

RTI Connex
Core Libraries
Release Notes

Version 7.0.0



Contents

Chapter 1 Introduction	1
Chapter 2 System Requirements	
2.1 Introduction	3
2.2 Supported Platforms	4
2.3 Requirements when Using Microsoft Visual Studio	6
2.4 Disk and Memory Usage	7
Chapter 3 Compatibility	
3.1 Wire Protocol Compatibility	8
3.2 Code and Configuration Compatibility	9
3.3 Extensible Types Compatibility	9
Chapter 4 What's Fixed in 7.0.0	
4.1 Fixes Related to Callbacks and Waitsets	10
4.1.1 Unsafe combinations of masks and Listeners may have led to segmentation fault	10
4.1.2 Failure calling DDS_Subscriber::get_datareaders in DDS_SubscriberListener::on_data_on_readers callback implementation	10
4.1.3 DDS_SubscriberListener::on_data_on_readers on a participant or subscriber not called when Listener installed after the entity is enabled	11
4.1.4 Unable to assign callback function for on_sample_removed event using Modern C++ API	11
4.1.5 Using certain callbacks at DomainParticipant or Publisher level may have led to segmentation fault	11
4.2 Fixes Related to Discovery	12
4.2.1 Unexpected memory growth when DataReader could not be matched with DataWriter due to unexpected error condition	12
4.2.2 Possible crash upon discovery of applications with unreachable locators	12
4.2.3 Communication problems with applications using shared memory on INTEGRITY systems	12
4.2.4 Types containing Typedefs were sent without the typedefs in discovery when using DynamicData	13
4.2.5 Unbounded memory growth in Spy when discovering multiple endpoints with the same Topics and types	13

4.2.6	Unnecessary discovery traffic related to IP mobility events on interfaces irrelevant to the transport	13
4.3	Fixes Related to Transports	14
4.3.1	Communication problems with applications using shared memory on INTEGRITY systems	14
4.3.2	Race condition could cause unbounded memory growth in TCP Transport Plugin	14
4.4	Fixes Related to Filtering and TopicQuery	15
4.4.1	Unnecessary repair traffic for DataWriters using TopicQueries and asynchronous publishing	15
4.4.2	Connex application using filtering feature may have crashed after running out of memory	15
4.4.3	Unnecessary sample filtering on a DataReader for samples already filtered by a DataWriter	15
4.4.4	Creation of a ContentFilteredTopic or reception of TopicQuery samples may have taken long time for complex types	16
4.4.5	Continuous creation of TopicQueries may have led to unnecessary memory fragmentation in OS memory allocator	16
4.4.6	rti::topic::find_registered_content_filters led to infinite recursion	16
4.4.7	Incorrect results for Unions when using DynamicData or Content Filters	16
4.4.8	Samples may have been unnecessarily filtered by Connex DataReader when DataWriter was from different DDS vendor	17
4.5	Fixes Related to Group Presentation	18
4.5.1	Application may not have received samples of coherent set when using GROUP access scope and TRANSIENT_LOCAL durability	18
4.5.2	Application may stop receiving samples from DataReaders using GROUP_PRESENTATION_QOS	18
4.5.3	Segmentation fault when using GROUP_PRESENTATION_QOS or HIGHEST_OFFERED_PRESENTATION_QOS and setting filter_redundant_samples to FALSE on DataReader	19
4.6	Fixes Related to XML Configuration	19
4.6.1	Parsing error loading XML configuration file containing a const whose expression refers to an enumerator	19
4.6.2	Discrepancy between range defined by schema and that defined by API	19
4.6.3	Parsing error loading XML configuration file with enum type containing enumerator whose value was an expression referring to a const	20
4.6.4	Parsing error loading an XML configuration file with enum type containing enumerator whose value was an expression	20
4.6.5	Type limits not checked for some attributes of XML types definition	20
4.6.6	Removed some elements in the XSD that were not supported internally but could be defined in XML	21
4.6.7	Builtin Discovery Plugins was not treated as a mask by the XSD file	22
4.6.8	Parsing error loading an XML configuration file with an enum type containing an enumerator whose value was an expression referring to another enumerator	22
4.7	Fixes Related to Vulnerabilities	23
4.7.1	Fixes related to Connex	23
4.7.2	Fixes related to third-party dependencies	23
4.8	Fixes Related to APIs	24
4.8.1	Input parameters to Property and DataTag helper functions do not have "const"	24

4.8.2	Standard 64-bit integer types are now supported (Modern C++ API)	25
4.8.3	Assigning DataWriter and DataReaderQos from a TopicQos caused a build error	25
4.8.4	Copy of SampleInfo::coherent_set_info field was not supported	25
4.8.5	In XML-based applications, generated IDL types did not take precedence over XML DynamicTypes (C# API)	26
4.8.6	Namespaces ignored when a type was explicitly registered in C# for XML-based applications	26
4.8.7	Corruption of LoanedDynamicData object when moved in some situations (Modern C++ API only)	26
4.8.8	Calling DynamicData::set_complex_member with an aliased type failed	27
4.8.9	Possible wrong results when adding Time or Duration objects that used very large numbers	27
4.8.10	Java API did not support RtpsReliableReaderProtocol_t.receive_window_size	28
4.9	Fixes Related to Crashes	28
4.9.1	Simultaneous deletion of an entity by multiple threads caused a crash when using Java	28
4.9.2	DataReader C++ application crashed if it received tampered sample with unsupported encapsulation ID	28
4.9.3	Segmentation fault after calling DomainParticipant::register_durable_subscription with a group containing a long role_name	29
4.9.4	Segmentation fault when application using MultiChannel ran out of memory	29
4.9.5	Application crashed when capturing traffic for a DomainParticipant created before enabling network capture	29
4.9.6	Possible crash when writing a sample	29
4.9.7	Potential crash during type registration if system ran out of memory	30
4.9.8	Segmentation fault after calling DomainParticipant::delete_durable_subscription with a group containing a long role_name	30
4.9.9	Potential crash or memory corruption if user application using thread-specific storage	30
4.10	Other Fixes	31
4.10.1	Serialization/deserialization of non-primitive sequences and arrays for XCDR2_DATA_REPRESENTATION did not follow Extensible Types specification	31
4.10.2	Possible hang when using best-effort writers and asynchronous publishing	31
4.10.3	Unnecessary sockets created during initialization of library	31
4.10.4	Various issues with RtpsReliableWriterProtocol_t::nack_suppression_duration	31
4.10.5	Possible error message printed during Entity disposal	32
4.10.6	Runtime error when using debug libraries for QNX x86 platform	32
4.10.7	Pushed samples may not have been received by reliable DataReader when DataWriter published Type that supports Zero Copy transfer over shared memory	33
4.10.8	Unbounded memory growth in Monitoring Library when creating and deleting endpoints	33
4.10.9	Unexpected behavior when two threads crashed at the same time on Windows systems	34
4.10.10	DataReaders setting reader_qos.protocol.expects_inline_qos to TRUE incorrectly matched with DataWriters	34
4.10.11	Source IP on Spy was not correct when DataWriters with same Topic were on different machines	34
4.10.12	Writer using durable writer history may not have blocked after send window filled up when disable pos-	35

itive ACKs was enabled	
4.10.13 Potential truncation of application-level acknowledgment response data	35
4.10.14 Error messages displayed that should not have been, when printing DataReaderQoS objects	35
4.10.15 Potential Valgrind invalid read when logging a message or enabling heap monitoring	36
4.10.16 Malformed IDL printed if multiple labels used for default case of a union	36
Chapter 5 Known Issues	
5.1 Known Issues with Usability	37
5.1.1 Cannot open USER_QOS_PROFILES.xml in rti_workspace/examples from Visual Studio	37
5.1.2 DataWriter's Listener callback on_application_acknowledgment() not triggered by late-joining DataReaders	37
5.1.3 HighThroughput and AutoTuning built-in QoS Profiles may cause communication failure when writing small samples	38
5.1.4 Memory leak if Foo::initialize() called twice	38
5.1.5 Wrong error code after timeout on write() from Asynchronous Publisher	38
5.1.6 Type Consistency enforcement disabled for structs with more than 10000 members	39
5.1.7 Escaping special characters in regular/filter expressions not supported in some cases	39
5.1.8 Calling take_snapshot within DDS Entity callbacks not supported	39
5.1.9 banish and subject_name APIs are unresolved in Modern C++ Windows dynamic libraries	40
5.2 Known Issues with Code Generation	40
5.2.1 Examples and generated code for Visual Studio 2017 and later may not compile (Error MSB8036)	40
5.3 Known Issues with Instance Lifecycle	40
5.3.1 Instance does not transition to ALIVE when "live" DataWriter detected	40
5.3.2 Persistence Service DataReaders ignore serialized key propagated with dispose updates	41
5.4 Known Issues with Reliability	41
5.4.1 DataReaders with different reliability kinds under Subscriber with GROUP_PRESENTATION_QOS may cause communication failure	41
5.5 Known Issues with Content Filters and Query Conditions	41
5.5.1 Writer-side filtering may cause missed deadline	41
5.5.2 filter_sample_* statistics in DDS_DataWriterProtocolStatus not updated correctly	41
5.6 Known Issues with TopicQueries	42
5.6.1 TopicQueries not supported with DataWriters configured to use batching or Durable Writer History	42
5.7 Known Issues with Transports	42
5.7.1 AppAck messages cannot be greater than underlying transport message size	42
5.7.2 DataReader cannot persist AppAck messages greater than 32767 bytes	43
5.7.3 Discovery with Connex Micro fails when shared memory transport enabled	43
5.7.4 Communication may not be reestablished in some IP mobility scenarios	43
5.7.5 Corrupted samples may be forwarded through Routing Service when using Zero-Copy transfer over shared memory	43

5.7.6 Network Capture does not support frames larger than 65535 bytes	44
5.8 Known Issues with FlatData	45
5.8.1 FlatData language bindings do not support automatic initialization of arrays of primitive values to non-zero default values	45
5.8.2 Flat Data: plain_cast on types with 64-bit integers may cause undefined behavior	45
5.8.3 FlatData in combination with payload encryption in RTI Security Plugins and/or compression will not save copies	45
5.9 Known Issues with Coherent Sets	46
5.9.1 Some coherent sets may be lost or reported as incomplete with batching configurations	46
5.9.2 Copy of SampleInfo::coherent_set_info field is not supported	46
5.9.3 Other known issues with coherent sets	46
5.10 Known Issues with Dynamic Data	46
5.10.1 Conversion of data by member-access primitives limited when converting to types that are not supported on all platforms	46
5.10.2 Types that contain bit fields not supported	47
5.11 Known Issues in RTI Monitoring Library	47
5.11.1 Problems with NDDS_Transport_Support_set_builtin_transport_property() if Participant Sends Monitoring Data	47
5.11.2 Participant's CPU and memory statistics are per application	47
5.11.3 XML-based entity creation nominally incompatible with static monitoring library	47
5.11.4 ResourceLimit channel_seq_max_length must not be changed	48
5.12 Known Issues with Installers	48
5.12.1 RTI Connex Micro 3.0.3 installation package currently compatible only with Connex 6.0.1 installer	48
5.13 Other Known Issues	48
5.13.1 Possible Valgrind still-reachable leaks when loading dynamic libraries	48
5.13.2 64-bit discriminator values greater than (2 ³¹ -1) or smaller than (-2 ³¹) supported only in Java, no other languages	49
5.13.3 Creating multiple DataReaders for the same Topic under the same Subscriber configured with Group Ordered Access is not supported	50
5.13.4 Segmentation fault when invoking Discovery Snapshot APIs in some cases	50
Chapter 6 Experimental Features	51

Chapter 1 Introduction

RTI® Connext® 7.0.0 is an Early Access release, based on release 6.1.1.

This document includes the following:

- [System Requirements \(Chapter 2 on page 3\)](#)
- [Compatibility \(Chapter 3 on page 8\)](#)
- [What's Fixed in 7.0.0 \(Chapter 4 on page 10\)](#)
- [Known Issues \(Chapter 5 on page 37\)](#)
- [Experimental Features \(Chapter 6 on page 51\)](#)

For an overview of new features in 7.0.0, see [RTI Connext Core Libraries What's New in 7.0.0](#).

Many readers will also want to look at additional documentation available online. In particular, RTI recommends the following:

- **Use the RTI Customer Portal** (<https://support.rti.com>) to download RTI software and contact RTI Support. The RTI Customer Portal requires a username and password. You will receive this in the email confirming your purchase. If you do not have this email, please contact license@rti.com. Resetting your login password can be done directly at the RTI Customer Portal.
- **The RTI Community Forum** (<https://community.rti.com>) provides a wealth of knowledge to help you use *Connext*, including:
 - Documentation, at <https://community.rti.com/documentation>
 - Best Practices,
 - Example code for specific features, as well as more complete use-case examples,
 - Solutions to common questions,

- A glossary,
 - Downloads of experimental software,
 - And more.
- Whitepapers and other articles are available from <http://www.rti.com/resources>.
 - Performance benchmark results for *Connex*t are published online at <http://www.rti.com/products/dds/benchmarks.html>. Updated results for new releases are typically published within two months after general availability of that release.

Chapter 2 System Requirements

2.1 Introduction

Connex requires a multi-threaded operating system. This section describes the supported host and target systems.

In this context, a host is the computer on which you will be developing a *Connex* application. A target is the computer on which the completed application will run. A host installation provides the *RTI Code Generator* tool (*rtiddsgen*), examples and documentation, as well as the header files required to build a *Connex* application for any architecture. You will also need a target installation, which provides the libraries required to build a *Connex* application for that particular target architecture.

Supported platforms, for all products in the *Connex* suite are listed in these tables:

- [Table 1 Supported Platforms for Compiler-Dependent Products on the next page.](#)
This table is for products that are compiled into your application.
- [Table 2 Supported Platforms for Connex Applications on page 5](#)
This table is for products that are ready to use, no compilation required.

Early Access releases are intended to showcase the latest *Connex* features; they support a smaller subset of platforms in comparison to LTS releases. The upcoming LTS release shall support a larger number of platforms.

Subsequent Early Access and LTS releases may not support all of the platforms supported in this release, or may support different versions of platforms supported in this release.

See the *RTI Connex Core Libraries Platform Notes* for more information on each platform.

2.2 Supported Platforms

Table 1 Supported Platforms for Compiler-Dependent Products

Platforms					Connex Professional						Connex Secure			Connex Anywhere	Add-ons
					Core Libraries	LBED [10]	Secured Transport	Infrastructure Services			Security Plugins				
OS	OS Version	CPU	Toolchain	RTI Architecture					TLS Support [1]	Persistence Service	Routing Service	Recording Service	with openssl [1]	for wolfSSL [8]	SDK [1]
Linux	Red Hat Enterprise Linux 8 Ubuntu 18.04, 20.04, 22.04 LTS	x64	gcc 7.3.0	x64Linux4gcc7.3.0	•	•	•	•	•	•	•		•	•	•
	Red Hat Enterprise Linux 7, 7.3, 7.5, 7.6 CentOS 7.0	x64	gcc 4.8.2	x64Linux3gcc4.8.2	•	•	•	•	•	•	•		•	•	•
	Ubuntu 18.04, 22.04 LTS	Arm v8 [2]	gcc 7.3.0	armv8Linux4gcc7.3.0	•		•		•	•	•			•	
Windows	Windows 10, 11 Server 2016 [3]	x64	VS 2017, 2019, 2022	x64Win64VS2017	•	•	•	•	•	•	•		•	•	•
	Windows 10 Windows Server 2012 R2, 2016	x64	VS 2015	x64Win64VS2015	•	•	•	•	•	•	•		•	•	•
macOS	macOS 10.13-10.15, 11, 12 (host and target) [9]	x64	clang 9.0-12.0, 13.0	x64Darwin17clang9.0	•		•	•	•	•	•		•	•	
	macOS 11, 12 (target only) [11]	Arm v8	clang 12.0, 13.0	arm64Darwin20clang12.0	•		•	•	•	•	•			•	
QNX	QNX Neutrino 7.1	Arm v8 [2]	qcc 8.3.0	armv8QNX7.1qcc_gpp8.3.0	•		•		•	•	•	•		•	
VxWorks	VxWorks 21.11 [9]	x64	llvm 12.0.1.1	x64Vx21.11llvm12.0.1.1 x64Vx21.11llvm12.0.1.1_rtp	•							• [4]		•	

Table 2 Supported Platforms for Connex Applications

Platforms			Connex Professional						Connex Anywhere
			Infrastructure Services	Tools					
OS	OS Version	CPU	Web Integration Service	Shapes Demo	Launcher	Monitor	Admin Console	System Designer	Cloud Discovery Service
Linux	Red Hat Enterprise Linux 8 Ubuntu 18.04, 20.04, 22.04 LTS	x64	•	•	•	•	•	• [5]	•
	Red Hat Enterprise Linux 7, 7.3, 7.5, 7.6 CentOS 7.0	x64	•	•	•	•	•	• [5]	•
	Ubuntu 18.04, 22.04 LTS	Arm v8 [2]	•						
Windows	Windows 10, 11 Windows Server 2016 [3]	x64	•	•	•	•	•	• [7]	•
	Windows 10 Windows Server 2012 R2, 2016	x64	•	•	•	•	•	• [7]	•
macOS	macOS 10.13-10.15, 11, 12 (host and target) [9]	x64	•	•	•	•	•	• [6]	•
	macOS 11, 12 (target only) [11]	Arm v8	•						
QNX	QNX Neutrino 7.1	Arm v8 [2]							
VxWorks	VxWorks 21.11 [9]	x64							

• = Supported

Blue = New compared to 6.1.1

1 Tested with openssl 1.1.1n unless stated otherwise

2 These libraries require a hardware FPU in the processor and are compatible with systems with hard-float libc

3 Per Microsoft, this should be compatible with Windows 10 IoT Enterprise with Windows native application

4 Tested with OpenSSL from VxWorks 7

5 Tested on Ubuntu 18.04 LTS only, with Chrome 77, Firefox 69, and Safari 12

6 Tested on macOS 10.14 only, with Chrome 77, Firefox 69, and Safari 12

7 Tested on Windows 10 only, with Chrome 77, Firefox 69 and Safari 12

8 Tested with wolfSSL 4.7

9 Future releases may support a different version

10 LBED = Limited Bandwidth Endpoint Discovery Plugin

11 Requires Rosetta@2 during installation, not required at runtime

2.3 Requirements when Using Microsoft Visual Studio

Note: Debug versions of applications and the various Visual C++ DLLs are not redistributable. Therefore, if you want to run debug versions, you must have the compiler installed.

When Using Visual Studio 2015 — Update 3 Redistributable Package Requirement

You must have the Visual C++ Redistributable for Visual Studio 2015 Update 3 installed on the machine where you are running an application linked with dynamic libraries. This includes C/C++ dynamically linked and all Java applications.

You can download the Visual C++ Redistributable for Visual Studio 2015 Update 3 from this Microsoft website: <https://www.microsoft.com/en-us/download/details.aspx?id=53840>.

When Using Visual Studio 2017 — Redistributable Package Requirement

You must have the Visual C++ Redistributable for Visual Studio 2017 installed on the machine where you are running an application linked with dynamic libraries. This includes C/C++ dynamically linked and all Java applications.

You can download the Visual C++ Redistributable for Visual Studio 2017 from this Microsoft website: <https://visualstudio.microsoft.com/vs/older-downloads/>. Then look in this section: "Redistributables and Build Tools" for Microsoft Visual C++ Redistributable for Visual Studio 2017".

When Using Visual Studio 2019 — Redistributable Package Requirement

You must have the Visual C++ Redistributable for Visual Studio 2019 installed on the machine where you are running an application linked with dynamic libraries. This includes C/C++ dynamically linked and all Java applications.

You can download the Visual C++ Redistributable for Visual Studio 2019 from this Microsoft website: <https://www.visualstudio.com/downloads/>. Then look in this section: "Other Tools and Frameworks" for Microsoft Visual C++ Redistributable for Visual Studio 2019".

When Using Visual Studio 2022 — Redistributable Package Requirement

You must have the Visual C++ Redistributable for Visual Studio 2022 installed on the machine where you are running an application linked with dynamic libraries. This includes C/C++ dynamically linked and all Java applications.

You can download the Visual C++ Redistributable for Visual Studio 2019 from this Microsoft website: <https://www.visualstudio.com/downloads/>. Then look in this section: "Other Tools, Frameworks, and Redistributables" for Microsoft Visual C++ Redistributable for Visual Studio 2022".

2.4 Disk and Memory Usage

Disk usage for a typical host-only installation is approximately 802 MB on Linux systems and 821 MB on Windows systems. Each additional architecture (host or target) requires an additional 498 MB on Linux systems and 609 MB on Windows systems.

We recommend that you have at least 256 MB RAM installed on your host development system. The target requirements are significantly smaller and they depend on the complexity of your application and hardware architecture.

Chapter 3 Compatibility

Below is basic compatibility information for this release.

Note: For backward-compatibility information between this and previous releases, see the *Migration Guide* on the RTI Community Portal (<https://community.rti.com/documentation>).

3.1 Wire Protocol Compatibility

*Connex*t communicates over the wire using the formal Real-Time Publish-Subscribe (RTPS) protocol. RTPS has been developed from the ground up with performance, interoperability and extensibility in mind. The RTPS protocol is an international standard managed by the OMG. The RTPS protocol has built-in extensibility mechanisms that enable new revisions to introduce new message types, extend the existing messages, or extend the Quality of Service settings in the product—without breaking interoperability.

RTPS 1.0 was introduced in 2001. The currently supported version is [OMG Real-Time Publish-Subscribe \(RTPS\) specification, version 2.5](#), although some features are not supported. Unsupported features currently are RTPS HeaderExtension, FilteredCountFlag in GAP Submessage, HeartbeatFrag Submessage, Checksum, and ALIVE_FILTERED instance state. RTI plans to maintain interoperability between middleware versions based on RTPS 2.1. For more details, see [Table 3.1 RTPS Versions](#).

[Table 3.1 RTPS Versions](#) shows RTPS versions supported for each *Connex*t release. In general, RTPS 2.1 and higher versions are interoperable, unless noted otherwise. RTPS 2.0 and RTPS 1.2 are incompatible with current (4.2e and later) versions of *Connex*t.

Although RTPS 2.1 and higher versions are generally interoperable, there may be specific wire protocol interoperability issues between *Connex*t releases. These issues are documented in the "Wire Protocol" section for your release, in the *Migration Guide* on the RTI Community Portal (<https://community.rti.com/documentation>). Wire protocol issues between 5.3.1 and previous releases are documented in the *RTI Connex*t Core Libraries Release Notes for release 5.3.1.

Table 3.1 RTPS Versions

Connex Release	RTPS Standard Version (a)	RTPS Protocol Version (b)
Connex 6 and above	2.5 (partial support)	2.3
Connex DDS 5.2 and 5.3	2.2	2.2
Connex DDS 4.5f - 5.1	2.1	2.1
Data Distribution Service 4.2e - 4.5e	2.1	2.1
Data Distribution Service 4.2c	2.0	2.0
Data Distribution Service 4.2b and lower	1.2	1.2

(a) Version number of the RTPS standards document, [OMG Real-Time Publish-Subscribe \(RTPS\) specification, version 2.5](#).

(b) RTPS wire protocol version number that *Connex* announces in messages it puts on the wire.

3.2 Code and Configuration Compatibility

The *Connex* core uses an API that is an extension of the [OMG Data Distribution Service \(DDS\) standard API, version 1.4](#). RTI strives to maintain API compatibility between versions, but will conform to changes in the OMG DDS standard.

The *Connex* core primarily consists of a library and a set of header files. In most cases, upgrading simply requires you to recompile your source using the new header files and link the new libraries. In some cases, minor modifications to your application code might be required; any such changes are noted in the *Migration Guide* on the RTI Community Portal (<https://community.rti.com/documentation>). The *Migration Guide* also indicates whether and how to regenerate code.

3.3 Extensible Types Compatibility

This release of *Connex* includes partial support for the [OMG 'Extensible and Dynamic Topic Types for DDS' specification, version 1.3](#) (DDS-XTypes) from the Object Management Group (OMG). This support allows systems to define data types in a more flexible way, and to evolve data types over time without giving up portability, interoperability, or the expressiveness of the DDS type system.

For information related to compatibility issues associated with the Extensible Types support, see the *Migration Guide* on the RTI Community Portal (<https://community.rti.com/documentation>). See also the *RTI Connex Core Libraries Extensible Types Guide* for a full list of the supported and unsupported extensible types features.

Chapter 4 What's Fixed in 7.0.0

This section describes bugs fixed in *Connex* 7.0.0. These fixes have been made since 6.1.1 was released.

4.1 Fixes Related to Callbacks and Waitsets

4.1.1 Unsafe combinations of masks and Listeners may have led to segmentation fault

When entities are created, a *Listener* may be provided by the user to receive calls when specified events occur. The events of interest are set using a **StatusKind** mask. If an event set in the **StatusKind** mask occurs, but no callback function has been assigned by the user, a null pointer dereference will occur. *Connex* checks for many of these errors and prevents the creation of entities when this error is present. However, some of these cases were not checked, allowing unsafe combinations of masks and *Listeners* to be used. This problem has been resolved. The new, stricter checking may cause entity creation errors when no errors were detected before.

[RTI Issue ID CORE-12610]

4.1.2 Failure calling `DDS_Subscriber::get_datareaders` in `DDS_SubscriberListener::on_data_on_readers` callback implementation

You may have seen the following errors when invoking `DDS_Subscriber::get_datareaders()` within the implementation of the `DDS_SubscriberListener::on_data_on_readers()` callback:

```
ERROR [0x01011B2D,0x8A450DE1,0xBAE5A2A0:0x80000009|SET GROUP LISTENER|GET READERS]
REDACursor_modifyReadWriteArea:!freeze read write area
ERROR [0x01011B2D,0x8A450DE1,0xBAE5A2A0:0x80000009|SET GROUP LISTENER|GET READERS]
PRESPsReaderGroup_getEA:!modify pres psReaderGroup
ERROR [0x01011B2D,0x8A450DE1,0xBAE5A2A0:0x80000009|SET GROUP LISTENER|GET READERS]
PRESPsReaderGroup_lock:!take semaphore
ERROR [0x01011B2D,0x8A450DE1,0xBAE5A2A0:0x80000009|SET GROUP LISTENER|GET READERS]
PRESPsReaderGroup_beginGetPsReaders:!get PRESPsReaderGroup_lock
```


4.1.3 DDS_SubscriberListener::on_data_on_readers on a participant or subscriber not called when

```
ERROR [0x01011B2D,0x8A450DE1,0xBAE5A2A0:0x80000009|SET GROUP LISTENER|GET READERS}
DDS_Subscriber_begin_get_datareadersI:ERROR: Failed to get PRESPsReaderGroup_
beginGetPsReaders
ERROR [0x01011B2D,0x8A450DE1,0xBAE5A2A0:0x80000009|SET GROUP LISTENER|GET READERS]
DDS_Subscriber_get_datareaders:ERROR: Failed to get DDS_Subscriber_begin_get_datareaders
```

In addition, when using the Traditional C++ API and the legacy .NET API, the application generated a segmentation fault after printing the error. The problem occurred only when:

- You installed a *Listener* on the *Subscriber* using the API `DDS_Subscriber::set_listener()` after the *Subscriber* was enabled.
- Or, you installed a *Listener* on the *DomainParticipant* using the API `DDS_Participant::set_listener()` after the *DomainParticipant* was enabled. This problem has been resolved.

[RTI Issue ID CORE-12316]

4.1.3 DDS_SubscriberListener::on_data_on_readers on a participant or subscriber not called when Listener installed after the entity is enabled

The callback `DDS_SubscriberListener::on_data_on_readers()` was not invoked when there was data available, if these two conditions were met:

- The *Listener* callback `on_data_on_readers()` was installed after the *Subscriber* or *DomainParticipant* implementing it was enabled.
- The *Listener* callback `on_data_available()` was not installed at any level (*DomainParticipant*, *Publisher*, or *DataReader*).

This problem has been resolved.

[RTI Issue ID CORE-12338]

4.1.4 Unable to assign callback function for on_sample_removed event using Modern C++ API

You may have been unable to assign a callback function for the `on_sample_removed` event using the Modern C++ API. Support for this callback has been added to the Modern C++ API in this release.

[RTI Issue ID CORE-12646]

4.1.5 Using certain callbacks at DomainParticipant or Publisher level may have led to segmentation fault

Handlers were not correctly implemented for the `on_instance_replaced()`, `on_sample_removed()`, `on_application_acknowledgment()`, and `on_service_request_accepted()` callbacks at the

DomainParticipant and *Publisher* levels. This could have led to segmentation faults when the corresponding events were enabled. This problem has been resolved.

[RTI Issue ID CORE-12647]

4.2 Fixes Related to Discovery

4.2.1 Unexpected memory growth when *DataReader* could not be matched with *DataWriter* due to unexpected error condition

Failing to match a *DataReader* with a *DataWriter* because of unexpected error conditions may have led to unexpected memory growth, because *Connex* may not have cleaned up the resources associated with the remote match completely. This problem has been resolved.

[RTI Issue ID CORE-8257]

4.2.2 Possible crash upon discovery of applications with unreachable locators

If an application used `DDS_STATUS_MASK_ALL` for a *DomainParticipant* or *Publisher Listener* and an unreachable locator was discovered, the application enabling the *Listener* may have crashed. An unreachable locator occurs most commonly when a Subscribing application uses a transport that the Publishing application does not use. For example, the Publishing application could use UDPv4 and the Subscribing application could use both UDPv4 and UDPv6.

More rarely, a crash may have occurred when a pre-5.2.0 Subscribing application used the shared memory transport and a 5.2.0+ Publishing application was not using the UDPv6 transport. A log message was generated if both participants were running on the same machine and this condition occurred. This condition was caused by a change to the way that transports are identified starting in version 5.2.0.

[RTI Issue ID CORE-11818]

4.2.3 Communication problems with applications using shared memory on INTEGRITY systems

If an application on an INTEGRITY platform used the shared-memory transport, the *Connex* libraries sometimes incorrectly assessed that a shared-memory segment was stale and could be reclaimed, when in fact it was not stale. This situation caused problems with communication between *DomainParticipants*, since information could be sent to a shared-memory segment that did not get dequeued by the intended recipient.

You may have seen error messages like these and the application may have hung while deleting the *DomainParticipant*:

```
<Target Output> ERROR RTIOsapiSharedMemoryBinarySemaphore_take:OS WaitForSemaphore ()
failure, error 0XD: ObjectClosed
<Target Output> ERROR NDSS_Transport_Shmem_receive_rEA:!take semaphore
```

4.2.4 Types containing Typedefs were sent without the typedefs in discovery when using DynamicData

```
<Target Output> ERROR RTIOsapiSharedMemoryBinarySemaphore_take:OS WaitForSemaphore ()  
failure, error 0X9: ObjectIsUseless
```

This problem has been resolved.

Incompatibility with 6.1.1 and prior releases:

The fix for this issue involved some changes that make shared-memory segments in applications incompatible with those in 6.1.1 (and earlier) versions.

[RTI Issue ID CORE-12097]

4.2.4 Types containing Typedefs were sent without the typedefs in discovery when using DynamicData

When an application was using a `DynamicDataReader` or `DynamicDataWriter` and using a type that contained a typedef, the type that was sent during endpoint discovery for that endpoint did not contain the typedef. While this did not cause any mismatches or communication failure, it did cause a number of issues that may have been noticeable depending on what other products you may have also been using.

See [4.2.5 Unbounded memory growth in Spy when discovering multiple endpoints with the same Topics and types below](#) for details about the specific issues that you may have encountered. The *RTI Admin Console Release Notes* and *RTI Routing Service Release Notes* also have related information. (See [ADMINCONSOLE-997](#) and [ROUTING-971](#), respectively.)

This issue has been resolved, meaning that the exact type definition that is registered with the participant, containing typedefs, is sent during discovery. This is a change in behavior from 6.0.0-based applications, which sent the type definitions without the typedef information.

[RTI Issue ID CORE-12107]

4.2.5 Unbounded memory growth in Spy when discovering multiple endpoints with the same Topics and types

Each time *DDS Spy* discovered an endpoint, it unnecessarily made a copy of the `TypeCode` that was associated with the endpoint's *Topic*, leading to unbounded memory growth. This issue has been fixed.

[RTI Issue ID CORE-12136]

4.2.6 Unnecessary discovery traffic related to IP mobility events on interfaces irrelevant to the transport

When there is a change on a network interface (an IP mobility event), a *Connex* application will update and resend its discovery information to include these changes. The transport can consider a change on an interface irrelevant (for example, changes on interfaces in the `deny_interfaces_list` of the transport). In this case, the new discovery messages are exactly the same as announced before, generating unnecessary discovery traffic that could affect the performance of the application.

This problem has been fixed. Now *Connex* only updates and resends its discovery information if there was a change on an interface relevant to the transport.

[RTI Issue ID CORE-12664]

4.3 Fixes Related to Transports

4.3.1 Communication problems with applications using shared memory on INTEGRITY systems

If an application on an INTEGRITY platform used the shared-memory transport, the *Connex* libraries sometimes incorrectly assessed that a shared-memory segment was stale and could be reclaimed, when in fact it was not stale.

This situation caused problems with communication between *DomainParticipants*, since information could be sent to a shared-memory segment that did not get dequeued by the intended recipient.

You may have seen error messages like these and the application may have hung while deleting the *DomainParticipant*:

```
<Target Output> ERROR RTIOsapiSharedMemoryBinarySemaphore_take:OS WaitForSemaphore ()  
failure, error 0XD: ObjectClosed  
<Target Output> ERROR NDDS_Transport_Shmem_receive_rEA:!take semaphore  
<Target Output> ERROR RTIOsapiSharedMemoryBinarySemaphore_take:OS WaitForSemaphore ()  
failure, error 0X9: ObjectIsUseless
```

This problem has been resolved.

Incompatibility with 6.1.1 and prior releases:

The fix for this issue involved some changes that make the shared memory segments incompatible with those in 6.1.1 (and earlier) versions.

[RTI Issue ID CORE-12097]

4.3.2 Race condition could cause unbounded memory growth in TCP Transport Plugin

Due to a race condition, the TCP Transport Plugin may have leaked memory when creating a new connection if the creation happened at the same time the *DomainParticipant* was being destroyed. The cause of the leak was the TCP Transport Plugin reallocating memory that was already released by the *DomainParticipant*. The race condition was unlikely to happen. However, in a system that frequently creates and destroys entities (and, therefore, TCP connections) and that runs for long enough, it may have lead to unbounded memory growth. The issue has been resolved.

[RTI Issue ID COREPLG-618]

4.4 Fixes Related to Filtering and TopicQuery

4.4.1 Unnecessary repair traffic for DataWriters using TopicQueries and asynchronous publishing

Samples that are sent in response to a TopicQuery are directed to the *DataReader* that created that TopicQuery. This means that those samples are only sent to the *DataReader* that made the request and have that *DataReader's* GUID attached to each sample in the sample's metadata. All other *DataReaders* receive GAP protocol messages, indicating to them that a given sequence number or set of sequence numbers is not meant for them.

Due to a defect, when a *DataReader* sent a NACK message requesting some TopicQuery samples to be repaired, if the requested sequence numbers included samples that were meant for a different *DataReader*, the *DataWriter* did not filter these samples and resend a GAP message. Instead, the *DataWriter* sent the *DataReader* samples that were not meant for it and the *DataReader* had to filter these samples out itself. As a result, the *DataReaders* may have received samples that should have been filtered out on the *DataWriter* side, leading to an increase in network traffic.

The problem only affected repair traffic. When a sample was filtered out by the *DataWriter* because it was directed to a different *DataReader*, the *DataWriter* sent a GAP protocol message to the *DataReader*. If the GAP message was lost, the *DataReader* NACKed for the sample; instead of sending a new GAP message, the *DataWriter* sent the sample. This problem has been resolved.

[RTI Issue ID CORE-12589]

4.4.2 Connex application using filtering feature may have crashed after running out of memory

In release 6.1.1.2, a *Connex* application using filtering features (that is, ContentFilteredTopic, QueryConditions, or TopicQuery) may have crashed after running out of memory. This problem has been resolved.

[RTI Issue ID CORE-12661]

4.4.3 Unnecessary sample filtering on a DataReader for samples already filtered by a DataWriter

When doing writer-side filtering, a late-joining *DataReader* using a ContentFilteredTopic may have spent unnecessary CPU cycles evaluating samples that pass the ContentFilteredTopic's expression. When using writer-side filtering, the filter evaluation is done by the *DataWriter* and it should not be necessary for the *DataReader* to do it again on samples that pass the filter expression. This problem, which only occurred for late-joining *DataReaders*, has been fixed.

[RTI Issue ID CORE-11084]

4.4.4 Creation of a ContentFilteredTopic or reception of TopicQuery samples may have taken long time for complex types

The creation of a ContentFilteredTopic or reception of TopicQuery samples, may have taken a long time for complex types. This issue has been resolved.

[RTI Issue ID CORE-12179]

4.4.5 Continuous creation of TopicQueries may have led to unnecessary memory fragmentation in OS memory allocator

In releases 6.0.x and 6.1.x, the continuous creation of TopicQueries may have led to unnecessary memory fragmentation in the OS memory allocator of the applications that receive the TopicQuery requests and dispatch responses. This issue may have resulted in an unexpected increase of the resident set size (RSS) memory of the application receiving and dispatching the TopicQueries compared to previous *Connex* releases. This problem has been fixed.

[RTI Issue ID CORE-12352]

4.4.6 rti::topic::find_registered_content_filters led to infinite recursion

The function `rti::topic::find_registered_content_filters()` was incorrectly implemented and would lead to infinite recursion and stack overflow in any application that called it. This problem has been resolved. This function returns the names of previously registered custom content filters. It is a little-used feature and does not affect the commonly used SQL content filter.

[RTI Issue ID CORE-12512]

4.4.7 Incorrect results for Unions when using DynamicData or Content Filters

When using a DynamicDataReader, samples containing a union may have had incorrect or invalid data after deserialization if the *DataReader's* type contained members that were not present in the *DataWriter's* type and those members had non-zero default values.

When using content filters, the filter results may have been incorrect if the type contained a union and the filter expression filtered on fields within the union that were present in the *DataReader's* type but were not present in the *DataWriter's* type and those members had non-zero default values.

For example, see this *DataWriterType*:

```
struct innerStructPub {
    short shortMember;
};
@mutable
union ComplexUnionTypePub switch(long) {
    case 0:
        long longMember;
    case 1:
```

4.4.8 Samples may have been unnecessarily filtered by Connex DataReader when DataWriter was

```
innerStructPub structMember;
};
```

and this *DataReaderType*:

```
struct innerStructSub {
    short shortMember;
    @default(5) long longMemberWithDefault;
};
@mutable
union ComplexUnionTypeSub switch(long) {
    case 0:
        long longMember;
    case 1:
        innerStructSub structMember;
};
```

In the above types, the member **longMemberWithDefault** is only present in the *DataReader's* type and has a default value of 5, so any sample that is received from the *DataWriter* should have this value set to 5 when read from the *DataReader's* queue. Instead, the value was incorrectly 0 when using *DynamicData*.

In addition, if this member was used as part of a content filter expression, a *DataReader* always used the value of 0 instead of 5 when evaluating a sample from a *DataWriter* using the *DataWriterType* which could lead to incorrect filter results. These issues have been fixed.

[RTI Issue ID CORE-12517]

4.4.8 Samples may have been unnecessarily filtered by Connex DataReader when DataWriter was from different DDS vendor

A Connex *DataReader* using a *ContentFilteredTopic* unnecessarily evaluated its filter on samples coming from a different vendor *DataWriter* that already marked the samples as passing the *DataReader* filter. This issue may have led to an increase in CPU utilization on the *DataReader* side, but it did not affect functional correctness or bandwidth utilization.

The problem occurred because *Connex* was not compliant with the way a filter signature is calculated according to the Section 9.6.4.1, *Content filter info (PID_CONTENT_FILTER_INFO)*, in the [Real-time Publish-Subscribe Protocol DDS Interoperability Wire Protocol \(DDSI-RTPSTM\) Specification version 2.5](#).

This problem has been resolved.

[RTI Issue ID CORE-12531]

4.5 Fixes Related to Group Presentation

4.5.1 Application may not have received samples of coherent set when using GROUP access scope and TRANSIENT_LOCAL durability

An application using **GROUP** access scope and **TRANSIENT_LOCAL** (or higher) durability may not have received the samples for some coherent sets, or it may have received the samples with delay.

Assume a coherent set '**CS1**' published by a set of *DataWriters* that are part of the same group. This coherent set was not provided to the application if all the following conditions were true:

1. The *DataReaders* receiving '**CS1**' matched with the *DataWriters* publishing '**CS1**' after the coherent set was published.
2. '**CS1**' did not contain samples for some of the *DataWriters* in the group, or the samples were removed after applying the Lifespan QoS Policy. If '**CS1**' contained at least one sample per *DataWriter* in the group, this problem did not occur.
3. The application did not publish a new coherent set after '**CS1**'; or, if it did, the new coherent set did not contain samples from at least one of the *DataWriters* that were missing samples from '**CS1**'.

If the third condition was not met, then the delivery of the coherent set would be delayed instead of the coherent set not being provided.

[RTI Issue ID CORE-12350]

4.5.2 Application may stop receiving samples from DataReaders using GROUP_PRESENTATION_QOS

An application may have stopped receiving samples from *DataReaders* that were part of a *Subscriber* using **GROUP_PRESENTATION_QOS** under the following scenario:

- The *Publisher's* group contained at least one keyed *DataWriter* and one unkeyed *DataWriter*.
- The *Subscriber's* group contained only keyed *DataReaders* or unkeyed *DataReaders*, but not both.

This problem has been resolved.

[RTI Issue ID CORE-12161]

4.5.3 Segmentation fault when using GROUP_PRESENTATION_QOS or HIGHEST_OFFERED_PRESENTATION_QOS and setting filter_redundant_samples to FALSE on DataReader

An application generated a segmentation fault if it created a *DataReader* with the following valid configuration:

- **subscriber_qos.presentation.access_scope = DDS_GROUP_PRESENTATION_QOS or DDS_HIGHEST_OFFERED_PRESENTATION_QOS**
- **datareader_qos.availability.max_data_availability_waiting_time = DDS_DURATION_ZERO**
- **datareader_qos.availability.max_endpoint_availability_waiting_time = DDS_DURATION_ZERO**
- **datareader_qos.property** contained **dds.data_reader.state.filter_redundant_samples** with the value “false”

This problem has been resolved by allowing the *DataReader* to be created.

[RTI Issue ID CORE-12771]

4.6 Fixes Related to XML Configuration

4.6.1 Parsing error loading XML configuration file containing a const whose expression refers to an enumerator

Connex failed to load an XML configuration file containing a const whose expression referred to an enumerator. For example:

```
<enum name="Enum1">
  <enumerator name="Enumerator1" value="1"/>
</enum>
<const name="Const1" type="int32" value="Enumerator1+1"/>
```

Loading this XML failed with an error similar to this:

```
DDS_XMLConst_initialize:Parse error at line 10: type 'Enum1' is not typedef
```

This problem has been fixed.

[RTI Issue ID CORE-5553]

4.6.2 Discrepancy between range defined by schema and that defined by API

There were discrepancies between the ranges defined by the schema files and those defined by the API for certain elements. This problem has been resolved. Now, validating an XML against the XSD should not fail when setting a value that is inside the range as defined by the API.

[RTI Issue ID CORE-7099]

4.6.3 Parsing error loading XML configuration file with enum type containing enumerator whose value was an expression referring to a const

Connexrt failed to load an XML configuration file with an enum type containing an enumerator whose value was an expression referring to a const. For example:

```
<const name="Const1" type="int32" value="10"/>
<enum name="Enum1">
  <enumerator name="Enumerator1" value="Const1"/>
</enum>
```

Loading this XML failed with an error similar to this:

```
DDS_XMLEnum_on_start_tag:Parse error at line xy: integer expected
```

This problem has been fixed.

[RTI Issue ID CORE-10060]

4.6.4 Parsing error loading an XML configuration file with enum type containing enumerator whose value was an expression

Connexrt failed to load an XML configuration file with an enum type containing an enumerator whose value was an expression. For example:

```
<enum name="Enum1">
  <enumerator name="Enumerator1" value="1 + 1"/>
</enum>
```

Loading this XML failed with an error similar to this:

```
DDS_XMLEnum_on_start_tag:Parse error at line xy: integer expected
```

This problem has been fixed.

[RTI Issue ID CORE-10269]

4.6.5 Type limits not checked for some attributes of XML types definition

When XML was used for defining types (for example, when using *DynamicData*), type limits were not checked for some attributes. If the specified value for any of the attributes was too large or too small, a variable overflow occurred, leading to undefined behavior.

This problem is fixed. Type limits are checked, throwing a meaningful error when they are not met.

The affected attributes were as follows:

- **value** in union's **caseDiscriminator**. Valid values should be between -2147483648 and 2147483647.
- **sequenceMaxLength**. Valid values should be between 0 and 2147483647. -1 (unbounded) is also allowed.
- **stringMaxLength**. Valid values should be between 0 and 2147483647. -1 (unbounded) is also allowed.
- **arrayDimensions**. Valid values should be between 1 and 4294967295.

[RTI Issue ID CORE-12181]

4.6.6 Removed some elements in the XSD that were not supported internally but could be defined in XML

The following elements were configurable in XML although internally they are not supported:

Publisher QoS:

- **presentation.drop_incomplete_coherent_set**
- **asynchronous_publisher.thread.cpu_list**
- **asynchronous_publisher.thread.cpu_rotation**
- **asynchronous_publisher.asynchronous_batch_thread.cpu_list**
- **asynchronous_publisher.asynchronous_batch_thread.cpu_rotation**
- **asynchronous_publisher.topic_query_publication_thread.cpu_list**
- **asynchronous_publisher.topic_query_publication_thread.cpu_rotation**

Participant QoS:

- **discovery_config.publication_reader.min_app_ack_response_keep_duration**
- **discovery_config.subscription_reader.min_app_ack_response_keep_duration**
- **discovery_config.asynchronous_publisher.thread.cpu_list**
- **discovery_config.asynchronous_publisher.thread.cpu_rotation**
- **discovery_config.asynchronous_publisher.disable_asynchronous_batch**
- **discovery_config.asynchronous_publisher.asynchronous_batch_thread**
- **discovery_config.asynchronous_publisher.disable_topic_query_publication**
- **discovery_config.asynchronous_publisher.topic_query_publication_thread**

EventQosPolicy:

- **thread.cpu_list**
- **thread.cpu_rotation**

DatabaseQosPolicy:

- **thread.cpu_list**
- **thread.cpu_rotation**

Those elements have been removed from the XSD and are no longer configurable in XML.

[RTI Issue ID CORE-12366]

4.6.7 Builtin Discovery Plugins was not treated as a mask by the XSD file

Because of a bug in the XML Schema Definition (XSD), if you specified more than one value for the **DiscoveryConfigQosPolicy::builtin_discovery_plugins** mask, your XML editor reported that the expression was not valid when it should have been.

For example, according to the XSD, this expression was not allowed:

```
<domain_participant_qos>
  <discovery_config>
    <builtin_discovery_plugins>SPDP|SEDP</builtin_discovery_plugins>
  </discovery_config>
</domain_participant_qos>
```

This issue has been fixed, and the XSD now accepts expressions containing more than one Builtin Discovery Plugin. This issue occurred only while editing XML files because of the schema. If you ran an application with the above configuration, it did not fail.

[RTI Issue ID CORE-12740]

4.6.8 Parsing error loading an XML configuration file with an enum type containing an enumerator whose value was an expression referring to another enumerator

Connex failed to load an XML configuration file with an enum type containing an enumerator whose value was an expression using another enumerator. For example:

```
<enum name="Enum1">
  <enumerator name="Enumerator1" value="0"/>
</enum>

<enum name="Enum2">
```

```
<enumerator name="Enumerator2" value="Enumerator1"/>
</enum>
```

Loading this XML would have failed with an error similar to this:

```
DDS_XMLEnum_on_start_tag:Parse error at line xy: integer expected
```

This problem has been fixed.

[RTI Issue ID CORE-12781]

4.7 Fixes Related to Vulnerabilities

4.7.1 Fixes related to Connex

This release fixes some potential vulnerabilities, including RTI Issue IDs CORE-12510 and CORE-12752.

4.7.2 Fixes related to third-party dependencies

This release fixes some potential vulnerabilities related to third-party dependencies, described below.

4.7.2.1 Potential crash or leak of sensitive information in Core Libraries XML parser due to vulnerabilities in Expat

The Core Libraries XML parser had a third-party dependency on Expat version 2.4.4, which is known to be affected by a number of publicly disclosed vulnerabilities.

These vulnerabilities have been fixed by upgrading Expat to the latest stable version, 2.4.8. See "Third-Party Software Upgrades" in [RTI Connex Core Libraries What's New in 7.0.0](#).

The impact on *Connex* applications of using the previous version varied depending on your *Connex* application configuration:

- With *Connex Secure* (enabling RTPS protection):
 - Exploitable through a compromised local file system containing malicious XML/DTD files.
 - Could lead to arbitrary code execution.
 - CVSS v3.1 Score: 8.4 HIGH
 - CVSS v3.1 Vector: [AV:L/AC:L/PR:N/UI:N/S:U/C:H/I:H/A:H](#)
- Without *Connex Secure*:
 - Exploitable through a compromised local file system containing malicious XML/DTD files.

- Remotely exploitable through malicious RTPS messages.
- Could lead to arbitrary code execution.
- CVSS v3.1 Score: 9.8 CRITICAL
- CVSS v3.1 Vector: [AV:N/AC:L/PR:N/UI:N/S:U/C:H/I:H/A:H](#)

[RTI Issue ID CORE-12872]

4.7.2.2 Potential memory corruption when using Zlib compression due to vulnerability in Zlib

The ZRTPS transport in the Core Libraries had a third-party dependency on Zlib version 1.2.11, which is known to be affected by a publicly disclosed vulnerability.

This vulnerability has been fixed by upgrading Zlib to the latest stable version, 1.2.12. See "Third-Party Software Upgrades" in [RTI Connex Core Libraries What's New in 7.0.0](#).

The impacts on *Connex* applications of using the previous version were as follows:

- Exploitable by triggering the compression of a sample containing a malicious payload.
- The application could crash.
- CVSS v3.1 Score: 7.5 HIGH
- CVSS v3.1 Vector: [AV:N/AC:L/PR:N/UI:N/S:U/C:N/I:N/A:H](#)

[RTI Issue ID CORE-12877]

4.8 Fixes Related to APIs

4.8.1 Input parameters to Property and DataTag helper functions do not have "const"

In the C API, the following functions were incorrectly missing a **const** before the **policy** parameter:

- **DDS_PropertyQosPolicyHelper_lookup_property()**
- **DDS_PropertyQosPolicyHelper_lookup_property_with_prefix()**
- **DDS_PropertyQosPolicyHelper_get_properties()**
- **DDS_DataTagQosPolicyHelper_lookup_tag()**

This problem has been fixed. The policies are now "const" because these functions do not change the policy.

[RTI Issue ID CORE-3166]

4.8.2 Standard 64-bit integer types are now supported (Modern C++ API)

Previous releases of the Modern C++ API had platform-specific definitions for 64-bit integers, defined in `rti::core::int64` and `rti::core::uint64`. This was required to support certain pre-C++11 platforms.

This release redefines those two types as `std::int64_t` and `std::uint64_t`.

[RTI Issue ID CORE-10913]

4.8.3 Assigning DataWriter and DataReaderQos from a TopicQos caused a build error

DataWriterQos and DataReaderQos could not be constructed from a TopicQos assignment. You may have seen a compiler error such as:

```
error: conversion from 'TEntityQos<rti::topic::qos::TopicQosImpl>' to
non-scalar type 'TEntityQos<rti::pub::qos::DataWriterQosImpl>' requested.
```

This problem has been resolved. Now this type of assignment works correctly. Any fields that are not in the TopicQos will use the default for the DataWriterQos or DataReaderQos.

[RTI Issue ID CORE-11185]

4.8.4 Copy of SampleInfo::coherent_set_info field was not supported

`SampleInfo::coherent_set_info` was not available when using take/read operations that did not loan the samples. The `SampleInfo::coherent_set_info` field was always set to NULL when you called the take/read operations that did not loan the samples. To get the `coherent_set_info` value, you had to use the read/take operations that loan the data.

In addition, the copy constructor and assignment operator in the Traditional C++ and Modern C++ APIs did not copy the `SampleInfo::coherent_set_info` field. This field was always set to NULL; it was your responsibility to make the copy and handle memory allocation and deletion for this field.

This problem has been fixed. If you work with the C API, starting with this release you will have to use the following functions to manipulate SampleInfo structures:

- `DDS_SampleInfo_initialize()`
- `DDS_SampleInfo_copy()`
- `DDS_SampleInfo_finalize()`

[RTI Issue ID CORE-11213]

4.8.5 In XML-based applications, generated IDL types did not take precedence over XML DynamicTypes (C# API)

In the C# API in previous releases, if a type was declared in XML as a dynamic type and also generated and registered by the application, the XML dynamic type took precedence. This led to the DataReaders or DataWriters using DynamicData instead of the generated C# user class. This behavior was unintuitive and inconsistent with the other language APIs. It has been resolved.

[RTI Issue ID CORE-11389]

4.8.6 Namespaces ignored when a type was explicitly registered in C# for XML-based applications

When a type was explicitly registered (this is only necessary to support generated IDL types with XML-Based Application Creation) as follows:

```
DomainParticipantFactory.RegisterType<A.B.Foo>()
```

The registered type name was to set to “Foo” instead of the expected “A::B::Foo”. In some situations, this may have stopped applications written in other languages to communicate with a C# application, if the regular algorithm of type matching was disabled.

[RTI Issue ID CORE-12074]

4.8.7 Corruption of LoanedDynamicData object when moved in some situations (Modern C++ API only)

Given a DynamicData sample, accessing a nested member within another nested member via **loan_value()** and then moving the latter may have corrupted the former. For example, given a sample such that "my_sample.a.b" is a member of a constructed type (struct or union):

```
DynamicData my_sample(my_dynamic_type);
LoanedDynamicData loan1 = my_sample.loan_value("a");
LoanedDynamicData loan2 = loan1.get().loan_value("b");
// The following corrupts loan2
LoanedDynamicData loan1_moved = std::move(loan1);
```

This may have affected applications that explicitly move-constructed a double-nested LoanedDynamicData or that otherwise indirectly called the move constructor in this situation (for example, by resizing a std::vector of LoanedDynamicData elements).

The LoanedDynamicData's move constructor and move-assignment operators have been fixed.

[RTI Issue ID CORE-12272]

4.8.8 Calling `DynamicData::set_complex_member` with an aliased type failed

Calling `DynamicData::set_complex_member()` with an aliased type failed. For example, given the following types:

```
struct Foo {
long x;
long y;
};
typedef Foo TypedefFoo;
struct MyType {
Foo my_inner_struct;
TypedefFoo my_typedef_struct;
};
```

The following code should have worked to set the `my_typedef_struct` member:

```
DDS_DynamicData *data = DDS_DynamicData_new(
    MyType_get_typecode(),
    &DDS_DYNAMIC_DATA_PROPERTY_DEFAULT);
DDS_DynamicData *inner_data = DDS_DynamicData_new(
    TypedefFoo_get_typecode(),
    &DDS_DYNAMIC_DATA_PROPERTY_DEFAULT);

// This call fails. If the above call used Foo_get_typecode instead then it would work

retcode = DDS_DynamicData_set_complex_member(data, "my_typedef_struct", 0, inner_data);
if (retcode != DDS_RETCODE_OK) {
    fprintf(stderr, "set_complex_member %d\n", retcode);
    return -1;
}
```

But instead, it failed with these errors:

```
DDS_DynamicData2_copy: Objects have different types. self type = TypedefFoo, other type =
TypedefFoo
DDS_DynamicData2_finalize_ex: finalizing object bound to a member, automatically unbinding
now.
DDS_DynamicData2_set_complex_member:ERROR: Failed to copy value
DDS_DynamicData2_unbind_complex_member:ERROR: Bad parameter: self has no bound member
DDS_DynamicData2_set_complex_member:!unbind complex member
```

This issue has been fixed. Now, using either the aliased type (`TypedefFoo` in our example) or the original type (`Foo` in our example) works to set a complex member using the `DynamicData` API.

[RTI Issue ID CORE-12273]

4.8.9 Possible wrong results when adding Time or Duration objects that used very large numbers

Adding Time or Duration objects could have previously produced wrong results when using very large numbers. Necessary checks are now in place to ensure that wrong results do not occur.

[RTI Issue ID CORE-12413]

4.8.10 Java API did not support RtpsReliableReaderProtocol_t.receive_window_size

This QoS setting was ignored by the Java API, and readers were always created with the default value (256). This problem has been resolved.

[RTI Issue ID CORE-12451]

4.9 Fixes Related to Crashes

4.9.1 Simultaneous deletion of an entity by multiple threads caused a crash when using Java

When two threads deleted an entity at the same time, in Java, this may have caused a crash with the following backtrace:

```
#7 0x00007f7c630dad3b in REDAWeakReference_getReferent (reference=0x78,
slNode=0x7f7c4407f988, frOut=0x0, tableWithStartedCursor=0x7f7c6452c000) at
WeakReference.c:144
#8 0x00007f7c630d2ff3 in REDACursor_gotoWeakReference (c=0x7f7c4407f950, fr=0x0, wr=0x78) at
Cursor.c:230
#9 0x00007f7c62d5ed46 in PRESPsService_destroyLocalEndpoint (me=0x7f7c64367cc0,
failReason=0x7f7cb0136fc0, group=0x7f7c64dbb340, endpoint=0x7f7c644f0e88,
worker=0x7f7c44015f70) at PsService.c:2130
#10 0x00007f7c62b6fc26 in PRESParticipant_destroyLocalEndpoint (me=0x7f7c64368a00,
failReason=0x7f7cb0136fc0, group=0x7f7c64dbb340, endpoint=0x7f7c644f0e88,
worker=0x7f7c44015f70) at Participant.c:5882
#11 0x00007f7c636fcc32 in DDS_DataReader_deleteI (reader=0x7f7c644f1070) at
DataReader.c:4250
#12 0x00007f7c6372667e in DDS_Subscriber_delete_datareader (self=0x7f7c64dbb620,
reader=0x7f7c644f1070) at Subscriber.c:1159
#13 0x00007f7c63daf24b in Java_com_rti_dds_subscription_SubscriberImpl_DDS_1Subscriber_
1delete_1datareader (env=0x7f7c781061f8, self_class=0x7f7cb0137148, self=140172244792864,
readerL=140172235575408) at SubscriberImpl.c:790
```

This issue has been resolved. Now one thread will remove the entity and the other thread will throw an exception with the error code **com.rti.dds.infrastructure.RETCODE_ALREADY_DELETED**.

[RTI Issue ID CORE-10768]

4.9.2 DataReader C++ application crashed if it received tampered sample with unsupported encapsulation ID

If a C++ application with a *DataReader* received a sample with a tampered or malformed encapsulation kind, a segmentation fault occurred when the *DataReader* attempted to deserialize the sample, leading to an application crash. This problem has been fixed.

[RTI Issue ID CORE-12356]

4.9.3 Segmentation fault after calling `DomainParticipant::register_durable_subscription` with a group containing a long `role_name`

An application using the API `DomainParticipant::register_durable_subscription()` may have experienced a segmentation fault if the `role_name` of the input group was NULL or had a length greater than 512 bytes. This problem has been fixed.

[RTI Issue ID CORE-12460]

4.9.4 Segmentation fault when application using `MultiChannel` ran out of memory

A *Connex* application using `MultiChannel` might have produced a segmentation fault if the system ran out of memory. This problem has been fixed.

[RTI Issue ID CORE-12493]

4.9.5 Application crashed when capturing traffic for a `DomainParticipant` created before enabling network capture

To capture network traffic, you must enable this feature before creating the *DomainParticipants* that will capture the traffic. Applications not satisfying this requirement crashed when starting, pausing, or resuming the capture.

This problem has been fixed. *Connex* will no longer crash in this situation, but will fail and log messages such as the following:

```
ERROR NDDS_Utility_start_network_capture_w_params_for_participant:!get network capture manager for DomainParticipant. Network capture must be enabled before creating the DomainParticipant

ERROR NDDS_Utility_start_network_capture_for_participant:!network capture could not be started for the participant

ERROR NDDS_Utility_run_network_capture_operation_for_all_participants:!failed to run network capture operation for participant

ERROR NDDS_Utility_start_network_capture_w_params:!error starting network capture for all participants

ERROR NDDS_Utility_start_network_capture:!start network capture for all participants. There was at least one participant that could not be started
```

[RTI Issue ID CORE-12511]

4.9.6 Possible crash when writing a sample

Due to an internal error, an application could crash when writing a sample using either a best-effort or reliable *DataWriter*. Before the crash, an error message in either of the following functions was printed:

```
* COMMENDBeWriterService_write
* COMMENDSrWriterService_write
```

This problem has been resolved.

[RTI Issue ID CORE-12561]

4.9.7 Potential crash during type registration if system ran out of memory

A crash may have occurred during type registration if the application ran out of memory. This problem has been resolved.

[RTI Issue ID CORE-12734]

4.9.8 Segmentation fault after calling `DomainParticipant::delete_durable_subscription` with a group containing a long `role_name`

An application using the API `DomainParticipant::register_durable_subscription()` may have experienced a segmentation fault if the `role_name` of the input group was NULL or had a length greater than 512 bytes. This problem has been fixed.

[RTI Issue ID CORE-12787]

4.9.9 Potential crash or memory corruption if user application using thread-specific storage

Starting with release 6.1.0, there was an issue that could lead to a potential crash or memory corruption if the user application was using thread-specific storage.

In particular, when using Activity Context or Heap Monitoring, a race condition could have been triggered upon creating a thread with the ThreadFactory at the same time the DomainParticipantFactory instance was initialized or finalized. When this race condition was triggered, *Connex* might have overwritten the user application's thread-specific storage, leading to memory corruption or crashes.

This issue is now fixed. If the race condition that led to the issue happens in an application, the following benign warning will be logged:

```
Unexpected RTIOsapiContextSupport_g_tssKey value. This could mean that this thread was
created at the same time you are destroying the DDSDomainParticipantFactory.
```

If that is the case, Activity Context and Heap Monitoring won't be available for that thread.

[RTI Issue ID CORE-12966]

4.10 Other Fixes

4.10.1 Serialization/deserialization of non-primitive sequences and arrays for XCDR2_DATA_REPRESENTATION did not follow Extensible Types specification

The serialization/deserialization of sequences and arrays with non-primitive members for XCDR2_DATA_REPRESENTATION did not follow the [OMG 'Extensible and Dynamic Topic Types for DDS' specification, version 1.3](#). This led to compatibility issues with other DDS implementations.

This problem has been fixed, although the new behavior is not enabled by default, in order to keep backward compatibility with previous *Connex* releases. You can configure a *DomainParticipant* to align with the specification by setting `dds.type_plugin.dheader_in_non_primitive_collections` to true in the *DomainParticipant*'s PROPERTY QoS Policy for all the *DomainParticipants* created by your *Connex* applications.

[RTI Issue ID CORE-12464]

4.10.2 Possible hang when using best-effort writers and asynchronous publishing

Due to an internal error, an application hung when using a best-effort writer and asynchronous publishing. Before the hang, the following error message was printed:

```
COMMENTBeWriterService_write:!retrieveJob  
This problem is now fixed.
```

[RTI Issue ID CORE-12562]

4.10.3 Unnecessary sockets created during initialization of library

The initialization of the *Connex* libraries always created a socket to obtain the IP address of the first valid interface of the machine. This IP address is used to generate identifiers when `DDS_DomainParticipantQos::wire_protocol::rtps_auto_id_kind` is `DDS_RTTPS_AUTO_ID_FROM_IP`, which is not the default value. Therefore, the creation of this socket was unnecessary in most cases. This problem has been solved, and now the socket is only created when it is needed.

[RTI Issue ID CORE-12587]

4.10.4 Various issues with `RtpsReliableWriterProtocol_t::nack_suppression_duration`

There were various issues with the `RtpsReliableWriterProtocol_t::nack_suppression_duration` QoS:

- NACKs were being incorrectly suppressed with asynchronous publishing or non-zero **min/max_nack_response_delay** if two NACK messages were received within the **nack_suppression_duration** window, even if they were NACKing for different sets of sequence numbers. The **nack_suppression_duration** is only meant to suppress NACKs with the same leading sequence number that are received within the **nack_suppression_duration** window. If two consecutive NACKs have different leading sequence numbers, this indicates that the reader is making progress and the second one should not be suppressed, regardless of the **nack_suppression_duration**. Incorrect suppression of NACKs was not an issue if **min/max_nack_response_delay** was zero and **PublishModeQosPolicy.kind** was **SYNCHRONOUS_PUBLISH_MODE_QOS**.
- If a NACK was received and suppressed due to the **nack_suppression_duration** before the previous NACK was responded to, then the NACK that had not been responded to yet, along with all NACKs for the duration of the **nack_suppression_duration**, were incorrectly suppressed. This problem did not occur if **min/max_nack_response_delay** was zero and **PublishModeQosPolicy.kind** was **SYNCHRONOUS_PUBLISH_MODE_QOS**.
- When **PublishModeQosPolicy.kind** was **ASYNCHRONOUS_PUBLISH_MODE_QOS**, if there were no GAP messages sent in response to a NACK, the **nack_suppression_duration** had no effect and NACKs were never suppressed. (GAP messages are sent to a *DataReader* to indicate that a sample or a set of samples are not meant for that *DataReader*. This can happen, for example, because the *DataWriter* has applied writer-side filtering or no longer has those samples in its queue.)

These issues have been resolved.

[RTI Issue ID CORE-12603]

4.10.5 Possible error message printed during Entity disposal

Upon the disposal of an entity, an error message from a callback associated with an event may have been printed. An excerpt of what the error may have looked like this:

```
ERROR [0x01013D3F,0x79453D76,0xA3558BB2:0x00000000]REMOVE_REMOTE_DR
0x01013D3F,0x79453D76,0xA3558BB2:0x80000007] OnReliableReaderActivityChangedCallback:An
exception was thrown: Omg.Dds.Core.DdsException: DDS operation failed:
  at Rti.Dds.NativeInterface.Helpers.ReturnCode.CheckResult(IntPtr result)
...
```

The disposal of entities has now been modified to ensure this error does not happen.

[RTI Issue ID CORE-12641]

4.10.6 Runtime error when using debug libraries for QNX x86 platform

When using the `i86QNX6.6qcc_cpp4.7.3` debug libraries, your application may have had a runtime error and hung. This was because the debug libraries included the symbol for a math function (“`isinf`”)

that was discontinued in QNX 6.3.

This problem has been resolved. The debug libraries now include “**isinf**” instead, which is supported.

A full list of the math functions that were discontinued in QNX 6.3 can be found here: http://www.qnx.com/developers/docs/6.6.0.update/index.html#com.qnx.doc.neutrino.lib_ref/topic/whats_new_630.html.

Note: QNX platforms on x86 are not supported in Connex 7.0.0.

[RTI Issue ID CORE-12695]

4.10.7 Pushed samples may not have been received by reliable DataReader when DataWriter published Type that supports Zero Copy transfer over shared memory

A reliable *DataReader* may not have received pushed samples from a *DataWriter* publishing a *Topic* on a type configured with the zero-copy transfer over shared memory **@transfer_mode(SHMEM_REF)**. This may have led to significant performance degradation because the *DataReader* has to continuously NACK the missing samples.

This problem only occurred when the following three conditions were true:

1. The *DataWriter* ran in a different host, or the *DataReader* did not have the builtin SHMEM transport enabled.
2. The *DataReader* used a ContentFilteredTopic, and the *DataWriter* did writer-side filtering, or the *DataReader* created TopicQueries.
3. The *DataWriter* was not configured to use an asynchronous publisher. This problem has been resolved.

[RTI Issue ID CORE-12775]

4.10.8 Unbounded memory growth in Monitoring Library when creating and deleting endpoints

Each time a *DataWriter* or *DataReader* is created in an application that has the RTI Monitoring Library enabled, a new instance is created in the *DataWriters* of the library. Since, by default, the maximum number of instances the *DataWriter* can handle is unlimited, and the instances of already deleted endpoints were not cleaned up automatically, unbounded memory growth was possible in the library's *DataWriters* when constantly creating and deleting endpoints in an application that had Monitoring Library enabled.

This problem has been fixed by setting the **writer_data_lifecycle::autopurge_disposed_instances_delay** QoS to **DDS_DURATION_ZERO** by default in the *DataWriters* for the Monitoring Library. That way, disposed instances will be instantly cleared.

[RTI Issue ID MONITOR-244]

4.10.9 Unexpected behavior when two threads crashed at the same time on Windows systems

When two threads crashed at the same time on Windows systems, *Connex* may have concurrently called the function **SymInitialize()** from **DbgHelp** from two crashing threads.

SymInitialize() is not thread safe, so the application may have run into unexpected behavior or memory corruption under this scenario.

This has been resolved, *Connex* no longer calls **SymInitialize()** from a crashing thread. Instead, **SymInitialize()** is now called during `DomainParticipantFactory` initialization.

[RTI Issue ID CORE-10066]

4.10.10 DataReaders setting `reader_qos.protocol.expects_inline_qos` to TRUE incorrectly matched with DataWriters

Connex DataWriters matched *DataReaders* that set `reader_qos.protocol.expects_inline_qos` to TRUE. This behavior was incorrect because *Connex DataWriters* do not support sending inline QoS, and they were not honoring the *DataReaders'* requests.

This issue has been fixed. The behavior has changed so that *DataWriters* will not match *DataReaders* that request inline QoS (i.e., that set `reader_qos.protocol.expects_inline_qos` to TRUE).

[RTI Issue ID CORE-10501]

4.10.11 Source IP on Spy was not correct when DataWriters with same Topic were on different machines

The source IP on Spy may not have been correct when *DataWriters* with the same Topic were on different machines. This issue has been fixed. Now the source IP is per Entity, not per Topic, and the output will look like this:

```
11:35:13 New reader from 10.200.130.20 : topic=""Example app"" type=""app""
11:35:18 New writer from 10.200.129.195 : topic=""Example app"" type=""app""
11:35:16 New writer from 10.200.130.3 : topic=""Example app"" type=""app""
11:42:58 New data from 10.200.129.195 : topic=""Example app"" type=""app""
11:42:58 New data from 10.200.130.3 : topic=""Example app"" type=""app""
11:43:00 New data from 10.200.129.195 : topic=""Example app"" type=""app""
11:43:00 New data from 10.200.130.3 : topic=""Example app"" type=""app""
```

[RTI Issue ID CORE-12169]

4.10.12 Writer using durable writer history may not have blocked after send window filled up when disable positive ACKs was enabled

In previous releases, a reliable *DataWriter* configuring a finite send window size may not have blocked when the send window filled up if all these conditions were met:

- *DataWriter* was configured to use durable writer history.
- *DataWriter* was configured to use disable positive ACKs.
- *DataWriter* was configured with **writer_qos.reliability.acknowledgment_kind** set to AUTO or EXPLICIT, or **writer_qos.availability.enable_required_subscriptions** was set to TRUE.

Because of this issue, the reliability protocol for the *DataWriter* may have been less efficient. This problem has been resolved.

[RTI Issue ID CORE-12225]

4.10.13 Potential truncation of application-level acknowledgment response data

Connex enforced a wrong maximum length for application-level acknowledgment response data. Specifically, *Connex* incorrectly allowed setting the DATA_READER_RESOURCE_LIMITS QoSPolicy's **max_app_ack_response_length** longer than the maximum serializable data, which resulted in the truncation of data when the length got close to 64kB.

This problem has been resolved: *Connex* now enforces a maximum length of 32kB for **max_app_ack_response_length** as part of *DataReader* QoS consistency checks, and it will log an error if you try to set **max_app_ack_response_length** longer than 32kB.

[RTI Issue ID CORE-12450]

4.10.14 Error messages displayed that should not have been, when printing DataReaderQoS objects

When printing *DataReaderQoS* objects, and the contained *DDSOwnershipQoSPolicy* or *DDS_TransportMulticastQoSPolicy* policies were printed, some error messages were displayed that should not have been. These error messages could have been safely ignored by an application. These error messages are no longer printed.

[RTI Issue ID CORE-12462]

4.10.15 Potential Valgrind invalid read when logging a message or enabling heap monitoring

When activity context was enabled in logging, or when heap monitoring was enabled, a Valgrind invalid read similar to the following one may have been reported:

```
==1344490== Invalid read of size 4
==1344490== at 0x4A3FA0A: RTIOsapiActivityContext_skipResourceGuid (ActivityContext.c:246)
==1344490== by 0x4A417B3: RTIOsapiActivityContext_getString (ActivityContext.c:820)
```

This issue has been resolved. The Valgrind invalid read error no longer appears.

[RTI Issue ID CORE-12537]

4.10.16 Malformed IDL printed if multiple labels used for default case of a union

The IDL produced by the C API's **DDS_TypeCode_print_IDL()** function (or the equivalent in other APIs) may have been malformed if multiple labels were assigned to the default case of a union. All of the labels were printed as "default: ", instead of their true value. This problem has been resolved.

[RTI Issue ID CORE-12624]

Chapter 5 Known Issues

Note: For an updated list of critical known issues, see the Critical Issues List on the RTI Customer Portal at <https://support.rti.com>.

5.1 Known Issues with Usability

5.1.1 Cannot open USER_QOS_PROFILES.xml in rti_workspace/examples from Visual Studio

When trying to open the **USER_QOS_PROFILES.xml** file from the resource folder of one of the provided examples, you may see the following error:

```
Could not find file : C:\Users\\Documents\rti_workspace\5.3.0\examples\connect_dds\c\
```

The problem is that the Visual Studio project is looking for the file in a wrong location (win32 folder).

You can open the file manually from here:

```
C:\Users\\Documents\rti_workspace\5.3.0\examples\connect_dds\c\
```

This issue does not affect the functionality of the example.

[RTI Issue ID CODEGENII-743]

5.1.2 DataWriter's Listener callback on_application_acknowledgment() not triggered by late-joining DataReaders

The *DataWriter's* listener callback **on_application_acknowledgment()** may not be triggered by late-joining *DataReaders* for a sample after the sample has been application-level acknowledged by all live *DataReaders* (no late-joiners).

If your application requires acknowledgment of message receipt by late-joiners, use the Request/Reply communication pattern with an Acknowledgment type (see the chapter "

Introduction to the Request-Reply Communication Pattern," in the *RTI Connext Core Libraries User's Manual*).

[RTI Issue ID CORE-5181]

5.1.3 HighThroughput and AutoTuning built-in QoS Profiles may cause communication failure when writing small samples

If you inherit from either the **BuiltinQosLibExp::Generic.StrictReliable.HighThroughput** or the **BuiltinQosLibExp::Generic.AutoTuning** built-in QoS profiles, your *DataWriters* and *DataReaders* will fail to communicate if you are writing small samples.

In *Connext* 5.1.0, if you wrote samples that were smaller than 384 bytes, you would run into this problem. In version 5.2.0 onward, you might experience this problem when writing samples that are smaller than 120 bytes.

This communication failure is due to an interaction between the batching QoS settings in the **Generic.HighThroughput** profile and the *DataReader's* **max_samples** resource limit, set in the **BuiltinQosLibExp::Generic.StrictReliable** profile. The size of the batches that the *DataWriter* writes are limited to 30,720 bytes (see **max_data_bytes**). This means that if you are writing samples that are smaller than 30,720/**max_samples** bytes, each batch will have more than **max_samples** samples in it. The *DataReader* cannot handle a batch with more than **max_samples** samples and the batch will be dropped.

There are a number of ways to fix this problem, the most straightforward of which is to overwrite the *DataReader's* **max_samples** resource limit. In your own QoS profile, use a higher value that accommodates the number of samples that will be sent in each batch. (Simply divide 30,720 by the size of your samples).

[RTI Issue ID CORE-6411]

5.1.4 Memory leak if Foo:initialize() called twice

Calling **Foo:initialize()** more than once will cause a memory leak.

[RTI Issue ID CORE-7678]

5.1.5 Wrong error code after timeout on write() from Asynchronous Publisher

When using an asynchronous publisher, if **write()** times out, it will mistakenly return **DDS_RETCODE_ERROR** instead of the correct code, **DDS_RETCODE_TIMEOUT**.

[RTI Issue ID CORE-2016, Bug # 11362]

5.1.6 Type Consistency enforcement disabled for structs with more than 10000 members

TypeObjects cannot be created from structs with more than 10000 members. Applications that publish or subscribe to such types may see errors like the following:

```
RTICdrStream_serializeNonPrimitiveSequence:sequence length (10005) exceeds maximum (10000)
RTICdrTypeObjectTypeLibraryElement_getTypeId:serialization error: Type
RTICdrTypeObject_fillType:!get TypeId
RTICdrTypeObject_assertTypeFromTypeCode:!create Structure Type
RTICdrTypeObject_createFromTypeCode:!create TypeObject
```

When the TypeObject can't be serialized, the type compatibility check between a reader and a writer falls back to exact type-name matching.

See the section "Verifying Type Consistency: Type Assignability" in the *RTI Connex Core Libraries Extensible Types Guide* for more information.

[RTI Issue ID CORE-8158]

5.1.7 Escaping special characters in regular/filter expressions not supported in some cases

Escaping special characters is not supported in expressions when using the following features:

- Partitions
- MultiChannel

Every occurrence of a backslash ('\') will be considered its own character and not a way to escape the character that follows. For example: 'A\?' does not match 'A?' because the first expression is considered an expression with three characters.

[RTI Issue ID CORE-11858]

5.1.8 Calling take_snapshot within DDS Entity callbacks not supported

This release adds the **take_snapshot** APIs to take discovery snapshots to help debug discovery issues. In this release, calling this API within a DDS Entity callback (e.g., **on_subscription_matched**) is not supported. As a result, calling any of the **take_snapshot** APIs within a DDS Entity callback may result in the application's hanging or even crashing.

[RTI Issue ID CORE-12959]

5.1.9 banish and subject_name APIs are unresolved in Modern C++ Windows dynamic libraries

The Modern C++ APIs **banish_ignored_participants**, **discovered_participant_subject_name**, and **discovered_participants_from_subject_name** are unresolved symbols in the **nddscpp2** Windows dynamic libraries. If you attempt to use them, you will get "LNK2019 unresolved external symbol" errors. This problem was introduced in 7.0.0 and has not been fixed.

[RTI Issue ID CORE-13053]

5.2 Known Issues with Code Generation

5.2.1 Examples and generated code for Visual Studio 2017 and later may not compile (Error MSB8036)

The examples provided with *Connex* and the code generated for Visual Studio 2017 and later will not compile out of the box if the Windows SDK version installed is not a specific number like 10.0.15063.0. If that happens, you will see the compilation error MSB8036. To compile these projects, select an installed version of Windows SDK from the Project menu -> Retarget solution.

Another option is to set the environment variable `RTI_VS_WINDOWS_TARGET_PLATFORM_VERSION` to the SDK version number. For example, set `RTI_VS_WINDOWS_TARGET_PLATFORM_VERSION` to 10.0.16299.0. (Note: the environment variable will not work if you have already retargeted the project via the Project menu.)

For further details, see the Windows chapter of the *RTI Connex Core Libraries Platform Notes*.

[RTI Issue ID CODEGENII-800]

5.3 Known Issues with Instance Lifecycle

5.3.1 Instance does not transition to ALIVE when "live" DataWriter detected

The "Data Distribution Service for Real-time Systems" specification allows transitioning an instance from the `NO_WRITERS` state to the `ALIVE` state when a "live" *DataWriter* writing the instance is detected. Currently, this state transition is not supported in *Connex*. The only way to transition an instance from `NO_WRITERS` to `ALIVE` state is by receiving a sample for the instance from one of the *DataWriters* publishing it.

Example:

1. A *DataWriter* writes a particular instance. The *DataReader* receives the sample. The *DataWriter* loses liveliness with the *DataReader*, making the instance transition from `ALIVE` to `NO_WRITERS`. The writer later becomes alive again, but it doesn't resume writing samples of the instance. In this case, the instance will stay in a `NO_WRITERS` state.

2. The *DataWriter* publishes a new sample for the instance. Only then does the instance state change on the *DataReader* from NO_WRITERS to ALIVE.

[RTI Issue ID CORE-3018]

5.3.2 Persistence Service DataReaders ignore serialized key propagated with dispose updates

Persistence Service DataReaders ignore the serialized key propagated with dispose updates. *Persistence Service DataWriters* cannot propagate the serialized key with dispose, and therefore ignore the `serialize_key_with_dispose` setting on the *DataWriter* QoS.

[RTI Issue ID PERSISTENCE-221]

5.4 Known Issues with Reliability

5.4.1 DataReaders with different reliability kinds under Subscriber with GROUP_PRESENTATION_QOS may cause communication failure

Creating a *Subscriber* with `PresentationQosPolicy.access_scope` GROUP_PRESENTATION_QOS and then creating *DataReaders* with different `ReliabilityQosPolicy.kind` values creates the potential for situations in which those *DataReaders* will not receive any data.

One such situation is when the *DataReaders* are discovered as late-joiners. In this case, samples are never delivered to the *DataReaders*. A workaround for this issue is to set the `AvailabilityQosPolicy.max_data_availability_waiting_time` to a finite value for each *DataReader*.

[RTI Issue ID CORE-7284]

5.5 Known Issues with Content Filters and Query Conditions

5.5.1 Writer-side filtering may cause missed deadline

If you are using a *ContentFilteredTopic* and you set the *Deadline* QoSPolicy, the deadline may be missed due to filtering by a *DataWriter*.

[RTI Issue ID CORE-1634, Bug # 10765]

5.5.2 filter_sample_* statistics in DDS_DataWriterProtocolStatus not updated correctly

The `filter_sample_*` statistics in the `DDS_DataWriterProtocolStatus` are not updated correctly. The values that you get after calling the following APIs may be smaller than the actual values:

- `DDS_DataWriter::get_datawriter_protocol_status`
- `DDS_DataWriter::get_matched_subscription_datawriter_protocol_status`
- `DDS_DataWriter::get_matched_subscription_datawriter_protocol_status_by_locator`

[RTI Issue ID CORE-5157]

5.6 Known Issues with TopicQueries

5.6.1 TopicQueries not supported with DataWriters configured to use batching or Durable Writer History

Getting TopicQuery data from a *DataWriter* configured to use batching or Durable Writer History is not supported.

[RTI Issue IDs CORE-7405, CORE-7406]

5.7 Known Issues with Transports

5.7.1 AppAck messages cannot be greater than underlying transport message size

A *DataReader* with **acknowledgment_kind** (in the ReliabilityQosPolicy) set to `DDS_APPLICATION_AUTO_ACKNOWLEDGMENT_MODE` or `DDS_APPLICATION_EXPLICIT_ACKNOWLEDGMENT_MODE` cannot send AppAck messages greater than the underlying transport message size.

If a *DataReader* tries to send an AppAck message greater than the transport message size, *Connex* will print the following error message:

```
COMMENDFacade_sendAppAck:!add APP_ACK to MIG
COMMENDSrReaderService_sendAppAck:!send APP_ACK
PRESPsService_onReaderAppAckSendEvent:!send acknowledgment
```

To recover from the above error, the *DataReader* must acknowledge samples until the size of the AppAck message goes below the transport message size threshold.

Why does an AppAck message increase its size? An AppAck message contains a list of sequence number intervals where each interval represents a set of consecutive sequence numbers that have been already acknowledged. As long as samples are acknowledged in order, the AppAck message will always have a single interval. However, when samples are acknowledged out of order, the number of intervals and the size of the AppAck will increase.

For more information, see the "Application Acknowledgment" section in the *RTI Connex Core Libraries User's Manual*.

[RTI Issue ID CORE-5329]

5.7.2 DataReader cannot persist AppAck messages greater than 32767 bytes

A *DataReader* using durable reader state, whose **acknowledgment_kind** (in the *ReliabilityQosPolicy*) is set to `DDS_APPLICATION_AUTO_ACKNOWLEDGMENT_MODE` or `DDS_APPLICATION_EXPLICIT_ACKNOWLEDGMENT_MODE`, cannot persist an AppAck message greater than 32767 bytes.

To recover from the previous error, the *DataReader* must acknowledge samples until the size of the AppAck message goes below the transport message size threshold.

For more information, see the section "Durable Reader State," in the *RTI Connext Core Libraries User's Manual*.

[RTI Issue ID CORE-5360]

5.7.3 Discovery with Connext Micro fails when shared memory transport enabled

Given a *Connext* application with the shared memory transport enabled, a *Connext* Micro 2.4.x application will fail to discover it. This is due to a bug in *Connext* Micro that prevents a received participant discovery message from being correctly processed. This bug will be fixed in a future release of *Connext* Micro. As a workaround, you can disable the shared memory transport in the *Connext* application and use UDPv4 instead.

[RTI Issue ID EDDY-1615]

5.7.4 Communication may not be reestablished in some IP mobility scenarios

If you have two *Connext* applications in different nodes and they change their IP address at the same time, they may not reestablish communication. This situation may happen in the following scenario:

- The applications see each other only from one single network.
- The IP address change happens at the same time in the network interface cards (NICs) that are in the network that is in common for both applications.
- The IP address change on one of the nodes happens before the arrival of the DDS discovery message propagating the address change from the other side.

[RTI Issue ID CORE-8260]

5.7.5 Corrupted samples may be forwarded through Routing Service when using Zero-Copy transfer over shared memory

When using *Zero Copy transfer over shared memory* together with *RTI Routing Service*, *Routing Service* avoids an additional copy of the data by passing a reference to the sample from the input to the

output of a route. If the sample is reused and rewritten by the original application *DataWriter* during the time between when the sample was received on the route input and copied into the route output buffer, the forwarded sample will contain the updated, and now invalid, values for the original sample.

This situation can be avoided in a few different ways, with various tradeoffs.

5.7.5.1 Use automatic application acknowledgment

Using automatic application acknowledgment (**acknowledgment_mode** = APPLICATION_AUTO_ACKNOWLEDGMENT in the Reliability QoS Policy) between the *Routing Service* input *DataReader* and its matching *DataWriters* will avoid the issue.

When using Zero Copy transfer over shared memory, *DataWriters* must loan samples using the **get_loan** API. Only samples that have been fully acknowledged will be returned by the **get_loan** API. This means that if automatic application acknowledgment is turned on, that only samples that the *Routing Service* has already copied and written to the route output will be available for reuse by the original *DataWriter*, because *Routing Service* does not return the loan on a sample until after it is forwarded to the route outputs.

The drawback to this approach is that it requires RELIABLE Reliability. In addition, application-level acknowledgments are not supported in *Connex Micro*, so this approach will not work if *Connex Micro* is the source of the Zero Copy samples.

5.7.5.2 Ensure that the number of available samples accounts for Routing Service processing time

Regardless of whether you are using *Routing Service*, it is important when using Zero Copy transfer over shared memory to size your resources so that your application can continue to write at the desired rate while the receiving applications receive and process the samples. If you are using *Routing Service* and cannot, or do not wish to, use automatic application acknowledgments, you must take into account the amount of time it will take to receive and forward a sample when setting **writer_loaned_sample_allocation** in the DATA_WRITER_RESOURCE_LIMITS QoS Policy and managing the samples in your application.

[RTI Issue ID CORE-10782]

5.7.6 Network Capture does not support frames larger than 65535 bytes

Network capture does not support frames larger than 65535 bytes. This limitation affects the TCP transport protocol if the **message_size_max** property is set to a value larger than the default one.

[RTI Issue ID CORE-11083]

5.8 Known Issues with FlatData

5.8.1 FlatData language bindings do not support automatic initialization of arrays of primitive values to non-zero default values

RTI FlatData™ language bindings do not support the automatic initialization of arrays of primitive values to non-zero default values, unless the primitive is an enumeration. It is possible to declare an alias to a primitive member with a default value using the `@default` annotation, and then to declare an array of that alias. For example:

```
@default(10)
typedef int32 myLongAlias;

struct MyType {
    myLongAlias myLongArray[25];
};
```

The default values of each member of the array in this case should be 10, but in FlatData they will all be set to 0.

[RTI Issue ID CORE-9176]

5.8.2 Flat Data: `plain_cast` on types with 64-bit integers may cause undefined behavior

The function `rti::flat::plain_cast` is allowed on FlatData samples containing `int64_t` members, but those members are not guaranteed to have an 8-byte alignment (a 4-byte alignment is guaranteed). Memory checkers such as Valgrind may report errors when accessing such members from the pointer returned by `plain_cast`.

[RTI Issue ID CORE-10092]

5.8.3 FlatData in combination with payload encryption in RTI Security Plugins and/or compression will not save copies

RTI FlatData™ language binding offers a reduced number of end-to-end copies when sending a sample (from four to two), providing improved latency for large data samples. (See the "FlatData Language Binding" section in the *RTI Connex Core Libraries User's Manual*.) When used with payload encryption and/or payload compression, however, there are no savings in the number of copies. (See the section "Interactions with *RTI Security Plugins* and Compression" in the "Using FlatData Language Binding" section of the *RTI Connex Core Libraries User's Manual*.) In future releases, other copies currently being made can potentially be optimized out in order to reduce the number of copies when using FlatData in combination with security and compression.

[RTI Issue ID CORE-11262]

5.9 Known Issues with Coherent Sets

5.9.1 Some coherent sets may be lost or reported as incomplete with batching configurations

If *Connex* 6.1.0 receives coherent sets from *Connex* 6.0.0 or lower using batching, coherent sets that are fully received and complete may be lost or marked as incomplete. (If the QoS **subscriber_qos.-presentation.drop_incomplete_coherent_set** is set to FALSE, then the samples marked as incomplete won't be dropped.)

[RTI Issue ID CORE-9691]

5.9.2 Copy of `SampleInfo::coherent_set_info` field is not supported

`SampleInfo::coherent_set_info` is not available when using take/read operations that do not loan the samples. The `SampleInfo::coherent_set_info` is always set to NULL when you call the take/read operations that do not loan the samples. To get the `coherent_set_info` value, make sure you use the read/-take operations that loan the data.

In addition, the copy constructor and assignment operator in the Traditional C++ and Modern C++ APIs do not copy the `SampleInfo::coherent_set_info` field. It is always set to NULL. It is your responsibility to make the copy and handle memory allocation and deletion for this field.

[RTI Issue ID CORE-11215]

5.9.3 Other known issues with coherent sets

Coherent sets are not propagated through *RTI Routing Service* [RTI Issue ID ROUTING-657].

Group coherent sets are not supported with ODBC writer history [RTI Issue ID CORE-9746].

Group coherent sets are not persisted by *RTI Persistence Service* [RTI Issue ID PERSISTENCE-191].

Group coherent sets cannot be stored or replayed with *RTI Recording Service* [RTI Issue ID RECORD-1083].

5.10 Known Issues with Dynamic Data

5.10.1 Conversion of data by member-access primitives limited when converting to types that are not supported on all platforms

The conversion of data by member-access primitives (`get_X()` operations) is limited when converting to types that are not supported on all platforms. This limitation applies when converting to a 64-bit int64 type (`get_longlong()` and `get_ulonglong()` operations) and a 128-bit long double type (`get_longdouble()`). These methods will always work for data members that are actually of the correct type, but will only support conversion from values that are stored as smaller types on a subset of platforms.

Conversion to 64-bit `int64s` from a 32-bit or smaller integer type is supported on all Windows and Linux architectures, and any additional 64-bit architectures. Conversion to 128-bit long doubles from a float or double is not supported.

[RTI Issue ID CORE-2986]

5.10.2 Types that contain bit fields not supported

Types that contain bit fields are not supported by `DynamicData`. Therefore, when `rtiddsspy` discovers any type that contains a bit field, `rtiddsspy` will print this message:

```
DDS_DynamicDataTypeSupport_initialize:type not supported (bitfield member)
```

[RTI Issue ID CORE-3949]

5.11 Known Issues in RTI Monitoring Library

5.11.1 Problems with `NDDS_Transport_Support_set_builtin_transport_property()` if Participant Sends Monitoring Data

If a *Connex* application uses the `NDDS_Transport_Support_set_builtin_transport_property()` API (instead of the `PropertyQoSPolicy`) to set built-in transport properties, it will not work with *Monitoring Library* if the user participant is used for sending all the monitoring data (the default settings). As a workaround, you can configure *Monitoring Library* to use another participant to publish monitoring data (using the property name `rti.monitor.config.new_participant_domain_id` in the `PropertyQoSPolicy`).

[RTI Issue ID MONITOR-222]

5.11.2 Participant's CPU and memory statistics are per application

The CPU and memory usage statistics published in the *DomainParticipant* entity statistics topic are per application instead of per *DomainParticipant*.

[RTI Issue ID CORE-7972]

5.11.3 XML-based entity creation nominally incompatible with static monitoring library

If setting the *DomainParticipant* QoS programmatically in the application is not possible (i.e., when using XML-based Application Creation), the monitoring `create` function pointer may still be provided via an XML profile by using the environment variable expansion functionality. The monitoring property within the *DomainParticipant* QoS profile in XML must be set as follows:

```
<domain_participant_qos>
  <property>
    <value>
      <element>
```

```

        <name>rti.monitor.library</name>
        <value>tmonitoring</value>
    </element>
    <element>
        <name>rti.monitor.create_function_ptr</name>
        <value>$(MONITORFUNC)</value>
    </element>
</value>
</property>
</domain_participant_qos>

```

Then in the application, before retrieving the DomainParticipantFactory, the environment variable must be set programmatically as follows:

```

...
sprintf(varString, "MONITORFUNC=%p", RTIDefaultMonitor_create);
int retVal = putenv(varString);
...
//DomainParticipantFactory must be created after env. variable setting

```

[RTI Issue ID CORE-5540]

5.11.4 ResourceLimit channel_seq_max_length must not be changed

The default value of `DDS_DomainParticipantResourceLimitsQosPolicy::channel_seq_max_length` can't be modified if a *DomainParticipant* is being monitored. If this QoS value is modified from its default value of 32, the monitoring library will fail.

[RTI Issue ID MONITOR-220]

5.12 Known Issues with Installers

5.12.1 RTI Connex Micro 3.0.3 installation package currently compatible only with Connex 6.0.1 installer

ConnexMicro 3.0.3 must be installed with *ConnexProfessional* release 6.0.1. It cannot be installed with release 6.1.0. *ConnexMicro* 3.0.3 can communicate with either release. Customers licensing *ConnexMicro* will be notified when a *ConnexMicro* release that is compatible with the 6.1.0 installer is available.

5.13 Other Known Issues

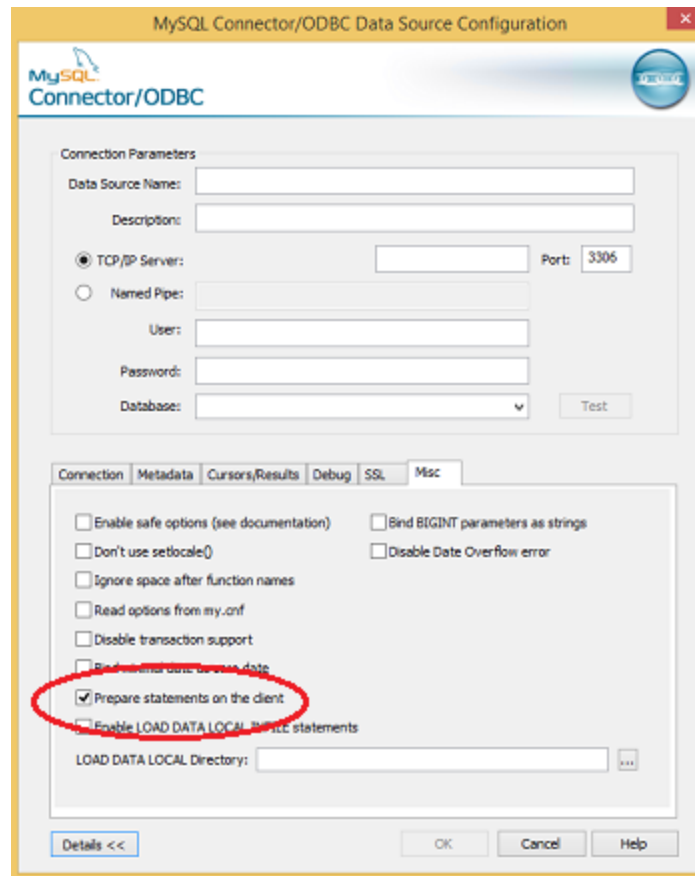
5.13.1 Possible Valgrind still-reachable leaks when loading dynamic libraries

If you load any dynamic libraries, you may see "still reachable" memory leaks in "dlopen" and "dlclose". These leaks are a result of a bug in Valgrind (<https://bug-s.launchpad.net/ubuntu/+source/valgrind/+bug/1160352>).

This issue affects the *Core Libraries*, *Security Plugins*, and *TLS Support*.

5.13.2 64-bit discriminator values greater than $(2^{31}-1)$ or smaller than (-2^{31}) supported only in Java,

[RTI Issue IDs CORE-9941, SEC-1026, and COREPLG-510]



5.13.2 64-bit discriminator values greater than $(2^{31}-1)$ or smaller than (-2^{31}) supported only in Java, no other languages

Unions with a 64-bit integer discriminator type containing discriminator values that cannot fit in a 32-bit value are not supported when using the following language bindings:

- C
- Traditional C++
- Modern C++
- New. NET
- DynamicData (regardless of the language)

They are also not supported with ContentFilteredTopics, regardless of the language binding.

Using label values greater than 32-bit may lead to receiving samples with invalid content or to filtering samples incorrectly.

For example, this is not supported:

```
union union_uint64 switch (uint64) {
    case 0x100000000:
        char m_char;
    case 0x200000000:
        int32 m_int32;
    case 0x300000000:
        string<5> m_string;
};
```

This is supported:

```
union union_uint64 switch (uint64) {
    case 1:
        char m_char;
    case 2:
        int32 m_int32;
    case 3:
        string<5> m_string;
};
```

[RTI Issue ID CORE-11437]

5.13.3 Creating multiple DataReaders for the same Topic under the same Subscriber configured with Group Ordered Access is not supported

Creating multiple *DataReaders* for the same *Topic* under the same *Subscriber* configured with PresentationQosPolicy `access_scope = GROUP` and `ordered_access = TRUE` is not supported. If you try to create a second reader in this situation, it will fail to be created and this error will be printed:

```
ERROR [0x0101E967,0x5C3A43B1,0x99D71EB7:0x80000309{Entity=Su,Domain=0}|CREATE DR WITH TOPIC FooTopic|LC:DISC]PRESpsService_createLocalEndpoint:NOT SUPPORTED | Creating more than one reader for the same topic within a single subscriber using GROUP presentation and ordered_access=true.
```

Instead, in this situation, you will need to use only one *DataReader*, or you will need to create a new *Subscriber* and *DataReader* in the same *DomainParticipant*.

[RTI Issue ID CORE-12448]

5.13.4 Segmentation fault when invoking Discovery Snapshot APIs in some cases

The `<Entity>::take_discovery_snapshot` APIs are not supported when invoked within a callback. Also, the `DataReader::take_discovery_snapshot` cannot be called on a *DataReader* using a *ContentFilteredTopic*. If you call the APIs in either of these ways, the application may produce a segmentation fault.

[RTI Issue ID CORE-13011]

Chapter 6 Experimental Features

This software may contain experimental features. These are used to evaluate potential new features and obtain customer feedback. They are not guaranteed to be consistent or supported and they should not be used in production.

In the API Reference HTML documentation, experimental APIs are marked with <<**experimental**>>.

The APIs for experimental features use the suffix **_exp** to distinguish them from other APIs. For example:

```
const DDS::TypeCode * DDS_DomainParticipant::get_typecode_exp(  
    const char * type_name);
```

Experimental features are also clearly noted as such in the *User's Manual* or *Getting Started Guide* for the component in which they are included.

Disclaimers:

- Experimental feature APIs may be only available in a subset of the supported languages and for a subset of the supported platforms.
- The names of experimental feature APIs will change if they become officially supported. At the very least, the suffix, **_exp**, will be removed.
- Experimental features may or may not appear in future product releases.
- Experimental features should not be used in production.

Please submit your comments and suggestions about experimental features to support@rti.com or via the RTI Customer Portal (<https://support.rti.com/>). Although the RTI Support team does not provide support for experimental features, you may be able to get help with experimental features from the RTI Community forum: <https://community.rti.com/>.

© 2022 Real-Time Innovations, Inc.
All rights reserved.
Printed in U.S.A. First printing.
September 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 solely under and subject to RTI's standard terms and conditions available at <https://www.rti.com/terms> and in accordance with your License Acknowledgement Certificate (LAC) and Maintenance and Support Certificate (MSC), except to the extent otherwise accepted in writing by a corporate officer of RTI.

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).

Notices

Deprecations and Removals

Any deprecations or removals noted in this document serve as notice under the Real-Time Innovations, Inc. Maintenance Policy #4220 and/or any other agreements by and between RTI and customer regarding maintenance and support of RTI's software.

Deprecated means that the item is still supported in the release, but will be removed in a future release. *Removed* means that the item is discontinued or no longer supported. By specifying that an item is deprecated in a release, RTI hereby provides customer notice that RTI reserves the right after one year from the date of such release and, with or without further notice, to immediately terminate maintenance (including without limitation, providing updates and upgrades) for the item, and no longer support the item, in a future release.

Early Access Software

“Real-Time Innovations, Inc. (“RTI”) licenses this Early Access release software (“Software”) to you subject to your agreement to all of the following conditions:

- (1) you may reproduce and execute the Software only for your internal business purposes, solely with other RTI software licensed to you by RTI under applicable agreements by and between you and RTI, and solely in a non-production environment;
- (2) you acknowledge that the Software has not gone through all of RTI’s standard commercial testing, and is not maintained by RTI’s support team;
- (3) the Software is provided to you on an “AS IS” basis, and RTI disclaims, to the maximum extent permitted by applicable law, all express and implied representations, warranties and guarantees, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, satisfactory quality, and non-infringement of third party rights;
- (4) any such suggestions or ideas you provide regarding the Software (collectively , “Feedback”), may be used and exploited in any and every way by RTI (including without limitation, by granting sub-licenses), on a non-exclusive, perpetual, irrevocable, transferable, and worldwide basis, without any compensation, without any obligation to report on such use, and without any other restriction or obligation to you; and
- (5) TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, IN NO EVENT WILL RTI BE LIABLE TO YOU FOR ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, EXEMPLARY OR PUNITIVE OR CONSEQUENTIAL DAMAGES OF ANY KIND, OR FOR LOST PROFITS, LOST DATA, LOST REPUTATION, OR COST OF COVER, REGARDLESS OF THE FORM OF ACTION WHETHER IN CONTRACT, TORT (INCLUDING WITHOUT LIMITATION, NEGLIGENCE), STRICT PRODUCT LIABILITY OR OTHERWISE, WHETHER ARISING OUT OF OR RELATING TO THE USE OR INABILITY TO USE THE SOFTWARE, EVEN IF RTI HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.”

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/>