

RTI Connex DDS

Core Libraries

Release Notes

Version 6.0.1



© 2019 Real-Time Innovations, Inc.
All rights reserved.
Printed in U.S.A. First printing.
November 2019.

Trademarks

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

Copy and Use Restrictions

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

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

The security features of this product include software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>).

Technical Support

Real-Time Innovations, Inc.

232 E. Java Drive

Sunnyvale, CA 94089

Phone: (408) 990-7444

Email: support@rti.com

Website: <https://support.rti.com/>

Contents

Chapter 1 Introduction	1
Chapter 2 System Requirements	
2.1 Supported Operating Systems	3
2.2 Requirements when Using Microsoft Visual Studio	5
2.3 Disk and Memory Usage	7
Chapter 3 Compatibility	
3.1 Wire Protocol Compatibility	8
3.2 Code and Configuration Compatibility	8
3.3 Extensible Types Compatibility	9
3.4 ODBC Database Compatibility	9
Chapter 4 What's Fixed in 6.0.1	
4.1 Fixes Related to Usability and Debuggability	10
4.1.1 Changed verbosity for "remote entity ignored by user" messages	10
4.1.2 Wrong GUID logged upon discovery sample rejection	10
4.1.3 DataWriter::set_listener() did not handle DDS_DATA_WRITER_INSTANCE_REPLACED status mask	10
4.1.4 Fixes to the printing of TypeCodes	11
4.1.5 Wrong return code or exception for DDS_DataWriter_get_matched_subscription_data and DDS_DataReader_get_matched_publication_data	11
4.1.6 Nanosecond field within source_timestamp was not validated in DataWriter::write_w_ params or DataWriter::write_w_timestamp	12
4.1.7 LogMessage's is_security_message flag did not work with formatted log messages	12
4.1.8 Using RTI DDS Spy from RTI Connexx Launcher did not print samples in some cases	12
4.1.9 Utilities summary printed NULL pointers, causing memory fault in some architectures	12
4.1.10 LogPrintFormat NDDS_CONFIG_LOG_PRINT_FORMAT_MINIMAL was not correct in Java API	13
4.1.11 Unexpected DDS_RETCODE_TIMEOUT write error when max instances was exceeded ..	13

4.2 Fixes Related to Transports	14
4.2.1 add_peer API returned RETCODE_DDS_OK when a hostname could not be resolved	14
4.2.2 Windows IP mobility changes not always detected	14
4.2.3 Shared memory locator ignored if it matched a multicast address	14
4.2.4 Malformed TRANSPORT_MULTICAST_MAPPING Qos caused a crash in the application	14
4.2.5 Network interface name not parsed properly on Windows	15
4.2.6 Errors parsing DTLS Transport properties were not reported	15
4.2.7 Slow or no communication when enabling TCP Transport's CRC	15
4.2.8 Potential rare crash upon TCP Transport shutdown	15
4.2.9 Potential rare hang during TCP Transport shutdown	16
4.3 Fixes Related to Reliability Protocol and Wire Representation	16
4.3.1 Invalid fragment size when using FlowController	16
4.3.2 ReliableReaderActivityChangedStatus may have had invalid content	16
4.3.3 Data fragmentation did not work for samples larger than 64k for certain transport configurations	17
4.3.4 DataReader reported incorrect sample lost and rejected when receiving coherent set	17
4.3.5 Inactivated non-progressing DataReaders were never sent data, preventing them from progressing	17
4.3.6 Unbounded memory growth on DataWriter when failing to process NACK_FRAGs from DataReaders	18
4.3.7 DataReaders may not receive samples from DataWriter when using FlowController	18
4.3.8 DataWriter RELIABLE_READER_ACTIVITY_CHANGED StatusMask not set when a matched DataReader was deleted	19
4.3.9 Invalid serialization of samples with types containing primitive members that require padding	19
4.3.10 Unexpected fragmentation in the ParticipantBuiltinTopicData DataWriter prevented discovery	20
4.4 Fixes Related to Content Filters	20
4.4.1 ContentFilteredTopic updates unnecessarily propagated when filter expression and parameters remained unchanged	20
4.4.2 Incorrect writer-side filtering by reliable DataWriters in some cases	21
4.5 Fixes Related to TopicQueries	21
4.5.1 DataReader::lookupInstance did not work with TopicQueries	21
4.5.2 Potential segmentation fault when using TopicQueries and enabling application-level acknowledgment	21
4.5.3 Precondition error when a DataReader issuing a TopicQuery was rematched	22
4.5.4 DataWriter did not dispatch TopicQueries for DataReaders that were inactive when the TopicQuery was received	22
4.6 Fixes Related to Dynamic Data and FlatData	23
4.6.1 FlatData: bad performance of PrimitiveSequenceBuilder::add_n	23
4.6.2 Memory leak while receiving encoded reassembled samples with Dynamic Data or FlatData language binding	23

4.6.3	Memory corruption when using hierarchical names to access members of sequences that were loaned from the DataReader	23
4.6.4	FlatData: possible compilation error building data for a type with a sequence	24
4.6.5	Accessing members using hierarchical member names may have failed if member names were subsets of each other	24
4.6.6	Unbounded memory growth on volatile DataWriter while publishing FlatData samples without a matched DataReader	24
4.6.7	Segmentation fault when using an unkeyed DynamicData DataReader with a content filter and writer-side filtering	25
4.6.8	Binding to an unset optional member caused some operations on the parent DynamicData object to fail	25
4.6.9	Possible data corruption or crash when using DynamicData and a type with inheritance	25
4.6.10	Missing parameter checking for several DynamicData APIs resulted in segmentation faults or incorrect return codes	26
4.6.11	Flat Data: Conversion warning on some platforms	27
4.7	Fixes Related to Modern C++ API	27
4.7.1	Functions <code>get_discovered_participants</code> and <code>get_discovered_participant_data</code> were in incorrect namespace	27
4.7.2	Possible compilation error due to clash with POSIX macro	27
4.7.3	Inaccessible <code>extensibility_kind</code> getter in <code>EnumType</code>	27
4.7.4	TopicQuery created with filter parameters leaked memory	27
4.7.5	Entities created from XML configuration were not deleted with <code>close()</code>	28
4.7.6	Closing a DataReader with TopicQueries may fail in some circumstances	28
4.7.7	Possible strict-aliasing warnings on some platforms	28
4.7.8	<code>dds::core::polymorphic_cast</code> didn't accept rvalues as its argument	29
4.7.9	Rare memory leak in DynamicType API	29
4.7.10	Possible leak creating a Replier with a listener	29
4.8	Fixes Related to Traditional C++ API	29
4.8.1	Extra semicolon warning when compiling with pedantic flag	29
4.8.2	Performance degradation in <code>write()</code> API when IDL file contained types annotated for Zero Copy transfer over shared memory	30
4.8.3	<code>read/take_replies</code> functions ignored <code>max_samples</code> argument	30
4.9	Fixes Related to XML Configuration	30
4.9.1	XML configuration of FlowController properties did not accept <code>DDS_LENGTH_UNLIMITED</code>	30
4.9.2	Elements within <code><domain_participant></code> were forced to have an order	30
4.9.3	Schema files were not compliant with DDS-XML spec	31
4.9.4	QoS policies not resolved correctly when inheriting from a <code><qos_profile></code> , which contained an <code><xxx_qos></code> specifying its own <code>base_name</code> attribute	32
4.10	Fixes Related to XML-Based Application Creation	32

4.10.1	Attribute base_name not allowed in <domain_participant> and <domain>	32
4.10.2	XSD validation failed if topic_ref or register_type_ref used some non-alphanumeric characters	33
4.11	Fixes Related to Micro Application Generator (MAG)	33
4.11.1	MAG failed to calculate the value of an Entity when inheriting from a profile in a different library	33
4.11.2	Generated code didn't compile when using data types in namespaces	34
4.11.3	Generated code didn't compile if publisher_name or subscriber_name not specified	34
4.11.4	Invalid generated code when the name value contained an invalid variable character	35
4.11.5	Incorrect QoS for DataWriters and DataReaders when inheriting from a base profile using a topic_ref	35
4.11.6	MAG failed when XML contained two topics with same name	36
4.11.7	MAG failed when base_name attribute was specified within <domain>	36
4.11.8	MAG failed to process XML file if a duration element contained DURATION_ZERO_SEC or DURATION_ZERO_NSEC	36
4.11.9	Command-line option -inputXml was not documented	37
4.11.10	MAG failed to generate code when an entity inherited from an empty entity or an entity whose values were the same as defaults used by Connex DDS Micro	37
4.11.11	Invalid QoS values for DataWriters and DataReaders when trying to inherit values using a topic_filter	37
4.11.12	NullPointerException from MAG when initial_peers not specified and builtin_transport was specified	38
4.11.13	MAG did not use same defaults as Connex DDS Micro for some values	38
4.11.14	Generated code did not compile when participant using dynamic participant, static endpoint (DPSE) discovery contained multiple remote publication and subscription elements	39
4.11.15	MAG failed to process XML file if thread mask contained more than one element	40
4.11.16	MAG failed to report error when XML file contained invalid use of LENGTH_UNLIMITED	41
4.11.17	XSD validation failed if topic_ref or register_type_ref used some non-alphanumeric characters	41
4.11.18	MAG calculated invalid values for QoS profile specified with -outputFinalQoS option	42
4.11.19	Invalid behavior when calculating final values for QoS profile specified with -outputFinalQoS option	42
4.12	Fixes Related to Monitoring Library	42
4.12.1	Unbounded generation of file handles on QNX platforms	42
4.12.2	Possible segmentation fault	42
4.12.3	Memory limit reached	43
4.12.4	Crash when an IDL type used FlatData or the XCDR2 data representation	43
4.13	Fixes Related to Vulnerabilities	43
4.14	Other Fixes	43
4.14.1	Crash when deserialized_type_object_dynamic_allocation_threshold set to 0	43
4.14.2	TypeCodeFactory and types XML parser allowed derived structure with different extensibility than its base structure	44

4.14.3	Json unicode escape sequences were not supported	44
4.14.4	Error in DomainParticipant creation on Windows platforms	44
4.14.5	Crash while accessing Subscription built-in topic data of a Zero Copy DataReader in Java	44
4.14.6	Connex DDS Micro Subscriber did not communicate with Connex DDS Professional Publisher using Zero Copy Transfer Mode	44
4.14.7	Full path name of Connex libraries was needed when using FindRTIConnexDDS.cmake	45
4.14.8	dds.sys_info.process_id property in DomainParticipantQos had negative value on some VxWorks platforms in kernel mode	45
Chapter 5 What's Fixed in 6.0.0		
5.1	Fixes Related to Discovery	46
5.1.1	Endpoint discovery initialization errors during participant creation left participant in inconsistent state	46
5.1.2	Writer/reader resource limits affected wrong builtin endpoints	46
5.1.3	Failure to send TypeObject when type had base with no members	47
5.1.4	Received incorrect QoS policy values through discovery when communicating with other vendors	47
5.2	Fixes Related to Reliability Protocol	48
5.2.1	VOLATILE DataReader may have received historical samples from DataWriter	48
5.2.2	Wrong RTPS GAP messages emitted by reliable DataWriters in some cases	48
5.2.3	Individual sample fragment losses may have triggered full sample repair	49
5.2.4	DataReader may have ignored some piggyback heartbeats, leading to performance degradation	49
5.2.5	max_bytes_per_nack_response was ignored when using asynchronous publication	49
5.2.6	Disabling positive ACKs may have caused segmentation fault in publishing application	49
5.2.7	Unexpected DDS_RETCODE_OUT_OF_RESOURCES error when writing a sample	49
5.3	Fixes Related to Instance Management and Lifecycle	50
5.3.1	Instances may not have transitioned to NOT_ALIVE_NO_WRITERS on DataReader matching with a DataWriter using MultiChannelQosPolicy	50
5.3.2	DomainParticipantFactory::get_instance always returned success	50
5.3.3	DataReaders may not have purged samples from instances in NOT_ALIVE_NO_WRITERS state when autopurge_nowriter_samples_delay was set to finite value	50
5.3.4	Instances in NOT_ALIVE_DISPOSED state may not have been purged from DataReader queue when autopurge_disposed_instances_delay was set to zero	51
5.3.5	A DataReader may have failed to calculate the keyhash for a sample containing zero-length strings	51
5.3.6	Incorrect warning reported while trying to purge disposed instances proactively	52
5.3.7	Purging disposed or unregistered instances based on source timestamp does not work when internal clock is set to monotonic	52
5.3.8	Possible leak upon application exit after using NDDSConfigVersion::get_instance() (Traditional C++ API only)	52
5.3.9	Unexpected errors when receiving a dispose sample for unbounded DDS_KeyedString BuiltinType topic	52

5.3.10	instance_replacement not applied correctly for durable DataWriters	53
5.3.11	Incorrect warning reported when replacing DISPOSE/ALIVE instance on a DataWriter	53
5.4	Fixes Related to Content Filters and Query Conditions	53
5.4.1	Possible crash in creation of a content filter for a type with an aliased base type	53
5.4.2	Reading samples by instance with QueryCondition returned no data when using TOPIC or GROUP PresentationQosPolicy access_scope	54
5.4.3	Content Filter issue when filtering on a member of the base type, for types with inheritance	54
5.5	Fixes Related to TopicQueries	54
5.5.1	DataWriter may have deadlocked if receiving "continuous" TopicQueries	54
5.5.2	Instances may not have transitioned to NOT_ALIVE_NO_WRITERS on DataReader matching with a DataWriter with TopicQuery enabled	54
5.5.3	Error when changing partition, group data, or topic data on a Publisher containing a DataWriter with TopicQuery enabled	55
5.5.4	Possible increasing memory and CPU usage in publishing applications using TopicQueries	55
5.5.5	Spurious log message related to TopicQuery has been removed	55
5.5.6	DataReader::getKey did not work with TopicQueries	56
5.5.7	TopicQuery samples that failed to be written a first time may have never been sent	56
5.6	Fixes Related to DynamicData	56
5.6.1	DynamicData had limitations with members larger than 65,535 bytes	56
5.6.2	Copying into a bound DynamicData object did not work	57
5.6.3	Corrupt DynamicData objects containing sequences with length 0	57
5.6.4	Binding to members of a sequence incorrectly created members in the DynamicData API	57
5.6.5	Error when unbinding from a union DynamicData object	57
5.6.6	Performance degradation when setting large sequences using the DynamicData API	58
5.6.7	DynamicData::is_member_key did not work for types using inheritance	58
5.6.8	DynamicData::set_complex_member did not work with aliased typecodes	58
5.6.9	DynamicData::set_string API did not accept NULL strings	58
5.6.10	Large memory allocation when binding to large sequences in the DynamicData API	59
5.6.11	DynamicData::from_cdr_buffer API did not resize DynamicData object	59
5.6.12	Error when unbinding from a DynamicData object	59
5.6.13	Accessing a member of an array or sequence by member ID failed for member IDs > 65535	59
5.6.14	DynamicData DataReader may have failed to correctly deserialize key for types containing strings ..	60
5.6.15	Missing DynamicData::clear_optional_member API in the .NET API	60
5.6.16	DynamicData::clear_optional_member API incorrectly returned error for unset optional members ..	60
5.6.17	DynamicData APIs did not check for an associated TypeCode	60
5.6.18	Setting members past a sequence's maximum bound was not prohibited in the DynamicData API ..	61
5.6.19	DynamicData::clear_all_members API did not work on bound DynamicData objects	61

5.7 Fixes Related to Transports	61
5.7.1 Communication between kernel and RTP Mode Participants with shared memory transport not working on 64-bit VxWorks 6 platforms	61
5.7.2 Network interface tracker may have reported non-existing changes	61
5.7.3 Possible continuous failure to send over shared memory transport	62
5.7.4 Race condition in shared memory transport led to cleanup failure	62
5.7.5 UDPv4/UDPv6 transport creation failed when setting send_socket_buffer_size or recv_socket_buffer_size to NDDS_TRANSPORT_UDPV4_SOCKET_BUFFER_SIZE_OS_DEFAULT or NDDS_TRANSPORT_UDPV6_SOCKET_BUFFER_SIZE_OS_DEFAULT	62
5.7.6 Unexpected shared memory locator when setting rtps_host_id to a value different than DDS_RTPS_AUTO_ID	63
5.7.7 Crash when network interface changed before transport was fully created	63
5.7.8 Wrong transport class name when logging 'No interfaces' warning	63
5.8 Fixes Related to Logging and Distributed Logger	64
5.8.1 Segmentation fault when simultaneously changing log files and writing to a log file	64
5.8.2 Unexpected "PRESCstReaderCollator_addInstanceEntry:exceeded max total instances" message at WARNING level	64
5.8.3 Segmentation fault when attempting to write to one file of a file set if that file failed to be opened ..	64
5.8.4 NDDS_Config_Logger_get_output_device() always returned NULL after DomainParticipant creation	64
5.8.5 Unable to configure Distributed Logger profile Using C++ DomainParticipantFactory	65
5.8.6 Possible segmentation fault when deleting options that were used to create Distributed Logger instance	65
5.8.7 Error code for setApplicationKind not captured in Distributed Logger C example	65
5.8.8 Logger settings from QoS File not set correctly when using Distributed Logger	66
5.8.9 "Max queue size reached, message will get lost" warning removed from DistributedLogger	66
5.9 Fixes Related to XML Configuration	66
5.9.1 Some XML example files were not compliant with target XSD schema	66
5.9.2 NDDS_QOS_PROFILES.xml not loaded from default location	66
5.9.3 is_default_participant_factory_profile="true" was ignored when <participant_factory_qos> only had the base_name attribute specified	67
5.9.4 Creating a typecode from an XML with a directive tag may have caused a crash	67
5.9.5 XML parser error if <domain_library> or <domain_participant_library> split into multiple tags with same name attribute	67
5.9.6 Connxt DDS XML parser failed to parse a const char with '\0' value	67
5.9.7 Value checks were not enforced for writer_qos.writer_resource_limits.instance_replacement field in DataWriterQos	68
5.9.8 Last, not first, value was retrieved when getting a QoS by specifying topic_filter	68
5.9.9 Unable to set WriterDataLifecycle::autopurge_disposed_instances_delay for builtin DataWriters via XML	69

5.10 Fixes Related to XML-Based Application Creation	69
5.10.1 Sequences defined in XML with sequenceMaxLength of -1 were not interpreted as unbounded	69
5.10.2 XML-Based Application Creation in Java didn't work with some TypeCodes	69
5.11 Fixes Related to OMG Specification Compliance	69
5.11.1 Connexst DDS not compliant with RTPS 2.2	69
5.11.2 Default GUID not compliant with RTPS specification	70
5.11.3 Manual by Participant Liveliness stopped working when communicating with old RTI Connexst Professional versions or non-RTI DDS implementations	70
5.11.4 NACK_FRAG message not compliant with RTPS specification	71
5.12 Fixes Related to Vulnerabilities	71
5.13 Fixes Related to Modern C++ API	71
5.13.1 Compilation failure when NDDS_USER_DLL_EXPORT defined on non-Windows platform	71
5.13.2 Downcasting a condition into a ReadCondition not supported	71
5.13.3 ReadCondition handler may not have been dispatched	71
5.13.4 Compilation error accessing a const LoanedSamples instance	72
5.13.5 Missing unregister_thread function	72
5.13.6 dds::sub::Sample creation or assignment from a LoanedSample failed when data was invalid	72
5.13.7 Possible symbol collision with Boost compiling the Modern C++ API	73
5.13.8 dds::topic::find could not return AnyTopic	73
5.13.9 Inconsistency between XML format and QosProvider::type() arguments	73
5.13.10 Missing accessors for Liveliness:: assertions_per_lease_duration	73
5.13.11 Entity Listeners may have missed some notifications	73
5.14 Other Fixes	74
5.14.1 strict-aliasing warnings when compiling code generated from an IDL with floats or enumerations ..	74
5.14.2 Unexpected sample loss notification on non-VOLATILE DataReader	74
5.14.3 RTPS messages with wrong alignment incorrectly accepted	74
5.14.4 Segmentation fault while looking up vendor-specific topics from a traditional C++ or .NET application	75
5.14.5 Memory leak when failing to create builtin endpoints	75
5.14.6 Segmentation Fault when Simultaneously Removing Remote Participant and Sending a Message to a Third Participant	75
5.14.7 Could not add property names that were prefixes of existing property names	75
5.14.8 Unbounded memory growth on DataWriter when enable_required_subscriptions set to true and DataReaders setting role name were created/destroyed continuously	76
5.14.9 Possible crash when applications used two extensible types that differed by one aliased primitive member	76
5.14.10 Errors reported when compiling a file including disc_rtps_impl.h with gcc -C	77
5.14.11 Unlikely segmentation fault when deleting a DataReader and using GROUP ordered access	77

5.14.12	RTIOSapiProcess_getId() returns incorrect PID in VxWorks Kernel Mode	77
5.14.13	Potential segmentation fault while unregistering a logger output device	77
5.14.14	Incorrect value for type_consistency in SubscriptionBuiltinTopicData in .NET API	78
5.14.15	DomainParticipant creation did not fail when using an unknown flow controller	78
5.14.16	Removed recursion in include files	78
5.14.17	max_blocking_time not honored for KEEP_LAST DataWriters	78
5.14.18	Use of -qosProfile argument in DDS Ping did not take topic_filter into consideration	79
5.14.19	Use of -qosProfile argument in DDS Spy did not take topic_filter into consideration	79
5.14.20	Rare race condition may have caused a keep-last DataWriter to time out in write()	79
5.14.21	DDS_PublisherQos_copy() in C API made a shallow copy	80
5.14.22	Unexpected sample losses with reason DDS_LOST_BY_REMOTE_WRITERS_PER_SAMPLE_LIMIT when using Group-Ordered access	80
5.14.23	Unexpected "!get remote writer queue" warning	80
5.14.24	DataReader may have incorrectly reported on_sample_lost() with a negative total_count	81
5.14.25	Unexpected warning printed when a Participant ignored itself	81
5.14.26	Potential crash when setting new_participant_domain_id in Java monitoring libraries	81
5.14.27	TypeCode.print_complete_IDL() failed if a type inherited through an alias	81
5.14.28	Possible error "Too many open files" specifying the discovery peers by a file	81
5.14.29	Error receiving batches on a DataWriter from a DataReader with a different endianness	82
5.14.30	Performance degradation when DataWriters sent HeartbeatFrag RTPS messages	82
5.14.31	Unexpected COMMENDSrReaderService_onSubmessage:!add NACK_FRAG error message	82
5.14.32	Visual Studio run-times copied to .NET project directory	82

Chapter 6 Known Issues

6.1	AppAck Messages Cannot be Greater than Underlying Transport Message Size	84
6.2	Cannot Open USER_QOS_PROFILES.xml in rti_workspace/examples from Visual Studio	84
6.3	DataReader Cannot Persist AppAck Messages Greater Than 32767 Bytes	85
6.4	DataReaders with Different Reliability Kinds Under Subscriber with GROUP_PRESENTATION_QOS may Cause Communication Failure	85
6.5	DataWriter's Listener Callback on_application_acknowledgment() not Triggered by Late-Joining DataReaders	86
6.6	Discovery with Connxt DDS Micro Fails when Shared Memory Transport Enabled	86
6.7	Examples and Generated Code for Visual Studio 2017 and later may not Compile (Error MSB8036)	86
6.8	HighThroughput and AutoTuning built-in QoS profiles may cause Communication Failure when Writing Small Samples	87
6.9	Memory Leak if Foo::initialize() Called Twice	87
6.10	Shared Memory Communication Requires Setting dds.transport.shmem.builtin.hostid in Transport Mobility Scenarios	87
6.11	TopicQueries not Supported with DataWriters Configured to Use Batching or Durable Writer History	88

6.12 Uninstalling on AIX Systems	88
6.13 Writer-Side Filtering May Cause Missed Deadline	88
6.14 Wrong Error Code After Timeout on write() from Asynchronous Publisher	88
6.15 Instance does not Transition to ALIVE when "live" DataWriter Detected	88
6.16 Communication may not be Reestablished in Some IP Mobility Scenarios	89
6.17 DomainParticipantFactoryQos in XML may not be Loaded	89
6.18 Multichannel and TopicQuery do not Work with Large Data in Some Cases	90
6.19 FlatData language bindings do not support automatic initialization of arrays of primitive values to non-zero default values	90
6.20 Potential application crash when receiving a sample on the service request channel	91
6.21 Known Issues with Dynamic Data	91
6.21.1 Conversion of data by member-access primitives limited when converting to types that are not supported on all platforms	91
6.21.2 Types that contain bit fields not supported	91
6.21.3 DynamicData language bindings do not support automatic initialization of arrays of primitive values to non-zero default values	91
6.22 Known Issues in RTI Monitoring Library	92
6.22.1 Problems with NDDS_Transport_Support_set_builtin_transport_property() if Participant Sends Monitoring Data	92
6.22.2 Participant's CPU and Memory Statistics are Per Application	92
6.22.3 XML-Based Entity Creation Nominally Incompatible with Static Monitoring Library	92
6.22.4 ResourceLimit channel_seq_max_length must not be Changed	93
Chapter 7 Experimental Features	94

Chapter 1 Introduction

RTI® Connex® DDS 6.0.1 is a maintenance release based on the feature release 6.0.0. This document describes fixes in 6.0.1. These enhancements have been made since 6.0.0.

This document includes the following:

- [System Requirements \(Chapter 2 on page 3\)](#)
- [Compatibility \(Chapter 3 on page 8\)](#)
- [What's Fixed in 6.0.1 \(Chapter 4 on page 10\)](#)
- [What's Fixed in 6.0.0 \(Chapter 5 on page 46\)](#)
- [Known Issues \(Chapter 6 on page 84\)](#)
- [Experimental Features \(Chapter 7 on page 94\)](#)

For an overview of new features in 6.0.1, see [RTI Connex DDS Core Libraries Whats New in 6.0.1](#).

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

- **Use the RTI Customer Portal** (<http://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 *Connex DDS*, 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 Supported Operating Systems

Connex DDS 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 DDS* 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 DDS* application for any architecture. You will also need a target installation, which provides the libraries required to build a *Connex DDS* application for that particular target architecture.

Connex DDS is available for the platforms in the following table.

Table 2.1 Supported Platforms

Operating System	Version
AIX®	AIX 7.1
Android™ (target only)	Android 5.0, 5.1, 9.0 (5.0 and 5.1 only available on request.)
INTEGRITY® (target only)	INTEGRITY 10.0.2, 11.0.4, 11.4.4 (10.0.2 and 11.0.4 only available on request.)
iOS® (target only)	iOS 8.2 (Only available on request.)
Linux® (Arm® CPU)	NI™ Linux 3 Raspbian Wheezy 7.0 Ubuntu® 16.04 LTS

Table 2.1 Supported Platforms

Operating System	Version
Linux (Intel® CPU)	CentOS™ 6.0, 6.2 - 6.4, 7.0 Red Hat® Enterprise Linux 6.0 - 6.5, 6.7, 6.8, 7.0, 7.3, 7.5, 8.0 Ubuntu 12.04 LTS, 14.04 LTS, 16.04 LTS, 18.04 LTS SUSE® Linux Enterprise Server 11 SP2, 11 SP3, 12 SP2 Wind River® Linux 7
LynxOS® (target only)	LynxOS 5.0
macOS®	macOS 10.12 - 10.14 (10.12 only available on request.)
QNX® (target only)	QNX Neutrino® 6.4.1, 6.5, 7.0 (6.4.1 and 6.5 only available on request.)
Solaris™	Solaris 2.10 (Only available on request.)
VxWorks® (target only)	VxWorks 6.9, 6.9.3.2, 6.9.4.2, 6.9.4.6, 7.0 (6.9 only available on request.) VxWorks 653 2.3 (Only available on request.)
Windows®	Windows 8, 8.1, 10 ^a Windows Server 2012 R2 Windows Server 2016

The following table lists additional target libraries for which RTI offers custom support. If you are interested in using one of these platforms, please contact your local RTI sales representative or email sales@rti.com.

Table 2.2 Custom Supported Platforms

Operating System	Version
INTEGRITY	INTEGRITY 5.0.11 on MPC 8349 CPU
INtime® for Windows	INtime 6.3 with Visual Studio 2017 on x86 CPU ^b

^aPer Microsoft, this should be compatible with Windows 10 IoT Enterprise with Windows native app.

^bTested on 64-bit Windows 10 operating system.

Table 2.2 Custom Supported Platforms

Operating System	Version
Linux	Debian® 7 on Arm v7 CPU
	Freescale™ 1.4 on QorIQ or P4040/P4080/P4081 CPU
	Red Hat Enterprise Linux 5.2 on x86 CPU
	RedHawk™ Linux 6.0 on x64 CPU
	RedHawk Linux 6.5 on x86 and x64 CPUs
	Wind River Linux 7 on Arm v7 CPU
	Wind River Linux 8 on PPC e6500 and Arm v7 CPUs
	Xilinx® 14.2 on Arm v7 CPU
	Yocto Project® 2.2 on 64-bit Arm v8 CPU Yocto Project 2.5 on Arm v7 CPU
QNX	QNX 6.5 on PPC e500v2 CPU
	QNX 6.6 on Arm v7 and x86 CPUs
	QNX 7.0 on Arm v7 CPU
VxWorks	VxWorks 6.9.3 on Arm v7 and PPC CPUs
	VxWorks 6.9.3.2 on MIPS CPU
	VxWorks 6.9.4.11 on Arm v7 CPU
	VxWorks 653 2.5.0.2 on PPC e500v2 CPU
	VxWorks 7.0 SR0540 on x86 CPU

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

2.2 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 2012 — Update 4 Redistributable Package Requirement

You must have the Visual C++ Redistributable for Visual Studio 2012 Update 4 installed on the machine where you are *running* an application linked with dynamic libraries. This includes dynamically linked C/C++ and all .NET and Java applications.

You can download Visual C++ Redistributable for Visual Studio 2012 Update 4 from this Microsoft website: <http://www.microsoft.com/en-ca/download/details.aspx?id=30679>

When Using Visual Studio 2013 — Redistributable Package Requirement

You must have Visual C++ Redistributable for Visual Studio 2013 installed on the machine where you are *running* an application linked with dynamic libraries. This includes C/C++ dynamically linked and all

.NET and Java applications.

You can download Visual C++ Redistributable for Visual Studio 2013 from this Microsoft website:

<https://www.microsoft.com/en-us/download/details.aspx?id=40784>

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 .NET and 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 .NET and 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

1. 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 .NET and 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".

2. *rtiddsgen* generated projects

Because Visual Studio 2019 is binary compatible with Visual Studio 2017, the Visual Studio 2017 libraries are used with Visual Studio 2019. There are two significant differences between the Visual Studio 2017 and Visual Studio 2019 libraries.

a. Platform Tools Set

The default Platform Tools set for Visual Studio 2017 is v141. The default Platform Tools set for Visual Studio 2019 is v142.

There are two ways to upgrade the Platform Tools set. Either open the solution file with Visual Studio and follow the upgrade prompts or on the Developer Studio command line, run: **devenv /upgrade <solutionFileName>**

b. Windows Target Platform Version

The default Windows Target Platform Version for Visual Studio 2017 is 8.1. Visual Studio 2017 requires an explicit number such as 10.0.17763.0 while Visual Studio 2019 will accept a more generalized version, such as 10.0.

There are two major ways to change the Windows Target Platform Version.

The first way is to open the solution file in Visual Studio 2019, right-click on the solution in the ‘Solution Explorer,’ then select the **retarget** option in the menu.

The second way is to edit the project file(s) with any text editor, including ‘**devenv /edit**’, and change the tag **<WindowsTargetPlatformVersion>xx.x.xxxxx.x</WindowsTargetPlatformVersion>** to match the version of your SDK.

RTI has introduced a workaround to the SDK versioning to simplify the process. Projects are generated with the target platform version using an environment variable like this:

<WindowsTargetPlatformVersion>\$(RTI_VS_WINDOWS_TARGET_PLATFORM_VERSION)</WindowsTargetPlatformVersion>

Simply define **RTI_VS_WINDOWS_TARGET_PLATFORM_VERSION** to the version you have installed. One simple way you can do that is to execute this command:

```
set RTI_VS_WINDOWS_TARGET_PLATFORM_VERSION=%WindowsSDKVersion:~0,-1%
```

A 3rd way, which does not modify the Visual Studio 2017 projects but allows you to build with Visual Studio 2019, is to use **msbuild** to override the project settings, like this:

```
msbuild /p:PlatformToolset=v142 /p:RTI_VS_WINDOWS_TARGET_PLATFORM_VERSION=%WindowsSDKVersion:~0,-1% <solutionFileName>
```

2.3 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 6.0.1 and previous releases, see the *Migration Guide* on the RTI Community Portal (<https://community.rti.com/documentation>).

3.1 Wire Protocol Compatibility

Connex DDS 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 current version is 2.3. RTI plans to maintain interoperability between middleware versions based on RTPS 2.x.

3.2 Code and Configuration Compatibility

The *Connex DDS* 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 DDS* 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 DDS* includes partial support for the "Extensible and Dynamic Topic Types for DDS" (DDS-XTypes) specification¹ from the Object Management Group (OMG), version 1.2. 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 DDS Core Libraries Getting Started Guide Addendum for Extensible Types* for a full list of the supported and unsupported extensible types features.

3.4 ODBC Database Compatibility

To use the Durable Writer History and Durable Reader State features, you must install a relational database such as MySQL.

To see if a specific architecture has been tested with the Durable Writer History and Durable Reader State features, see the *RTI Connex DDS Core Libraries Platform Notes*. To see what databases are supported, see the *RTI Connex DDS Core Libraries Getting Started Guide Addendum for Database Setup*.

¹ See the specification at <http://www.omg.org/spec/DDS-XTypes/>.

Chapter 4 What's Fixed in 6.0.1

Release 6.0.1 is a maintenance release based on the feature release 6.0.0. This section describes bugs fixed in 6.0.1. These fixes have been made since 6.0.0.

4.1 Fixes Related to Usability and Debuggability

4.1.1 Changed verbosity for "remote entity ignored by user" messages

The verbosity of "remote entity ignored by user" warning messages has been changed from WARNING to STATUS_REMOTE. Since these messages inform you of expected behavior, they should not be warnings.

[RTI Issue ID CORE-7197]

4.1.2 Wrong GUID logged upon discovery sample rejection

When a discovery sample is rejected, *Connex DDS* logs a warning message similar to the following:

```
DISCSimpleEndpointDiscoveryPlugin_subscriptionOnSampleRejected: 4c7  
{101E1C6,AF1B4A43,CD78BE4,80000104}; reason 1, total 1, delta 1
```

Previously, the GUID listed in that message was not correctly populated (it either was wrong or set to all zeros).

This issue is now resolved. Now the GUID is logged with the correct value.

[RTI Issue ID CORE-7448]

4.1.3 `DataWriter::set_listener()` did not handle `DDS_DATA_WRITER_INSTANCE_REPLACED` status mask

The specified callback was not called when set for the `DDS_DATA_WRITER_INSTANCE_REPLACED` status on a *DataWriter* using `set_listener()`.

This problem did not occur if the listener was set when the *DataWriter* was created.

This problem has been resolved.

[RTI Issue ID CORE-7551]

4.1.4 Fixes to the printing of TypeCodes

The following bugs in the **DDS_TypeCode_print_IDL** API have been resolved:

- Printing a TypeCode that contained an array resulted in the dimensions of the array being printed in the wrong location (e.g., `long[3] myLongArray` instead of `long myLongArray[3]`). [RTI Issue ID CORE-6676]
- A union with a boolean type as the discriminator was incorrectly printed. Instead of printing "TRUE" and "FALSE", "1" and "0" were printed. [RTI Issue ID CORE-9455]
- Union types with an enum as their discriminator were incorrectly printed when using the **DDS_TypeCode_print_IDL** API (and its equivalent API in other languages). Instead of printing the textual enum label, the associated numerical value was printed. [RTI Issue ID CORE-9426]
- A struct type that used inheritance incorrectly had "public" printed before its members. [RTI Issue ID CORE-7554]

Additionally, the following changes have been made to the behavior of the API:

- The APIs will now print annotations using the prefix annotation (e.g., `@key` and `@extensibility` instead of `//@key` and `//@extensibility`).
- If the values of the annotations `@min`, `@max`, and `@default` are equal to the defaults, they are no longer printed.
- By default, the ordinal value of an enum is only printed if it is non-default (this behavior is configurable using the **print_ordinals** field of the **DDS_TypeCodePrintFormatProperty** structure). A default value is defined as 0 for the first element in the enum, and one more than the previous ordinal for all other ordinals.

[RTI Issue ID CORE-9206]

4.1.5 Wrong return code or exception for **DDS_DataWriter_get_matched_subscription_data** and **DDS_DataReader_get_matched_publication_data**

In release 6.0.0, **DDS_DataWriter_get_matched_subscription_data** and **DDS_DataReader_get_matched_publication_data** incorrectly returned `DDS_RETCODE_ERROR` instead of `DDS_RETCODE_PRECONDITION_NOT_MET` when the instance handle did not correspond to any

matched endpoint. This problem also affected APIs that use exceptions instead of return codes (Modern C++, Java, and .NET). This problem has been fixed.

[RTI Issue ID CORE-9232]

4.1.6 Nanosecond field within `source_timestamp` was not validated in `DataWriter::write_w_params` or `DataWriter::write_w_timestamp`

The `DDS_Time_t::nanosec` field was not validated when calling `DataWriter::write_w_params` or `DataWriter::write_w_timestamp`.

This problem has been resolved; if the nanosec field exceeds 999999999 (i.e., one less than the number of nanoseconds in a second), these functions now return `DDS_RETCODE_BAD_PARAMETER`.

[RTI Issue ID CORE-9272]

4.1.7 `LogMessage`'s `is_security_message` flag did not work with formatted log messages

When the log message had certain `LogPrintFormats` enabled, such as timestamp or thread name, the `LogMessage`'s flag `is_security_message` was not set properly. This issue has been resolved. The flag `is_security_message` is now evaluated correctly for formatted messages.

[RTI Issue ID CORE-9617]

4.1.8 Using RTI DDS Spy from RTI Connex Launcher did not print samples in some cases

In release 6.0.0, using *RTI DDS Spy* from *RTI Launcher* did not print samples (using the "Print samples" option) while receiving data in XCDR encoding version 2 (XCDR2). This problem has been resolved.

[RTI Issue ID CORE-9630]

4.1.9 Utilities summary printed NULL pointers, causing memory fault in some architectures

In release 6.0.0, when running the utilities *RTI DDS Ping*, *DDS Spy*, or *RTI Connex DDS Prototyper* with a verbosity of 3 or more, you may have experienced a memory fault.

This issue occurred when the utility's configuration summary was shown, and some fields had NULL pointers. When running the utility with verbosity 3 or more, the utility prints a summary of what is configured—for example, the domain or the name of the configuration file. If one of these values pointed to NULL—that is, it was not verified before printing it—the utility failed because it tried to print the value pointing to a NULL value. As a result, some compilers may have printed NULL; others may have failed, exiting with a memory fault.

This problem has been resolved. Now the utility checks if a value is not null: if it is not null, the utility prints the value; if it is null, the utility prints "(not set)".

[RTI Issue ID CORE-9720]

4.1.10 LogPrintFormat NDDS_CONFIG_LOG_PRINT_FORMAT_MINIMAL was not correct in Java API

In the Java API, the LogPrintFormat NDDS_CONFIG_LOG_PRINT_FORMAT_MINIMAL was supposed to print only message number and method name; however, it logged the actual message and the method name.

This issue has now been resolved. The LogPrintFormat NDDS_CONFIG_LOG_PRINT_FORMAT_MINIMAL now prints the message number and method name.

[RTI Issue ID CORE-9749]

4.1.11 Unexpected DDS_RETCODE_TIMEOUT write error when max instances was exceeded

The call to the **DataWriter::write** operation may have failed unexpectedly with a DDS_RETCODE_TIMEOUT error when you tried to write a sample for a new instance that could have replaced an existing instance according to the instance replacement policy.

For example, if the instance replacement was set DDS_DISPOSED_INSTANCE_REPLACEMENT, the **DataWriter::write** operation for a new instance may have failed with DDS_RETCODE_TIMEOUT error even though there were DISPOSED instances with all the samples acknowledged that could have been replaced.

When this problem occurred, you would have seen the following log messages:

```
DL Warning: : WriterHistoryOdbcPlugin_createResources:writer history full
DL Warning: : PRESWriterHistoryDriver_addWrite:!history full
DL Error: : PERSISTENCEServiceTopic_onDataAvailable:!write
DL Warning: : PRESReaderWriter_blockWriter:max blocking time expired
DL Error: : PERSISTENCEServiceTopic_writeSample:!write
DL Warning: : PRESWriterHistoryDriver_registerInstance:!history full
DL Warning: : PRESWriterHistoryDriver_addDispose:!history full
```

This problem has been fixed.

[RTI Issue ID CORE-9863]

4.2 Fixes Related to Transports

4.2.1 `add_peer` API returned `RETCODE_DDS_OK` when a hostname could not be resolved

When adding a peer with a hostname instead of an IP address, the *DomainParticipant*'s `add_peer` API returned `RETCODE_DDS_OK` even if the hostname could not be resolved. This issue has been resolved. Now adding a peer with a hostname that cannot be resolved produces an error.

[RTI Issue ID CORE-8959]

4.2.2 Windows IP mobility changes not always detected

In some scenarios, new network interfaces were not detected on Windows hosts. This issue has been solved, but now a network interface may be detected before it is ready to join a multicast group. A new UDPv4/UDPv6 transport property, `join_multicast_group_timeout`, has been introduced to control how long *Connex DDS* will retry to join a multicast group with the new network interface before withdrawing.

[RTI Issue ID CORE-9217]

4.2.3 Shared memory locator ignored if it matched a multicast address

A shared memory locator that matched a multicast address pattern was incorrectly treated as invalid. As a result, shared memory transport may have failed to communicate or *DomainParticipant* creation may have even failed, if the MAC address of the first network interface in the host running *Connex DDS* matched certain problematic values. Note that the problematic values could not be predictably found within a specific address range, because the locator that shared memory uses is computed as an md5 hash of the actual address; therefore, it was specific values that matched and caused the problem.

This behavior has been fixed. Now shared memory transport works as expected independently of the MAC addresses of the network interfaces available in the host running *Connex DDS*.

[RTI Issue ID CORE-9732]

4.2.4 Malformed `TRANSPORT_MULTICAST_MAPPING` QoS caused a crash in the application

A malformed `TRANSPORT_MULTICAST_MAPPING` QoS without the `topic_expression` field could lead to a crash in the application. This issue has been fixed, and the malformed QoS is now reported as inconsistent.

[RTI Issue ID CORE-9742]

4.2.5 Network interface name not parsed properly on Windows

In some scenarios, the name of a network interface was not properly parsed on Windows. This incorrect parsing could lead to two different scenarios: the allow/deny interface list properties' not filtering the expected interfaces or *Connex DDS* not being able to use the network interfaces. This problem has been fixed. Now the name of the network interfaces is parsed correctly on Windows.

[RTI Issue ID CORE-9794]

4.2.6 Errors parsing DTLS Transport properties were not reported

Errors in the parsing of properties related to the DTLS Transport may have not been reported. This problem has been fixed. Now the errors in the parsing of the properties are reported correctly.

[RTI Issue ID COREPLG-471]

4.2.7 Slow or no communication when enabling TCP Transport's CRC

Enabling TCP Transport's CRC feature may have resulted in a slow communication, or even in no communication at all. When this issue was triggered, the following (or similar) messages were logged to the output:

```
NDDS_Transport_TCPv4_Connection_checkCrc:failed to validate checksum
NDDS_Transport_TCPv4_receive_rEA:received corrupted user data message, skipping it...
NDDS_Transport_TCPv4_Connection_nonBlockingRead:wrong msg signature: recv=00000100,
expected=dd54dd55
NDDS_Transport_TCPv4_Connection_checkCrc:failed to validate checksum
NDDS_Transport_TCPv4_receive_rEA:received corrupted user data message, skipping it...
NDDS_Transport_TCPv4_Connection_nonBlockingRead:wrong msg signature: recv=00000100,
expected=dd54dd55
NDDS_Transport_TCPv4_Connection_checkCrc:failed to validate checksum
NDDS_Transport_TCPv4_receive_rEA:received corrupted user data message, skipping it...
NDDS_Transport_TCPv4_Connection_nonBlockingRead:wrong msg signature: recv=00000100,
expected=dd54dd55
```

This problem, which was more likely to be triggered when sending large samples, is now resolved.

[RTI Issue ID COREPLG-475]

4.2.8 Potential rare crash upon TCP Transport shutdown

There was an issue in TCP Transport that may have provoked a crash upon TCP Transport shutdown (it may have occurred during *DomainParticipant* shutdown). Specifically, the issue was triggered by a rare race condition where the processing of a new TCP connection overlapped in time with TCP Transport destruction.

This issue is now resolved.

[RTI Issue ID COREPLG-501]

4.2.9 Potential rare hang during TCP Transport shutdown

In rare situations, TCP Transport may have run into an inconsistent state during its shutdown. When this issue was triggered, the following message have logged continuously, preventing TCP Transport from shutting down:

```
NDDS_Transport_TCPv4_delete_cEA: unexpected situation: serverCloseConnection: unknown
connection type
```

This issue is now fixed.

[RTI Issue ID COREPLG-502]

4.3 Fixes Related to Reliability Protocol and Wire Representation

4.3.1 Invalid fragment size when using FlowController

When using a FlowController, the fragment size for a sample should be either (a) `bytes_per_token` or (b) the smallest of the `message_size_max` values across all transports installed with the *DataWriter*, whichever is less.

In previous releases, however, if the serialized size of the sample was smaller than or equal to the maximum message size of the transport with which the sample was sent, minus some protocol overhead (512 bytes by default), *Connex DDS* chose the maximum message size of the transport as the fragment size, ignoring `bytes_per_token`, which may be smaller.

This behavior may have led to a bursty behavior of the FlowController that resulted in not sending fragments every period but sending them only when enough tokens were accumulated to send the full sample.

This problem has been resolved.

If you need to go back to the old behavior before this fix, please contact RTI Support (support@rti.com).

[RTI Issue ID CORE-9287]

4.3.2 ReliableReaderActivityChangedStatus may have had invalid content

The `ReliableReaderActivityChangedStatus` retrieved from a *DataWriter* may have had invalid content in some cases. For example, the counters could have gone negative or they may simply have contained invalid information. This issue only affected the value of the status and did not affect functional correctness.

[RTI Issue ID CORE-9349]

4.3.3 Data fragmentation did not work for samples larger than 64k for certain transport configurations

If a *DomainParticipant's DataWriter* wrote samples with a serialized size larger than 64k, those samples may have failed to be delivered to any *DataReader*. Specifically, this problem was triggered for samples when both of the following were true:

- The serialized sample size was larger than the minimum value of **message_size_max** across all of the transports within the *DomainParticipant* (that is, it was a size such that it triggered data fragmentation).
- The minimum value of **message_size_max** across all of the transports was higher than 64k.

This problem is now solved. Now samples with a serialized size larger than 64k will be delivered in this scenario. Note that the maximum fragment size on the wire is 65535 bytes, so if fragmentation is triggered, that will be the maximum fragment size (minus some bytes to count for protocol overhead), independently of the configured **message_size_max**. To help identify this scenario, a new log message (with log level STATUS_LOCAL) has been added (in the following example, the minimum value of **message_size_max** across all of the transports is 130560 bytes):

```
[D0064|ENABLE] COMMENDLocalWriterRO_init:[Local Participant: 1016fa7 2db6acae 9f0c9410 | Local Endpoint 200c2] Local writer will fragment samples if their serialized size is greater than 130560 bytes. When fragmenting samples, local writer will send a maximum of 65023 bytes per fragment.
```

[RTI Issue ID CORE-9429]

4.3.4 DataReader reported incorrect sample lost and rejected when receiving coherent set

In release 6.0.0, there was an issue that provoked *DataReaders* receiving a coherent set to incorrectly increment the sample lost and sample rejected count statistics, and to incorrectly call the **onSampleLost()** and **onSampleRejected()** callbacks. When this issue was triggered, the following warnings were logged:

```
PRESCstReaderCollator_addCollatorEntryToPolled:!assert instance  
PRESCstReaderCollator_assertInstance:!keyhash not available, dropping the sample...
```

This issue is now resolved. Receiving a coherent set no longer results in reporting a sample lost and rejected.

[RTI Issue ID CORE-9576]

4.3.5 Inactivated non-progressing DataReaders were never sent data, preventing them from progressing

A *DataReader* that had been marked as inactive may never have been sent data again, even after reconnecting and requesting it. This problem happened only if **DDS_RtpsReliableWriterProtocol_**

t::inactivate_nonprogressing_readers was set to **true** and one of the following was true:

- The *DataWriter* that was communicating with the *DataReader* was not matched with any active *DataReaders*.
- The *DataReader* was a late-joining *DataReader* that had not received all historical data before being marked as inactive, and the *DataWriter* was matched with at least one other active *DataReader*.

This issue has been resolved.

[RTI Issue ID CORE-9585]

4.3.6 Unbounded memory growth on DataWriter when failing to process NACK_FRAGs from DataReaders

When fragmentation of data samples is required, a NACK_FRAG is a message that is sent from a *DataReader* to a *DataWriter* to inform the *DataWriter* about specific fragment numbers that the *DataReader* is still missing. There was an incorrect unbounded memory growth on a *DataWriter* when it received NACK_FRAGs for samples but correctly failed to process the NACK_FRAGs (e.g., due to inconsistencies between the contents of the NACK_FRAG and information about the sample). This problem has been fixed by removing the unbounded memory growth.

[RTI Issue ID CORE-9273]

4.3.7 DataReaders may not receive samples from DataWriter when using FlowController

A *DataReader* may not have received samples from a *DataWriter*, or it may have received the samples after a long delay, if the *DataWriter* was configured to use an RR or HPF FlowController. This issue occurred only when the following conditions were met:

- There were other *DataReaders* matching with the *DataWriter* or the *DataReader* announced multiple locators but was reachable only in a subset of these locators.
- The number of tokens distributed every period was not enough to provide at least one token to a reachable locator for the *DataReader*.

This problem has been resolved.

[RTI Issue ID CORE-9293]

4.3.8 DataWriter RELIABLE_READER_ACTIVITY_CHANGED StatusMask not set when a matched DataReader was deleted

The *DataWriter* StatusMask for RELIABLE_READER_ACTIVITY_CHANGED status was not set when a matched *DataReader* was deleted from the *DataWriter* records. This problem could also have happened if there was a loss of liveness with the remote *DomainParticipant* containing the *DataReader*.

This problem has been resolved.

[RTI Issue ID CORE-9418]

4.3.9 Invalid serialization of samples with types containing primitive members that require padding

In 6.0.0, the serialization of samples with a type containing a nested complex type with primitive members that require padding may have failed. This means that a *DataReader* may have received an invalid value for a sample.

Example:

```
@nested struct Struct_3 {
    float m1;
    long long m2;
    short m3;
};

@nested struct Struct_2 {
    Struct_3 m1;
};

struct Struct_1 {
    Struct_2 m1;
};
```

In the above example, Struct_3 is nested, and there is padding between m1 (4-byte aligned) and m2 (8-byte aligned) of 4 bytes.

This problem affected DynamicData and the generated code for the following languages: C, C++, C++03, and C++11.

For generated code, a potential workaround to this problem was to generate code with a value of 1 for the **-optimization**, but this may have had performance implications.

This problem has been resolved.

[RTI Issue ID CORE-9777]

4.3.10 Unexpected fragmentation in the ParticipantBuiltinTopicData DataWriter prevented discovery

In 6.0.0, if the *DomainParticipant* discovery information "data(p)" sample was bigger than the smallest **message_size_max** value across all installed transports, minus 512 bytes, the sample was incorrectly fragmented on the builtin *DataWriter*'s side. Additionally, since the builtin *ParticipantBuiltinTopicData DataReader* does not support fragmentation, it failed to reassemble these fragmented messages, and discovery did not complete.

This behavior has been fixed by disabling fragmentation of *DomainParticipant* discovery samples on the *DataWriter* side (since it's not supported on the *DataReader* side). As a result, it is now possible to send "data(p)" samples with a size up to the smallest **message_size_max** value across all installed transports, which was the existing behavior in 5.3.1 and earlier.

Now, if the participant discovery information "data(p)" sample is bigger than the smallest **message_size_max** value across all installed transports, you will see the following message:

```
COMMENDAnonWriterService_checkFragmentationSupport:!fragment data. Not supported by this writer
```

Otherwise, if the participant discovery information "data(p)" sample, plus the overhead of the RTPS message encapsulating the DATA(p) sample, is bigger than the **message_size_max**, you will see a message similar to the following:

```
MIGGenerator_addData:[Participant: 10172c3 8d0be83b 7092ba51] Message size max ('950') (for at least one of the installed transports) is too small for propagating the participant discovery information. Please contact support@rti.com for more information.
```

If you see either of these error messages, the solution is to increase the transport **message_size_max** property or decrease the size of the *DomainParticipant* discovery information "data(p)" sample.

[RTI Issue ID CORE-9872]

4.4 Fixes Related to Content Filters

4.4.1 ContentFilteredTopic updates unnecessarily propagated when filter expression and parameters remained unchanged

If you updated a filter expression and/or parameters, but decided to use the same values as before, remote participants were unnecessarily notified about the update, even though the expression did not change. This behavior may have resulted in excessive bandwidth usage.

This issue has been resolved. Now remote participants are notified of *ContentFilteredTopic* updates only when the filter expression or parameters have actually changed their values.

[RTI Issue ID CORE-9412]

4.4.2 Incorrect writer-side filtering by reliable DataWriters in some cases

When a *DataWriter* is matched to a *DataReader* using a *ContentFilteredTopic*, the *DataWriter* issues GAP messages for samples that do not pass the filter. Writer-side filtering by a reliable *DataWriter* was incorrect in the following two scenarios:

- If the *DataReader* changed its filter, then all unacknowledged samples in the *DataWriter* cache were announced via heartbeat messages to the *DataReader*. Samples in the *DataWriter* cache that didn't pass the old filter but pass the new filter were published to the *DataReader*. This behavior was incorrect, because samples were already filtered when they were initially published and they should not be sent to the *DataReader* now. (Although they pass the filter "now," they didn't pass it then, so they should not be sent.)
- If a new *DataReader* was created on the same locator as the existing *DataReader* and it matched with the *DataWriter*, then all unacknowledged samples in the *DataWriter* cache were announced to all matched *DataReaders* in the locator via heartbeat messages. If the old *DataReader* changed its filter, then it could receive previously filtered samples if they passed the new filter. (But, as in the first scenario, they didn't pass before, so they should not be sent.)

This problem has been resolved via GAP messages issued by the *DataWriter* to the matched *DataReaders* indicating the previously filtered samples.

[RTI Issue ID CORE-9707]

4.5 Fixes Related to TopicQueries

4.5.1 DataReader::lookupInstance did not work with TopicQueries

The **DataReader::lookupInstance** API may have returned a NIL instance handle for an instance that was known to a *DataReader*, if samples for that instance were only in the *DataReader's* queue in response to a *TopicQuery*.

This problem has been resolved. The API now searches for the instance in the queues associated with live data as well as *TopicQuery* data.

[RTI Issue ID CORE-9661]

4.5.2 Potential segmentation fault when using TopicQueries and enabling application-level acknowledgment

A subscribing application may have produced a segmentation fault or printed the following precondition errors:

```
PRESsReaderQueueRemoteWriterQueue_getAcknowledgementState:!precondition: "self == 0 ||
appAckOut = 0"
PRESsService_readerSampleListenerOnGetAppAck:!get acknowledgement state (no key)
```

```
COMMENDSrReaderService_onAckPeriodicEvent:!request app ack
PRESPsReaderQueueRemoteWriterQueue_getAcknowledgementState:!precondition: "self == 0 ||
appAckOut = 0"
PRESPsService_readerSampleListenerOnGetAppAck:!get acknowledgement state (no key)
COMMENDSrReaderService_onAckPeriodicEvent:!request app ack
```

The problem may have occurred if the application created a *DataReader* that enabled application-level acknowledgment (by setting **reader_qos.reliability.acknowledgment_kind**), and the *DataReader* matched with a *DataWriter* that enabled TopicQuery support (by setting **writer_qos.topic_query_dispatch.enable** to TRUE).

This problem has been resolved.

[RTI Issue ID CORE-9778]

4.5.3 Precondition error when a DataReader issuing a TopicQuery was rematched

The following sequence of operations may have led to a precondition error with debug libraries and an incorrect value for the DDS_ServiceRequestAcceptedStatus for release libraries:

```
DataReader (DR) and DataWriter (DW) match
DR creates TopicQuery and DW dispatches it
DR stops matching (e.g. partition change)
DR matches again (partition reverted)
```

The precondition error was something like this:

```
rEvt0006afff78e Mx0D:PsRemoteTopicQuery.c:1497:RTI0x2000022:!precondition:
"matchedTopicQueryAlreadyExists"
```

This problem has been fixed.

[RTI Issue ID CORE-9869]

4.5.4 DataWriter did not dispatch TopicQueries for DataReaders that were inactive when the TopicQuery was received

Under the following sequence of events, a *DataWriter* would have never dispatched a TopicQuery for a *DataReader*:

1. The DataWriter matched with the DataReader.
2. The DataReader became inactive after `writer_qos.protocol.rtps_reliable_writer.max_heartbeat_retries` Heartbeats (HBs) due to the lack of response to HBs (default behavior) or to not making progress on the NACK messages (non default - requires setting `writer_qos.protocol.rtps_reliable_writer.inactivate_nonprogressing_readers` set to TRUE).
3. The DataWriter received a TopicQuery from the DataReader.

4. The `DataReader` became active after starting to respond to HBs and/or making progress on the NACK messages.
5. The `DataWriter` incorrectly did not dispatch the `TopicQuery`.

This problem has been fixed.

[RTI Issue ID CORE-9875]

4.6 Fixes Related to Dynamic Data and FlatData

4.6.1 FlatData: bad performance of `PrimitiveSequenceBuilder::add_n`

In release 6.0.0, the function `PrimitiveSequenceBuilder::add_n(int count)`, which adds a number of elements to a primitive sequence in a `FlatData™` sample, initialized the new elements to zero. This initialization added unnecessary latency to the sample-building process.

This problem has been resolved by avoiding the initialization in that function, and by adding a separate overload (`PrimitiveSequenceBuilder::add_n(int count, T value)`) that initializes the elements.

[RTI Issue ID CORE-9363]

4.6.2 Memory leak while receiving encoded reassembled samples with Dynamic Data or FlatData language binding

In release 6.0.0, a memory leak was observed while receiving fragments of DDS samples via `DynamicData` or `FlatData` language binding. The leak occurred when these two features were used with data protection of the serialized payload provided by *RTI Security Plugins*. This memory leak was exposed in RTI infrastructure services and tools, since these use `DynamicData`. Affected services and tools included *RTI Routing Service*, *Recording Service*, *Queuing Service*, *Database Integration Service*, *Web Integration Service*, *DDS Toolkit (for LabVIEW™)*, *DDS SPY*, *Connector*, and *Spreadsheet Add-in for Microsoft® Excel®*.

This problem has been resolved.

[RTI Issue ID CORE-9498]

4.6.3 Memory corruption when using hierarchical names to access members of sequences that were loaned from the `DataReader`

In release 6.0.0, there may have been memory corruption when using hierarchical names to access members of a sequence in a `DynamicData` object that was loaned from the *DataReader*. For example:

```
retCode = DDS_DynamicData_get_long (
    DDS_DynamicDataSeq_get_reference (&dataSeq, 0),
    &longOut,
    "sequenceOfStructs[0].myLong",
```

```
DDS_DYNAMIC_DATA_MEMBER_ID_UNSPECIFIED);
```

This issue could present itself through a crash, exceptions, or incorrect data the next time that *DynamicData* object was used to loan a sample from the *DataReader* queue. This issue has been fixed.

[RTI Issue ID CORE-9512]

4.6.4 FlatData: possible compilation error building data for a type with a sequence

In release 6.0.0, applications using the member function **FinalSequenceBuilder::add_n(count)** failed to compile when FlatData language binding was used with the Traditional C++ API. The member function **add_next()** could be used instead, but **add_n** may achieve better performance. This problem affected types containing sequences of final structs.

Applications using FlatData language binding with the Modern C++ API were not affected by this problem.

This problem has been resolved.

[RTI Issue ID CORE-9519]

4.6.5 Accessing members using hierarchical member names may have failed if member names were subsets of each other

In release 6.0.0, accessing a member in a *DynamicData* object using a hierarchical member name failed if the member name was a subset of the previously accessed member name. For example, accessing member "nested_struct.x" after accessing member "nested_struct_array[1].y" would have failed because "nested_struct" is a subset of "nested_struct_array".

This issue has been resolved.

[RTI Issue ID CORE-9609]

4.6.6 Unbounded memory growth on volatile DataWriter while publishing FlatData samples without a matched DataReader

In release 6.0.0, *Connex DDS* may have failed to release the memory used for a *DataWriter*'s FlatData-loaned samples until the *DataWriter* was deleted, resulting in an unbounded memory growth. This unbounded memory growth was observed when a volatile *DataWriter* published FlatData samples when the writer was not matched to any *DataReader*. This leak was not observed when the *DataWriter* was configured to use FlatData language binding along with Zero Copy transfer over shared memory. This leak has been fixed.

[RTI Issue ID CORE-9618]

4.6.7 Segmentation fault when using an unkeyed DynamicData DataReader with a content filter and writer-side filtering

In release 6.0.0, a *DataReader* may have received samples that did not match the content filter that was being used by the *DataReader*, even if the *DataWriter* was performing writer-side filtering. This could happen for two reasons:

- The *DataReader* changed its filter after the *DataWriter* wrote the sample and the sample passed the *DataReader*'s previous filter.
- There was another *DataReader* at the same locator that should have received the sample.

An unkeyed DynamicData *DataReader* that received a sample due to the second reason above would have segfaulted while attempting to return the sample resources after dropping the filtered sample.

This problem has been resolved.

[RTI Issue ID CORE-9653]

4.6.8 Binding to an unset optional member caused some operations on the parent DynamicData object to fail

In release 6.0.0, binding to an unset optional member caused some operations to fail if the optional member was not unbound from before the operation was executed. The operations that would fail were **DynamicData::equals**, **DynamicData::print**, and **DynamicData::write**. These operations would behave as though the unset optional member was set. For example, the print operation would print the unset optional member, and the write operation would send it on the wire.

This behavior was most likely to be observed when using hierarchical naming because using hierarchical names cause members in the DynamicData object to be bound implicitly. These members are not unbound until a different member in the DynamicData object is accessed.

This issue has been resolved.

[RTI Issue ID CORE-9685]

4.6.9 Possible data corruption or crash when using DynamicData and a type with inheritance

In release 6.0.0, when using DynamicData and a type with inheritance, there was a possibility for data corruption or a crash when receiving the data. This error would happen if the inherited base type contained one or more members of the following types:

- A sequence of complex members (structs, sequences, unions, array)
- An optional member of a complex type

4.6.10 Missing parameter checking for several DynamicData APIs resulted in segmentation faults or

The complex members mentioned above had to themselves contain at least one of:

- a string
- a sequence (of any kind)
- an optional member

In addition, if the complex member was mutable, any member of the complex member that was not explicitly sent would not be initialized during deserialization to the correct default value if the default was something other than 0 (for example, through the use of the `@default` annotation or an enumeration with a non-zero default enumerator).

For example, the following type may have exhibited this behavior because the `BaseType` contains a sequence of complex members that contain a string:

```
struct ComplexMember {
    string myString;
};

struct BaseType {
    sequence<ComplexMember> myComplexSequence;
};

struct DerivedType : BaseType {
    long myLong;
};
```

This problem has been resolved.

[RTI Issue ID CORE-9821]

4.6.10 Missing parameter checking for several DynamicData APIs resulted in segmentation faults or incorrect return codes

Calling the following functions resulted in a segmentation fault if the `self` parameter was `NULL`:

- `DynamicDataTypeSupport::unregister_type`
- `DynamicDataTypeSupport::get_data_type`
- `DynamicDataTypeSupport::get_type_name`

`DynamicDataTypeSupport::register_type` incorrectly returned `RETCODE_ERROR` instead of `RETCODE_BAD_PARAMETER` when a `NULL` value was passed in for any of the parameters.

These issues have been resolved.

[RTI Issue ID CORE-9832]

4.6.11 Flat Data: Conversion warning on some platforms

Applications using FlatData types may have seen a warning such as the following:

```
warning: conversion to 'unsigned int' from 'long unsigned int' may alter its value
[-Wconversion]
    enum { value = primitive_lc_code_helper<sizeof(T)>::single };
```

This warning has been resolved.

[RTI Issue ID CORE-9874]

4.7 Fixes Related to Modern C++ API

4.7.1 Functions `get_discovered_participants` and `get_discovered_participant_data` were in incorrect namespace

These functions were defined in the `rti::domain` namespace, but, as standard functions, they need to be in the `dds::domain` namespace.

The functions have been moved to the correct namespace, `dds::domain`, but they are still accessible from their previous namespace to preserve application backward compatibility.

[RTI Issue ID CORE-7552]

4.7.2 Possible compilation error due to clash with POSIX macro

The enumerator `dds::core::xtypes::TypeKind::MAP_TYPE` may have clashed with a POSIX macro called `MAP_TYPE`, causing the compilation on some UNIX-based systems to fail.

This problem has been fixed by renaming the enumerator. In any case, applications shouldn't use it, since *ConnexT DDS* doesn't support IDL map types.

[RTI Issue ID CORE-8527]

4.7.3 Inaccessible `extensibility_kind` getter in `EnumType`

The member function `extensibility_kind()` in `dds::core::xtypes::EnumType` that allowed retrieving the extensibility kind of an enum type was not accessible. The setter didn't have this problem.

This problem has been resolved.

[RTI Issue ID CORE-8806]

4.7.4 `TopicQuery` created with filter parameters leaked memory

The implementation of the constructor for `TopicQuery` had a memory leak.

```
TopicQuery (dds::sub::AnyDataReader reader, const TopicQuerySelection &selection)
    Creates a TopicQuery for a given DataReader.
```

The memory leak occurred because the implementation made a deep copy of the filter parameters into a `DDS_StringSeq` for passing down into the C layer, without finalizing them.

This problem has now been resolved by fixing the constructor to finalize the allocated `DDS_StringSeq` once the `TopicQuery` is created by the C layer.

[RTI Issue ID CORE-9358]

4.7.5 Entities created from XML configuration were not deleted with close()

Entities such as `dds::sub::Subscriber` or `dds::pub::DataWriter` can be explicitly deleted with their member function `close()`. Otherwise, they're deleted when their reference count goes down to zero.

However, if these entities were created from an XML configuration file (using `create_participant_from_config()`), `close()` didn't delete them. They were only deleted when the *DomainParticipant* that contained them was deleted (explicitly or not).

This problem has been resolved. Now `close()` deletes entities that have been created from an XML file.

[RTI Issue ID CORE-9483]

4.7.6 Closing a DataReader with TopicQueries may fail in some circumstances

Explicitly closing a *DataReader* with its member function `close()` may have failed if the special TopicQueries `TopicQuery::SelectAll()` or `TopicQuery::UseReaderContentFilter()` were created for that *DataReader*.

The automatic destruction of the *DataReader* didn't exhibit this problem.

This problem has been resolved. Now `close()` will not fail under these circumstances.

[RTI Issue ID CORE-9487]

4.7.7 Possible strict-aliasing warnings on some platforms

Some compilers may have produced the following warning while compiling an application using the Modern C++ API:

```
warning: dereferencing type-punned pointer will break strict-aliasing rules [-Wstrict-aliasing]
return reinterpret_cast<value_type*&>(seq._contiguous_buffer);
```

This code has been fixed, and the warning no longer appears.

[RTI Issue ID CORE-9500]

4.7.8 dds::core::polymorphic_cast didn't accept rvalues as its argument

The following expression didn't compile because the signature of `polymorphic_cast` couldn't receive an rvalue:

```
auto topic = dds::core::polymorphic_cast<Topic<Foo>>(reader.topic_description());
```

It had to be rewritten to use an lvalue:

```
TopicDescription<Foo> topic_description = reader.topic_description();  
auto topic = dds::core::polymorphic_cast<Topic<Foo>>(topic_description);
```

This problem has been resolved. Now `polymorphic_cast` can receive an rvalue as shown in the first code example.

[RTI Issue ID CORE-9556]

4.7.9 Rare memory leak in DynamicType API

In release 6.0.0, a memory leak affected the copy constructor and copy-assignment operator of the class `dds::core::xtypes::Member` and only affected string members that were annotated with the `@default` annotation.

This memory leak has been resolved.

[RTI Issue ID CORE-9673]

4.7.10 Possible leak creating a Replier with a listener

The Replier constructor that receives a listener created an object in the heap that was never rereleased. Note the leak didn't happen if the Replier was created without a listener, or if the listener was set after creation.

This memory leak has been resolved.

[RTI Issue ID REQREPLY-69]

4.8 Fixes Related to Traditional C++ API

4.8.1 Extra semicolon warning when compiling with pedantic flag

When compiling an application using the Traditional C++ API and using the `"-pedantic"` compiler flag, you may have seen a warning related to an extra semicolon in `<NDDSHOME>/include/ndds/dds_c/dds_c_sequence.h`. This problem has been resolved.

[RTI Issue ID CORE-9258]

4.8.2 Performance degradation in write() API when IDL file contained types annotated for Zero Copy transfer over shared memory

In release 6.0.0, if a Foo and a Bar type were defined in an IDL file where the Bar type was annotated with `@transfer_mode(SHMEM_REF)`, then the performance of `FooDataWriter::write()` operation was degraded in the Traditional C++ API. This performance degradation was visible when the Foo type was small. This problem has been resolved.

[RTI Issue ID CORE-9693]

4.8.3 read/take_replies functions ignored max_samples argument

The functions `Requester::read/take_replies(max_samples, related_id)` ignored the `max_samples` argument and ran as if `max_samples` was always 1.

This problem has been resolved.

[RTI Issue ID REQREPLY-65]

4.9 Fixes Related to XML Configuration

4.9.1 XML configuration of FlowController properties did not accept DDS_LENGTH_UNLIMITED

When configuring a FlowController in XML QoS using the *DomainParticipant's* Property QoS policy, the XML did not accept the value `DDS_LENGTH_UNLIMITED` for the FlowController properties, as shown in the following example:

```
<element>
  <name>dds.flow_controller.token_bucket.custom_flowcontroller.token_bucket.tokens_leaked_per_period</name>
  <value>DDS_LENGTH_UNLIMITED</value>
</element>
```

This issue has been resolved. The `DDS_LENGTH_UNLIMITED` value is now accepted.

[RTI Issue ID CORE-9253]

4.9.2 Elements within <domain_participant> were forced to have an order

In release 6.0.0, the XSD definition file (https://community.rti.com/schema/6.0.0/rti_dds_profiles.xsd), enforced a specific ordering for tags within the `<domain_participant>` tag in the `<domain_participant_library>` tag.

The following XML example was valid in 5.3.1, but not in 6.0.0:

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

```

xsi:noNamespaceSchemaLocation="https://community.rti.com/schema/6.0.0/rti_dds_
profiles.xsd"> <!-- QoS Profiles -->
  <qos_library name="QosLibrary">
    ...
  </qos_library> <!-- Domains -->
  <domain_library name="DomainLibrary">
    <!-- Shape -->
    <domain name="Shape" domain_id="0">
      ...
    </domain> <!-- Participants -->
  <domain_participant_library name="DomainParticipantLibrary">
    <!-- Shape -->
    <domain_participant name="ShapePublisher" domain_ref="DomainLibrary::Shape">
      <participant_qos base_name="QosLibrary::Base">
        ...
      </participant_qos>
      <publisher name="Publisher">
        ...
      </publisher>
    </domain_participant>
    <domain_participant name="ShapeSubscriber" domain_ref="DomainLibrary::Shape">
      <participant_qos base_name="QosLibrary::Base">
        ...
      </participant_qos>
      <subscriber name="Subscriber">
        ...
      </subscriber>
    </domain_participant>
  </domain_participant_library>
</dds>

```

The validation of this example using the above-mentioned XSD failed because the `<publisher>` tag was supposed to appear before the `<participant_qos>` tag.

This problem has been resolved. Now any element can appear in any order. See the XSD embedded documentation for more details.

[RTI Issue ID CORE-9306]

4.9.3 Schema files were not compliant with DDS-XML spec

The following changes have been made to the schema files (`rti_dds_profiles.xsd`, `rti_dds_qos_profiles.xsd`, `rti_dds_qos_profiles.xsd` and their included files) to make them compliant with the DDS-XML spec:

- Documented the elements that are not part of the XML spec as RTI extensions
- Updated the top-level complex type inside of the `rti_dds_profiles.xsd` file to allow empty `domain_participant_library` elements
- Removed the version attribute files from the XML files generated by *RTI Code Generator*
- Removed the default values

[RTI Issue ID CORE-9375]

4.9.4 QoS policies not resolved correctly when inheriting from a <qos_profile>, which contained an <xxx_qos> specifying its own base_name attribute

In release 6.0.0, the implementation of QoS Profile composition introduced a regression. Consider this example:

```
<qos_profile name="Parent">
  <datawriter_qos name="DW_InParent" base_name="BuiltinQosLibExp::Generic.StrictReliable">
    <!-- Values specified by DW_InParent -->
  </datawriter_qos>
</qos_profile>

<qos_profile name="Child">
  <datawriter_qos name="DW_IncorrectValues" base_name="Parent" />
</qos_profile>
```

In this example, "DW_IncorrectValues" points to "Parent," which is not a <datawriter_qos> but a <qos_profile>. Therefore, *Connex DDS* should inherit values from the following elements, in order:

1. **BuiltinQosLibExp::Generic.StrictReliable** (and its parents, and their parents if any)
2. Values set in XML by DW_InParent

There was a problem, however, with the way *Connex DDS* populated parents in this scenario. As a result, *Connex DDS* only copied values set in XML by "DW_InParent," ignoring the QoS parent value **BuiltinQosLibExp::Generic.StrictReliable** specified in the **base_name** attribute, giving incorrect results.

This problem has been resolved by fixing the internal logic to populate parent values correctly. Now when the inherited <qos_profile> contains an <xxx_qos> having its own base_name attribute value, *Connex DDS* finds and returns the correct QoS values.

[RTI Issue ID CORE-9376]

4.10 Fixes Related to XML-Based Application Creation

4.10.1 Attribute base_name not allowed in <domain_participant> and <domain>

The XSD definition file (https://community.rti.com/schema/6.0.0/rti_dds_profiles.xsd) was missing the base_name attribute in the <domain_participant> and <domain> tag. As a result, if you used the base_name attribute, the XSD validation failed with an error like the following:

4.10.2 XSD validation failed if topic_ref or register_type_ref used some non-alphanumeric characters

```
[main] ERROR com.rti.micro.appgen.MicroAppGen - cvc-complex-type.3.2.2: Attribute 'base_name' is not allowed to appear in element 'domain_participant'.
```

(For information about the base_name attribute, see the "Participant Library" section in the *RTI Connext DDS Core Libraries XML-Based Application Creation Getting Started Guide*.)

This problem has been resolved. Now the base_name attribute is allowed in the <domain_participant> and <domain> tag.

[RTI Issue ID CORE-9359]

4.10.2 XSD validation failed if topic_ref or register_type_ref used some non-alphanumeric characters

In release 6.0.0, the XSD validation of an XML application creation file failed if there were *DataWriter* topic_ref and register_type_ref attributes using non-alphanumeric characters such as '#'. The only strings allowed were ones matching this pattern: ((?[a-zA-Z0-9_])+) This problem has been fixed. Now all characters are allowed.

[RTI Issue ID CORE-9484]

4.11 Fixes Related to Micro Application Generator (MAG)

This section describes fixes to bugs in Micro Application Generator (MAG). This feature is described in the *XML-Based Application Creation Getting Started Guide*.

4.11.1 MAG failed to calculate the value of an Entity when inheriting from a profile in a different library

MAG failed to derive the value of an Entity if the **base_name** pointed to a profile in a different library. For example, the following didn't work:

```
<qos_library name="BaseLibrary">
  <qos_profile name="BaseProfile">
    <datareader_qos>
      ...
    </datareader_qos>
  </qos_profile>
</qos_library>
<qos_library name="QosLibrary">
  <qos_profile name="Test">
    <datareader_qos base_name="BaseLibrary::BaseProfile"/>
  </qos_profile>
</qos_library>
```

The above example caused the following error:

```
[main] ERROR com.rti.micro.appgen.MicroAppGen - Failed to calculate the system model.
java.lang.Exception: Unable to find QoS library/profile 'QosLibrary::BaseLibrary'.
```

This problem has been resolved. Now MAG can derive an Entity's final value when inheriting from a profile in a different library.

[RTI Issue ID MAG-47]

4.11.2 Generated code didn't compile when using data types in namespaces

The code generated by MAG didn't compile when the types were inside namespaces. This issue has been resolved and now MAG supports types inside namespaces. In addition, there is a new *rtiddsmag* command-line option, **-namespace**, which needs to be used when the code generated by *rtiddsgen* uses the **-namespace** flag.

[RTI Issue ID MAG-48]

4.11.3 Generated code didn't compile if publisher_name or subscriber_name not specified

The code generated by MAG did not compile if the **publisher_qos/subscriber_qos** used by the one of the participants defined in the XML file did not specify its **publisher_name/subscriber_name**. For example:

```
<domain_participant_library name="Library">
  <domain_participant name="Participant" domain_ref="...">
    <subscriber name="Sub">
      ...
      <subscriber_qos>
      </subscriber_qos>
    </subscriber>
  </domain_participant>
</domain_participant_library>
```

The generated code contained an invalid comma after `RTI_MANAGEMENT_QOS_POLICY_DEFAULT`:

```
#define RTI_APP_GEN__S_QOS_Library_Participant_Sub_subscriber \
{ \
  DDS_ENTITY_FACTORY_QOS_POLICY_DEFAULT, \
  RTI_MANAGEMENT_QOS_POLICY_DEFAULT, \
  DDS_SUBSCRIBERQOS_APPGEN_INITIALIZER \
}
```

When it should have generated this:

```
#define RTI_APP_GEN__S_QOS_Library_Participant_Sub_subscriber \
{ \
  DDS_ENTITY_FACTORY_QOS_POLICY_DEFAULT, \
  RTI_MANAGEMENT_QOS_POLICY_DEFAULT \
  DDS_SUBSCRIBERQOS_APPGEN_INITIALIZER \
}
```

The comma should only be added when the name is set:

```
<domain_participant_library name="Library">
  <domain_participant name="Participant" domain_ref="...">
    <subscriber name="Sub">
      ...
      <subscriber_qos>
        <subscriber_name>
          <name>MySub</name>
        </subscriber_name>
      </subscriber_qos>
    </subscriber>
  </domain_participant>
</domain_participant_library>
```

This problem has been resolved. Now MAG adds the comma only if the **publisher_name/subscriber_name** is set in the XML file.

[RTI Issue ID MAG-49]

4.11.4 Invalid generated code when the name value contained an invalid variable character

MAG generated invalid code when a name value contained an invalid variable character. For example:

```
<domain_participant_library name="HelloWorldAppLibrary">
<domain_participant name="HelloWorldDPSEPubDP-"
```

The above resulted in the following invalid variable:

```
#define RTI_APP_GEN_DP_HelloWorldAppLibrary-HelloWorldDPSEPubDP- \
```

This problem has been resolved. Now invalid variable characters are converted to '_' in the generated code. For example:

```
#define RTI_APP_GEN_DP_HelloWorldAppLibrary_HelloWorldDPSEPubDP_ \
```

[RTI Issue ID MAG-53]

4.11.5 Incorrect QoS for DataWriters and DataReaders when inheriting from a base profile using a topic_ref

MAG derived incorrect QoS values for a *DataWriter* or *DataReader* when they inherited from a base profile using a **topic_ref**. For example:

```
<qos_library name="ChatQosLib">
  <qos_profile name="Pattern.Status">
    <datawriter_qos name="test">
      <history>
        <depth>32</depth>
      </history>
    </datawriter_qos>
  </qos_profile>
  <qos_profile name="User.Service.Chat">
```

```

    <datawriter_qos>
      <history>
        <depth>12</depth>
      </history>
    </datawriter_qos>
    <datawriter_qos topic_filter="Chat*" base_name="Pattern.Status::test"/>
  </qos_profile>
</qos_library>

<data_writer name="HelloWorldWriter" topic_ref="HelloWorldTopic">
  <datawriter_qos base_name="ChatQosLib::User.Service.Chat"/>
</data_writer>

```

The `DataWriterQos` mistakenly inherited values from the second `datawriter_qos` within `User-Service.Chat` because `topic_ref` was not taken into account. The resulting QoS should have inherited from the first `datawriter_qos` within `User.Service.Chat` because "HelloWorldTopic" doesn't match with the `topic_filter` "Chat*".

This problem has been resolved. Now if the `topic_filter` is not specified, MAG uses the `topic_ref` when deriving the QoS values for a `DataWriter` or `DataReader`.

[RTI Issue ID MAG-57]

4.11.6 MAG failed when XML contained two topics with same name

MAG failed when the XML contained two topics with the same name, even though they were in different domains. This problem has been resolved. Now MAG allows duplicate topic names, as long as they are in different domains.

[RTI Issue ID MAG-59]

4.11.7 MAG failed when base_name attribute was specified within <domain>

MAG reported the following error when a domain tried to inherit from another domain:

```
[main] ERROR com.rti.micro.appgen.MicroAppGen - cvc-complex-type.3.2.2: Attribute 'base_name'
is not allowed to appear in element 'domain'.
```

The schema used by MAG to validate the XML file didn't contain the `base_name` attribute. This problem has been resolved and now domain inheritance is supported by MAG.

[RTI Issue ID MAG-60]

4.11.8 MAG failed to process XML file if a duration element contained DURATION_ZERO_SEC or DURATION_ZERO_NSEC

MAG failed to process the XML file if any Duration element contained `DURATION_ZERO_SEC` or `DURATION_ZERO_NSEC`. In this case, MAG reported the following error:


```
10:55:08.250 [main] ERROR com.rti.micro.appgen.MicroAppGen - Failed to parse referenced file:
...
10:55:08.250 [main] DEBUG com.rti.micro.appgen.MicroAppGen -
java.lang.Exception: Invalid values found while creating the Model.
```

This problem has been resolved. Now MAG will properly process XML files containing `DURATION_ZERO_SEC` or `DURATION_ZERO_NSEC` (transforming those values to 0) on any `Duration` element.

[RTI Issue ID MAG-61]

4.11.9 Command-line option `-inputXml` was not documented

The command-line option for specifying an input XML file was not documented. This problem has been resolved. Now the optional parameter `-inputXml` is in the documentation and appears in the output when you run `rtiddsmag -help`.

[RTI Issue ID MAG-62]

4.11.10 MAG failed to generate code when an entity inherited from an empty entity or an entity whose values were the same as defaults used by Connex DDS Micro

MAG failed to generate code when an entity inherited from an empty entity or from an entity whose values were the same as the defaults used by *Connex DDS Micro*. This problem has been resolved.

[RTI Issue ID MAG-67]

4.11.11 Invalid QoS values for DataWriters and DataReaders when trying to inherit values using a `topic_filter`

MAG derived the wrong value for a `DataWriterQos` or `DataReaderQos` when it inherited from a base profile using a `topic_filter`.

The `topic_filter` was not taken into account:

If an intermediate profile didn't contain a valid candidate, e.g:

```
<qos_library name="QosLibrary">
  <qos_profile name="BaseProfile">
    <datawriter_qos topic_filter="Chat*">
      ...
    </datawriter_qos>
    <datareader_qos topic_filter="Chat*">
      ...
    </datareader_qos>
  </qos_profile>

  <qos_profile name="MyDerivedQoSProfile" base_name="BaseProfile">
    </qos_profile>
```

4.11.12 NullPointerException from MAG when initial_peers not specified and builtin_transport was

```
<qos_profile name="MyQoSProfile">
  <datareader_qos name="MyReader" topic_filter="Chat1" base_
name="QosLibrary::MyDerivedQoSProfile"/>
</qos_profile>
</qos_library>
```

MyReader should have inherited its values from the *DataReader* in BaseProfile, but it didn't because the logic ignored the original **topic_filter** value if one of the base profiles (in this case, MyDerivedQoSProfile) didn't contain a valid candidate (MyDerivedQoSProfile doesn't contain any **datareader_qos**).

If the logic found only one valid candidate, e.g:

```
<qos_library name="QosLibrary">
  <qos_profile name="BaseProfile">
    <datareader_qos topic_filter="Chat*">
      ...
    </datareader_qos>
  </qos_profile>

  <qos_profile name="MyQoSProfile">
    <datareader_qos name="MyReader" base_name="QosLibrary::BaseProfile"/>
  </qos_profile>
</qos_library>
```

MyReader should have not inherited from the *DataReader* inside the BaseProfile because the **topic_filter** didn't match.

These two issues have been resolved.

[RTI Issue IDs MAG-71, MAG-74]

4.11.12 NullPointerException from MAG when initial_peers not specified and builtin_transport was specified

MAG threw a NullPointerException when the **initial_peers** were not specified and the **builtin_transport** was specified in the XML file. This problem has been resolved.

[RTI Issue ID MAG-75]

4.11.13 MAG did not use same defaults as Connex DDS Micro for some values

MAG should use the same default values as *Connex DDS Micro*. However, MAG used different defaults for the following:

- **DPDE_DiscoveryPluginProperty::builtin_writer_heartbeats_per_max_samples**
- **UDP_InterfaceFactoryProperty::max_message_size**
- **UDP_InterfaceFactoryProperty::recv_thread::options**

- `DDS_DataWriterQos::reliability::kind`
- `DDS_DataWriterQos::transfer_mode::shmem_ref_settings::enable_data_consistency_check`

This problem has been resolved.

[RTI Issue ID MAG-76]

4.11.14 Generated code did not compile when participant using dynamic participant, static endpoint (DPSE) discovery contained multiple remote publication and subscription elements

The code generated by MAG contained duplicate variable names for the remote publication and subscription elements. This happened because the names of the remote entities were not taken into account. For example:

```
<!-- Participant Library -->
<domain_participant_library name="HelloWorldAppLibrary1">
  <domain_participant name="HelloWorldDPSEPubDP"
    domain_ref="HelloWorldLibrary::HelloWorldDomain">
    <publisher name="HelloWorldDPSEPub">
      <data_writer topic_ref="HelloWorldTopic" name="HelloWorldDPSEDW">
        <datawriter_qos base_name="QosLibrary::DPSEProfile"/>
      </data_writer>
    </publisher>
    <participant_qos base_name="QosLibrary::DPSEProfile"/>
  </domain_participant>
</domain_participant_library>
<domain_participant name="HelloWorldDPSESubDP"
  domain_ref="HelloWorldLibrary::HelloWorldDomain">
  <subscriber name="HelloWorldDPSESub">
    <data_reader topic_ref="HelloWorldTopic" name="HelloWorldDPSEDR">
      <datareader_qos base_name="QosLibrary::DPSEProfile"/>
    </data_reader>
  </subscriber>
  <participant_qos base_name="QosLibrary::DPSEProfile"/>
</domain_participant>
</domain_participant_library>
<domain_participant_library name="HelloWorldAppLibrary2">
  <domain_participant name="HelloWorldDPSEPubDP2"
    domain_ref="HelloWorldLibrary::HelloWorldDomain">
    <publisher name="HelloWorldDPSEPub">
      <data_writer topic_ref="HelloWorldTopic" name="HelloWorldDPSEDW">
        <datawriter_qos base_name="QosLibrary::DPSEProfile"/>
      </data_writer>
    </publisher>
    <participant_qos base_name="QosLibrary::DPSEProfile"/>
  </domain_participant>
</domain_participant_library>
<domain_participant name="HelloWorldDPSESubDP2"
  domain_ref="HelloWorldLibrary::HelloWorldDomain">
  <subscriber name="HelloWorldDPSESub">
    <data_reader topic_ref="HelloWorldTopic" name="HelloWorldDPSEDR">
```

```

        <datareader_qos base_name="QosLibrary::DPSEProfile"/>
    </data_reader>
</subscriber>
    <participant_qos base_name="QosLibrary::DPSEProfile"/>
</domain_participant>
</domain_participant_library>

```

For the HelloWorldDPSESubDP2 participant, MAG found two remote publishers: the publisher defined in the HelloWorldDPSEPubDP2 participant and the publisher defined in HelloWorldDPSEPubDP. This caused MAG to generate two variables with the same name to store this information:

```

const struct APPGEN_RemotePublicationModel
HelloWorldAppLibrary1_HelloWorldDPSESubDP_remote_publishers[1] =
{
    RTI_APP_GEN_RPD_HelloWorldAppLibrary1_HelloWorldDPSESubDP_HelloWorldAppLibrary1_
HelloWorldDPSEPubDP_HelloWorldDPSEPub_HelloWorldDPSEDW
};
const struct APPGEN_RemotePublicationModel
HelloWorldAppLibrary1_HelloWorldDPSESubDP_remote_publishers[1] =
{
    RTI_APP_GEN_RPD_HelloWorldAppLibrary1_HelloWorldDPSESubDP_HelloWorldAppLibrary2_
HelloWorldDPSEPubDP1_HelloWorldDPSEPub_HelloWorldDPSEDW
};

```

The correct generated code should have been:

```

const struct APPGEN_RemotePublicationModel
HelloWorldAppLibrary1_HelloWorldDPSESubDP_HelloWorldAppLibrary1_HelloWorldDPSEPubDP_remote_
publishers[1] =
{
    RTI_APP_GEN_RPD_HelloWorldAppLibrary1_HelloWorldDPSESubDP_HelloWorldAppLibrary1_
HelloWorldDPSEPubDP_HelloWorldDPSEPub_HelloWorldDPSEDW
};

const struct APPGEN_RemotePublicationModel
HelloWorldAppLibrary1_HelloWorldDPSESubDP_HelloWorldAppLibrary2_HelloWorldDPSEPubDP2_remote_
publishers[1] =
{
    RTI_APP_GEN_RPD_HelloWorldAppLibrary1_HelloWorldDPSESubDP_HelloWorldAppLibrary2_
HelloWorldDPSEPubDP2_HelloWorldDPSEPub_HelloWorldDPSEDW
};

```

This problem has been resolved. Now the remote publication and subscription variable names are different.

[RTI Issue ID MAG-80]

4.11.15 MAG failed to process XML file if thread mask contained more than one element

MAG failed to process the XML file if a thread mask contained more than one element, such as the following valid QoS snippet:

4.11.16 MAG failed to report error when XML file contained invalid use of LENGTH_UNLIMITED

```
<receiver_pool>
  <thread>
    <mask>STDIO | FLOATING_POINT</mask>
  </thread>
</receiver_pool>
```

In this case, MAG reported the following error:

```
10:42:20.580 [main] ERROR com.rti.micro.appgen.utils.ConverterUtils
- STDIO | FLOATING_POINT is not a valid value for recv_thread.options in Micro,
only values inside of the range [STDIO, SUSPEND_ENABLE, PRIORITY_ENFORCE,
CANCEL_ASYNCHRONOUS, MASK_DEFAULT, FLOATING_POINT, REALTIME_PRIORITY] are supported.
```

This problem has been resolved. Now MAG can parse XML files that have more than one element in a thread mask.

[RTI Issue ID MAG-88]

4.11.16 MAG failed to report error when XML file contained invalid use of LENGTH_UNLIMITED

MAG failed to report an error if it found LENGTH_UNLIMITED being used for an element that doesn't support it for *Connex DDS Micro*.

For example, LENGTH_UNLIMITED is not a valid value for **DDS_ResourceLimitsQosPolicy::max_samples** in *Connex DDS Micro*, but it is valid value for *Connex DDS Professional*. There are several elements for which LENGTH_UNLIMITED is not supported when generating code for *Connex DDS Micro*. Consult the XSD files in *<installation directory>/resource/schema/<version>/definitions* for details.

This problem has been resolved. Now MAG will report an error if it detects an element using LENGTH_UNLIMITED when it is not supported.

[RTI Issue ID MAG-90]

4.11.17 XSD validation failed if topic_ref or register_type_ref used some non-alphanumeric characters

MAG reported the following error if there were *DataWriter* **topic_ref** and **register_type_ref** attributes using non-alphanumeric characters such as '#':

```
[main] ERROR com.rti.micro.appgen.MicroAppGen - cvc-pattern-valid: Value 'HelloWorldTopic#' is
not facet-valid with respect to pattern '((:)?[a-zA-Z0-9_.])+' for type 'elementReference'.
```

The schema used by MAG to validate the XML file only allowed strings matching this pattern: ((:)?[a-zA-Z0-9_.])+ This problem has been fixed. Now all characters are allowed.

[RTI Issue ID CORE-9484]

4.11.18 MAG calculated invalid values for QoS profile specified with -outputFinalQoS option

In some cases, MAG calculated invalid values for the QoS profile specified with the **-outputFinalQoS** option. This problem has been resolved. Now MAG properly calculates the final values for the specified QoS when using the **-outputFinalQoS** option.

[RTI Issue ID MAG-97]

4.11.19 Invalid behavior when calculating final values for QoS profile specified with -outputFinalQoS option

When using the **-outputFinalQoS** option, MAG incorrectly checked to see if the final values for the specified profile were supported by *Connex DDS Micro*. This problem has been resolved. Now MAG doesn't check if the QoS values are supported by *Connex DDS Micro* when using the **-outputFinalQoS** option.

[RTI Issue ID MAG-99]

4.12 Fixes Related to Monitoring Library

4.12.1 Unbounded generation of file handles on QNX platforms

On QNX platforms, the Monitoring Library may have periodically opened file handles and never released them. Eventually this situation led to the following errors:

```
RTIOsapiInterfaces_getIPv4Interfaces:OS getifaddrs() failure, error 0X18: Too many open files
RTIOsapiInterfaceTracker_updateInterfacesUnsafe:!get interfaces failed
RTIDefaultMonitorProcess_getMemoryUsage:!get open /proc/self/as (QNX)
RTIDefaultMonitorProcess_getProcessInfo:!memory usage
```

This issue has been resolved.

[RTI Issue ID MONITOR-251]

4.12.2 Possible segmentation fault

There was a rare race condition when using Monitoring Libraries that may have led to a segmentation fault. This issue was more likely to occur if the DDS application using the monitoring libraries created and deleted entities often. This problem has been fixed.

Note: This problem was reported as fixed in MONITOR-246, in release 6.0.0; however, MONITOR-246 did not fix the problem with topics. Now, when creating and deleting topics often, the application will also not crash.

[RTI Issue ID MONITOR-252]

4.12.3 Memory limit reached

On QNX platforms, you may have received the following output when using Monitoring Library:

```
RTIDefaultMonitorProcess_getMemoryUsage:!get Number of entries greater than the limit, you
should consider increasing such limit
RTIDefaultMonitorProcess_getProcessInfo:!memory usage
```

This issue occurred because some memory was statically reserved to allocate the data structure necessary to measure memory usage. As the application grew, however, the structure should have also grown, but didn't.

This problem has been fixed. Now the structure is reserved dynamically and can measure bigger applications.

[RTI Issue ID MONITOR-254]

4.12.4 Crash when an IDL type used FlatData or the XCDR2 data representation

In release 6.0.0, applications that enabled the Monitoring Library and created a *DataReader* or a *DataWriter* for a type annotated with `@allowed_data_representation(XCDR2)` or `@language_binding(FLAT_DATA)` crashed.

This problem has been resolved.

[RTI Issue ID MONITOR-260]

4.13 Fixes Related to Vulnerabilities

This release fixes some potential vulnerabilities, including RTI Issue IDs CORE-9264, CORE-9307, CORE-9459, CORE-9651, SEC-931, and SEC-938.

4.14 Other Fixes

4.14.1 Crash when `deserialized_type_object_dynamic_allocation_threshold` set to 0

In release 6.0.0, *Connex DDS* applications that set `deserialized_type_object_dynamic_allocation_threshold` to zero (in the DomainParticipantQos's ResourceLimits QosPolicy) may have crashed if `endpoint_type_object_lb_serialization_threshold` in the DomainParticipantQos's DiscoveryConfig QosPolicy was set to a value other than -1.

This problem is fixed. Now applications no longer crash when using this configuration.

[RTI Issue ID CORE-9532]

4.14.2 TypeCodeFactory and types XML parser allowed derived structure with different extensibility than its base structure

The TypeCodeFactory allowed creating a TypeCode for structure 'A' where the extensibility for 'A' was different than the extensibility of the base structure 'B', which 'A' inherits.

Likewise, the builtin XML parser used by infrastructure services and XML application creation allowed providing a type definition for structure 'A' where the extensibility for 'A' was different than the extensibility of the base structure 'B', which 'A' inherits.

This problem has been fixed. These type definitions are no longer allowed.

[RTI Issue ID CORE-9552]

4.14.3 Json unicode escape sequences were not supported

JSON parser could not correctly convert the Unicode escape sequences (e.g., `\u000A`) to a *Connex DDS* Wchar element. This problem has been resolved.

[RTI Issue ID CORE-9649]

4.14.4 Error in DomainParticipant creation on Windows platforms

DomainParticipant creation failed on Windows platforms in some situations. The error happened when trying to retrieve the MAC address of an interface, when a buffer was not large enough to keep the list of interfaces. This error has been resolved; now the *DomainParticipant* will be created successfully in this situation.

[RTI Issue ID CORE-9664]

4.14.5 Crash while accessing Subscription built-in topic data of a Zero Copy DataReader in Java

In release 6.0.0, a crash occurred when using the Java API to retrieve Subscription built-in topic data from a *DataReader* that uses Zero Copy transfer over shared memory. This problem has been resolved.

[RTI Issue ID CORE-9798]

4.14.6 Connex DDS Micro Subscriber did not communicate with Connex DDS Professional Publisher using Zero Copy Transfer Mode

A *Connex DDS Micro* Subscriber did not receive samples from a *Connex DDS Professional* Publisher when the type of the Topic was marked with the annotation `@transfer_mode(SHMEM_REF)`. This problem has been resolved.

[RTI Issue ID CORE-9919]

4.14.7 Full path name of Connex libraries was needed when using FindRTIConnexDDS.cmake

Connex DDS did not provide information about the functionality of its libraries using an soname. As a result, build frameworks, such as CMake, that try to use the soname of a library to link against it, could not get that information. Therefore, CMake used the complete path to the library during the call to the linker. This behavior hardcoded the path to the library in the final application and forced developers to redistribute the library in the exact path used in the build machine.

This problem has been resolved. *Connex DDS* now provides the parameter `IMPORTED_NO_SONAME` set to `TRUE` for the imported targets.

[RTI Issue ID INSTALL-438]

4.14.8 `dds.sys_info.process_id` property in `DomainParticipantQos` had negative value on some VxWorks platforms in kernel mode

On some VxWorks platforms in kernel mode, the `dds.sys_info.process_id` property in the *DomainParticipant's* `PropertyQosPolicy` was populated with a negative value due to an incorrect format specifier. The problem has been resolved by using the `"%llu"` specifier for the unsigned long long type.

[RTI Issue ID PLATFORMS-1640]

Chapter 5 What's Fixed in 6.0.0

Release 6.0.0 is a general access release based on the maintenance release 5.3.1¹. This section describes bugs fixed in the Core Libraries in 6.0.0.

For an overview of new features and improvements in the Core Libraries in 6.0.0, please see the separate *What's New* document for 6.0.0.

5.1 Fixes Related to Discovery

5.1.1 Endpoint discovery initialization errors during participant creation left participant in inconsistent state

If an error occurred during a participant creation's endpoint discovery initialization, the initialization may have incorrectly reported success. In this situation, the participant was left in an inconsistent state. Trying to enable the participant would have failed with errors similar to the following:

```
[D0064|CREATE Participant|D0064|ENABLE]DISCSimpleEndpointDiscoveryPlugin_
enable:!precondition: me == NULL || worker == NULL
[D0064|CREATE Participant|D0064|ENABLE]DDS_DomainParticipantDiscovery_enableI:!enable
simple endpoint discovery plugin
[D0064|CREATE Participant|D0064|ENABLE]DDS_DomainParticipant_enableI:!enable discovery
```

This problem is now resolved.

[RTI Issue ID CORE-8249]

5.1.2 Writer/reader resource limits affected wrong builtin endpoints

Certain members of the DomainParticipantResourceLimitsQosPolicy affected the wrong endpoint-discovery builtin endpoints. For example, **writer_property_list_max_length** determined the max-

¹For What's Fixed in 5.3.1, see the *RTI Connex DDS Core Libraries Release Notes* provided with 5.3.1.

imum number of properties that the SubscriptionBuiltinTopicDataWriter could serialize when it sent discovery information about *DataReaders*. Consequently, the following scenario would fail:

1. Set **reader_property_list_max_length** to one.
2. Set **writer_property_list_max_length** to zero.
3. Create a reader with one property.
4. Create no writers.
5. The reader's discovery information would fail to be sent.

The following members of DomainParticipantResourceLimitsQosPolicy had this kind of problem:

- **writer_property_list_max_length**
- **reader_property_list_max_length**
- **writer_user_data_max_length**
- **reader_user_data_max_length**
- **publisher_group_data_max_length**
- **subscriber_group_data_max_length**

This problem has been resolved by making writer resource limits affect the PublicationBuiltinTopic endpoints and reader resource limits affect the SubscriptionBuiltinTopic endpoints.

[RTI Issue ID CORE-8411]

5.1.3 Failure to send TypeObject when type had base with no members

Connex DDS applications publishing or subscribing to a topic whose type definition had a base type with no members may have failed to send the TypeCode during endpoint discovery.

This problem may have affected only systems running *Connex DDS* 5.0.0+ applications along with pre-5.0.0 applications, or RTI tools that required the availability of the type information to work. Systems running only 5.0.0 and above relied on the TypeObject, which was not affected by this problem.

[RTI Issue ID CORE-8526]

5.1.4 Received incorrect QoS policy values through discovery when communicating with other vendors

During the discovery process, if a remote participant omitted a QoS policy because it was set to the default value, *Connex DDS* may have set that QoS policy to a default value not matching the OMG DDS and RTPS specifications. This issue affected the following QoS policies:

- DDS_ParticipantBuiltinTopicData::lease_duration
- DDS_ParticipantBuiltinTopicData::reachability_lease_duration
- DDS_PublicationBuiltinTopicData::reliability
- DDS_SubscriptionBuiltinTopicData::reliability
- DDS_PublicationBuiltinTopicData::ownership_strength

This issue only affected communication with other vendors, since *Connex DDS* never omits the propagation of these QoS policies.

This issue is now fixed: the above-mentioned QoS policies that are not received from other vendors are now set to the default value defined in the OMG DDS and RTPS specifications.

[RTI Issue ID CORE-8882]

5.2 Fixes Related to Reliability Protocol

5.2.1 VOLATILE *DataReader* may have received historical samples from *DataWriter*

A VOLATILE *DataReader* may have received some historical samples from a matching *DataWriter* after being created.

Specifically, the *DataReader* would have received all the samples from the *DataWriter* queue that were not acknowledged by all the matching *DataReaders* at the moment the new *DataReader* was discovered.

This problem has been resolved. Now a VOLATILE *DataReader* does not receive historical samples from a *DataWriter*.

[RTI Issue ID CORE-4923]

5.2.2 Wrong RTPS GAP messages emitted by reliable *DataWriters* in some cases

A reliable *DataWriter* may have sent invalid GAP messages to a VOLATILE reliable *DataReader* after receiving a preemptive ACK or after receiving a NACK for samples that the *DataReader* had already received. These messages did not cause the *DataReader* to misbehave, but they could have led to issues when interoperating with other DDS vendors.

This issue, which was incorrectly marked as fixed in 5.3.0, is now resolved. A reliable *DataWriter* now sends correctly formatted GAPs in this case.

[RTI Issue ID CORE-7836]

5.2.3 Individual sample fragment losses may have triggered full sample repair

In some cases, losing individual fragments for a sample may have triggered the repair of the whole sample, as opposed to a repair of the missing fragments. This incorrect behavior resulted in excessive bandwidth usage.

This problem has been resolved. Now individual fragment losses are always repaired individually.

[RTI Issue ID CORE-8354]

5.2.4 DataReader may have ignored some piggyback heartbeats, leading to performance degradation

A *DataReader* may have ignored some piggyback heartbeats from a *DataWriter* when the *DataReader* used a content-filtered topic that led to writer-side filtering or when the *DataReader* created TopicQueries. By not processing the piggyback heartbeats, the *DataReader* may not have sent ACKNACK messages to the *DataWriter*. This problem may have caused the send window to fill up and the *DataWriter* to unnecessarily block on a **write()** call. This problem has been resolved.

[RTI Issue ID CORE-8908]

5.2.5 max_bytes_per_nack_response was ignored when using asynchronous publication

A *DataWriter* setting `publish_mode` to `DDS_ASYNCCHRONOUS_PUBLISH_MODE_QOS` ignored the protocol setting `max_bytes_per_nack_response`.

This meant that the number of samples that were sent in response to a NACK message from a *DataReader* was not limited based on the `max_bytes_per_nack_response` setting. This issue has been resolved.

[RTI Issue ID CORE-8973]

5.2.6 Disabling positive ACKs may have caused segmentation fault in publishing application

Disabling positive ACKs by setting `protocol.disable_positive_acks` to `TRUE` on the *DataWriter* or *DataReader* may have led to a rare segmentation fault. This problem has been fixed.

[RTI Issue ID CORE-9144]

5.2.7 Unexpected DDS_RETCODE_OUT_OF_RESOURCES error when writing a sample

The **DataWriter::write operation** may have failed with an unexpected `DDS_RETCODE_OUT_OF_RESOURCES` error when the *DataWriter* had enough resources to accept the sample that was being

written.

This error occurred only when:

- Communication was reliable.
- The *DataWriter* set `<max_samples>` to a finite value.
- Some of the *DataReaders* matching with the *DataWriter* used a `ContentFilteredTopic`.

This problem has been fixed.

[RTI Issue ID CORE-9190]

5.3 Fixes Related to Instance Management and Lifecycle

5.3.1 Instances may not have transitioned to NOT_ALIVE_NO_WRITERS on DataReader matching with a DataWriter using MultiChannelQosPolicy

Some instances may not have transitioned to NOT_ALIVE_NO_WRITERS state when a *DataReader* stopped matching with a *DataWriter* configured to use MultiChannel by setting `writer_qos.multi_channel`. This problem might have occurred only for instances that were not published in all channels. This problem has been resolved.

[RTI Issue ID CORE-8458]

5.3.2 DomainParticipantFactory::get_instance always returned success

In release 5.3.0, the operation `DomainParticipantFactory::get_instance` always returned success even if there was an error executing it. This problem has been resolved.

[RTI Issue ID CORE-8518]

5.3.3 DataReaders may not have purged samples from instances in NOT_ALIVE_NO_WRITERS state when autopurge_nowriter_samples_delay was set to finite value

A *DataReader* may not have purged samples from instances that transitioned to NOT_ALIVE_NO_WRITERS state when `reader_qos.reader_data_lifecycle.autopurge_nowriter_samples_delay` was set to a finite value.

This problem only occurred when *DataWriters* did not unregister instances explicitly by calling the unregister operation. For example:

- When a *DataWriter* was shut down gracefully or ungracefully without calling `unregister` for the instances that were in its cache.

5.3.4 Instances in NOT_ALIVE_DISPOSED state may not have been purged from DataReader queue

- When a *DataWriter* became incompatible with a *DataReader* because of a QoS change (for example, a partition).

This issue affected 5.2.x and 5.3.x releases. This problem has been fixed.

[RTI Issue ID CORE-8522]

5.3.4 Instances in NOT_ALIVE_DISPOSED state may not have been purged from DataReader queue when autopurge_disposed_instances_delay was set to zero

Instances in NOT_ALIVE_DISPOSED state without samples may not have been removed from a *DataReader* queue configured with `reader_qos.reader_data_lifecycle.autopurge_disposed_instances_delay` set to zero.

This problem was likely to occur in the following scenarios:

- Scenarios in which the *DataReader* received instances from redundant *Routing Services*.
- Scenarios in which the *DataReader* received instances from *Persistence Service*.
- Scenarios in which destination order was set to BY_SOURCE_TIMESTAMP.

Because the instances were not removed from the queue, unbounded memory growth could have occurred when `reader_qos.resource_limits.max_instances` was configured to UNLIMITED, or sample rejection may have occurred with reason SAMPLES_PER_INSTANCE_LIMIT when `reader_qos.resource_limits.max_instances` had a finite value.

This problem has been fixed.

[RTI Issue ID CORE-8538]

5.3.5 A DataReader may have failed to calculate the keyhash for a sample containing zero-length strings

A *DataReader* may have failed to calculate the keyhash for a sample containing zero-length strings. This would occur only if the strings were not a part of the key themselves and were declared before at least one of the members of the type that was part of the key.

A *DataReader* calculates the keyhash for samples coming from *DataWriters* with `DataWriterProtocolQosPolicy.disable_inline_keyhash` set to TRUE. If keyhash calculation fails, a sample may be incorrectly added to the wrong instance and the correct instance will not be found in the *DataReader*.

This issue has been resolved.

[RTI Issue ID CORE-8603]

5.3.6 Incorrect warning reported while trying to purge disposed instances proactively

When `autopurge_dispose_instances_delay` is not `DURATION_INFINITE`, *Connex DDS* attempts to purge disposed instances proactively, independently of hitting resource limits. The warning “Writer-HistoryMemoryPlugin_dropFullyAkedDisposedInstance:unregistered instances not fully acked” was logged during this process; however, this warning is valid only when resource limits are hit, not when instances are proactively removed. This problem has been resolved.

[RTI Issue ID CORE-8629]

5.3.7 Purging disposed or unregistered instances based on source timestamp does not work when internal clock is set to monotonic

This issue applied only to release 5.3.0.8 when the internal clock for the participant was set to monotonic. In this case, setting the property `dds.data_writer.history.source_timestamp_based_autopurge_instances_delay` to 1 on a *DataWriter* in order to purge disposed or unregistered instances using the source timestamp, did not work when the purging period was finite. (The purging period is configurable using the QoS values `writer_data_lifecycle.autopurge_disposed_instances_delay` and `writer_data_lifecycle.autopurge_unregistered_instances_delay`.)

This problem has been fixed.

[RTI Issue ID CORE-8637]

5.3.8 Possible leak upon application exit after using `NDDSCfgVersion::get_instance()` (Traditional C++ API only)

Tools such as Valgrind may have reported a memory leak upon application exit (not a recurring leak) when `NDDSCfgVersion::get_instance()` was called. This leak has been resolved.

[RTI Issue ID CORE-8686]

5.3.9 Unexpected errors when receiving a dispose sample for unbounded `DDS_KeyedString BuiltinType` topic

When a reader received a dispose sample without a keyhash, with a serialized key for an unbounded `DDS_KeyedString` builtin type topic, the following errors were logged and the dispose sample was not received:

```
DDS_KeyedStringPlugin_get_serialized_sample_size:value cannot be NULL
PRESCstReaderCollator_addRegisteredInstanceEntry:!serialize key
PRESCstReaderCollator_addInstanceEntry:!add registered instance
DDS_KeyedStringPlugin_get_serialized_sample_size:value cannot be NULL
PRESCstReaderCollator_addRegisteredInstanceEntry:!serialize key
PRESCstReaderCollator_addInstanceEntry:!add registered instance
```


This problem is now resolved: the dispose sample is now received and no errors are logged.

[RTI Issue ID CORE-8747]

5.3.10 instance_replacement not applied correctly for durable DataWriters

If you set <max_instances> to a finite value, <instance_replacement> may not have been applied correctly. That is, when <max_instances> was exceeded, the *DataWriter* may have replaced an instance that did not meet the replacement criteria defined in <instance_replacement>.

For example, if you set the instance replacement to DDS_DISPOSED_INSTANCE_REPLACEMENT, when <max_instances> was exceeded the *DataWriter* may have chosen for replacement an instance or multiple instances that were not in the DISPOSED state.

This problem has been resolved.

[RTI Issue ID CORE-8829]

5.3.11 Incorrect warning reported when replacing DISPOSE/ALIVE instance on a DataWriter

When a *DataWriter* was configured with an instance_replacement different than DDS_UNREGISTERED_INSTANCE_REPLACEMENT, *Connex DDS* incorrectly printed the following warning while trying to replace an instance, even if the replacement was successful:

```
WriterHistoryMemoryPlugin_dropEmptyAndFullyAkedUnregisteredInstance:no unregistered instances
```

This problem has been resolved.

[RTI Issue ID CORE-8902]

5.4 Fixes Related to Content Filters and Query Conditions

5.4.1 Possible crash in creation of a content filter for a type with an aliased base type

Applications creating an SQL content filter for a topic-type that inherits from an aliased type crashed in all the languages' APIs except for the C API.

For example, the creation of a *DataReader* with a ContentFilteredTopic based on the following IDL type "Foo" caused this problem (except in the C API):

```
struct Base { long x; };  
typedef Base BaseAlias;  
struct Foo : BaseAlias { long y; };
```

This problem has been resolved.

[RTI Issue ID CORE-8462]

5.4.2 Reading samples by instance with QueryCondition returned no data when using TOPIC or GROUP PresentationQosPolicy access_scope

An application using the PresentationQosPolicy `access_scope` TOPIC or GROUP while `ordered_access = TRUE` was not able to retrieve data using the `DataReader::take_next_instance_w_condition`, `DataReader::read_next_instance_w_condition`, `DataReader::take_instance_w_condition`, and `DataReader::read_instance_w_condition` APIs when the condition being used was a *QueryCondition*. In this case, these APIs returned `DDS_RETCODE_NO_DATA` when in fact there was data matching the requested instance and condition. (The data could be retrieved successfully using other `read` or `take` APIs.) This issue has been resolved.

[RTI Issue ID CORE-8508]

5.4.3 Content Filter issue when filtering on a member of the base type, for types with inheritance

When using a Content Filter on a type with inheritance, if the filter expression required evaluating a member of the base type, the result of the filter was incorrect. This issue has been fixed. Now the filter expression is applied properly, and the result is correct.

[RTI Issue ID CORE-9035]

5.5 Fixes Related to TopicQueries

5.5.1 DataWriter may have deadlocked if receiving "continuous" TopicQueries

A race condition may have caused a thread to deadlock when a *DataWriter* dispatched a continuous TopicQuery. This problem stopped the *DataWriter* from operating normally.

This deadlock could happen only if `topic_query_dispatch.samples_per_period`, in the *DataWriterQos*, was set to unlimited (the default value), and the *DataWriter* received a TopicQuery with continuous selection kind (the default was history snapshot, not continuous).

This problem affected release 5.3.0.8 only and has been fixed in this release.

[RTI Issue ID CORE-8448]

5.5.2 Instances may not have transitioned to NOT_ALIVE_NO_WRITERS on DataReader matching with a DataWriter with TopicQuery enabled

Some instances may not have transitioned to `NOT_ALIVE_NO_WRITERS` state when a *DataReader* stopped matching with a *DataWriter* with TopicQuery enabled (`writer_qos.topic_query_dispatch.enable = TRUE`). For example, this problem could have occurred if the *DataWriter* was unmatched due to a partition change. This problem has been fixed.

[RTI Issue ID CORE-8452]

5.5.3 Error when changing partition, group data, or topic data on a Publisher containing a *DataWriter* with TopicQuery enabled

Changing a partition or the group data on a Publisher containing a *DataWriter* with TopicQuery enabled (`writer_qos.topic_query_dispatch.enable = TRUE`) caused the following error:

```
PRESPsService_assertRemoteEndpoint:!assert pres psRemoteWriter  
PRESPsService_notifyOfGroupDataOrPartitionChange:!assert remote endpoint
```

The issue also occurred with a slightly different error message when the topic data for the Topic associated with the Publisher's *DataWriter* was updated.

This problem has been fixed.

[RTI Issue ID CORE-8453]

5.5.4 Possible increasing memory and CPU usage in publishing applications using TopicQueries

A publishing application that enables TopicQuery may have been subject to increasing CPU and memory usage.

This growth was in linear proportion to the preceding number of TopicQueries times the preceding number of samples selected, but growing with respect to a sample only for as long as such a sample remained present in the *DataWriter*'s history. That is, removing a sample from the history would reset the contribution to this growth attributed to that sample.

The *DataWriter* keeps state for each TopicQuery it receives. This state can be deleted when all the samples that the TopicQuery selects have been delivered and acknowledged; however, some state information was not removed until the samples were replaced (for example, when updating an instance and exceeding the history depth).

This problem has been resolved.

[RTI Issue ID CORE-8817]

5.5.5 Spurious log message related to TopicQuery has been removed

The following log message may have been incorrectly displayed with "local" verbosity:

```
PRESPsService_dispatchMatchingTopicQueries:ignoring TopicQuery for this DataWriter because it  
doesn't enable them
```

This problem happened when a *DataWriter* discovered a new *DataReader*, regardless of the existence of a TopicQuery.

The log message will no longer be displayed in that situation. It will only be displayed when a *DataWriter* that doesn't enable TopicQueries receives a TopicQuery.

[RTI Issue ID CORE-8820]

5.5.6 `DataReader::getKey` did not work with TopicQueries

The `DataReader::getKey` API may have returned `DDS_RETCODE_BAD_PARAMETER` for an instance that was known to a *DataReader* if samples for that instance were only in the *DataReader's* queue in response to a TopicQuery.

This problem has been resolved. The API now searches for the instance in the queues associated with live data as well as TopicQuery data.

[RTI Issue ID CORE-8953]

5.5.7 TopicQuery samples that failed to be written a first time may have never been sent

If the write operation that publishes TopicQuery samples failed (for example, due to a timeout because of blocking), the sample that failed to be written may never have been sent. This error would happen in one of two cases, given a sample with sequence number n that failed to be written:

- The next sample in the writer queue that matched the TopicQuery expression had a sequence number $> n + 1$.
- The sample that failed to be written was the last TopicQuery sample in the writer's queue, and the last sample in the writer's queue at the time of the write error had a sequence number $> n + 1$.

This issue has been resolved. Now if the internal write call for a TopicQuery sample fails, the sample will be sent during the next TopicQuery publication period.

[RTI Issue ID CORE-9188]

5.6 Fixes Related to DynamicData

5.6.1 DynamicData had limitations with members larger than 65,535 bytes

DynamicData did not support resizing or out-of-order assignment of members that were longer than 65,535 bytes. Doing one of these two operations would result in the following error:

```
sparingly stored member exceeds 65535 bytes
```

There are no longer any restrictions regarding a member's size when using any of the DynamicData APIs in any order.

[RTI Issue ID CORE-3177]

5.6.2 Copying into a bound DynamicData object did not work

Using the `DynamicData::copy` operation did not work properly when copying into a `DynamicData` object that was bound to another object. The operation may have caused corruption of the destination object.

This issue has been resolved.

[RTI Issue ID CORE-3385]

5.6.3 Corrupt DynamicData objects containing sequences with length 0

In some cases, a `DynamicData` object may have been corrupted if a sequence member with length 0 was set in the object. This would only occur with sequences of primitives with a size greater than 4: `double`, `long long`, `unsigned long long`, `long double`. The corruption would have been noticeable when set and get operations performed on the `DynamicData` object returned errors or incorrect values.

This issue has been resolved.

[RTI Issue ID CORE-5161]

5.6.4 Binding to members of a sequence incorrectly created members in the DynamicData API

Binding to member `n` of a sequence in a `DynamicData` object that did not previously exist incorrectly created members 0 through `n-1` of the sequence, even when the member was unbound without setting any values.

This issue has been resolved. Now, if a member of a sequence is bound and unbound without calling any set APIs, the length of the sequence that was bound to remains unchanged.

[RTI Issue ID CORE-5800]

5.6.5 Error when unbinding from a union DynamicData object

In some rare situations, unbinding from a union member may have returned `RETCODE_ERROR` with the following error message:

```
DDS_DynamicData_unbind_complex_member:internal error 1 trying to to stream
```

This issue has been resolved.

[RTI Issue ID CORE-6657]

5.6.6 Performance degradation when setting large sequences using the DynamicData API

If a DynamicData object's type contained any optional members, setting a large sequence (tens or hundreds of thousands of elements) using any of the **DynamicData::set_*_seq** APIs may have taken a very long time, on the order of minutes. The sequence itself did not have to be an optional member in order for this issue to occur.

This issue has been resolved.

[RTI Issue ID CORE-6979]

5.6.7 DynamicData::is_member_key did not work for types using inheritance

The **DynamicData::is_member_key** API did not work correctly for members that were part of a base type of the DynamicData object's type. The API may have returned incorrect results for whether or not a member is a key.

This issue has been resolved.

[RTI Issue ID CORE-7505]

5.6.8 DynamicData::set_complex_member did not work with aliased typecodes

The **DynamicData::set_complex_member** API failed with `RETCODE_ERROR` and an error message similar to the following:

```
DDS_DynamicData_set_complex_member:type mismatch for field aField (id=1)
```

if the complex member being set had a TypeCode with kind `TK_ALIAS`.

This issue has been resolved.

[RTI Issue ID CORE-7561]

5.6.9 DynamicData::set_string API did not accept NULL strings

The **DynamicData::set_string** API did not accept a NULL string. This behavior was not parallel with generated types, which allowed setting and sending NULL strings.

This issue has been fixed. The **DynamicData::set_string** API now accepts a NULL string. A string with length 0 containing only the null-terminator will be set in the DynamicData object.

[RTI Issue ID CORE-8163]

5.6.10 Large memory allocation when binding to large sequences in the DynamicData API

There was a large memory allocation when calling the `DynamicData::bind_complex_member` API on a large sequence of non-primitive members. More concretely, the allocation was (4 * max length of the sequence) bytes.

This issue has been resolved. Binding to a sequence no longer requires this memory allocation.

[RTI Issue ID CORE-8358]

5.6.11 DynamicData::from_cdr_buffer API did not resize DynamicData object

The `DynamicData::from_cdr_buffer` API did not resize the DynamicData buffer. This meant that the API failed if the same DynamicData object was used in repeated calls to the API in which the CDR buffer was larger than in previous calls. The following errors were printed in these situations:

```
DDS_DynamicDataTypePlugin_parametrized_cdr_to_cdr:deserialization error: insufficient space
DDS_DynamicDataTypePlugin_parametrized_cdr_to_cdr:error converting from extended CDR to CDR
DDS_DynamicDataTypePlugin_deserialize:error converting from extended CDR to CDR
DDS_DynamicData_from_cdr_buffer:deserialization error: buffer
DDS_DynamicData_from_cdr_buffer error 1
```

This issue has been resolved. The `DynamicData::from_cdr_buffer` will resize the DynamicData object's internal buffer if required to accommodate the deserialized sample.

[RTI Issue ID CORE-8394]

5.6.12 Error when unbinding from a DynamicData object

In rare situations, unbinding from a member may have returned `RETCODE_ERROR` with the following error message:

```
DDS_DynamicData_unbind_complex_member:internal error 1 trying to assert complex member
```

This issue has been resolved.

[RTI Issue ID CORE-8386]

5.6.13 Accessing a member of an array or sequence by member ID failed for member IDs > 65535

Using any of the `DynamicData::get_*` APIs on a sequence or array and providing a member ID greater than 65535 failed with `DDS_RETCODE_NO_DATA`, even if the member existed.

This issue has been resolved.

[RTI Issue ID CORE-8561]

5.6.14 DynamicData DataReader may have failed to correctly deserialize key for types containing strings

If a DynamicData *DataReader* received a keyed sample with a zero-length string, it may have failed to correctly deserialize the key. This problem occurred if the string was not part of the key and came before at least one of the key members in the type definition. When this issue occurred, the following error would be printed:

```
"DDS_DynamicDataUtility_skip_compact_type:stream error trying to access string"
```

This issue has been resolved.

[RTI Issue ID CORE-8604]

5.6.15 Missing DynamicData::clear_optional_member API in the .NET API

In previous releases, the **DynamicData::clear_optional_member** API was missing from the .NET API. This issue has been resolved.

[RTI Issue ID CORE-9077]

5.6.16 DynamicData::clear_optional_member API incorrectly returned error for unset optional members

Calling the **DynamicData::clear_optional_member** API on an unset optional member returned **PRECONDITION_NOT_MET** and produced the following error message:

```
DDS_DynamicData_clear_optional_member:cannot clear non-optional members
```

This issue has been resolved. The **DynamicData::clear_optional_member** now returns **RETCODE_OK** when called on an already unset optional member.

[RTI Issue ID CORE-9092]

5.6.17 DynamicData APIs did not check for an associated TypeCode

The DynamicData APIs were not consistent in checking for a TypeCode associated with the DynamicData object before trying to execute the operation. This inconsistency resulted in unclear error messages and inconsistent return codes.

This issue has been resolved. DynamicData APIs that require the DynamicData object to have an associated TypeCode now check for a TypeCode and return **RETCODE_PRECONDITION_NOT_MET** when there is no TypeCode associated with the DynamicData object.

[RTI Issue ID CORE-9093]

5.6.18 Setting members past a sequence's maximum bound was not prohibited in the DynamicData API

Setting a member past a sequence's maximum bound was not prohibited in the DynamicData API. An attempt to set a member past the sequence's maximum bound now returns an error.

[RTI Issue ID CORE-9174]

5.6.19 DynamicData::clear_all_members API did not work on bound DynamicData objects

The **DynamicData::clear_all_members** API did not work on bound DynamicData objects. If you called **DynamicData::clear_all_members** on a bound DynamicData object, unbound that object, and then bound the object again to the same member, all of the previous values were still set, as if the **DynamicData::clear_all_members** API had never been called.

Note: In the Modern C++ API, binding is equivalent to loaning a DynamicData object.

This issue has been resolved.

[RTI Issue ID CORE-9175]

5.7 Fixes Related to Transports

5.7.1 Communication between kernel and RTP Mode Participants with shared memory transport not working on 64-bit VxWorks 6 platforms

Applications could not communicate between kernel and RTP mode using the shared memory transport on 64-bit VxWorks 6 systems. (This is not an issue for 64-bit VxWorks 7 systems.) This problem has been resolved.

As a result of this fix, applications built with older *Connex DDS* releases cannot communicate with applications built with this release if the communication uses shared memory on 64-bit VxWorks 6 systems and occurs between kernel and RTP mode. (This configuration is not a common use case.)

[RTI Issue ID CORE-8171]

5.7.2 Network interface tracker may have reported non-existing changes

In some configurations in which several network interfaces share the same IP address, the network interface tracker may have reported non-existing changes on the interfaces. This problem has been resolved.

[RTI Issue ID CORE-8205]

5.7.3 Possible continuous failure to send over shared memory transport

When using the shared memory transport and rapidly creating and deleting *DomainParticipants*, it was possible for a *DataWriter* or *DataReader* to continuously fail to send content. The following log message was generated with the `NDDS_CONFIG_LOG_VERBOSITY_STATUS_ALL` verbosity and `NDDS_CONFIG_LOG_CATEGORY_COMMUNICATION` category:

```
RTINetioSender_send: send failed: locator:
shmem://0000:0702:0007:0000:0000:0000:0000:10414
```

(The specific locator numbers may differ.)

This problem has been resolved.

[RTI Issue ID CORE-8295]

5.7.4 Race condition in shared memory transport led to cleanup failure

A race condition in the shared memory transport occurred when two processes attempted to create a *DomainParticipant* with the same domain ID and participant ID at the same time. This race condition led to cleanup problems when one of the *DomainParticipants* was being deleted. It would generate these error messages:

```
NDDS_Transport_Shmem_destroy_recvresource_rrEA:!take mutex
NDDS_Transport_Shmem_destroy_recvresource_rrEA:!give mutex
RTIOsapiSharedMemoryMutex_delete:OS semctl() failure, error 0X16: Invalid argument
```

This race condition has been fixed. These errors no longer appear.

[RTI Issue ID CORE-8534]

5.7.5 UDPv4/UDPv6 transport creation failed when setting `send_socket_buffer_size` or `recv_socket_buffer_size` to `NDDS_TRANSPORT_UDPV4_SOCKET_BUFFER_SIZE_OS_DEFAULT` or `NDDS_TRANSPORT_UDPV6_SOCKET_BUFFER_SIZE_OS_DEFAULT`

In release 5.3.0, there was an issue provoking UDPv4 and UDPv6 transport creation to fail if any of the `send_socket_buffer_size` or `recv_socket_buffer_size` properties was set to either `NDDS_TRANSPORT_UDPV4_SOCKET_BUFFER_SIZE_OS_DEFAULT` or `NDDS_TRANSPORT_UDPV6_SOCKET_BUFFER_SIZE_OS_DEFAULT`.

This problem is now resolved.

[RTI Issue ID CORE-8585]

5.7.6 Unexpected shared memory locator when setting `rtps_host_id` to a value different than `DDS RTPS_AUTO_ID`

In *Connex DDS 5.3.0*, the generated shared memory locator was incorrect when the `rtps_host_id` value was different than `DDS RTPS_AUTO_ID`. As a result, 5.3.0 may have failed to communicate over shared memory with previous versions of the product.

Starting with *Connex DDS 6.0.0*, the shared memory locator is generated as follows:

- If the `dds.transport.shmem.builtin.host_id` property is defined, the shared memory locator is derived from this value.
- If the Wire Protocol `rtps_auto_id_kind` is `DDS RTPS_AUTO_ID_FROM_IP` and `rtps_host_id` is different than `DDS RTPS_AUTO_ID`, the shared memory locator is derived from the `rtps_host_id` value.
- In all other cases, the shared memory locator is derived from the MAC address, IP address, or UUID value depending on the `rtps_auto_id_kind` value.

[RTI Issue ID CORE-8738]

5.7.7 Crash when network interface changed before transport was fully created

A crash may have occurred in IP Mobility scenarios if a network interface changed before the transport creation was completed. This problem has been resolved. Now when the network interface tracker starts, it can properly handle changes in the interfaces.

[RTI Issue ID CORE-8869]

5.7.8 Wrong transport class name when logging 'No interfaces' warning

The logging of a warning due to the absence of valid interfaces in TCP Transport, in `TCPv4_LAN` or `TLSv4_LAN` mode, triggered the following two messages:

```
NDDS_Transport_UDP_get_class_name_cEA:!family parameter not valid
NDDS_Transport_IP_selectValidInterfaces:WARNING: No interfaces for transport (null) match
allowed patterns
```

This behavior was wrong: it did not properly show the transport class, and, although the message said `WARNING`, it actually logged an exception. This problem is now fixed: the message now properly shows the transport class, and it logs a warning instead of an exception.

[RTI Issue ID COREPLG-444]

5.8 Fixes Related to Logging and Distributed Logger

5.8.1 Segmentation fault when simultaneously changing log files and writing to a log file

The functions `NDDS_Config_Logger_set_output_file_set()`, `NDDS_Config_Logger_set_output_file()`, `NDDS_Config_Logger_set_output_file_name()`, and `DDS_DomainParticipantFactory_set_qos()` were not thread-safe. If any of these functions was called to change the logging output file(s) while a separate thread was writing a log message, the writing thread may have crashed with a segmentation fault. This problem has been resolved. These functions are now thread-safe.

[RTI Issue ID CORE-8710]

5.8.2 Unexpected "PRESCstReaderCollator_addInstanceEntry:exceeded max total instances" message at WARNING level

You may have seen the following message when using the `NDDS_CONFIG_LOG_VERBOSITY_WARNING` verbosity level:

```
PRESCstReaderCollator_addInstanceEntry:exceeded max total instances.
```

This message should not be a warning, since exceeding `max_total_instances` is expected when `reader_qos.resource_limits.max_instances` is finite, `reader_qos.resource_limits.max_total_instances` is AUTO (default), and the *DataReader* receives more than `reader_qos.resource_limits.max_instances`.

The verbosity level at which this message is printed has been changed to `NDDS_CONFIG_LOG_VERBOSITY_STATUS_ALL`.

[RTI Issue ID CORE-8822]

5.8.3 Segmentation fault when attempting to write to one file of a file set if that file failed to be opened

After an application called the `NDDS_Config_Logger_set_output_file_set()` function, if one of the files failed to be opened (let's call it file N), the first attempt to write to that file failed gracefully (which was correct behavior), but the second attempt to write to that file resulted in a segmentation fault. This problem has been resolved. Now, the second attempt to write will be to file N+1 rather than file N.

[RTI Issue ID CORE-9181]

5.8.4 `NDDS_Config_Logger_get_output_device()` always returned NULL after DomainParticipant creation

`NDDS_Config_Logger_get_output_device()` always returned NULL after the first *DomainParticipant* was created. This problem only affected 5.3.x releases and has been fixed. `NDDS_Config_Logger_get_output_device()` will now return the correct device.

[RTI Issue ID CORE-9202]

5.8.5 Unable to configure Distributed Logger profile Using C++ DomainParticipantFactory

C++ *Distributed Logger* was not using the C++ DomainParticipantFactory internally, but using the C factory. This led to issues like these:

- The operations `RTI_DLOptions::setQosLibrary` and `RTI_DLOptions::setQosProfile` could not be used to configure a QoS profile for *RTI Distributed Logger* while using the C++ API. This problem only occurred when the application using *Distributed Logger* loaded its XML configuration using the QoS values `DomainParticipantFactoryQos.profile.url_profile` or `DomainParticipantFactoryQos.profile.string_profile`.
- If *Distributed Logger* used its own DomainParticipant, the threads associated with this DomainParticipant were not created using the ThreadFactory installed on the C++ DomainParticipantFactory with the API `DDS_DomainParticipantFactory::set_thread_factory`.
- Because *Distributed Logger* was using the C DomainParticipantFactory when creating its own DomainParticipant, you may have been forced to finalize this DomainParticipantFactory (using `DDS_DomainParticipantFactory_finalize_instance`), in addition to the C++ DomainParticipantFactory to avoid a memory leak.

This problem has been resolved.

[RTI Issue ID DISTLOG-155]

5.8.6 Possible segmentation fault when deleting options that were used to create Distributed Logger instance

An application using *Distributed Logger* may have issued a segmentation fault if the options used to create it were deleted before the *Distributed Logger* instance was finalized. This problem has been resolved.

[RTI Issue ID DISTLOG-157]

5.8.7 Error code for setApplicationKind not captured in Distributed Logger C example

The error code was not captured in the *Distributed Logger* C example, when the options were configured with an application kind. This problem has been resolved.

[RTI Issue ID DISTLOG-159]

5.8.8 Logger settings from QoS File not set correctly when using Distributed Logger

If your application created a Distributed Logger instance before loading the QoS profiles, the verbosity settings specified in the QoS file were overwritten with the default settings. This problem has been resolved.

[RTI Issue ID DISTLOG-175]

5.8.9 "Max queue size reached, message will get lost" warning removed from DistributedLogger

Distributed Logger used to print a warning when a message could not be added to the internal messages queue because the queue was full. This warning has been removed. The warning is no longer necessary because there is a new field in `com::rti::dl::LogMessage` called `messageId` that can be used to detect if messages are lost.

[RTI Issue ID DISTLOG-194]

5.9 Fixes Related to XML Configuration

5.9.1 Some XML example files were not compliant with target XSD schema

The following XML example files were not compliant with their target XSD schemas:

- MONITORING_DEMO.xml
- RTIDDSPING_QOS_PROFILES.example.xml
- RTIDDSPY_QOS_PROFILES.example.xml
- RTI_RECORDING_SERVICE.xml
- USER_RECORDING_SERVICE.xml
- USER_ROUTING_SERVICE.xml

This problem has been resolved; these example files are now compliant.

[RTI Issue ID CORE-8259]

5.9.2 NDDS_QOS_PROFILES.xml not loaded from default location

Connex DDS did not load a file with the name `NDDS_QOS_PROFILES.xml` from the default location of `resource/xml`. This problem has been resolved. Now, if such a file exists in that location, *Connex DDS* will automatically load the file.

[RTI Issue ID CORE-8279]

5.9.3 is_default_participant_factory_profile="true" was ignored when <participant_factory_qos> only had the base_name attribute specified

If a <participant_factory_qos> didn't have any tags set in XML and specified just a base_name, the metadata information that indicated what fields were set in XML for the parent weren't copied as a part of the single inheritance code path. This led to problems loading the correct value when the **is_default_participant_factory_profile="true"** attribute was used for a <qos_profile> that contained this <participant_factory_qos>. The reason was that other functions relied on the metadata to detect if something other than the default value was being set.

This problem has been resolved in this release, for both single and multiple QoS Profile inheritance via XML. The metadata is now copied during inheritance.

[RTI Issue ID CORE-8364]

5.9.4 Creating a typecode from an XML with a directive tag may have caused a crash

Creating a typecode from an XML file may have caused a segmentation fault if the tag "directive" was used. This problem occurred when calling the function **DDS_TypeCodeFactory_create_tc_from_xml_file**. This problem has been resolved.

[RTI Issue ID CORE-8493]

5.9.5 XML parser error if <domain_library> or <domain_participant_library> split into multiple tags with same name attribute

The *Connex* DDS XML parser reported an error if a <domain_library> or <domain_participant_library> appeared more than once with the same value in its name attribute. This behavior has changed. Now, a domain or participant library can be split into multiple <domain_library> or <domain_participant_library> tags, respectively.

The resulting behavior is equivalent to defining a single library containing all the elements specified in each library, with the same value in the name attribute.

[RTI Issue ID CORE-8294]

5.9.6 Connex DDS XML parser failed to parse a const char with '\0' value

The Connex DDS XML parser failed with the following error when parsing a const char '\0':

```
DDS_ExpressionEvaluator_get_next_token:expression parse error at column 2: invalid scape character
DDS_ExpressionEvaluator_evaluate:expression parse error at column 2: empty expression
DDS_XMLConst_evaluate_expression:Parse error at line 3: error evaluating const expression
DDS_XMLConst_initialize:error evaluating const expression
DDS_XMLConst_new:!init XML const object
```

```
RTIXMLParser_onStartTag:Parse error at line 3: Error processing tag 'const'
```

Now the parser can parse this character.

[RTI Issue ID CORE-8802]

5.9.7 Value checks were not enforced for `writer_qos.writer_resource_limits.instance_replacement` field in `DataWriterQos`

The *RTI Connext DDS XML* parser did not produce an error if the value specified for `writer_qos.writer_resource_limits.instance_replacement` wasn't one of the standard values specified in the documentation. (In the API Reference HTML documentation, search for **DDS_DataWriter-ResourceLimitsInstanceReplacementKind** for a list of the accepted values.) The parser now produces an error if the value is not one of these values.

[RTI Issue ID CORE-8810]

5.9.8 Last, not first, value was retrieved when getting a QoS by specifying `topic_filter`

A problem occurred when a Profile had multiple QoSes of the same type, with the `topic_filter` attribute either not specified (NULL) or set to `**`. In this case, API calls (such as **DDS_DomainParticipantFactory_get_datareader_qos_from_profile_w_topic_name()**) retrieved the wrong QoS values when passing those values to the appropriate function calls (such as **DDS_XMLQosProfile_get_datareader_dds_qos_filtered()**) in fallback scenarios only.

The fallback scenarios are the following:

- When a `topic_filter` argument is set to NULL, a lookup API matches against the first QoS type in the Profile with the `topic_filter` attribute set to NULL. When such a QoS does not exist, the fallback scenario is for the API to return the first matching QoS type (in the case of duplicates) with `topic_filter` set to `**`.
- Likewise, when a `topic_filter` argument is set to a valid string value, the lookup API matches against the first QoS type in the Profile with the `topic_filter` attribute set to a matching pattern value. If such a QoS does not exist, the fallback scenario is for the API to return the first matching QoS type (in the case of duplicates) with `topic_filter` set to NULL.

In both of these fallback scenarios, the problem was that the API was not returning the first matching QoS type, but the last one.

This problem is now resolved. In both cases, the API now correctly returns the first matching QoS type.

[RTI Issue ID CORE-9024]

5.9.9 Unable to set `WriterDataLifecycle::autopurge_disposed_instances_delay` for builtin `DataWriters` via XML

Due to an error in the shipped DTD, it was not possible to set `WriterDataLifecycle::autopurge_disposed_instances_delay` for the builtin `DataWriters` via XML. Attempting to do so would result in errors during participant creation similar to the following:

```
[CREATE Participant] RTIXMLParser_validateOnStartTag:Parse error at line 113: Unexpected tag 'autopurge_disposed_instances_delay'
[CREATE Participant] RTIXMLParser_parseFromFile_ex:Parse error in file 'USER_QOS_PROFILES.xml'
```

This issue has been resolved.

[RTI Issue ID CORE-9082]

5.10 Fixes Related to XML-Based Application Creation

5.10.1 Sequences defined in XML with `sequenceMaxLength` of -1 were not interpreted as unbounded

If a type defined in XML had its sequences or strings set to -1 in the attributes `sequenceMaxLength` or `stringMaxLength`, the resulting sequence or string length was interpreted as being of a default size rather than unbounded. This problem has been resolved. The resulting length is now interpreted as unbounded $((2^{31})-1$ bytes).

[RTI Issue ID CORE-8494]

5.10.2 XML-Based Application Creation in Java didn't work with some `TypeCodes`

Types larger than 65,535 bytes and types that use extensible type features only available in `TypeObject` weren't properly serialized when using XML-Based Application Creation in Java. This issue caused `DomainParticipant` creation to produce a `BAD_TYPECODE` exception for large types, and could have prevented communication between Java applications and applications using a different language binding for types using extensibility features. This issue has been resolved.

[RTI Issue ID CORE-CORE-8957]

5.11 Fixes Related to OMG Specification Compliance

5.11.1 Connex DDS not compliant with RTPS 2.2

Previous releases of `Connex DDS` were not compliant with the OMG RTPS 2.2 specification. In particular, the GUID announced by default was not compliant with the RTPS 2.2 specification (as described in [5.11.2 Default GUID not compliant with RTPS specification on the next page](#)).

This release is now fully compliant with RTPS 2.2. In addition, the RTPS wire protocol version that *Connex DDS* announces in messages it puts on the wire has been updated to 2.3, since this release is also fully compliant with the RTPS 2.3 wire protocol (described in the RTPS 2.2 and DDS Security 1.1 specifications).

Note that this release also supports some of new RTPS 2.4 wire protocol features. One example is the ability to announce QoS for builtin endpoints (for more information about this, please refer to [5.11.3 Manual by Participant Liveliness stopped working when communicating with old RTI Connex Professional versions or non-RTI DDS implementations below](#)).

[RTI Issue ID CORE-6902]

5.11.2 Default GUID not compliant with RTPS specification

The default value of `DDS_DomainParticipantQos::wire_protocol::rtps_auto_id_kind` was `DDS RTPS_AUTO_ID_FROM_IP`. This value caused the `rtps_host_id` to be the IP address by default, which was not compliant with this statement in the RTPS specification:

To comply with this specification, implementations of the RTPS protocol shall set the first two bytes of the `guidPrefix` to match their assigned `vendorId`.

This problem has been resolved by changing the default value of `rtps_auto_id_kind` to `DDS RTPS_AUTO_ID_FROM_UUID`. The UUID always has the first two bytes equal to the RTI RTPS `vendorId`. Note that `DDS RTPS_AUTO_ID_FROM_UUID` is currently the only value that will result in a specification-compliant GUID.

[RTI Issue ID CORE-8033]

5.11.3 Manual by Participant Liveliness stopped working when communicating with old RTI Connex Professional versions or non-RTI DDS implementations

In release 5.2 and above, Manual by Participant Liveliness may have stopped working when communicating with remote participants in *Connex DDS Professional* 5.1 or below or when communicating with non-RTI remote participants. Note that this issue was only triggered when setting `DiscoveryConfig Qos's participant_message_reader_reliability_kind` to `BEST_EFFORT_RELIABILITY` (default). In particular, the problem was provoked by a wrong configuration of the Reliability QoS used when communicating with remote liveliness builtin endpoints.

This problem is now resolved by implementing the RTPS 2.3 specification's `BuiltinEndpointQos_t` propagation mechanism, which ensures the proper QoS is selected when communicating with remote liveliness builtin endpoints. (RTPS 2.3 defines version 2.4 of the wire protocol.)

[RTI Issue ID CORE-8762]

5.11.4 NACK_FRAG message not compliant with RTPS specification

The NACK_FRAG message was not compliant with the format specified in the RTPS specification. This may have led to interoperability issues with other DDS vendors.

This problem has been fixed.

[RTI Issue ID CORE-9193]

5.12 Fixes Related to Vulnerabilities

This release fixes some potential vulnerabilities, including RTI Issue IDs CORE-8581, CORE-8645, CORE-8868, and CORE-9038.

5.13 Fixes Related to Modern C++ API

5.13.1 Compilation failure when NDDS_USER_DLL_EXPORT defined on non-Windows platform

The preprocessor definition NDDS_USER_DLL_EXPORT is only intended for Windows systems and should be ignored on other platforms. But that was not the case in *Connex DDS 5.3.0*, where this situation could have caused a compilation failure in the Modern C++ API. This problem has been resolved.

[RTI Issue ID CORE-8176]

5.13.2 Downcasting a condition into a ReadCondition not supported

In the Modern C++ API, a Reference Type acting as a base class (for example, `dds::core::Entity`) can be downcasted to a child type (for example, `dds::domain::DomainParticipant`) using `dds::core::polymorphic_cast` (analogous to `std::dynamic_pointer_cast`).

That was the intended behavior with `dds::core::Condition` and `dds::sub::ReadCondition`, as well as `dds::sub::QueryCondition`; however, the following code did not compile:

```
dds::core::Condition c = ...;
auto rc = dds::core::polymorphic_cast<dds::sub::ReadCondition>(c);
auto qc = dds::core::polymorphic_cast<dds::sub::QueryCondition>(rc);
```

This problem has been resolved, and the previous casts now work.

[RTI Issue ID CORE-8347]

5.13.3 ReadCondition handler may not have been dispatched

One-argument condition handlers used to create *ReadConditions* (or *QueryConditions*) were not called by `WaitSet::dispatch()`.

Other types of *Conditions* were not affected, and no-argument handlers worked in all cases.

This problem has been resolved.

[RTI Issue ID CORE-8352]

5.13.4 Compilation error accessing a const LoanedSamples instance

Attempting to access a const LoanedSamples may have caused a compilation error. For example, the following code didn't compile, although it should be legal:

```
void print_samples(const LoanedSamples<Foo>& samples)
{
    for (auto&& s : samples) {
        std::cout << s.data() << std::endl;
    }
}
```

This problem has been resolved. Now it is possible to iterate through a const LoanedSamples using its `const_iterator`.

[RTI Issue ID CORE-8423]

5.13.5 Missing unregister_thread function

Release 5.2.0 added the function `unregister_thread()`, which allows releasing *Connex DDS* thread-local memory. The Modern C++ API did not provide this function.

This release adds the function `rti::core::unregister_thread()` and the utility type `rti::core::UnregisterThreadOnExit`, which calls `unregister_thread` in its destructor.

[RTI Issue ID CORE-8425]

5.13.6 dds::sub::Sample creation or assignment from a LoanedSample failed when data was invalid

The creation or assignment of a `dds::sub::Sample` from a LoanedSample (the value type of LoanedSamples) threw an exception if the LoanedSample's data was invalid. This was not the expected behavior, since Sample can hold invalid data.

This problem has been resolved, and now this is the behavior:

- The constructor `Sample<T>(const LoanedSample<T>& ls)` will default-construct the sample data when `ls.info().valid()` is false.
- The assignment operator `operator=(const LoanedSample<T>& ls)` will not modify the sample data (only the sample info) when `ls.info().valid()` is false.

[RTI Issue ID CORE-8446]

5.13.7 Possible symbol collision with Boost compiling the Modern C++ API

An application using both Boost and the Modern C++ API may have failed to compile due to a symbol collision. The problem happened only if the application included (directly or indirectly) the file `boost/detail/iterator.hpp`. This problem has been resolved.

[RTI Issue ID CORE-8536]

5.13.8 `dds::topic::find` could not return `AnyTopic`

The function `dds::topic::find` can be used to return a typed topic (`Topic<Foo>`), but was intended to allow returning an `AnyTopic`. In previous releases, `dds::topic::find<AnyTopic>(…)` didn't compile.

This problem has been resolved.

[RTI Issue ID CORE-8618]

5.13.9 Inconsistency between XML format and `QosProvider::type()` arguments

The `QosProvider` allows loading `DynamicTypes` from an XML document. The function to retrieve them, `type()`, required two arguments: a library name and a type name. But the XML format to define types doesn't require a library name.

This inconsistency has been resolved by adding an overload whose only argument is the type name. See an example in the Modern C++ API Reference HTML documentation, under "Modules > Programming How-To's > DynamicType and DynamicData Use Cases > Create a DynamicType from an XML description."

[RTI Issue ID CORE-8805]

5.13.10 Missing accessors for `Liveliness::assertions_per_lease_duration`

In the Modern C++ API, the type `dds::core::policy::Liveliness` did not provide an accessor to the extension field `assertions_per_lease_duration`.

This problem has been resolved, and a getter and a setter are now provided:

```
dds::core::policy::Liveliness liveliness;
liveliness.extensions().assertions_per_lease_duration(3);
assert(liveliness.extensions().assertions_per_lease_duration() == 3);
```

[RTI Issue ID CORE-8833]

5.13.11 Entity Listeners may have missed some notifications

A race condition during the creation of an *Entity* (*DomainParticipant*, *DataReader*, etc.) with a *Listener* may have caused the *Listener* to miss a notification for a status change that occurred right before the

Entity's constructor finished. This problem has been resolved.

[RTI Issue ID CORE-8905]

5.14 Other Fixes

5.14.1 strict-aliasing warnings when compiling code generated from an IDL with floats or enumerations

When compiling code generated from an IDL with floats or enumerations using the GCC compiler flag **-Wstrict-aliasing**, you would see strict-aliasing warnings that contained the following:

```
warning: dereferencing type-punned pointer will break strict-aliasing rules
note: in expansion of macro 'RTICdrStream_serializeFloat'
```

This problem has been resolved. These compiler warnings no longer appear.

[RTI Issue ID CORE-6778]

5.14.2 Unexpected sample loss notification on non-VOLATILE DataReader

A new *DataReader* that you have just created and enabled, that is configured with non-VOLATILE durability, may have incorrectly reported sample losses upon matching with a non-VOLATILE *DataWriter*.

This issue occurred when the *DataWriter* was configured with **finite autopurge_unregistered_instances_delay** or **autopurge_unregistered_instances_delay** and some of the instances were removed due to these delays before the *DataReader* was matched with the *DataWriter*.

This problem has been fixed.

[RTI Issue ID CORE-7841]

5.14.3 RTPS messages with wrong alignment incorrectly accepted

RTPS messages and submessages with wrong alignment may have been incorrectly accepted by *Connex* DDS.

This problem is now resolved. By default these messages with wrong alignment are now dropped, and the following message is logged:

```
MIGInterpreter_parse:submessage not aligned to 4
```

Note this behavior can be changed by setting the following property on the Participant PropertyQos to true: **dds.participant.use_45d_compatible_alignment_enforcement**.

[RTI Issue ID CORE-8060]

5.14.4 Segmentation fault while looking up vendor-specific topics from a traditional C++ or .NET application

When monitoring libraries were enabled in a traditional C++ or .NET application, the use of **lookup_topicdescription()** or **find_topic()** to get a monitoring topic led to a segmentation fault. This problem has been resolved. These functions now successfully retrieve the monitoring topic.

Similar problems existed for the *Distributed Logger* topics and the *Security Plugins* Builtin Logging topic. These problems have also been resolved.

[RTI Issue ID CORE-8256]

5.14.5 Memory leak when failing to create builtin endpoints

Failure by *Connex DDS* to create certain builtin endpoints resulted in a memory leak. The affected endpoints were the *DataReaders* and *DataWriters* for the publication, subscription, and participant message builtin topics. These leaks have been fixed.

[RTI Issue ID CORE-8367]

5.14.6 Segmentation Fault when Simultaneously Removing Remote Participant and Sending a Message to a Third Participant

A rare segmentation fault may have occurred in the following scenario:

- Three DomainParticipants: P1, P2, P3.
- Delete P1.
- P2 receives an announcement (via ParticipantBuiltinTopicData) from P1 that it has been deleted. P2 starts removing P1 as a remote participant.
- Simultaneously, P2 sends a message to P3. For example, it could be an ACKNACK from P2's DataReader to P3's DataWriter.
- While P2 is sending the message, P2 may have crashed with a bad pointer access.

This problem has been resolved.

[RTI Issue ID CORE-8464]

5.14.7 Could not add property names that were prefixes of existing property names

If you wanted to introduce certain properties in the PropertyQosPolicy that had application-specific meaning, there were certain property names that could not be used because they were the prefix of an existing *Connex DDS* property name.

One example of an unusable property name was "d" because "d" is a prefix of "dds", which is a prefix of many existing property names. If "d" belonged to a *DataWriter*, then the value of "d" would have been misinterpreted as the value of "dds.sample_assignability.accept_unknown_enum_value".

This problem has been resolved. You may now use "d" or any other property name you want.

[RTI Issue ID CORE-8525]

5.14.8 Unbounded memory growth on DataWriter when enable_required_subscriptions set to true and DataReaders setting role name were created/destroyed continuously

In previous releases there may have been an unbounded memory growth on a *DataWriter* when:

- **writer_qos.availability.enable_required_subscriptions** was set to true.
- You created/deleted matching *DataReaders* where **reader_qos.subscription_name.role_name** was set to a value other than NULL.

This problem has been fixed.

[RTI Issue ID CORE-8650]

5.14.9 Possible crash when applications used two extensible types that differed by one aliased primitive member

The algorithm that verifies the compatibility of two extensible types may have crashed when one of the two types added an aliased primitive member at the end.

For example, an application publishing or subscribing T1 and trying to match with an application using T2 would crash:

```
struct T1 {
    long x;
};

typedef long LongAlias;
struct T2 {
    long x;
    LongAlias y;
};
```

This problem affected types with extensible extensibility (the default). It did not affect types with final or mutable extensibility.

This problem has been resolved.

[RTI Issue ID CORE-8516]

5.14.10 Errors reported when compiling a file including disc_rtps_impl.h with gcc -C

On UNIX-like systems, errors were reported when compiling a file that included disc_rtps_impl.h with the -C option. The errors reported by the gcc compiler are now fixed, and the code now compiles without issue.

[RTI Issue ID CORE-8575]

5.14.11 Unlikely segmentation fault when deleting a DataReader and using GROUP ordered access

An unlikely segmentation fault may have occurred when deleting a *DataReader* that was part of a *Subscriber*, when all of the following was true:

- **subscriber_qos.presentation.access_scope** was set to DDS_GROUP_PRESENTATION_QOS or DDS_HIGHEST_OFFERED_PRESENTATION_QOS
- **subscriber_qos.presentation.ordered_access** was set to TRUE
- the *DataReader* matched a *DataWriter* that was part of a Publisher, where publisher_qos.presentation.access_scope was set to DDS_GROUP_PRESENTATION_QOS

This problem has been resolved.

[RTI Issue ID CORE-8615]

5.14.12 RTIOsapiProcess_getId() returns incorrect PID in VxWorks Kernel Mode

In previous releases, calling **RTIOsapiProcess_getId** in Vxworks Kernel Mode would return different PIDs when called from different threads (tasks). This problem is resolved. **RTIOsapiProcess_getId()** will now always return the same TASK_ID, regardless of the thread that called it.

[RTI Issue ID CORE-8634]

5.14.13 Potential segmentation fault while unregistering a logger output device

The function **NDDS_Config_Logger_set_output_device()** is not thread-safe. For example, when calling this function to unregister a device, it is possible for a different thread to be logging a message using the device. This situation can lead to a segmentation fault in the function **ADVLOGLoggerDeviceMgr_logMessageLNOOP()**. This problem has been fixed. **NDDS_Config_Logger_set_output_device()** is now thread-safe.

[RTI Issue ID CORE-8671]

5.14.14 Incorrect value for `type_consistency` in `SubscriptionBuiltinTopicData` in .NET API

When using the .NET API, the value for `type_consistency` within `SubscriptionBuiltinTopicData` was incorrect. It did not reflect the underlying native value. This problem has been resolved.

[RTI Issue ID CORE-8678]

5.14.15 `DomainParticipant` creation did not fail when using an unknown flow controller

Using an unknown flow controller to configure one of the built-in *Topics* did not cause the participant creation to fail. For example:

```
<discovery_config>
<publication_writer_publish_mode> <flow_controller_name>dds.flow_controller.token_
bucket.UNKNOWN</flow_controller_name>
  <kind>ASYNCHRONOUS_PUBLISH_MODE_QOS</kind>
</publication_writer_publish_mode>
</discovery_config>
```

With the above configuration, the `DomainParticipantFactory::create_participant` operation logged the following error, but it still succeeded:

```
DDS_PublishModeQosPolicy_to_presentation_qos_policy: flow controller name 'dds.flow_
controller.token_bucket.UNKNOWN' not found
```

This problem has been resolved. Now the participant creation fails.

[RTI Issue ID CORE-8722]

5.14.16 Removed recursion in include files

Some static code analysis tools detected that there was a recursion in the following header files:

1. `rtxml_extension.h` included `rtxml_object.h`
2. `rtxml_object.h` included `rtxml_object_impl.h`
3. `rtxml_object_impl.h` included `rtxml_extension.h`

Although there are preprocessor guards protecting against multiple inclusion, this recursion has been fixed.

[RTI Issue ID CORE-8739]

5.14.17 `max_blocking_time` not honored for `KEEP_LAST` `DataWriters`

A *DataWriter* might have not honored the `ReliabilityQosPolicy`'s `max_blocking_time` when the `HistoryQoS`'s `kind` was set to `KEEP_LAST` and the `max_send_window_size` was less than the

HistoryQoSPolicy's **depth**. This problem has been resolved: now **max_blocking_time** is always honored.

[RTI Issue ID CORE-8753]

5.14.18 Use of -qosProfile argument in DDS Ping did not take topic_filter into consideration

In *RTI DDS Ping* (*rtiddsping*), the **-qosProfile** command-line argument can be used to specify a Profile within an XML QoS file. *rtiddsping* is used for checking reachability between two DDS nodes. It also allows you to specify the topic name for those published samples using the **-topicName** command-line argument.

When using the **-qosProfile** command line argument, if the Profile contained QoSes with `topic_filters` specified, the expectation was that the utility would use the appropriate QoS (*DataReader* or *DataWriter* for the **-subscriber** or **-publisher** command-line argument respectively) based on the **-topicName** argument or the default name 'PingTopic'. But this wasn't the case in previous releases.

This problem has been resolved in this release. Now *rtiddsping* correctly uses a QoS taking the `topic_filter` into consideration.

[RTI Issue ID CORE-8837]

5.14.19 Use of -qosProfile argument in DDS Spy did not take topic_filter into consideration

In *RTI DDS Spy* (*rtiddsspy*), the **-qosProfile** command-line argument can be used to specify a Profile within an XML QoS file. Since *rtiddsspy* snoops for any topic names being published on a given domain ID, if the Profile contains QoSes with `topic_filters` specified, the expectation is that the utility will use the appropriate QoS based on the snooped topic's name. But this wasn't the case, up through release 5.3.1.

This problem has been resolved in this release. Now *rtiddsspy* correctly uses a QoS taking the `topic_filter` into consideration.

[RTI Issue ID CORE-8838]

5.14.20 Rare race condition may have caused a keep-last DataWriter to time out in write()

The **write()** operation in a reliable *DataWriter* configured with keep-last history QoS shouldn't time out.

A race condition may have caused it to time out if the blocking time was very small.

This problem has been resolved.

[RTI Issue ID CORE-8845]

5.14.21 DDS_PublisherQos_copy() in C API made a shallow copy

DDS_PublisherQos_copy() in the *Connex DDS C* API made a shallow copy for all the `DDS_ThreadSettings_t` it contained within the `asynchronous_publisher` members:

- `asynchronous_batch_thread`
- `thread`
- `topic_query_publication_thread`

Because of this shallow copy, if you made a copy of the `DDS_PublisherQos` from an original value and finalized the copy by calling **finalize()**, accessing the original structure caused a crash, since the `DDS_LongSeq` corresponding to `cpu_list` would have been cleaned up due to the nature of shallow copy. The crash occurred because `DDS_ThreadSettings_t` contained a member of type `DDS_LongSeq`, which, being a sequence, has just pointers to the actual values on the heap. Since other language wrappers were calling the C API, this problem was present in all languages.

Now, **DDS_PublisherQos_copy()** makes a deep copy by replicating the heap values and not just the pointers, in all languages.

[RTI Issue ID CORE-8855]

5.14.22 Unexpected sample losses with reason `DDS_LOST_BY_REMOTE_WRITERS_PER_SAMPLE_LIMIT` when using Group-Ordered access

DataReaders of applications using Group-Ordered Access may have seen unexpected samples losses with lost reason `DDS_LOST_BY_REMOTE_WRITERS_PER_SAMPLE_LIMIT`.

This problem has been fixed.

[RTI Issue ID CORE-8871]

5.14.23 Unexpected "!get remote writer queue" warning

In previous releases, when installing a listener on the built-in *Topics*, you may have observed the following warning:

```
PRESsReader_deleteRemoteWriterQueue:!get remote writer queue
```

The verbosity of the message has been changed to `NDDS_CONFIG_LOG_VERBOSITY_STATUS_LOCAL`.

[RTI Issue ID CORE-8903]

5.14.24 DataReader may have incorrectly reported on_sample_lost() with a negative total_count

When receiving a piggyback heartbeat around the time of matching with a remote *DataWriter*, a *DataReader* may have incorrectly reported **on_sample_lost()**. Moreover, the `total_count` of the `DDS_SampleLostStatus` was a negative number. This problem occurred between a *DataReader* with volatile durability and a *DataWriter* with batching enabled. This problem has been resolved.

[RTI Issue ID CORE-8921]

5.14.25 Unexpected warning printed when a Participant ignored itself

In previous releases, when a Participant ignored itself using the API `DDS_DomainParticipant_ignore_participant`, *Connex DDS* printed the following unexpected warning:

```
PRESParticipant_removeRemoteParticipant:!goto key remoteParticipant
```

This issue has been fixed. This warning is no longer printed.

[RTI Issue ID CORE-8922]

5.14.26 Potential crash when setting new_participant_domain_id in Java monitoring libraries

In release 5.3.0, a crash may have occurred when using Java monitoring libraries with a **new_participant_domain_id** different than the default. This issue is now resolved: the Java monitoring libraries no longer crash when setting **new_participant_domain_id**.

[RTI Issue ID CORE-8940]

5.14.27 TypeCode.print_complete_IDL() failed if a type inherited through an alias

`TypeCode.print_complete_IDL()` failed by producing an `IllegalStateException` if the base type was specified to be an alias (a type created through a typedef). This issue has been resolved.

[RTI Issue ID CORE-8986]

5.14.28 Possible error "Too many open files" specifying the discovery peers by a file

Specifying the discovery peers in a file with the well-known name `NDDS_DISCOVERY_PEERS` may have triggered the error "Too many open files". This error occurred because the file descriptor was not properly closed. This problem has been resolved. Now the file descriptor is closed after it is used.

[RTI Issue ID CORE-9122]

5.14.29 Error receiving batches on a DataWriter from a DataReader with a different endianness

A *DataReader* may not have received samples published by a *DataWriter* where batching is enabled. If the endianness of the platform in which the *DataReader* was running was different than the endianness of the platform in which the *DataWriter* was running, the *DataReader* printed deserialization errors. This problem has been resolved (samples are now received).

[RTI Issue ID CORE-9128]

5.14.30 Performance degradation when DataWriters sent HeartbeatFrag RTPS messages

You may have observed performance degradation in large data scenarios in which a *DataWriter* sends DataFrag RTPS messages.

The problem occurred only when the *DataWriter* publishing the DataFrag messages also sent HeartbeatFrag messages. This problem happened only with third-party vendor *DataWriters* or when using a *Connex DDS Micro* 3.0.0 (or higher) *DataWriter* with *Connex DDS Professional* 5.3.0 or lower. (The problem does not occur when using *Micro* 3.0.0 with *Connex DDS* 6.0.0.)

This problem has been fixed.

[RTI Issue ID CORE-9211]

5.14.31 Unexpected COMMENDSrReaderService_onSubmessage:!add NACK_FRAG error message

You may have seen this error message on DataReaders receiving RTPS data fragments (DataFrag):

```
COMMENDSrReaderService_onSubmessage:!add NACK_FRAG
```

This error occurred only when the *DataWriter* publishing the DataFrag messages also sent HeartbeatFrag messages. The error occurred only with third-party vendor *DataWriters* or when using a *Connex DDS Micro* 3.0.0 (or higher) *DataWriter* with *Connex DDS Professional* 5.3.0 or lower. (The problem does not occur when using *Micro* 3.0.0 with *Connex DDS* 6.0.0.)

This problem has been fixed.

[RTI Issue ID CORE-9213]

5.14.32 Visual Studio run-times copied to .NET project directory

Every Windows machine needs to use the redistributable libraries installed on that particular system. Previously, some incompatible redistributable libraries were mistakenly copied into the final application binary directory. Those incompatible libraries in the application's binary directory were used instead of the system

redistributable libraries. In this release, those libraries are no longer copied into the application's binary directory. This allows the correct system ones to be used instead.

[RTI Issue ID PLATFORMS-1316]

Chapter 6 Known Issues

6.1 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 DDS* 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 Section 6.3.12, Application Acknowledgment, in the *RTI Connex DDS Core Libraries User's Manual*.

[RTI Issue ID CORE-5329]

6.2 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\rtd_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

This issue does not affect the functionality of the example.

[RTI Issue ID CODEGENII-743]

6.3 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 DDS Core Libraries User's Manual*.

[RTI Issue ID CORE-5360]

6.4 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]

6.5 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 DDS Core Libraries User's Manual*).

[RTI Issue ID CORE-5181]

6.6 Discovery with Connext DDS Micro Fails when Shared Memory Transport Enabled

Given a *Connext DDS* 6.0.1 application with the shared memory transport enabled, a *Connext DDS* Micro 2.4.x application will fail to discover it. This is due to a bug in *Connext DDS* Micro that prevents a received participant discovery message from being correctly processed. This bug will be fixed in a future release of *Connext DDS* Micro. As a workaround, you can disable the shared memory transport in the *Connext DDS* application and use UDPv4 instead.

[RTI Issue ID EDDY-1615]

6.7 Examples and Generated Code for Visual Studio 2017 and later may not Compile (Error MSB8036)

The examples provided with *Connext DDS* 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 Connext DDS Core Libraries Platform Notes*.

[RTI Issue ID CODEGENII-800]

6.8 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 *Connex DDS* 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 $30,720/\mathbf{max_samples}$ 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]

6.9 Memory Leak if Foo::initialize() Called Twice

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

[RTI Issue ID CORE-7678]

6.10 Shared Memory Communication Requires Setting **dds.transport.shmem.builtin.hostid** in Transport Mobility Scenarios

On some platforms, to use the shared memory transport in a transport mobility scenario, you will also need to set the **dds.transport.shmem.builtin.hostid** property in the *DomainParticipant's* Property QoS policy. Use this property to assign a unique *hostid* to the transport. In this release, that unique *hostid* (a 32-bit integer) must be generated by the user and must be the same for all applications running on the same host.

For instance, if you want to two *Connex DDS* applications to communicate using shared memory and one application is started while no NICs are enabled, and after that you enable a NIC and start another *Connex DDS* application, your applications will not communicate by default on certain platforms. To make both applications communicate, you need to set the property **dds.transport.shmem.builtin.hostid** to the same value in both applications.

Affected platforms: AIX, Solaris.

[RTI Issue ID CORE-8040]

6.11 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]

6.12 Uninstalling on AIX Systems

To uninstall *Connex DDS* on an AIX system: if the original installation is on an NFS drive, the uninstaller will hang and fail to completely uninstall the product. As a workaround, you can remove the installation with this command:

```
rm -rf $INSTALL_PATH/rti_connex_dds-6.0.1
```

[RTI Issue ID INSTALL-323]

6.13 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]

6.14 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]

6.15 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 DDS*. 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 liveness 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]

6.16 Communication may not be Reestablished in Some IP Mobility Scenarios

If you have two *Connex DDS* 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]

6.17 DomainParticipantFactoryQos in XML may not be Loaded

The *DomainParticipantFactoryQos* set in XML may not be loaded in the Modern C++ API. This happens when a *QosProvider* is accessed before any DDS entities have been created. The most straightforward workaround for this issue is to set the *DomainParticipantFactoryQos* in code as follows:

```
dds::core::QosProvider qos_provider = dds::core::QosProvider("USER_QOS_PROFILES.xml",
" MyQosLibrary::MyQosProfile");

dds::domain::qos::DomainParticipantFactoryQos factory_qos;
factory_qos->entity_factory.autoenable_created_entities(false);

// Set the DomainParticipantFactoryQos
dds::domain::DomainParticipant::participant_factory_qos(factory_qos);

// Continue with application code
```

```
dds::domain::DomainParticipant participant(domain_id, qos_provider.participant_qos());
```

[RTI Issue ID CORE-6846]

6.18 Multichannel and TopicQuery do not Work with Large Data in Some Cases

In releases 5.3.0.16 and 5.3.1, RTI Issue ID CORE-8422 ("MultiChannel and TopicQuery did not Work with Large Data") was documented as fixed:

A MultiChannel *DataWriter* or a *DataWriter* configured to dispatch TopicQueries (by setting **writer_qos.topic_query_dispatch.enable** to true) was not able to send large data (that is, samples bigger than the transport **message_size_max**). The **write()** call would fail with **RETCODE_ERROR**. This problem has been resolved.

This problem, however, was not fully resolved, due to another bug, described in RTI Issue ID CORE-9335. With a MultiChannel *DataWriter*, some of the samples published on the channels may not be received by a *DataReader*. In addition, with both MultiChannel *DataWriters* and *DataWriters* configured to dispatch TopicQueries, the original identity of the samples published by these *DataWriters* may be wrong. The original identity is accessible using the fields **original_publication_virtual_sequence_number** and **original_publication_virtual_guid** in *SampleInfo* when a sample is received.

Large data support for Multichannel and TopicQuery has been disabled in *Connex DDS* 6.0.1. If you need this functionality enabled as it was in 6.0.0, please contact RTI support (support@rti.com).

[RTI Issue ID CORE-9335]

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

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 long 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]

6.20 Potential application crash when receiving a sample on the service request channel

The reception of a sample over the service request channel (for instance, a TopicQuery request) may cause an application to crash. The crash may happen using non-C libraries if the related *DomainParticipant* has installed a listener that sets the **onDataOnReaders** status on its mask.

[RTI Issue ID CORE-9891]

6.21 Known Issues with Dynamic Data

6.21.1 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 long long 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 long longs from a 32-bit or smaller integer type is supported on all Windows, Solaris, and Linux architectures, and any additional 64-bit architectures. Conversion to 128-bit long doubles from a float or double is only supported on Solaris SPARC architectures.

[RTI Issue ID CORE-2986]

6.21.2 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]

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

DynamicData 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 long myLongAlias;

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

```
};
```

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

[RTI Issue ID CORE-9176]

6.22 Known Issues in RTI Monitoring Library

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

If a *Connex DDS* 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 work-around, 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]

6.22.2 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]

6.22.3 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:

```
<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>
</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]

6.22.4 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]

Chapter 7 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/>).