. If the array
starts in a memory address aligned to 4, there would be padding between the first and the
second element of the array
```

```
struct Dimension {
    int32 height;
    int16 width;
    int16 padding;
}; // Inlinable
```

For serialization and deserialization purposes, *Code Generator* will consider an inlinable structure (according to the previous rules) as a primitive array where the alignment of the primitive type corresponds to the alignment of the first member of the structure. A member with type 'MyStruct' will be serialized with a single copy (memcpy) invocation.

When *Code Generator* serializes the members of a data structure, it will also try to coalesce the serialization of consecutive primitive members into a single copy operation if possible. *Code Generator* only

applies this optimization when the alignment of the next member is equal to or smaller than the alignment of the current member.

```
struct Dimension {
    int16 height;
    int32 width;
}; // Coalescing not possible because the alignment of width 4 is greater than the alignment of
height 2

struct Dimension {
    int32 width;
    int16 height;
}; // Coalescing is possible because the alignment of width 4 is greater than the alignment of
height 2
```

7.2.4 Guidelines

As a rule of thumb, to take advantage of optimization level 2 for types containing only primitive types:


- Order the members in descending alignment order (this will help with copy coalescing).
- For XCDR version 2 encapsulation, use `@final` extensibility if your types will not evolve. For XCDR version 1 encapsulation, use `@final` or `@appendable` if possible (this will help with inline expansion).
- If you use `ContentFilteredTopics`, it is recommended that fields that appear in the filter expression are placed at the beginning of the type.

Chapter 8 Boosting Performance

If you need to invoke *Code Generator* multiple times with different parameters and/or type files, there will be a performance penalty derived from loading the Java Virtual Machine (JVM) and compiling the velocity templates every time you invoke *Code Generator*. If this is your scenario, you can run *Code Generator* in server mode to avoid doing this process multiple times. Or, if you would like to reduce the JVM startup and execution time, use the JVM optimization provided with *Code Generator*, for a reduced time in code generation. These two options (server mode and JVM optimization) cannot be used together.

8.1 Using Server Mode

One way to boost performance is to run *Code Generator* in server mode. Server mode runs a native executable that opens a TCP connection to a server instance of the code generator that is spawned the first time the executable is run, as depicted below:



To invoke *Code Generator* in server mode, use the script `rtiddsgen_server(.bat)`, which is in the `scripts` directory.

The default port *Code Generator* server attaches to is 14662. If you want to modify the port *Code Generator* server is attached to, use the `-n_serverport <port>` argument. Note that the *Code Generator* server can use up to three ports; make sure you have two free ports after the specified one.

The *Code Generator* server offers a log-to-file option that you can enable by using the argument `-n_logfilepath <log directory>`. In the specified log directory, the *Code Generator* server will create a file named `CodeGenServerLog<portNumber>.txt` containing all logging messages from the server side.

Code Generator server comes with builtin timeouts, some of which you can change:

- When *Code Generator* is used in server mode, JVM is loaded a single time when the server is started; the velocity templates are also compiled a single time. The server will wait up to 5 seconds for *Code Generator* to initialize. You can change this value by specifying the number of milliseconds using the argument `-n_connectiontimeout<time in milliseconds>`.
- The *Code Generator* server will automatically stop if it is not used (that is, if it does not receive any calls) for a certain amount of time. The default value is 20 seconds; you can change this by editing the `rtiddsgen_server` script and adjusting the value of the argument `-n_servertimeout <time in milliseconds>`.
- *Code Generator* server sends a handshake message to the client after accepting the connection. In the client, there is a timeout when waiting for the handshake message. If this message is not received

in the client in a short amount of time (10 seconds, an internal value that cannot be changed), the *Code Generator* client will time out. This timeout might mean that there was another application running in the port.

Notes:

- Mixing different versions of *Code Generator* server is not supported. See *Limitations* in the [RTI Code Generator Release Notes](#).
- *Code Generator* server cannot be parallelized. Each execution of *Code Generator* server is attached to a port where it receives requests, and it can only generate code for one request at a time. Therefore, if you try to send multiple requests simultaneously, *Code Generator* server will process them sequentially.

8.2 Using JVM Optimization

Java Virtual Machine (JVM) offers different options that help improve its performance. RTI applies some of these options when creating the *Code Generator* JVM, resulting in an improvement in the *Code Generator* execution time. Because these options are non-standard, they are disabled by default. If you are using the JRE shipped with *Connex* DDS, you can enable these options; otherwise, enabling these options is at your own risk.

The options applied to *Code Generator* JVM are as follows:

- -XX:+TieredCompilation
- -XX:TieredStopAtLevel=1
- -XX:CICompilerCount=4
- -Xverify:none
- -XX:+UseParallelGC
- -XX:+OptimizeStringConcat
- -XX:CompileThreshold=5000

For more information on these options, see the JDK documentation:

<https://docs.oracle.com/en/java/javase/11/>. (Currently, you can find most of the options described in the Tools Reference, in "Main Tools to Create and Build Applications" > "java").

To enable the performance improvement associated with the *Code Generator* JVM options, set the `RTIDDSGEN_JVM_OPTIMIZATION` environment variable to true. To disable the improvement, unset the environment variable. (By default, it is not set.)