

# RTI Connex Micro

User's Manual

Version 4.2.0



# Contents

<b>1</b>	<b>Getting Started</b>	<b>1</b>
1.1	Before You Get Started . . . . .	1
1.1.1	What is RTI Connex Micro? . . . . .	1
1.1.2	Downloading Connex Micro . . . . .	2
1.1.3	Installing Connex Micro . . . . .	2
	Installing with a GUI . . . . .	3
	Installing from a command line . . . . .	3
	Paths mentioned in documentation . . . . .	4
1.1.4	Checking what is installed . . . . .	6
1.1.5	Where do I get more help? . . . . .	6
1.2	Hello World with Publish/Subscribe . . . . .	6
1.2.1	Introduction to DataWriters, DataReaders, and Topics . . . . .	7
1.2.2	Create a DataWriter and DataReader . . . . .	8
	Set environment variables . . . . .	8
	Generate an example application . . . . .	8
	Modify application . . . . .	11
	Modify publisher . . . . .	11
	Modify subscriber . . . . .	13
	Compile your changes . . . . .	14
	Run the applications . . . . .	14
	Taking it further . . . . .	16
1.2.3	Troubleshooting . . . . .	18
	Why aren't my applications communicating? . . . . .	18
	Why does the DataReader miss the first samples? . . . . .	18
1.2.4	Next steps . . . . .	18
<b>2</b>	<b>Developing Applications</b>	<b>19</b>
2.1	Prepare Your Development Environment . . . . .	19
2.1.1	Set environment variables . . . . .	19
2.1.2	Add required preprocessor flags . . . . .	19
2.1.3	Link applications and libraries . . . . .	20
2.2	Define a Data Type . . . . .	21
2.3	Generate Type Support Code with rtiddsgen . . . . .	22
2.4	Create an Application . . . . .	23
2.4.1	Registry Configuration . . . . .	23
2.5	Configure UDP Transport . . . . .	25
2.6	Create DomainParticipant, Topic, and Type . . . . .	26

2.6.1	Register Type . . . . .	27
2.6.2	Create Topic of Registered Type . . . . .	28
2.6.3	DPSE Discovery: Assert Remote Participant . . . . .	28
2.7	Create Publisher . . . . .	29
2.8	Create DataWriter . . . . .	29
2.8.1	DPSE Discovery: Assert Remote Subscription . . . . .	30
2.8.2	Writing Samples . . . . .	31
2.9	Create Subscriber . . . . .	31
2.10	Create DataReader . . . . .	32
2.10.1	DPSE Discovery: Assert Remote Publication . . . . .	34
2.10.2	Receiving Samples . . . . .	35
2.10.3	Filtering Samples . . . . .	37
<b>3</b>	<b>User's Manual . . . . .</b>	<b>38</b>
3.1	Data Types . . . . .	38
3.1.1	Introduction to the Type System . . . . .	39
	Sequences . . . . .	40
	Strings and Wide Strings . . . . .	41
	Extensible Types (XTypes) 1.2 Compatibility . . . . .	42
3.1.2	Creating User Data Types with IDL . . . . .	43
3.1.3	Working with DDS Data Samples . . . . .	43
3.2	DDS Entities . . . . .	44
3.3	Sending Data . . . . .	45
3.3.1	Preview: Steps to Sending Data . . . . .	45
3.3.2	Publishers . . . . .	46
3.3.3	DataWriters . . . . .	46
3.3.4	Publisher/Subscriber QosPolicies . . . . .	47
3.3.5	DataWriter QosPolicies . . . . .	47
3.4	Receiving Data . . . . .	47
3.4.1	Preview: Steps to Receiving Data . . . . .	47
3.4.2	Subscribers . . . . .	49
3.4.3	DataReaders . . . . .	50
3.4.4	Using DataReaders to Access Data (Read & Take) . . . . .	50
3.4.5	Subscriber QosPolicies . . . . .	50
3.4.6	DataReader QosPolicies . . . . .	50
3.5	DDS Domains . . . . .	50
3.5.1	Fundamentals of DDS Domains and DomainParticipants . . . . .	50
3.5.2	Discovery Announcements . . . . .	52
3.6	Transports . . . . .	53
3.6.1	Introduction . . . . .	53
3.6.2	Transport Registration . . . . .	54
3.6.3	Transport Addresses . . . . .	55
3.6.4	Transport Port Number . . . . .	55
3.6.5	RTPS . . . . .	56
	Registration of RTPS . . . . .	56
	Overriding the Builtin RTPS Checksum Functions . . . . .	57
	Example . . . . .	58
3.6.6	INTRA Transport . . . . .	60

	Registering the INTRA Transport . . . . .	61
	Reliability and Durability . . . . .	61
	Threading Model . . . . .	61
3.6.7	Shared Memory Transport (SHMEM) . . . . .	62
	Registering the SHMEM Transport . . . . .	62
	Threading Model . . . . .	64
	SHMEM Configuration . . . . .	65
	Caveats . . . . .	66
3.6.8	Zero Copy v2 Transport . . . . .	67
	Generate example and type support files . . . . .	67
	Initialize the Zero Copy v2 transport . . . . .	68
	Register the Zero Copy v2 transport . . . . .	68
	Enable transports . . . . .	69
	Sample management . . . . .	70
	Configuration . . . . .	71
	Using multiple Zero Copy v2 transport instances . . . . .	72
3.6.9	UDP Transport . . . . .	72
	Registering the UDP Transport . . . . .	73
	Threading Model . . . . .	74
	UDP Configuration . . . . .	75
	UDP Transformations . . . . .	80
3.6.10	NETIO Datagram Transport . . . . .	109
	Registering a Datagram Interface . . . . .	109
	Addressing a Datagram Transport . . . . .	111
	Datagram UDP Setup . . . . .	111
	User Interface . . . . .	113
3.7	Discovery . . . . .	114
3.7.1	What is Discovery? . . . . .	114
	Simple Participant Discovery . . . . .	115
	Simple Endpoint Discovery . . . . .	115
3.7.2	Configuring Participant Discovery Peers . . . . .	116
	peer_desc_string . . . . .	116
3.7.3	Configuring Initial Peers and Adding Peers . . . . .	117
3.7.4	Discovery Plugins . . . . .	117
	Dynamic Discovery Plugin . . . . .	118
	Static Discovery Plugin . . . . .	118
3.7.5	DomainParticipant Discovery by Name . . . . .	121
3.7.6	Queueing Discovery Messages . . . . .	121
3.8	User Discovery Data . . . . .	122
3.8.1	Introduction . . . . .	122
3.8.2	Resource Limits . . . . .	122
3.8.3	Propagating User Discovery Data . . . . .	123
3.8.4	Accessing User Discovery Data . . . . .	124
3.8.5	QoS Policies . . . . .	126
	USER_DATA . . . . .	126
	TOPIC_DATA . . . . .	126
	GROUP_DATA . . . . .	127
3.9	Partitions . . . . .	127

3.9.1	Introduction . . . . .	127
3.9.2	Rules for PARTITION matching . . . . .	129
3.9.3	Pattern matching for PARTITION names . . . . .	129
	Regular Expression Matching . . . . .	130
3.9.4	Example . . . . .	131
	Location-based partitions . . . . .	131
	Access-control group partitions . . . . .	132
3.9.5	Properties . . . . .	133
3.9.6	Resource limits . . . . .	133
	Configuring for runtime allocation . . . . .	133
	Configuring for preallocated memory . . . . .	134
3.10	Content Filtering . . . . .	134
3.10.1	Creating a Content Filter . . . . .	135
3.10.2	Where Filtering is Applied—Publishing vs. Subscribing Side . . . . .	138
3.10.3	Interoperability . . . . .	139
	With Connex Professional . . . . .	139
	With Connex Cert . . . . .	140
3.10.4	SQL Filter Expression Notation . . . . .	140
	Example SQL filter expressions . . . . .	140
	SQL grammar . . . . .	142
	Token expressions . . . . .	143
	Type compatibility in the Predicate . . . . .	144
	Enumerations . . . . .	144
3.11	Generating Type Support . . . . .	145
3.11.1	IDL type definition . . . . .	145
3.11.2	Generating type support . . . . .	145
	C . . . . .	145
	C++ . . . . .	146
	Notes on command-line options . . . . .	146
	Generated type support files . . . . .	146
3.11.3	Using custom data-types in Connex Micro applications . . . . .	146
3.11.4	Customizing generated code . . . . .	147
3.11.5	Unsupported Features of rtiddsgen with Connex Micro . . . . .	148
3.12	Threading Model . . . . .	148
3.12.1	Introduction . . . . .	148
3.12.2	Architectural Overview . . . . .	148
3.12.3	Threading Model . . . . .	149
	OSAPI Threads . . . . .	149
	UDP Transport Threads . . . . .	150
	General Thread Configuration . . . . .	151
3.12.4	Critical Sections . . . . .	151
	Calling DDS APIs from listeners . . . . .	151
3.13	Batching . . . . .	152
3.13.1	Overview . . . . .	152
3.13.2	Interoperability . . . . .	152
3.13.3	Performance . . . . .	152
3.13.4	Example Configuration . . . . .	153
3.14	Message Integrity Checking . . . . .	154

3.14.1	RTPS checksum . . . . .	154
3.14.2	Configurations . . . . .	155
	Selecting a checksum algorithm . . . . .	155
	Configuring the DDS DomainParticipant . . . . .	155
3.14.3	Participant discovery and compatibility . . . . .	156
3.14.4	Interoperability with Connex Professional . . . . .	157
3.15	Sending Large Data . . . . .	157
3.15.1	Overview . . . . .	158
3.15.2	Configuration of Large Data . . . . .	159
3.15.3	Limitations . . . . .	159
3.16	Zero Copy Transfer . . . . .	160
3.16.1	Compatibility . . . . .	161
3.16.2	Overview . . . . .	162
3.16.3	Getting started . . . . .	162
	Writing samples . . . . .	163
	Reading samples . . . . .	164
3.16.4	Synchronizing samples . . . . .	165
	Zero Copy v1 synchronization . . . . .	165
	Zero Copy v2 synchronization . . . . .	166
3.16.5	Caveats . . . . .	167
3.17	FlatData Language Binding . . . . .	167
3.17.1	Overview . . . . .	167
3.17.2	Getting Started . . . . .	167
3.17.3	Further Information . . . . .	168
3.18	Application Generation Using XML . . . . .	168
3.18.1	Defining an Application in XML . . . . .	168
	Important Points . . . . .	169
3.18.2	Generating the Application from XML . . . . .	169
	Micro Application Generator (MAG) Tool . . . . .	169
	Generating the Application with MAG . . . . .	171
	MAG Command-Line Options . . . . .	172
	Integrating Generated Files into Your Application's Build . . . . .	175
3.18.3	Creating the Application . . . . .	175
	Generate Type-Support Code from the Type Definition . . . . .	175
	Generate DDS Entities from the System Definition . . . . .	177
	Examining the application . . . . .	177
	Call API to Create DomainParticipant . . . . .	178
	Retrieve Entities by Name . . . . .	178
3.18.4	Examine example XML configuration files . . . . .	178
	Type Definition . . . . .	179
	Domain Definition . . . . .	180
	DomainParticipant Definition . . . . .	181
	QoS Definition . . . . .	186
	Transport and Discovery Configuration . . . . .	199
	Flow Controllers . . . . .	204
	Static Discovery . . . . .	208
	Lightweight Security Plugin . . . . .	215
	Content filtering . . . . .	217

3.18.5	Errors Caused by Invalid Configurations and QoS . . . . .	223
3.19	Building Against FACE Conformance Libraries . . . . .	225
3.19.1	Requirements . . . . .	226
	Connex Micro Source Code . . . . .	226
	FACE Conformance Tools . . . . .	226
	CMake . . . . .	226
3.19.2	FACE Golden Libraries . . . . .	226
	Building the FACE Golden Libraries . . . . .	226
3.19.3	Building the Connex Micro Source . . . . .	226
3.20	Working With Sequences . . . . .	228
3.20.1	Introduction . . . . .	228
3.20.2	Working with Sequences . . . . .	228
	Overview . . . . .	228
	Working with IDL Sequences . . . . .	230
	Working with Application Defined Sequences . . . . .	231
3.21	Debugging . . . . .	232
3.21.1	Overview . . . . .	232
3.21.2	Configuring Logging . . . . .	233
3.21.3	Log Message Kinds . . . . .	234
3.21.4	Interpreting Log Messages and Error Codes . . . . .	234
3.22	Example Applications . . . . .	235
3.22.1	Generating examples . . . . .	235
	Default example . . . . .	235
	Custom example . . . . .	236
	Descriptions of generated examples . . . . .	237
	How to compile the generated examples . . . . .	238
	How to run the generated examples . . . . .	239
3.22.2	Provided examples . . . . .	240
3.23	Lightweight Security Plugin . . . . .	240
3.23.1	Overview . . . . .	240
3.23.2	Downloading and Installing the Lightweight Security Plugin . . . . .	241
	Package contents . . . . .	241
3.23.3	Getting Started . . . . .	241
	Link application and libraries . . . . .	242
	Enable lightweight security in your application . . . . .	242
	Specify the plugin in the DomainParticipant QoS . . . . .	243
	Configure the plugin . . . . .	243
	Create a protected DomainParticipant . . . . .	244
3.23.4	Configuring the Lightweight Security Plugin . . . . .	244
3.23.5	Technical Considerations . . . . .	246
	RTPS transport only . . . . .	246
	Maximum Transmission Unit (MTU) . . . . .	246
3.23.6	Interoperability . . . . .	246
	Additional Authenticated Data (AAD) . . . . .	246
	Passphrase ID limitations . . . . .	247
3.24	Memory Management . . . . .	247
3.24.1	Resource limits . . . . .	247
3.24.2	Dynamic memory allocation . . . . .	247

3.24.3	DDS resource limits . . . . .	248
3.24.4	Discovery plugin resource limits . . . . .	248
3.24.5	Type support resource limits . . . . .	249
<b>4</b>	<b>Platform Notes</b>	<b>250</b>
4.1	Introduction . . . . .	250
4.1.1	Library types . . . . .	250
	Integrated libraries . . . . .	251
	Split libraries . . . . .	251
4.1.2	Library descriptions . . . . .	252
4.1.3	Build profiles . . . . .	253
4.1.4	Supported libraries by platform . . . . .	254
4.1.5	Supported transports by platform . . . . .	257
4.2	FreeRTOS Platforms . . . . .	259
4.2.1	Port overview . . . . .	259
4.2.2	How to configure lwIP and FreeRTOS . . . . .	260
4.2.3	How the PIL was built for FreeRTOS . . . . .	271
4.2.4	Building the PSL from source for FreeRTOS platforms . . . . .	273
4.2.5	Building FreeRTOS applications with Connex Micro . . . . .	274
4.2.6	System tick rollovers . . . . .	276
4.3	Linux Platforms . . . . .	277
4.3.1	How the PIL was built for Linux platforms . . . . .	277
4.3.2	Building the PSL from source for Linux platforms . . . . .	279
4.4	macOS Platforms . . . . .	282
4.4.1	How the PIL was built for macOS platforms . . . . .	282
4.4.2	Building the PSL from source for macOS platforms . . . . .	284
4.5	QNX Platforms . . . . .	288
4.5.1	How the PIL was built for QNX platforms . . . . .	288
4.5.2	Building the PSL from source for QNX platforms . . . . .	290
4.6	Windows Platforms . . . . .	292
4.6.1	How the PIL was built for Windows platforms . . . . .	292
4.6.2	Building the PSL from source for Windows platforms . . . . .	293
<b>5</b>	<b>Building Connex Micro</b>	<b>295</b>
5.1	Connex Micro Platforms . . . . .	295
5.2	Building Connex Micro for Common Platforms . . . . .	295
5.2.1	Setting up the build environment . . . . .	296
	The host environment . . . . .	296
	The target environment . . . . .	297
5.2.2	Building the PSL . . . . .	297
	Building the PSL with rtime-make . . . . .	297
	Building the PSL with CMake . . . . .	299
5.2.3	Building the source . . . . .	300
	Building with rtime-make . . . . .	301
	Building with CMake . . . . .	302
5.2.4	Cross-compiling Connex Micro . . . . .	305
5.3	Building Connex Micro with Compatibility for Connex Cert . . . . .	305
5.3.1	Compiling with compatibility for Connex Cert . . . . .	306



5.3.2	Compiling applications with compatibility for Connex Cert . . . . .	307
<b>6</b>	<b>Working with Connex Micro and Connex Professional</b>	<b>308</b>
6.1	Supported DDS Features . . . . .	308
6.1.1	DDS Entity Support . . . . .	308
6.1.2	DDS QoS Policy Support . . . . .	309
6.2	Development Environment . . . . .	310
6.3	Non-standard APIs . . . . .	310
6.4	QoS Policies . . . . .	310
6.5	Standard APIs . . . . .	310
6.6	IDL Files . . . . .	310
6.7	Interoperability . . . . .	311
6.7.1	Discovery . . . . .	311
6.7.2	Transports . . . . .	311
6.8	Connex Tools . . . . .	311
6.8.1	Admin Console . . . . .	311
6.8.2	Distributed Logger . . . . .	312
6.8.3	LabVIEW . . . . .	312
6.8.4	Monitor . . . . .	313
6.8.5	Recording Service . . . . .	313
	RTI Recorder . . . . .	313
	RTI Replay . . . . .	313
	RTI Converter . . . . .	313
6.8.6	Wireshark . . . . .	313
6.8.7	Persistence Service . . . . .	314
6.8.8	Application Generation Using XML . . . . .	314
<b>7</b>	<b>API Reference</b>	<b>315</b>
<b>8</b>	<b>Release Notes</b>	<b>316</b>
8.1	Compatibility . . . . .	316
8.2	Supported Platforms and Programming Languages . . . . .	316
8.3	What's New in 4.2.0 . . . . .	318
8.3.1	Reduced core library size . . . . .	318
8.3.2	Secure RTPS communication with the Lightweight Security Plugin . . . . .	318
8.3.3	Optimize network bandwidth usage with content filtering . . . . .	319
8.3.4	Improve system robustness with DomainParticipant rediscovery . . . . .	319
8.3.5	Develop secure applications in XML with MAG . . . . .	319
8.3.6	Develop applications with content filtering in XML with MAG . . . . .	319
8.3.7	Enable DomainParticipant rediscovery in XML with MAG . . . . .	320
8.3.8	Enable endpoint discovery message queueing in XML with MAG . . . . .	320
8.3.9	Third-party and open source software changes . . . . .	320
8.4	What's Fixed in 4.2.0 . . . . .	320
8.4.1	Discovery . . . . .	321
	[Critical] Lost INFO_DST messages with GUIDPREFIX_UNKNOWN . . . . .	321
	[Major] DomainParticipants stopped sending announcements to initial peer locator after a discovered DomainParticipant went away . . . . .	321
	[Major] Possible mismatched endpoint messages during discovery . . . . .	321

	[Major] DomainParticipants continued to send DATA(p) messages to nonexistent remote participants . . . . .	321
	[Major] DataReader failed to acknowledge samples when it received more instances than it could store . . . . .	322
	[Minor] Participants using DPDE sent one extra announcement message . . .	322
	[Minor] Failed to delete a participant before sending at least one participant announcement . . . . .	322
8.4.2	Reliability Protocol and Wire Representation . . . . .	322
	[Major] Asynchronous flow controller may have stopped sending data . . . .	322
	[Major] Reliable DataReader may have stopped receiving data . . . . .	322
8.4.3	APIs (C or Traditional C++) . . . . .	323
	[Major] DDS_DomainParticipantFactory_finalize_instance failed if INTRA transport had been unregistered . . . . .	323
	[Minor] Maximum blocking time is limited to 25 days . . . . .	323
8.4.4	Generated Code (C, Traditional C++, and Modern C++) . . . . .	323
	[Major] Failed to serialize a union with only a default label . . . . .	323
8.4.5	Hangs . . . . .	323
	[Critical] Application may have frozen during discovery of a remote participant	323
8.4.6	Memory Leaks/Growth . . . . .	324
	[Minor] Memory leak in static discovery examples . . . . .	324
	[Major] Resources not freed when calling unregister_instance on a reliable Zero Copy DataWriter . . . . .	324
8.4.7	Platform and Build Changes . . . . .	324
	[Minor] Linker errors occurred when compiling HelloWorld examples for Windows systems . . . . .	324
8.4.8	Shipped Examples . . . . .	324
	[Trivial] Failure to compile example generated for MAG . . . . .	324
8.4.9	Interoperability . . . . .	325
	[Major] Zero Copy communication failed when applications were started from different directories . . . . .	325
8.4.10	Other . . . . .	325
	[Critical] Potential loss of queued samples when using asynchronous communication . . . . .	325
	[Critical] Missed DATA_FRAG messages with multiple fragments . . . . .	325
	[Major] Micro Application Generator (MAG) failed to start on macOS platforms	325
	[Major] DataWriter with finite timestamp failed to send data if date exceeded January 2038 . . . . .	326
	[Major] Connex Micro failed to parse malformed DATA_FRAG submessages	326
	[Major] Connex Micro would not work if year exceeded 2038 . . . . .	326
	[Minor] Multiple endpoints did not match properly when using Zero Copy transfer . . . . .	326
	[Minor] SYSTEM_RESOURCE_LIMITS.max_components could not be changed . . . . .	326
	[Trivial] rtime-make did not support building multiple targets on Windows systems and failed to report some CMake errors . . . . .	326
8.5	Previous Releases . . . . .	327
8.5.1	What's New in 4.1.0 . . . . .	327

	Platform-independent code is now separate from OS and network stack inte-	
	gration . . . . .	327
	Transfer large data samples quickly with Zero Copy v2 . . . . .	328
	Enable and configure Zero Copy transfer with MAG . . . . .	328
	Enhance data reliability by detecting and discarding corrupted RTPS messages	328
	Develop more reliable applications with MAG . . . . .	329
	Guarantee compatibility with Connex Professional with MAG when using	
	the Shared Memory Transport . . . . .	329
	Improve control of data distribution to multicast addresses with new UDP	
	transport options . . . . .	329
	Develop applications with new UDP transport options with MAG . . . . .	330
	Build Connex Micro libraries conveniently with symlinks . . . . .	330
8.5.2	What's Fixed in 4.1.0 . . . . .	330
	Discovery . . . . .	330
	Usability . . . . .	331
	Transports . . . . .	331
	Reliability Protocol and Wire Representation . . . . .	332
	APIs (C or Traditional C++) . . . . .	332
	XML Configuration . . . . .	333
	Crashes . . . . .	333
	Hangs . . . . .	334
	Memory Leaks/Growth . . . . .	335
	Data Corruption . . . . .	335
	Interoperability . . . . .	335
	Other . . . . .	336
8.5.3	What's New in 4.0.1 . . . . .	338
	Enable or disable padding bits with PROPERTY QoS policy in DomainPar-	
	ticipant and Data Writer . . . . .	338
	Generate examples with new template options for Code Generator . . . . .	338
8.5.4	What's Fixed in 4.0.1 . . . . .	339
	Discovery . . . . .	339
	Usability . . . . .	339
	APIs (C or Traditional C++) . . . . .	340
	Generated Code (C, Traditional C++, and Modern C++) . . . . .	341
	Crashes . . . . .	341
	Data Corruption . . . . .	342
	Interoperability . . . . .	344
8.5.5	What's New in 4.0.0 . . . . .	344
	Enhanced performance for asynchronous DataWriters . . . . .	344
	Further control which entities communicate with each other using new Par-	
	tition QoS policy . . . . .	345
	Store additional entity-related information that is passed between applica-	
	tions during discovery using new User/Topic/Group Data QoS policies	345
	Verify that locally created participant GUIDs are unique within a Domain-	
	ParticipantFactory . . . . .	345
	Micro Application Generator (MAG) . . . . .	345
8.5.6	What's Fixed in 4.0.0 . . . . .	348
	Discovery . . . . .	348

	Serialization and Deserialization . . . . .	349
	Usability . . . . .	352
	Transports . . . . .	352
	Reliability Protocol and Wire Representation . . . . .	352
	Logging . . . . .	353
	Performance and Scalability . . . . .	353
	APIs (C or Traditional C++) . . . . .	354
	Generated Code (C, Traditional C++, and Modern C++) . . . . .	355
	Micro Application Generator . . . . .	356
	OMG Specification Compliance . . . . .	358
	Interoperability . . . . .	358
	Vulnerabilities . . . . .	359
	Other . . . . .	360
8.6	Known Issues . . . . .	360
8.6.1	Samples cannot be recovered if subscribing application fails to return loan .	360
8.6.2	Connex Micro does not work with wide-string characters in the network interface name . . . . .	361
8.6.3	64-bit discriminator values greater than $(2^{31}-1)$ or smaller than $(-2^{31})$ not supported . . . . .	361
8.6.4	NaN and INF float and doubles are not detected and will not cause errors .	361
8.6.5	Ungracefully terminated QNX processes using SHMEM transport prevents startup of new processes due to unclosed POSIX semaphores . . . . .	361
8.6.6	Flow Controllers require RTOS . . . . .	362
8.6.7	LatencyBudget is not part of the DataReaderQos or DataWriterQos policy .	362
8.6.8	Porting Guide not included . . . . .	362
8.6.9	Platform Independent Library toolchain dependencies . . . . .	362
8.7	Experimental Features . . . . .	362
<b>9</b>	<b>Benchmarks</b>	<b>364</b>
<b>10</b>	<b>Copyrights</b>	<b>365</b>
<b>11</b>	<b>Third-Party and Open Source Software</b>	<b>367</b>
11.1	Connex Micro Libraries . . . . .	367
11.1.1	fnmatch . . . . .	367
11.1.2	crc32c.c . . . . .	368
11.1.3	MD5 . . . . .	369
11.2	RTI Code Generator (rtiddsgen) . . . . .	369
11.2.1	ANTLR . . . . .	369
11.2.2	Apache Commons Lang . . . . .	370
11.2.3	Apache Log4j 2 . . . . .	370
11.2.4	Apache Velocity . . . . .	371
11.2.5	Simple Logging Facade for Java (SLF4J) . . . . .	371
11.2.6	Gson . . . . .	371
11.3	Micro Application Generator (rtiddsmag) . . . . .	372
11.3.1	Apache Commons CLI . . . . .	372
11.3.2	Apache Commons Lang . . . . .	372
11.3.3	Apache Log4j 2 . . . . .	372

11.3.4	Apache Velocity . . . . .	372
11.3.5	Extended StAX API . . . . .	372
11.3.6	Fast Infoset . . . . .	373
11.3.7	Istack Common Utility Code Runtime . . . . .	374
11.3.8	JavaBeans Activation Framework API . . . . .	374
11.3.9	Javax Annotation API . . . . .	374
11.3.10	JAXB API . . . . .	374
11.3.11	JAXB Runtime . . . . .	374
11.3.12	Simple Logging Facade for Java (SLF4J) . . . . .	374
11.3.13	TXW2 . . . . .	375
11.4	Lightweight Security Plugin . . . . .	375
11.4.1	OpenSSL . . . . .	375
11.5	Appendix . . . . .	376
11.5.1	Apache License version 2.0, January 2004 ( <a href="http://www.apache.org/licenses/">http://www.apache.org/licenses/</a> )	376
11.5.2	Common Development and Distribution License (CDDL) Version 1.1 . . . .	380
<b>12</b>	<b>Appendix</b>	<b>387</b>
12.1	Package Contents . . . . .	387
12.1.1	Host bundle . . . . .	387
12.1.2	Target bundle . . . . .	388
12.1.3	Lightweight Security Plugin host bundle . . . . .	389
12.1.4	Lightweight Security Plugin target bundle . . . . .	389
12.2	Directory Structure . . . . .	390
<b>13</b>	<b>Contact Support</b>	<b>391</b>
<b>14</b>	<b>Join the Community</b>	<b>392</b>

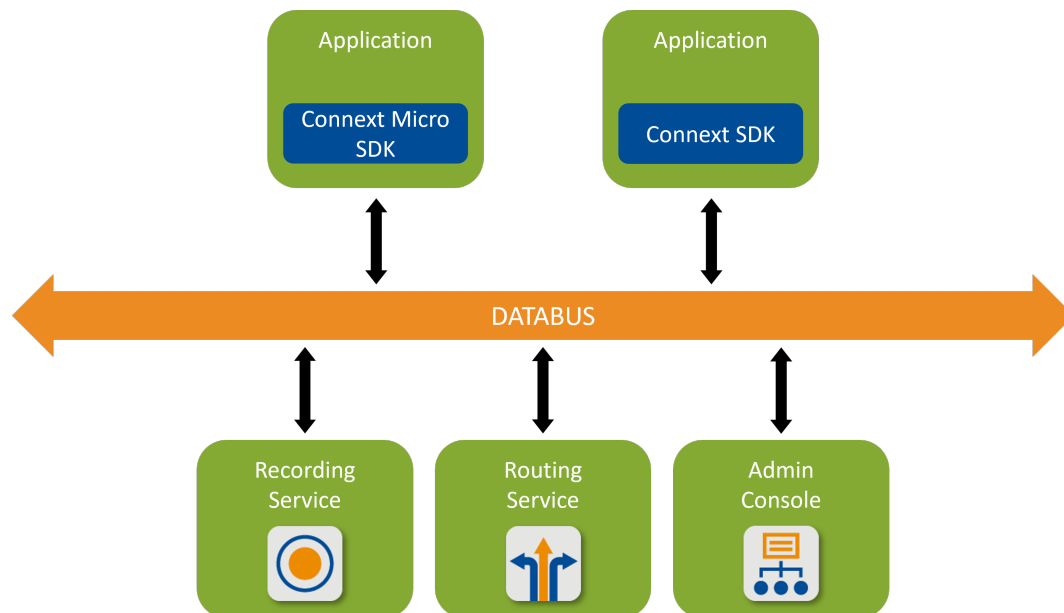
# Chapter 1

## Getting Started

Welcome to *RTI® Connex® Micro*! This section is for first-time users, and will walk you through the basics of *Connex Micro*. We will cover what *Connex Micro* is, how to download and install it, and how to run your first “Hello world!” application.

### 1.1 Before You Get Started

#### 1.1.1 What is RTI Connex Micro?



*RTI® Connex® Micro®* is part of the *RTI® Connex®* connectivity framework for building distributed applications using a shared databus. *Connex Micro* provides C and C++ libraries and tools for resource-constrained embedded systems (that is, systems with limited memory and CPU power) to connect to the *Connex* databus, allowing them to exchange data without compromising performance or reliability. It interoperates seamlessly with other *Connex* applications, tools, and infrastructure services (with some exceptions as noted in this documentation).

RTI also offers *Connex Cert*, which provides C libraries and tools for safety-critical systems. *Connex Micro* includes C libraries that are comparable to *Connex Cert*'s feature set; you can use these to develop applications that can be migrated to *Connex Cert* with minimal changes. We describe this in detail in *Building Connex Micro with Compatibility for Connex Cert*. For more information on *Connex Cert*, please contact RTI.

**Warning:** Note that *Connex Micro* is **not** safety certifiable. The *Connex Cert*-comparable C libraries included with *Connex Micro* are intended for development purposes **only**.

### 1.1.2 Downloading Connex Micro

Once you have purchased *Connex Micro*, you can download the package files from [support.rti.com](http://support.rti.com). *Connex Micro* is provided in the following RTI package files (`.rtipkg`):

- `rti_connex_dds_micro-<version>-host.rtipkg`
- `rti_connex_dds_micro-<version>-target-<architecture>.rtipkg`

Download the host package with `<version>` 4.2.0, and a target package with `<version>` 4.2.0 and `<architecture>` matching your CPU and compiler. If you are unsure which `<architecture>` package to download for your platform, refer to the *Supported Platforms and Programming Languages* section and consult the **RTI Architecture Abbreviations** column in the Supported Platforms (PIL) table.

---

**Tip:** On the **Downloads** page, the **Info** button lists which versions of *Connex Professional* and *Connex Drive* are compatible with your *Connex Micro* release. See also the *Compatibility* section.

---

### 1.1.3 Installing Connex Micro

You can install *Connex Micro* from either the *RTI Launcher* GUI or a command line.

---

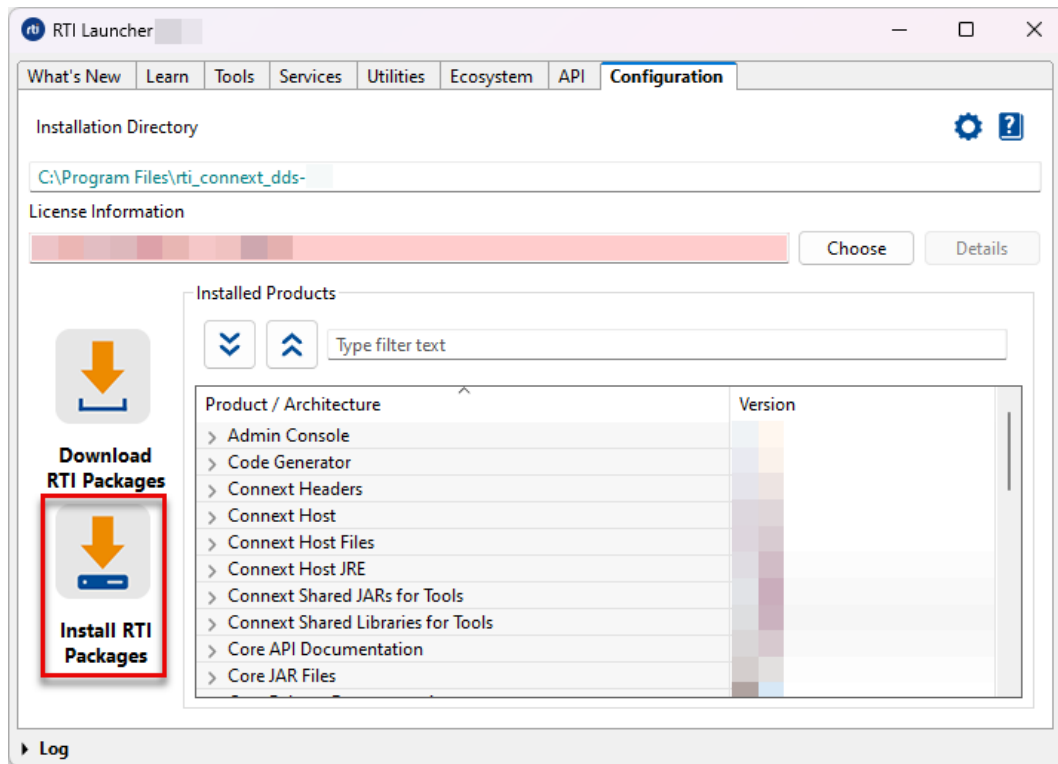
**Note:** You must first install *Connex Professional* or *Connex Drive* before installing *Connex Micro* packages. Check the *Compatibility* section for the right versions of these products to use, then refer to the [RTI Connex Installation Guide](#) or [Installing Connex Drive](#) for installation instructions.

---

## Installing with a GUI

After you install *Connext Professional*, you'll see an `rti_connext_dds-<version>` directory. In that directory, you'll have a tool called *RTI Launcher*.

To install the *Connext Micro* packages from the *Launcher* tool, open the Configuration tab, and select "Install RTI Packages." This will open a dialog that allows you to select one or more **.rtipkg** files that you would like to install.



Select the host and target packages in the dialog, then click **Install**.

Once installed, you will find a directory called `rti_connext_dds_micro-<version>` in the `rti_connext_dds-<version>` directory.

## Installing from a command line

To install *Connext Micro* from the command line, type:

Linux

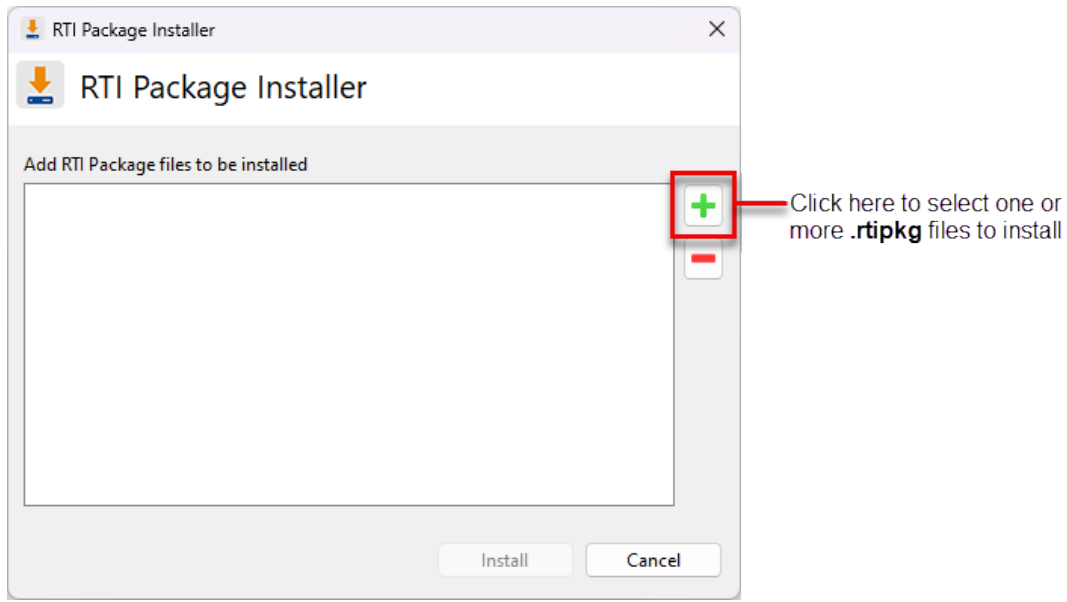
```
$ rti_connext_dds-<version>/bin/rtipkginstall <path to rtipkg>
```

macOS

```
$ rti_connext_dds-<version>/bin/rtipkginstall <path to rtipkg>
```

Windows





```
> rti_connext_dds-<version>\bin\rtipkginstall.bat <path to rtipkg>
```

Run this command for both the host and target `.rtipkg` files.

### Paths mentioned in documentation

This documentation refers to the following directories, depending on your operating system:

Linux

- `$NDDSHOME` This refers to the installation directory for *Connext*.
- `$RTIMEHOME` This refers to the installation directory for *Connext Micro*.

The default installation paths are:

- Non-root user:
  - *Connext*: `/home/<your user name>/rti_connext_dds-<version>`
  - *Connext Micro*: `/home/<your user name>/rti_connext_dds-<version>/rti_connext_dds_micro-<version>`
- Root user:
  - *Connext*: `/opt/rti_connext_dds-<version>`
  - *Connext Micro*: `/opt/rti_connext_dds-<version>/rti_connext_dds_micro-<version>`

`$NDDSHOME` is an environment variable set to the installation path for *Connext*. `$RTIMEHOME` is an environment variable set to the installation path for *Connext Micro*.

macOS

- `$NDDSHOME` This refers to the installation directory for *Connext*.
- `$RTIMEHOME` This refers to the installation directory for *Connext Micro*.

The default installation paths are:

- *Connext*: `/Applications/rti_connext_dds-<version>`
- *Connext Micro*: `/Applications/rti_connext_dds-<version>/rti_connext_dds_micro-<version>`

`$NDDSHOME` is an environment variable set to the installation path for *Connext*. `$RTIMEHOME` is an environment variable set to the installation path for *Connext Micro*.

Windows

- `%NDDSHOME%` This refers to the installation directory for *Connext*.
- `%RTIMEHOME%` This refers to the installation directory for *Connext Micro*.

The default installation paths are:

- User without Administrator privileges:
  - *Connext*: `<your home directory>\rti_connext_dds-<version>`
  - *Connext Micro*: `<your home directory>\rti_connext_dds-<version>\rti_connext_dds_micro-<version>`
- User with Administrator privileges:
  - *Connext*: `"C:\Program Files\rti_connext_dds-<version>"`
  - *Connext Micro*: `"C:\Program Files\rti_connext_dds-<version>\rti_connext_dds_micro-<version>"`

`%NDDSHOME%` is an environment variable set to the installation path for *Connext*. `%RTIMEHOME%` is an environment variable set to the installation path for *Connext Micro*.

---

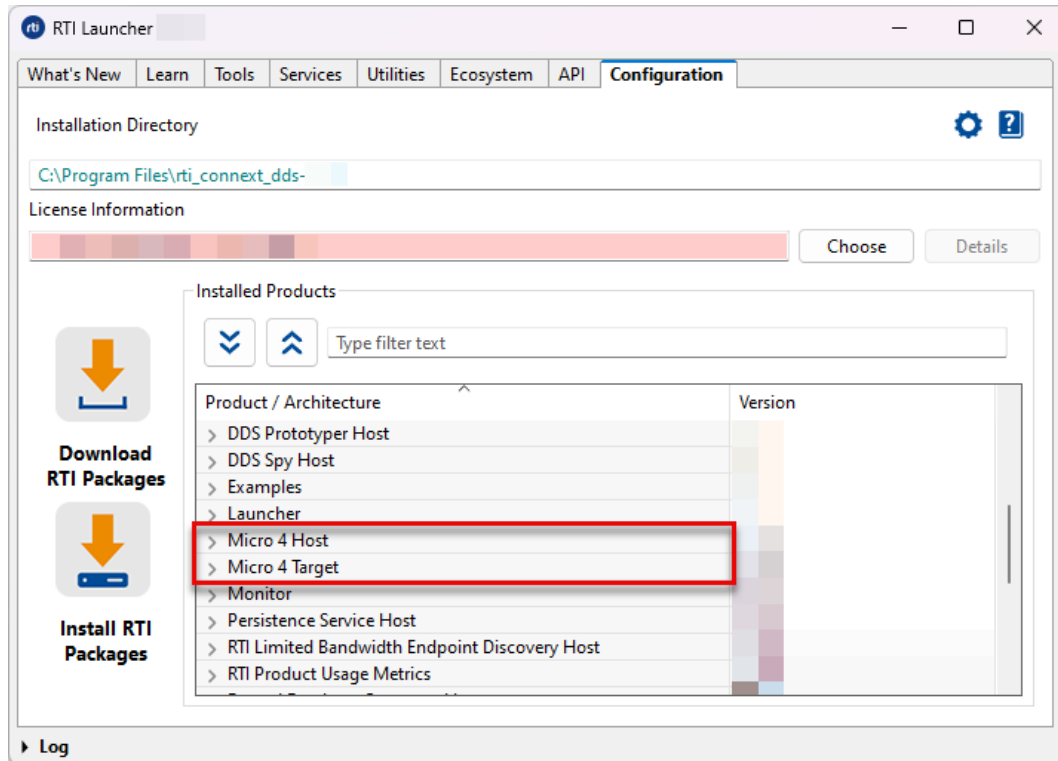
**Note:** When using a command prompt to enter a command that includes the path `C:\Program Files` (or any directory name that has a space), enclose the path in quotation marks. For example: `"C:\Program Files\rti_connext_dds-version\bin\rtilauncher.bat"`. Or if you have defined the `NDDSHOME` environment variable: `"%NDDSHOME%\bin\rtilauncher.bat"`.

---

Sometimes this documentation uses `<NDDSHOME>` to refer to the installation path for *Connext*, or `<RTIMEHOME>` for *Connext Micro*. Whenever you see `<NDDSHOME>` or `<RTIMEHOME>` used in a path, replace it with `$NDDSHOME` or `$RTIMEHOME` for Linux or macOS systems, with `%NDDSHOME%` or `%RTIMEHOME%` for Windows systems, or with your installation path.

### 1.1.4 Checking what is installed

To confirm that you have installed *Connext Micro*, you can use the *Launcher* tool. Once in *Launcher*, open the Configuration tab:



If you do not see both the *Connext Micro* host and target listed, you are missing one of the *Connext Micro* packages in your installation.

### 1.1.5 Where do I get more help?

Additional documentation and user forums can be found on [community.rti.com](https://community.rti.com).

Continue to *Hello World with Publish/Subscribe* to start learning about the capabilities and features of *Connext Micro*.

## 1.2 Hello World with Publish/Subscribe

In this section we are going to learn about the Publish/Subscribe model using a simple “Hello World” application. You will learn about the *RTI Code Generator* and how to create a *DataWriter* and a *DataReader* to publish and subscribe to data.

Publish/Subscribe is a communications model where data producers “publish” data and data consumers “subscribe” to data. These publishers and subscribers don’t need to know about each other ahead of time; they discover each other dynamically at runtime. The data they share is described by a “topic,” and publishers and subscribers send and receive data only for the topics they are

interested in. In this pattern, many publishers may publish the same topic, and many subscribers may subscribe to the same topic. Subscribers receive data from all of the publishers that they share a topic with. Publishers send data directly to subscribers, with no need for a broker or centralized application to mediate communications.

### 1.2.1 Introduction to DataWriters, DataReaders, and Topics

*Connext Micro* conforms to the [OMG Data Distribution Service \(DDS\)](#) connectivity standard. In DDS, the objects that actually publish data are called *DataWriters*, and the objects that subscribe to data are *DataReaders*. *DataWriters* and *DataReaders* are associated with a single *Topic* object that describes that data. (DDS also has *Publisher* and *Subscriber* objects, but we will talk about them later.) An application typically has a combination of *DataWriters* and *DataReaders*.

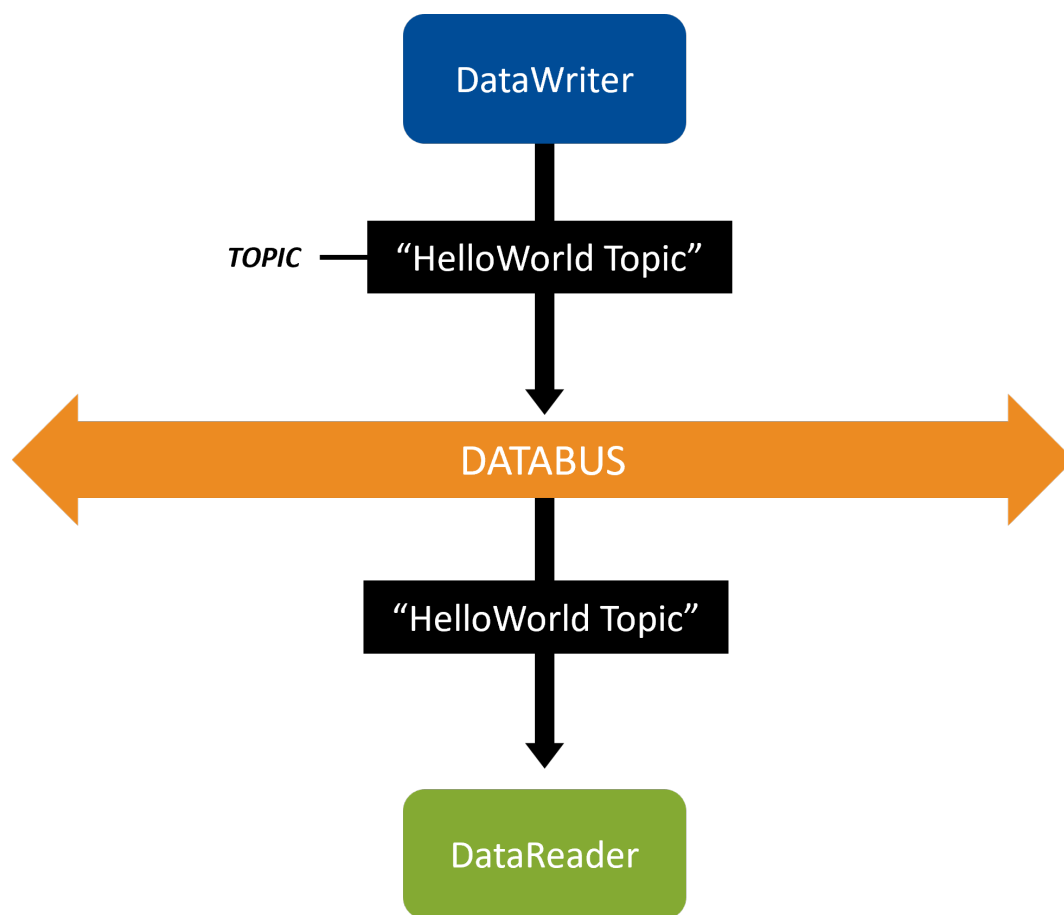


Figure 1.1: *DataWriters* write data and *DataReaders* read data of a *Topic*.

*Connext Micro* is responsible for discovering *DataWriters* and *DataReaders* in a system, checking if they have a matching *Topic* (and compatible quality of service, which we will discuss later) and then providing the communication between those *DataWriters* and *DataReaders*. Logically, this means you can visualize your applications as having *DataWriters* and *DataReaders* that connect to a “databus,” because your applications are not specifying exactly which other applications they communicate with—they only specify which *Topics* they read from and write to the databus, and

*Connext Micro* sets up the communication. Note that there is no “databus” object or *Connext Micro* service deployed in your system—this is simply a logical way to visualize systems in which you don’t have to configure each communication path.

### 1.2.2 Create a DataWriter and DataReader

We are going to start with a simple “Hello World” application in C to show how to use the *Code Generator*, and how to create a *DataWriter* and a *DataReader*.

---

**Tip:** By the end of this exercise, a publishing application will send data, and a subscribing application will receive and print it to the console.

---

#### Set environment variables

Before we begin, we need to set some environment variables:

1. Set the `RTIMEHOME` environment variable to the installation directory path for *RTI Connext Micro*.

If you installed *Connext Professional* with default settings, *RTI Connext Micro* will be here: `<path_to_connext_dds_installation>/rti_connext_dds-<version>/rti_connext_micro-<version>`. If you copied *RTI Connext Micro* to another place, set `RTIMEHOME` to point to that location.

2. Add `<RTIMEHOME>/rtiddsgen/scripts` to your path environment variable folder.
3. Add `<RTIMEHOME>/resource/scripts` to your path environment variable folder.

#### Generate an example application

We’ll use the *RTI Code Generator* (also referred to as *rtiddsgen*) included with *Connext Micro* to generate a DDS example application. To do that, we’ll start with a type definition file as input.

#### Create a type definition file

Create a **HelloWorld.idl** file with the following type definition:

```
// Hello world!
struct HelloWorld {

    // String with maximum length of 256 characters
    string<256> msg;

};
```

In this file, the data **HelloWorld** is described in the language-independent Interface Definition Language (IDL). IDL allows you to declare data types used for communication. The RTI *Code Generator* (*rtiddsgen*) translates from this language-independent data type into code specific for your programming language. The generated code serializes and deserializes your data into and out of a network format.

## Run Code Generator

**Tip:** Before running *rtiddsgen*, make sure you added `rti_connext_dds_micro-<version>/rtiddsgen/scripts` to your path environment variable folder as described in *Set environment variables*. Run the following command to check:

```
rtiddsgen -version
```

You should see the following output:

```
INFO com.rti.ndds.nddsgen.Main rtiddsgen version <version>
INFO com.rti.ndds.nddsgen.Main Connext DDS templates: XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-
↳XXXX
INFO com.rti.ndds.nddsgen.Main Connext Micro templates: XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-
↳XXXX-XXXX
INFO com.rti.ndds.nddsgen.Main Done
```

If you don't see **Connext Micro templates:** in the output, you might be using a version of *rtiddsgen* from *Connext Professional*. Make sure you added *rtiddsgen* to the environment path from `rti_connext_dds_micro-<version>/rtiddsgen/scripts`.

From a terminal or command prompt, change directories to where you created **HelloWorld.idl** and run the following command:

Linux

```
$ rtiddsgen -example -language C HelloWorld.idl
```

macOS

```
$ rtiddsgen -example -language C HelloWorld.idl
```

Windows

```
> rtiddsgen -example -language C -ppDisable HelloWorld.idl
```

`-ppDisable` disables the preprocessor. It is necessary for running *rtiddsgen* on a Windows system if the preprocessor is not in your path. You can only use `-ppDisable` if your IDL does not contain any C preprocessor directives, as here—otherwise you must add the preprocessor to your path.

## Overview of generated and example code

The code you just generated includes the files in Table 1.1.

Table 1.1: Generated and Example Code Files

Files	Description
<ul style="list-style-type: none"> <li>• HelloWorld.h</li> <li>• HelloWorld.c</li> <li>• HelloWorldPlugin.h</li> <li>• HelloWorldPlugin.c</li> <li>• HelloWorldSupport.h</li> <li>• HelloWorldSupport.c</li> </ul>	<p>The C definition of your data type, and the code used to serialize and deserialize it (convert it to a format for the network). This is the type-specific code that will be used in your real application. These files were generated when you ran <i>rtiddsgen</i>.</p>
<ul style="list-style-type: none"> <li>• HelloWorldApplication.h</li> <li>• HelloWorldApplication.c</li> <li>• HelloWorld_publisher.c</li> <li>• HelloWorld_subscriber.c</li> </ul>	<p>Example application code. It includes code for two applications you can read and modify. These will compile into separate applications, <b>HelloWorld_publisher</b> and <b>HelloWorld_subscriber</b>. These files were generated when you ran <i>rtiddsgen</i>.</p>
<ul style="list-style-type: none"> <li>• CMakeLists.txt</li> </ul>	<p>Used to compile your application with CMake. This file was generated when you ran <i>rtiddsgen</i>.</p>
<ul style="list-style-type: none"> <li>• README.txt</li> </ul>	<p>Instructions for how to open and modify the files, compile, and run the example. This file was generated when you ran <i>rtiddsgen</i>.</p>

## Modify application

1. In your example application directory, open `HelloWorldApplication.c` in your IDE of choice.

This snippet shows how to create a *Topic* (with a name and data type):

```
sprintf(application->topic_name, "Example HelloWorld");
application->topic = DDS_DomainParticipant_create_topic(
    application->participant,
    application->topic_name,
    application->type_name,
    &DDS_TOPIC_QOS_DEFAULT,
    NULL,
    DDS_STATUS_MASK_NONE);
if (application->topic == NULL)
{
    printf("topic == NULL\n");
    goto done;
}
```

2. Change the *Topic* name from “Example HelloWorld” to “HelloWorld Topic”:

```
sprintf(application->topic_name, "HelloWorld Topic");
application->topic = DDS_DomainParticipant_create_topic(
    application->participant,
    application->topic_name,
    application->type_name,
    &DDS_TOPIC_QOS_DEFAULT,
    NULL,
    DDS_STATUS_MASK_NONE);
if (application->topic == NULL)
{
    printf("topic == NULL\n");
    goto done;
}
```

## Modify publisher

1. Open `HelloWorld_publisher.c` in your IDE of choice.

The following snippet shows how to write a HelloWorld update using the *DataWriter*’s write method:

```
for (i = 0; (application->count <= 0) || (i < application->count); ++i)
{
    /* TODO set sample attributes here */

    retcode = HelloWorldDataWriter_write(
        hw_datawriter,
        sample,
```

(continues on next page)



(continued from previous page)

```

        &DDS_HANDLE_NIL);
    if (retcode != DDS_RETCODE_OK)
    {
        printf("Failed to write sample\n");
    }
    else
    {
        printf("Written sample %d\n", (i+1));
    }

    OSAPI_Thread_sleep((RTI_UINT32)application->sleep_time);
}

```

2. Add the following code to send the message “Hello world!” with a count:

```

for (i = 0; (application->count <= 0) || (i < application->count); ++i)
{
    /* TODO set sample attributes here */

    snprintf(sample->msg, 128, "Hello World (%d) ! \n", i);

    retcode = HelloWorldDataWriter_write(
        hw_datawriter,
        sample,
        &DDS_HANDLE_NIL);
    if (retcode != DDS_RETCODE_OK)
    {
        printf("Failed to write sample\n");
    }
    else
    {
        printf("Written sample %d\n", (i+1));
    }

    OSAPI_Thread_sleep((RTI_UINT32)application->sleep_time);
}

```

Recall that your “HelloWorld Topic” describes your data. This *Topic* is associated with the data type **HelloWorld**, which is defined in the IDL file (see *Run Code Generator*). The data type **HelloWorld** contains a string field named **msg**. In this step, you have just added code to set a value for the **msg** field. Now, when the *DataWriter* writes data, the **msg** field in the data will contain the string “Hello world! 1”, “Hello world! 2”, etc.

---

## Definition

A **sample** is a single update to a *Topic*, such as “Hello world (1) !”. Every time an application calls `write()`, it is “writing a sample.” Every time an application receives data, it is “receiving a sample.”

Note that samples don’t necessarily overwrite each other. For example, if you set up a

**RELIABLE** Quality of Service (QoS)<sup>1</sup> with a History **kind** of KEEP\_ALL, all samples will be saved and accumulate.

## Modify subscriber

1. Open HelloWorld\_subscriber.c in your IDE of choice.

The following snippet shows how to process a sample:

```
for (i = 0; i < HelloWorldSeq_get_length(&sample_seq); ++i)
{
    sample_info = DDS_SampleInfoSeq_get_reference(&info_seq, i);

    if (sample_info->valid_data)
    {
        sample = HelloWorldSeq_get_reference(&sample_seq, i);
        printf("\nValid sample received\n");
        *total_samples += 1;

        /* TODO read and process sample attributes here */
        (void)sample;
    }
    else
    {
        printf("\nSample received\n\tINVALID DATA\n");
    }
}
```

2. Add the following code to display the received messages:

```
for (i = 0; i < HelloWorldSeq_get_length(&sample_seq); ++i)
{
    sample_info = DDS_SampleInfoSeq_get_reference(&info_seq, i);

    if (sample_info->valid_data)
    {
        sample = HelloWorldSeq_get_reference(&sample_seq, i);
        printf("\nValid sample received\n");
        *total_samples += 1;

        /* TODO read and process sample attributes here */
        printf("%s\n", sample->msg);
    }
    else
    {
        printf("\nSample received\n\tINVALID DATA\n");
    }
}
```

<sup>1</sup> For more information on QoS policies, refer to [QoS Policies](#) in the API Reference.

## Compile your changes

Now that you have made changes to the example code, compile the code with your modifications. Run the following command from the example directory:

Linux

```
$ rtime-make --config Debug --build --target x86_64leElfgcc12.3.0-Linux5 --source-dir . -
↳G "Unix Makefiles" --delete
```

macOS

```
$ rtime-make --config Debug --build --target x86_64leMachOclang15.0-Darwin23 --source-
↳dir . -G "Unix Makefiles" --delete
```

Windows

```
> rtime-make.bat --config Debug -A x64 --target self --name x86_64lePEvs2017-Win10 --
↳build --source-dir .
```

After running the command, you should see two applications in the `objs/<architecture>` or `objs/<architecture>/Debug` directory:

- HelloWorld\_publisher
- HelloWorld\_subscriber

## Run the applications

**Note:** By default, this example uses two interfaces to receive samples. Since your installation might use different names for these interfaces (which would prevent communication), we'll use the option `-udp_intf <interface name>` to specify them in the following commands. You can use command `ifconfig` (on Linux or macOS systems) or `ipconfig` (on Windows systems) to know the name of all available interfaces.

1. From within the `objs/<architecture>` or `objs/<architecture>/Debug` directory, enter the following full path:

Linux

```
$ objs/x86_64leElfgcc12.3.0-Linux5/Debug/HelloWorld_publisher -udp_intf <interface_
↳name>
```

macOS

```
$ objs/x86_64leMachOclang15.0-Darwin23/Debug/HelloWorld_publisher -udp_intf
↳<interface name>
```

Windows

```
> objs\x86_64lePEvs2017-Win10\Debug\HelloWorld_publisher.exe -udp_intf <interface_
↵name>
```

You should see this in the window for the publisher:

```
Written sample 1
Written sample 2
Written sample 3
Written sample 4
...

```

2. Open another command prompt window, and from within the example directory, enter the following full path:

Linux

```
$ objs/<architecture>/Debug/HelloWorld_subscriber -udp_intf <interface name>
```

macOS

```
$ objs/<architecture>/Debug/HelloWorld_subscriber -udp_intf <interface name>
```

Windows

```
> objs\<architecture>\Debug\HelloWorld_subscriber.exe -udp_intf <interface name>
```

You should see this in the window for the subscriber:

```
Valid sample received
Hello World (1) !

Subscriber sleeping for 1000 msec...

Valid sample received
Hello World (2) !

Subscriber sleeping for 1000 msec...

Valid sample received
Hello World (3) !

Subscriber sleeping for 1000 msec...

Valid sample received
Hello World (4) !
...

```

The `Hello world <count> !` line is the data being sent by the *DataWriter*. If the *DataWriter* weren't communicating with the *DataReader*, you would just see the `Subscriber sleeping for 1000 msec...` lines and not the “msg” lines. (The subscribing application prints the

“timed out” lines after the `WaitSet` times out while waiting for data, then it prints the “msg:” lines when it receives data from the *DataWriter*.)

## Taking it further

Under the hood, the publishing and subscribing applications are doing a lot of work:

Before communication starts, the *DataWriter* and *DataReader* discover each other and check that they have the same *Topic* name, compatible data types, and compatible QoS. After discovery, the *DataWriter* sends data directly to the *DataReader*, with no message broker required.

When you run the applications on the same machine, by default they communicate over shared memory. If you run one on another machine, they communicate over the network using UDP.

## Start up multiple publishing or subscribing applications

Try starting up multiple publishing or subscribing applications, and you will see that they will also send or receive data. (Remember to run from the example directory.)

```

Written sample 1
Matched a subscriber
Matched a subscriber
Written sample 2
Written sample 3
Written sample 4
Written sample 5
Written sample 6
Written sample 7
Written sample 8
Written sample 9
Written sample 10
Written sample 11
Written sample 12
Written sample 13
Written sample 14
Written sample 15

Valid sample received
Hello World (1) !

Valid sample received
Hello World (11) !

Subscriber sleeping for 1000 msec...

Valid sample received
Hello World (2) !

Valid sample received
Hello World (12) !

Written sample 1
Matched a subscriber
Matched a subscriber
Written sample 2
Written sample 3
Written sample 4
Written sample 5
Written sample 6
Written sample 7
Written sample 8
Written sample 9
Written sample 10
Written sample 11
Written sample 12
Written sample 13
Written sample 14
Written sample 15

Valid sample received
Hello World (1) !

Valid sample received
Hello World (11) !

Subscriber sleeping for 1000 msec...

Valid sample received
Hello World (2) !

Valid sample received
Hello World (12) !

Subscriber sleeping for 1000 msec...

```

Figure 1.2: Two applications publishing, and two subscribing, to the same *Topic*. Notice that each subscriber is receiving data from both publishers.

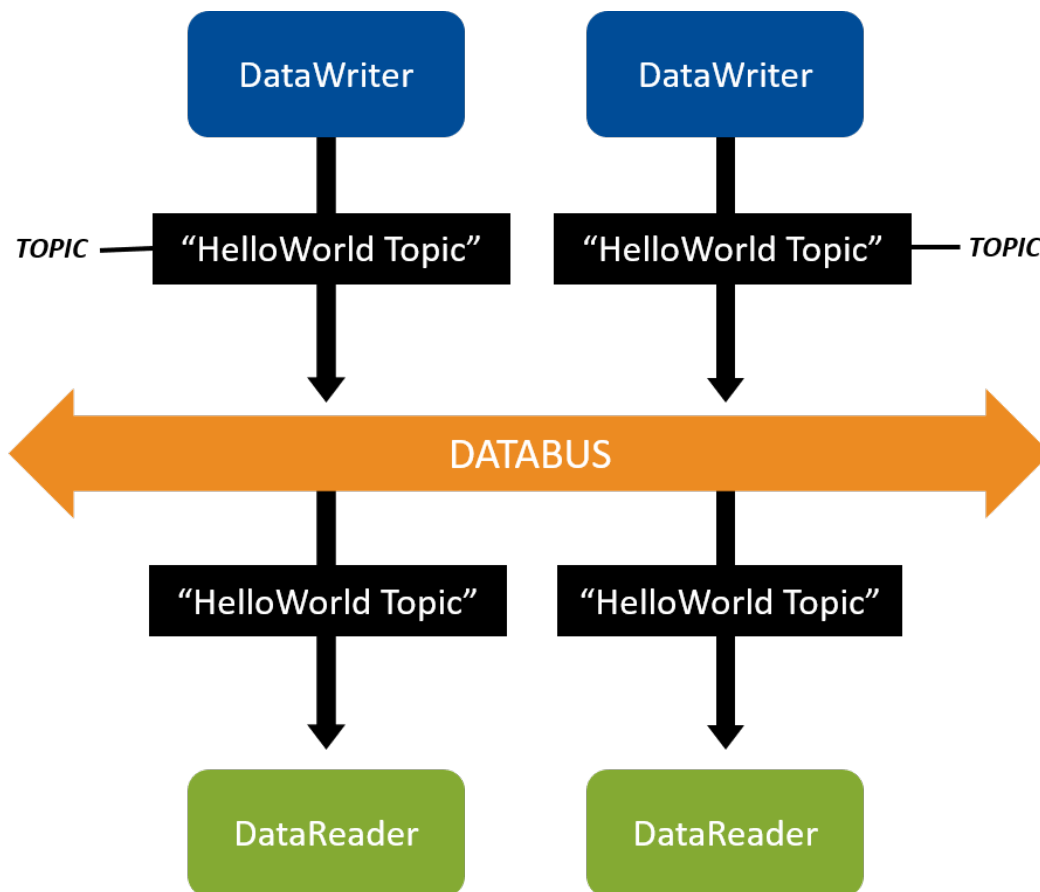


Figure 1.3: Two applications publishing, and two subscribing, to the same *Topic*.

### 1.2.3 Troubleshooting

#### Why aren't my applications communicating?

If you are running both applications and they aren't communicating (you don't see the `Hello world <count> !` lines shown at the end of *Run the applications*), here are some things to check:

- Make sure you specified the correct interfaces for both your applications with the `-udp_intf` option, as described in *Run the applications*.
- Check if your interface supports local loopback. If it doesn't, try using an interface that does, or run the application between two host machines and change the peer on each one to be the host address of the other machine.
- See if your machine uses a firewall. You may need to disable your firewall or configure it to allow multicast traffic for communication to be established.

#### Why does the `DataReader` miss the first samples?

Discovery is not an instantaneous event. It takes some time for the discovery process between applications to complete. The *DataWriter* and *DataReader* must discover each other before they can start communicating. Therefore, if you send data immediately after creating the *Connext Micro* entities, *DataReaders* will not receive the first few samples because they are still in-process of discovering the *DataWriters* and vice versa. This is true even when the *DataWriters* and *DataReaders* are reliable, because the [Reliability QoS Policy](#) on its own does not guarantee delivery to *DataReaders* that have not been discovered yet.

You can overcome this behavior with the [Durability QoS Policy](#), which can be set to deliver historical samples (that *DataWriters* already sent) to late-joining *DataReaders*. The *DataReaders* will then receive the first samples they originally missed.

### 1.2.4 Next steps

Congratulations! You've written your first *Connext Micro* application. You've experienced a quick overview of the development process from defining a data type and using the *RTI Code Generator*, to building an example application.

Continue to *Developing Applications* for a deeper look at the process of developing your own *Connext Micro* applications.

## Chapter 2

# Developing Applications

This section describes how to write *Connex Micro* applications. It covers preparing your development environment, defining data types, generating support code for your data types, and creating the entities that publish and subscribe to data.

For a deeper dive into *Connex Micro*'s features, refer to the *User's Manual*.

### 2.1 Prepare Your Development Environment

This section describes how to set up your development environment for *Connex Micro* applications, such as the required environment variables, compilers, compiler definitions, and libraries.

#### 2.1.1 Set environment variables

The `RTIMEHOME` environment variable must be set to the installation directory path for *RTI Connex Micro*. If you installed *RTI Connex* with default settings, *RTI Connex Micro* will be here: `<path_to_connex_dds_installation>/rti_connex_dds-<version>/rti_connex_micro-<version>`. If you copied *RTI Connex Micro* to another place, set `RTIMEHOME` to point to that location.

#### 2.1.2 Add required preprocessor flags

All *Connex Micro* applications require the following preprocessor defines:

```
-IRTIMEHOME/include  
-IRTIMEHOME/include/rti_me
```

Add the following preprocessor defines, according to your platform and compiler:

Windows

Using MSVSCC:



```
-DOSAPI_CC_DEF_H=osapi/osapi_cc_msvsc.h
-DRTI_WIN32
```

macOS

Using clang:

```
-DOSAPI_CC_DEF_H=osapi/osapi_cc_clang.h
-DRTI_DARWIN
```

Linux

Using GCC:

```
-DOSAPI_CC_DEF_H=osapi/osapi_cc_gcc.h
-DRTI_LINUX
```

QNX

Using QCC:

```
-DOSAPI_CC_DEF_H=osapi/osapi_cc_qcc.h
-DRTI_QNX
```

### 2.1.3 Link applications and libraries

Add the library path for both the PIL and PSL to the linker's search path:

- RTIMEHOME/lib/<arch>/ (PIL)
- RTIMEHOME/lib/<arch>-<PSL>/ (PSL)

---

**Note:** When executing executables that are linked with the *Connex Micro* shared libraries, you must add the path to the PIL architecture directory to the runtime linker's search path.

---

To link a C application, the libraries are required in the following order:

- RTIMEHOME/lib/<arch>/
  1. `rtdi_me_appgen` (if using the *Application Generation Using XML* plugin)
  2. `rtdi_me_netiosdm` (if using *Zero Copy Transfer* with the *Shared Memory Transport (SHMEM)*)
  3. `rtdi_me_netiozcopy` (if using the *Zero Copy v2 Transport*)
  4. `rtdi_me_netioshmem` (if using the *Shared Memory Transport (SHMEM)*)
  5. `rtdi_me_discdpde` (if using DPDE)
  6. `rtdi_me_discdpse` (if using DPSE)
  7. `rtdi_me_ddsfilter` (if using *Content Filtering*)

8. `rti_me_ddsxtypes` (if using X-Types)
9. `rti_me_rhsm`, `rti_me_whsm`, and `rti_me` (always required)
- `RTIMEHOME/lib/<arch>-<PSL>/`
  10. `rti_me_netioysl` (when building with a Platform Independent Library)
  11. `rti_me_ospsl` (when building with a Platform Independent Library)

To link a C++ application, the libraries are required in the following order:

- `RTIMEHOME/lib/<arch>-<PSL>/`
  1. `rti_me_appgen` (if using the *Application Generation Using XML* plugin)
  2. `rti_me_ddsxtypes` (if using X-Types)
  3. `rti_me_cpp`, `rti_me_netiosdm`, `rti_me_netiozcopy`, `rti_me_netioshmem`, `rti_me_discdpde`, `rti_me_discdpse`, `rti_me_ddsfilter`<sup>1</sup>, `rti_me_rhsm`, `rti_me_whsm`, and `rti_me` (always required)
- `RTIMEHOME/lib/<arch>-<PSL>/`
  4. `rti_me_netioysl_cpp` (when building with a Platform Independent Library)
  5. `rti_me_netioysl` (when building with a Platform Independent Library)
  6. `rti_me_ospsl` (when building with a Platform Independent Library)

## 2.2 Define a Data Type

To distribute data using *Connext Micro*, you must first define a data type, then run the *rtiddsgen* utility. This utility will generate the type-specific support code that *Connext Micro* needs and the code that makes calls to publish and subscribe to that data type.

*Connext Micro* accepts types definitions in Interface Definition Language (IDL) format.

For instance, the HelloWorld examples provided with *Connext Micro* use this simple type, which contains a string “msg” with a maximum length of 128 chars:

```
struct HelloWorld {
    long id; //@key
    string<128> msg;
    sequence<octet, 1000> image;
};
```

For more details, see *Data Types* in the *User's Manual*.

<sup>1</sup> Only required if using *Content Filtering*.

## 2.3 Generate Type Support Code with rtiddsgen

You will provide your IDL as an input to *rtiddsgen*. *rtiddsgen* supports code generation for the following standard types:

- `octet`, `char`, `wchar`
- `short`, `unsigned short`
- `long`, `unsigned long`
- `long long`, `unsigned long long`, `float`
- `double`, `long double`
- `boolean`
- `string`
- `struct`
- `array`
- `enum`
- `wstring`
- `sequence`
- `union`
- `typedef`
- `value type`

*rtiddsgen* is in `<your_top_level_dir>/rti_connext_dds-7.3.0/rti_connext_micro-4.0.1/rtiddsgen/scripts`.

To generate support code for data types in a file called `HelloWorld.idl`:

```
rtiddsgen -micro -language C -replace HelloWorld.idl
```

Run `rtiddsgen -help` to see all available options. For the options used here:

- The `-micro` option is necessary to generate support code specific to *Connext Micro*, as *rtiddsgen* is also capable of generating support code for *Connext*, and the generated code for the two are different.
- The `-language` option specifies the language of the generated code. *Connext Micro* supports C and C++ (with `-language C++`).
- The `-replace` option specifies that the new generated code will replace, or overwrite, any existing files with the same name.

*rtiddsgen* generates the following files for an input file `HelloWorld.idl`:

- **`HelloWorld.h` and `HelloWorld.c`.** Operations to manage a sample of the type, and a DDS sequence of the type.

- **HelloWorldPlugin.h and HelloWorldPlugin.c.** Implements the type-plugin interface defined by *Connext Micro*. Includes operations to serialize and deserialize a sample of the type and its DDS instance keys.
- **HelloWorldSupport.h and HelloWorldSupport.c.** Support operations to generate a type-specific *DataWriter* and *DataReader*, and to register the type with a DDS *DomainParticipant*.

This release of *Connext Micro* supports a new way to generate support code for IDL Types that will generate a *TypeCode* object containing information used by an interpreter to serialize and deserialize samples. Prior to this release, the code for serialization and deserialization was generated for each type. To disable generating code to be used by the interpreter, use the `-interpreted 0` command-line option to generate code. This option generates code in the same way as was done in previous releases.

For more details, see *Generating Type Support* in the *User's Manual*.

## 2.4 Create an Application

The rest of this guide will walk you through the steps to create an application and will provide example code snippets. It assumes that you have defined your types (see *Define a Data Type*) and have used *rtiddsgen* to generate their support code (see *Generate Type Support Code with rtiddsgen*).

### 2.4.1 Registry Configuration

The [DomainParticipantFactory](#), in addition to its standard role of creating and deleting *DomainParticipants*, contains the RT Registry that a new application registers with some necessary components.

The *Connext Micro* architecture defines a run-time (RT) component interface that provides a generic framework for organizing and extending functionality of an application. An RT component is created and deleted with an RT component factory. Each RT component factory must be registered within an RT registry in order for its components to be usable by an application.

*Connext Micro* automatically registers components that provide necessary functionality. These include components for DDS *Writers* and *Readers*, the RTPS protocol, and the UDP transport.

In addition, every DDS application must register three components:

- **Writer History.** Queue of written samples of a *DataWriter*. Must be registered with the name “wh”.
- **Reader History.** Queue of received samples of a *DataReader*. Must be registered with the name “rh”.
- **Discovery (DPDE or DPSE).** Discovery component. Choose either dynamic (DPDE) or static (DPSE) endpoint discovery.

Example source:

- Get the RT Registry from the [DomainParticipantFactory](#) singleton:

```
DDS_DomainParticipantFactory *factory = NULL;
RT_Registry_T *registry = NULL;

factory = DDS_DomainParticipantFactory_get_instance();
registry = DDS_DomainParticipantFactory_get_registry(factory);
```

- Register the Writer History and Reader History components with the registry:

```
/* Register Writer History */
if (!RT_Registry_register(registry, "wh",
                        WHSM_HistoryFactory_get_interface(), NULL, NULL))
{
    /* failure */
}

/* Register Reader History */
if (!RT_Registry_register(registry, "rh",
                        RHSM_HistoryFactory_get_interface(), NULL, NULL))
{
    /* failure */
}
```

Only one discovery component can be registered, either DPDE or DPSE. Each has its own properties that can be configured upon registration.

- Register [DPDE](#) for dynamic participant, dynamic endpoint discovery:

```
struct DPDE_DiscoveryPluginProperty discovery_plugin_properties =
    DPDE_DiscoveryPluginProperty_INITIALIZER;

/* Configure properties */
discovery_plugin_properties.participant_liveliness_assert_period.sec = 5;
discovery_plugin_properties.participant_liveliness_assert_period.nanosec = 0;
discovery_plugin_properties.participant_liveliness_lease_duration.sec = 30;
discovery_plugin_properties.participant_liveliness_lease_duration.nanosec = 0;

/* Register DPDE with updated properties */
if (!RT_Registry_register(registry,
                        "dpde",
                        DPDE_DiscoveryFactory_get_interface(),
                        &discovery_plugin_properties._parent,
                        NULL))
{
    /* failure */
}
```

- Register [DPSE](#) for dynamic participant, static endpoint discovery:

```
struct DPSE_DiscoveryPluginProperty discovery_plugin_properties =
    DPSE_DiscoveryPluginProperty_INITIALIZER;
```

(continues on next page)

(continued from previous page)

```

/* Configure properties */
discovery_plugin_properties.participant_liveliness_assert_period.sec = 5;
discovery_plugin_properties.participant_liveliness_assert_period.nanosec = 0;
discovery_plugin_properties.participant_liveliness_lease_duration.sec = 30;
discovery_plugin_properties.participant_liveliness_lease_duration.nanosec = 0;

/* Register DPSE with updated properties */
if (!RT_Registry_register(registry,
                        "dpse",
                        DPSE_DiscoveryFactory_get_interface(),
                        &discovery_plugin_properties._parent,
                        NULL))
{
    printf("failed to register dpse\n");
    goto done;
}

```

For more information, see the *Application Generation Using XML* section in the User's Manual.

## 2.5 Configure UDP Transport

You may need to configure the UDP transport component that is pre-registered by *RTI Connext Micro*. To change the properties of the UDP transport, first the UDP component has be unregistered, then the properties have to be updated, and finally the component must be re-registered with the updated properties.

Example code:

- Unregister the pre-registered UDP component:

```

/* Unregister the pre-registered UDP component */
if (!RT_Registry_unregister(registry, "_udp", NULL, NULL))
{
    /* failure */
}

```

- Configure UDP transport properties:

```

struct UDP_InterfaceFactoryProperty *udp_property = NULL;

udp_property = (struct UDP_InterfaceFactoryProperty *)
    malloc(sizeof(struct UDP_InterfaceFactoryProperty));
if (udp_property != NULL)
{
    *udp_property = UDP_INTERFACE_FACTORY_PROPERTY_DEFAULT;

    /* allow_interface: Names of network interfaces allowed to send/receive.
     * Allow one loopback (lo) and one NIC (eth0).
     */
}

```

(continues on next page)

(continued from previous page)

```

    REDA_StringSeq_set_maximum(&udp_property->allow_interface,2);
    REDA_StringSeq_set_length(&udp_property->allow_interface,2);

    *REDA_StringSeq_get_reference(&udp_property->allow_interface,0) = DDS_String_
    ↪dup("lo");
    *REDA_StringSeq_get_reference(&udp_property->allow_interface,1) = DDS_String_
    ↪dup("eth0");
}
else
{
    /* failure */
}

```

- Re-register UDP component with updated properties:

```

if (!RT_Registry_register(registry, "_udp",
                        UDP_InterfaceFactory_get_interface(),
                        (struct RT_ComponentFactoryProperty*)udp_property, NULL))
{
    /* failure */
}

```

For more details, see the *Transports* section in the User's Manual.

## 2.6 Create DomainParticipant, Topic, and Type

A [DomainParticipantFactory](#) creates *DomainParticipants*, and a *DomainParticipant* itself is the factory for creating *Publishers*, *Subscribers*, and *Topics*.

When creating a *DomainParticipant*, you may need to customize [DomainParticipantQos](#), notably for:

- **Resource limits.** Default resource limits are set at minimum values.
- **Initial peers.**
- **Discovery.** The name of the registered discovery component (“dpde” or “dpse”) must be assigned to [DiscoveryQosPolicy](#)’s name.
- **Participant Name.** Every *DomainParticipant* is given the same default name. Must be unique when using DPSE discovery.

Example code:

- Create a *DomainParticipant* with configured [DomainParticipantQos](#):

```

DDS_DomainParticipant *participant = NULL;
struct DDS_DomainParticipantQos dp_qos =
    DDS_DomainParticipantQos_INITIALIZER;

/* DDS domain of DomainParticipant */

```

(continues on next page)

(continued from previous page)

```

DDS_Long domain_id = 0;

/* Name of your registered Discovery component */
if (!RT_ComponentFactoryId_set_name(&dp_qos.discovery.discovery.name, "dpde"))
{
    /* failure */
}

/* Initial peers: use only default multicast peer */
DDS_StringSeq_set_maximum(&dp_qos.discovery.initial_peers,1);
DDS_StringSeq_set_length(&dp_qos.discovery.initial_peers,1);
*DDS_StringSeq_get_reference(&dp_qos.discovery.initial_peers,0) =
    DDS_String_dup("239.255.0.1");

/* Resource limits */
dp_qos.resource_limits.max_destination_ports = 32;
dp_qos.resource_limits.max_receive_ports = 32;
dp_qos.resource_limits.local_topic_allocation = 1;
dp_qos.resource_limits.local_type_allocation = 1;
dp_qos.resource_limits.local_reader_allocation = 1;
dp_qos.resource_limits.local_writer_allocation = 1;
dp_qos.resource_limits.remote_participant_allocation = 8;
dp_qos.resource_limits.remote_reader_allocation = 8;
dp_qos.resource_limits.remote_writer_allocation = 8;

/* Participant name */
strcpy(dp_qos.participant_name.name, "Participant_1");

participant =
    DDS_DomainParticipantFactory_create_participant(factory,
                                                    domain_id,
                                                    &dp_qos,
                                                    NULL,
                                                    DDS_STATUS_MASK_NONE);

if (participant == NULL)
{
    /* failure */
}

```

### 2.6.1 Register Type

Your data types that have been generated from IDL need to be registered with the *DomainParticipants* that will be using them. Each registered type must have a unique name, preferably the same as its IDL defined name.

```

DDS_ReturnCode_t retcode;

retcode = DDS_DomainParticipant_register_type(participant,
                                              "HelloWorld",
                                              HelloWorldTypePlugin_get());

```

(continues on next page)



(continued from previous page)

```
if (retcode != DDS_RETCODE_OK)
{
    /* failure */
}
```

## 2.6.2 Create Topic of Registered Type

DDS *Topics* encapsulate the types being communicated, and you can create *Topics* for your type once your type is registered.

A topic is given a name at creation (e.g. “Example HelloWorld”). The type associated with the *Topic* is specified with its registered name.

```
DDS_Topic *topic = NULL;

topic = DDS_DomainParticipant_create_topic(participant,
                                           "Example HelloWorld",
                                           "HelloWorld",
                                           &DDS_TOPIC_QOS_DEFAULT,
                                           NULL,
                                           DDS_STATUS_MASK_NONE);

if (topic == NULL)
{
    /* failure */
}
```

## 2.6.3 DPSE Discovery: Assert Remote Participant

DPSE Discovery relies on the application to specify the other, or remote, *DomainParticipants* that its local *DomainParticipants* are allowed to discover. Your application must call a [DPSE](#) API for each remote participant to be discovered. The API takes as input the name of the remote participant.

```
/* Enable discovery of remote participant with name Participant_2 */
retcode = DPSE_RemoteParticipant_assert(participant, "Participant_2");
if (retcode != DDS_RETCODE_OK)
{
    /* failure */
}
```

For more information, see the *DDS Domains* section in the User’s Manual.

## 2.7 Create Publisher

A publishing application needs to create a DDS *Publisher* and then a *DataWriter* for each *Topic* it wants to publish.

In *Connext Micro*, [PublisherQos](#) in general contains no policies that need to be customized, while [DataWriterQos](#) does contain several customizable policies.

- Create *Publisher*:

```
DDS_Publisher *publisher = NULL;
publisher = DDS_DomainParticipant_create_publisher(participant,
                                                    &DDS_PUBLISHER_QOS_DEFAULT,
                                                    NULL,
                                                    DDS_STATUS_MASK_NONE);

if (publisher == NULL)
{
    /* failure */
}
```

For more information, see the *Sending Data* section in the User's Manual.

## 2.8 Create DataWriter

```
DDS_DataWriter *datawriter = NULL;
struct DDS_DataWriterQos dw_qos = DDS_DataWriterQos_INITIALIZER;
struct DDS_DataWriterListener dw_listener = DDS_DataWriterListener_INITIALIZER;

/* Configure writer Qos */
dw_qos.protocol.rtps_object_id = 100;
dw_qos.reliability.kind = DDS_RELIABLE_RELIABILITY_QOS;
dw_qos.resource_limits.max_samples_per_instance = 2;
dw_qos.resource_limits.max_instances = 2;
dw_qos.resource_limits.max_samples =
    dw_qos.resource_limits.max_samples_per_instance * dw_qos.resource_limits.max_
↪instances;
dw_qos.history.depth = 1;
dw_qos.durability.kind = DDS_VOLATILE_DURABILITY_QOS;
dw_qos.protocol.rtps_reliable_writer.heartbeat_period.sec = 0;
dw_qos.protocol.rtps_reliable_writer.heartbeat_period.nanosec = 250000000;

/* Set enabled listener callbacks */
dw_listener.on_publication_matched = HelloWorldPublisher_on_publication_matched;

datawriter =
    DDS_Publisher_create_datawriter(publisher,
                                    topic,
                                    &dw_qos,
                                    &dw_listener,
                                    DDS_PUBLICATION_MATCHED_STATUS);
```

(continues on next page)

(continued from previous page)

```

if (datawriter == NULL)
{
    /* failure */
}

```

The `DataWriterListener` has its callbacks selectively enabled by the DDS status mask. In the example, the mask has set the `on_publication_matched` status, and accordingly the `DataWriterListener` has its `on_publication_matched` assigned to a callback function.

```

void HelloWorldPublisher_on_publication_matched(void *listener_data,
                                              DDS_DataWriter * writer,
                                              const struct DDS_
↳ PublicationMatchedStatus *status)
{
    /* Print on match/unmatch */
    if (status->current_count_change > 0)
    {
        printf("Matched a subscriber\n");
    }
    else
    {
        printf("Unmatched a subscriber\n");
    }
}

```

### 2.8.1 DPSE Discovery: Assert Remote Subscription

A publishing application using `DPSE` discovery must specify the other *DataReaders* that its *DataWriters* are allowed to discover. Like the API for asserting a remote participant, the `DPSE` API for asserting a remote subscription must be called for each remote *DataReader* that a *DataWriter* may discover.

Whereas asserting a remote participant requires only the remote *Participant*'s name, asserting a remote subscription requires more configuration, as all QoS policies of the subscription necessary to determine matching must be known and thus specified.

```

struct DDS_SubscriptionBuiltinTopicData rem_subscription_data =
    DDS_SubscriptionBuiltinTopicData_INITIALIZER;

/* Set Reader's protocol.rtps_object_id */
rem_subscription_data.key.value[DDS_BUILTIN_TOPIC_KEY_OBJECT_ID] = 200;

rem_subscription_data.topic_name = DDS_String_dup("Example HelloWorld");
rem_subscription_data.type_name = DDS_String_dup("HelloWorld");

rem_subscription_data.reliability.kind = DDS_RELIABLE_RELIABILITY_QOS;

retcode = DPSE_RemoteSubscription_assert(participant,
                                         "Participant_2",

```

(continues on next page)

(continued from previous page)

```

        &rem_subscription_data,
        HelloWorld_get_key_kind(HelloWorldTypePlugin_
↪get(),
                                NULL)));
if (retcode != DDS_RETCODE_OK)
{
    /* failure */
}

```

## 2.8.2 Writing Samples

Within the generated type support code are declarations of the type-specific *DataWriter*. For the HelloWorld type, this is the HelloWorldDataWriter.

Writing a HelloWorld sample is done by calling the write API of the HelloWorldDataWriter.

```

HelloWorldDataWriter *hw_datawriter;
DDS_ReturnCode_t retcode;
HelloWorld *sample = NULL;

/* Create and set sample */
sample = HelloWorld_create();
if (sample == NULL)
{
    /* failure */
}
sprintf(sample->msg, "Hello World!");

/* Write sample */
hw_datawriter = HelloWorldDataWriter_narrow(datawriter);

retcode = HelloWorldDataWriter_write(hw_datawriter, sample, &DDS_HANDLE_NIL);
if (retcode != DDS_RETCODE_OK)
{
    /* failure */
}

```

For more information, see the *Sending Data* section in the User's Manual.

## 2.9 Create Subscriber

A subscribing application needs to create a DDS *Subscriber* and then a *DataReader* for each *Topic* to which it wants to subscribe.

In *Connext Micro*, [SubscriberQos](#) in general contains no policies that need to be customized, while [DataReaderQos](#) does contain several customizable policies.

```

DDS_Subscriber *subscriber = NULL;
subscriber = DDS_DomainParticipant_create_subscriber(participant,
                                                    &DDS_SUBSCRIBER_QOS_DEFAULT,
                                                    NULL,
                                                    DDS_STATUS_MASK_NONE);

if (subscriber == NULL)
{
    /* failure */
}

```

For more information, see the *Receiving Data* section in the User's Manual.

## 2.10 Create DataReader

```

DDS_DataReader *datareader = NULL;
struct DDS_DataReaderQos dr_qos = DDS_DataReaderQos_INITIALIZER;
struct DDS_DataReaderListener dr_listener = DDS_DataReaderListener_INITIALIZER;

/* Configure Reader Qos */
dr_qos.protocol.rtps_object_id = 200;
dr_qos.resource_limits.max_instances = 2;
dr_qos.resource_limits.max_samples_per_instance = 2;
dr_qos.resource_limits.max_samples =
    dr_qos.resource_limits.max_samples_per_instance * dr_qos.resource_limits.max_
↳ instances;
dr_qos.reader_resource_limits.max_remote_writers = 10;
dr_qos.reader_resource_limits.max_remote_writers_per_instance = 10;
dr_qos.history.depth = 1;
dr_qos.durability.kind = DDS_VOLATILE_DURABILITY_QOS;
dr_qos.reliability.kind = DDS_RELIABLE_RELIABILITY_QOS;

/* Set listener callbacks */
dr_listener.on_data_available = HelloWorldSubscriber_on_data_available;
dr_listener.on_subscription_matched = HelloWorldSubscriber_on_subscription_matched;

datareader = DDS_Subscriber_create_datareader(subscriber,
                                              DDS_Topic_as_topicdescription(topic),
                                              &dr_qos,
                                              &dr_listener,
                                              DDS_DATA_AVAILABLE_STATUS | DDS_
↳ SUBSCRIPTION_MATCHED_STATUS);
if (datareader == NULL)
{
    /* failure */
}

```

The [DataReaderListener](#) has its callbacks selectively enabled by the DDS status mask. In the example, the mask has set the [DDS\\_SUBSCRIPTION\\_MATCHED\\_STATUS](#) and [DDS\\_DATA\\_AVAILABLE\\_STATUS](#) statuses, and accordingly the [DataReaderListener](#) has its [on\\_subscription\\_matched](#) and [on\\_data\\_available](#) assigned to callback functions.

```

void HelloWorldSubscriber_on_subscription_matched(void *listener_data,
                                                DDS_DataReader * reader,
                                                const struct DDS_
↳SubscriptionMatchedReader *status)
{
    if (status->current_count_change > 0)
    {
        printf("Matched a publisher\n");
    }
    else
    {
        printf("Unmatched a publisher\n");
    }
}

```

```

void HelloWorldSubscriber_on_data_available(void* listener_data,
                                           DDS_DataReader* reader)
{
    HelloWorldDataReader *hw_reader = HelloWorldDataReader_narrow(reader);
    DDS_ReturnCode_t retcode;
    struct DDS_SampleInfo *sample_info = NULL;
    HelloWorld *sample = NULL;

    struct DDS_SampleInfoSeq info_seq =
        DDS_SEQUENCE_INITIALIZER(struct DDS_SampleInfo);
    struct HelloWorldSeq sample_seq =
        DDS_SEQUENCE_INITIALIZER(HelloWorld);

    const DDS_Long TAKE_MAX_SAMPLES = 32;
    DDS_Long i;

    retcode = HelloWorldDataReader_take(hw_reader,
        &sample_seq, &info_seq, TAKE_MAX_SAMPLES,
        DDS_ANY_SAMPLE_STATE, DDS_ANY_VIEW_STATE, DDS_ANY_INSTANCE_STATE);

    if (retcode != DDS_RETCODE_OK)
    {
        printf("failed to take data: %d\n", retcode);
        goto done;
    }

    /* Print each valid sample taken */
    for (i = 0; i < HelloWorldSeq_get_length(&sample_seq); ++i)
    {
        sample_info = DDS_SampleInfoSeq_get_reference(&info_seq, i);

        if (sample_info->valid_data)
        {
            sample = HelloWorldSeq_get_reference(&sample_seq, i);
            printf("\nSample received\n\tmsg: %s\n", sample->msg);
        }
        else

```

(continues on next page)

(continued from previous page)

```

    {
        printf("not valid data\n");
    }
}

HelloWorldDataReader_return_loan(hw_reader, &sample_seq, &info_seq);

done:
    HelloWorldSeq_finalize(&sample_seq);
    DDS_SampleInfoSeq_finalize(&info_seq);
}

```

### 2.10.1 DPSE Discovery: Assert Remote Publication

A subscribing application using [DPSE](#) discovery must specify the other *DataWriters* that its *DataReaders* are allowed to discover. Like the API for asserting a remote participant, the [DPSE](#) API for asserting a remote publication must be called for each remote *DataWriter* that a *DataReader* may discover.

```

struct DDS_PublicationBuiltinTopicData rem_publication_data =
    DDS_PublicationBuiltinTopicData_INITIALIZER;

/* Set Writer's protocol.rtps_object_id */
rem_publication_data.key.value[DDS_BUILTIN_TOPIC_KEY_OBJECT_ID] = 100;

rem_publication_data.topic_name = DDS_String_dup("Example HelloWorld");
rem_publication_data.type_name = DDS_String_dup("HelloWorld");

rem_publication_data.reliability.kind = DDS_RELIABLE_RELIABILITY_QOS;

retcode = DPSE_RemotePublication_assert(participant,
                                        "Participant_1",
                                        &rem_publication_data,
                                        HelloWorld_get_key_kind(HelloWorldTypePlugin_
↪get(),
                                        NULL));
if (retcode != DDS_RETCODE_OK)
{
    /* failure */
}

```

Asserting a remote publication requires configuration of all QoS policies necessary to determine matching.

## 2.10.2 Receiving Samples

Accessing received samples can be done in a few ways:

- **Polling.** Do read or take within a periodic polling loop.
- **Listener.** When a new sample is received, the [DataReaderListener](#)'s `on_data_available` is called. Processing is done in the context of the middleware's receive thread. See the above `HelloWorldSubscriber_on_data_available` callback for example code.
- **Waitset.** Create a waitset, attach it to a status condition with the `data_available` status enabled, and wait for a received sample to trigger the waitset. Processing is done in the context of the user's application thread. (Note: the code snippet below is taken from the shipped `HelloWorld_dpde_waitset` example).

```
DDS_WaitSet *waitset = NULL;
struct DDS_Duration_t wait_timeout = { 10, 0 }; /* 10 seconds */
DDS_StatusCondition *dr_condition = NULL;
struct DDS_ConditionSeq active_conditions =
    DDS_SEQUENCE_INITIALIZER(struct DDS_ConditionSeq);

if (!DDS_ConditionSeq_initialize(&active_conditions))
{
    /* failure */
}

if (!DDS_ConditionSeq_set_maximum(&active_conditions, 1))
{
    /* failure */
}

waitset = DDS_WaitSet_new();
if (waitset == NULL )
{
    /* failure */
}

dr_condition = DDS_Entity_get_statuscondition(DDS_DataReader_as_entity(datareader));

retcode = DDS_StatusCondition_set_enabled_statuses(dr_condition,
                                                    DDS_DATA_AVAILABLE_STATUS);

if (retcode != DDS_RETCODE_OK)
{
    /* failure */
}

retcode = DDS_WaitSet_attach_condition(waitset,
                                        DDS_StatusCondition_as_condition(dr_condition));

if (retcode != DDS_RETCODE_OK)
{
    /* failure */
}
```

(continues on next page)



(continued from previous page)

```

retcode = DDS_WaitSet_wait(waitset, active_conditions, &wait_timeout);

switch (retcode) {
    case DDS_RETCODE_OK:
    {
        /* This WaitSet only has a single condition attached to it
         * so we can implicitly assume the DataReader's status condition
         * to be active (with the enabled DATA_AVAILABLE status) upon
         * successful return of wait().
         * If more than one conditions were attached to the WaitSet,
         * the returned sequence must be examined using the
         * commented out code instead of the following.
         */

        HelloWorldSubscriber_take_data(HelloWorldDataReader_narrow(datareader));

        /*
        DDS_Long active_len = DDS_ConditionSeq_get_length(&active_conditions);
        for (i = active_len - 1; i >= 0; --i)
        {
            DDS_Condition *active_condition =
                *DDS_ConditionSeq_get_reference(&active_conditions, i);

            if (active_condition ==
                DDS_StatusCondition_as_condition(dr_condition))
            {
                total_samples += HelloWorldSubscriber_take_data(
                    HelloWorldDataReader_narrow(datareader));
            }
            else if (active_condition == some_other_condition)
            {
                do_something_else();
            }
        }
        */
        break;
    }
    case DDS_RETCODE_TIMEOUT:
    {
        printf("WaitSet_wait timed out\n");
        break;
    }
    default:
    {
        printf("ERROR in WaitSet_wait: retcode=%d\n", retcode);
        break;
    }
}

```

### 2.10.3 Filtering Samples

*DataReaders* support filtering data in their QoS; see *Content Filtering* for more details. Alternatively, a [DataReaderListener](#) in *Connext Micro* provides callbacks in order to perform application-level filtering on a per-sample basis.

- **on\_before\_sample\_deserialize.** Through this callback, a received sample is presented to the application before it has been deserialized or stored in the *DataReader*'s queue.
- **on\_before\_sample\_commit.** Through this callback, a received sample is presented to the application after it has been deserialized but before it has been stored in the *DataReader*'s queue.

You control the callbacks' `sample_dropped` parameter; upon exiting either callback, the *DataReader* will drop the sample if `sample_dropped` is true. Consequently, dropped samples are not stored in the *DataReader*'s queue and are not available to be read or taken.

Neither callback is associated with a DDS Status. Rather, each is enabled when assigned, to a non-NULL callback.

NOTE: Because it is called after the sample has been deserialized, [on\\_before\\_sample\\_commit](#) provides an additional [sample\\_info](#) parameter, containing some of the usual sample information that would be available when the sample is read or taken.

The HelloWorld\_dpde example's subscriber has this [on\\_before\\_sample\\_commit](#) callback:

```
DDS_Boolean HelloWorldSubscriber_on_before_sample_commit(
    void *listener_data,
    DDS_DataReader *reader,
    const void *const sample,
    const struct DDS_SampleInfo *const sample_info,
    DDS_Boolean *dropped)
{
    HelloWorld *hw_sample = (HelloWorld *)sample;

    /* Drop samples with even-numbered count in msg */
    HelloWorldSubscriber_filter_sample(hw_sample, dropped);

    if (*dropped)
    {
        printf("\nSample filtered, before commit\n\tDROPPED - msg: %s\n",
            hw_sample->msg);
    }

    return DDS_BOOLEAN_TRUE;
}

...

dr_listener.on_before_sample_commit =
    HelloWorldSubscriber_on_before_sample_commit;
```

For more information, see the *Receiving Data* section in the User's Manual.

# Chapter 3

## User's Manual

### 3.1 Data Types

How data is stored or laid out in memory can vary from language to language, compiler to compiler, operating system to operating system, and processor to processor. This combination of language/compiler/operating system/processor is called a *platform*. Any modern middleware must be able to take data from one specific platform (for example, C/gcc.7.3.0/Linux®/PPC) and transparently deliver it to another (for example, C/gcc.7.3.0/Linux/Arm® v8). This process is commonly called *serialization/deserialization*, or *marshalling/demarshalling*.

*Connexrt Micro* data samples sent on the same *Connexrt Micro* topic share a data type. This type defines the fields that exist in the DDS data samples and what their constituent types are. The middleware stores and propagates this meta-information separately from the individual DDS data samples, allowing it to propagate DDS samples efficiently while handling byte ordering and alignment issues for you.

To publish and/or subscribe to data with *Connexrt Micro*, you will carry out the following steps:

1. Select a type to describe your data and use the *RTI Code Generator* to define a type at compile-time using a language-independent description language.

The *RTI Code Generator* accepts input in the following formats:

- **OMG IDL.** This format is a standardized component of the DDS specification. It describes data types with a C++-like syntax. A link to the latest specification can be found here: <https://www.omg.org/spec/IDL>.
- **XML in a DDS-specific format.** This XML format is terser, and therefore easier to read and write by hand, than an XSD file. It offers the general benefits of XML-extensibility and ease of integration, while fully supporting DDS-specific data types and concepts. A link to the latest specification, including a description of the XML format, can be found here: <https://www.omg.org/spec/DDS-XTypes/>.
- **XSD format.** You can describe data types with XML schemas (XSD). A link to the latest specification, including a description of the XSD format, can be found here: <https://www.omg.org/spec/DDS-XTypes/>.

Define a type programmatically at run time.

This method may be appropriate for applications with dynamic data description needs: applications for which types change frequently or cannot be known ahead of time.

2. Register your type with a logical name.
3. Create a *Topic* using the type name you previously registered.

If you've chosen to use a built-in type instead of defining your own, you will use the API constant corresponding to that type's name.

4. Create one or more *DataWriters* to publish your data and one or more *DataReaders* to subscribe to it.

The concrete types of these objects depend on the concrete data type you've selected, in order to provide you with a measure of type safety.

Whether publishing or subscribing to data, you will need to know how to create and delete DDS data samples and how to get and set their fields. These tasks are described in [the section on Working with DDS Data Samples](#) in the [RTI Connex DDS Core Libraries User's Manual](#) (available [here](#) if you have Internet access).

### 3.1.1 Introduction to the Type System

A *user data type* is any custom type that your application defines for use with *RTI Connex Micro*. It may be a structure, a union, a value type, an enumeration, or a typedef (or language equivalents).

Your application can have any number of user data types. They can be composed of any of the primitive data types listed below or of other user data types.

Only structures, unions, and value types may be read and written directly by *Connex Micro*; enums, typedefs, and primitive types must be contained within a structure, union, or value type. In order for a *DataReader* and *DataWriter* to communicate with each other, the data types associated with their respective Topic definitions must be identical.

- octet, char, wchar
- short, unsigned short
- long, unsigned long
- long long, unsigned long long
- float
- double, long double
- boolean
- enum (with or without explicit values)
- bounded string and wstring

The following type-building constructs are also supported:

- module (also called a package or namespace)

- pointer
- array of primitive or user type elements
- bounded sequence of elements—a sequence is a variable-length ordered collection, such as a vector or list
- typedef
- union
- struct
- value type, a complex type that supports inheritance and other object-oriented features

To use a data type with *Connext Micro*, you must define that type in a way the middleware understands and then register the type with the middleware. These steps allow *Connext Micro* to serialize, deserialize, and otherwise operate on specific types. They will be described in detail in the following sections.

## Sequences

A sequence contains an ordered collection of elements that are all of the same type. The operations supported in the sequence are documented in the [C API Reference](#) and [C++ API Reference](#) HTML documentation.

Elements in a sequence are accessed with their index, just like elements in an array. Indices start at zero in all APIs. Unlike arrays, however, sequences can grow in size. A sequence has two sizes associated with it: a physical size (the “maximum”) and a logical size (the “length”). The physical size indicates how many elements are currently allocated by the sequence to hold; the logical size indicates how many valid elements the sequence actually holds. The length can vary from zero up to the maximum. Elements cannot be accessed at indices beyond the current length.

A sequence must be declared as bounded. A sequence’s “bound” is the maximum number of elements that the sequence can contain at any one time. A finite bound is very important because it allows *RTI Connext Micro* to preallocate buffers to hold serialized and deserialized samples of your types; these buffers are used when communicating with other nodes in your distributed system.

By default, any unbounded sequences found in an IDL file will be given a default bound of 100 elements. This default value can be overwritten using *RTI Code Generator*’s **-sequenceSize** command-line argument (see [the Command-Line Arguments chapter](#) in the *RTI Code Generator User’s Manual*, available [here](#) if you have Internet access).

## Strings and Wide Strings

*Connext Micro* supports both strings consisting of single-byte characters (the IDL string type) and strings consisting of wide characters (IDL wstring). The wide characters supported by *Connext Micro* are large enough to store two-byte Unicode/UTF16 characters.

Like sequences, strings must be bounded. A string’s “bound” is its maximum length (not counting the trailing NULL character in C and C++).

In C and Traditional C++, strings are mapped to `char*`.

By default, any unbounded string found in an IDL file will be given a default bound of 255 elements. This default value can be overwritten using *RTI Code Generator*’s `-stringSize` command-line argument (see [the Command-Line Arguments chapter](#) in the *RTI Code Generator User’s Manual*, available [here](#) if you have Internet access).

## IDL String Encoding

The “Extensible and Dynamic Topic Types for DDS specification” (<https://www.omg.org/spec/DDS-XTypes/>) standardizes the default encoding for strings to UTF-8. This encoding shall be used as the wire format. Language bindings may use the representation that is most natural in that particular language. If this representation is different from UTF-8, the language binding shall manage the transformation to/from the UTF-8 wire representation.

As an extension, *Connext Micro* offers ISO\_8859\_1 as an alternative string wire encoding.

This section describes the encoding for IDL strings across different languages in *Connext Micro* and how to configure that encoding.

- C, Traditional C++

IDL strings are mapped to a NULL-terminated array of `DDS_Char` (`char*`). Users are responsible for using the right character encoding (UTF-8 or ISO\_8859\_1) when populating the string values. This applies to all generated code, `DynamicData`, and Built-in data types. The middleware does not transform from the language binding encoding to the wire encoding.

## IDL Wide Strings Encoding

The “Extensible and Dynamic Topic Types for DDS specification” (<https://www.omg.org/spec/DDS-XTypes/>) standardizes the default encoding for wide strings to UTF-16. This encoding shall be used as the wire format.

When the data representation is Extended CDR version 1, wide-string characters have a size of 4 bytes on the wire with UTF-16 encoding. When the data representation is Extended CDR version 2, wide-string characters have a size of 2 bytes on the wire with UTF-16 encoding.

Language bindings may use the representation that is most natural in that particular language. If this representation is different from UTF-16, the language binding shall manage the transformation to/from the UTF-16 wire representation.

- C, Traditional C++

IDL wide strings are mapped to a NULL-terminated array of `DDS_Wchar` (`DDS_Wchar*`). `DDS_Wchar` is an unsigned 2-byte integer. Users are responsible for using the right character encoding (UTF-16) when populating the wide-string values. This applies to all generated code, `DynamicData`, and Built-in data types. *Connext Micro* does not transform from the language binding encoding to the wire encoding.

## Sending Type Information on the Network

*Connext Professional* can send type information on the network using a concept called type objects. A type object is a description of a type suitable to network transmission, and is commonly used by tools to visualize data from any application.

However, please note that *Connext Micro* does not support sending type information on the network. Instead, tools can load type information from XML files generated from IDL using *rtiddsgen*. Please refer [here](#) for more information.

## Extensible Types (XTypes) 1.2 Compatibility

*Connext Micro* supports the “Extensible and Dynamic Topic Types for DDS” (DDS-XTypes) specification from the Object Management Group (OMG), version 1.2 (<https://www.omg.org/spec/DDS-XTypes/1.2>) with the following limitations:

- Extended Common Data Representation (CDR) encoding version 1 (XCDR) and Extended CDR encoding version 2 (XCDR2) are supported by default.
- If *RTI Code Generator* (*rtiddsgen*) is used with the option **-interpreted 0**, support for XTypes is disabled and only plain CDR is supported (CDRv1 final types).
- *Connext Micro* does not send type information.
- *Connext Micro* does not perform type-compatibility checking based on the type information, only the type-name. This means that advanced XTypes 1.2 features cannot be supported, such as:
  - Type equivalence
  - String-length matching and truncation
  - Sequence-length matching and truncation

### 3.1.2 Creating User Data Types with IDL

You can create user data types in a text file using IDL (Interface Description Language). IDL is programming-language independent, so the same file can be used to generate code in C and Traditional C++. *RTI Code Generator* parses the IDL file and automatically generates all the necessary routines and wrapper functions to bind the types for use by *Connext Micro* at run time. You will end up with a set of required routines and structures that your application and *Connext Micro* will use to manipulate the data.

Please refer to [the section on Creating User Data Types with IDL](#) in the [RTI Connext DDS Core Libraries User's Manual](#) for more information (available [here](#) if you have Internet access).

Note: Not all features in *RTI Code Generator* are supported when generating code for *Connext Micro*, see *Unsupported Features of rtiddsgen with Connext Micro*.

### 3.1.3 Working with DDS Data Samples

You should now understand how to define and work with data types. Now that you have chosen one or more data types to work with, this section will help you understand how to create and manipulate objects of those types.

#### In C:

You create and delete your own objects from factories, just as you create *Connext Micro* objects from factories. In the case of user data types, the factory is a singleton object called the type support. Objects allocated from these factories are deeply allocated and fully initialized.

```
/* In the generated header file: */
struct MyData {
    char* myString;
};
/* In your code: */
MyData* sample = MyDataTypeSupport_create_data();
char* str = sample->myString; /*empty, non-NULL string*/
/* ... */
MyDataTypeSupport_delete_data(sample);
```

#### In Traditional C++:

Without the **-constructor option**, you create and delete objects using the TypeSupport factories.

```
MyData* sample = MyDataTypeSupport::create_data();
char* str = sample->myString; // empty, non-NULL string
// ...
MyDataTypeSupport::delete_data(sample);
```

Please refer to [the section on Working with DDS Data Samples](#) in the [RTI Connext DDS Core Libraries User's Manual](#) for more information (available [here](#) if you have Internet access).



## 3.2 DDS Entities

The main classes extend an abstract base class called a *DDS Entity*. Every *DDS Entity* has a set of associated events known as statuses and a set of associated Quality of Service Policies (QoS Policies). In addition, a *Listener* may be registered with the *Entity* to be called when status changes occur. *DDS Entities* may also have attached *DDS Conditions*, which provide a way to wait for status changes. *Figure 4.1: Overview of DDS Entities* presents an overview in a UML diagram.

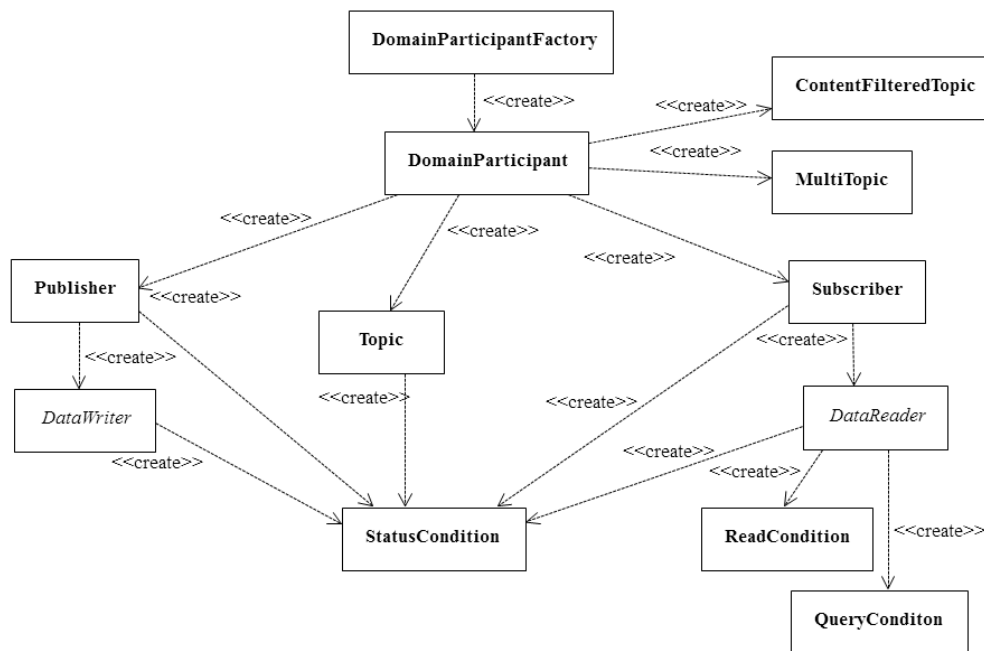


Figure 3.1: Overview of DDS Entities

Please note that *RTI Connext Micro* does not support the following:

- **MultiTopic**
- **ContentFilteredTopic**
- **ReadCondition**
- **QueryConditions**

For a general description of *DDS Entities* and their operations, please refer to [the DDS Entities chapter](#) in the [RTI Connext DDS Core Libraries User's Manual](#) (available [here](#) if you have Internet access). Note that *RTI Connext Micro* does not support all APIs and QoS Policies; please refer to the [C API Reference](#) and [C++ API Reference](#) documentation for more information.

## 3.3 Sending Data

This section discusses how to create, configure, and use *Publishers* and *DataWriters* to send data. It describes how these *Entities* interact, as well as the types of operations that are available for them.

The goal of this section is to help you become familiar with the *Entities* you need for sending data. For up-to-date details such as formal parameters and return codes on any mentioned operations, please see the [C API Reference](#) and [C++ API Reference](#) documentation.

### 3.3.1 Preview: Steps to Sending Data

To send DDS samples of a data instance:

1. Create and configure the required *Entities*:
  - a. Create a *DomainParticipant*.
  - b. Register user data types with the *DomainParticipant*. For example, the **'FooDataType'**.
  - c. Use the *DomainParticipant* to create a *Topic* with the registered data type.
  - d. Use the *DomainParticipant* to create a *Publisher*.
  - e. Use the *Publisher* or *DomainParticipant* to create a *DataWriter* for the *Topic*.
  - f. Use a type-safe method to cast the generic *DataWriter* created by the *Publisher* to a type-specific *DataWriter*. For example, **'FooDataWriter'**. Optionally, register data instances with the *DataWriter*. If the *Topic*'s user data type contain key fields, then registering a data instance (data with a specific key value) will improve performance when repeatedly sending data with the same key. You may register many different data instances; each registration will return an instance handle corresponding to the specific key value. For non-keyed data types, instance registration has no effect.
2. Every time there is changed data to be published:
  - a. Store the data in a variable of the correct data type (for instance, variable **'Foo'** of the type **'FooDataType'**).
  - b. Call the **FooDataWriter's write()** operation, passing it a reference to the variable **'Foo'**.
    - For non-keyed data types or for non-registered instances, also pass in **DDS\_HANDLE\_NIL**.
    - For keyed data types, pass in the instance handle corresponding to the instance stored in **'Foo'**, if you have registered the instance previously. This means that the data stored in **'Foo'** has the same key value that was used to create instance handle.
  - c. The **write()** function will take a snapshot of the contents of **'Foo'** and store it in *Connext DDS* internal buffers from where the DDS data sample is sent under the criteria set by the *Publisher's* and *DataWriter's* QoS Policies. If there are matched *DataReaders*, then

the DDS data sample will have been passed to the physical transport plug-in/device driver by the time that **write()** returns.

### 3.3.2 Publishers

An application that intends to publish information needs the following *Entities*: *DomainParticipant*, *Topic*, *Publisher*, and *DataWriter*. All *Entities* have a corresponding specialized *Listener* and a set of QoS Policies. A *Listener* is how *Connext DDS* notifies your application of status changes relevant to the *Entity*. The QoS Policies allow your application to configure the behavior and resources of the *Entity*.

- A *DomainParticipant* defines the DDS domain in which the information will be made available.
- A *Topic* defines the name under which the data will be published, as well as the type (format) of the data itself.
- An application writes data using a *DataWriter*. The *DataWriter* is bound at creation time to a *Topic*, thus specifying the name under which the *DataWriter* will publish the data and the type associated with the data. The application uses the *DataWriter*'s **write()** operation to indicate that a new value of the data is available for dissemination.
- A *Publisher* manages the activities of several *DataWriters*. The *Publisher* determines when the data is actually sent to other applications. Depending on the settings of various QoS Policies of the *Publisher* and *DataWriter*, data may be buffered to be sent with the data of other *DataWriters* or not sent at all. By default, the data is sent as soon as the *DataWriter*'s **write()** function is called.

You may have multiple *Publishers*, each managing a different set of *DataWriters*, or you may choose to use one *Publisher* for all your *DataWriters*.

### 3.3.3 DataWriters

To create a *DataWriter*, you need a *DomainParticipant*, *Publisher*, and a *Topic*.

You need a *DataWriter* for each *Topic* that you want to publish. For more details on all operations, see the [C API Reference](#) and [C++ API Reference](#) documentation.

For more details on creating, deleting, and setting up *DataWriters*, see replace:: [the DataWriters section](#) in the [RTI Connext DDS Core Libraries User's Manual](#) (available [here](#) if you have Internet access).

### 3.3.4 Publisher/Subscriber QosPolicies

Please refer to the [C API Reference](#) and [C++ API Reference](#) for details on supported QosPolicies.

### 3.3.5 DataWriter QosPolicies

Please refer to the [C API Reference](#) and [C++ API Reference](#) for details on supported QosPolicies.

## 3.4 Receiving Data

This section discusses how to create, configure, and use *Subscribers* and *DataReaders* to receive data. It describes how these objects interact, as well as the types of operations that are available for them.

The goal of this section is to help you become familiar with the *Entities* you need for receiving data. For up-to-date details such as formal parameters and return codes on any mentioned operations, please see the [C API Reference](#) and [C++ API Reference](#) documentation.

**Warning:** *Connext Micro DataReaders* cannot match with or receive data from *Connext DataWriters* that are configured to send compressed data. See the *Interoperability* section for more information.

### 3.4.1 Preview: Steps to Receiving Data

There are three ways to receive data:

- Your application can explicitly check for new data by calling a *DataReader's* **read()** or **take()** operation. This method is also known as *polling for data*.
- Your application can be notified asynchronously whenever new DDS data samples arrive—this is done with a *Listener* on either the *Subscriber* or the *DataReader*. *RTI Connext Micro* will invoke the *Listener's* callback routine when there is new data. Within the callback routine, user code can access the data by calling **read()** or **take()** on the *DataReader*. This method is the way for your application to receive data with the least amount of latency.
- Your application can wait for new data by using *Conditions* and a *WaitSet*, then calling **wait()**. *Connext Micro* will block your application's thread until the criteria (such as the arrival of DDS samples, or a specific status) set in the *Condition* becomes true. Then your application resumes and can access the data with **read()** or **take()**.

The *DataReader's* **read()** operation gives your application a copy of the data and leaves the data in the *DataReader's* receive queue. The *DataReader's* **take()** operation removes data from the receive queue before giving it to your application.

**To prepare to receive data, create and configure the required Entities:**

1. Create a *DomainParticipant*.

2. Register user data types with the *DomainParticipant*. For example, the **'FooDataType'**.
3. Use the *DomainParticipant* to create a *Topic* with the registered data type.
4. Use the *DomainParticipant* to create a *Subscriber*.
5. Use the *Subscriber* or *DomainParticipant* to create a *DataReader* for the *Topic*.
6. Use a type-safe method to cast the generic *DataReader* created by the *Subscriber* to a type-specific *DataReader*. For example, **'FooDataReader'**.

Then use one of the following mechanisms to receive data.

- To receive DDS data samples by polling for new data:
  - Using a **FooDataReader**, use the **read()** or **take()** operations to access the DDS data samples that have been received and stored for the *DataReader*. These operations can be invoked at any time, even if the receive queue is empty.
- To receive DDS data samples asynchronously:
  - Install a *Listener* on the *DataReader* or *Subscriber* that will be called back by an internal *Connext Micro* thread when new DDS data samples arrive for the *DataReader*.
- 1. Create a *DDSDataReaderListener* for the *FooDataReader* or a *DDSSubscriberListener* for *Subscriber*. In C++ you must derive your own *Listener* class from those base classes. In C, you must create the individual functions and store them in a structure.

If you created a *DDSDataReaderListener* with the **on\_data\_available()** callback enabled: **on\_data\_available()** will be called when new data arrives for that **DataReader**.

If you created a *DDSSubscriberListener* with the **on\_data\_on\_readers()** callback enabled: **on\_data\_on\_readers()** will be called when data arrives for any *DataReader* created by the *Subscriber*.

2. Install the *Listener* on either the *FooDataReader* or *Subscriber*.

For the *DataReader*, the *Listener* should be installed to handle changes in the **DATA\_AVAILABLE** status.

For the *Subscriber*, the *Listener* should be installed to handle changes in the **DATA\_ON\_READERS** status.

3. Only 1 *Listener* will be called back when new data arrives for a *DataReader*.

*Connext Micro* will call the *Subscriber's Listener* if it is installed. Otherwise, the *DataReader's Listener* is called if it is installed. That is, the **on\_data\_on\_readers()** operation takes precedence over the **on\_data\_available()** operation.

If neither *Listeners* are installed or neither *Listeners* are enabled to handle their respective statuses, then *Connext Micro* will not call any user functions when new data arrives for the *DataReader*.

4. In the **on\_data\_available()** method of the *DDSDataReaderListener*, invoke **read()** or **take()** on the *FooDataReader* to access the data.

If the `on_data_on_readers()` method of the *DDSSubscriberListener* is called, the code can invoke `read()` or `take()` directly on the *Subscriber's DataReaders* that have received new data. Alternatively, the code can invoke the *Subscriber's* `notify_datareaders()` operation. This will in turn call the `on_data_available()` methods of the *DataReaderListeners* (if installed and enabled) for each of the *DataReaders* that have received new DDS data samples.

#### To wait (block) until DDS data samples arrive:

1. Use the *DataReader* to create a *StatusCondition* that describes the DDS samples for which you want to wait. For example, you can specify that you want to wait for never-before-seen DDS samples from *DataReaders* that are still considered to be 'alive.'
2. Create a *WaitSet*.
3. Attach the *StatusCondition* to the *WaitSet*.
4. Call the *WaitSet's* `wait()` operation, specifying how long you are willing to wait for the desired DDS samples. When `wait()` returns, it will indicate that it timed out, or that the attached Condition become true (and therefore the desired DDS samples are available).
5. Using a **FooDataReader**, use the `read()` or `take()` operations to access the DDS data samples that have been received and stored for the *DataReader*.

### 3.4.2 Subscribers

An application that intends to subscribe to information needs the following *Entities*: *DomainParticipant*, *Topic*, *Subscriber*, and *DataReader*. All *Entities* have a corresponding specialized *Listener* and a set of QosPolicies. The *Listener* is how *RTI Connext Micro* notifies your application of status changes relevant to the *Entity*. The QosPolicies allow your application to configure the behavior and resources of the *Entity*.

- The *DomainParticipant* defines the DDS domain on which the information will be available.
- The *Topic* defines the name of the data to be subscribed, as well as the type (format) of the data itself.
- The *DataReader* is the *Entity* used by the application to subscribe to updated values of the data. The *DataReader* is bound at creation time to a *Topic*, thus specifying the named and typed data stream to which it is subscribed. The application uses the *DataWriter's* `read()` or `take()` operation to access DDS data samples received for the *Topic*.
- The *Subscriber* manages the activities of several *DataReader* entities. The application receives data using a *DataReader* that belongs to a *Subscriber*. However, the *Subscriber* will determine when the data received from applications is actually available for access through the *DataReader*. Depending on the settings of various QosPolicies of the *Subscriber* and *DataReader*, data may be buffered until DDS data samples for associated *DataReaders* are also received. By default, the data is available to the application as soon as it is received.

For more information on creating and deleting *Subscribers*, as well as setting QosPolicies, see [the Subscribers section](#) in the [RTI Connext DDS Core Libraries User's Manual](#) (available [here](#) if you have Internet access).

### 3.4.3 DataReaders

To create a *DataReader*, you need a *DomainParticipant*, a *Topic*, and a *Subscriber*. You need at least one *DataReader* for each *Topic* whose DDS data samples you want to receive.

For more details on all operations, see the [C API Reference](#) and [C++ API Reference](#) HTML documentation.

### 3.4.4 Using DataReaders to Access Data (Read & Take)

For user applications to access the data received for a *DataReader*, they must use the type-specific derived class or set of functions in the [C API Reference](#). Thus for a user data type ‘**Foo**’, you must use methods of the **FooDataReader** class. The type-specific class or functions are automatically generated if you use *RTI Code Generator*.

### 3.4.5 Subscriber QosPolicies

Please refer to the [C API Reference](#) and [C++ API Reference](#) for details on supported QosPolicies.

### 3.4.6 DataReader QosPolicies

Please refer to the [C API Reference](#) and [C++ API Reference](#) for details on supported QosPolicies.

## 3.5 DDS Domains

This section discusses how to use *DomainParticipants*. It describes the types of operations that are available for them and their QosPolicies.

The goal of this section is to help you become familiar with the objects you need for setting up your *RTI Connext Micro* application. For specific details on any mentioned operations, see the [C API Reference](#) and [C++ API Reference](#) documentation.

### 3.5.1 Fundamentals of DDS Domains and DomainParticipants

*DomainParticipants* are the focal point for creating, destroying, and managing other *RTI Connext Micro* objects. A DDS *domain* is a logical network of applications: only applications that belong to the same DDS *domain* may communicate using *Connext Micro*. A DDS *domain* is identified by a unique integer value known as a domain ID. An application participates in a DDS domain by creating a *DomainParticipant* for that domain ID.

As seen in *Figure 4.2: Relationship between Applications and DDS Domains*, a single application can participate in multiple DDS domains by creating multiple *DomainParticipants* with different domain IDs. *DomainParticipants* in the same DDS domain form a logical network; they are isolated from *DomainParticipants* of other DDS domains, even those running on the same set of physical computers sharing the same physical network. *DomainParticipants* in different DDS domains will

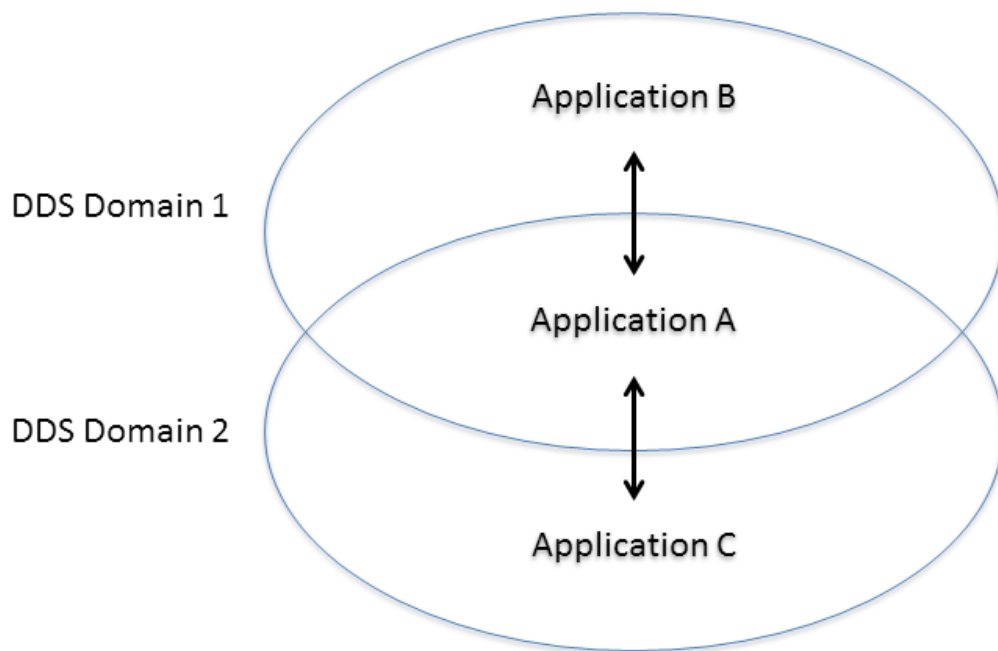


Figure 3.2: Relationship between Applications and DDS Domains

Applications can belong to multiple DDS domains—*A* belongs to DDS domains 1 and 2. Applications in the same DDS domain can communicate with each other, such as *A* and *B*, or *A* and *C*. Applications in different DDS domains, such as *B* and *C*, are not even aware of each other and will not exchange messages.



never exchange messages with each other. Thus, a DDS domain establishes a “virtual network” linking all *DomainParticipants* that share the same domain ID.

An application that wants to participate in a certain DDS domain will need to create a *DomainParticipant*. As seen in Figure 4.3: *DDS Domain Module*, a *DomainParticipant* object is a container for all other *Entities* that belong to the same DDS domain. It acts as factory for the *Publisher*, *Subscriber*, and *Topic* entities. (As seen in *Sending Data* and *Receiving Data*, in turn, *Publishers* are factories for *DataWriters* and *Subscribers* are factories for *DataReaders*.) *DomainParticipants* cannot contain other *DomainParticipants*.

Like all *Entities*, *DomainParticipants* have QosPolicies and *Listeners*. The *DomainParticipant* entity also allows you to set ‘default’ values for the QosPolicies for all the entities created from it or from the entities that it creates (*Publishers*, *Subscribers*, *Topics*, *DataWriters*, and *DataReaders*).

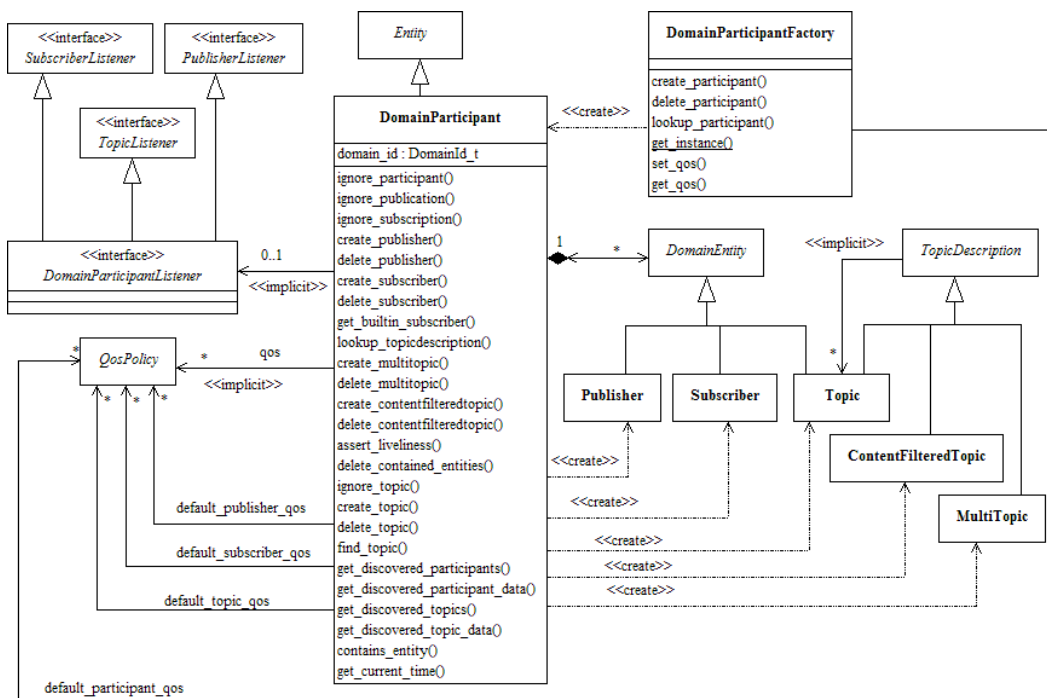


Figure 3.3: DDS Domain Module

Note: MultiTopics are not supported.

### 3.5.2 Discovery Announcements

Each *DomainParticipant* announces information about itself, such as which locators other *DomainParticipants* must use to communicate with it. A locator is an address that consists of an address kind, a port number, and an address. Four locator types are defined:

- A **unicast meta-traffic locator**. This locator type is used to identify where unicast discovery messages shall be sent. A maximum of four locators of this type can be specified.
- A **multicast meta-traffic locator**. This locator type is used to identify where multicast discovery messages shall be sent. A maximum of four locators of this type can be specified.

- A **unicast user-traffic locator**. This locator type is used to identify where unicast user-traffic messages shall be sent. A maximum of four locators of this type can be specified.
- A **multicast user-traffic locator**. This locator type is used to identify where multicast user-traffic messages shall be sent. A maximum of four locators of this type can be specified.

It is important to note that a maximum of *four* locators of *each* kind can be sent in a *DomainParticipant* discovery message.

The locators in a *DomainParticipant*'s discovery announcement is used for two purposes:

- It informs other *DomainParticipants* where to send their discovery announcements to this *DomainParticipants*.
- It informs the *DataReaders* and *DataWriters* in other *DomainParticipants* where to send data to the *DataReaders* and *DataWriters* in this *DomainParticipant* unless a *DataReader* or *DataWriter* specifies its own locators.

If a *DataReader* or *DataWriter* specifies their own locators, only user-traffic locators can be specified, then the exact same rules apply as for the *DomainParticipant*.

This document uses *address* and *locator* interchangeably. An address corresponds to the port and address part of a locator. The same address may exist as different kinds, in which case they are unique.

For more details about the discovery process, see the *Discovery* section.

## 3.6 Transports

### 3.6.1 Introduction

*RTI Connext Micro* has a pluggable-transports architecture. The core of *Connext Micro* is transport agnostic—it does not make any assumptions about the actual transports used to send and receive messages. Instead, *Connext Micro* uses an abstract “transport API” to interact with the transport plugins that implement that API. A transport plugin implements the abstract transport API, and performs the actual work of sending and receiving messages over a physical transport.

In *Connext Micro* a Network Input/Output (NETIO) interface is a software layer that may send and/or receive data from a higher and/or lower level locally, as well as communicate with a peer. A transport is a NETIO interface that is at the lowest level of the protocol stack. For example, the UDP NETIO interface is a transport.

A transport can send and receive on addresses as defined by the concrete transport. For example, the *Connext Micro* UDP transport can listen to and send to UDPv4 ports and addresses. In order to establish communication between two transports, the addresses that the transport can listen to must be determined and announced to other *DomainParticipants* that want to communicate with it. This document describes how the addresses are reserved and how these addresses are used by the DDS layer in *Connext Micro*.

While the NETIO interface is not limited to DDS, the rest of this document is written in the context of how *Connext Micro* uses the NETIO interfaces as part of the DDS implementation.

### 3.6.2 Transport Registration

*RTI Connext Micro* supports different transports and transports must be registered with *RTI Connext Micro* before they can be used. A transport must be given a name when it is registered and this name is later used when configuring discovery and user-traffic. A transport name cannot exceed 7 UTF-8 characters.

The following example registers the UDP transport with *RTI Connext Micro* and makes it available to all DDS applications within the same memory space. Please note that each DDS applications creates its *own* instance of a transport. Resources are *not* shared between instances of a transport unless stated.

For example, to register two UDP transports with the names `myudp1` and `myudp2`, the following code is required:

```
DDS_DomainParticipantFactory *factory;
RT_Registry_T *registry;
struct UDP_InterfaceFactoryProperty udp_property;

factory = DDS_DomainParticipantFactory_get_instance();
registry = DDS_DomainParticipantFactory_get_registry(factory);

/* Set UDP properties */
if (!RT_Registry_register(registry,"myudp1",
                        UDP_InterfaceFactory_get_interface(),
                        &udp_property._parent._parent,NULL))
{
    return error;
}

/* Set UDP properties */
if (!RT_Registry_register(registry,"myudp2",
                        UDP_InterfaceFactory_get_interface(),
                        &udp_property._parent._parent,NULL))
{
    return error;
}
```

Before a DomainParticipant can make use of a registered transport, it must enable it for use within the DomainParticipant. This is done by setting the [TransportQoS](#). For example, to enable only `myudp1`, the following code is required (error checking is not shown for clarity):

```
DDS_StringSeq_set_maximum(&dp_qos.transports.enabled_transports,1);
DDS_StringSeq_set_length(&dp_qos.transports.enabled_transports,1);
*DDS_StringSeq_get_reference(&dp_qos.transports.enabled_transports,0) =
    REDA_String_dup("myudp1");
```

To enable both transports:

```
DDS_StringSeq_set_maximum(&dp_qos.transports.enabled_transports,2);
DDS_StringSeq_set_length(&dp_qos.transports.enabled_transports,2);
*DDS_StringSeq_get_reference(&dp_qos.transports.enabled_transports,0) =
```

(continues on next page)

(continued from previous page)

```

                                REDA_String_dup("myudp1");
*DDS_StringSeq_get_reference(&dp_qos.transports.enabled_transports,0) =
                                REDA_String_dup("myudp2");

```

Before enabled transports may be used for communication in *Connext Micro*, they must be registered and added to the [DiscoveryQos](#) and [UserTrafficQos](#) policies. Please see the section on *Discovery* for details.

### 3.6.3 Transport Addresses

Address reservation is the process to determine which locators should be used in the discovery announcement. Which transports and addresses to be used is determined as described in *Discovery*.

When a *DomainParticipant* is created, it calculates a port number and tries to reserve this port on all addresses available in *all* the transports based on the registration properties. If the port cannot be reserved on all transports, then it release the port on *all* transports and tries again. If no free port can be found the process fails and the *DomainParticipant* cannot be created.

The number of locators which can be announced is limited to *only* the first *four* for each type across *all* transports available for each policy. If more than four are available of any kind, these are *ignored*. This is by design, although it may be changed in the future. The order in which the locators are read is also not known, thus the four locators which will be used are not deterministic.

To ensure that *all* the desired addresses and *only* the desired address are used in a transport, follow these rules:

- Make sure that no more than four unicast addresses and four multicast addresses can be returned across *all* transports for discovery traffic.
- Make sure that no more than four unicast addresses and four multicast addresses can be returned across *all* transports for user traffic.
- Make sure that no more than four unicast addresses and four multicast addresses can be returned across *all* transports for user-traffic, for *DataReader* and *DataWriter* specific locators, and that they do *not* duplicate any of the *DomainParticipant*'s locators.

### 3.6.4 Transport Port Number

The port number of a locator is not directly configurable. Rather, it is configured indirectly by the [DDS\\_WireProtocolQosPolicy](#) ([rtps\\_well\\_known\\_ports](#)) of the *DomainParticipant*'s QoS, where a well-known, interoperable RTPS port number is assigned.

### 3.6.5 RTPS

The RTPS transport encapsulates user-data in RTPS messages and parses received RTPS messages for user-data. This chapter describes how to configure RTPS.

#### Registration of RTPS

RTPS is automatically registered when a *DDS\_DomainParticipantFactory* is initialized with *DDS\_DomainParticipantFactory\_get\_instance()*. In order to change the RTPS configuration, it is necessary to first unregister it from the participant factory, set the properties, and then register RTPS with the new properties. This process is identical to other plugins in *Connext Micro*, such as the UDP transport and discovery plugins.

The following code shows the steps:

```
int main(int argc, char *argv)
{
    struct RTPS_InterfaceFactoryProperty *rtps_property = NULL;
    DDS_DomainParticipantFactory *factory = NULL;
    RT_Registry_T *registry = NULL;
    struct RTPS_InterfaceFactoryProperty *rtps_property = NULL;

    /* get the Domain Participant factory and registry*/
    factory = DDS_DomainParticipantFactory_get_instance();

    registry = DDS_DomainParticipantFactory_get_registry
               (DDS_DomainParticipantFactory_get_instance());

    /* unregister the RTPS transport */
    if (!RT_Registry_unregister(registry, NETIO_DEFAULT_RTPS_NAME,
                               NULL, NULL))
    {
        printf("failed to unregister rtps\n");
        return 0;
    }

    rtps_property = (struct RTPS_InterfaceFactoryProperty *)
        malloc(sizeof(struct RTPS_InterfaceFactoryProperty));

    if (rtps_property == NULL)
    {
        printf("failed to allocate rtps properties\n");
        return 0;
    }

    /* Set the new properties and register RTPS again */

    if (!RT_Registry_register(registry, NETIO_DEFAULT_RTPS_NAME,
                             RTPS_InterfaceFactory_get_interface(),
                             (struct RT_ComponentFactoryProperty*)rtps_property,
```

(continues on next page)

(continued from previous page)

```

        NULL))
    {
        printf("failed to register rtps\n");
        return 0;
    }

    DDS_DomainParticipantFactory_create_participant(
        factory, domain_id, &dp_qos, NULL, DDS_STATUS_MASK_NONE);
}

```

Please note that the RTPS properties *must* be valid for the *entire* life-cycle of the participant factory because RTPS *does not* make an internal copy. This saves memory when properties are stored in preallocated memory (for example in ROM).

### Overriding the Builtin RTPS Checksum Functions

Some applications may require specialized functions to guarantee message integrity or may have special hardware that supports faster checksum calculations. *Connext Micro* provides a way for users to override the builtin checksum functions. Note that if a different checksum is calculated it may prevent interoperability with other DDS implementations.

#### Checksum function definition

A checksum function must define a structure of the following type:

```

typedef struct RTPS_ChecksumClass
{
    RTPS_ChecksumClassId_T class_id;
    void *context;
    RTPS_CalculateChecksum_T calculate_checksum;
} RTPS_ChecksumClass_T;

```

The type has three members:

1. class\_id - The class ID must be:
  - RTPS\_CHECKSUM\_CLASSID\_BUILTIN32 for the 32-bit checksum.
  - RTPS\_CHECKSUM\_CLASSID\_BUILTIN64 for the 64-bit checksum.
  - RTPS\_CHECKSUM\_CLASSID\_BUILTIN128 for the 128-bit checksum.
2. context - An opaque object for you to provide context for this function. This context will be passed to the *calculate\_checksum* every time it is called.
3. checksum\_calculate - The function pointer to the checksum function. The function is defined as

```
typedef RTI_BOOL
(*RTPS_ChecksumCalculate_T)(void *context,
                           const struct REDA_Buffer *buf,
                           RTI_UINT32 buf_length,
                           RTPS_Checksum_T *checksum);
```

- context: *Connex Micro* will pass in the context as defined in the class.
- buf: An array of REDA\_Buffer. Each REDA\_Buffer includes a pointer and size of the buffer.
- buf\_length: The size of the array.

RTPS\_Checksum\_T checksum: This is the out parameter of this function. It is a union defined as follows:

```
typedef union RTPS_Checksum
{
    RTI_UINT32 checksum32;
    RTI_UINT64 checksum64;
    RTI_UINT8 checksum128[16];
} RTPS_Checksum_T;
```

Please note the following *important* information regarding the output values:

1. The number returned in checksum32 is assumed to be in *host order* endianness.
2. The number returned in checksum64 is assumed to be in *host order* endianness.
3. checksum128 is treated as an octet array.

## Example

Below is an example implementation of a custom CRC-32 function using the Intel intrinsic functions. It shows the QoS that needs to be set, as well as how to override the builtin checksum function.

```
RTI_BOOL
CrcClassTest_custom_crc32_other(void *context,
                                const struct REDA_Buffer *buf,
                                unsigned int buf_length,
                                union RTPS_CrcChecksum *checksum)
{
    RTI_UINT32 crc = 0;
    unsigned char *data = (unsigned char *) buf[0].pointer;
    RTI_UINT32 length = buf[0].length;
    int k;

    UNUSED_ARG(k);
    UNUSED_ARG(context);
    UNUSED_ARG(buf_length);

    for (k = 0; k < length; k++)
```

(continues on next page)

(continued from previous page)

```

{
    crc = _mm_crc32_u8(crc, data[k]);
}

checksum->checksum32 = crc;

return RTI_TRUE;
}

int main(int argc, char *argv)
{
    struct DDS_DomainParticipantQos dp_qos =
        DDS_DomainParticipantQos_INITIALIZER;
    struct RTPS_InterfaceFactoryProperty *rtps_property = NULL;
    DDS_DomainParticipantFactory *factory = NULL;
    RT_Registry_T *registry = NULL;
    struct RTPS_InterfaceFactoryProperty *rtps_property = NULL;

    /* Instantiate a RTPS_CrcClass for your custom function*/
    struct RTPS_ChecksumClass custom_crc32 =
    {
        RTPS_CHECKSUM_CLASSID_BUILTIN32, /*class_id*/
        NULL, /*context*/
        CrcClassTest_custom_crc32_other /*Custom function*/
    };

    /* get the Domain Participant factory and registry*/
    factory = DDS_DomainParticipantFactory_get_instance();

    registry = DDS_DomainParticipantFactory_get_registry
        (DDS_DomainParticipantFactory_get_instance());

    /* unregister the RTPS transport */
    if (!RT_Registry_unregister(registry, NETIO_DEFAULT_RTPS_NAME,
                               NULL, NULL))
    {
        printf("failed to unregister rtps\n");
        return 0;
    }

    rtps_property = (struct RTPS_InterfaceFactoryProperty *)
        malloc(sizeof(struct RTPS_InterfaceFactoryProperty));

    if (rtps_property == NULL)
    {
        printf("failed to allocate rtps properties\n");
        return 0;
    }
}

```

(continues on next page)



(continued from previous page)

```

/* the rtps property takes the structure with the custom
 * function
 */

*rtps_property = RTPS_INTERFACE_FACTORY_DEFAULT;
rtps_property->checksum.allow_builtin_override = RTI_TRUE;
rtps_property->checksum.builtin_checksum32_class = custom_crc32;

/* register the RTPS transport */
if (!RT_Registry_register(registry, NETIO_DEFAULT_RTPS_NAME,
    RTPS_InterfaceFactory_get_interface(),
    (struct RT_ComponentFactoryProperty*)rtps_property,
    NULL))
{
    printf("failed to register rtps\n");
    return 0;
}

/* modify the domain participant qos */
dp_qos.protocol.compute_crc = DDS_BOOLEAN_TRUE;
dp_qos.protocol.check_crc = DDS_BOOLEAN_TRUE;
dp_qos.protocol.require_crc = DDS_BOOLEAN_TRUE;
dp_qos.protocol.computed_crc_kind = DDS_CHECKSUM_BUILTIN32;
dp_qos.protocol.allowed_crc_mask = DDS_CHECKSUM_BUILTIN32;

/* use the qos and the factory to create a participant */

DDS_DomainParticipantFactory_create_participant(
    factory, domain_id,&dp_qos, NULL,DDS_STATUS_MASK_NONE);
}

```

### 3.6.6 INTRA Transport

The builtin intra participant transport (INTRA) is a transport that bypasses RTPS and reduces the number of data-copies from three to one for data published by a *DataWriter* to a *DataReader* within the same participant. When a sample is published, it is copied directly to the data reader's cache (if there is space). This transport is used for communication between *DataReaders* and *DataWriters* created within the same participant by default.

Please refer to *Threading Model* for important details regarding application constraints when using this transport.

## Registering the INTRA Transport

The builtin INTRA transport is a *RTI Connext Micro* component that is automatically registered when the `DomainParticipantFactory_get_instance()` method is called. By default, data published by a *DataWriter* is sent to all *DataReaders* within the same participant using the INTRA transport.

In order to prevent the INTRA transport from being used it is necessary to remove it as a transport and a user-data transport. The following code shows how to only use the builtin UDP transport for user-data.

```
struct DDS_DomainParticipantQos dp_qos =
    DDS_DomainParticipantQos_INITIALIZER;

REDA_StringSeq_set_maximum(&dp_qos.transports.enabled_transports,1);
REDA_StringSeq_set_length(&dp_qos.transports.enabled_transports,1);
*REDA_StringSeq_get_reference(&dp_qos.transports.enabled_transports,0) =
    REDA_String_dup(NETIO_DEFAULT_UDP_NAME);

/* Use only unicast for user-data traffic. */
REDA_StringSeq_set_maximum(&dp_qos.user_traffic.enabled_transports,1);
REDA_StringSeq_set_length(&dp_qos.user_traffic.enabled_transports,1);
*REDA_StringSeq_get_reference(&dp_qos.user_traffic.enabled_transports,0) =
    REDA_String_dup("_udp://");
```

Note that the INTRA transport is never used for discovery traffic internally. It is not possible to disable matching of *DataReaders* and *DataWriters* within the same participant.

## Reliability and Durability

Because a sample sent over INTRA bypasses the RTPS reliability and DDS durability queue, the [Reliability](#) and [Durability](#) Qos policies are *not* supported by the INTRA transport. However, by creating all the *DataReaders* before the *DataWriters* durability is not required.

## Threading Model

The INTRA transport does not create any threads. Instead, a *DataReader* receives data over the INTRA transport in the context of the *DataWriter*'s *send thread*.

This model has two *important limitations*:

- Because a *DataReader*'s `on_data_available()` listener is called in the context of the *DataWriter*'s send thread, a *DataReader* may potentially process data at a different priority than intended (the *DataWriter*'s). While it is generally not recommended to process data in a *DataReader*'s `on_data_available()` listener, it is particularly important *to not do so* when using the INTRA transport. Instead, use a DDS WaitSet or a similar construct to wake up a separate thread to process data.
- Because a *DataReader*'s `on_data_available()`

- listener is called in the context of the *DataWriter*'s send thread, any method called in the `on_data_available()` listener is done in the context of the *DataWriter*'s stack. Calling a *DataWriter* `write()` in the callback could result in an infinite call stack. Thus, it is recommended *not* to call in this listener any *Connext Micro* APIs that write data.

### 3.6.7 Shared Memory Transport (SHMEM)

This section describes the optional builtin *RTI Connext Micro* SHMEM transport and how to configure it.

Shared Memory Transport (SHMEM) is an optional transport that can be used in *Connext Micro*. It is part of a standalone library that can be optionally linked in.

The SHMEM Transport also allows *Connext Micro* to transmit data samples without copying them internally. For an overview of this feature, see *Zero Copy Transfer*.

Currently, *Connext Micro* supports the following functionality:

- Unicast
- Configuration of the shared memory receive queues

#### Registering the SHMEM Transport

The builtin SHMEM transport is a *Connext Micro* component that needs to be registered before a *DomainParticipant* can be created with the ability to send data across shared memory. Unlike the UDP Transport, this transport is not automatically registered. Register the transport using the code snippet below:

```
#include "netio_shmem/netio_shmem.h"

...

{
    DDS_DomainParticipantFactory *factory = NULL;
    RT_Registry_T *registry = NULL;
    struct NETIO_SHMEMInterfaceFactoryProperty shmem_property = NETIO_
    ↪SHMEMInterfaceFactoryProperty_INITIALIZER;
    struct DDS_DomainParticipantQos dp_qos = DDS_DomainParticipantQos_INITIALIZER;

    /* Optionally configure the transport settings */
    shmem_property.received_message_count_max = ...
    shmem_property.receive_buffer_size = ...
    shmem_property.message_size_max = ...

    factory = DDS_DomainParticipantFactory_get_instance();

    registry = DDS_DomainParticipantFactory_get_registry(factory);
    if (!RT_Registry_register(
        registry,
        "_shmem",
```

(continues on next page)

(continued from previous page)

```

        NETIO_SHMEMInterfaceFactory_get_interface(),
        (struct RT_ComponentFactoryProperty*)&shmem_property,
        NULL))
    {
        /* ERROR */
    }

    /* Enable the transport on a Domain Participant */
    DDS_StringSeq_set_maximum(&dp_qos.transports.enabled_transports,1);
    DDS_StringSeq_set_length(&dp_qos.transports.enabled_transports,1);
    *DDS_StringSeq_get_reference(&dp_qos.transports.enabled_transports,0) = DDS_String_
↪dup("_shmem");

    DDS_StringSeq_set_maximum(&dp_qos.discovery.enabled_transports,1);
    DDS_StringSeq_set_length(&dp_qos.discovery.enabled_transports,1);
    *DDS_StringSeq_get_reference(&dp_qos.discovery.enabled_transports,0) = DDS_String_
↪dup("_shmem://");

    DDS_StringSeq_set_maximum(&dp_qos.user_traffic.enabled_transports,1);
    DDS_StringSeq_set_length(&dp_qos.user_traffic.enabled_transports,1);
    *DDS_StringSeq_get_reference(&dp_qos.user_traffic.enabled_transports,0) = DDS_String_
↪dup("_shmem://");

    DDS_StringSeq_set_maximum(&dp_qos.discovery.initial_peers,1);
    DDS_StringSeq_set_length(&dp_qos.discovery.initial_peers,1);
    *DDS_StringSeq_get_reference(&dp_qos.discovery.initial_peers,0) = DDS_String_dup("_
↪shmem://");

    ...

    /* Explicitly unregister the shared memory transport before clean up */
    if (!RT_Registry_unregister(
        registry,
        "_shmem",
        NULL,
        NULL)
    {
        /* ERROR */
    }
}

```

The above snippet will register a transport with the default settings. To configure it, change the individual configurations as described in *SHMEM Configuration*.

When a component is registered, the registration takes the properties and a listener as the 3rd and 4th parameters. The registration of the shared memory component will make a copy of the properties configurable within a shared memory transport. There is currently no support for passing in a listener as the 4th parameter.

It should be noted that the SHMEM transport can be registered with any name, but all transport QoS policies and initial peers must refer to this name. If a transport is referred to and it does not exist, an error message is logged.

While it is possible to register multiple SHMEM transports, it is not possible to use multiple SHMEM transports within the same participant. The reason is that SHMEM port allocation is not synchronized between transports.

## Threading Model

The SHMEM transport creates one receive thread for each unique SHMEM receive address and port. Thus, by default two SHMEM threads are created:

- A unicast receive thread for discovery data
- A unicast receive thread for user data

Each receive thread will create a shared memory segment that will act as a message queue. Other *DomainParticipants* will send RTPS message to this message queue.

This message queue has a fixed size and can accommodate a fixed number of messages (*received\_message\_count\_max*) each with a maximum payload size of (*message\_size\_max*). The total size of the queue is configurable with (*receive\_buffer\_size*).

## Configuring SHMEM Receive Threads

All threads in the SHMEM transport share the same thread settings. It is important to note that all the SHMEM properties must be set before the SHMEM transport is registered. *Connext Micro* preregisters the SHMEM transport with default settings when the [DomainParticipantFactory](#) is initialized. To change the SHMEM thread settings, use the following code.

```
struct SHMEM_InterfaceFactoryProperty shmem_property = NETIO_SHMEMInterfaceFactoryProperty_
↪INITIALIZER

shmem_property.recv_thread_property.options = ...;

/* The stack-size is platform dependent, it is passed directly to the OS */
shmem_property.recv_thread_property.stack_size = ...;

/* The priority is platform dependent, it is passed directly to the OS */
shmem_property.recv_thread_property.priority = ...;

if (!RT_Registry_register(registry, "_shmem",
                        SHMEM_InterfaceFactory_get_interface(),
                        (struct RT_ComponentFactoryProperty*)&shmem_property,
                        NULL))
{
    /* ERROR */
}
```

## SHMEM Configuration

All the configuration of the SHMEM transport is done via the struct SHMEM\_InterfaceFactoryProperty structure:

```
struct NETIO_SHMEMInterfaceFactoryProperty
{
    struct NETIO_InterfaceFactoryProperty _parent;
    /* Max number of received message sizes that can be residing
       inside the shared memory transport concurrent queue
    */
    RTI_INT32 received_message_count_max;
    /* The size of the receive socket buffer */
    RTI_INT32 receive_buffer_size;
    /* The maximum size of the message which can be received */
    RTI_INT32 message_size_max;
    /* Thread properties for each receive thread created by this
       NETIO interface.
    */
    struct OSAPI_ThreadProperty recv_thread_property;
};
```

### received\_message\_count\_max

The number of maximum RTPS messages that can be inside a receive thread's receive buffer. By default this is 64.

### receive\_buffer\_size

The size of the message queue residing inside a shared memory region accessible from different processes. The default size is  $((received\_message\_count\_max * message\_size\_max) / 4)$ .

### message\_size\_max

The size of an RTPS message that can be sent across the shared memory transport. By default this number is 65536.

### recv\_thread\_property

The recv\_thread field is used to configure all the receive threads. Please refer to *Threading Model* for details.

### **pro\_minimum\_compatibility\_version**

The minimum version of *Connext Professional* with which to guarantee compatibility when using shared memory. This only needs to be specified if `dds.transport.minimum_compatibility_version` has been specified in *Connext Professional* and compatibility with *Connext Micro* is required. The default value is `DDS_PRODUCTVERSION_UNKNOWN`.

See [Capturing Shared Memory Traffic in the Core Libraries User's Manual](#) for more information on `dds.transport.minimum_compatibility_version` in *Connext Professional*.

### **Caveats**

#### **Leftover shared memory resources**

*Connext Micro* implements the shared memory transport and utilizes shared memory semaphores that can be used concurrently by processes. *Connext Micro* implements a shared memory mutex from a shared memory semaphore. If an application exits ungracefully, then the shared memory mutex may be left in a state that prevents it from being used. This can occur because the *Connext Micro* Shared Memory Transport tries to re-use and clean up and leftover segments as a result of an applications ungraceful termination. If ungraceful termination occurs, the leftover shared memory mutexes need to be cleaned up either manually or by restarting the system.

The same applies to shared memory semaphores. If an application exists ungracefully, there can be leftover shared memory segments.

### **Darwin and Linux systems**

In the case of Darwin and Linux systems which use SysV semaphores, you can view any leftover shared memory segments using `ipcs -a`. They can be removed using the `ipcrm` command. Shared memory keys used by *Connext Micro* are in the range of 0x00400000. For example:

- `ipcs -m | grep 0x004`

The shared semaphore keys used by *Connext Micro* are in the range of 0x800000; the shared memory mutex keys are in the range of 0xb00000. For example:

- `ipcs -m | grep 0x008`
- `ipcs -m | grep 0x00b`

## QNX systems

QNX® systems use POSIX® APIs to create shared memory segments or semaphores. The shared memory segment resources are located in `/dev/shmem` and the shared memory mutex and semaphores are located in `/dev/sem`.

To view any leftover shared memory segments when no *Connext Micro* applications are running:

- `ls /dev/shmem/RTIOsapi*`
- `ls /dev/sem/RTIOsapi*`

To clean up the shared memory resources, remove the files listed.

## Windows and VxWorks systems

On Windows and VxWorks® systems, once all the processes that are attached to a shared memory segment, shared memory mutex, or shared memory semaphores are terminated (either gracefully or ungracefully), the shared memory resources will be automatically cleaned up by the operating system.

### 3.6.8 Zero Copy v2 Transport

The Zero Copy v2 transport enables *RTI Connext Micro* to share data samples between publishers and subscribers without serializing, transmitting, or deserializing the samples. For an overview of this feature and its utility, see *Zero Copy Transfer*.

This section outlines the basic steps required to enable the Zero Copy v2 transport in an application. All the example code shown below is taken from a Zero Copy v2 application that you can generate using *rtiddsgen* (see *Example Applications* for more details).

#### Generate example and type support files

First, identify types that require Zero Copy transfer and annotate them with the `@transfer_mode(SHMEM_REF)` annotation. See the example IDL file below:

```
@transfer_mode(SHMEM_REF)
struct HelloWorld {
    @key long id;
    long data[100];
};
```

*rtiddsgen* generates additional *TypePlugin* code when a type is annotated with `@transfer_mode(SHMEM_REF)` in the IDL files. This code allows a *DataWriter* and a *DataReader* to communicate using a reference to the sample in shared memory.



---

**Note:** Zero Copy v2 for *Connext Micro* only supports contiguous data types; this means that fixed-size arrays are supported, but sequences and strings are not.

---

Next, generate the type support and example files with the following command:

```
rtiddsgen -micro -example -exampleTemplate zcv2 -language C HelloWorld.idl
```

The generated files will appear in the same directory as the type file.

### Initialize the Zero Copy v2 transport

Before the Zero Copy v2 transport can be used, it must be initialized. This must be done before creating a *DomainParticipant* and after registering the *DataReader* and *DataWriter* history plugins. The order is important because the Zero Copy v2 transport will perform actions on the history components during initialization.

The following example code from **HelloWorldApplication.c** demonstrates how to initialize the Zero Copy v2 transport with `NDDS_Transport_ZeroCopy_initialize()`:

```
if (!NDDS_Transport_ZeroCopy_initialize(registry, NULL, NULL))
{
    printf("failed to initialize zero copy\n");
    /* handle error */
}
```

### Register the Zero Copy v2 transport

The Zero Copy v2 transport needs a notification mechanism to notify *DataReaders* when a *DataWriter* has published data samples. RTI provides a default notification mechanism based on POSIX. You can also implement your own custom notification mechanism, but doing so is beyond the scope of this documentation; for more information, contact [support@rti.com](mailto:support@rti.com).

The default provided by RTI is a POSIX implementation of the notification mechanism. This mechanism is based on a monitor implemented in shared memory. In this documentation, we will assume that you are using the default implementation unless otherwise noted.

Once the Zero Copy transport is initialized, configure the `notif` interface factory with the `ZCOPY_NotifInterfaceFactoryProperty` property. This property has three fields you need to set:

- `max_samples_per_notif`: The number of samples processed per notification. By default this value is 1. Note that a high value may starve other threads from progressing.
- `user_intf`: This is the implementation of your chosen notification mechanism. It is populated automatically if you are using the default implementation.
- `user_property`: Any properties associated with your chosen notification mechanism. *Connext Micro* treats this as an opaque pointer.

When using the default mechanism provided by RTI, the `user_property` mentioned above is resolved to `ZCOPY_NotifMechanismProperty`. Both of these properties are required to configure the transport.

See the following example from `HelloWorldApplication.c`:

```
struct ZCOPY_NotifInterfaceFactoryProperty notif_prop;
struct ZCOPY_NotifMechanismProperty notif_mech_prop;
notif_mech_prop.intf_addr = 0;
notif_prop.user_property = &notif_mech_prop;
notif_prop.max_samples_per_notif = 1;
```

For more information on these properties, see the *Configuration* section.

Finally, call `ZCOPY_NotifMechanism_register()` (a utility function on the default notification mechanism) to register the Zero Copy v2 transport with the default notification mechanism. This makes it available for use. The following example registers a `notif` with the name `NETIO_DEFAULT_NOTIF_NAME`:

```
if (!ZCOPY_NotifMechanism_register(registry, NETIO_DEFAULT_NOTIF_NAME, &notif_prop))
{
    printf("failed to register notif\n");
    goto done;
}
```

## Enable transports

With the specific notification mechanism registered, you can enable the Zero Copy and UDP transports for the *DomainParticipant*. Consider the following example code:

```
if (!DDS_StringSeq_set_maximum(&dp_qos.transports.enabled_transports, 2))
{
    printf("failed to set transports.enabled_transports maximum\n");
    goto done;
}
if (!DDS_StringSeq_set_length(&dp_qos.transports.enabled_transports, 2))
{
    printf("failed to set transports.enabled_transports length\n");
    goto done;
}
/* UDP and Notif are enabled*/
*DDS_StringSeq_get_reference(&dp_qos.transports.enabled_transports, 0) =
    DDS_String_dup(NETIO_DEFAULT_NOTIF_NAME);
*DDS_StringSeq_get_reference(&dp_qos.transports.enabled_transports, 1) =
    DDS_String_dup(NETIO_DEFAULT_UDP_NAME);

/* Discovery takes place over UDP */
DDS_StringSeq_set_maximum(&dp_qos.discovery.enabled_transports, 1);
DDS_StringSeq_set_length(&dp_qos.discovery.enabled_transports, 1);
*DDS_StringSeq_get_reference(&dp_qos.discovery.enabled_transports, 0) =
    DDS_String_dup("_udp://");
```

(continues on next page)

(continued from previous page)

```

/* User data uses Notif */
DDS_StringSeq_set_maximum(&dp_qos.user_traffic.enabled_transports, 1);
DDS_StringSeq_set_length(&dp_qos.user_traffic.enabled_transports, 1);
*DDS_StringSeq_get_reference(&dp_qos.user_traffic.enabled_transports, 0) =
    DDS_String_dup("notif://");

```

**Note:** The *UDP Transport* or *Shared Memory Transport (SHMEM)* must be registered while using Zero Copy v2 transfer because DDS Discovery requires one of them in order to function (see *Discovery* for more details).

## Sample management

When using the Zero Copy v2 transport, each *DataWriter* manages a pool of samples, and the application must obtain samples from this pool using [get\\_loan\(\)](#). We can see this in the following example:

```

hw_datawriter = HelloWorldDataWriter_narrow(datawriter);
retcode = HelloWorldDataWriter_get_loan(hw_datawriter, &sample);
if (retcode != DDS_RETCODE_OK)
{
    printf("ERROR: Failed to loan sample\n");
}
retcode = HelloWorldDataWriter_write(hw_datawriter, sample, &DDS_HANDLE_NIL);
if (retcode != DDS_RETCODE_OK)
{
    printf("ERROR: Failed to write to sample\n");
}

```

As seen above, the *DataWriter* **must** get a loan before each write call; it cannot write a loaned sample multiple times. The *DataWriter* does not need to explicitly return any loan to the pool, since this is managed by the middleware. However, if a loaned sample will not be written, it can be discarded with [discard\\_loan\(\)](#).

**Warning:** It is not possible to write a sample that has not been obtained with [get\\_loan\(\)](#).

**Note:** A Zero Copy-enabled *DataWriter* can also send samples using other transports (such as UDPv4) to non-Zero Copy *DataReaders*. When a *DataWriter* uses both the Zero Copy v2 transport and a transport which uses serialized data (such as UDP), the same sample is sent over all transports.

This may adversely affect performance, since the sample must be serialized for network transmission even if it is in shared memory. For best performance, you should consider an architecture where

a *DataWriter* matches either with Zero Copy-enabled or non Zero Copy-enabled *DataReaders*, but not both.

---

On the *DataReader* side, Zero Copy v2 application code is identical to subscribing applications not using Zero Copy.

When a *DataReader* calls [read\(\)](#) or [take\(\)](#) and receives samples, it is being given samples that are loaned from the *DataWriter*'s pool. Thus, failing to return the loan when the sample is no longer needed will deplete the available samples in the pool, eventually causing calls to [get\\_loan\(\)](#) to fail.

## Configuration

*Connext Micro* Zero Copy v2 includes some properties unique to its functionality. The following properties are always required:

- [max\\_samples\\_per\\_notif](#)
- [user\\_intf](#)<sup>1</sup>
- [user\\_property](#)<sup>2</sup>

The following properties are required if you are using the default implementation of the notification mechanism for Zero Copy v2. These are essentially a default set of user-defined properties; if you are using your own notification mechanism, you can set your own user-defined properties as needed.

- [user\\_property.intf\\_addr](#)
- [user\\_property.thread\\_prop](#)
  - [user\\_property.thread\\_prop.stack\\_size](#)
  - [user\\_property.thread\\_prop.priority](#)
  - [user\\_property.thread\\_prop.options](#)
- [user\\_property.max\\_receive\\_ports](#)
- [user\\_property.max\\_routes](#)

The following additional properties are only required if you are using your own notification mechanism for Zero Copy v2, not the default implementation.

- [user\\_intf.create\\_instance](#)
- [user\\_intf.delete\\_instance](#)
- [user\\_intf.get\\_route\\_table](#)
- [user\\_intf.reserve\\_address](#)
- [user\\_intf.release\\_address](#)
- [user\\_intf.resolve\\_address](#)

---

<sup>1</sup> This property is only required if you choose to implement your own notification mechanism and not use the default implementation provided by RTI.

<sup>2</sup> Resolves to [ZCOPY\\_NotifMechanismProperty](#) when using the default notification mechanism.

- [user\\_intf.add\\_route](#)
- [user\\_intf.delete\\_route](#)
- [user\\_intf.bind](#)
- [user\\_intf.unbind](#)
- [user\\_intf.send](#)
- [user\\_intf.notify\\_rcv\\_port](#)
- [user\\_intf.create\\_instance](#)

### Using multiple Zero Copy v2 transport instances

The platform-independent Zero Copy v2 transport supports multiple instances, provided that the user-defined, platform-specific implementation of the `notif` interface implements a way to uniquely identify each instance. In this case, each Zero Copy v2 transport instance should be registered with uniquely different names and properties.

When multiple instances of the Zero Copy v2 transport exist, individual *DataReaders* and *DataWriters* can be configured to use a specific instance of the Zero Copy v2 transport. This configuration is done in the entity's [enabled\\_transports](#) QoS configuration. For more information, see *Transport Registration*.

### 3.6.9 UDP Transport

This section describes the builtin *RTI Connext Micro* UDP transport and how to configure it.

The builtin UDP transport (UDP) is a fairly generic UDPv4 transport. *Connext Micro* supports the following functionality:

- Unicast
- Multicast
- Automatic detection of available network interfaces
- Manual configuration of network interfaces
- Allow/Deny lists to select which network interfaces can be used to receive data
- Simple NAT configuration
- Configuration of receive threads

---

**Note:** *Connext Micro* supports up to four network interfaces at once for each of the following:

- Unicast user-data
- Multicast user-data
- Unicast discovery data

- Multicast discovery data

## Registering the UDP Transport

The builtin UDP transport is a *Connext Micro* component that is automatically registered when the `DDS_DomainParticipantFactory_get_instance()` method is called. To change the UDP configuration, it is necessary to first unregister the transport as shown below:

```
DDS_DomainParticipantFactory *factory = NULL;
RT_Registry_T *registry = NULL;

factory = DDS_DomainParticipantFactory_get_instance();
registry = DDS_DomainParticipantFactory_get_registry(factory);

/* The builtin transport does not return any properties (3rd param) or
 * listener (4th param)
 */
if (!RT_Registry_unregister(registry, "_udp", NULL, NULL))
{
    /* ERROR */
}
```

When a component is registered, the registration takes the properties and a listener as the 3rd and 4th parameters. In general, it is up to the caller to manage the memory for the properties and the listeners. There is no guarantee that a component makes a copy.

The following code-snippet shows how to register the UDP transport with new parameters.

```
struct UDP_InterfaceFactoryProperty *udp_property = NULL;

/* Allocate a property structure for the heap, it must be valid as long
 * as the component is registered
 */
udp_property = (struct UDP_InterfaceFactoryProperty *)
    malloc(sizeof(struct UDP_InterfaceFactoryProperty));
if (udp_property != NULL)
{
    *udp_property = UDP_INTERFACE_FACTORY_PROPERTY_DEFAULT;

    /* Only allow network interface "eth0" to be used;
     */
    REDA_StringSeq_set_maximum(&udp_property->allow_interface, 1);
    REDA_StringSeq_set_length(&udp_property->allow_interface, 1);

    *REDA_StringSeq_get_reference(&udp_property->allow_interface, 0) =
        REDA_String_dup("eth0");

    /* Register the transport again, using the builtin name
     */
    if (!RT_Registry_register(registry, "_udp",
```

(continues on next page)

(continued from previous page)

```

        UDP_InterfaceFactory_get_interface(),
        (struct RT_ComponentFactoryProperty*)udp_property,
        NULL))
    {
        /* ERROR */
    }
}
else
{
    /* ERROR */
}

```

It should be noted that the UDP transport can be registered with any name, but all transport QoS policies and initial peers must refer to this name. If a transport is referred to and it does not exist, an error message is logged.

It is possible to register multiple UDP transports with a [DomainParticipantFactory](#). It is also possible to use different UDP transports within the same *DomainParticipant* when multiple network interfaces are available (either physical or virtual).

When UDP transformations are enabled, this feature is always enabled and determined by the *allow\_interface* and *deny\_interface* lists. If any of the lists are non-empty the UDP transports will bind each receive socket to the specific interfaces.

When UDP transformations are not enabled, this feature is determined by the value of the *enable\_interface\_bind*. If this value is set to **RTI\_TRUE** and the *allow\_interface* and/or *deny\_interface* properties are non-empty, the receive sockets are bound to specific interfaces.

## Threading Model

The UDP transport creates one receive thread for each unique UDP receive address and port. Thus, by default, three UDP threads are created:

- A multicast receive thread for discovery data (assuming multicast is available and enabled)
- A unicast receive thread for discovery data
- A unicast receive thread for user data

Additional threads may be created depending on the transport configuration for a *DomainParticipant*, *DataReader*, and *DataWriter*. The UDP transport creates threads based on the following criteria:

- Each unique unicast port creates a new thread
- Each unique multicast address *and* port creates a new thread

For example, if a *DataReader* specifies its own multicast receive address, a new receive thread will be created.

## Configuring UDP Receive Threads

All threads in the UDP transport share the same thread settings. It is important to note that all the UDP properties must be set before the UDP transport is registered. *Connext Micro* preregisters the UDP transport with default settings when the [DomainParticipantFactory](#) is initialized. To change the UDP thread settings, use the following code.

```
struct UDP_InterfaceFactoryProperty *udp_property = NULL;
struct UDP_InterfaceFactoryProperty udp_property =
    UDP_INTERFACE_FACTORY_PROPERTY_DEFAULT;

/* Allocate a property structure for the heap, it must be valid as long
 * as the component is registered
 */
udp_property = (struct UDP_InterfaceFactoryProperty *)
    malloc(sizeof(struct UDP_InterfaceFactoryProperty));
*udp_property = UDP_INTERFACE_FACTORY_PROPERTY_DEFAULT;

/* Please refer to OSAPI_ThreadOptions for possible options */
udp_property->recv_thread.options = ...;

/* The stack-size is platform dependent, it is passed directly to the OS */
udp_property->recv_thread.stack_size = ....

/* The priority is platform dependent, it is passed directly to the OS */
udp_property->recv_thread.priority = ....

if (!RT_Registry_register(registry, "_udp",
    UDP_InterfaceFactory_get_interface(),
    (struct RT_ComponentFactoryProperty*)udp_property,
    NULL))
{
    /* ERROR */
}
```

## UDP Configuration

You can configure the UDP transport via the [UDP\\_InterfaceFactoryProperty](#). The following fields are available:

### allow\_interface

The *allow\_interface* string sequence determines which interfaces are allowed to be used for communication. Each string element is the name of a network interface, such as “en0” or “eth1”.

If this sequence is empty, all interface names pass the allow test. The default value is empty. Thus, all interfaces are allowed.



**deny\_interface**

The *deny\_interface* string sequence determines which interfaces are not allowed to be used for communication. Each string element is the name of a network interface, such as “en0” or “eth1”.

If this sequence is empty, the test is false. That is, the interface is allowed. Note that the deny list is checked *after* the allow list. Thus, if an interface appears in both, it is denied. The default value is empty, thus no interfaces are denied.

**max\_send\_buffer\_size**

The *max\_send\_buffer\_size* is the maximum size of the send socket buffer and it *must* be at least as big as the largest sample. Typically, this buffer should be a multiple of the maximum number of samples that can be sent at any given time. The default value is 256KB.

**max\_receive\_buffer\_size**

The *max\_receive\_buffer\_size* is the maximum size of the receive socket buffer and it *must* be at least as big as the largest sample. Typically, this buffer should be a multiple of the maximum number of samples that can be received at any given time. The default value is 256KB.

**max\_message\_size**

The *max\_message\_size* is the maximum size of the message which can be received, including any packet overhead. The default value is 65507 bytes.

**multicast\_ttl**

The *multicast\_ttl* is the Multicast Time-To-Live (TTL). This value is only used for multicast. It limits the number of hops a packet can pass through before it is dropped by a router. The default value is 1.

**nat**

*Connex Micro* supports firewalls with NAT. However, this feature has limited use and only supports translation between a private and public IP address. UDP ports are not translated. Furthermore, because *Connex Micro* does not support any hole punching technique or WAN server, this feature is only useful when the private and public address mapping is static and known in advance. For example, to test between an Android emulator and the host, the following configuration can be used:

```

UDP_NatEntrySeq_set_maximum(&udp_property->nat,2);
UDP_NatEntrySeq_set_length(&udp_property->nat,2);

/* Translate the local emulator eth0 address 10.10.2.f:7410 to
 * 127.0.0.1:7410. This ensures that the address advertised by the
 * emulator to the host machine is the host's loopback interface, not
 * the emulator's host interface
 */
UDP_NatEntrySeq_get_reference(&udp_property->nat,0)->
    local_address.kind = NETIO_ADDRESS_KIND_UDPv4;
UDP_NatEntrySeq_get_reference(&udp_property->nat,0)->
    local_address.port = 7410;
UDP_NatEntrySeq_get_reference(&udp_property->nat,0)->
    local_address.value.ipv4.address = 0x0a00020f;

UDP_NatEntrySeq_get_reference(&udp_property->nat,0)->
    public_address.kind = NETIO_ADDRESS_KIND_UDPv4;
UDP_NatEntrySeq_get_reference(&udp_property->nat,0)->
    public_address.port = 7410;
UDP_NatEntrySeq_get_reference(&udp_property->nat,0)->
    public_address.value.ipv4.address = 0x7f000001;

/* Translate the local emulator eth0 address 10.10.2.f:7411 to
 * 127.0.0.1:7411. This ensures that the address advertised by the
 * emulator to the host machine is the host's loopback interface
 */
UDP_NatEntrySeq_get_reference(&udp_property->nat,1)->
    local_address.kind = NETIO_ADDRESS_KIND_UDPv4;
UDP_NatEntrySeq_get_reference(&udp_property->nat,1)->
    local_address.port = 7411;
UDP_NatEntrySeq_get_reference(&udp_property->nat,1)->
    local_address.value.ipv4.address = 0x0a00020f;

UDP_NatEntrySeq_get_reference(&udp_property->nat,1)->
    public_address.kind = NETIO_ADDRESS_KIND_UDPv4;
UDP_NatEntrySeq_get_reference(&udp_property->nat,1)->
    public_address.port = 7411;
UDP_NatEntrySeq_get_reference(&udp_property->nat,1)->
    public_address.value.ipv4.address = 0x7f000001;

```

## if\_table

The *if\_table* provides a method to manually configure which interfaces are available for use; for example, when using IP stacks that do not support reading interface lists. The following example shows how to manually configure the interfaces.

```

/* The arguments to the UDP_InterfaceTable_add_entry functions are:
 * The if_table itself
 * The network address of the interface
 * The netmask of the interface

```

(continues on next page)

(continued from previous page)

```

* The name of the interface
* Interface flags. Valid flags are:
*   UDP_INTERFACE_INTERFACE_UP_FLAG           - The interface is UP
*   UDP_INTERFACE_INTERFACE_MULTICAST_FLAG - The interface supports multicast
*/
if (!UDP_InterfaceTable_add_entry(&udp_property->if_table,
                                0x7f000001, 0xff000000, "loopback",
                                UDP_INTERFACE_INTERFACE_UP_FLAG |
                                UDP_INTERFACE_INTERFACE_MULTICAST_FLAG))
{
    /* Error */
}

```

**multicast\_interface**

The *multicast\_interface* may be used to select a particular network interface to be used to send multicast packets. The default value is any interface (that is, the OS selects the interface).

**is\_default\_interface**

The *is\_default\_interface* flag is used to indicate that this *Connex Micro* network transport shall be used if no other transport is found. The default value is **RTI\_TRUE**.

**disable\_auto\_interface\_config**

Normally, the UDP transport will try to read out the interface list (on platforms that support it). Setting *disable\_auto\_interface\_config* to **RTI\_TRUE** will prevent the UDP transport from reading the interface list.

**multicast\_loopback\_disabled**

The [multicast\\_loopback\\_disabled](#) field controls whether *Connex Micro* puts multicast packets onto the loopback interface.

**recv\_thread**

The *recv\_thread* field is used to configure all the receive threads. Please refer to *Threading Model* for details.

**enable\_interface\_bind**

When this is set to **TRUE** the UDP transport binds each receive port to a specific interface when the *allow\_interface/deny\_interface* lists are non-empty. This allows multiple UDP transports to be used by a single *DomainParticipant* at the expense of an increased number of threads. This property is ignored when transformations are enabled and the *allow\_interface/deny\_interface* lists are non-empty.

**disable\_multicast\_bind**

The [disable\\_multicast\\_bind](#) field controls whether *Connex Micro* will bind to a multicast address receive address (if set to 0) or bind to ANY multicast address (if set to 1).

**disable\_multicast\_interface\_select**

The [disable\\_multicast\\_interface\\_select](#) field controls whether *Connex Micro* will use the *multicast\_interface* (if specified), the *allow\_interface/deny\_interface* (if specified and *multicast\_interface* is not specified) to select the interfaces used for sending to multicast addresses. If set to 1, *Connex Micro* will not select any interface.

**source\_rules**

Rules for how to transform received UDP payloads based on the source address.

**destination\_rules**

Rules for how to transform sent UDP payloads based on the destination address.

**transform\_udp\_mode**

Determines how regular UDP is supported when transformations are supported. When transformations are enabled the default value is **UDP\_TRANSFORM\_UDP\_MODE\_DISABLED**.

**transform\_locator\_kind**

The locator to use for locators that have transformations. When transformation rules have been enabled, they are announced as a vendor specific locator. This property overrides this value.

NOTE: Changing this value may prevent communication.

**UDP Transformations**

The UDP transform feature enables custom transformation of incoming and outgoing UDP payloads based on transformation rules between a pair of source and destination IP addresses. Some examples of transformations are encrypted data or logging.

This section explains how to implement and use transformations in an application and is organized as follows:

- *Overview*
- *Enabling UDP Transformation*
- *Creating a Transformation Library*
- *Creating Transformation Rules*
- *Interoperability*
- *Error Handling*
- *Example Code*
- *Examples*
- *OS Configuration*

**Overview**

The UDP transformation feature enables custom transformation of incoming and outgoing UDP payloads. For the purpose of this section, a UDP payload is defined as a sequence of octets sent or received as a single UDP datagram excluding UDP headers – typically UDP port numbers – and trailers, such as the optional used checksum.

An outgoing payload is the UDP payload passed to the network stack. The transformation feature allows a custom transformation of this payload just before it is sent. The UDP transport receives payloads to send from an upstream layer. In *Connex Micro* this layer is typically RTPS, which creates payloads containing one or more RTPS messages. The transformation feature enables transformation of the entire RTPS payload before it is passed to the network stack.

The same RTPS payload may be sent to one or more locators. A locator identifies a destination address, such as an IPv4 address, a port, such as a UDP port, and a transport kind. The address and port are used by the UDP transport to reach a destination. However, only the destination address is used to determine which transformation to apply.

An incoming payload is the UDP payload received from the network stack. The transformation feature enables transformation of the UDP payload received from the network stack *before* it is passed to the upstream interface, typically RTPS. The UDP transport only receives payloads destined for one of its network interface addresses, but may receive UDP payloads destined for many different ports. The transformation does not take a port into account, only the source address. In *Connext Micro* the payload is typically a RTPS payload containing one or more RTPS messages.

UDP transformations are registered with *Connext Micro* and used by the UDP transport to determine how to transform payloads based on a source or destination address. Please refer to *Creating a Transformation Library* for details on how to implement transformations and *Creating Transformation Rules* for how to add rules.

Transformations are local resources. There is no exchange between different UDP transports regarding what a transformation does to a payload. This is considered a-priori knowledge and depends on the implementation of the transformation. Any negotiation of e.g. keys must be handled before the UDP transport is registered. Thus, if a sender and receiver do not apply consistent rules, they may not be able to communicate, or incorrect data may result. Note that while information is typically in the direction from a *DataWriter* to a *DataReader*, a reliable *DataReader* also send protocol data to a *DataWriter*. These messages are also transformed.

## Network Interface Selection

When a *DomainParticipant* is created, it first creates an instance of each transport configured in the [DDS\\_DomainParticipantQos::transports](#) QoS policy. Thus, each UDP transport registered with *Connext Micro* must have a unique name (up to 7 characters). Each registered transport can be configured to use all or some of the available interfaces using the *allow\_interface* and *deny\_interface* properties. The registered transports may now be used for either discovery data (specified in [DomainParticipantQos::discovery](#)), user\_traffic (specified in [DomainParticipantQos::user\\_traffic](#)) or both. The *DomainParticipant* also queries the transport for which addresses it is capable of sending to.

When a participant creates multiple instances of the UDP transport, it is important that instances use non-overlapping networking interface resources.

## Data Reception

Which transport to use for discovery data is determined by the [DomainParticipantQos::discovery](#) QoS policy. For each transport listed, the *DomainParticipant* reserves a network address to listen to. This network address is sent as part of the discovery data and is used by other *DomainParticipants* as the address to send discovery data for this *DomainParticipant*. Because a UDP transformation only looks at source and destination addresses, if different transformations are needed for discovery and user-data, different UDP transport registrations must be used and hence different network interfaces.

## Data Transmission

Which address to send data to is based on the locators received as part of discovery and the peer list.

Received locators are analyzed and a transport locally registered with a *DomainParticipant* is selected based on the locator kind, address and mask. The first matching transport is selected. If a matching transport is not found, the locator is discarded.

NOTE: A transport is not a matching criteria at the same level as a QoS policy. If a discovered entity requests user data on a transport that doesn't exist, it is not unmatched.

The peer list, as specified by the application, is a list of locators to send participant discovery announcements to. If the transport to use is not specified, e.g. "udp1@192.168.1.1", but instead "192.168.1.1", then all transports that understand this address will send to it. Thus, in this case the latter is used, and two different UDP transports are registered; they will both send to the same address. However, one transport may send transformed data and the other may not depending on the destination address.

## Enabling UDP Transformation

The UDP transformation feature is not enabled by default. In order to use transformations, you must do the following:

1. Compile *Connex Micro* with the Cmake flag `-DRTIME_UDP_ENABLE_TRANSFORMS=TRUE`, and;
2. Compile your application with the `-DUDP_ENABLE_TRANSFORMS=1` C preprocessor flag.

---

**Note:** In order for UDP transformations to interoperate with a UDP transport without transformations enabled, you must specify the following UDP property:

```
udp_property->transform_udp_mode = UDP_TRANSFORM_UDP_MODE_ENABLED;
```

---

## Creating a Transformation Library

The transformation library is responsible for creating and performing transformations. Note that a library is a logical concept and does not refer to an actual library in an operating system. A library in this context is a collection of routines that together creates, manages, and performs transformations. How these routines are compiled and linked with an application using *Connex Micro* is out of scope of this section.

The transformation library must be registered with *Connex Micro*'s run-time and must implement the required interfaces. This ensures proper life-cycle management of transformation resources as well as clear guidelines regarding concurrency and memory management.

From *Connex Micro*'s run-time point of view, the transformation library must implement methods so that:

- A library can be initialized.
- A library can be instantiated.
- An instance of the library performs and manages transformations.

The first two tasks are handled by *Connext Micro*'s run-time factory interface which is common for all libraries managed by *Connext Micro*. The third task is handled by the transformation interface, which is specific to UDP transformations.

The following describes the relationship between the different interfaces:

- A library is initialized once when it is registered with *Connext Micro*.
- A library is finalized once when it is unregistered from *Connext Micro*.
- Multiple library instances can be created. If a library is used twice, for example registered with two different transports, two different library contexts are created using the factory interface. *Connext Micro* assumes that concurrent access to two different instances is allowed.
- Different instances of the library can be deleted independently. An instance is deleted using the factory interface.
- A library instance creates specific source or destination transformations. Each transformation is expected to transform a payload to exactly one destination or from one source.

The following relationship is true between the UDP transport and a UDP transformation library:

- Each registered UDP transport may make use of one or more UDP transformation libraries.
- A DDS *DomainParticipant* creates one instance of each registered UDP transport.
- Each instance of the UDP transport creates one instance of each enabled transformation library registered with the UDP transport.
- Each Transformation rule created by the UDP transport creates one send or one receive transformation.

## Creating Transformation Rules

Transformation rules decide how a payload should be transformed based on either a source or destination address. Before a UDP transport is registered, it must be configured with the transformation libraries to use, as well as which library to use for each source and destination address. For each UDP payload sent or received, an instance of the UDP transport searches for a matching source or destination rule to determine which transformation to apply.

The transformation rules are added to the [UDP\\_InterfaceFactoryProperty](#) before registration takes place.

If no transformation rules have been configured, all payloads are treated as regular UDP packets.

If no send rules have been asserted, the payload is sent as is. If all outgoing messages are to be transformed, a single entry is sufficient (address = 0, mask = 0).

If no receive rules have been asserted, it is passed upstream as is. If all incoming messages are to be transformed, a single entry is sufficient (address = 0, mask = 0).



If no matching rule is found, the packet is dropped and an error is logged.

NOTE: `UDP_InterfaceFactoryProperty` is immutable after the UDP transport has been registered.

## Interoperability

When the UDP transformations has enabled at least one transformation, it will only inter-operate with another UDP transport which also has at least one transformation.

UDP transformations does not interoperate with *RTI Connext Professional*.

## Error Handling

The transformation rules are applied on a local basis and correctness is based on configuration. It is not possible to detect that a peer participant is configured for different behavior and errors cannot be detected by the UDP transport itself. However, the transformation interface can return errors which are logged.

## Example Code

Example Header file `MyUdpTransform.h`:

```
#ifndef MyUdpTransform_h
#define MyUdpTransform_h

#include "rti_me_c.h"
#include "netio/netio_udp.h"
#include "netio/netio_interface.h"

struct MyUdpTransformFactoryProperty
{
    struct RT_ComponentFactoryProperty _parent;
};

extern struct RT_ComponentFactoryI*
MyUdpTransformFactory_get_interface(void);

extern RTI_BOOL
MyUdpTransformFactory_register(RT_Registry_T *registry,
                              const char *const name,
                              struct MyUdpTransformFactoryProperty *property);

extern RTI_BOOL
MyUdpTransformFactory_unregister(RT_Registry_T *registry,
                                 const char *const name,
                                 struct MyUdpTransformFactoryProperty **);

#endif
```

Example Source file MyUdpTransform.c:

```

/*ce
 * \file
 * \defgroup UDPTransformExampleModule MyUdpTransform
 * \ingroup UserManuals_UDPTransform
 * \brief UDP Transform Example
 *
 * \details
 *
 * The UDP interface is implemented as a NETIO interface and NETIO interface
 * factory.
 */

/*ce \addtogroup UDPTransformExampleModule
 * @f
 */
#include <stdio.h>

#include "MyUdpTransform.h"

/*ce
 * \brief The UDP Transformation factory class
 *
 * \details
 * All Transformation components must have a factory. A factory creates one
 * instance of the component as needed. In the case of UDP transformations,
 * \rttime creates one instance per UDP transport instance.
 */
struct MyUdpTransformFactory
{
    /*ce
     * \brief Base-class. All \rttime Factories must inherit from RT_ComponentFactory.
     */
    struct RT_ComponentFactory _parent;

    /*ce
     * \brief A pointer to the properties of the factory.
     *
     * \details
     *
     * When a factory is registered with \rttime it can be registered with
     * properties specific to the component. However \rttime does not
     * make a copy ( that would require additional methods). Furthermore, it
     * may not be desirable to make a copy. Instead, this decision is
     * left to the implementer of the component. \rttime does not access
     * any custom properties.
     */
    struct MyUdpTransformFactoryProperty *property;
};

/*ce

```

(continues on next page)

(continued from previous page)

```

* \brief The custom UDP transformation class.
*
* \details
* The MyUdpTransformFactory creates one instance of this class for each
* UDP interface created. In this example one packet buffer (NETIO_Packet_T),
* is allocated and a buffer to hold the transformed data (\ref buffer)
*
* Only one transformation can be done at a time and it is synchronous. Thus,
* it is sufficient with one buffer to transform input and output per
* instance of the MyUdpTransform.
*/
struct MyUdpTransform
{
    /*ce
    * \brief Base-class. All UDP transforms must inherit from UDP_Transform
    */
    struct UDP_Transform _parent;

    /*ce \brief A reference to its own factory, if properties must be accessed
    */
    struct MyUdpTransformFactory *factory;

    /*ce \brief NETIO_Packet to hold a transformed payload.
    *
    * \details
    *
    * \rttime uses a NETIO_Packet_T to abstract data payload and this is
    * what is being passed between the UDP transport and the transformation.
    * The transformation must convert a payload into a NETIO_Packet. This
    * is done with NETIO_Packet_initialize_from. This function saves all
    * state except the payload buffer.
    */
    NETIO_Packet_T packet;

    /*ce \brief The payload to assign to NETIO_Packet_T
    *
    * \details
    *
    * A transformation cannot do in-place transformations because the input
    * buffer may be sent multiple times (for example due to reliability).
    * A transformation instance can only transform one buffer at a time
    * (send or receive). The buffer must be large enough to hold a transformed
    * payload. When the the transformation is created it receives a
    * \ref UDP_TransformProperty. This property has the max send and
    * receive buffers for transport and can be used to size the buffer.
    * Please refer to \ref UDP_InterfaceFactoryProperty::max_send_message_size
    * and \ref UDP_InterfaceFactoryProperty::max_message_size.
    */
    char *buffer;

    /*ce \brief The maximum length of the buffer. NOTE: The buffer must

```

(continues on next page)

(continued from previous page)

```

    * be 1 byte larger than the largest buffer.
    */
    RTI_SIZE_T max_buffer_length;
};

/*ce \brief Forward declaration of the interface implementation
*/
static struct UDP_TransformI MyUdpTransform_fv_Intf;

/*ce \brief Forward declaration of the interface factory implementation
*/
static struct RT_ComponentFactoryI MyUdpTransformFactory_fv_Intf;

/*ce \brief Method to create an instance of MyUdpTransform
*
* \param[in] factory The factory creating this instance
* \param[in] property Generic UDP_Transform properties
*
* \return A pointer to MyUdpTransform on sucess, NULL on failure.
*/
RTI_PRIVATE struct MyUdpTransform*
MyUdpTransform_create(struct MyUdpTransformFactory *factory,
                     const struct UDP_TransformProperty *const property)
{
    struct MyUdpTransform *t;

    OSAPI_Heap_allocate_struct(&t, struct MyUdpTransform);
    if (t == NULL)
    {
        return NULL;
    }

    /* All component instances must initialize the parent using this
    * call.
    */
    RT_Component_initialize(&t->_parent._parent,
                          &MyUdpTransform_fv_Intf._parent,
                          0,
                          (property ? &property->_parent : NULL),
                          NULL);

    t->factory = factory;

    /* Allocate a buffer that is the larger of the send and receive
    * size.
    */
    t->max_buffer_length = property->max_receive_message_size;
    if (property->max_send_message_size > t->max_buffer_length )
    {
        t->max_buffer_length = property->max_send_message_size;
    }
}

```

(continues on next page)

(continued from previous page)

```

    /* Allocate 1 extra byte */
    OSAPI_Heap_allocate_buffer(&t->buffer,t->max_buffer_length+1,
                              OSAPI_ALIGNMENT_DEFAULT);

    if (t->buffer == NULL)
    {
        OSAPI_Heap_free_struct(t);
        t = NULL;
    }

    return t;
}

/*ce \brief Method to delete an instance of MyUdpTransform
 *
 * \param[in] t Transformation instance to delete
 */
RTI_PRIVATE void
MyUdpTransform_delete(struct MyUdpTransform *t)
{
    OSAPI_Heap_free_buffer(t->buffer);
    OSAPI_Heap_free_struct(t);
}

/*ce \brief Method to create a transformation for an destination address
 *
 * \details
 *
 * For each asserted destination rule a transform is created by the transformation
 * instance. This method determines how a UDP payload is transformed before
 * it is sent to an address that matches destination & netmask.
 *
 * \param[in] udptf      UDP Transform instance that creates the transformation
 * \param[out] context   Pointer to a transformation context
 * \param[in] destination Destination address for the transformation
 * \param[in] netmask    The netmask to apply to this destination.
 * \param[in] user_data  The user_data the rule was asserted with
 * \param[in] property   UDP transform specific properties
 * \param[out] ec        User defined error code
 *
 * \return RTI_TRUE on success, RTI_FALSE on failure
 */
RTI_PRIVATE RTI_BOOL
MyUdpTransform_create_destination_transform(
    UDP_Transform_T *const udptf,
    void **const context,
    const struct NETIO_Address *const destination,
    const struct NETIO_Netmask *const netmask,
    void *user_data,
    const struct UDP_TransformProperty *const property,

```

(continues on next page)

(continued from previous page)

```

RTI_INT32 *const ec)
{
    struct MyUdpTransform *self = (struct MyUdpTransform*)udptf;
    UNUSED_ARG(self);
    UNUSED_ARG(destination);
    UNUSED_ARG(user_data);
    UNUSED_ARG(property);
    UNUSED_ARG(ec);
    UNUSED_ARG(netmask);

    /* Save the user-data to determine which transform to apply later */
    *context = (void*)user_data;

    return RTI_TRUE;
}

/*ce \brief Method to delete a transformation for an destination address
 *
 *
 * \param[in]  udptf      UDP Transform instance that created the transformation
 * \param[out] context    Pointer to a transformation context
 * \param[in]  destination Destination address for the transformation
 * \param[in]  netmask    The netmask to apply to this destination.
 * \param[out] ec         User defined error code
 *
 * \return RTI_TRUE on success, RTI_FALSE on failure
 */
RTI_PRIVATE RTI_BOOL
MyUdpTransform_delete_destination_transform(UDP_Transform_T *const udptf,
                                           void *context,
                                           const struct NETIO_Address *const destination,
                                           const struct NETIO_Netmask *const netmask,
                                           RTI_INT32 *const ec)
{
    UNUSED_ARG(udptf);
    UNUSED_ARG(context);
    UNUSED_ARG(destination);
    UNUSED_ARG(ec);
    UNUSED_ARG(netmask);

    return RTI_TRUE;
}

/*ce \brief Method to create a transformation for an source address
 *
 * \details
 *
 * For each asserted source rule a transform is created by the transformation
 * instance. This method determines how a UDP payload is transformed when
 * it is received from an address that matches source & netmask.
 *

```

(continues on next page)

(continued from previous page)

```

* \param[in]  udptf      UDP Transform instance that creates the transformation
* \param[out] context    Pointer to a transformation context
* \param[in]  source     Destination address for the transformation
* \param[in]  netmask    The netmask to apply to this destination.
* \param[in]  user_data  The user_data the rule was asserted with
* \param[in]  property    UDP transform specific properties
* \param[out] ec         User defined error code
*
* \return RTI_TRUE on success, RTI_FALSE on failure
*/
RTI_PRIVATE RTI_BOOL
MyUdpTransform_create_source_transform(UDP_Transform_T *const udptf,
                                     void **const context,
                                     const struct NETIO_Address *const source,
                                     const struct NETIO_Netmask *const netmask,
                                     void *user_data,
                                     const struct UDP_TransformProperty *const property,
                                     RTI_INT32 *const ec)
{
    struct MyUdpTransform *self = (struct MyUdpTransform*)udptf;
    UNUSED_ARG(self);
    UNUSED_ARG(source);
    UNUSED_ARG(user_data);
    UNUSED_ARG(property);
    UNUSED_ARG(ec);
    UNUSED_ARG(netmask);

    *context = (void*)user_data;

    return RTI_TRUE;
}

/*ce \brief Method to delete a transformation for an source address
*
*
* \param[in]  udptf      UDP Transform instance that created the transformation
* \param[out] context    Pointer to a transformation context
* \param[in]  source     Source address for the transformation
* \param[in]  netmask    The netmask to apply to this destination.
* \param[out] ec         User defined error code
*
* \return RTI_TRUE on success, RTI_FALSE on failure
*/
RTI_PRIVATE RTI_BOOL
MyUdpTransform_delete_source_transform(UDP_Transform_T *const udptf,
                                     void *context,
                                     const struct NETIO_Address *const source,
                                     const struct NETIO_Netmask *const netmask,
                                     RTI_INT32 *const ec)
{
    UNUSED_ARG(udptf);

```

(continues on next page)

(continued from previous page)

```

    UNUSED_ARG(context);
    UNUSED_ARG(source);
    UNUSED_ARG(ec);
    UNUSED_ARG(netmask);

    return RTI_TRUE;
}

/*ce \brief Method to transform data based on a source address
 *
 * \param[in]  udptf      UDP_Transform_T that performs the transformation
 * \param[in]  context    Reference to context created by \ref MyUdpTransform_create_
 * \param[in]  source     Source address for the transformation
 * \param[in]  in_packet  The NETIO packet to transform
 * \param[out] out_packet The transformed NETIO packet
 * \param[out] ec         User defined error code
 *
 * \return RTI_TRUE on success, RTI_FALSE on failure
 */
RTI_PRIVATE RTI_BOOL
MyUdpTransform_transform_source(UDP_Transform_T *const udptf,
                               void *context,
                               const struct NETIO_Address *const source,
                               const NETIO_Packet_T *const in_packet,
                               NETIO_Packet_T **out_packet,
                               RTI_INT32 *const ec)
{
    struct MyUdpTransform *self = (struct MyUdpTransform*)udptf;
    char *buf_ptr,*buf_end;
    char *from_buf_ptr,*from_buf_end;
    UNUSED_ARG(context);
    UNUSED_ARG(source);

    *ec = 0;

    /* Assigned the transform buffer to the outgoing packet
     * saving state from the incoming packet. In this case the
     * outgoing length is the same as the incoming. How to buffer
     * is filled in is of no interest to \runtime. All it cares about is
     * where it starts and where it ends.
     */
    if (!NETIO_Packet_initialize_from(
        &self->packet,in_packet,
        self->buffer,self->max_buffer_length,
        0,NETIO_Packet_get_payload_length(in_packet)))
    {
        return RTI_FALSE;
    }

    *out_packet = &self->packet;

```

(continues on next page)



(continued from previous page)

```

    buf_ptr = NETIO_Packet_get_head(&self->packet);
    buf_end = NETIO_Packet_get_tail(&self->packet);
    from_buf_ptr = NETIO_Packet_get_head(in_packet);
    from_buf_end = NETIO_Packet_get_tail(in_packet);

    /* Perform a transformation based on the user-data */
    while (from_buf_ptr < from_buf_end)
    {
        if (context == (void*)1)
        {
            *buf_ptr = ~(*from_buf_ptr);
        }
        else if (context == (void*)2)
        {
            *buf_ptr = (*from_buf_ptr)+1;
        }

        ++buf_ptr;
        ++from_buf_ptr;
    }

    return RTI_TRUE;
}

/*ce \brief Method to transform data based on a destination address
 *
 * \param[in] udptf      UDP_Transform_T that performs the transformation
 * \param[in] context    Reference to context created by \ref MyUdpTransform_create_
 * \param[in] destination Source address for the transformation
 * \param[in] in_packet  The NETIO packet to transform
 * \param[out] packet_out The transformed NETIO packet
 * \param[out] ec        User defined error code
 *
 * \return RTI_TRUE on success, RTI_FALSE on failure
 */
RTI_PRIVATE RTI_BOOL
MyUdpTransform_transform_destination(UDP_Transform_T *const udptf,
                                    void *context,
                                    const struct NETIO_Address *const destination,
                                    const NETIO_Packet_T *const in_packet,
                                    NETIO_Packet_T **packet_out,
                                    RTI_INT32 *const ec)
{
    struct MyUdpTransform *self = (struct MyUdpTransform*)udptf;
    char *buf_ptr,*buf_end;
    char *from_buf_ptr,*from_buf_end;
    UNUSED_ARG(context);
    UNUSED_ARG(destination);

```

(continues on next page)

(continued from previous page)

```

*ec = 0;

if (!NETIO_Packet_initialize_from(
    &self->packet,in_packet,
    self->buffer,8192,
    0,NETIO_Packet_get_payload_length(in_packet)))
{
    return RTI_FALSE;
}

*out_packet = &self->packet;

buf_ptr = NETIO_Packet_get_head(&self->packet);
buf_end = NETIO_Packet_get_tail(&self->packet);
from_buf_ptr = NETIO_Packet_get_head(in_packet);
from_buf_end = NETIO_Packet_get_tail(in_packet);

while (from_buf_ptr < from_buf_end)
{
    if (context == (void*)1)
    {
        *buf_ptr = ~(*from_buf_ptr);
    }
    else if (context == (void*)2)
    {
        *buf_ptr = (*from_buf_ptr)-1;
    }

    ++buf_ptr;
    ++from_buf_ptr;
}

return RTI_TRUE;
}

/*ce \brief Definition of the transformation interface
*/
RTI_PRIVATE struct UDP_TransformI MyUdpTransform_fv_Intf =
{
    RT_COMPONENTI_BASE,
    MyUdpTransform_create_destination_transform,
    MyUdpTransform_create_source_transform,
    MyUdpTransform_transform_source,
    MyUdpTransform_transform_destination,
    MyUdpTransform_delete_destination_transform,
    MyUdpTransform_delete_source_transform
};

/*ce \brief Method called by \runtime to create an instance of transformation
*/
MUST_CHECK_RETURN RTI_PRIVATE RT_Component_T*

```

(continues on next page)

(continued from previous page)

```

MyUdpTransformFactory_create_component(struct RT_ComponentFactory *factory,
                                     struct RT_ComponentProperty *property,
                                     struct RT_ComponentListener *listener)
{
    struct MyUdpTransform *t;
    UNUSED_ARG(listener);

    t = MyUdpTransform_create(
        (struct MyUdpTransformFactory*)factory,
        (struct UDP_TransformProperty*)property);

    return &t->_parent._parent;
}

/*ce \brief Method called by \rttime to delete an instance of transformation
*/
RTI_PRIVATE void
MyUdpTransformFactory_delete_component(
    struct RT_ComponentFactory *factory,
    RT_Component_T *component)
{
    UNUSED_ARG(factory);

    MyUdpTransform_delete((struct MyUdpTransform*)component);
}

/*ce \brief Method called by \rttime when a factory is registered
*/
MUST_CHECK_RETURN RTI_PRIVATE struct RT_ComponentFactory*
MyUdpTransformFactory_initialize(struct RT_ComponentFactoryProperty* property,
                                struct RT_ComponentFactoryListener *listener)
{
    struct MyUdpTransformFactory *fac;
    UNUSED_ARG(property);
    UNUSED_ARG(listener);

    OSAPI_Heap_allocate_struct(&fac, struct MyUdpTransformFactory);

    fac->_parent._factory = &fac->_parent;
    fac->_parent.intf = &MyUdpTransformFactory_fv_Intf;
    fac->property = (struct MyUdpTransformFactoryProperty*)property;

    return &fac->_parent;
}

/*ce \brief Method called by \rttime when a factory is unregistered
*/
RTI_PRIVATE void
MyUdpTransformFactory_finalize(struct RT_ComponentFactory *factory,
                              struct RT_ComponentFactoryProperty **property,
                              struct RT_ComponentFactoryListener **listener)

```

(continues on next page)

(continued from previous page)

```

{
    struct MyUdpTransformFactory *fac =
        (struct MyUdpTransformFactory*)factory;

    UNUSED_ARG(property);
    UNUSED_ARG(listener);

    if (listener != NULL)
    {
        *listener = NULL;
    }

    if (property != NULL)
    {
        *property = (struct RT_ComponentFactoryProperty*)fac->property;
    }

    OSAPI_Heap_free_struct(factory);

    return;
}

/*ce \brief Definition of the factory interface
*/
RTI_PRIVATE struct RT_ComponentFactoryI MyUdpTransformFactory_fv_Intf =
{
    UDP_INTERFACE_INTERFACE_ID,
    MyUdpTransformFactory_initialize,
    MyUdpTransformFactory_finalize,
    MyUdpTransformFactory_create_component,
    MyUdpTransformFactory_delete_component,
    NULL,
    NULL
};

struct RT_ComponentFactoryI*
MyUdpTransformFactory_get_interface(void)
{
    return &MyUdpTransformFactory_fv_Intf;
}

/*ce \brief Method to register this transformation in a registry
*/
RTI_BOOL
MyUdpTransformFactory_register(RT_Registry_T *registry,
                               const char *const name,
                               struct MyUdpTransformFactoryProperty *property)
{
    return RT_Registry_register(registry, name,
                                MyUdpTransformFactory_get_interface(),
                                &property->parent, NULL);
}

```

(continues on next page)

(continued from previous page)

```

}

/*ce \brief Method to unregister this transformation from a registry
*/
RTI_BOOL
MyUdpTransformFactory_unregister(RT_Registry_T *registry,
    const char *const name,
    struct MyUdpTransformFactoryProperty **property)
{
    return RT_Registry_unregister(registry, name,
        (struct RT_ComponentFactoryProperty**)property,
        NULL);
}

/*! @} */

```

Example configuration of rules:

```

#include <stdio.h>
#include <stdlib.h>
#include <string.h>

#include "common.h"

void
MyAppApplication_help(char *appname)
{
    printf("%s [options]\n", appname);
    printf("options:\n");
    printf("-h                - This text\n");
    printf("-domain <id>       - DomainId (default: 0)\n");
    printf("-udp_intf <intf>    - udp interface (no default)\n");
    printf("-peer <address>     - peer address (no default)\n");
    printf("-count <count>      - count (default -1)\n");
    printf("-sleep <ms>         - sleep between sends (default 1s)\n");
    printf("\n");
}

struct MyAppApplication*
MyAppApplication_create(const char *local_participant_name,
    const char *remote_participant_name,
    DDS_Long domain_id, char *udp_intf, char *peer,
    DDS_Long sleep_time, DDS_Long count)
{
    DDS_ReturnCode_t retcode;
    DDS_DomainParticipantFactory *factory = NULL;
    struct DDS_DomainParticipantFactoryQos dpf_qos =
        DDS_DomainParticipantFactoryQos_INITIALIZER;
    struct DDS_DomainParticipantQos dp_qos =
        DDS_DomainParticipantQos_INITIALIZER;

```

(continues on next page)

(continued from previous page)

```

DDS_Boolean success = DDS_BOOLEAN_FALSE;
struct MyAppApplication *application = NULL;
RT_Registry_T *registry = NULL;
struct UDP_InterfaceFactoryProperty *udp_property = NULL;
struct DPDE_DiscoveryPluginProperty discovery_plugin_properties =
    DPDE_DiscoveryPluginProperty_INITIALIZER;
UNUSED_ARG(local_participant_name);
UNUSED_ARG(remote_participant_name);

/* Uncomment to increase verbosity level:
   OSAPILog_set_verbosity(OSAPI_LOG_VERBOSITY_WARNING);
*/
application = (struct MyAppApplication *)malloc(sizeof(struct MyAppApplication));

if (application == NULL)
{
    printf("failed to allocate application\n");
    goto done;
}

application->sleep_time = sleep_time;
application->count = count;

factory = DDS_DomainParticipantFactory_get_instance();

if (DDS_DomainParticipantFactory_get_qos(factory, &dpf_qos) != DDS_RETCODE_OK)
{
    printf("failed to get number of components\n");
    goto done;
}

dpf_qos.resource_limits.max_components = 128;

if (DDS_DomainParticipantFactory_set_qos(factory, &dpf_qos) != DDS_RETCODE_OK)
{
    printf("failed to increase number of components\n");
    goto done;
}

registry = DDS_DomainParticipantFactory_get_registry(
    DDS_DomainParticipantFactory_get_instance());

if (!RT_Registry_register(registry, DDSHST_WRITER_DEFAULT_HISTORY_NAME,
    WHSM_HistoryFactory_get_interface(), NULL, NULL))
{
    printf("failed to register wh\n");
    goto done;
}

if (!RT_Registry_register(registry, DDSHST_READER_DEFAULT_HISTORY_NAME,
    RHSM_HistoryFactory_get_interface(), NULL, NULL))

```

(continues on next page)

(continued from previous page)

```

{
    printf("failed to register rh\n");
    goto done;
}

if (!MyUdpTransformFactory_register(registry, "T0", NULL))
{
    printf("failed to register T0\n");
    goto done;
}

if (!MyUdpTransformFactory_register(registry, "T1", NULL))
{
    printf("failed to register T0\n");
    goto done;
}

/* Configure UDP transport's allowed interfaces */
if (!RT_Registry_unregister(registry, NETIO_DEFAULT_UDP_NAME, NULL, NULL))
{
    printf("failed to unregister udp\n");
    goto done;
}

udp_property = (struct UDP_InterfaceFactoryProperty *)
               malloc(sizeof(struct UDP_InterfaceFactoryProperty));
if (udp_property == NULL)
{
    printf("failed to allocate udp properties\n");
    goto done;
}
*udp_property = UDP_INTERFACE_FACTORY_PROPERTY_DEFAULT;

/* For additional allowed interface(s), increase maximum and length, and
   set interface below:
   */
udp_property->max_send_message_size = 16384;
udp_property->max_message_size = 32768;

if (udp_intf != NULL)
{
    REDA_StringSeq_set_maximum(&udp_property->allow_interface, 1);
    REDA_StringSeq_set_length(&udp_property->allow_interface, 1);
    *REDA_StringSeq_get_reference(&udp_property->allow_interface, 0) =
        DDS_String_dup(udp_intf);
}

/* A rule that says: For payloads received from 192.168.10.* (netmask is
   * 0xffffffff), apply transformation T0.
   */

```

(continues on next page)

(continued from previous page)

```

if (!UDP_TransformRules_assert_source_rule(
    &udp_property->source_rules,
    0xc0a80ae8, 0xffffffff00, "T0", (void*)2))
{
    printf("Failed to assert source rule\n");
    goto done;
}

/* A rule that says: For payloads sent to 192.168.10.* (netmask is
 * 0xffffffff00), apply transformation T0.
 */
if (!UDP_TransformRules_assert_destination_rule(
    &udp_property->destination_rules,
    0xc0a80ae8, 0xffffffff00, "T0", (void*)2))
{
    printf("Failed to assert source rule\n");
    goto done;
}

/* A rule that says: For payloads received from 192.168.20.* (netmask is
 * 0xffffffff00), apply transformation T1.
 */
if (!UDP_TransformRules_assert_source_rule(
    &udp_property->source_rules,
    0xc0a81465, 0xffffffff00, "T1", (void*)1))
{
    printf("Failed to assert source rule\n");
    goto done;
}

/* A rule that says: For payloads received from 192.168.20.* (netmask is
 * 0xffffffff00), apply transformation T1.
 */
if (!UDP_TransformRules_assert_destination_rule(
    &udp_property->destination_rules,
    0xc0a81465, 0xffffffff00, "T1", (void*)1))
{
    printf("Failed to assert source rule\n");
    goto done;
}

if (!RT_Registry_register(registry, NETIO_DEFAULT_UDP_NAME,
    UDP_InterfaceFactory_get_interface(),
    (struct RT_ComponentFactoryProperty*)udp_property, NULL))
{
    printf("failed to register udp\n");
    goto done;
}

DDS_DomainParticipantFactory_get_qos(factory, &dpf_qos);
dpf_qos.entity_factory.autoenable_created_entities = DDS_BOOLEAN_FALSE;

```

(continues on next page)



(continued from previous page)

```

DDS_DomainParticipantFactory_set_qos(factory, &dpf_qos);

if (peer == NULL)
{
    peer = "127.0.0.1"; /* default to loopback */
}

if (!RT_Registry_register(registry,
                          "dpde",
                          DPDE_DiscoveryFactory_get_interface(),
                          &discovery_plugin_properties._parent,
                          NULL))
{
    printf("failed to register dpde\n");
    goto done;
}

if (!RT_ComponentFactoryId_set_name(&dp_qos.discovery.discovery.name, "dpde"))
{
    printf("failed to set discovery plugin name\n");
    goto done;
}

DDS_StringSeq_set_maximum(&dp_qos.discovery.initial_peers, 1);
DDS_StringSeq_set_length(&dp_qos.discovery.initial_peers, 1);
*DDS_StringSeq_get_reference(&dp_qos.discovery.initial_peers, 0) = DDS_String_
↪dup(peer);

DDS_StringSeq_set_maximum(&dp_qos.discovery.enabled_transports, 1);
DDS_StringSeq_set_length(&dp_qos.discovery.enabled_transports, 1);

/* Use network interface 192.168.10.232 for discovery. T0 is used for
 * discovery
 */
*DDS_StringSeq_get_reference(&dp_qos.discovery.enabled_transports, 0) = DDS_String_
↪dup("_udp://192.168.10.232");

DDS_StringSeq_set_maximum(&dp_qos.user_traffic.enabled_transports, 1);
DDS_StringSeq_set_length(&dp_qos.user_traffic.enabled_transports, 1);

/* Use network interface 192.168.20.101 for user-data. T1 is used for
 * this interface.
 */
*DDS_StringSeq_get_reference(&dp_qos.user_traffic.enabled_transports, 0) = DDS_String_
↪dup("_udp://192.168.20.101");

/* if there are more remote or local endpoints, you need to increase these limits */
dp_qos.resource_limits.max_destination_ports = 32;
dp_qos.resource_limits.max_receive_ports = 32;
dp_qos.resource_limits.local_topic_allocation = 1;
dp_qos.resource_limits.local_type_allocation = 1;

```

(continues on next page)

(continued from previous page)

```

dp_qos.resource_limits.local_reader_allocation = 1;
dp_qos.resource_limits.local_writer_allocation = 1;
dp_qos.resource_limits.remote_participant_allocation = 8;
dp_qos.resource_limits.remote_reader_allocation = 8;
dp_qos.resource_limits.remote_writer_allocation = 8;

application->participant =
    DDS_DomainParticipantFactory_create_participant(factory, domain_id,
                                                    &dp_qos, NULL,
                                                    DDS_STATUS_MASK_NONE);

if (application->participant == NULL)
{
    printf("failed to create participant\n");
    goto done;
}

sprintf(application->type_name, "HelloWorld");
retcode = DDS_DomainParticipant_register_type(application->participant,
                                              application->type_name,
                                              HelloWorldTypePlugin_get());

if (retcode != DDS_RETCODE_OK)
{
    printf("failed to register type: %s\n", "test_type");
    goto done;
}

sprintf(application->topic_name, "HelloWorld");
application->topic =
    DDS_DomainParticipant_create_topic(application->participant,
                                       application->topic_name,
                                       application->type_name,
                                       &DDS_TOPIC_QOS_DEFAULT, NULL,
                                       DDS_STATUS_MASK_NONE);

if (application->topic == NULL)
{
    printf("topic == NULL\n");
    goto done;
}

success = DDS_BOOLEAN_TRUE;

done:

if (!success)
{
    if (udp_property != NULL)
    {
        free(udp_property);
    }
}

```

(continues on next page)

(continued from previous page)

```

        free(application);
        application = NULL;
    }

    return application;
}

DDS_ReturnCode_t
MyAppApplication_enable(struct MyAppApplication * application)
{
    DDS_Entity *entity;
    DDS_ReturnCode_t retcode;

    entity = DDS_DomainParticipant_as_entity(application->participant);

    retcode = DDS_Entity_enable(entity);
    if (retcode != DDS_RETCODE_OK)
    {
        printf("failed to enable entity\n");
    }

    return retcode;
}

void
MyAppApplication_delete(struct MyAppApplication *application)
{
    DDS_ReturnCode_t retcode;
    RT_Registry_T *registry = NULL;

    retcode = DDS_DomainParticipant_delete_contained_entities(application->participant);
    if (retcode != DDS_RETCODE_OK)
    {
        printf("failed to delete contained entities (retcode=%d)\n", retcode);
    }

    if (DDS_DomainParticipant_unregister_type(application->participant,
        application->type_name) != HelloWorldTypePlugin_get())
    {
        printf("failed to unregister type: %s\n", application->type_name);
        return;
    }

    retcode = DDS_DomainParticipantFactory_delete_participant(
        DDS_DomainParticipantFactory_get_instance(),
        application->participant);

    if (retcode != DDS_RETCODE_OK)
    {
        printf("failed to delete participant: %d\n", retcode);
        return;
    }
}

```

(continues on next page)

(continued from previous page)

```

}

registry = DDS_DomainParticipantFactory_get_registry(
    DDS_DomainParticipantFactory_get_instance());

if (!RT_Registry_unregister(registry, "dpde", NULL, NULL))
{
    printf("failed to unregister dpde\n");
    return;
}
if (!RT_Registry_unregister(registry, DDSHST_READER_DEFAULT_HISTORY_NAME, NULL, ↵
↵NULL))
{
    printf("failed to unregister rh\n");
    return;
}
if (!RT_Registry_unregister(registry, DDSHST_WRITER_DEFAULT_HISTORY_NAME, NULL, ↵
↵NULL))
{
    printf("failed to unregister wh\n");
    return;
}

free(application);

DDS_DomainParticipantFactory_finalize_instance();
}

```

## Examples

The following examples illustrate how this feature can be used in a system with a mixture of different types of UDP transport configurations.

For the purpose of the examples, the following terminology is used:

- Plain communication – No transformations have been applied.
- Transformed User Data – Only the user-data is transformed, discovery is plain.
- Transformed Discovery – Only the discovery data is transformed, user-data is plain.
- Transformed Data – Both discovery and user-data are transformed. Unless stated otherwise the transformations are different.

A transformation  $T_n$  is a transformation such that an outgoing payload transformed with  $T_n$  can be transformed back to its original state by applying  $T_n$  to the incoming data.

A network interface can be either physical or virtual.

## Plain Communication Between 2 Nodes

In this system two Nodes, A and B, are communicating with plain communication. Node A has one interface, a0, and Node B has one interface, b0.

Node A:

- Register the UDP transport Ua with allow\_interface = a0.
- DomainParticipantQos.transports.enabled\_transports = "Ua"
- DomainParticipantQos.discovery.enabled\_transports = "Ua://"
- DomainParticipantQos.user\_data.enabled\_transports = "Ua://"

Node B:

- Register the UDP transport Ub with allow\_interface = b0.
- DomainParticipantQos.transports.enabled\_transports = "Ub"
- DomainParticipantQos.discovery.enabled\_transports = "Ub://"
- DomainParticipantQos.user\_data.enabled\_transports = "Ub://"

## Transformed User Data Between 2 Nodes

In this system two Nodes, A and B, are communicating with transformed user data. Node A has two interfaces, a0 and a1, and Node B has two interfaces, b0 and b1. Since each node has only one peer, a single transformation is sufficient.

Node A:

- Add a destination transformation T0 to Ua0, indicating that all sent data is transformed with T0.
- Add a source transformation T1 to Ua0, indicating that all received data is transformed with T1.
- Register the UDP transport Ua0 with allow\_interface = a0.
- Register the UDP transport Ua1 with allow\_interface = a1.
- No transformations are registered with Ua1.
- DomainParticipantQos.transports.enabled\_transports = "Ua0","Ua1"
- DomainParticipantQos.discovery.enabled\_transports = "Ua1://"
- DomainParticipantQos.user\_traffic.enabled\_transports = "Ua0://"

Node B:

- Add a destination transformation T1 to Ub0, indicating that all sent data is transformed with T1.

- Add a source transformation T0 to Ub0, indicating that all received data is transformed with T0.
- Register the UDP transport Ub0 with `allow_interface = b0`.
- Register the UDP transport Ub1 with `allow_interface = b1`.
- No transformations are registered with Ub1.
- `DomainParticipantQos.transports.enabled_transports = "Ub0","Ub1"`
- `DomainParticipantQos.discovery.enabled_transports = "Ub1:/"`
- `DomainParticipantQos.user_traffic.enabled_transports = "Ub0:/"`

Ua0 and Ub0 perform transformations and are used for user-data. Ua1 and Ub1 are used for discovery and no transformations takes place.

### Transformed Discovery Data Between 2 Nodes

In this system two Nodes, A and B, are communicating with transformed user data. Node A has two interfaces, a0 and a1, and Node B has two interfaces, b0 and b1. Since each node has only one peer, a single transformation is sufficient.

Node A:

- Add a destination transformation T0 to Ua0, indicating that all sent data is transformed with T0.
- Add a source transformation T1 to Ua0, indicating that all received data is transformed with T1.
- Register the UDP transport Ua0 with `allow_interface = a0`.
- Register the UDP transport Ua1 with `allow_interface = a1`.
- No transformations are registered with Ua1.
- `DomainParticipantQos.transports.enabled_transports = "Ua0","Ua1"`
- `DomainParticipantQos.discovery.enabled_transports = "Ua0:/"`
- `DomainParticipantQos.user_data.enabled_transports = "Ua1:/"`

Node B:

- Add a destination transformation T1 to Ub0, indicating that all sent data is transformed with T1.
- Add a source transformation T0 to Ub0, indicating that all received data is transformed with T0.
- Register the UDP transport Ub0 with `allow_interface = b0`.
- Register the UDP transport Ub1 with `allow_interface = b1`.
- No transformations are registered with Ub1.

- DomainParticipantQos.transports.enabled\_\_transports = "Ub0","Ub1"
- DomainParticipantQos.discovery.enabled\_\_transports = "Ub0://"
- DomainParticipantQos.user\_\_data.enabled\_\_transports = "Ub1://"

Ua0 and Ub0 perform transformations and are used for discovery. Ua1 and Ub1 are used for user-data and no transformation takes place.

### Transformed Data Between 2 Nodes (same transformation)

In this system two Nodes, A and B, are communicating with transformed data using the same transformation for user and discovery data. Node A has one interface, a0, and Node B has one interface, b0.

Node A:

- Add a destination transformation T0 to Ua0, indicating that all sent data is transformed with T0.
- Add a source transformation T1 to Ua0, indicating that all received data is transformed with T1.
- Register the UDP transport Ua0 with allow\_\_interface = a0.
- DomainParticipantQos.transports.enabled\_\_transports = "Ua0"
- DomainParticipantQos.discovery.enabled\_\_transports = "Ua0://"
- DomainParticipantQos.user\_\_data.enabled\_\_transports = "Ua0://"

Node B:

- Add a destination transformation T1 to Ub0, indicating that all sent data is transformed with T1.
- Add a source transformation T0 to Ub0, indicating that all received data is transformed with T0.
- Register the UDP transport Ub0 with allow\_\_interface = b0.
- DomainParticipantQos.transports.enabled\_\_transports = "Ub0"
- DomainParticipantQos.discovery.enabled\_\_transports = "Ub0://"
- DomainParticipantQos.user\_\_data.enabled\_\_transports = "Ub0://"

Ua0 and Ub0 performs transformations and are used for discovery and for user-data.

### Transformed Data Between 2 Nodes (different transformations)

In this system two Nodes, A and B, are communicating with transformed data using different transformations for user and discovery data. Node A has two interfaces, a0 and a1, and Node B has two interfaces, b0 and b1.

Node A:

- Add a destination transformation T0 to Ua0, indicating that all sent data is transformed with T0.
- Add a source transformation T1 to Ua0, indicating that all received data is transformed with T1.
- Add a destination transformation T2 to Ua1, indicating that all sent data is transformed with T2.
- Add a source transformation T3 to Ua1, indicating that all received data is transformed with T3.
- Register the UDP transport Ua0 with `allow_interface = a0`.
- Register the UDP transport Ua1 with `allow_interface = a1`.
- `DomainParticipantQos.transports.enabled_transports = "Ua0","Ua1"`
- `DomainParticipantQos.discovery.enabled_transports = "Ua0://"`
- `DomainParticipantQos.user_data.enabled_transports = "Ua1://"`

Node B:

- Add a destination transformation T1 to Ub0, indicating that all sent data is transformed with T1.
- Add a source transformation T0 to Ub0, indicating that all received data is transformed with T0.
- Add a destination transformation T3 to Ub1, indicating that all sent data is transformed with T3.
- Add a source transformation T2 to Ub1, indicating that all received data is transformed with T2.
- Register the UDP transport Ub0 with `allow_interface = b0`.
- Register the UDP transport Ub1 with `allow_interface = b1`.
- `DomainParticipantQos.transports.enabled_transports = "Ub0","Ub1"`
- `DomainParticipantQos.discovery.enabled_transports = "Ub0://"`
- `DomainParticipantQos.user_data.enabled_transports = "Ub1://"`

Ua0 and Ub0 perform transformations and are used for discovery. Ua1 and Ub1 perform transformations and are used for user-data.



## OS Configuration

In systems with several network interfaces, *Connex Micro* cannot ensure which network interface should be used to send a packet. Depending on the UDP transformations configured, this might be a problem.

To illustrate this problem, let's assume a system with two nodes, A and B. Node A has two network interfaces, a0 and a1, and Node B has two network interfaces, b0 and b1. In this system, Node A is communicating with Node B using a transformation for discovery and a different transformation for user data.

Node A:

- Add a destination transformation T0 to Ua0, indicating that sent data to b0 is transformed with T0.
- Add a source transformation T1 to Ua0, indicating that received data from b0 is transformed with T1.
- Add a destination transformation T2 to Ua1, indicating that sent data to b1 is transformed with T2.
- Add a source transformation T3 to Ua1, indicating that received data from b1 is transformed with T3.
- Register the UDP transport Ua0 with `allow_interface = a0`.
- Register the UDP transport Ua1 with `allow_interface = a1`.
- `DomainParticipantQos.transports.enabled_transports = "Ua0","Ua1"`
- `DomainParticipantQos.discovery.enabled_transports = "Ua0://"`
- `DomainParticipantQos.user_data.enabled_transports = "Ua1://"`

Node B:

- Add a destination transformation T1 to Ub0, indicating that sent data to a0 is transformed with T1.
- Add a source transformation T0 to Ub0, indicating that received data from a0 transformed with T0.
- Add a destination transformation T3 to Ub1, indicating that sent data to a1 is transformed with T3.
- Add a source transformation T2 to Ub1, indicating that received data from a1 transformed with T2.
- Register the UDP transport Ub0 with `allow_interface = b0`.
- Register the UDP transport Ub1 with `allow_interface = b1`.
- `DomainParticipantQos.transports.enabled_transports = "Ub0","Ub1"`
- `DomainParticipantQos.discovery.enabled_transports = "Ub0://"`
- `DomainParticipantQos.user_data.enabled_transports = "Ub1://"`

Node A sends a discovery packet to Node B to interface b0. This packet will be transformed using T0 as specified by Node A's configuration. When this packet is received in Node B, it will be transformed using either T0 or T2 depending on the source address. Node's A OS will use a0 or a1 to send this packet but *Connext Micro* cannot ensure which one will be used. In case the OS sends the packet using a1, the wrong transformation will be applied in Node B.

Some systems have the possibility to configure the source address that should be used when a packet is sent. In POSIX systems, the command `ip route add <string> dev <interface>` can be used.

By typing the command `ip route add < b0 ip >/32 dev a0` in Node A, the OS will send all packets to Node B's b0 IP address using interface a0. This would ensure that the correct transformation is applied in Node B. The same should be done to ensure that user data is sent with the right address `ip route add < b1 ip >/32 dev a1`. Of course, similar configuration is needed in Node B.

### 3.6.10 NETIO Datagram Transport

This section describes the built-in *Connext Micro* Datagram transport and how to configure it.

The built-in Datagram transport (DGRAM) is a generic transport plugin service.

DGRAM is part of the *RTI Connext Micro* core library that is compiled for a specific CPU architecture with a specific compiler. However, the DGRAM transport does not include integration with any particular network stack. Instead, the DGRAM transport provides a simplified interface which can integrate with a variety of different networking technologies.

The DGRAM plugin supports transmission and reception of RTPS messages over a connectionless network link. Note that while the DGRAM transport itself has no knowledge of the underlying network stack, the DGRAM API does not include API related to establishing connections, such as TCP.

#### Registering a Datagram Interface

DGRAM is a *Connext Micro* component that can be registered with *Connext Micro* with `NETIO_DGRAM_InterfaceFactory_register()` as shown below:

The factory gets the registry. The registry registers the Datagram.

```
DDS_DomainParticipantFactory *factory = NULL;
RT_Registry_T *registry = NULL;

factory = DDS_DomainParticipantFactory_get_instance();
registry = DDS_DomainParticipantFactory_get_registry(factory);
```

When a component is registered, the registration takes the DGRAM interface as the 3rd parameter and the properties as the 4th parameter. In general, it is up to the caller to manage the memory for the properties and ensure they are valid as long as the DGRAM transport is registered. There is no guarantee that a component makes a copy.

The DGRAM Interface is a component interfaces the *Connex Micro* core library'. The user is responsible for implementing the `NETIO_DGRAM_InterfaceI`` which integrates with a specific network technology. This struct must be compliant with the `NETIO_DGRAM_InterfaceI` structure.

```
/* Create the DGRAM User Interface property struct */
struct MyDgramInterfaceProperty
{
    RTI_INT32 a_property;
    struct UTEST_Context *setting;
} MyDgramInterfaceProperty = {10,NULL};

/* Example operation */
struct NETIO_Interface*
MyDgramInterface_create_instance(NETIO_Interface_T *upstream,void *property)
{
    /* Perform operations */
    ...
    return myInterface;
}
...

/* Create the DGRAM Interface struct where each member points to it's
 * respective operation */
RTI_PRIVATE struct NETIO_DGRAM_InterfaceI MyDgramInterface =
{
    MyDgramInterface_create_instance,
    MyDgramInterface_get_interface_list,
    MyDgramInterface_release_address,
    MyDgramInterface_resolve_address_udp4,
    MyDgramInterface_send,
    MyDgramInterface_get_route_table,
    MyDgramInterface_bind_address
};
```

The following code snippet shows how to register the DGRAM Interface with new parameters. The Datagram needs to register the DGRAM Interface with a property that has the interface to call:

```
/* Register the transport again, using the builtin name
 */
if (!NETIO_DGRAM_InterfaceFactory_register(registry,
    "name",
    &MyDgramInterface,
    &MyDgramInterfaceProperty))
{
    /* ERROR */
}
```

It should be noted that the Datagram transport can be registered with any name, but all transport QoS policies and initial peers must refer to this name. If a transport is referred to and it does not exist, an error message will be logged.

## Addressing a Datagram Transport

The interface may also set the enabled transports to receive data as follows:

```
struct DDS_DomainParticipantQos dp_qos =
    DDS_DomainParticipantQos_INITIALIZER;

/* Datagram enable transport xyz. A second transport can be added
 * by setting the enabled_transports value to 2 and adding a second
 * transport name. enabled_transport indicates what addresses the entity is
 * listening on.
 */
DDS_StringSeq_set_maximum(&dp_qos.transports.enabled_transports,1);
DDS_StringSeq_set_length(&dp_qos.transports.enabled_transports,1);
*DDS_StringSeq_get_reference(&dp_qos.transports.enabled_transports,0) =
    DDS_String_dup("xyz");

/* Receive discovery traffic on xyz */
DDS_StringSeq_set_maximum(&dp_qos.discovery.enabled_transports,1);
DDS_StringSeq_set_length(&dp_qos.discovery.enabled_transports,1);
*DDS_StringSeq_get_reference(&dp_qos.discovery.enabled_transports,0) =
    DDS_String_dup("xyz://");

/* Receive user-data traffic on xyz. */
DDS_StringSeq_set_maximum(&dp_qos.user_traffic.enabled_transports,1);
DDS_StringSeq_set_length(&dp_qos.user_traffic.enabled_transports,1);
*DDS_StringSeq_get_reference(&dp_qos.user_traffic.enabled_transports,0) =
    DDS_String_dup("xyz://");
```

An address may setup peers to send messages over this interface. For example, interface xyz may set its initial peers as:

```
/* Send discovery data on address 0x0A00020F*/
DDS_StringSeq_set_maximum(&dp_qos.discovery.initial_peers,1);
DDS_StringSeq_set_length(&dp_qos.discovery.initial_peers,1);
*DDS_StringSeq_get_reference(&dp_qos.discovery.initial_peers,0) =
    DDS_String_dup("0x0A00020F");
```

## Datagram UDP Setup

The built-in Datagram transport can support UDP integration.

The registering the built-in Datagram transport for UDP registers differently than the generic Datagram component. Use `UDP_Interface_register()` with UDP properties to create the datagram instance for UDP.

```
struct UDP_InterfaceFactoryProperty udp_property =
    UDP_InterfaceFactoryProperty_INITIALIZER;

/* To enable sharing this property must be set to RTI_FALSE */
udp_property->enable_interface_bind = RTI_TRUE;
```

(continues on next page)

(continued from previous page)

```

/* */
REDA_StringSeq_set_maximum(&udp_property->allow_interface,1);
REDA_StringSeq_set_length(&udp_property->allow_interface,1);
*REDA_StringSeq_get_reference(&udp_property->allow_interface,0) =
    REDA_String_dup(intf);

/* To enable multicast operations the multicast flag and multicast_interface
 * property must be set */
if (is_multicast)
{
    flags |= UDP_INTERFACE_INTERFACE_MULTICAST_FLAG;
    udp_property->multicast_interface = DDS_String_dup(intf);
}
else
{
    /* Set the mutlicast interface to NULL when not used*/
    udp_property->multicast_interface = NULL;
}

/* Add an available interface for UDP */
if (!UDP_InterfaceTable_add_entry(&udp_property->if_table,
    address, netmask, intf_name, flags))

{
    /* error */
}

/* Buffer properties */
udp_property->max_send_buffer_size = MAX_SEND_BUFFER_SIZE;
udp_property->max_receive_buffer_size = MAX_RECV_BUFFER_SIZE;
udp_property->max_message_size = MAX_RECV_BUFFER_SIZE;

/* Register the datagram */
if(!UDP_Interface_register(registry, "_udp", udp_property))
{
    /* error */
}

```

Enabled transports can be configured with “\_udp://”. This will use all interfaces. Enabling a UDP is similar to generic addressing:

```

struct DDS_DomainParticipantQos dp_qos =
    DDS_DomainParticipantQos_INITIALIZER;

DDS_StringSeq_set_maximum(&dp_qos.transports.enabled_transports,1);
DDS_StringSeq_set_length(&dp_qos.transports.enabled_transports,1);
DDS_StringSeq_set_maximum(&dp_qos.discovery.enabled_transports,1);
DDS_StringSeq_set_length(&dp_qos.discovery.enabled_transports,1);
DDS_StringSeq_set_maximum(&dp_qos.user_traffic.enabled_transports,1);
DDS_StringSeq_set_length(&dp_qos.user_traffic.enabled_transports,1);

```

(continues on next page)

(continued from previous page)

```

/* This only requires the transport name */
*DDS_StringSeq_get_reference(&dp_qos.transports.enabled_transports,0) =
    DDS_String_dup("_udp");

/* _udp:// indicates to use all available locators */
*DDS_StringSeq_get_reference(&dp_qos.discovery.enabled_transports,0) =
    DDS_String_dup("_udp://");

/* _udp://10.10.0.1 would indicate to use only that address */
*DDS_StringSeq_get_reference(&dp_qos.user_traffic.enabled_transports,0) =
    DDS_String_dup("_udp://");

```

### Datagram Shared flag

*RTI Connext Micro* uses Locators to specify transport addresses to send and receive data. A Locator consists of a kind, port, and a transport address. The kind indicates the type of transport, such as UDPv4, the port is used to reach a DDS DomainParticipant and the address is used to reach the destination transport. *RTI Connext Micro* can work with two different types of transports, one that uses shared ports and one that does not.

When a transport uses shared ports it means it does not matter which transport address a message was received on, only the port matters. For example, if a computer has two network interfaces A and B and is listening for messages on port P, it does not matter if the message is received on A or B. That is, as long as the message is received on any network interface capable of receiving on port P, the message is accepted.

When a transport does not use shared ports it means it does matter which transport address a message was received on. For example, if a computer has two network interfaces A and B and is listening for messages on port P, but has only specified that the A should receive on port P, then messages received on interface B and port P are ignored.

*RTI Connext Micro* support this flag on per *RTI Connext Micro* transport basis. It is important to note that when a message is accepted, it is routed to all relevant DDS datareaders and datawriters. Thus, this feature cannot be used to control that some DDS topics should only be accepted when received on a specific transport interface. However, this feature could be useful to allow different DDS DomainParticipants to use the same port, but with different network interfaces.

### User Interface

[NETIO\\_DGRAM\\_InterfaceFactory\\_register\(\)](#) registers a user interface structure that is passed in via `user_intf`. The *DomainParticipant* utilizes these functions for network operations, such as creating a Datagram interface instance and getting the interface list.

Table 3.1: Structure for the User Interface

Interface Attribute	Description
<code>create_instance</code>	Creates an instance of the <code>NETIO_DGRAM</code> interface.
<code>delete_instance</code>	Deletes an instance of the <code>NETIO_DGRAM</code> interface.
<code>get_interface_list</code>	Reads the available interfaces from the <code>NETIO_DGRAM</code> interface.
<code>release_address</code>	Instructs the <code>NETIO_DGRAM</code> interface to stop listening for messages on the source address.
<code>resolve_address</code>	Instructs the <code>NETIO_DGRAM</code> interface to determine if the address string is valid.
<code>send</code>	Instructs the <code>NETIO_DGRAM</code> interface to send a message.
<code>get_route_table</code>	Instructs the <code>NETIO_DGRAM</code> interface <code>netio_intf</code> to return a sequence of address and netmask pairs that this interface can send to.
<code>bind_address</code>	Instructs the <code>NETIO_DGRAM</code> interface to listen for messages on the source address.

## 3.7 Discovery

This section discusses the implementation of discovery plugins in *RTI Connext Micro*. For a general overview of discovery in *RTI Connext Micro*, see *What is Discovery?*.

*Connext Micro* discovery traffic is conducted through transports. Please see the *Transports* section for more information about registering and configuring transports.

### 3.7.1 What is Discovery?

Discovery is the behind-the-scenes way in which *RTI Connext Micro* objects (*DomainParticipants*, *DataWriters*, and *DataReaders*) on different nodes find out about each other. Each *DomainParticipant* maintains a database of information about all the active *DataReaders* and *DataWriters* that are in the same DDS domain. This database is what makes it possible for *DataWriters* and *DataReaders* to communicate. To create and refresh the database, each application follows a common discovery process.

This section describes the default discovery mechanism known as the Simple Discovery Protocol, which includes two phases: *Simple Participant Discovery* and *Simple Endpoint Discovery*.

The goal of these two phases is to build, for each *DomainParticipant*, a complete picture of all the entities that belong to the remote participants that are in its peers list. The peers list is the list of nodes with which a participant may communicate. It starts out the same as the `initial_peers` list that you configure in the [DISCOVERY](#) QosPolicy. If the `accept_unknown_peers` flag in that same QosPolicy is `TRUE`, then other nodes may also be added as they are discovered; if it is `FALSE`, then the peers list will match the `initial_peers` list, plus any peers added using the *DomainParticipant's* `add_peer()` operation.

The following section discusses how *Connext Micro* objects on different nodes find out about each other using the default Simple Discovery Protocol (SDP). It describes the sequence of messages

that are passed between *Connext Micro* on the sending and receiving sides.

The discovery process occurs automatically, so you do not have to implement any special code. For more information about advanced topics related to Discovery, please refer to [the Discovery chapter](#) in the [RTI Connext DDS Core Libraries User's Manual](#) (available [here](#) if you have Internet access).

### Simple Participant Discovery

This phase of the Simple Discovery Protocol is performed by the Simple Participant Discovery Protocol (SPDP).

During the Participant Discovery phase, *DomainParticipants* learn about each other. The *DomainParticipant's* details are communicated to all other *DomainParticipants* in the same DDS domain by sending participant declaration messages, also known as participant *DATA* submessages. The details include the *DomainParticipant's* unique identifying key (GUID or Globally Unique ID described below), transport locators (addresses and port numbers), and QoS. These messages are sent on a periodic basis using best-effort communication.

*Participant DATAs* are sent periodically to maintain the liveliness of the *DomainParticipant*. They are also used to communicate changes in the *DomainParticipant's* QoS. Only changes to QoS Policies that are part of the *DomainParticipant's* built-in data need to be propagated.

When receiving remote participant discovery information, *RTI Connext Micro* determines if the local participant matches the remote one. A 'match' between the local and remote participant occurs only if the local and remote participant have the same Domain ID and Domain Tag. This matching process occurs as soon as the local participant receives discovery information from the remote one. If there is no match, the discovery *DATA* is ignored, resulting in the remote participant (and all its associated entities) not being discovered.

When a *DomainParticipant* is deleted, a participant *DATA (delete)* submessage with the *DomainParticipant's* identifying GUID is sent.

The GUID is a unique reference to an entity. It is composed of a GUID prefix and an Entity ID. By default, the GUID prefix is calculated from the IP address and the process ID. The entityID is set by *Connext Micro* (you may be able to change it in a future version).

Once a pair of remote participants have discovered each other, they can move on to the Endpoint Discovery phase, which is how *DataWriters* and *DataReaders* find each other.

### Simple Endpoint Discovery

This phase of the Simple Discovery Protocol is performed by the Simple Endpoint Discovery Protocol (SEDP).

During the Endpoint Discovery phase, *RTI Connext Micro* matches *DataWriters* and *DataReaders*. Information (GUID, QoS, etc.) about your application's *DataReaders* and *DataWriters* is exchanged by sending publication/subscription declarations in *DATA* messages that we will refer to as *publication DATAs* and *subscription DATAs*. The Endpoint Discovery phase uses reliable communication.



These declaration or DATA messages are exchanged until each *DomainParticipant* has a complete database of information about the participants in its peers list and their entities. Then the discovery process is complete and the system switches to a steady state. During steady state, *participant DATAs* are still sent periodically to maintain the liveness status of participants. They may also be sent to communicate QoS changes or the deletion of a *DomainParticipant*.

When a remote *DataWriter/DataReader* is discovered, *Connext Micro* determines if the local application has a matching *DataReader/DataWriter*. A ‘match’ between the local and remote entities occurs only if the *DataReader* and *DataWriter* have the same *Topic*, same data type, and compatible QoS Policies. Furthermore, if the *DomainParticipant* has been set up to ignore certain *DataWriters/DataReaders*, those entities will not be considered during the matching process.

This ‘matching’ process occurs as soon as a remote entity is discovered, even if the entire database is not yet complete: that is, the application may still be discovering other remote entities.

A *DataReader* and *DataWriter* can only communicate with each other if each one’s application has hooked up its local entity with the matching remote entity. That is, both sides must agree to the connection.

Please refer to [the section on Discovery Implementation](#) in the [RTI Connext DDS Core Libraries User’s Manual](#) for more details about the discovery process (available [here](#) if you have Internet access).

### 3.7.2 Configuring Participant Discovery Peers

An *RTI Connext Micro DomainParticipant* must be able to send participant discovery announcement messages for other *DomainParticipants* to discover itself, and it must receive announcements from other *DomainParticipants* to discover them.

To do so, each *DomainParticipant* will send its discovery announcements to a set of locators known as its peer list, where a peer is the transport locator of one or more potential other *DomainParticipants* to discover.

#### peer\_desc\_string

A peer descriptor string of the [initial\\_peers](#) string sequence conveys the interface and address of the locator to which to send, as well as the indices of participants to which to send. For example:

```
DDS_StringSeq_set_maximum(&dp_qos.discovery.initial_peers, 3);
DDS_StringSeq_set_length(&dp_qos.discovery.initial_peers, 3);

*DDS_StringSeq_get_reference(&dp_qos.discovery.initial_peers, 0) =
    DDS_String_dup("_udp://239.255.0.1");

*DDS_StringSeq_get_reference(&dp_qos.discovery.initial_peers, 1) =
    DDS_String_dup("[1-4]@_udp://10.10.30.101");

*DDS_StringSeq_get_reference(&dp_qos.discovery.initial_peers, 2) =
    DDS_String_dup("[2]@_udp://10.10.30.102");
```

The peer descriptor format is:

```
[index@] [interface://] address
```

Remember that every *DomainParticipant* has a participant index that is unique within a DDS domain. The participant index (also referred to as the participant ID), together with the DDS domain ID, is used to calculate the network port on which *DataReaders* of that participant will receive messages. Thus, by specifying the participant index, or a range of indices, for a peer locator, that locator becomes a port to which messages will be sent only if addressed to the entities of a particular *DomainParticipant*. Specifying indices restricts the number of participant announcements sent to a locator where other *DomainParticipants* exist and, thus, should be considered to minimize network bandwidth usage.

In the above example, the first peer, “\_udp://239.255.0.1,” has the default UDPv4 multicast peer locator. Note that there is no [index@] associated with a multicast locator.

The second peer, “[1-4]@\_udp://10.10.30.101,” has a unicast address. It also has indices in brackets, [1-4]. These represent a range of participant indices, 1 through 4, to which participant discovery messages will be sent.

Lastly, the third peer, “[2]@\_udp://10.10.30.102,” is a unicast locator to a single participant with index 2.

### 3.7.3 Configuring Initial Peers and Adding Peers

[DiscoveryQosPolicy\\_initial\\_peers](#) is the list of peers a *DomainParticipant* sends its participant announcement messages, when it is enabled, as part of the discovery process.

[DiscoveryQosPolicy\\_initial\\_peers](#) is an empty sequence by default, so while [DiscoveryQosPolicy\\_enabled\\_transports](#) by default includes the DDS default loopback and multicast (239.255.0.1) addresses, [initial\\_peers](#) must be configured to include them.

Peers can also be added to the list, before and after a *DomainParticipant* has been enabled, by using [DomainParticipant\\_add\\_peer](#).

The *DomainParticipant* will start sending participant announcement messages to the new peer as soon as it is enabled.

### 3.7.4 Discovery Plugins

When a *DomainParticipant* receives a participant discovery message from another *DomainParticipant*, it will engage in the process of exchanging information of user-created *DataWriter* and *DataReader* endpoints.

*RTI Connext Micro* provides two ways of determining endpoint information of other *DomainParticipants*: *Dynamic Discovery Plugin* and *Static Discovery Plugin*.

## Dynamic Discovery Plugin

Dynamic endpoint discovery uses builtin discovery *DataWriters* and *DataReader* to exchange messages about user created *DataWriter* and *DataReaders*. A *DomainParticipant* using dynamic participant, dynamic endpoint (DPDE) discovery will have a pair of builtin *DataWriters* for sending messages about its own user created *DataWriters* and *DataReaders*, and a pair of builtin *DataReaders* for receiving messages from other *DomainParticipants* about their user created *DataWriters* and *DataReaders*.

Given a *DomainParticipant* with a user *DataWriter*, receiving an endpoint discovery message for a user *DataReader* allows the *DomainParticipant* to get the type, topic, and QoS of the *DataReader* that determine whether the *DataReader* is a match. When a matching *DataReader* is discovered, the *DataWriter* will include that *DataReader* and its locators as destinations for its subsequent writes.

---

**Note:** *RTI Connext* uses the acronyms SPDP and SEDP to distinguish between the two phases of Simple Discovery: participant and endpoint phases (see [Discovery in the Core Libraries User's Manual](#)). *RTI Connext Micro* uses the acronyms DPSE and DPDE to distinguish between the static and dynamic endpoint discovery plugins available in *RTI Connext Micro*. The DPSE plugin implements the SPDP protocol and DPDE implements the SPDP and SEDP protocol.

---

## Static Discovery Plugin

Static endpoint discovery uses function calls to statically assert information about remote endpoints belonging to remote *DomainParticipants*. An application with a *DomainParticipant* using dynamic participant, static endpoint (DPSE) discovery has control over which endpoints belonging to particular remote *DomainParticipants* are discoverable.

Whereas dynamic endpoint-discovery can establish matches for all endpoint-discovery messages it receives, static endpoint-discovery establishes matches only for the endpoint that have been asserted programmatically.

With DPSE, a user needs to know *a priori* the configuration of the entities that will need to be discovered by its application. The user must know the names of all *DomainParticipants* within the DDS domain and the exact QoS of the remote *DataWriters* and *DataReaders*.

---

**Note:** *RTI Connext* uses the acronyms SPDP and SEDP to distinguish between the two phases of Simple Discovery: participant and endpoint phases (see [Discovery in the Core Libraries User's Manual](#)). *RTI Connext Micro* uses the acronyms DPSE and DPDE to distinguish between the static and dynamic endpoint discovery plugins available in *RTI Connext Micro*. The DPSE plugin implements the SPDP protocol and DPDE implements the SPDP and SEDP protocol.

---

Please refer to the [C API Reference](#) and [C++ API Reference](#) for the following remote entity assertion APIs:

- [DPSE\\_RemoteParticipant\\_assert](#)

- [DPSE\\_RemotePublication\\_assert](#)
- [DPSE\\_RemoteSubscription\\_assert](#)

## Remote Participant Assertion

Given a local *DomainParticipant*, static discovery requires first the names of remote *DomainParticipants* to be asserted, in order for endpoints on them to match. This is done by calling [DPSE\\_RemoteParticipant\\_assert](#) with the name of a remote *DomainParticipant*. The name must match the name contained in the participant discovery announcement produced by that *DomainParticipant*. This has to be done reciprocally between two *DomainParticipants* so that they may discover one another.

For example, a *DomainParticipant* has entity name “participant\_1”, while another *DomainParticipant* has name “participant\_2.” participant\_1 should call [DPSE\\_RemoteParticipant\\_assert](#)(“participant\_2”) in order to discover participant\_2. Similarly, participant\_2 must also assert participant\_1 for discovery between the two to succeed.

```
/* participant_1 is asserting (remote) participant_2 */
retcode = DPSE_RemoteParticipant_assert(participant_1,
                                         "participant_2");

if (retcode != DDS_RETCODE_OK) {
    printf("participant_1 failed to assert participant_2\n");
    goto done;
}
```

## Remote Publication and Subscription Assertion

Next, a *DomainParticipant* needs to assert the remote endpoints it wants to match that belong to an already asserted remote *DomainParticipant*. The endpoint assertion function is used, specifying an argument which contains all the QoS and configuration of the remote endpoint. Where [DPDE](#) gets remote endpoint QoS information from received endpoint-discovery messages, in [DPSE](#), the remote endpoint’s QoS must be configured locally. With remote endpoints asserted, the *DomainParticipant* then waits until it receives a participant discovery announcement from an asserted remote *DomainParticipant*. Once received that, all endpoints that have been asserted for that remote *DomainParticipant* are considered discovered and ready to be matched with local endpoints.

Assume participant\_1 contains a *DataWriter*, and participant\_2 has a *DataReader*, both communicating on topic HelloWorld. participant\_1 needs to assert the *DataReader* in participant\_2 as a remote subscription. The remote subscription data passed to the operation must match exactly the QoS actually used by the remote *DataReader*:

```
/* Set participant_2's reader's QoS in remote subscription data */
rem_subscription_data.key.value[DDS_BUILTIN_TOPIC_KEY_OBJECT_ID] = 200;
rem_subscription_data.topic_name = DDS_String_dup("Example HelloWorld");
rem_subscription_data.type_name = DDS_String_dup("HelloWorld");
rem_subscription_data.reliability.kind = DDS_RELIABLE_RELIABILITY_QOS;
```

(continues on next page)

(continued from previous page)

```

/* Assert reader as a remote subscription belonging to (remote) participant_2 */
retcode = DPSE_RemoteSubscription_assert(participant_1,
                                         "participant_2",
                                         &rem_subscription_data,
                                         HelloWorld_get_key_kind(HelloWorldTypePlugin_
↳get(), NULL));
if (retcode != DDS_RETCODE_OK)
{
    printf("failed to assert remote subscription\n");
    goto done;
}

```

Reciprocally, participant\_2 must assert participant\_1's *DataWriter* as a remote publication, also specifying matching QoS parameters:

```

/* Set participant_1's writer's QoS in remote publication data */
rem_publication_data.key.value[DDS_BUILTIN_TOPIC_KEY_OBJECT_ID] = 100;
rem_publication_data.key.value.topic_name = DDS_String_dup("Example HelloWorld");
rem_publication_data.key.value.type_name = DDS_String_dup("HelloWorld");
rem_publication_data.key.value.reliability.kind = DDS_RELIABLE_RELIABILITY_QOS;

/* Assert writer as a remote publication belonging to (remote) participant_1 */
retcode = DPSE_RemotePublication_assert(participant_2,
                                         "participant_1",
                                         &rem_publication_data,
                                         HelloWorld_get_key_kind(HelloWorldTypePlugin_
↳get(), NULL));
if (retcode != DDS_RETCODE_OK)
{
    printf("failed to assert remote publication\n");
    goto done;
}

```

When participant\_1 receives a participant discovery message from participant\_2, it is aware of participant\_2, based on its previous assertion, and it knows participant\_2 has a matching *DataReader*, also based on the previous assertion of the remote endpoint. It therefore establishes a match between its *DataWriter* and participant\_2's *DataReader*. Likewise, participant\_2 will match participant\_1's *DataWriter* with its local *DataRead*, upon receiving one of participant\_1's participant discovery messages.

Note, with [DPSE](#), there is no runtime check of QoS consistency between *DataWriters* and *DataReaders*, because no endpoint discovery messages are exchanged. This makes it extremely important that users of [DPSE](#) ensure that the QoS set for a local *DataWriter* and *DataReader* is the same QoS being used by another *DomainParticipant* to assert it as a remote *DataWriter* or *DataReader*.

### 3.7.5 DomainParticipant Discovery by Name

If a *DomainParticipant* is restarted after an ungraceful shutdown, other *DomainParticipants* will not remove the *DomainParticipant*'s discovery information until its `lease_duration` has expired. Typically, you can reduce the `lease_duration` or increase the resource limits to avoid running out of resources to discover restarted *DomainParticipants*. However, this may not be a feasible solution, since it increases network load and memory usage.

To mitigate these issues, *Connext Micro* can be configured to discover *DomainParticipants* by name instead of by GUID. This speeds up the discovery process because the restarted *DomainParticipant* can immediately replace the old *DomainParticipant* without waiting for the old one's `lease_duration` to expire.

To enable discovery by name for a *DomainParticipant*, do the following:

1. Set a globally unique `participant_name` in the *DomainParticipant*'s [DomainParticipantQos](#).
2. Set the `enable_participant_discovery_by_name` property to `DDS_BOOLEAN_TRUE` in the *DomainParticipant*'s [DISCOVERY](#) QosPolicy.

The *DomainParticipant* can now be rediscovered smoothly by other remote *DomainParticipants* after a restart.

**Attention:** When using discovery by name, a *DomainParticipant* must be restarted with a different GUID from its previous GUID.

**Attention:** To avoid undefined behavior while using discovery by name, make sure that each *DomainParticipant* in a given domain has a unique `participant_name`.

### 3.7.6 Queueing Discovery Messages

During a system restart or partial restart, the discovery process also restarts as *DomainParticipants* come back up. During this period, it is possible that each *DomainParticipant* is in a different phase. For example:

- Some *DomainParticipants* may have terminated gracefully, but have not started up again.
- Some *DomainParticipants* may have been abruptly taken off-line; e.g., during a power cycle.
- Some *DomainParticipants* may have already restarted.
- Some *DomainParticipants* may not have restarted yet, or will not restart.

During this period it is possible for a *DomainParticipant* to temporarily exceed its resource limits. For example, a *DomainParticipant* that was not restarted may store discovery information for a *DomainParticipant* that was ungracefully shut down because its lease duration has not expired, and **also** discover the restarted *DomainParticipant*.

*Connext Micro* will normally discard a discovery message if it receives an endpoint discovery message and lacks sufficient resources to store the endpoint. *Connext Micro* will (by default) acknowledge the message as received, but discard the discovery information. This can lead to a mismatch in discovery states between *DomainParticipants*. However, as lease durations expire on restarted *DomainParticipants*, resources become available to discover new *DomainParticipants*.

To mitigate this temporary lack of resources, you can set the `DomainParticipantQos.discovery.enable_endpoint_discovery_queue` property can be set to `DDS_BOOLEAN_TRUE`. Endpoint discovery messages will then instead be queued for later processing when resources become available.

**Attention:** This new field assumes that there are sufficient resources available for discovery information and that the lack of resources is temporary (such as during a system restart). Setting this value to `DDS_BOOLEAN_TRUE` without sufficient resources may cause undefined behavior.

## 3.8 User Discovery Data

### 3.8.1 Introduction

User Discovery Data is a feature of *Connext Micro* that provides areas where your application can store additional information related to DDS *Entities*. How this information is used will be up to user code; *Connext Micro* distributes this information to other applications as part of the discovery process, but *Connext Micro* does not interpret the information. Use cases are usually application-to-application identification, authentication, authorization, and encryption.

There are three User Discovery Data QoS policies:

- [USER\\_DATA](#): associated with *DomainParticipants*, *DataWriters*, and *DataReaders*.
- [TOPIC\\_DATA](#): associated with a *Topic*.
- [GROUP\\_DATA](#): associated with a *Publisher* or *Subscriber*.

**Warning:** These QoS policies must be specified when an entity is created and cannot be modified at runtime.

### 3.8.2 Resource Limits

Before these QoS policies can be used, the [DomainParticipantResourceLimitsQosPolicy](#) must be configured to set the maximum length of each kind of data which will be used. These settings are listed below:

- `participant_user_data_max_length`
- `topic_data_max_length`
- `publisher_group_data_max_length`



- `subscriber_group_data_max_length`
- `writer_user_data_max_length`
- `reader_user_data_max_length`

These policy settings limit the length of the data field, and must be configured to the same values for all *DomainParticipants* in the same DDS domain. Attempting to create an entity with user discovery data larger than the corresponding QoS setting will cause creation of that entity to fail. Similarly, discovering remote entities will fail if those entities have user discovery data larger than the QoS policy setting.

Memory usage by the discovery data will be directly affected by the maximum length of each type of data because *Connext Micro* will pre-allocate all of the memory necessary to store received data. Setting the maximum length appropriately will limit the memory usage of this feature.

*Connext Micro* also adds settings to further optimize memory usage via the `max_count` resource limit options, listed below:

- `participant_user_data_max_count`
- `topic_data_max_count`
- `publisher_group_data_max_count`
- `subscriber_group_data_max_count`
- `writer_user_data_max_count`
- `reader_user_data_max_count`

These options limit the number of unique data of a given type. Data from local and discovered entities both contribute toward this limit.

These maximums will default to `DDS_SIZE_AUTO`, in which case *Connext Micro* will generate an appropriate value based on other QoSes to ensure that discovery will always succeed. However, if you know the maximum number of unique data that will exist in the domain, setting these options accordingly can reduce the total amount of memory allocated.

**Warning:** Once the `max_count` limit of unique data has been reached for a given entity type, discovery of entities with additional unique data will fail.

### 3.8.3 Propagating User Discovery Data

When using *Dynamic Discovery Plugin* (DPDE), the information associated with all entities is automatically passed between applications during discovery using builtin topics. When using *Static Discovery Plugin* (DPSE), only the `USER_DATA` associated with a participant will be automatically passed between applications; for remote *DataReaders* and *DataWriters*, the associated `USER_DATA`, `TOPIC_DATA`, or `GROUP_DATA` must be asserted with remote publications and subscriptions. (See *Accessing User Discovery Data* below.)

For example, to assert `USER_DATA` associated with a remote subscription in DPSE:



```

struct DDS_SubscriptionBuiltinTopicData rem_subscription_data =
    DDS_SubscriptionBuiltinTopicData_INITIALIZER;

/* Set Reader's protocol.rtps_object_id */
rem_subscription_data.key.value[DDS_BUILTIN_TOPIC_KEY_OBJECT_ID] = 200;

rem_subscription_data.topic_name = DDS_String_dup("Example HelloWorld");
rem_subscription_data.type_name = DDS_String_dup("HelloWorld");

rem_subscription_data.reliability.kind = DDS_RELIABLE_RELIABILITY_QOS;

/* Set USER_DATA */
DDS_OctetSeq_set_maximum(&rem_subscription_data.user_data.value, 2);
DDS_OctetSeq_set_length(&rem_subscription_data.user_data.value, 2);
*DDS_OctetSeq_get_reference(&rem_subscription_data.user_data.value, 0) = 0xAA;
*DDS_OctetSeq_get_reference(&rem_subscription_data.user_data.value, 1) = 0xBB;

retcode = DPSE_RemoteSubscription_assert(participant,
                                         "Participant_2",
                                         &rem_subscription_data,
                                         HelloWorld_get_key_kind(HelloWorldTypePlugin_
↪get(),
                                         NULL));

if (retcode != DDS_RETCODE_OK)
{
    /* failure */
}

```

### 3.8.4 Accessing User Discovery Data

Whether using DPDE or DPSE, the [USER\\_DATA](#), [TOPIC\\_DATA](#), and [GROUP\\_DATA](#) is propagated with the information about remote *DomainParticipants*, *DataWriters*, and *DataReaders*. For *DomainParticipants*, the associated [USER\\_DATA](#) can be accessed through [ParticipantBuiltinTopicData](#). For *DataWriters* and *DataReaders*, the associated [USER\\_DATA](#), [TOPIC\\_DATA](#), and [GROUP\\_DATA](#) can be accessed through the [PublicationBuiltinTopicData](#) and [SubscriptionBuiltinTopicData](#) respectively.

For *DomainParticipants*, the discovery information for discovered participants can be accessed through the `get_discovered` APIs:

- [DDS\\_DomainParticipant\\_get\\_discovered\\_participants](#)
- [DDS\\_DomainParticipant\\_get\\_discovered\\_participant\\_data](#)

For *DataReaders* and *DataWriters*, the information on matched entities can be retrieved through the `get_matched` APIs:

- [DDS\\_DataWriter\\_get\\_matched\\_subscriptions](#)
- [DDS\\_DataWriter\\_get\\_matched\\_subscription\\_data](#)
- [DDS\\_DataReader\\_get\\_matched\\_publications](#)

- [DDS\\_DataReader\\_get\\_matched\\_publication\\_data](#)

Note that these APIs will perform a copy into the provided data sample. If the provided sample does not have enough memory to store the data, additional memory will be allocated to fit it. This memory can instead be pre-allocated with the corresponding `initialize_from_qos` function in C. If a sample is pre-allocated based on the configured QoS, then no additional memory will need to be allocated to perform the copy.

The C++ API does not support `initialize_from_qos`. The default constructor initializes memory to the maximum value, except for [USER\\_DATA](#), [TOPIC\\_DATA](#), and [GROUP\\_DATA](#), which it sets to an empty sequence. The overloaded constructor accepts a *DomainParticipant* and initializes the memory according to the resource limits, including memory for [USER\\_DATA](#), [TOPIC\\_DATA](#), and [GROUP\\_DATA](#).

For example, to retrieve information on discovered participants:

```
struct DDS_InstanceHandleSeq handles = DDS_SEQUENCE_INITIALIZER;
struct DDS_InstanceHandle handle;
struct DDS_DomainParticipantQos dp_qos =
    DDS_DomainParticipantQos_INITIALIZER;
struct DDS_ParticipantBuiltinTopicData dp_data =
    DDS_ParticipantBuiltinTopicData_INITIALIZER;

/* Pre-allocate memory for discovery data based on participant's QoS */
retcode = DDS_DomainParticipant_get_qos(participant, &dp_qos);
if (retcode != DDS_RETCODE_OK)
{
    /* failure */
}
if (!DDS_ParticipantBuiltinTopicData_initialize_from_qos(&dp_data, &dp_qos))
{
    /* failure */
}

/* Get instance handles of discovered participants */
retcode = DDS_DomainParticipant_get_discovered_participants(participant, &handles);
if (retcode != DDS_RETCODE_OK)
{
    /* failure */
}

/* For each handle, get the discovery data */
for (DDS_Long j = 0; j < DDS_InstanceHandleSeq_get_length(&handles); ++j)
{
    handle = DDS_InstanceHandleSeq_get_reference(&handles, j);
    if (handle == NULL)
    {
        /* failure */
    }

    retcode = DDS_DomainParticipant_get_discovered_participant_data(participant,
                                                                    &dp_data,
                                                                    handle);
}
```

(continues on next page)

(continued from previous page)

```

if (retcode != DDS_RETCODE_OK)
{
    /* failure */
}
else
{
    /* Discovered participant USER_DATA can be accessed in dp_data.user_data */
}
}

```

### 3.8.5 QoS Policies

#### USER\_DATA

This QoS Policy provides an area where your application can store additional information related to a *DomainParticipant*, *DataWriter*, or *DataReader*.

You will need to access the value of [USER\\_DATA](#) through [ParticipantBuiltinTopicData](#), [PublicationBuiltinTopicData](#) or [SubscriptionBuiltinTopicData](#). (See *Accessing User Discovery Data*.)

The structure for the [USER\\_DATA](#) QoSPolicy includes just one field, as seen in Table 3.2. The field is a sequence of octets that translates to a contiguous buffer of bytes whose contents and length are set by the user. The maximum size for the data is set in the [DomainParticipantResourceLimitsQoSPolicy](#). (See *Resource Limits*.)

Table 3.2: DDS\_UserDataQoSPolicy

Type	Field Name	Description
DDS_OctetSeq	value	Empty by default

#### TOPIC\_DATA

This QoS Policy provides an area where your application can store additional information related to the *Topic*.

Currently, [TOPIC\\_DATA](#) of the associated *Topic* is only propagated with the information that declares a *DataWriter* or *DataReader*. Thus, you will need to access the value of [TOPIC\\_DATA](#) through [PublicationBuiltinTopicData](#) or [SubscriptionBuiltinTopicData](#). (See *Accessing User Discovery Data*.)

The structure for the [TOPIC\\_DATA](#) QoSPolicy includes just one field, as seen in Table 3.3. The field is a sequence of octets that translates to a contiguous buffer of bytes whose contents and length are set by the user. The maximum size for the data is set in the [DomainParticipantResourceLimitsQoSPolicy](#). (See *Resource Limits*.)

Table 3.3: DDS\_TopicDataQosPolicy

Type	Field Name	Description
DDS_OctetSeq	value	Empty by default

## GROUP\_DATA

This Qos Policy provides an area where your application can store additional information related to the *Publisher* and *Subscriber*.

Currently, [GROUP\\_DATA](#) of the associated *Publisher* or *Subscriber* is only propagated with the information that declares a *DataWriter* or *DataReader*. Thus, you will need to access the value of [GROUP\\_DATA](#) through [PublicationBuiltinTopicData](#) or [SubscriptionBuiltinTopicData](#). (See *Accessing User Discovery Data*)

The structure for the [TOPIC\\_DATA](#) QosPolicy includes just one field, as seen in Table 3.4. The field is a sequence of octets that translates to a contiguous buffer of bytes whose contents and length are set by the user. The maximum size for the data is set in the [DomainParticipantResourceLimitsQosPolicy](#). (See *Resource Limits*)

Table 3.4: DDS\_GroupDataQosPolicy

Type	Field Name	Description
DDS_OctetSeq	value	Empty by default

## 3.9 Partitions

### 3.9.1 Introduction

The [PARTITION](#) QoS provides a way to control which *Entities* will match—and thus communicate with—which other *Entities*. It can be used to prevent *Entities* that would have otherwise matched from talking to each other. Much in the same way that only applications within the same DDS domain will communicate with each other, only *Entities* that belong to the same partition can talk to each other.

The [PARTITION](#) QoS applies to *Publishers* and *Subscribers*. *DataWriters* and *DataReaders* belong to the partitions as set in the QoS of the *Publishers* and *Subscribers* that created them.

The [PARTITION](#) QoS consists of a set of partition names that identify the partitions of which the *Entity* is a member. These names can be concrete (e.g., ExamplePartition) or regular expression strings (e.g, Example\*), and two *Entities* are considered to be in the same partition if one of the *Entities* has a concrete partition name matching one of the concrete or regular expression partition names of the other *Entity* (see *Pattern matching for PARTITION names*). By default, *DataWriters* and *DataReaders* (through their *Publisher/Subscriber* parents), belong to a single partition whose name is the empty string, “”.

Conceptually, each partition name can be thought of as defining a “visibility plane” within the DDS domain. *DataWriters* will make their data available on all of the visibility planes that correspond

to their *Publisher's* partition names, and the *DataReaders* will see the data that is placed on all of the visibility planes that correspond to their *Subscriber's* partition names.

Figure 3.4 illustrates the concept of PARTITION QoS at the *Publisher* and *Subscriber* level. In this figure, all *DataWriters* and *DataReaders* belong to the same DDS domain ID and *Domain-Participant* partition, and they use the same *Topic*. *DataWriter1* is configured to belong to three partitions: partition\_A, partition\_B, and partition\_C. *DataWriter2* belongs to partition\_C and partition\_D.

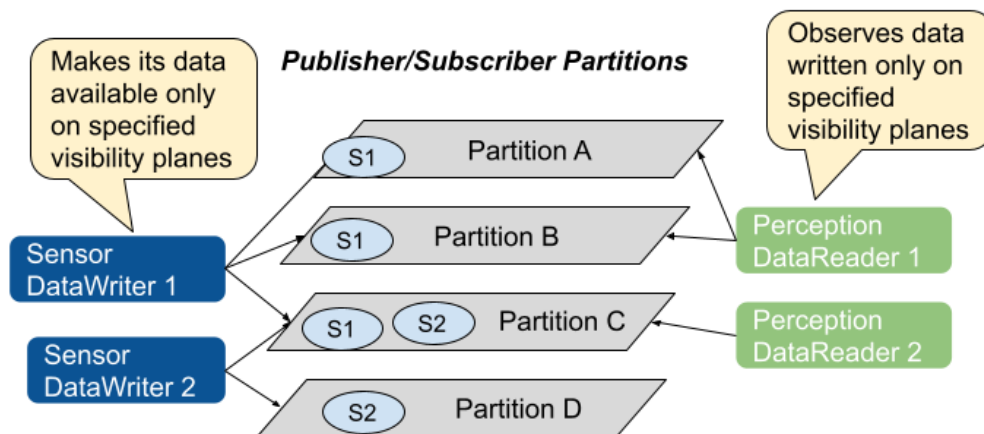


Figure 3.4: Controlling Visibility of Data with PARTITION QoS

Similarly, *DataReader1* is configured to belong to partition\_A and partition\_B, and *DataReader2* belongs only to partition\_C. Given this topology, the data written by *DataWriter1* is visible in partitions A, B, and C. The oval tagged with the number “S1” represents one DDS data sample written by *DataWriter1*.

Similarly, the data written by *DataWriter2* is visible in partitions C and D. The oval tagged with the number “S2” represents one DDS data sample written by *DataWriter2*.

The result is that the data written by *DataWriter1* will be received by both *DataReader1* and *DataReader2*, but the data written by *DataWriter2* will only be visible by *DataReader2*.

*Publishers* and *Subscribers* always belong to a partition. By default, *Publishers* and *Subscribers* belong to a single partition whose name is the empty string, “”. If you set the **PARTITION** QoS to be an empty set, *Connext Micro* will assign the *Publisher* or *Subscriber* to the default partition, “”. Thus, for the example above, without using the **PARTITION** QoS on any of the entities, *DataReaders* 1 and 2 would have received all data samples written by *DataWriters* 1 and 2.

### 3.9.2 Rules for PARTITION matching

The **PARTITION** QoSPolicy associates a set of partition names with the entity (*Publisher* or *Subscriber*). The partition names are concrete names (e.g., ExamplePartition) or regular expression strings (e.g., Example\*).

With regard to the **PARTITION** QoS, a *DataWriter* will communicate with a *DataReader* if and only if the following conditions apply:

1. The *DataWriter* and *DataReader* belong to *DomainParticipants* bound to the same DDS domain ID.
2. The *DataWriter* and *DataReader* have matching *Topics*. That is, each is associated with a *Topic* with the same name and compatible data type.
3. The QoS offered by the *DataWriter* is compatible with the QoS requested by the *DataReader*.

Matching partition names is done by string pattern matching, and partition names are case-sensitive.

---

**Note:** Failure to match partitions (on *Publisher* or *Subscriber*) is not considered an incompatible QoS and does not trigger any listeners or change any status conditions.

---

### 3.9.3 Pattern matching for PARTITION names

You may add strings that are regular expressions to the **PARTITION** QoSPolicy. A **PARTITION.name** is a regular expression if it contains any of the following unescaped special characters: \*, ?, [, ], !, or ^. The **PARTITION.name** strings can be “concrete” names or regular expression strings; a **PARTITION.name** element that is a regular expression will only match against concrete strings found in a **PARTITION.name** element of a different *Entity*’s **PARTITION** QoSPolicy.

If a **PARTITION** QoS only contains regular expressions, then the *Entity* will be assigned automatically to the default partition with the empty string name (“”). Thus, a **PARTITION** QoS that only contains the string \* matches another *Entity*’s **PARTITION** QoS that also only contains the string \*, not because the regular expression strings are identical, but because they both belong to the default “” partition.

For more on regular expressions, see *Regular Expression Matching* below.

Two *Entities* are considered to have a partition in common if the sets of partitions associated with them have:

- At least one concrete partition name in common
- A regular expression in one *Entity* that matches a concrete partition name in another *Entity*

The programmatic representation of the **PARTITION** QoS is shown in Table 3.5. The QoSPolicy contains the single string sequence, name. Each element in the sequence can be a concrete name or a regular expression. The *Entity* will be assigned to the default “” partition if the sequence is empty, or if the sequence contains only regular expressions.

Table 3.5: DDS\_PartitionQosPolicy

Type	Field Name	Description
DDS_StringSeq	name	Empty by default. There can be up to 64 names, with a maximum of 256 characters (including the NUL terminator), summed across all names.

You can have one long partition string of 256 chars, or multiple shorter strings that add up to 256 or fewer characters. For example, you can have one string of 4 chars and one string of 252 chars.

### Regular Expression Matching

The SQL expression format provided by *Connext Micro* supports the relational operator **MATCH**. It may only be used with string fields. The right-hand operator is a string pattern, which specifies a template that the left-hand field must match.

**MATCH** is case-sensitive. The following characters have special meaning, unless escaped by the escape character: , \ / ? \* [ ] - ^ ! \ %.

The pattern allows limited “wild card” matching under the rules in Table 3.6.

The syntax is similar to the POSIX® fnmatch syntax. (See <http://www.opengroup.org/onlinepubs/000095399/functions/fnmatch.html>.)

Table 3.6: Wild Card Matching

Character	Meaning
,	<b>NOT SUPPORTED</b> A , separates a list of alternate patterns. The field string is matched if it matches one or more of the patterns.
/	<b>NOT SUPPORTED</b> A / in the pattern string matches a / in the field string. It separates a sequence of mandatory substrings.
?	A ? in the patterns tring matches any single non-special characters in the field string.
*	A * in the pattern string matches 0 or more non-special characters in the field string.
%	<b>NOT SUPPORTED</b> This special character is used to designate filter expression parameters.
	Escape character for special characters.
[charlist]	Matches any one of the characters in charlist.
[!charlist] or [^charlist]	Matches any one of the characters <i>not</i> in charlist.
[s-e]	Matches any character from <b>s</b> to <b>e</b> , inclusive.
[!s-e] or [^s-e]	Matches any character <i>not</i> in the interval <b>s</b> to <b>e</b> .

**Note:** To use special characters as regular characters in regular expressions, you must escape

them using the character `\`. For example, `A[` is considered a malformed expression and the result is undefined.

### 3.9.4 Example

The **PARTITION** QosPolicy is useful to control which *DataWriters* can communicate with which *DataReaders* and vice versa—even if all of the *DataWriters* and *DataReaders* are for the same *Topic*. This facility is useful for creating temporary separation groups among *Entities* that would otherwise be connected to and exchange data each other.

The code below illustrates how to set the **PARTITION** QosPolicy on a *Publisher*:

```
struct DDS_PublisherQos pub_qos = DDS_PublisherQos_INITIALIZER;
    DDS_StringSeq_set_maximum(&pub_qos.partition.name,2);
    DDS_StringSeq_set_length(&pub_qos.partition.name,2);
    *DDS_StringSeq_get_reference(&pub_qos.partition.name,0) = DDS_String_dup("partition1
↪");
    *DDS_StringSeq_get_reference(&pub_qos.partition.name,1) = DDS_String_dup("partition2
↪");

publisher = DDS_DomainParticipant_create_publisher(application->participant,&pub_qos,
↪NULL,DDS_STATUS_MASK_NONE);
    if (publisher == NULL)
    {
        ...
    }
```

Using partitions, connectivity can be controlled based on location-based partitioning, access-control groups, or a combination of these and other application-defined criteria. We will examine some of these options via concrete examples.

#### Location-based partitions

Assume you have a set of *Topics* in a traffic management system such as “TrafficAlert,” “AccidentReport,” and “CongestionStatus.” You may want to control the visibility of these *Topics* based on the actual location to which the information applies. You can do this by placing the *Publisher* in a partition that represents the area to which the information applies. This can be done using a string that includes the city, state, and country, such as “USA/California/Santa Clara.” A *Subscriber* can then choose whether it wants to see the alerts in a single city, the accidents in a set of states, or the congestion status across the US. Some concrete examples are shown in Table 3.7.



Table 3.7: Example of Using Location-Based Partitions

Publisher Partitions	Subscriber Partitions	Result
Specify a single partition name using the pattern: “<country>/<state>/<city>”	Specify multiple partition names, one per region of interest	Limits the visibility of the data to <i>Subscribers</i> that express interest in the geographical region.
“USA/California/Santa Clara”	( <i>Subscriber</i> partition is irrelevant here.)	Send only information for Santa Clara, California.
(Publisher partition is irrelevant here.)	“USA/California/Santa Clara”	Receive only information for Santa Clara, California.
(Publisher partition is irrelevant here.)	“USA/California/Santa Clara” “USA/California/Sunnyvale”	Receive information for Santa Clara or Sunnyvale, California.
(Publisher partition is irrelevant here.)	“USA/California/*” “USA/Nevada/*”	Receive information for California or Nevada.
(Publisher partition is irrelevant here.)	“USA/California/*” “USA/Nevada/Reno” “USA/Nevada/Las Vegas”	Receive information for California and two cities in Nevada.

### Access-control group partitions

Suppose you have an application where access to the information must be restricted based on reader membership to access-control groups. You can map this group-controlled visibility to partitions by naming all the groups (e.g., executives, payroll, financial, general-staff, consultants, external-people) and assigning the *Publisher* to the set of partitions that represents which groups should have access to the information. The *Subscribers* specify the groups to which they belong, and the partition-matching behavior will ensure that the information is only distributed to *Subscribers* belonging to the appropriate groups. Some concrete examples are shown in Table 3.8

Table 3.8: Example of Access-Control Group Partitions

Publisher Partitions	Subscriber Partitions	Result
Specify several partition names, one per group that is allowed access:	Specify multiple partition names, one per group to which the <i>Subscriber</i> belongs.	Limits the visibility of the data to <i>Subscribers</i> that belong to the access-groups specified by the <i>Publisher</i> .
“payroll” “financial”	( <i>Subscriber</i> partition is irrelevant here.)	Makes information available only to <i>Subscribers</i> that have access to either financial or payroll information.
(Publisher partition is irrelevant here.)	“executives” “financial”	Gain access to information that is intended for executives or people with access to the finances.

A slight variation of this pattern could be used to confine the information based on security levels.

### 3.9.5 Properties

This QoSPolicy cannot be modified at runtime.

Strictly speaking, this QoSPolicy does not have request-offered semantics, although it is matched between *DataWriters* and *DataReaders*, and communication is established only if there is a match between partition names.

### 3.9.6 Resource limits

Before this QoS policy can be used, you must configure the following [DomainParticipantResourceLimitsQoSPolicy](#) fields:

- `max_partitions`: sets the maximum number of partitions for each [PARTITION](#) QoS.
- `max_partition_cumulative_characters`: sets the maximum number of characters (per *DomainParticipant*) that can be used for the sum-total length of all partition names. Note that the NUL terminator in each string contributes to the character count.
- `max_partition_string_size`: sets the maximum number of characters that can be used for each partition name. This can be set to a value greater than 0 or `DDS_LENGTH_UNLIMITED`.
- `max_partition_string_allocation`: sets the maximum total memory allocated to partition names across all *DomainParticipants*. This can be set to a value greater than 0 or `DDS_LENGTH_UNLIMITED`.

---

**Note:** All applications in the DDS domain must have the same resource limit values in order to communicate. For example, if two applications have different values, and one application sets the [PARTITION](#) QoSPolicy to hold more partitions or longer names than set by another application, the matching *Entities* between the two applications will not connect. This is similar to the restrictions for the *GROUP\_DATA*, *USER\_DATA*, and *TOPIC\_DATA* QoS Policies.

---

These fields collectively determine how your application manages partition memory. The subsections below explain how to configure your [DomainParticipantResourceLimitsQoSPolicy](#) for different behaviors.

#### Configuring for runtime allocation

This configuration allows memory for `PARTITION.name` strings to be allocated and freed during runtime. Each `PARTITION.name` string can be of any size; however, the sum-total string length of all partition names is still limited by `max_partition_cumulative_characters`.

For this behavior, set the following:

- Set `max_partition_string_size` to `DDS_LENGTH_UNLIMITED`.
- Set `max_partition_string_allocation` to `DDS_LENGTH_UNLIMITED`.

## Configuring for preallocated memory

Preallocating memory gives you greater control over memory utilization. There are two possible configurations for preallocated memory: reusable and non-reusable.

### Reusable preallocated memory

In this configuration, memory for `PARTITION.name` strings is preallocated and will never be freed during operation. However, memory will be reused for `PARTITION` names that are added, internally deleted, and no longer needed. For example, if the application creates a Publisher with a unique `PARTITION` name instance and then deletes it, the application will reuse the memory that was storing the unique name (unless there are other uses of that name).

For this behavior, set the following:

- Set `max_partition_string_size` to a value greater than 0.
- Set `max_partition_string_allocation` to a value greater than 0. This value must be large enough to store every instance of each `PARTITION` name that will be created or discovered.

Note that each `PARTITION` name will take up memory equal to `max_partition_string_size`, regardless of the actual string length.

### Non-reusable preallocated memory

In this configuration, memory for `PARTITION.name` strings is preallocated and will never be freed or reused.

For this behavior, set the following:

- Set `max_partition_string_size` to `DDS_LENGTH_UNLIMITED`.
- Set `max_partition_string_allocation` to a value greater than 0. This value must be large enough to store every instance of each `PARTITION` name that is created and discovered.

Note that each `PARTITION` name will take up memory equal to its exact string size.

## 3.10 Content Filtering

**Attention:** This is an experimental feature in *Connext Micro* 4.2.0. It is not guaranteed to be consistent or supported and should not be used in production.

Refer to *Experimental Features* for more information.

Content filtering provides a mechanism for *Connext Micro* to filter out data samples for *DataReaders* based on their contents. The user-specified filtering properties consist of a filter expression and a set of parameters:

- The filter expression evaluates a logical expression on the data samples. The filter expression is similar to the WHERE clause in an SQL expression.
- The parameters are strings that give values to the ‘parameters’ in the filter expression. There must be one parameter string for each parameter in the filter expression.

Content filtering in *Connext Micro* functions similarly to [Filtering Data](#) in *Connext Professional*, except that *Connext Micro* does not use ContentFilteredTopics. Instead, it functions by applying a QoS policy (DDS\_ContentFilterQosPolicy) to a *DataReader*. Some limitations apply to *SQL Filter Expression Notation* as well.

### 3.10.1 Creating a Content Filter

This section will walk you through how to set up content filtering for a *DataReader* in your application.

---

**Note:** When using the C++ API, register the content filtering library using the same method as the C API.

---

1. Register the DDS\_FilterPluginFactory as shown below:

```
#include "dds_filter/dds_filter.h"

if (DDS_FilterPluginFactory_register() != DDS_RETCODE_OK)
{
    /* Error handling */
}
```

This enables the content filtering feature for your application. It must be registered before a *DomainParticipant* is created in order for it to use the feature.

2. *Optional.* You can configure content filtering support for a given *DomainParticipant*, as shown in the snippet below:

```
struct DDS_DomainParticipantQos dp_qos = DDS_DomainParticipantQos_INITIALIZER;

/* To disable content filtering support for a specific participant,
 * clear the filter plugin name in the DDS_FilterQosPolicy
 */
RT_ComponentFactoryId_clear(&dp_qos.filter.name);

/* If you know that a DomainParticipant will only need to support filtering
 * for a limited number of entities, then you can reduce memory usage
 * by specifying the maximum number of unique filter expressions which
 * can be supported at one time.
 */
dp_qos.filter.resource_limits.filter_expression_max_count = ...;

/* The DDS_FilterQosPolicy has default limits for the maximum length of a filter_
↪ expression,
```

(continues on next page)

(continued from previous page)

```

    * the maximum number of parameters per expression, and the maximum length of a
    ↪filter
    * parameter. Here you can raise these limits if needed to support your use case or
    ↪they
    * can be reduced to limit memory usage.
    */
    dp_qos.filter.resource_limits.filter_expression_max_length = ...;
    dp_qos.filter.resource_limits.filter_parameter_max_count_per_expression = ...;
    dp_qos.filter.resource_limits.filter_parameter_max_length = ...;

    /* To reduce memory usage at the cost of network bandwidth, writer-side
    * filtering can also be disabled for all DataWriters created by this participant.
    */
    dp_qos.filter.disable_writer_filtering = DDS_BOOLEAN_TRUE;

    /* Optionally, configure the maximum number of predicates per SQL expression. For
    * longer expressions, it may need to be increased from the default of 4. For
    * simpler expressions, it can be reduced to limit memory usage.
    */
    dp_qos.sql_predicate_max_count_per_expression = 4;

    ... /* Create the DomainParticipant with the given Qos */

```

The `DDS_DomainParticipantQos` contains a `DDS_FilterQosPolicy` for configuring properties of the content filtering feature for a given *DomainParticipant*. By default, if the `FilterPluginFactory` is registered, every *DomainParticipant* within the application will support content filtering for all of its child entities. Through the `FilterQosPolicy`, you can change this default behavior to disable support or limit memory usage.

### 3. Create a *DataReader* with a content filter.

To enable content filtering on a *DataReader*, you must configure the `DDS_ContentFilterQosPolicy` in the *DataReader*'s QoS policy. A content filter can either be configured when the *DataReader* is created or it can be added/updated by updating the QoS of the *DataReader*.

Generally, *Connext Micro* does not support changing the QoS of DDS entities after they are enabled. The `ContentFilterQosPolicy` is an exception in that it can be updated after a *DataReader* is enabled by using `DDS_DataReader_set_qos()`. However, it will fail to update if any other fields of the QoS do not match those specified when the *DataReader* was created.

The contents of the `DDS_ContentFilterQosPolicy` are as follows:

Field	Type	Description
<code>filter_name</code>	DDS_String	String name to uniquely identify this content filter within a <i>DomainParticipant</i> . For interoperability with <i>Connext Professional</i> , it is equivalent to the <code>filter_name</code> when configuring a content filter with XML application creation.
<code>filter_class_name</code>	DDS_String	String name that identifies the type of filter expression. <i>Connext Micro</i> currently only supports the “DDSSQL” filter class.
<code>filter_expression</code>	DDS_String	String logical expression on the contents of the <i>DataReader</i> ’s <i>Topic</i> . If the expression evaluates to TRUE, a DDS sample is received; otherwise, it is discarded. The notation for this expression depends on the filter that you are using (specified by the <code>filter_name</code> parameter). See <i>SQL Filter Expression Notation</i> .
<code>expression_parameters</code>	DDS_StringSeq	String sequence of filter expression parameters. Each parameter corresponds to a positional argument in the filter expression: element 0 corresponds to positional argument 0, element 1 to positional argument 1, and so forth.

For example, to create a *DataReader* with a content filter:

```

struct DDS_DataReaderQos dr_qos = DDS_DataReaderQos_INITIALIZER;

/* Configure the content filter as desired */
dr_qos.content_filter.filter_name = DDS_String_dup("MyFilter");
dr_qos.content_filter.filter_class_name = DDS_String_dup(DDS_SQL_CONTENT_FILTER_
→CLASS);
dr_qos.content_filter.filter_expression = DDS_String_dup("a < 10");

... /* Create the DataReader with the given Qos */

```

---

**Note:** Content filtering can also be configured via *Application Generation Using XML*.

---

If your *DataReader* fails to be created or updated, check if any of the following is true:

- Your *DataReader* specifies a content filter, but content filtering is not enabled on its parent *DomainParticipant*.
- The filter expression exceeds one of the filter resource limits in the `DDS_FilterQosPolicy`.
- The total number of unique filter expressions within the *DomainParticipant* would exceed the `filter_expression_max_count`.

### 3.10.2 Where Filtering is Applied—Publishing vs. Subscribing Side

Content filtering can be applied on either side of your distributed application. We refer to filtering applied to a *DataReader* as *reader-side filtering*, while filtering applied for a *DataReader* by a *DataWriter* is called *writer-side filtering*.

Once a *DataReader* configures a filter in its QoS, it will only receive samples that pass the filter. Whether or not that filter is applied by the reader or the writer depends on whether the *DataWriter* supports writer-side filtering. If the writer supports writer-side filtering, then filtered samples can skip being sent over the network. If the *DataWriter* does not support writer-side filtering, then the *DataReader* will continue to only receive samples that pass its filter, but the *DataWriter* will send all samples over the network.

Therefore, writer-side filtering can significantly reduce bandwidth by only sending samples over the network that pass the *DataReader*'s filter. To perform writer-side filtering, the *DataWriter* obtains the filter expression and parameters from the *DataReader* during discovery.

A *DataWriter* will automatically filter DDS data samples for a *DataReader* if all of the following are true:

1. The content filtering feature is enabled on the parent *DomainParticipant*.
2. `dp_qos.filter.disable_writer_filtering` is set to `DDS_BOOLEAN_FALSE` on the parent *DomainParticipant*.
3. The *DataWriter* is filtering for no more than `writer_resource_limits.max_remote_reader_filters` *DataReaders* at the same time.
  - a. There is a resource-limit on the *DataWriter* called `writer_resource_limits.max_remote_reader_filters`. This value can be from `[0, (231)-1]` or `DDS_LENGTH_UNLIMITED` (default). 0 means do not filter any *DataReader*; 1 to `(231)-1` means that the *DataWriter* will filter for up to the specified number of *DataReaders*; `DDS_LENGTH_UNLIMITED` means that the *DataWriter* will filter for up to `writer_resource_limits.max_remote_readers`. The *DataWriter* will store the filter result per sample per filtered *DataReader*.
  - b. If a *DataWriter* is filtering `max_remote_reader_filters` *DataReaders* at the same time and a new filtered *DataReader* is discovered, then the newly created *DataReader* (`max_remote_reader_filters + 1`) will not be filtered by the *DataWriter* (but the *DataReader* will still apply the filter itself).

Even if one of the first `(max_remote_reader_filters)` *DataReaders* is deleted, the already-created *DataReader* (`max_remote_reader_filters + 1`) will still not be filtered. However, any subsequently created *DataReaders* will be filtered as long as the number of *DataReaders* currently being filtered is not more than `writer_resource_limits.max_remote_reader_filters`.

4. There are no more than four matching *DataReaders* in the same locator (transport destination, for example: IP address + port).

---

**Note:** *Connext Micro* supports limited writer-side filtering if there are more than four matching *DataReaders* in the same locator. The middleware will not send any sample to a

locator if the sample is filtered out by all the *DataReaders* receiving samples on that locator. However, if there are one or more *DataReaders* to which the sample has to be sent, some *DataReaders* on the locator will perform reader-side filtering for the incoming sample.

---

5. The *DataWriter* has infinite liveliness.
6. The *DataWriter* is not using Asynchronous Publication. Either the *DataWriter*'s PUBLISH\_MODE kind is set to DDS\_SYNCRHONOUS\_PUBLISHER\_MODE\_QOS, or the kind is set to DDS\_AUTO\_PUBLISHER\_MODE\_QOS and the data type does not require fragmentation.
7. Support for the specific filter class is registered in the *DomainParticipant* of the *DataWriter*.  
SQL filter support is enabled by default when filtering is enabled on a participant.
8. The specific sample being filtered was written by the *DataWriter* after matching with the *DataReader*.

If using a DURABILITY QoS setting greater than volatile, then samples written before the *DataWriter* matched with the *DataReader* may be available to the *DataReader* but will not have filtering applied by the *DataWriter*.

9. The sample doesn't exceed any of the expression-specific limits or the filter\_expression\_max\_count within the *DomainParticipant*.

If any of the above criteria is not met, the *DataReader* will perform reader-side content filtering.

### 3.10.3 Interoperability

Enabling a content filter on a *DataReader* has no impact on the interoperability of that *DataReader* with any other *DataWriter*. However, writer-side filtering is not supported by all other DDS products.

#### With Connext Professional

*Connext Micro*'s implementation of content filtering is compliant with the RTPS specification, allowing other DDS implementations (such as *Connext Professional*) to apply writer-side filtering for a *Connext Micro* application. Similarly, a *Connext Micro* application can apply writer-side filtering for a *Connext Professional* application as long as it meets the criteria listed in *Where Filtering is Applied—Publishing vs. Subscribing Side*.

*Connext Micro* supports a subset of the SQL filter syntax supported by *Connext Professional*. Any valid SQL filter on a *Connext Micro DataReader* can be applied by a *Connext Professional DataWriter*. However, if a *Connext Professional DataReader* specifies a filter with syntax not supported by *Connext Micro*, then *Connext Micro* will be unable to apply writer-side filtering for that *DataReader*. Refer to *SQL Filter Expression Notation* for more details.



## With Connext Cert

Content filtering is NOT supported when using the CERT profile or older versions of *Connext Micro*.

### 3.10.4 SQL Filter Expression Notation

Content filtering uses logical SQL filter expressions, similar to the **WHERE** clause in SQL. The following sections provide more information.

**Attention:** Compared to ContentFilteredTopics in *Connext Professional*, SQL filter expressions in *Connext Micro* have the following limitations:

- Filters can only be applied to the top-level members of a struct (nested structs and top-level unions are not supported).
- You cannot filter on the following members: strings, sequences, arrays, unions, and optional members.
- The XTypes library (`rti_me_ddsxtypes`) is required and a type must be compiled with the interpreter. See *Extensible Types (XTypes) 1.2 Compatibility*.
- Limited syntax; parentheses are not supported.

Future *Connext Micro* releases may remove these limitations.

#### Example SQL filter expressions

Assume that you have a *Topic* with two floats, X and Y, which are the coordinates of an object moving inside a rectangle measuring 200 x 200 units. This object moves quite a bit, generating lots of DDS samples that you are not interested in. Instead you only want to receive DDS samples outside the middle of the rectangle, as seen in Figure 3.5. That is, you want to filter out data points in the gray box.

The filter expression would look like this (remember the expression is written so that DDS samples that we *do* want will *pass*):

```
X < 50 OR X > 150 OR Y < 50 OR Y > 150 OR X > 150 AND Y < 50 OR X > 150 AND Y > 150
```

**Note:** This release of *Connext Micro* does not support parentheses in SQL expressions. However, AND takes logical precedence over OR.

Suppose you would like the ability to adjust the coordinates that are considered outside the acceptable range (changing the size of the gray box). You can achieve this by changing the whole filter expression or by using filter parameters. The expression can be written using filter parameters as follows:




Figure 3.5: Filtering Example

```
X < %0 AND Y < %2 OR X < %0 AND Y > %3 OR X > %1 AND Y < %2 OR X > %1 AND Y > %3
```

See the [String](#) and [Sequence Support](#) sections of the API Reference for more information.

In C++, the filter parameters could be assigned like this:

```
dr_qos.content_filter.expression_parameters.maximum(4);
dr_qos.content_filter.expression_parameters.length(4);
dr_qos.content_filter.expression_parameters[0] = DDS_String_dup("50");
dr_qos.content_filter.expression_parameters[1] = DDS_String_dup("150");
dr_qos.content_filter.expression_parameters[2] = DDS_String_dup("50");
dr_qos.content_filter.expression_parameters[3] = DDS_String_dup("150");
```

## SQL grammar

This section describes the subset of SQL syntax, in Backus-Naur Form (BNF), that you can use to form filter expressions.

We use the following notational conventions:

- *NonTerminals* are typeset in italics.
- ‘Terminals’ are quoted and typeset in a fixed-width font. They are written in upper case in most cases in the BNF-grammar below, but should be case insensitive.
- **TOKENS** are typeset in bold.
- The notation (element // ‘,’) represents a non-empty, comma-separated list of elements.

```
FilterExpression ::= Condition
Condition ::= Predicate
            | Condition 'AND' Condition
            | Condition 'OR' Condition
            | 'NOT' Condition
Predicate ::= ComparisonPredicate
            | BetweenPredicate
ComparisonPredicate ::= FIELDNAME RelOp Parameter
                    | Parameter RelOp FIELDNAME
                    | FIELDNAME RelOp FIELDNAME
BetweenPredicate ::= FIELDNAME 'BETWEEN' Range
                  | FIELDNAME 'NOT BETWEEN' Range
RelOp ::= '=' | '>' | '>=' | '<' | '<=' | '<>'
Range ::= Parameter 'AND' Parameter
Parameter ::= INTEGERVALUE
            | CHARVALUE
            | FLOATVALUE
            | ENUMERATEDVALUE
            | PARAMETER
```

## Token expressions

This section describes the syntax and meaning of the tokens used in SQL grammar.

**FIELDNAME:** A reference to a field in the data structure. Defined as any series of characters ‘a’, ..., ‘z’, ‘A’, ..., ‘Z’, ‘0’, ..., ‘9’, ‘\_’ but may not start with a digit. Only top-level fields of a data structure can be referenced. The names of the fields are those specified in the IDL definition of the corresponding structure, which may or may not match the fieldnames that appear on the language-specific mapping of the structure. FIELDNAME must resolve to a primitive IDL type; that is boolean, octet, int16, uint16, int32, uint32, int64, uint64, float, double, char, or enum.

```
FIELDNAME: LETTER (PART_LETTER)*
```

where LETTER: [ “A”-“Z”, “\_”, “a”-“z” ] PART\_LETTER: [ “A”-“Z”, “\_”, “a”-“z”, “0”-“9” ]

Primitive IDL types referenced by FIELDNAME are treated as different types in predicate according to the following table:

Predicate Data Type	IDL Type
BOOLEANVALUE	boolean
INTEGERVALUE	octet, int16, uint16, int32, uint32, int64, uint64
FLOATVALUE	float, double
CHARVALUE	char
ENUMERATEDVALUE	enum

**INTEGERVALUE:** Any series of digits, optionally preceded by a plus or minus sign, representing a decimal integer value within the range of the system. ‘L’ or ‘I’ must be used for int64 (long long), otherwise int32 (long) is assumed. A hexadecimal number is preceded by 0x and must be a valid hexadecimal expression.

```
INTEGERVALUE : ([ "+", "-" ]? ([ "0"-"9" ])+ [ ("L", "l") ]? | ([ "+", "-" ]? ["0x", "0X"] ([ "0"-"9"
→ ", "A"-"F", "a"-"f" ])+ [ ("L", "l") ]?)
```

**CHARVALUE:** A single character enclosed between single quotes.

```
CHARVALUE : "'" (~["'"])?''
```

**FLOATVALUE:** Any series of digits, optionally preceded by a plus or minus sign and optionally including a floating point (‘.’). ‘F’ or ‘f’ must be used for float, otherwise double is assumed. A power-of-ten expression may be postfixed, which has the syntax en or En, where n is a number, optionally preceded by a plus or minus sign.

```
FLOATVALUE : ([ "+", "-" ]? ([ "0"-"9" ])* ( "." )? ([ "0"-"9" ])+ (EXPONENT)? [ ("F", "f") ]?)
```

where EXPONENT: [ “e”, “E” ] ([ “+”, “-” ])? ([ “0”-“9” ])+

**ENUMERATEDVALUE:** A reference to a value declared within an enumeration. Enumerated values consist of the name of the enumeration label enclosed in single quotes. The name used for the enumeration label must correspond to the label names specified in the IDL definition of the enumeration.

```
ENUMERATEDVALUE : "'" ["A" - "Z", "a" - "z"] ["A" - "Z", "a" - "z", "_", "0" - "9"]* "'"
```

**BOOLEANVALUE:** Can either be **TRUE** or **FALSE**, and is case insensitive.

```
BOOLEANVALUE : ["TRUE", "FALSE"]
```

**PARAMETER:** Takes the form %*italic*:*n*, where *n* represents a natural number (zero included) smaller than 100. It refers to the (*n* + 1)th argument in the given context. This argument can only be in primitive type value format. It cannot be a FIELDNAME.

```
PARAMETER : "%" (["0"-"9"])+
```

### Type compatibility in the Predicate

As seen in Table 3.9, only certain combinations of type comparisons are valid in the Predicate.

Table 3.9: Valid Type Comparisons

	Boolean	Integer	Float	Char	Enum
Boolean	YES	✓			
Integer		✓	✓		✓ <sup>1</sup>
Float		✓	✓		
Char				✓	<sup>2</sup>
Enum		✓ <sup>1</sup>		<sup>2</sup>	✓

### Enumerations

A filter expression can use enumeration values, such as GREEN, instead of the numerical value. For example, if *x* is an enumeration of GREEN, YELLOW and RED, the following expressions are valid:

```
"x = 'GREEN' "
"x < 'RED' "
```

<sup>1</sup> An int32 (long) constant is compatible with an enumerated variable, but an enumerated constant is not compatible with an integer variable. For example, “MyEnum=1” is valid, but “MyInteger='EnumValue'” is not.

<sup>2</sup> Whether or not a constant value in single quotes is interpreted as a char or an enum depends on the type of the variable it is being compared with. For example, in “MyChar='A'”, A is interpreted as a char constant. In “MyEnum='A'”, A is interpreted as an enum value and must be a defined label for that enumerated type.

## 3.11 Generating Type Support

For *Connext Micro* to publish and subscribe to topics of user-defined types, the types have to be defined and programmatically registered with *Connext Micro*. Rather than have users manually implement each new type, *Connext Micro* provides the RTI Code Generator (*rtiddsgen*) utility for automatically generating type support code.

### 3.11.1 IDL type definition

*rtiddsgen* for *Connext Micro* accepts types defined in IDL. The HelloWorld examples included with *Connext Micro* use the following HelloWorld.idl:

```
struct HelloWorld {  
    string<128> msg;  
};
```

For further reference, see [the section on Creating User Data Types with IDL](#) in the [RTI Connext DDS Core Libraries User's Manual](#) (available [here](#) if you have Internet access).

### 3.11.2 Generating type support

Before running *rtiddsgen*, some environment variables must be set:

- `RTIMEHOME` sets the path of the *Connext Micro* installation directory
- `RTIMEARCH` sets the platform architecture (e.g. `i86Linux2.6gcc4.4.5` or `i86Win32VS2010`)
- `JREHOME` sets the path for a Java JRE

Note that a JRE is shipped with *Connext Professional* on platforms supported for the execution of *rtiddsgen* (Linux, Windows, and macOS). It is not necessary to set `JREHOME` on these platforms, unless a specific JRE is preferred.

## C

Run *rtiddsgen* from the command line to generate C language type-support for a UserType.idl (and replace any existing generated files):

```
> cd $rti_connext_micro_root/rtiddsgen/scripts  
> rtiddsgen -micro -language C -replace UserType.idl
```

## C++

Run *rtiddsgen* from the command line to generate C++ language type-support for a *UserType.idl* (and replace any existing generated files):

```
> cd $rti_connext_micro_root/rtiddsgen/scripts
> rtiddsgen -micro -language C++ -replace UserType.idl
```

### Notes on command-line options

In order to target *Connext Micro* when generating code with *rtiddsgen*, the `-micro` option must be specified on the command line.

To list all command-line options specifically supported by *rtiddsgen* for *Connext Micro*, enter:

```
> cd $rti_connext_micro_root/rtiddsgen/scripts
> rtiddsgen -micro -help
```

Existing users might notice that that previously available options, `-language microC` and `-language microC++`, have been replaced by `-micro -language C` and `-micro -language C++`, respectively. It is still possible to specify `microC` and `microC++` for backwards compatibility, but users are advised to switch to using the `-micro` command-line option along with other arguments.

### Generated type support files

*rtiddsgen* will produce the following header and source files for each IDL file passed to it:

- *UserType.h* and *UserType.c(xx)* implement creation/initialization and deletion of a single sample and a sequence of samples of the type (or types) defined in the IDL description.
- *UserTypePlugin.h* and *UserTypePlugin.c(xx)* implement the pluggable type interface that *Connext Micro* uses to serialize and deserialize the type.
- *UserTypeSupport.h* and *UserTypeSupport.c(xx)* define type-specific *DataWriters* and *DataReaders* for user-defined types.

### 3.11.3 Using custom data-types in Connext Micro applications

A *Connext Micro* application must first of all include the generated headers. Then it must register the type with the *DomainParticipant* before a topic of that type can be defined. For an example *HelloWorld* type, the following code registers the type with the participant and then creates a topic of that type:

```
#include "HelloWorldPlugin.h"

/* ... */

retcode = DDS_DomainParticipant_register_type(application->participant,
```

(continues on next page)

(continued from previous page)

```

                                "HelloWorld",
                                HelloWorldTypePlugin_get());
if (retcode != DDS_RETCODE_OK)
{
    /* Log an error */
    goto done;
}

application->topic =
    DDS_DomainParticipant_create_topic(application->participant,
                                       "Example HelloWorld",
                                       "HelloWorld",
                                       &DDS_TOPIC_QOS_DEFAULT, NULL,
                                       DDS_STATUS_MASK_NONE);

if (application->topic == NULL)
{
    /* Log an error */
    goto done;
}

```

See the full HelloWorld examples for further reference.

### 3.11.4 Customizing generated code

*rtiddsgen* allows *Connext Micro* users to select whether they want to generate code to subscribe to and/or publish a custom data-type. When generating code for subscriptions, only those parts of code dealing with deserialization of data and the implementation of a typed *DataReader* endpoint are generated. Conversely, only those parts of code addressing serialization and the implementation of a *DataWriter* are considered when generating publishing code.

Control over these options is provide by two command-line arguments:

- **-reader** generates code for deserializing custom data-types and creating *DataReaders* from them.
- **-writer** generates code for serializing custom data-types and creating *DataWriters* from them.

If neither of these two options are supplied to *rtiddsgen*, they will both be considered active and code for both *DataReaders* and *DataWriters* will be generated. If only one of the two options is supplied to *rtiddsgen*, only that one is enabled. If both options are supplied, both are enabled.



### 3.11.5 Unsupported Features of rtiddsgen with Connext Micro

*Connext Micro* supports a subset of the features and options in *rtiddsgen*. Use *rtiddsgen -micro -help* to see the list of features supported by *rtiddsgen* for *Connext Micro*.

## 3.12 Threading Model

### 3.12.1 Introduction

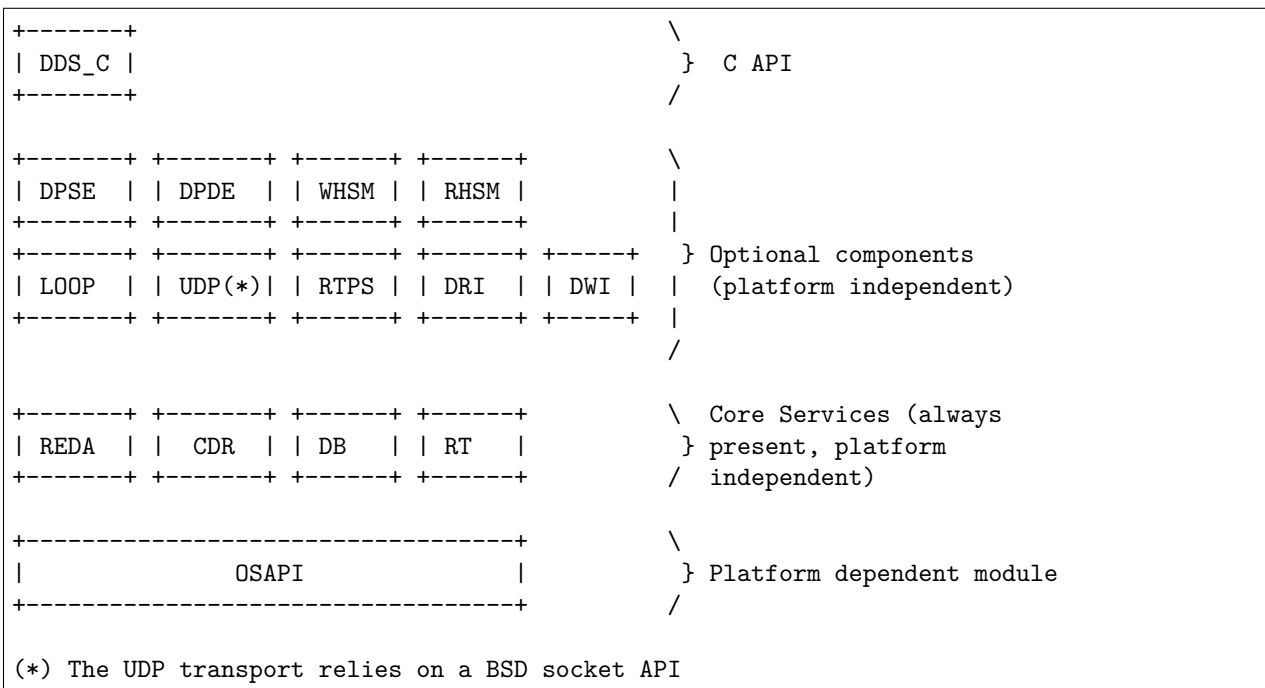
This section describes the threading model, the use of critical sections, and how to configure thread parameters in *RTI Connext Micro*. Please note that the information contained in this document applies to application development using *Connext Micro*.

This section includes:

- *Architectural Overview*
- *Threading Model*
- *UDP Transport Threads*

### 3.12.2 Architectural Overview

*RTI Connext Micro* consists of a core library and a number of components. The core library provides a porting layer, frequently used data-structures and abstractions, and the DDS API. Components provide additional functionality such as UDP communication, DDS discovery plugins, DDS history caches, etc.



### 3.12.3 Threading Model

*RTI Connext Micro* is architected in a way that makes it possible to create a port of *Connext Micro* that uses no threads, for example on platforms with no operating system. Thus, the following discussion can only be guaranteed to be true for *Connext Micro* libraries from RTI.

#### OSAPI Threads

The *Connext Micro* OSAPI layer creates one thread per OS process. This thread manages all the *Connext Micro* timers, such as deadline and liveliness timers. This thread is created by the *Connext Micro* OSAPI System when the `OSAPI_System_initialize()` function is called. When the *Connext Micro* DDS API is used `DomainParticipantFactory_get_instance()` calls this function once.

#### Configuring OSAPI Threads

The timer thread is configured through the `OSAPI_SystemProperty` structure and any changes must be made before `OSAPI_System_initialize()` is called. In *Connext Micro*, `DomainParticipantFactory_get_instance()` calls `OSAPI_System_initialize()`. Thus, if it is necessary to change the system timer thread settings, it must be done before `DomainParticipantFactory_get_instance()` is called the first time.

Please refer to `OSAPI_Thread` for supported thread options. Note that not all options are supported by all platforms.

```
struct OSAPI_SystemProperty sys_property = OSAPI_SystemProperty_INITIALIZER;

if (!OSAPI_System_get_property(&sys_property))
{
    /* ERROR */
}

/* Please refer to OSAPI_ThreadOptions for possible options */
sys_property.timer_property.thread.options = ....;

/* The stack-size is platform dependent, it is passed directly to the OS */
sys_property.timer_property.thread.stack_size = ....

/* The priority is platform dependent, it is passed directly to the OS */
sys_property.timer_property.thread.priority = ....

if (!OSAPI_System_set_property(&sys_property))
{
    /* ERROR */
}
```

## UDP Transport Threads

Of the components that RTI provides, only the UDP component creates threads. The UDP transport creates one receive thread for each unique UDP receive address and port. Thus, three UDP threads are created by default:

- A multicast receive thread for discovery data (assuming multicast is available and enabled)
- A unicast receive thread for discovery data
- A unicast receive thread for user-data

Additional threads may be created depending on the transport configuration for a *DomainParticipant*, *DataReader* and *DataWriter*. The UDP transport creates threads based on the following criteria:

- Each unique unicast port creates a new thread
- Each unique multicast address *and* port creates a new thread

For example, if a *DataReader* specifies its own multicast receive address a new receive thread will be created.

## Configuring UDP Receive Threads

All threads in the UDP transport share the same thread settings. It is important to note that all the UDP properties must be set before the UDP transport is registered. *Connext Micro* pre-registers the UDP transport with default settings when the [DomainParticipantFactory](#) is initialized. To change the UDP thread settings, use the following code.

```
RT_Registry_T *registry = NULL;
DDS_DomainParticipantFactory *factory = NULL;
struct UDP_InterfaceFactoryProperty *udp_property = NULL;

factory = DDS_DomainParticipantFactory_get_instance();

udp_property = (struct UDP_InterfaceFactoryProperty *)
    malloc(sizeof(struct UDP_InterfaceFactoryProperty));
*udp_property = UDP_INTERFACE_FACTORY_PROPERTY_DEFAULT;

registry = DDS_DomainParticipantFactory_get_registry(factory);

if (!RT_Registry_unregister(registry, "_udp", NULL, NULL))
{
    /* ERROR */
}

/* Please refer to OSAPI_ThreadOptions for possible options */
udp_property->recv_thread.options = ...;

/* The stack-size is platform dependent, it is passed directly to the OS */
udp_property->recv_thread.stack_size = ...
```

(continues on next page)

(continued from previous page)

```
/* The priority is platform dependent, it is passed directly to the OS */
udp_property->recv_thread.priority = ....

if (!RT_Registry_register(registry, "_udp",
                        UDP_InterfaceFactory_get_interface(),
                        (struct RT_ComponentFactoryProperty*)udp_property,
                        NULL))
{
    /* ERROR */
}
```

## General Thread Configuration

The *Connext Micro* architecture consists of a number of components and layers, and each layer and component has its own properties. It is important to remember that the layers and components are configured independently of each other, as opposed to configuring everything through DDS. This design makes it possible to relatively easily swap out one part of the library for another.

All threads created based on *Connext Micro* OSAPI APIs use the same [OSAPI\\_ThreadProperty](#) structure.

### 3.12.4 Critical Sections

*RTI Connext Micro* may create multiple threads, but from an application point of view there is only a single critical section protecting all DDS resources. Note that although *Connext Micro* may create multiple mutexes, these are used to protect resources in the OSAPI layer and are thus not relevant when using the public DDS APIs.

## Calling DDS APIs from listeners

When DDS is executing in a listener, it holds a critical section. Thus it is important to return as quickly as possible to avoid stalling network I/O.

There are no deadlock scenarios when calling *Connext Micro* DDS APIs from a listener. However, there are no checks on whether or not an API call will cause problems, such as deleting a participant when processing data in [on\\_data\\_available](#) from a reader within the same participant.

## 3.13 Batching

This section is organized as follows:

- *Overview*
- *Interoperability*
- *Performance*
- *Example Configuration*

### 3.13.1 Overview

Batching refers to a mechanism that allows *RTI Connext Micro* to collect multiple user data DDS samples to be sent in a single network packet, to take advantage of the efficiency of sending larger packets and thus increase effective throughput.

*Connext Micro* supports receiving batches of user data DDS samples, but does not support any mechanism to collect and send batches of user data.

Receiving batches of user samples is transparent to the application, which receives the samples as if the samples had been received one at a time. Note though that the reception sequence number refers to the sample sequence number, not the RTPS sequence number used to send RTPS messages. The RTPS sequence number is the batch sequence number for the entire batch.

A *Connext Micro DataReader* can receive both batched and non-batched samples.

For a more detailed explanation, please refer to [the BATCH QosPolicy section](#) in the [RTI Connext DDS Core Libraries User's Manual](#) (available [here](#) if you have Internet access).

### 3.13.2 Interoperability

*RTI Connext Professional* supports both sending and receiving batches, whereas *RTI Connext Micro* supports only receiving batches. Thus, this feature primarily exists in *Connext Micro* to interoperate with *RTI Connext* applications that have enabled batching. An *Connext Micro DataReader* can receive both batched and non-batched samples.

### 3.13.3 Performance

The purpose of batching is to increase throughput when writing small DDS samples at a high rate. In such cases, throughput can be increased several-fold, approaching much more closely the physical limitations of the underlying network transport.

However, collecting DDS samples into a batch implies that they are not sent on the network immediately when the application writes them; this can potentially increase latency. But, if the application sends data faster than the network can support, an increased proportion of the network's available bandwidth will be spent on acknowledgements and DDS sample resends. In this case, reducing that overhead by turning on batching could decrease latency while increasing throughput.

### 3.13.4 Example Configuration

This section includes several examples that explain how to enable batching in *RTI Connext Professional*. For more detailed and advanced configuration, please refer to the [RTI Connext DDS Core Libraries User's Manual](#).

- This configuration ensures that a batch will be sent with a maximum of 10 samples:

```
<datawriter_qos>
  <publication_name>
    <name>HelloWorldDataWriter</name>
  </publication_name>
  <batch>
    <enable>true</enable>
    <max_samples>10</max_samples>
  </batch>
</datawriter_qos>
```

- This configuration ensures that a batch is automatically flushed after the delay specified by `max_flush_delay`. The delay is measured from the time the first sample in the batch is written by the application:

```
<datawriter_qos>
  <publication_name>
    <name>HelloWorldDataWriter</name>
  </publication_name>
  <batch>
    <enable>true</enable>
    <max_flush_delay>
      <sec>1</sec>
      <nanosec>0</nanosec>
    </max_flush_delay>
  </batch>
</datawriter_qos>
```

- The following configuration ensures that a batch is flushed automatically when `max_data_bytes` is reached (in this example 8192).

```
<datawriter_qos>
  <publication_name>
    <name>HelloWorldDataWriter</name>
  </publication_name>
  <batch>
    <enable>true</enable>
    <max_data_bytes>8192</max_data_bytes>
  </batch>
</datawriter_qos>
```

Note that `max_data_bytes` does not include the metadata associated with the batch samples.

Batches must contain whole samples. If a new batch is started and its initial sample causes the serialized size to exceed `max_data_bytes`, *RTI Connext Professional* will send the sample in a single batch.

## 3.14 Message Integrity Checking

*Connext Micro* uses the DDS-I/RTPS protocol for communication between DDS applications, and RTPS messages are sent and received by a transport. When an RTPS message is sent across a communication link, such as Ethernet, it is possible that some bits may change value. These errors may cause communication failures or incorrect data to be received. In order to *detect* these types of errors, transports such as UDP often include a checksum to validate the integrity of the data: a sender adds the checksum to the transmitted data and the receiver validates that the calculated checksum for the received data matches the checksum received from the sender. If the checksums are different, a data corruption has occurred.

By default, *Connext Micro* relies on the underlying transport, such as UDP, to handle data integrity checking. However, the underlying transport may not provide sufficient integrity checking, or may itself introduce errors that *Connext Micro* must be able to detect regardless of the transport.

In order to address both of these scenarios for *any* transport, *Connext Micro* supports RTPS message integrity checking by adding a checksum to the RTPS message itself. This chapter describes the setup and default options to access this feature.

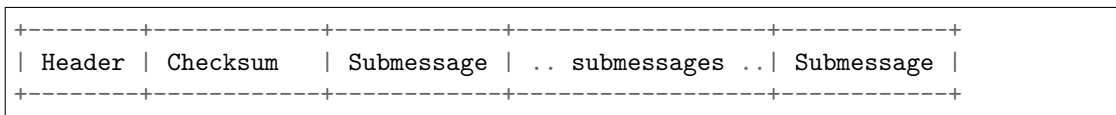
For information on how to write custom checksum functions, please refer to *RTPS*.

### 3.14.1 RTPS checksum

*Connext Micro* implements checksum validation on a complete RTPS message. A typical RTPS message without a checksum has the following structure:



When the message integrity checking feature is enabled, the structure of the RTPS message changes as illustrated below:



The sender calculates the checksum for the entire message with a checksum field set to 0 and places the result in the checksum field.

The receiver saves the the received checksum, sets the received checksum field to 0, and calculates the checksum for the entire message. It then compares the calculated checksum with the received checksum. If the checksums differ, the entire RTPS message is considered corrupted.

Note that the checksum is used *only* for error *detection* and *not* for error *correction*.

### 3.14.2 Configurations

You can configure your application to define which algorithms to use and validate as well as the requirements enforced by the participant when communicating with other participants using the [DDS\\_WireProtocolQosPolicy](#).

Configuring the message integrity checking consists of the two parts:

1. Selecting the checksum algorithm.
2. Configuring how a participant applies the checksums.

#### Selecting a checksum algorithm

*Connext Micro* supports three built-in algorithms and can be configured to use any of the following algorithms:

1. DDS\_CHECKSUM\_BUILTIN32: CRC-32 As defined by ISO/IEC 13239:2002.
2. DDS\_CHECKSUM\_BUILTIN64: CRC-64 As defined by ISO/IEC 13239:2002.
3. DDS\_CRC\_BUILTIN128: MD5 Message Digest

The CRC functions have the following properties:

Checksum	Polynom	Initial Value	Input Reflected	Output Reflected	XOR Value
CRC-32	0x04c11db7	$2^{32} - 1$	true	true	$2^{32} - 1$
CRC-64	0x1b	$2^{64} - 1$	true	true	$2^{64} - 1$

Please refer to *RTPS* for information on how to implement custom checksum functions.

#### Configuring the DDS DomainParticipant

The RTPS message integrity feature is configured in the [DDS\\_WireProtocolQosPolicy](#) for a participant. This QoS determines which RTPS checksum should be allowed, and if checksums should be sent and/or validated.

The following three fields determine how a participant uses RTPS checksums:

- **compute\_crc** - This configures the participant to send a checksum with each RTPS message. Which checksum to send is determined by **computed\_crc\_kind**.
- **check\_crc** - This configures the participant to verify the checksum in each received RTPS message if the checksum is present. If the checksum is valid, the message is accepted; otherwise, the message is dropped. If a message is received *without* a checksum, it is accepted and processed.
- **require\_crc** - This configures the participant to *require* that a checksum is present in the receiving packet. Messages without a checksum are dropped without further processing. Note that this option is orthogonal to the **check\_crc** options. This option only requires that a



checksum is included, it does not validate it. To validate and only accept messages with a checksum, *both* `check_crc` and `require_crc` must be `true`.

The following two fields determine which checksums are used:

- `computed_crc_kind` - The checksum type to include in each RTPS message when `compute_crc` is `true`.
- `allowed_crc_mask` - A mask of all checksum algorithms that the participant can verify. This allows the participant to receive messages from other participants with a different `computed_crc_kind`. A participant will ignore a participant that is sending a checksum that it cannot validate.

For example, the following snippet shows how to configure the participant to:

- Send all messages (except the participant announcements; see *Participant discovery and compatibility* below) with `DDS_CHECKSUM_BUILTIN64`.
- Accept `DDS_CHECKSUM_BUILTIN32`, `DDS_CHECKSUM_BUILTIN64`, and `DDS_CHECKSUM_BUILTIN128` algorithms.

```
struct DDS_DomainParticipantQos dp_qos =
    DDS_DomainParticipantQos_INITIALIZER;

dp_qos.protocol.computed_crc_kind = DDS_CHECKSUM_BUILTIN64;

dp_qos.protocol.allowed_crc_mask = DDS_CHECKSUM_BUILTIN32
    | DDS_CHECKSUM_BUILTIN64
    | DDS_CHECKSUM_BUILTIN128;
```

### 3.14.3 Participant discovery and compatibility

*Connext Micro* ensures that participants establish communication with each other only when they have compatible checksum configurations. If `compute_crc` is `true`, all messages sent from the participant are protected by a checksum. Since each participant can use a different type of checksum, a mechanism is required to ensure that participants are compatible during discovery.

To bootstrap this mechanism, all participant announcements (if `compute_crc` is set to `true`) include a checksum of type `DDS_CHECKSUM_BUILTIN32`. The participant announcement carries information about the `computed_crc_kind` (the checksum kind used by the participant) and the `allowed_crc_mask` (the checksum kinds understood by the participant), and whether or not the participant requires a checksum for each RTPS message (if `require_crc` is set to `true`). Please note that messages with `DDS_CHECKSUM_BUILTIN32` checksum are *always* accepted to enable discovering new participants.

For a Participant (A) to match with another Participant (B), the `computed_crc_kind` of Participant (B) must be a strict subset of the `allowed_crc_mask` of Participant (A) and vice versa. If Participant (B) does not send a checksum (`compute_crc` is set to `false`), it can only match Participant (A) if it does not set `require_crc` to `true`.

### 3.14.4 Interoperability with Connext Professional

*Connext Micro* supports two different kinds of RTPS submessages for CRC 32-bit checksums:

- The standard CRC 32-bit checksum in the RTPS header extension, as defined by the [OMG](#).
- The legacy CRC 32-bit checksum submessage used by older versions of *Connext Professional*.

*Connext Micro* will understand and accept either kind of received submessage. However, it may be necessary to change the transmit mode of *Connext Micro* to enable interoperability with older versions of *Connext Professional* and allow *Connext Professional* to validate the checksum.

The following two transmit modes are available:

- **RTPS\_CRC\_TXMODE\_OMG**: uses the standard method as defined by the [OMG](#). This is the default mode. The checksums sent by *Connext Micro* may not be understood by older versions of *Connext Professional* and cause *Connext Professional* to treat the message as if it does not include a checksum.
- **RTPS\_CRC\_TXMODE\_RTICRC32**: uses the legacy CRC32 mode. This mode sets the `computed_crc_kind` to `DDS_CRC_BUILTIN32`. The checksum sent by *Connext Micro* will be understood by older versions of *Connext Professional*. Use this option only if there is a *Connext Professional* application in your system which requires the legacy CRC 32-bit checksum.

**Warning:** *Connext Professional* does not support the `require_crc` field in the [DDS\\_Wire-ProtocolQosPolicy](#).

## 3.15 Sending Large Data

*Connext Micro* supports transmission and reception of data types that exceed the maximum message size of a transport. This section describes the behavior and the configuration options.

This section includes:

- *Overview*
- *Configuration of Large Data*
- *Limitations*

### 3.15.1 Overview

*Connext Micro* supports transmission and reception of data samples that exceed the maximum message size of a transport. For example, UDP supports a maximum user payload of 65507 bytes. In order to send samples larger than 65507 bytes, *Connext Micro* must split the sample into multiple UDP payloads.

When a sample larger than the transport size is sent, *Connext Micro* splits the sample into fragments and starts sending fragments based on a flow-control algorithm. A bandwidth-allocation parameter on the *DataWriter* and the scheduling rate determine how frequently and how much data can be sent each period.

When a sample is received in multiple fragments, the receiver reassembles each fragment into a complete serialized sample. The serialized data is then deserialized and made available to the user as regular data.

When working with large data, it is important to keep the following in mind:

- Fragmentation is always enabled.
- Fragmentation is per *DataWriter*.
- Flow-control is per *DataWriter*. It is important to keep this in mind since in *RTI Connext Professional* the flow-controller works across all *DataWriters* in the same publisher.
- Fragmentation is on a per sample basis. That is, two samples of the same type may lead to fragmentation of one sample, but not the other. The application is never exposed to fragments.
- It is the *DataWriters* that determine the fragmentation size. Different *DataWriters* can use different fragmentation sizes for the same type.
- All fragments must be received before a sample can be reconstructed. When using best-effort, if a fragment is lost, the entire sample is lost. When using reliability, a fragment that is not received may be resent. If a fragment is no longer available, the entire sample is dropped.
- If one of the DDS **write()** APIs is called too fast when writing large samples, *Connext Micro* may run out of resources. This is because the sample may take a long time to send and resources are not freed until the complete sample has been sent.

It is important to distinguish between the following concepts:

- Fragmentation by *Connext Micro*.
- Fragmentation by an underlying transport, e.g., IP fragmentation when UDP datagrams exceed about 1488 bytes.
- The maximum transmit message size of the sender. This is the maximum size of any payload going over the transport.
- The maximum transport transmit buffer size of the sender. This is the maximum number of bytes that can be stored by the transport.
- The maximum receive message size of a receiver. This is the maximum size of a single payload on a transport.

- The maximum receive buffer size of a receiver. This is the maximum number of bytes that can be received.

### 3.15.2 Configuration of Large Data

For a general overview of writing large data, please refer to these sections in the [RTI Connext DDS Core Libraries User's Manual](#):

- [the ASYNCHRONOUS\\_PUBLISHER QoSPolicy section](#) (available [here](#) if you have Internet access)
- [the FlowControllers section](#) (available [here](#) if you have Internet access)

NOTE: *Connext Micro* only supports the default FlowController.

Asynchronous publishing is handled by a separate thread that runs at a fixed rate. The rate and properties of this thread can be adjusted in the `OSAPI_SystemProperty` and the following fields before `DomainParticipantFactory_get_instance()` is called.

```
struct OSAPI_SystemProperty sys_property = OSAPI_SystemProperty_INITIALIZER;
DDS_DomainParticipantFactory *factory = NULL;

if (!OSAPI_System_get_property(&sys_property))
{
    /* error */
}

sys_property.task_scheduler.thread.stack_size = ....
sys_property.task_scheduler.thread.options = ....
sys_property.task_scheduler.thread.priority = ....
sys_property.task_scheduler.rate = rate in nanosec;

if (!OSAPI_System_set_property(&sys_property))
{
    /* error */
}

factory = DDS_DomainParticipantFactory_get_instance();

....
```

### 3.15.3 Limitations

The following are known limitations and issues with Large Data support:

- It is not possible to disable fragmentation support.
- The scheduler thread accuracy is based on the operating system.

### 3.16 Zero Copy Transfer

Zero Copy transfer enables *RTI Connext Micro* to transmit data samples without copying them internally, similar to [Zero Copy Transfer Over Shared Memory](#) in *Connext Professional*. This offers several benefits, including higher throughput of user data, reduced latency between DDS endpoints (compared to other transports that send serialized data, such as UDP), and decoupling sample size from latency. This is particularly useful in applications with large sample sizes, such as image or lidar point cloud data.

At a high level, Zero Copy transfer works by a *DataWriter* and *DataReader* accessing the same shared memory; see Figure 3.6 below. A Zero Copy-enabled *DataWriter* creates a structure in a shared memory region and allocates samples from its pool to shared memory. When samples are published by the *DataWriter*, matching *DataReaders* are notified via a transport that new samples are available in shared memory. The *DataReader* then accesses the samples in shared memory using the standard DDS read/take APIs. Note that the *DataReader* and *DataWriter* must be co-located—that is, within the same operating system instance.




Figure 3.6: Zero Copy Transfer in *Connext Micro*

There are two methods for performing Zero Copy transfer in *Connext Micro*, which we refer to as Zero Copy v1 and v2 in this documentation. The main difference between these two methods is the transport used to notify *DataReaders* that new samples are available:

- Zero Copy v1: *Shared Memory Transport (SHMEM)*
- Zero Copy v2: *Zero Copy v2 Transport*

When choosing between the two versions, consider primarily whether you plan to interoperate *Connext Micro* with *Connext Cert* or *Connext Professional*. Zero Copy v1 is compatible with *Connext Professional*, while Zero Copy v2 is compatible with *Connext Cert*. Other functional differences between v1 and v2 are described in *Compatibility*.

### 3.16.1 Compatibility

Zero Copy v1 and v2 are API-compatible, meaning the API is identical whether you are using the Shared Memory transport or Zero Copy v2 transport. However, there are some important functional differences between the two versions:

1. **Interoperability:** Zero Copy v1 is compatible with the [Zero Copy Transfer Over Shared Memory](#) feature in *Connext Professional*, but not with *Connext Cert*. Zero Copy v2 is not compatible with *Connext Professional*, but is compatible with the Zero Copy transfer feature in *Connext Cert*.

**Warning:** Zero Copy transfer is not supported on all versions of *Connext Cert*. Consult the documentation in your *Connext Cert* installation for more information.

2. **Sample synchronization:** Zero Copy v2 synchronizes samples between the *DataWriter* and the *DataReader*. This ensures that a *DataReader* cannot access a sample while it is being modified by the *DataWriter*. As a result, the *DataReader* does not need to call the API `FooDataReader_is_data_consistent` to verify whether the sample has been changed, because the API always returns TRUE. In contrast, Zero Copy v1 does not provide any consistency checks, making it necessary for you to call `FooDataReader_is_data_consistent`.

---

**Note:** In Zero Copy v2, a slow *DataReader* may still miss samples if the *DataWriter* overwrites them. These missed samples will not be delivered to the reader. This is similar to Zero Copy v1, where samples can be missed, but v1 offers no guarantees of consistency between *DataWriter* and *DataReader*.

---

3. **Volatile DataReaders:** When using Zero Copy v2, a volatile *DataReader* receives all the historical samples available in the *DataWriter*'s queue. Zero Copy v1 uses the `depth` that is set for historical samples.
4. **Sample acknowledgement:** Zero Copy v2 does not support sample acknowledgments. The transport is inherently reliable, meaning that the *DataWriter*'s samples are immediately available to the *DataReader*. However, samples may be removed from the *DataWriter*'s cache when the queue reaches its `history.depth`, even if the reader has not accessed them.
5. **KEEP\_ALL support:** Zero Copy v2 does not support the `KEEP_ALL` policy. Zero Copy v1 supports `KEEP_ALL`.
6. **Version priority:** If both Zero Copy v1 and v2 are enabled, *Connext Micro* will prioritize Zero Copy v1 over v2.

The following table summarizes the key differences between the two versions:

Table 3.10: Zero Copy Compatibility

Feature	Zero Copy v1	Zero Copy v2
Compatibility with <i>Connext Professional</i>	✓	
Compatibility with <i>Connext Cert</i>		✓ <sup>1</sup>
Sample synchronization		✓
Sample acknowledgement	✓	
Support for KEEP_ALL	✓	
Transfer of discovery data	✓ <sup>2</sup>	

### 3.16.2 Overview

Zero Copy samples reside in a shared memory region accessible from multiple processes. When creating a [FooDataWriter](#) that supports Zero Copy transfer of user samples, a sample must be created with a new non-DDS API ([FooDataWriter\\_get\\_loan\(\)](#)). This will return a pointer A\* to a sample Foo that lies inside a shared memory segment. A reference to this sample will be sent to a receiving [FooDataReader](#) across the shared memory. This [FooDataReader](#) will attach to a shared memory segment, and a pointer B\* to sample Foo will be presented to the user. Because the two processes share different memory spaces, A\* and B\* will be different but they will point to the same place in RAM.

This feature requires using new RTI DDS Extension APIs:

- [FooDataWriter\\_get\\_loan\(\)](#)
- [FooDataWriter\\_discard\\_loan\(\)](#)
- [FooDataReader\\_is\\_data\\_consistent\(\)](#)

### 3.16.3 Getting started

To enable Zero Copy transfer (either v1 or v2), follow these steps:

1. Annotate your type with the `@transfer_mode(SHMEM_REF)` annotation.

Currently, variable-length types (strings and sequences) are not supported for types using this transfer mode when a type is annotated with the PLAIN language binding (e.g., `@language_binding(PLAIN)` in IDL).

```
@transfer_mode(SHMEM_REF)
struct HelloWorld
{
    long id;
    char raw_image_data[1024 * 1024]; // 1 MB
};
```

<sup>1</sup> Only if Zero Copy v2 is supported on the corresponding version of *Connext Cert*.

<sup>2</sup> Discovery data is sent over the Shared Memory Transport (SHMEM), but it is not zero copied over.

2. Register the *Shared Memory Transport (SHMEM)* OR the *Zero Copy v2 Transport*. References will be sent across the chosen transport.

---

**Note:** If both transports are registered, *Connext Micro* will prioritize the Shared Memory transport (SHMEM) over the Zero Copy v2 transport as described in *Compatibility*.

---

**Warning:** If neither transport is registered AND your type is annotated with `@transfer_mode(SHMEM_REF)` (which can occur when using your annotated type with a version of *Connext Micro* that doesn't support Zero Copy), two things will happen:

1. *Connext Micro* will not create a shared memory region for data samples, and;
2. Calls to `FooDataWriter_get_loan()` will fail with `PRECONDITION_NOT_MET`.

However, the *DataWriter* will still be created and can be used to send samples without using Zero Copy transfer.

3. Create a `FooDataWriter` for the above type.
4. Get a loan on a sample using `FooDataWriter_get_loan()`.
5. Write a sample using `FooDataWriter_write()`.

For more information, see the example `HelloWorld_zero_copy`, or generate an example for a type annotated with `@transfer_mode(SHMEM_REF)`:

```
rtiddsgen -example -micro -language C HelloWorld.idl
```

## Writing samples

The following code illustrates how to write samples annotated with `@transfer_mode(SHMEM_REF)`:

```
for (int i = 0; i < 10; i++)
{
    Foo* sample;
    DDS_ReturnCode_t dds_rc;
    /* NEW API
       IMPORTANT - call get_loan each time when writing a NEW sample
    */
    dds_rc = FooDataWriter_get_loan(hw_datawriter, &sample);

    if (dds_rc != DDS_RETCODE_OK)
    {
        printf("Failed to get a loan\n");
        return -1;
    }

    /* After this function returns with DDS_RETCODE_OK,
```

(continues on next page)



(continued from previous page)

```
    * the middleware owns the sample
    */
    dds_rc = FooDataWriter_write(hw_datawriter, sample, &DDS_HANDLE_NIL);
}
```

## Reading samples

The following code illustrates how to read samples annotated with `@transfer_mode(SHMEM_REF)`:

```
DDS_ReturnCode_t dds_rc;
dds_rc = FooDataReader_take(...)

/* process sample here */
/* NEW API
    IMPORTANT - is_data_consistent will always return true when ZC v2 is being used
*/

dds_rc = FooDataReader_is_data_consistent(hw_reader,
                                          &is_data_consistent,
                                          sample, sample_info);

if (dds_rc == DDS_RETCODE_OK)
{
    if (is_data_consistent)
    {
        /* Sample is consistent. Processing of sample is valid */
    }
    else
    {
        /* Sample is NOT consistent. Any processing of the sample should
        * be discarded and considered invalid.
        */
    }
}
```

### 3.16.4 Synchronizing samples

Zero Copy v1 and v2 handle sample synchronization differently. The following sections explain these differences in detail.

#### Zero Copy v1 synchronization

In Zero Copy v1, no synchronization exists between the sender (*DataWriter*) and the receiver (*DataReader*) for Zero Copy samples. This means that a sample's content can be invalidated before the receiver has a chance to read it.

For example, consider creating a best-effort *DataWriter* with `max_samples = 1`. When the *DataWriter* is initialized, the middleware pre-allocates a pool of `max_samples + 1` (2) samples in a shared memory region. These samples are loaned to the *DataWriter* when calling `FooDataWriter_get_loan()`.

The following code illustrates this:

```
DDS_ReturnCode_t ddsrsrc;
Foo* sample;

ddsrc = FooDataWriter_get_loan(dw, &sample); /* returns pointer to sample 1 */
sample->value = 10000;
ddsrc = FooDataWriter_write(datawriter, sample, &DDS_HANDLE_NIL);
/*
 * As this is a best-effort writer, the middleware immediately makes
 * this sample available for reuse by another FooDataWriter_get_loan(...) call.
 */

ddsrc = FooDataWriter_get_loan(dw, &sample); /* returns pointer to sample 2 */
sample->value = 20000;
ddsrc = FooDataWriter_write(datawriter, sample, &DDS_HANDLE_NIL);
/*
 * Again, the sample is made available immediately for reuse.
 */

/*
 * At this point, the sample may have been received by the DataReader
 * but not yet presented to the user.
 */

ddsrc = FooDataWriter_get_loan(dw, &sample); /* returns pointer to sample 1 */
/*
 * sample->value will now contain 10000 because the sample is reused
 * from a pool with only 2 buffers.
 *
 * Additionally, references to both sample 1 and sample 2 might already
 * have been received by the middleware on the *DataReader* side and stored
 * in its internal cache. However, these samples may not yet have been delivered
 * to the application. If sample->value is modified to 999 at this point,
 * a subsequent call to *read()* or *take()* from the Subscriber will return 999,
```

(continues on next page)

(continued from previous page)

```

* not the expected 10000. This happens because both the Publisher and Subscriber
* share the same memory region.
*
* Use `FooDataReader_is_data_consistent` to verify sample consistency and avoid
* this issue.
*
* Note: A sample becomes invalidated right after `FooDataWriter_get_loan(dw, &sample)`
* is completed. If the sample address has already been written to and has not yet
* been read by the receiver, the previously written data is invalidated.
*/

ddsrc = FooDataWriter_write(datawriter, sample, &DDS_HANDLE_NIL);

```

### Zero Copy v2 synchronization

Zero Copy v2 provides synchronization between the *DataWriter* and *DataReader*. Since the queue size is limited, samples are reused once the queue is full. However, if a *DataWriter* modifies a sample before the *DataReader* has accessed it, that sample will not be presented to the user. Additionally, samples currently being read by the *DataReader* are locked, preventing the *DataWriter* from accessing them.

Consider the following example with `max_samples = 1` (internally, the middleware will allocate 2 samples):

```

ddsrc = FooDataWriter_get_loan(dw, &sample); /* returns pointer to sample 1 */
sample->value = 10000;
ddsrc = FooDataWriter_write(datawriter, sample, &DDS_HANDLE_NIL);

ddsrc = FooDataWriter_get_loan(dw, &sample); /* returns pointer to sample 2 */
sample->value = 20000;
ddsrc = FooDataWriter_write(datawriter, sample, &DDS_HANDLE_NIL);

/* Both samples are now available to the user, but the Reader may not have accessed them
↳yet. */

ddsrc = FooDataWriter_get_loan(dw, &sample); /* returns pointer to sample 1 */
sample->value = 30000;
ddsrc = FooDataWriter_write(datawriter, sample, &DDS_HANDLE_NIL);

```

If the *DataReader* takes all the samples:

```

FooDataReader_take(reader,
    &sample_seq, &info_seq,
    DDS_LENGTH_UNLIMITED,
    DDS_ANY_SAMPLE_STATE,
    DDS_ANY_VIEW_STATE,
    DDS_ANY_INSTANCE_STATE);

```

It will only receive two samples with values of 20000 and 30000, respectively.

If both the samples are currently being accessed by the user by calling `FooDataReader_read()` or `FooDataReader_take()`, they will be locked. Any attempt to call `FooDataWriter_get_loan()` on the *DataWriter* will result in an `OUT_OF_RESOURCES` error.

### 3.16.5 Caveats

- After you call `FooDataWriter_write()`, the middleware takes ownership of the sample. It is no longer safe to make any changes to the sample that was written. If, for whatever reason, you call `FooDataWriter_get_loan()` but never write the sample, you must call `FooDataWriter_discard_loan()` to return the sample back to the `FooDataWriter`. Otherwise, subsequently calling `FooDataWriter_get_loan()` may fail, because the `FooDataWriter` has no samples to loan.
- The current maximum supported sample size is a little under the maximum value of a signed 32-bit integer. For that reason, do not use any samples greater than 2000000000 bytes.

## 3.17 FlatData Language Binding

This section is organized as follows:

- *Overview*
- *Getting Started*
- *Further Information*

### 3.17.1 Overview

*RTI Connext Micro* supports the FlatData<sup>TM</sup> language binding in the same manner as *RTI Connext*. However, *Connext Micro* only supports the FlatData language binding for traditional C++ APIs, whereas *RTI Connext* also supports it for the Modern C++ API. The FlatData language binding is not supported for the C language binding.

### 3.17.2 Getting Started

The best way to start is to generate an example by creating an example IDL file `HelloWorld.idl` containing the following IDL type:

```
@final
@language_binding(FLAT_DATA)
struct HelloWorld
{
    long a;
}
```

Next, run:

```
rtiddsgen -example -micro -language C++ HelloWorld.idl
```

### 3.17.3 Further Information

For more details about this feature, please see [the FlatData Language Binding section](#) in the [RTI Connex DDS Core Libraries User's Manual](#) (available [here](#) if you have Internet access).

For details on how to build and read a FlatData sample, see [FlatData](#).

## 3.18 Application Generation Using XML

*RTI Connex Micro*'s Application Generation feature enables you to specify an application in XML. It simplifies and accelerates application development by enabling the creation of DDS *Entities* (and registration of the factories) used in an application by compiling an XML configuration file, linking the result to an application, and calling a single API. Once created, all *Entities* can be retrieved from the application code using standard “lookup\_by\_name” operations so that they can be used to read and write data. UDP transport, DPDE (Dynamic Participant Dynamic Endpoint), and DPSE (Dynamic Participant Static Endpoint) discovery configuration can also be configured as needed. C or C++ source code is generated from the XML configuration and compiled with the application.

Once you have your XML file definition, you must use the Micro Application Generator (MAG) tool to load the XML file definition into *Connex Micro*. MAG is needed because *Connex Micro* does not include an XML parser (this would significantly increase code size and amount of memory needed). MAG generates C source code from the XML configuration that you must then compile with the application. The generated C source code contains the same information as the XML configuration file. The generated C source code can be used from both the [C API Reference](#) and [C++ API Reference](#).

The *Connex Micro* Application Generation is enabled by default in this release when compiling with `rtime-make`. However, future releases may disable the feature by default. Thus, it is advised to always compile with the *Connex Micro* Application Generation feature enabled (`-DRTIME_DDS_ENABLE_APPGEN=1` to CMake).

### 3.18.1 Defining an Application in XML

Each *Entity* configured in the XML file is given a name. This name is used to retrieve the entities at runtime using the *Connex Micro* API.

In the XML file, you need to distinguish between two names:

- Configuration name: The name of a specific *Entity*'s configuration. It is given by the name attribute of the corresponding element.
- Entity name in the *Entity*'s QoS: The Entity name in the *Entity*'s QoS.

At runtime, the *Entity* will be created using the Entity name in the *Entity*'s QoS; the configuration name will be used if this is an empty string.

The attribute multiplicity indicates that a set of *Entities* should be created from the same configuration. Since each *Entity* must have a unique name, the system will automatically append a number to the Entity name in the *Entity*'s QoS (or, if it is an empty string, the configuration name)

to obtain the Entity name. For example, if we specified a multiplicity of “N”, then for each index “i” between 0 and N-1, the system will assign Entity names according to the table below:

Entity Name	Index: i
“configuration_name”	0
“configuration_name#i”	[1,N-1]

That is, the *Entity* name followed by the token “#” and an index.

See *Generate Type-Support Code from the Type Definition* for an example XML file.

### Important Points

Applications can create a *RTI Connex Micro Entity* from a *DomainParticipant* configuration described in the XML configuration file. All the *Entities* defined by such *DomainParticipant* configuration are created automatically as part of the creation. In addition, multiple *DomainParticipant* configurations may be defined within a single XML configuration file.

All the *Entities* created from a *DomainParticipant* configuration are automatically assigned an entity name. *Entities* can be retrieved via “lookup\_by\_name” operations specifying their name. Each *Entity* stores its own name in the QoS policies of the *Entity* so that they can be retrieved locally (via a lookup) and communicated via discovery.

A configuration file is not tied to the application that uses it. Different applications may run using the same configuration file. A single file may define multiple *DomainParticipant* configurations. Normally, a single application will instantiate only one *DomainParticipant*, but an application can instantiate as many *DomainParticipants* as needed.

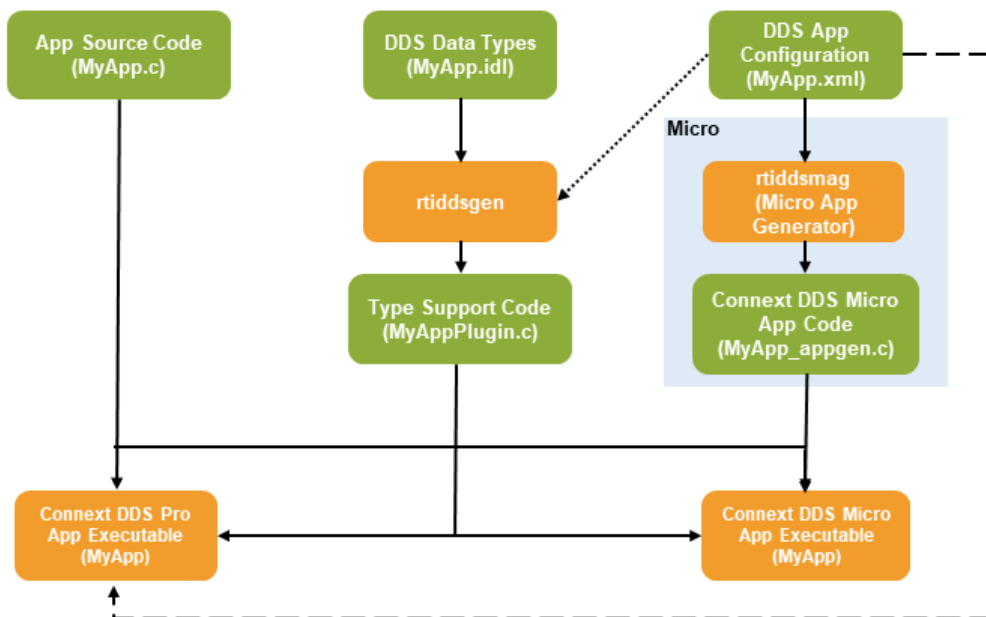
Changes in the XML configuration file require regenerating the C/C++ source code and recompiling the application.

### 3.18.2 Generating the Application from XML

*Connex Micro* comes with a tool, the Micro Application Generator (MAG). This tool is used to generate supporting files to create XML-defined applications at runtime.

#### Micro Application Generator (MAG) Tool

Micro Application Generator (MAG) is a required tool to configure *Connex Micro* applications by generating code from an XML configuration file; it creates DDS entities and registers all the components needed for a *Connex Micro*-based application. MAG can process your own XML configuration file, or it can process an *XML-Based Application Creation* file that you created for *RTI Connex Professional*.



*Connext Micro* Application Generation, in combination with MAG, enables two important use cases:

- Users who may eventually develop with *Connext Micro*, but who haven't determined their final platform, can prototype applications on a generic platform and validate that the QoS and DDS Entity configuration is within scope of what *Connext Micro* supports. MAG ignores fields in the XML file that *Connext Micro* doesn't use (and produces an error for the few fields it cannot ignore; see "Unsupported values" in *Errors Caused by Invalid Configurations and QoS*).
- Users who want to develop directly with *Connext Micro* can simplify their development efforts through shared XML files that can be configuration-managed. This reduces the burden on system integrators who want to configure *Connext Micro* systems without having to manually code in static configurations.

The main features of MAG are:

- Generates code for the languages supported by *Connext Micro*: C and C++.
- Automatically configures the remote entities that are needed to communicate with applications that use static discovery.
- Automatically tries to use the default values used by *Connext Micro*, to reduce the size of the generated code.
- Optimizes the components used by your application. By default, MAG generates code that will unregister transports that your application is not using.
- Ignores fields and transports not supported by *Connext Micro* (any fields or transports not described in the API Reference) and raises errors for things it can't ignore. See *Errors Caused by Invalid Configurations and QoS*.

**Note:**

- MAG has been tested with Java 17.0.6, which is included in the *Connext Professional* installation.
  - MAG does not support customizable templates. (It doesn't support the functionality described in [Customizing the Generated Code](#) in the *Code Generator User's Manual*.)
- 

## Generating the Application with MAG

### Running MAG

To run the MAG tool, use the following command:

For example, on a Windows system:

```
<RTIMEHOME>\rtiddsmag\scripts\rtiddsmag.bat -language C -referencedFile HelloWorldQos.  
↪.xml HelloWorld.xml
```

For example, on a Linux or macOS system:

```
<RTIMEHOME>/rtiddsmag/scripts/rtiddsmag -language C -referencedFile HelloWorldQos.xml  
↪HelloWorld.xml
```

Please refer to *MAG Command-Line Options* for valid command-line options.

### Generated Files

The following table shows the files that MAG creates for an example XML file, **HelloWorld.xml** (which contains the application definition) and a referenced file, **HelloWorldQos.xml** (which contains the QoS definition). This second file is optional; you can define the QoS in the application file.

---

**Note:** Changes in the XML configuration file require regenerating the C/C++ source code and recompiling the application.

---



Table 3.11: C and C++ Files Created for Example HelloWorld.xml

Generated Files	Description
HelloWorldAppgen.h (C and C++)	Generated code for each DDS <i>Entity</i> and its run-time components.
HelloWorldAppgen.c (C and C++)	Generated code for each Entity Model; also contains the values of each array used in the header file.
HelloWorldAppgen_plugin.h (C++ only)	Header file that contains the declarations of all the wrappers.
HelloWorldAppgen_plugin.cxx (C++ only)	A wrapper for the <code>_get()</code> call ( <code>get_plugin_type</code> ): <pre>struct DDS_TypePluginI *HelloWorldPlugin_get_cpp(void) {     return HelloWorldPlugin_get(); }</pre>

**Warning:** You should not modify the generated code. MAG will overwrite your modifications when it regenerates the C/C++ code from XML if the `-replace` argument is used.

## MAG Command-Line Options

The following table shows the options available when using *rtiddsmag* to generate code for *Connext Micro* applications.

Table 3.12: Command-Line Options for *rtiddsmag*

Option	Description
<code>-d &lt;outdir&gt;</code>	Generates the output in the specified directory. By default, MAG will generate files in the directory where the XML file is located.
<code>-dontAddLocations</code>	Use this flag to avoid adding the input file location of fields into the generated files. By default (when this flag is not used), MAG will add the location where an entity was defined in the XML file. The location will be placed above the definition of that entity in the generated code.
<code>-dontOptimizeSE</code>	Use this flag to avoid static endpoint discovery optimization. Then MAG will include all <i>DataWriters</i> and <i>DataReaders</i> when calculating the remote entities. By default (when this option is not used) MAG will optimize the number of remote entities by only including Data Writers and <i>DataReaders</i> that use the same Topic in the remote model.

continues on next page

Table 3.12 – continued from previous page

Option	Description
<code>-dontUpdateResourceLimits</code>	Use this flag to avoid automatically updating the resource limit settings for the <code>DomainParticipantFactory</code> , <i>DomainParticipants</i> , <i>DataReaders</i> , and <i>DataWriters</i> . <b>Note:</b> The use of this flag for the <code>DomainParticipantFactory</code> is currently not supported. By default (when this flag is not used), MAG will update the resource limits so it will at least be able to support the entities defined in the XML file. If your applications communicate with more remote entities than the ones specified in the XML file, you might need to manually update them.
<code>-dontUseDefaultValues</code>	Use this flag to avoid automatically generating code using default QoS policy values when possible. By default (when this flag is not used), MAG will check whether the values that are set in every element of the QoS policies for each entity are the same as the defaults used by <i>Connex Micro</i> . If that's the case, the generated code will contain the default values for those policies, instead of the values set by the user.
<code>-dpdeName &lt;name&gt;</code>	Specifies the name used by MAG when registering a DPDE discovery plugin. By default, this name is <code>dpde</code> . <i>rtiddsmag</i> will also append a <i>DomainParticipant</i> number to the default name for each unique configuration.
<code>-dpseName &lt;name&gt;</code>	Specifies the name used by MAG when registering a DPSE discovery plugin. By default, this name is <code>dpse</code> . <i>rtiddsmag</i> will also append a <i>DomainParticipant</i> number to the default name for each unique configuration.
<code>-help</code>	Prints out the command-line options for MAG.
<code>-idlFile &lt;file&gt;</code>	Specifies the IDL file name used by <i>rtiddsgen</i> to generate the code. This value is used by MAG to specify the Plugin header generated by <i>rtiddsgen</i> . By default, MAG uses the name of the XML file.
<code>-inputXml &lt;file&gt;</code>	Specifies the XML configuration file used to generate code. The XML configuration file can be passed directly to MAG without using the <code>-inputXml</code> option, by default MAG knows that any argument with no option is the input file.
<code>-language &lt;C C++&gt;</code>	Specifies the language to use for the generated files. The default language is C.
<code>-onlyValidate</code>	Causes MAG to just validate the input file. It will not generate any code.

continues on next page

Table 3.12 – continued from previous page

Option	Description
<code>-outputFinalQoS</code> <code>&lt;QosLibrary::QosProfile&gt;</code>	Use this flag to display the final values of the specified QoS profile after applying inheritance. Although MAG currently doesn't generate code to set the QoS for <i>Connex Micro</i> , using this flag will determine the final values in the profile after applying inheritance. For complex XML files, with multiple levels of inheritance, it might be a challenge to determine the final QoS values. Using this flag simplifies the process. <b>Note:</b> This option does not check whether or not the final values are supported by <i>Connex Micro</i> .
<code>-pskName &lt;name&gt;</code>	Change the name of the pre-shared key (PSK) component for the <i>Lightweight Security Plugin</i> . The default name is <code>psk</code> . <i>rtiddsmag</i> will also append a <i>DomainParticipant</i> number to the default name for each unique configuration.
<code>-referencedFile &lt;file&gt;</code>	Specifies a file which is referenced from the one being used to generate code. In general, it is recommended to split the application definition from the QoS definition. This way, the QoS can be shared among various applications. <b>Note:</b> Can be repeated. <code>-referencedFile &lt;file1&gt;</code> <code>-referencedFile &lt;file2&gt; ...</code>
<code>-replace</code>	Use this flag to overwrite existing generated files.
<code>-shmemName &lt;name&gt;</code>	Specifies the name used by MAG when registering a shared memory (SHMEM) transport plugin. By default, this name is <code>shmem</code> . <i>rtiddsmag</i> will also append a <i>DomainParticipant</i> number to the default name for each unique configuration.
<code>-udpName &lt;name&gt;</code>	Specifies the name used by MAG when registering a UDP transport plugin. By default, this name is <code>udp</code> . <i>rtiddsmag</i> will also append a <i>DomainParticipant</i> number to the default name for each unique configuration.
<code>-verbosity [1-4]</code>	Sets the MAG verbosity: 1: Exceptions. 2: Exceptions and warnings. 3: Exceptions, warnings, and information. <b>(Default)</b> 4: Exceptions, warnings, information, and debug.
<code>-version</code>	Displays the version of MAG being used, such as 4.0.1.
<code>-zCopyName &lt;zcopy name&gt;</code>	Specifies the name MAG uses when registering a Zero Copy v2 transport plugin. By default, this name is <code>zcopy</code> . <i>rtiddsmag</i> will also append a <i>DomainParticipant</i> number to the default name for each unique configuration.

## Integrating Generated Files into Your Application's Build

Integrating the generated files into your application is as easy as including the generated files **HelloWorldAppgen.h** and **HelloWorldAppgen.c** in your application. If your application uses C++, you will also need to include **HelloWorldAppgen\_plugin.h** and **HelloWorldAppgen\_plugin.cxx**.

Then you can create entities using the standard `DDS_DomainParticipantFactory_create_participant_from_cor` operation and retrieve all the entities from your application code using the standard `lookup_<entity>_by_name()` operations, such as `lookup_datawriter_by_name()`. For details on these operations, see the [DomainParticipantFactory](#) module in the *Connext Micro* API reference HTML documentation.

### 3.18.3 Creating the Application

We'll create an example application using *rtiddsgen* as seen in *Data Types* in the *User's Manual* and learn from the generated files.

#### Generate Type-Support Code from the Type Definition

The first step is to describe the data type in a programming language-neutral manner. Three languages are supported by *RTI Code Generator*: XML, IDL, and XSD. These three languages provide equivalent type-definition capabilities, so you can choose whichever one you prefer. You can even transform between one of these three languages and another with *RTI Code Generator*. That said, since the rest of the configuration files use XML, it is often more convenient to also use XML to describe the data types, so they can be shared or moved to other XML configuration files.

Create a file called **HelloWorld.xml** that contains the following XML content:

```
<types>
  <const name="MAX_NAME_LEN" type="long" value="64"/>
  <const name="MAX_MSG_LEN" type="long" value="128"/>
  <struct name="HelloWorld">
    <member name="sender" type="string" stringMaxLength="MAX_NAME_LEN" key="true"/>
    <member name="message" type="string" stringMaxLength="MAX_MSG_LEN"/>
    <member name="count" type="long"/>
  </struct>
</types>
```

We'll use this data type for the following example. The data associated with the *HelloWorld Topic* consists of two strings and a numeric counter:

1. The first string contains the name of the sender of the message. This field is marked as the "key" since it signals the identity of the data-object.
2. The second string contains a message.
3. The third field is a simple counter, which the application increments with each message.

Once the type has been defined, we use *rtiddsgen* to generate the code for the *HelloWorld* data type.

We will generate a DPDE example. From your command shell, change to the directory where you created **HelloWorld.xml** and type:

To generate code with **rtiddsgen**:

- On a Windows system:

```
<RTIMEHOME>\rtiddsgen\scripts\rtiddsgen.bat -example -exampleTemplate mag/dpde -  
↪language C HelloWorld.xml -replace
```

- On a Linux or macOS system:

```
<RTIMEHOME>/rtiddsgen/scripts/rtiddsgen -example -exampleTemplate mag/dpde -  
↪language C HelloWorld.xml -replace
```

After running *rtiddsgen*, you will see **HelloWorldQos.xml** and the following files and their associated header files in the **HelloWorld\_mag\_dpde** directory:

- **HelloWorld.c**
- **HelloWorldPlugin.c**
- **HelloWorldSupport.c**
- **HelloWorldAppgen.c**
- **HelloWorld\_publisher.c**
- **HelloWorld\_subscriber.c**
- **HelloWorldApplication.c**

The most notable files are **HelloWorld.h** and **HelloWorldPlugin.h**:

- **HelloWorld.h** contains the declaration of the C structure, built according to the specification in the XML file:

```
#define MAX_NAME_LEN (64L)  
#define MAX_MSG_LEN (128L)  
typedef struct HelloWorld  
{  
    DDS_String    sender;  
    DDS_String    message;  
    DDS_Long      count;  
} HelloWorld;
```

- **HelloWorldPlugin.h** contains the `TypePlugin_get()` function that MAG will use when generating the code to create all the DDS entities:

```
NDDSUSERDllExport extern struct NDDS_Type_Plugin*  
    HelloWorldTypePlugin_get(void);
```

## Generate DDS Entities from the System Definition

This step uses *rtiddsmag* to generate code to support the creation of DDS entities using Application Generation in *Connext Micro*.

*rtiddsmag* supports C and C++. We will generate the DPDE example.

---

**Note:** Type code doesn't need to exist when running *rtiddsmag*. However, we'll be using the HelloWorldQos.xml and HelloWorld.xml files generated from the *Generate Type-Support Code from the Type Definition* section.

---

From your command shell, change to the directory where you created **HelloWorld.xml** and type:

### To generate code with *rtiddsmag*:

- On a Windows system:

```
<RTIMEHOME>\rtiddsmag\scripts\rtiddsmag.bat -language C -referencedFile
↪HelloWorldQos.xml HelloWorld.xml
```

- On a Linux or macOS system:

```
<RTIMEHOME>/rtiddsmag/scripts/rtiddsmag -language C -referencedFile HelloWorldQos.
↪xml HelloWorld.xml
```

We will examine the content of the generated files in the next section.

## Examining the application

*rtiddsmag* generates two applications: **HelloWorld\_publisher**, which writes the *Topic*, HelloWorldTopic, and **HelloWorld\_subscriber**, which subscribes to that *Topic*.

*rtiddsmag* also generates the following examples from the DPSE and the DPDE directories:

- **Domain Participant “HelloWorldDPDEPubDP”**

This application defines a publisher which uses DPDE discovery.

The application has one named “HelloWorldDPDEPubDP”, one named “HelloWorldDPDEPub”, and one named “HelloWorldDPDEDW” which uses topic name “Example HelloWorld”. The application registers one type with name “HelloWorld” and defines one with name “Example HelloWorld” which uses the type “HelloWorld”.

- **Domain Participant “HelloWorldDPDESubDP”**

This application defines a subscriber which uses DPDE discovery.

The application has one named “HelloWorldDPDESubDP”, one named “HelloWorldDPDESub”, and one named “HelloWorldDPDEDR” which uses topic name “Example HelloWorld”. The application registers one type with name “HelloWorld” and defines one with name “Example HelloWorld” which uses the type “HelloWorld”.

## Call API to Create DomainParticipant

Call the API `create_participant_from_config()` to create the application from the generated XML definition. This API creates the *DomainParticipant*; all applications start with the *DomainParticipant*. This API receives the configuration name and creates all the *Entities* defined by that configuration.

## Retrieve Entities by Name

After *DomainParticipant* creation, you can retrieve the defined *Entities* by using `lookup_by_name()` operations.

### 3.18.4 Examine example XML configuration files

An entire **HelloWorld.xml** file is shown below. Let's review its content to see how this scenario was constructed. The main sections in the file are:

- *Type Definition*
- *Domain Definition*
- *DomainParticipant Definition*

```
<?xml version="1.0"?>
<dds xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:noNamespaceSchemaLocation="http://community.rti.com/schema/current/rti_dds_profiles.
  ↪xsd">
  <!-- Type Definition -->
  <types>
    <const name="MAX_NAME_LEN" type="int32" value="64"/>
    <const name="MAX_MSG_LEN" type="int32" value="128"/>
    <struct name="HelloWorld">
      <member name="sender" type="string" stringMaxLength="MAX_NAME_LEN" key="true
      ↪"/>
      <member name="message" type="string" stringMaxLength="MAX_MSG_LEN"/>
      <member name="count" type="int32"/>
    </struct>
  </types>
  <!-- Domain Library -->
  <domain_library name="HelloWorldLibrary">
    <domain name="HelloWorldDomain" domain_id="0">
      <register_type name="HelloWorldType" type_ref="HelloWorld">
      </register_type>
      <topic name="HelloWorldTopic" register_type_ref="HelloWorldType">
        <registered_name>HelloWorldTopic</registered_name>
      </topic>
    </domain>
  </domain_library>
  <!-- Participant Library -->
  <domain_participant_library name="HelloWorldAppLibrary">
```

(continues on next page)

(continued from previous page)

```

<domain_participant name="HelloWorldDPDEPubDP"
  domain_ref="HelloWorldLibrary::HelloWorldDomain">
  <publisher name="HelloWorldDPDEPub">
    <data_writer topic_ref="HelloWorldTopic" name="HelloWorldDPDEDW">
      <datawriter_qos base_name="QosLibrary::DPDEProfile"/>
    </data_writer>
  </publisher>
  <domain_participant_qos base_name="QosLibrary::DPDEProfile"/>
</domain_participant>
<domain_participant name="HelloWorldDPDESubDP"
  domain_ref="HelloWorldLibrary::HelloWorldDomain">
  <subscriber name="HelloWorldDPDESub">
    <data_reader topic_ref="HelloWorldTopic" name="HelloWorldDPDEDR">
      <datareader_qos base_name="QosLibrary::DPDEProfile"/>
    </data_reader>
  </subscriber>
  <domain_participant_qos base_name="QosLibrary::DPDEProfile"/>
</domain_participant>
<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>
  <domain_participant_qos base_name="QosLibrary::DPSEProfile"/>
</domain_participant>
<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>
  <domain_participant_qos base_name="QosLibrary::DPSEProfile"/>
</domain_participant>
</domain_participant_library>
</dds>

```

## Type Definition

*rtiddsmag* doesn't use the types section of the XML file to generate any code. This section is used by *rtiddsgen* to generate the code to support the direct use of the structure 'HelloWorld' from application code (see *Generate Type-Support Code from the Type Definition*).

```

<types>
  <const name="MAX_NAME_LEN" type="int32" value="64"/>
  <const name="MAX_MSG_LEN" type="int32" value="128"/>
  <struct name="HelloWorld">
    <member name="sender" type="string" stringMaxLength="MAX_NAME_LEN" key="true"/>

```

(continues on next page)



(continued from previous page)

```

    <member name="message" type="string" stringMaxLength="MAX_MSG_LEN"/>
    <member name="count" type="int32"/>
  </struct>
</types>

```

## Domain Definition

The domain section defines the system's *Topics* and their corresponding data types. To define a *Topic*, the associated data type must be registered with the domain, giving it a registered type name. The registered type name is used to refer to that data type within the domain when the *Topic* is defined.

In this example, the configuration file registers the previously defined HelloWorld type under the name HelloWorldType. Then it defines a *Topic* named HelloWorldTopic, which is associated with the registered type, referring to its registered name, HelloWorldType. The value used in `TypePlugin_get` depends on how the registration of the data-type is configured inside the domain:

1. If a `<register_type>` tag is specified *without* a `type_ref` attribute, the value of `TypePlugin_get` is generated from the `<register_type>` tag plus the string "Plugin\_get".
2. If a `<register_type>` tag is specified *with* a `type_ref` attribute, the value of `TypePlugin_get` is generated from that attribute plus the string "TypePlugin\_get". Our example has `type_ref` = "HelloWorld", so the value of `TypePlugin_get` will be HelloWorldTypePlugin\_get.

```

<!-- Domain Library -->
<domain_library name="HelloWorldLibrary">
  <domain name="HelloWorldDomain" domain_id="0">
    <register_type name="HelloWorld" type_ref="HelloWorld">
    </register_type>

    <topic name="HelloWorldTopic" register_type_ref="HelloWorld">
      <registered_name>HelloWorldTopic</registered_name>
    </topic>
  </domain>
</domain_library>

```

`rtiddsmag` generates the following code for each entity that uses this *Topic*:

- HelloWorldAppgen.c

```

const struct APPGEN_TypeRegistrationModel
HelloWorldAppLibrary_HelloWorldDPDEPubDP_type_registrations[1] =
{
  {
    "HelloWorldType", /* registered_type_name */
    HelloWorldTypePlugin_get /* get_type_plugin */
  }
};
const struct APPGEN_TopicModel
HelloWorldAppLibrary_HelloWorldDPDEPubDP_topics[1] =

```

(continues on next page)

(continued from previous page)

```

{
    {
        "HelloWorldTopic", /* topic_name */
        "HelloWorldType", /* type_name */
        DDS_TopicQos_INITIALIZER /* topic_qos*/
    }
};

```

These two structures are used in the *DomainParticipant* definition, where they will be registered by *Connext Micro* when calling the Micro Application Generation API.

- **HelloWorldAppgen.h**

```

extern const struct APPGEN_TypeRegistrationModel
HelloWorldAppLibrary_HelloWorldDPDEPubDP_type_registrations[1];

extern const struct APPGEN_TopicModel
HelloWorldAppLibrary_HelloWorldDPDEPubDP_topics[1];

#define RTI_APP_GEN__DP_HelloWorldAppLibrary_HelloWorldDPDEPubDP \
{ \
    ...
    1UL, /* type_registration_count */ \
    HelloWorldAppLibrary_HelloWorldDPDEPubDP_type_registrations, /* type_
↪registrations*/ \
    1UL, /* topic_count */ \
    HelloWorldAppLibrary_HelloWorldDPDEPubDP_topics, /* topics */ \
    ...
}

```

---

**Note:** *Connext Micro* automatically registers the types that *rtiddsmag* generates. This means the content inside the Domain definition must match the types generated by *rtiddsgen*.

---

## DomainParticipant Definition

The *DomainParticipant* section defines the *DomainParticipants* in the system and the *DataWriters* and *DataReaders* that each *DomainParticipant* has. *DomainParticipants* are defined within the `<domain_participant_library>` tag.

Each *DomainParticipant*:

- Has a unique name (within the library) which will be used later by the application that creates it.
- Is associated with a domain, which defines the `domain_id`, *Topics*, and the data types the *DomainParticipant* will use.
- Defines the *Publishers* and *Subscribers* within the *DomainParticipant*. *Publishers* contain *DataWriters*, *Subscribers* contain *DataReaders*.

- Defines the set of *DataReaders* it will use to read data. Each *DataReader* has a QoS and a unique name which can be used from application code to retrieve it.
- Defines the set of *DataWriters* it will use to write data. Each *DataWriter* has a QoS and a unique name which can be used from application code to retrieve it.
- Optionally, the *DomainParticipants*, *Publishers*, *Subscribers*, *DataWriters*, and *DataReaders* can specify a QoS profile that will be used to configure them.

The example below defines four *DomainParticipants*, two of them (HelloWorldDPDEPubDP and HelloWorldDPDESubDP) use Dynamic Participant/Dynamic Endpoint (DPDE) and the other two (HelloWorldDPSEPubDP and HelloWorldDPSESubDP) use Dynamic Participant/Static Endpoint (DPSE) discovery:

```
<!-- Participant Library -->
<domain_participant_library name="HelloWorldAppLibrary">
  <domain_participant name="HelloWorldDPDEPubDP"
    domain_ref="HelloWorldLibrary::HelloWorldDomain">
    <publisher name="HelloWorldDPDEPub">
      <data_writer topic_ref="HelloWorldTopic" name="HelloWorldDPDEDW">
        <datawriter_qos base_name="QosLibrary::DPDEProfile"/>
      </data_writer>
    </publisher>
    <domain_participant_qos base_name="QosLibrary::DPDEProfile"/>
  </domain_participant>

  <domain_participant name="HelloWorldDPDESubDP"
    domain_ref="HelloWorldLibrary::HelloWorldDomain">
    <subscriber name="HelloWorldDPDESub">
      <data_reader topic_ref="HelloWorldTopic" name="HelloWorldDPDEDR">
        <datareader_qos base_name="QosLibrary::DPDEProfile"/>
      </data_reader>
    </subscriber>
    <domain_participant_qos base_name="QosLibrary::DPDEProfile"/>
  </domain_participant>

  <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>
    <domain_participant_qos base_name="QosLibrary::DPSEProfile"/>
  </domain_participant>

  <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>
```

(continues on next page)

(continued from previous page)

```

    <domain_participant_qos base_name="QosLibrary::DPSEProfile"/>
  </domain_participant>
</domain_participant_library>

```

Examining the XML, we see that:

- Each *DomainParticipant* is bound to the Domain, HelloWorldLibrary::HelloWorldDomain.
- The two *DomainParticipants* that use DPDE as their discovery mechanism inherit from the profile QosLibrary::DPDELibrary, while the other two that use DPSE as their discovery mechanism inherit from QosLibrary::DPSELibrary.
- Each *DomainParticipant* contains a single *Publisher* or *Subscriber*, which in turn contains a single *DataWriter* or *DataReader* that inherits from QosLibrary::DPDELibrary or QosLibrary::DPSELibrary, depending on the discovery mechanism used by its *DomainParticipant*.
- Each *DataWriter* writes the *Topic* HelloWorldTopic, which is defined in the domain HelloWorldLibrary::HelloWorldDomain. Each *DataReader* reads the same *Topic*.

Since both Dynamic *DomainParticipants* (those which are using DPDE as their discovery mechanism) are in the same domain and the *DataWriter* writes the same *Topic* that the *DataReader* reads, the two *DomainParticipants* will communicate. This also applies to both static participants (those which are using DPSE as their discovery mechanism); the only difference is that *rtiddsmag* will generate extra code to configure the remote entities (for details, see *Static Discovery*).

Let's look at the content of a *DomainParticipant* definition to explain the code generated by *rtiddsmag*.

```

<domain_participant name="HelloWorldDPDEPubDP"
  domain_ref="HelloWorldLibrary::HelloWorldDomain">
  <publisher name="HelloWorldDPDEPub">
    <data_writer topic_ref="HelloWorldTopic" name="HelloWorldDPDEDW">
      <datawriter_qos base_name="QosLibrary::DPDEProfile"/>
    </data_writer>
  </publisher>
  <domain_participant_qos base_name="QosLibrary::DPDEProfile"/>
</domain_participant>

```

*rtiddsmag* generates the code needed to register each component used by this *DomainParticipant* and unregister those components that are not being used. In our example, for each *DomainParticipant*, *rtiddsmag* registers the discovery transport, *dpde* or *dpse*; registers the UDP transport used by each *DomainParticipant* (since they use the same configuration, only one UDP transport configuration is generated); and unregisters the default UDP and INTRA transports, since they are not being used (these two are the only ones that can be unregistered by *rtiddsmag*).

It also creates the code for each entity. In this case, it generates the code needed to create:

- A *Publisher* named HelloWorldDPDEPub
- A *DataWriter* named HelloWorldDPDEDW
- A *DomainParticipant* named HelloWorldDPDEPubDP

- The QoS used by this *DomainParticipant* (see *QoS Definition*)

## HelloWorldAppgen.c

```

const struct ComponentFactoryUnregisterModel
HelloWorldAppLibrary_HelloWorldDPDEPubDP_unregister_components[2] =
{
    {
        "_udp", /* NETIO_DEFAULT_UDP_NAME */
        NULL, /* udp struct RT_ComponentFactoryProperty** */
        NULL /* udp struct RT_ComponentFactoryListener** */
    },
    {
        "_intra", /* NETIO_DEFAULT_INTRA_NAME */
        NULL, /* _intra struct RT_ComponentFactoryProperty** */
        NULL /* _intra struct RT_ComponentFactoryListener** */
    }
};

struct DPDE_DiscoveryPluginProperty
HelloWorldAppLibrary_HelloWorldDPDEPubDP_dpde[1] =
{
    RTI_APP_GEN___dpde__HelloWorldAppLibrary_HelloWorldDPDEPubDP_dpde1
};

struct UDP_InterfaceFactoryProperty
HelloWorldAppLibrary_HelloWorldDPDEPubDP_udp4[1] =
{
    RTI_APP_GEN___udp4__HelloWorldAppLibrary_HelloWorldDPDEPubDP_udp1
};

const struct ComponentFactoryRegisterModel
HelloWorldAppLibrary_HelloWorldDPDEPubDP_register_components[2] =
{
    {
        "dpde1", /* register_name */
        DPDE_DiscoveryFactory_get_interface, /* register_intf */
        &HelloWorldAppLibrary_HelloWorldDPDEPubDP_dpde[0]._parent, /* register_property_
↪ */
        NULL /* register_listener */
    },
    {
        "udp1", /* register_name */
        UDP_InterfaceFactory_get_interface, /* register_intf */
        &HelloWorldAppLibrary_HelloWorldDPDEPubDP_udp4[0]._parent._parent, /* register_
↪ property */
        NULL /* register_listener */
    }
};

...

const struct APPGEN_DataWriterModel
HelloWorldAppLibrary_HelloWorldDPDEPubDP_publisher_HelloWorldDPDEPub_data_writers[1] =
{

```

(continues on next page)

(continued from previous page)

```

{
    "HelloWorldDPDEDW", /* name */
    1UL, /* multiplicity */
    "HelloWorldTopic", /* topic_name */
    RTI_APP_GEN___DW_QOS_HelloWorldAppLibrary_HelloWorldDPDEPubDP_HelloWorldDPDEPub_
↪HelloWorldDPDEDW /* writer_qos */
}
};
const struct APPGEN_PublisherModel
HelloWorldAppLibrary_HelloWorldDPDEPubDP_publishers[1] =
{
    {
        "HelloWorldDPDEPub", /* name */
        1UL, /* multiplicity */
        DDS_PublisherQos_INITIALIZER, /* publisher_qos */
        1UL, /* writer_count */
        HelloWorldAppLibrary_HelloWorldDPDEPubDP_publisher_HelloWorldDPDEPub_data_
↪writers /* data_writers */
    }
};

```

## HelloWorldAppgen.h

```

extern struct DPDE_DiscoveryPluginProperty HelloWorldAppLibrary_HelloWorldDPDEPubDP_
↪dpde[1];
extern struct UDP_InterfaceFactoryProperty HelloWorldAppLibrary_HelloWorldDPDEPubDP_
↪udpv4[1];
extern const struct ComponentFactoryUnregisterModel
    HelloWorldAppLibrary_HelloWorldDPDEPubDP_unregister_components[2];
extern const struct ComponentFactoryRegisterModel
    HelloWorldAppLibrary_HelloWorldDPDEPubDP_register_components[2];

#define RTI_APP_GEN___DPF_HelloWorldAppLibrary_HelloWorldDPDEPubDP \
{ \
    2UL, /* unregister_count */ \
    HelloWorldAppLibrary_HelloWorldDPDEPubDP_unregister_components, /* unregister_
↪components */ \
    2UL, /* register_count */ \
    HelloWorldAppLibrary_HelloWorldDPDEPubDP_register_components, /* register_components_
↪*/ \
    RTI_APP_GEN___DPF_QOS_QosLibrary_DefaultProfile /* factory_qos */ \
}

extern const struct APPGEN_TypeRegistrationModel
    HelloWorldAppLibrary_HelloWorldDPDEPubDP_type_registrations[1];
extern const struct APPGEN_TopicModel HelloWorldAppLibrary_HelloWorldDPDEPubDP_topics[1];
extern const struct APPGEN_PublisherModel
    HelloWorldAppLibrary_HelloWorldDPDEPubDP_publishers[1];

#define RTI_APP_GEN___DP_HelloWorldAppLibrary_HelloWorldDPDEPubDP \
{ \

```

(continues on next page)

(continued from previous page)

```

    "HelloWorldDPDEPubDP", /* name */ \
    RTI_APP_GEN__DPF_HelloWorldAppLibrary_HelloWorldDPDEPubDP, /* domain_participant_
↪factory */ \
    RTI_APP_GEN__DP_QOS_HelloWorldAppLibrary_HelloWorldDPDEPubDP, /* domain_participant_
↪qos */ \
    OL, /* domain_id */ \
    1UL, /* type_registration_count */ \
    HelloWorldAppLibrary_HelloWorldDPDEPubDP_type_registrations, /* type_registrations */
↪ \
    1UL, /* topic_count */ \
    HelloWorldAppLibrary_HelloWorldDPDEPubDP_topics, /* topics */ \
    1UL, /* publisher_count */ \
    HelloWorldAppLibrary_HelloWorldDPDEPubDP_publishers, /* publishers */ \
    OUL, /* subscriber_count */ \
    NULL, /* subscribers */ \
    OUL, /* remote_participant_count */ \
    NULL, /* remote_participants */ \
    OUL, /* custom_flow_controller_count */ \
    NULL /* custom_flow_controllers */ \
}

```

## QoS Definition

The defined DDS Entities have an associated QoS Policy, which can be defined in a separate file such as **HelloWorldQos.xml** or within the System XML file.

For more information on how to configure DDS Entities in an XML file, see [Configuring QoS with XML](#) (if you have internet access).

See the entire file below. Then we will examine the file section by section, showing the code generated by *rtiddsmag* for the DPSE example.

```

<?xml version="1.0"?>
<dds xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:noNamespaceSchemaLocation="http://community.rti.com/schema/current/rti_dds_profiles.
↪xsd">
  <qos_library name="QosLibrary">
    <qos_profile name="DefaultProfile" is_default_participant_factory_profile="true">

      <!-- Participant Factory Qos -->
      <participant_factory_qos>
        <entity_factory>
          <autoenable_created_entities>>false</autoenable_created_entities>
        </entity_factory>
      </participant_factory_qos>

      <!-- Participant Qos -->
      <domain_participant_qos>
        <discovery>
          <accept_unknown_peers>>false</accept_unknown_peers>

```

(continues on next page)

(continued from previous page)

```

<initial_peers>
  <element>127.0.0.1</element>
  <element>239.255.0.1</element>
</initial_peers>
<enabled_transports>
  <element>udp4</element>
</enabled_transports>
<multicast_receive_addresses>
  <element>udp4://127.0.0.1</element>
  <element>udp4://239.255.0.1</element>
</multicast_receive_addresses>
</discovery>
<default_unicast>
  <value>
    <element>
      <transports>
        <element>udp4</element>
      </transports>
    </element>
  </value>
</default_unicast>
<transport_built_in>
  <mask>UDPv4</mask>
</transport_built_in>
<resource_limits>
  <local_writer_allocation>
    <max_count>1</max_count>
  </local_writer_allocation>
  <local_reader_allocation>
    <max_count>1</max_count>
  </local_reader_allocation>
  <local_publisher_allocation>
    <max_count>1</max_count>
  </local_publisher_allocation>
  <local_subscriber_allocation>
    <max_count>1</max_count>
  </local_subscriber_allocation>
  <local_topic_allocation>
    <max_count>1</max_count>
  </local_topic_allocation>
  <local_type_allocation>
    <max_count>1</max_count>
  </local_type_allocation>
  <remote_participant_allocation>
    <max_count>8</max_count>
  </remote_participant_allocation>
  <remote_writer_allocation>
    <max_count>8</max_count>
  </remote_writer_allocation>
  <remote_reader_allocation>
    <max_count>8</max_count>
  </remote_reader_allocation>

```

(continues on next page)



(continued from previous page)

```

        </remote_reader_allocation>
        <max_receive_ports>32</max_receive_ports>
        <max_destination_ports>32</max_destination_ports>
    </resource_limits>
</domain_participant_qos>
<!-- DataWriter Qos -->
<datawriter_qos>
    <history>
        <depth>32</depth>
    </history>
    <resource_limits>
        <max_instances>2</max_instances>
        <max_samples>64</max_samples>
        <max_samples_per_instance>32</max_samples_per_instance>
    </resource_limits>
    <reliability>
        <kind>RELIABLE_RELIABILITY_QOS</kind>
    </reliability>
    <protocol>
        <rtps_reliable_writer>
            <heartbeat_period>
                <nanosec>250000000</nanosec>
                <sec>0</sec>
            </heartbeat_period>
        </rtps_reliable_writer>
    </protocol>
<!-- transports -->
<unicast>
    <value>
        <element>
            <transports>
                <element>udp4</element>
            </transports>
        </element>
    </value>
</unicast>
</datawriter_qos>
<!-- DataReader Qos -->
<datareader_qos>
    <history>
        <depth>32</depth>
    </history>
    <resource_limits>
        <max_instances>2</max_instances>
        <max_samples>64</max_samples>
        <max_samples_per_instance>32</max_samples_per_instance>
    </resource_limits>
    <reliability>
        <kind>RELIABLE_RELIABILITY_QOS</kind>
    </reliability>
    <reader_resource_limits>

```

(continues on next page)

(continued from previous page)

```

        <max_remote_writers>10</max_remote_writers>
        <max_remote_writers_per_instance>10</max_remote_writers_per_instance>
    </reader_resource_limits>
    <!-- transports -->
    <unicast>
        <value>
            <element>
                <transports>
                    <element>udpv4</element>
                </transports>
            </element>
        </value>
    </unicast>
    <multicast>
        <value>
            <element>
                <receive_address>127.0.0.1</receive_address>
                <transports>
                    <element>udpv4</element>
                </transports>
            </element>
        </value>
    </multicast>
</datareader_qos>
</qos_profile>

<qos_profile name="DPDEProfile" base_name="DefaultProfile">
    <domain_participant_qos>
        <discovery_config>
            <builtin_discovery_plugins>SDP</builtin_discovery_plugins>
        </discovery_config>
    </domain_participant_qos>
</qos_profile>

<qos_profile name="DPSEProfile" base_name="DefaultProfile">
    <domain_participant_qos>
        <discovery_config>
            <builtin_discovery_plugins>DPSE</builtin_discovery_plugins>
        </discovery_config>
    </domain_participant_qos>
</qos_profile>
</qos_library>
</dds>

```

**Note:** *rtiddsmag* only generates code for the QoS policies used by at least one entity, unless the QoS profile has either of the default flags `is_default_participant_factory_profile` or `is_default_qos` set to true.

## DomainParticipant Factory QoS

*rtiddsmag* only generates code for the `<participant_factory_qos>` in the `<qos_profile>` that has the flag `is_default_participant_factory_profile` set to true. The log verbosity can also be configured by using `<verbosity>` inside `<logging>`. For example:

```
<!-- Participant Factory Qos -->
<participant_factory_qos>
  <entity_factory>
    <autoenable_created_entities>false</autoenable_created_entities>
  </entity_factory>
  <resource_limits>
    <max_participants>4</max_participants>
    <max_components>20</max_components>
  </resource_limits>
</participant_factory_qos>
```

*rtiddsmag* generates the following code:

### HelloWorldAppgen.h

```
#define RTI_APP_GEN__DPF_QOS_QosLibrary_DefaultProfile \
{ \
  { /* entity_factory */ \
    DDS_BOOLEAN_FALSE /* autoenable_created_entities */ \
  }, \
  { /* resource_limits */ \
    4L, /* max_participants */ \
    20L /* max_components */ \
  } \
}
#define RTI_APP_GEN__DPF>HelloWorldAppLibrary>HelloWorldDPDEPubDP \
{ \
  ...,
  RTI_APP_GEN__DPF_QOS_QosLibrary_DefaultProfile /* factory_qos */ \
}
```

## DomainParticipant QoS

The example defines a base profile named `DefaultProfile`, which contains the base QoSs used by each *DomainParticipant*. You can see the content of the *DomainParticipant* QoS below.

```
<domain_participant_qos>
  <discovery>
    <accept_unknown_peers>false</accept_unknown_peers>
    <initial_peers>
      <element>127.0.0.1</element>
      <element>239.255.0.1</element>
    </initial_peers>
    <enabled_transports>
```

(continues on next page)

(continued from previous page)

```

        <element>udp4</element>
    </enabled_transports>
    <multicast_receive_addresses>
        <element>udp4://127.0.0.1</element>
        <element>udp4://239.255.0.1</element>
    </multicast_receive_addresses>
</discovery>
<default_unicast>
    <value>
        <element>
            <transports>
                <element>udp4</element>
            </transports>
        </element>
    </value>
</default_unicast>
<transport_builtin>
    <mask>UDPV4</mask>
</transport_builtin>
<resource_limits>
    <local_writer_allocation>
        <max_count>1</max_count>
    </local_writer_allocation>
    <local_reader_allocation>
        <max_count>1</max_count>
    </local_reader_allocation>
    <local_publisher_allocation>
        <max_count>1</max_count>
    </local_publisher_allocation>
    <local_subscriber_allocation>
        <max_count>1</max_count>
    </local_subscriber_allocation>
    <local_topic_allocation>
        <max_count>1</max_count>
    </local_topic_allocation>
    <local_type_allocation>
        <max_count>1</max_count>
    </local_type_allocation>
    <remote_participant_allocation>
        <max_count>8</max_count>
    </remote_participant_allocation>
    <remote_writer_allocation>
        <max_count>8</max_count>
    </remote_writer_allocation>
    <remote_reader_allocation>
        <max_count>8</max_count>
    </remote_reader_allocation>
    <max_receive_ports>32</max_receive_ports>
    <max_destination_ports>32</max_destination_ports>
</resource_limits>
</domain_participant_qos>

```

This *DomainParticipant* is then inherited by two different profiles, which set up the discovery mechanism:

```
<domain_participant_qos>
  <discovery_config>
    <builtin_discovery_plugins>SDP</builtin_discovery_plugins>
  </discovery_config>
</domain_participant_qos>
<domain_participant_qos>
  <discovery_config>
    <builtin_discovery_plugins>DPSE</builtin_discovery_plugins>
  </discovery_config>
</domain_participant_qos>
```

*rtiddsmag* generates the following code for each *DomainParticipant* whose QoS inherits from any of the previous ones, adding those values that are specified in the XML configuration file (which is not the case in our example).

### HelloWorldAppgen.c

```
const char *const HelloWorldAppLibrary_HelloWorldDPDEPubDP_initial_peers[2] =
{
    "127.0.0.1",
    "239.255.0.1"
};
const char *const HelloWorldAppLibrary_HelloWorldDPDEPubDP_discovery_enabled_
↳ transports[3] =
{
    "udp1://",
    "udp1://127.0.0.1",
    "udp1://239.255.0.1"
};
const char *const HelloWorldAppLibrary_HelloWorldDPDEPubDP_transport_enabled_
↳ transports[1] =
{
    "udp1"
};
const char *const HelloWorldAppLibrary_HelloWorldDPDEPubDP_user_traffic_enabled_
↳ transports[1] =
{
    "udp1://"
};
```

### HelloWorldAppgen.h

```
extern const char *const HelloWorldAppLibrary_HelloWorldDPDEPubDP_initial_peers[2];
extern const char *const HelloWorldAppLibrary_HelloWorldDPDEPubDP_discovery_enabled_
↳ transports[3];
extern const char *const HelloWorldAppLibrary_HelloWorldDPDEPubDP_transport_enabled_
↳ transports[1];
extern const char *const HelloWorldAppLibrary_HelloWorldDPDEPubDP_user_traffic_enabled_
↳ transports[1];
```

(continues on next page)

(continued from previous page)

```

#define RTI_APP_GEN___DP_QOS>HelloWorldAppLibrary>HelloWorldDPDEPubDP \
{ \
    { /* entity_factory */ \
        DDS_BOOLEAN_TRUE /* autoenable_created_entities */ \
    }, \
    { /* discovery */ \
        REDA_StringSeq_INITIALIZER_W_LOAN(HelloWorldAppLibrary>HelloWorldDPDEPubDP_
↪initial_peers, 2, 2), /* initial_peers */ \
        REDA_StringSeq_INITIALIZER_W_LOAN(HelloWorldAppLibrary>HelloWorldDPDEPubDP_
↪discovery_enabled_transports, 3, 3), /* enabled_transports */ \
        { \
            { { "dpde1" } }, /* RT_ComponentFactoryId_INITIALIZER */ \
            NDDS_Discovery_Property_INITIALIZER \
        }, /* discovery_component */ \
        DDS_BOOLEAN_FALSE /* accept_unknown_peers */ \
    }, \
    { /* resource_limits */ \
        1L, /* local_writer_allocation */ \
        1L, /* local_reader_allocation */ \
        1L, /* local_publisher_allocation */ \
        1L, /* local_subscriber_allocation */ \
        1L, /* local_topic_allocation */ \
        1L, /* local_type_allocation */ \
        8L, /* remote_participant_allocation */ \
        8L, /* remote_writer_allocation */ \
        8L, /* remote_reader_allocation */ \
        32L, /* matching_writer_reader_pair_allocation */ \
        32L, /* matching_reader_writer_pair_allocation */ \
        32L, /* max_receive_ports */ \
        32L, /* max_destination_ports */ \
        65536, /* unbound_data_buffer_size */ \
        500UL, /* shm_ref_transfer_mode_max_segments */ \
        0L, /* participant_user_data_max_length */ \
        DDS_SIZE_AUTO, /* participant_user_data_max_count */ \
        0L, /* topic_data_max_length */ \
        DDS_SIZE_AUTO, /* topic_data_max_count */ \
        0L, /* publisher_group_data_max_length */ \
        DDS_SIZE_AUTO, /* publisher_group_data_max_count */ \
        0L, /* subscriber_group_data_max_length */ \
        DDS_SIZE_AUTO, /* subscriber_group_data_max_count */ \
        0L, /* writer_user_data_max_length */ \
        DDS_SIZE_AUTO, /* writer_user_data_max_count */ \
        0L, /* reader_user_data_max_length */ \
        DDS_SIZE_AUTO, /* reader_user_data_max_count */ \
        64L, /* max_partitions */ \
        256L, /* max_partition_cumulative_characters */ \
        DDS_LENGTH_UNLIMITED, /* max_partition_string_size */ \
        DDS_LENGTH_UNLIMITED /* max_partition_string_allocation */ \
    }, \
    DDS_ENTITY_NAME_QOS_POLICY_DEFAULT, \
    DDS_WIRE_PROTOCOL_QOS_POLICY_DEFAULT, \

```

(continues on next page)

(continued from previous page)

```

{ /* transports */ \
  REDA_StringSeq_INITIALIZER_W_LOAN>HelloWorldAppLibrary>HelloWorldDPDEPubDP_
↪transport_enabled_transports, 1, 1) /* enabled_transports */ \
}, \
{ /* user_traffic */ \
  REDA_StringSeq_INITIALIZER_W_LOAN>HelloWorldAppLibrary>HelloWorldDPDEPubDP_user_
↪traffic_enabled_transports, 1, 1) /* enabled_transports */ \
}, \
  DDS_TRUST_QOS_POLICY_DEFAULT, \
  DDS_PROPERTY_QOS_POLICY_DEFAULT, \
  DDS_USER_DATA_QOS_POLICY_DEFAULT \
}

```

## Publisher QoS

Our example doesn't specify any value for *Publisher* QoS, however *rtiddsmag* would generate code if it was specified.

## DataWriter QoS

The example defines a base profile named *DefaultProfile*, which contains the base QoSs used by each *DomainParticipant*. You can see the content of the *DataWriter* QoS below.

```

<!-- DataWriter Qos -->
<datawriter_qos>
  <history>
    <depth>32</depth>
  </history>
  <resource_limits>
    <max_instances>2</max_instances>
    <max_samples>64</max_samples>
    <max_samples_per_instance>32</max_samples_per_instance>
  </resource_limits>
  <reliability>
    <kind>RELIABLE_RELIABILITY_QOS</kind>
  </reliability>
  <protocol>
    <rtps_reliable_writer>
      <heartbeat_period>
        <nanosec>250000000</nanosec>
        <sec>0</sec>
      </heartbeat_period>
    </rtps_reliable_writer>
  </protocol>
  <!-- transports -->
  <unicast>
    <value>
      <element>

```

(continues on next page)

(continued from previous page)

```

        <transports>
            <element>udp4</element>
        </transports>
    </element>
</value>
</unicast>
</datawriter_qos>

```

*rtiddsmag* generates the following code:

### HelloWorldAppgen.c

```

const char *const
HelloWorldAppLibrary_HelloWorldDPDEPubDP_HelloWorldDPDEPub_HelloWorldDPDEDW_transport_
↪enabled_transports[1] =
{
    "udp1://"
};

```

### HelloWorldAppgen.h

```

extern const char *const
HelloWorldAppLibrary_HelloWorldDPDEPubDP_HelloWorldDPDEPub_HelloWorldDPDEDW_transport_
↪enabled_transports[1];

#define RTI_APP_GEN___DW_QOS_HelloWorldAppLibrary_HelloWorldDPDEPubDP_HelloWorldDPDEPub_
↪HelloWorldDPDEDW \
{ \
    DDS_DEADLINE_QOS_POLICY_DEFAULT, \
    DDS_LIVELINESS_QOS_POLICY_DEFAULT, \
    { /* history */ \
        DDS_KEEP_LAST_HISTORY_QOS, /* kind */ \
        32L /* depth */ \
    }, \
    { /* resource_limits */ \
        64L, /* max_samples */ \
        2L, /* max_instances */ \
        32L /* max_samples_per_instance */ \
    }, \
    DDS_OWNERSHIP_QOS_POLICY_DEFAULT, \
    DDS_OWNERSHIP_STRENGTH_QOS_POLICY_DEFAULT, \
    DDS_LATENCY_BUDGET_QOS_POLICY_DEFAULT, \
    { /* reliability */ \
        DDS_RELIABLE_RELIABILITY_QOS, /* kind */ \
        { /* max_blocking_time */ \
            0L, /* sec */ \
            100000000L /* nanosec */ \
        } \
    }, \
    DDS_DURABILITY_QOS_POLICY_DEFAULT, \
    DDS_DESTINATION_ORDER_QOS_POLICY_DEFAULT, \
    DDS_TRANSPORT_ENCAPSULATION_QOS_POLICY_DEFAULT, \
}

```

(continues on next page)



(continued from previous page)

```

DDS_DATA_REPRESENTATION_QOS_POLICY_DEFAULT, \
{ /* protocol */ \
  DDS_RTPS_AUTO_ID, /* rtps_object_id */ \
  { /* rtps_reliable_writer */ \
    { /* heartbeat_period */ \
      0L, /* sec */ \
      2500000000L /* nanosec */ \
    }, \
    1L, /* heartbeats_per_max_samples */ \
    DDS_LENGTH_UNLIMITED, /* max_send_window */ \
    DDS_LENGTH_UNLIMITED, /* max_heartbeat_retries */ \
    { /* first_write_sequence_number */ \
      0, /* high */ \
      1 /* low */ \
    } \
  }, \
  DDS_BOOLEAN_TRUE /* serialize_on_write */ \
}, \
DDS_TYPESUPPORT_QOS_POLICY_DEFAULT, \
{ /* transports */ \
  REDA_StringSeq_INITIALIZER_W_LOAN>HelloWorldAppLibrary>HelloWorldDPDEPubDP_
↳HelloWorldDPDEPub>HelloWorldDPDEDW_transport_enabled_transports, 1, 1) /* enabled_
↳transports */ \
}, \
RTI_MANAGEMENT_QOS_POLICY_DEFAULT, \
DDS_DATAWRITERRESOURCE_LIMITS_QOS_POLICY_DEFAULT, \
DDS_PUBLISH_MODE_QOS_POLICY_DEFAULT, \
DDS_USER_DATA_QOS_POLICY_DEFAULT, \
DDS_DATAWRITERQOS_TRUST_INITIALIZER \
DDS_DATAWRITERQOS_APPGEN_INITIALIZER \
NULL, \
DDS_DataWriterTransferModeQosPolicy_INITIALIZER \
}

```

## Subscriber QoS

Our example doesn't specify any value for Subscriber QoS, however *rtiddsmag* would generate code if it was specified.

## DataReader QoS

The example defines a base profile named *DefaultProfile*, which contains the base QoSs used by each *DomainParticipant*. You can see the content of the *DataReader* QoS below.

```

<!-- DataReader QoS -->
<datareader_qos>
  <history>
    <depth>32</depth>

```

(continues on next page)

(continued from previous page)

```

</history>
<resource_limits>
  <max_instances>2</max_instances>
  <max_samples>64</max_samples>
  <max_samples_per_instance>32</max_samples_per_instance>
</resource_limits>
<reliability>
  <kind>RELIABLE_RELIABILITY_QOS</kind>
</reliability>
<reader_resource_limits>
  <max_remote_writers>10</max_remote_writers>
  <max_remote_writers_per_instance>10</max_remote_writers_per_instance>
</reader_resource_limits>
<!-- transports -->
<unicast>
  <value>
    <element>
      <transports>
        <element>udp4</element>
      </transports>
    </element>
  </value>
</unicast>
<multicast>
  <value>
    <element>
      <receive_address>127.0.0.1</receive_address>
      <transports>
        <element>udp4</element>
      </transports>
    </element>
  </value>
</multicast>
</datareader_qos>

```

*rtiddsmag* generates the following code:

### HelloWorldAppgen.c

```

const char *const
HelloWorldAppLibrary_HelloWorldDPDESubDP_HelloWorldDPDESub_HelloWorldDPDEDR_transport_
↪enabled_transports[2] =
{
  "udp1://",
  "udp1://127.0.0.1"
};

```

### HelloWorldAppgen.h

```

extern const char *const
HelloWorldAppLibrary_HelloWorldDPDESubDP_HelloWorldDPDESub_HelloWorldDPDEDR_transport_
↪enabled_transports[2];

```

(continues on next page)

(continued from previous page)

```

#define RTI_APP_GEN___DR_QOS>HelloWorldAppLibrary>HelloWorldDPDESubDP>HelloWorldDPDESub_
↳HelloWorldDPDEDR \
{ \
    DDS_DEADLINE_QOS_POLICY_DEFAULT, \
    DDS_LIVELINESS_QOS_POLICY_DEFAULT, \
    { /* history */ \
        DDS_KEEP_LAST_HISTORY_QOS, /* kind */ \
        32L /* depth */ \
    }, \
    { /* resource_limits */ \
        64L, /* max_samples */ \
        2L, /* max_instances */ \
        32L /* max_samples_per_instance */ \
    }, \
    DDS_OWNERSHIP_QOS_POLICY_DEFAULT, \
    DDS_LATENCY_BUDGET_QOS_POLICY_DEFAULT, \
    { /* reliability */ \
        DDS_RELIABLE_RELIABILITY_QOS, /* kind */ \
        { /* max_blocking_time */ \
            0L, /* sec */ \
            0L /* nanosec */ \
        } \
    }, \
    DDS_DURABILITY_QOS_POLICY_DEFAULT, \
    DDS_DESTINATION_ORDER_QOS_POLICY_DEFAULT, \
    DDS_TRANSPORT_ENCAPSULATION_QOS_POLICY_DEFAULT, \
    DDS_DATA_REPRESENTATION_QOS_POLICY_DEFAULT, \
    DDS_TYPESUPPORT_QOS_POLICY_DEFAULT, \
    DDS_DATA_READER_PROTOCOL_QOS_POLICY_DEFAULT, \
    { /* transports */ \
        REDA_StringSeq_INITIALIZER_W_LOAN(HelloWorldAppLibrary>HelloWorldDPDESubDP_
↳HelloWorldDPDESub>HelloWorldDPDEDR_transport_enabled_transports, 2, 2) /* enabled_
↳transports */ \
    }, \
    { /* reader_resource_limits */ \
        10L, /* max_remote_writers */ \
        10L, /* max_remote_writers_per_instance */ \
        1L, /* max_samples_per_remote_writer */ \
        1L, /* max_outstanding_reads */ \
        DDS_NO_INSTANCE_REPLACEMENT_QOS, /* instance_replacement */ \
        4L, /* max_routes_per_writer */ \
        DDS_MAX_AUTO, /* max_fragmented_samples */ \
        DDS_MAX_AUTO, /* max_fragmented_samples_per_remote_writer */ \
        DDS_SIZE_AUTO /* shmем_ref_transfer_mode_attached_segment_allocation */ \
    }, \
    RTI_MANAGEMENT_QOS_POLICY_DEFAULT, \
    DDS_USER_DATA_QOS_POLICY_DEFAULT, \
    DDS_DATAREADERQOS_TRUST_INITIALIZER \
    DDS_DATAREADERQOS_APPGEN_INITIALIZER \
    NULL \
}

```

## Topic QoS

Our example doesn't specify any value for Topic QoS; however, *rtiddsmag* would generate code if it were specified.

## Transport and Discovery Configuration

*rtiddsmag* creates the code necessary to configure each one of the available transports used by *Connext Micro* (UDP, SHMEM, and Zero Copy v2) and the discovery mechanism (Dynamic and Static discovery). It also generates the name automatically for each component regardless of if it is a transport or discovery; for this *rtiddsmag* will add a *DomainParticipant* number at the end of its name, only if that configuration is not used by any other *DomainParticipant*:

- UDP Transport: `udp + participant_number`.
- SHMEM Transport: `shmem + participant_number`.
- Zero Copy v2 Transport: `zcopy + participant_number`.
- DPDE: `dpde + participant_number`.
- DPSE: `dpse + participant_number`.

These names can be changed by using the `...Name` options described in *MAG Command-Line Options*.

---

### Note:

- *rtiddsmag* will only create the transport configuration based on the strongly typed XML elements in the schema. *rtiddsmag* **will not** use the values in the property tag to configure the transport.
  - If the length of one of these names exceeds the maximum length, *rtiddsmag* will throw an error.
- 

The following configuration specifies dynamic discovery:

```
<domain_participant_qos>
  <discovery_config>
    <builtin_discovery_plugins>SDP</builtin_discovery_plugins>
  </discovery_config>
</domain_participant_qos>
```

## HelloWorldAppgen.h

```
#define RTI_APP_GEN___dpde__HelloWorldAppLibrary_HelloWorldDPDEPubDP_dpde1 \
{ \
  RT_ComponentFactoryProperty_INITIALIZER, /* _parent */ \
  { /*participant_liveliness_assert_period */ \
    30L, /* sec */ \
```

(continues on next page)

(continued from previous page)

```

        OL /* nanosec */ \
    }, \
    { /*participant_liveliness_lease_duration */ \
        100L, /* sec */ \
        OL /* nanosec */ \
    }, \
    5, /* initial_participant_announcements */ \
    { /*initial_participant_announcement_period */ \
        1L, /* sec */ \
        OL /* nanosec */ \
    }, \
    DDS_BOOLEAN_FALSE, /* cache_serialized_samples */ \
    DDS_LENGTH_AUTO, /* max_participant_locators */ \
    4, /* max_locators_per_discovered_participant */ \
    8, /* max_samples_per_builtin_endpoint_reader */ \
    DDS_LENGTH_UNLIMITED, /* builtin_writer_max_heartbeat_retries */ \
    { /*builtin_writer_heartbeat_period */ \
        OL, /* sec */ \
        100000000L /* nanosec */ \
    }, \
    1L /* builtin_writer_heartbeats_per_max_samples */ \
    DDS_PARTICIPANT_MESSAGE_READER_RELIABILITY_KIND_INITIALIZER \
}

#define RTI_APP_GEN___DP_QOS>HelloWorldAppLibrary>HelloWorldDPDEPubDP \
{ \
    ...
    { /* discovery */ \
        REDA_StringSeq_INITIALIZER_W_LOAN(HelloWorldAppLibrary>HelloWorldDPDEPubDP_
↪initial_peers, 2, 2), /* initial_peers */ \
        REDA_StringSeq_INITIALIZER_W_LOAN(HelloWorldAppLibrary>HelloWorldDPDEPubDP_
↪discovery_enabled_transports, 3, 3), /* enabled_transports */ \
        { \
            { { "dpde1" } }, /* RT_ComponentFactoryId_INITIALIZER */ \
            NDDS_Discovery_Property_INITIALIZER \
        }, /* discovery_component */ \
        DDS_BOOLEAN_FALSE /* accept_unknown_peers */ \
    }, \
    ...
}

```

**Note:**

- *rtiddsmag* will generate an error if the list of available transports for the *DomainParticipant*, *DataWriter*, and *DataReader* contains a transport alias that is not part of the `transport_builtin` mask.
- *rtiddsmag* will not generate code for the SHMEM or UDPv4 transport if it is not specified in the `transport_builtin` mask.
- UDP transformation is not supported in XML.

When using the transport alias to specify the `enabled_transports` for the discovery *DomainParticipant*, *DataWriter* or *DataReader*, you could use the transport names for the built-in transport plugins: `shmem`, `udp4`, and `zcopy`. *rtiddsmag* will automatically modify this alias to match the new one with the *DomainParticipant* number at the end of the name.

The Zero Copy v2 transport is configured differently than the other transports. It cannot be configured through the `transport_builtin` element in XML, and it cannot be enabled with the `transport_builtin` mask. Instead, it can be configured through properties in the XML file. The following properties are required by *rtiddsmag* to configure the Zero Copy v2 transport:

- `dds.transport.micro.zero_copy.max_samples_per_notif`
- `dds.transport.micro.zero_copy.user_intf`<sup>1</sup>
- `dds.transport.micro.zero_copy.user_property`<sup>2</sup>

The following additional properties are only required if you are using the default implementation of the notification mechanism for Zero Copy v2; see *Register the Zero Copy v2 transport* for more information. When configuring the `user_intf` in Zero Copy v2, you must define all or none of the values. If only some of them are defined, MAG will report an error and the generated code will not work with *Connext Micro*.

- `dds.transport.micro.zero_copy.user_property_intf_addr`
- `dds.transport.micro.zero_copy.user_property.thread_prop`
  - `dds.transport.micro.zero_copy.user_property.thread_prop.stack_size`
  - `dds.transport.micro.zero_copy.user_property.thread_prop.priority`
  - `dds.transport.micro.zero_copy.user_property.thread_prop.options`
- `dds.transport.micro.zero_copy.user_property.max_receive_ports`
- `dds.transport.micro.zero_copy.user_property.max_routes`

The following additional properties are only required if you are using your own notification mechanism for Zero Copy v2, NOT the default implementation.

- `dds.transport.micro.zero_copy.user_intf.create_instance`
- `dds.transport.micro.zero_copy.user_intf.delete_instance`
- `dds.transport.micro.zero_copy.user_intf.get_route_table`
- `dds.transport.micro.zero_copy.user_intf.reserve_address`
- `dds.transport.micro.zero_copy.user_intf.release_address`
- `dds.transport.micro.zero_copy.user_intf.resolve_address`
- `dds.transport.micro.zero_copy.user_intf.add_route`

---

<sup>1</sup> This property is only required if you choose to implement your own notification mechanism and not use the default implementation provided by RTI.

<sup>2</sup> Resolves to `ZCOPY_NotifMechanismProperty` when using the default notification mechanism.

- `dds.transport.micro.zero_copy.user_intf.delete_route`
- `dds.transport.micro.zero_copy.user_intf.bind`
- `dds.transport.micro.zero_copy.user_intf.unbind`
- `dds.transport.micro.zero_copy.user_intf.send`
- `dds.transport.micro.zero_copy.user_intf.notify_recv_port`
- `dds.transport.micro.zero_copy.user_intf.create_instance`

The following code is an example of how to configure the Zero Copy v2 transport in XML and the resulting code that *rtiddsmag* generates:

```
<property>
  <value>
    <element>
      <name>dds.micro.zero_copy.enable</name>
      <value>true</value>
    </element>
    <element>
      <name>dds.micro.zero_copy.max_samples_per_notif</name>
      <value>256</value>
    </element>
    <element>
      <name>dds.micro.zero_copy.user_property_intf_addr</name>
      <value>125</value>
    </element>
  </value>
</property>
```

### HelloWorldAppgen.c

```
struct ZCOPY_NotifMechanismProperty zcopy1_user_property = RTI_APP_GEN___zcopy__
↳ HelloWorldAppLibrary_HelloWorldDPDEPubDP_zcopy1_NOTIF_USER_PROPERTY;
struct ZCOPY_NotifLoaderFactoryProperty HelloWorldAppLibrary_HelloWorldDPDEPubDP_
↳ zcopy[1] =
{
  RTI_APP_GEN___zcopy__HelloWorldAppLibrary_HelloWorldDPDEPubDP_zcopy1
};
const struct ComponentFactoryRegisterModel HelloWorldAppLibrary_HelloWorldDPDEPubDP_
↳ register_components[3] =
{
  /* ... */
  {
    "zcopy1_", /* register_name */
    ZCOPY_Loader_get_interface, /* register_intf */
    &HelloWorldAppLibrary_HelloWorldDPDEPubDP_zcopy[0]._parent._parent._parent, /*
↳ register_property */
    NULL /* register_listener */
  }
};
```

### HelloWorldAppgen.h

```

#define RTI_APP_GEN__zcopy__HelloWorldAppLibrary_HelloWorldDPDEPubDP_zcopy1_NOTIF_USER_
↳PROPERTY \
{ \
    125U, /* intf_addr */ \
    OSAPI_ThreadProperty_INITIALIZER, \
    2U, /* max_receive_ports */ \
    32U /* max_routes */ \
}

extern struct ZCOPY_NotifMechanismProperty zcopy1_user_property;

#define RTI_APP_GEN__zcopy__HelloWorldAppLibrary_HelloWorldDPDEPubDP_zcopy1_PROPERTY \
{ \
    NETIO_InterfaceFactoryProperty_INITIALIZER, \
    256L, /* max_samples_per_notif */ \
    NULL, /* user_intf */ \
    &zcopy1_user_property /* user_property */ \
}

#define RTI_APP_GEN__zcopy__HelloWorldAppLibrary_HelloWorldDPDEPubDP_zcopy1 \
{ \
    RTI_APP_GEN__zcopy__HelloWorldAppLibrary_HelloWorldDPDEPubDP_zcopy1_PROPERTY, /* _
↳parent */ \
    "zcopy1" /* notif_transport_name */ \
}

```

## UDP Transport Configuration

*rtiddsmag* supports configuring the following properties via the [PROPERTY](#) QoS policy for the *DomainParticipant*:

- [disable\\_multicast\\_bind](#)
- [multicast\\_loopback\\_disable](#)
- [disable\\_multicast\\_interface\\_select](#)

Refer to *UDP Configuration* for more information on these properties.



## Shared Memory Transport (SHMEM) Configuration

*rtiddsmag* supports configuring the `dds.transport.minimum_compatibility_version` property, which you can set via the [PROPERTY](#) QoS policy for the *DomainParticipant*. Refer to *SHMEM Configuration* for more information on `dds.transport.minimum_compatibility_version`.

## Discovery Configuration

*rtiddsmag* supports configuring the following [DISCOVERY](#) QoS policy fields:

- `dds.micro.discovery.enable_participant_discovery_by_name`
- `dds.micro.discovery.enable_endpoint_discovery_queue`

However, you must configure these fields via the [PROPERTY](#) QoS policy instead of [DISCOVERY](#), as shown in the example code below:

```
<property>
  <value>
    <element>
      <name>dds.micro.discovery.enable_participant_discovery_by_name</name>
      <value>true</value>
    </element>
    <element>
      <name>dds.micro.discovery.enable_endpoint_discovery_queue</name>
      <value>true</value>
    </element>
  </value>
</property>
```

For more information on these fields, refer to *DomainParticipant Discovery by Name* and *[Major] Possible mismatched endpoint messages during discovery*.

## Flow Controllers

*rtiddsmag* creates code which will be used by *Connext Micro* to create a custom flow controller. The custom flow controller is configured through properties in the XML file. Let's see an example of how to configure a custom flow controller named `custom_flowcontroller` and the code that *rtiddsmag* generates:

```
<domain_participant_qos>
  ...
  <property>
    <value>
      <element>
        <name>
dds.flow_controller.token_bucket.custom_flowcontroller.token_bucket.max_tokens
        </name>
        <value>2</value>
      </element>
    </value>
  </property>
</domain_participant_qos>
```

(continues on next page)

(continued from previous page)

```

        </element>
        <element>
            <name>
dds.flow_controller.token_bucket.custom_flowcontroller.token_bucket.tokens_added_per_
↪period
            </name>
            <value>2</value>
        </element>
        <element>
            <name>
dds.flow_controller.token_bucket.custom_flowcontroller.token_bucket.tokens_leaked_per_
↪period
            </name>
            <!-- The value -1 means LENGTH_UNLIMITED -->
            <value>-1</value>
        </element>
        <element>
            <name>
dds.flow_controller.token_bucket.custom_flowcontroller.token_bucket.period.sec
            </name>
            <value>0</value>
        </element>
        <element>
            <name>
dds.flow_controller.token_bucket.custom_flowcontroller.token_bucket.period.nanosec
            </name>
            <value>100000000</value>
        </element>
        <element>
            <name>
dds.flow_controller.token_bucket.custom_flowcontroller.token_bucket.bytes_per_token
            </name>
            <value>1024</value>
        </element>
    </value>
</property>
</domain_participant_qos>

<datawriter_qos>
    <publish_mode>
        <flow_controller_name>
            dds.flow_controller.token_bucket.custom_flowcontroller
        </flow_controller_name>
        <kind>ASYNCHRONOUS_PUBLISH_MODE_QOS</kind>
        <priority>12</priority>
    </publish_mode>
</datawriter_qos>

```

### HelloWorldAppgen.c

```
const struct APPGEN_CustomFlowControllerModel
```

(continues on next page)

(continued from previous page)

```

HelloWorldAppLibrary_HelloWorldDPDEPubDP_flow_controllers[1] =
{
    {
        "custom_flowcontroller", /* name */
        RTI_APP_GEN___FC_P_QOS_HelloWorldAppLibrary_HelloWorldDPDEPubDP_custom_
↪flowcontroller /* flow_controller_property */
    }
};

```

## HelloWorldAppgen.h

```

#define
RTI_APP_GEN___FC_P_QOS_HelloWorldAppLibrary_HelloWorldDPDEPubDP_custom_flowcontroller \
{ \
    NETIO_FlowControllerProperty_INITIALIZER, \
    DDS_EDF_FLOW_CONTROLLER_SCHED_POLICY, /* scheduling_policy */ \
    { /* token_bucket */ \
        2L, /* max_tokens */ \
        2L, /* tokens_added_per_period */ \
        -1L, /* tokens_leaked_per_period */ \
        { /* period */ \
            0L, /* sec */ \
            1000000000L /* nanosec */ \
        }, \
        1024L /* bytes_per_token */ \
    }, \
    DDS_BOOLEAN_FALSE /* is_vendor_specific */ \
}

#define
RTI_APP_GEN___DW_QOS_HelloWorldAppLibrary_HelloWorldDPDEPubDP_HelloWorldDPDEPub_
↪HelloWorldDPDEDW \
{ \
    ...
    { /* publish_mode */ \
        DDS_ASYNCHRONOUS_PUBLISH_MODE_QOS, /* max_remote_readers */ \
        "custom_flowcontroller", /* flow_controller_name */ \
        12L /* priority */ \
    }, \
    ...
}

extern const struct APPGEN_CustomFlowControllerModel
HelloWorldAppLibrary_HelloWorldDPDEPubDP_flow_controllers[1];

#define RTI_APP_GEN___DP_HelloWorldAppLibrary_HelloWorldDPDEPubDP \
{ \
    ...
    1UL, /* custom_flow_controller_count */ \
    HelloWorldAppLibrary_HelloWorldDPDEPubDP_flow_controllers /* custom_flow_controllers_
↪*/ \
}

```

The three built-in Flow Controllers are also supported by *rtiddsmag*:

- `DEFAULT_FLOW_CONTROLLER_NAME`
- `FIXED_RATE_FLOW_CONTROLLER_NAME`
- `ON_DEMAND_FLOW_CONTROLLER_NAME`

The generated code is slightly different when any of these three built-in Flow Controllers are configured, as there is no need to generate code to register the Flow Controller.

```
<datawriter_qos>
  <publish_mode>
    <flow_controller_name>DEFAULT_FLOW_CONTROLLER_NAME</flow_controller_name>
    <kind>ASYNCHRONOUS_PUBLISH_MODE_QOS</kind>
    <priority>12</priority>
  </publish_mode>
</datawriter_qos>
```

## HelloWorldAppgen.h

```
#define
RTI_APP_GEN__DW_QOS>HelloWorldAppLibrary>HelloWorldDPDEPubDP>HelloWorldDPDEPub_
↪HelloWorldDPDEDW \
{ \
  ...
  { /* publish_mode */ \
    DDS_ASYNCHRONOUS_PUBLISH_MODE_QOS, /* max_remote_readers */ \
    "DDS_DEFAULT_FLOW_CONTROLLER_NAME", /* flow_controller_name */ \
    12L /* priority */ \
  }, \
  ...
}

#define RTI_APP_GEN__DP>HelloWorldAppLibrary>HelloWorldDPDEPubDP \
{ \
  ...
  OUL, /* custom_flow_controller_count */ \
  NULL /* custom_flow_controllers */ \
}
```

**Note:** A flow controller is only used by Micro when the `publish_mode` kind is set to either `ASYNCHRONOUS_PUBLISH_MODE_QOS` or `AUTOMATIC_PUBLISH_MODE_QOS`.

## Static Discovery

*rtiddsmag* iterates through each *DomainParticipant* definition in the XML configuration file, creating the remote entities that are needed to communicate with applications that use static discovery, and updating the *object\_id* of each *DataWriter* or *DataReader* involved if they don't have a valid value or they are using the default value.

Let's see an example of two applications that use static discovery and how *rtiddsmag* generates the necessary code that will be asserted by *Connext Micro* to communicate with both applications:

```
<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>
  <domain_participant_qos base_name="QosLibrary::DPSEProfile"/>
</domain_participant>

<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>
  <domain_participant_qos base_name="QosLibrary::DPSEProfile"/>
</domain_participant>
```

For these two *DomainParticipants*, *rtiddsmag* will update the *rtps\_object\_id* for the *DataWriter* and *DataReader*, since they didn't have any values set in the XML file. You can see this in the following snippet from **HelloWorldAppgen.h**:

```
#define
RTI_APP_GEN__DW_QOS_HelloWorldAppLibrary_HelloWorldDPSEPubDP_HelloWorldDPSEPub_
↪HelloWorldDPSEDW \
{ \
  ...
  { /* protocol */ \
    1UL, /* rtps_object_id */ \
    { /* rtps_reliable_writer */ \
      { /* heartbeat_period */ \
        0L, /* sec */ \
        250000000UL /* nanosec */ \
      }, \
      1L, /* heartbeats_per_max_samples */ \
      DDS_LENGTH_UNLIMITED, /* max_send_window */ \
      DDS_LENGTH_UNLIMITED, /* max_heartbeat_retries */ \
      { /* first_write_sequence_number */ \
        0, /* high */ \
        1 /* low */ \
      } \
    } \
  } \
}
```

(continues on next page)

(continued from previous page)

```

        } \
    }, \
    DDS_BOOLEAN_TRUE /* serialize_on_write */ \
}, \
...
}

#define
RTI_APP_GEN___DR_QOS>HelloWorldAppLibrary>HelloWorldDPSESubDP>HelloWorldDPSESub_
↪HelloWorldDPSEDR \
{ \
    ...
    { /* protocol */ \
        2UL /* rtps_object_id */ \
    }, \
    ...
}

```

*rtiddsmag* will also generate the remote *DomainParticipants*, *DataWriters*, and *DataReaders* that need to be asserted in order for endpoints to match:

### HelloWorldAppgen.c

```

const struct APPGEN_RemoteSubscriptionModel
HelloWorldAppLibrary>HelloWorldDPSEPubDP_remote_subscribers[1] =
{
    RTI_APP_GEN__RSD>HelloWorldAppLibrary>HelloWorldDPSEPubDP>HelloWorldAppLibrary_
↪HelloWorldDPSESubDP>HelloWorldDPSESub>HelloWorldDPSEDR
};

const struct APPGEN_RemoteParticipantModel
HelloWorldAppLibrary>HelloWorldDPSEPubDP_remote_participants[1] =
{
    {
        "HelloWorldDPSESubDP", /* name */
        0UL, /* remote_publisher_count */
        NULL, /* remote_publishers */
        1UL, /* remote_subscriber_count */
        HelloWorldAppLibrary>HelloWorldDPSEPubDP_remote_subscribers /* remote_
↪subscribers */
    }
};

const struct APPGEN_RemotePublicationModel
HelloWorldAppLibrary>HelloWorldDPSESubDP_remote_publishers[1] =
{
    RTI_APP_GEN__RPD>HelloWorldAppLibrary>HelloWorldDPSESubDP>HelloWorldAppLibrary_
↪HelloWorldDPSEPubDP>HelloWorldDPSEPub>HelloWorldDPSEDW
};

const struct APPGEN_RemoteParticipantModel
HelloWorldAppLibrary>HelloWorldDPSESubDP_remote_participants[1] =

```

(continues on next page)

(continued from previous page)

```

{
    {
        "HelloWorldDPSEPubDP", /* name */
        1UL, /* remote_publisher_count */
        HelloWorldAppLibrary_HelloWorldDPSESubDP_remote_publishers, /* remote_publishers
↪ */
        0UL, /* remote_subscriber_count */
        NULL /* remote_subscribers */
    }
};

```

## HelloWorldAppgen.h

```

#define RTI_APP_GEN__RSD_HelloWorldAppLibrary_HelloWorldDPSEPubDP_HelloWorldAppLibrary_
↪HelloWorldDPSESubDP_HelloWorldDPSESub_HelloWorldDPSEDR \
{ \
    { /* subscription_data */ \
        { \
            { 0, 0, 0, 2 } /* key */ \
        }, \
        { \
            { 0, 0, 0, 0 } /* participant_key */ \
        }, \
        "HelloWorldTopic", /* topic_name */ \
        "HelloWorldType", /* type_name */ \
        DDS_DEADLINE_QOS_POLICY_DEFAULT, \
        DDS_OWNERSHIP_QOS_POLICY_DEFAULT, \
        DDS_LATENCY_BUDGET_QOS_POLICY_DEFAULT, \
        { /* reliability */ \
            DDS_RELIABLE_RELIABILITY_QOS, /* kind */ \
            { /* max_blocking_time */ \
                0L, /* sec */ \
                0L /* nanosec */ \
            } \
        }, \
        DDS_LIVELINESS_QOS_POLICY_DEFAULT, \
        DDS_DURABILITY_QOS_POLICY_DEFAULT, \
        DDS_DESTINATION_ORDER_QOS_POLICY_DEFAULT, \
        DDS_SEQUENCE_INITIALIZER, \
        DDS_SEQUENCE_INITIALIZER, \
        DDS_DATA_REPRESENTATION_QOS_POLICY_DEFAULT \
        DDS_TRUST_SUBSCRIPTION_DATA_INITIALIZER \
    }, \
    HelloWorldTypePlugin_get /* get_type_plugin */ \
}
extern const struct APPGEN_RemoteSubscriptionModel HelloWorldAppLibrary_
↪HelloWorldDPSEPubDP_remote_subscribers[1];
extern const struct APPGEN_RemoteParticipantModel HelloWorldAppLibrary_
↪HelloWorldDPSEPubDP_remote_participants[1];

#define RTI_APP_GEN__DP_HelloWorldAppLibrary_HelloWorldDPSEPubDP \

```

(continues on next page)

(continued from previous page)

```

{ \
    "HelloWorldDPSEPubDP", /* name */ \
    RTI_APP_GEN__DPF_HelloWorldAppLibrary_HelloWorldDPSEPubDP, /* domain_participant_
↪factory */ \
    RTI_APP_GEN__DP_QOS_HelloWorldAppLibrary_HelloWorldDPSEPubDP, /* domain_participant_
↪qos */ \
    OL, /* domain_id */ \
    1UL, /* type_registration_count */ \
    HelloWorldAppLibrary_HelloWorldDPSEPubDP_type_registrations, /* type_registrations */
↪ \
    1UL, /* topic_count */ \
    HelloWorldAppLibrary_HelloWorldDPSEPubDP_topics, /* topics */ \
    1UL, /* publisher_count */ \
    HelloWorldAppLibrary_HelloWorldDPSEPubDP_publishers, /* publishers */ \
    OUL, /* subscriber_count */ \
    NULL, /* subscribers */ \
    1UL, /* remote_participant_count */ \
    HelloWorldAppLibrary_HelloWorldDPSEPubDP_remote_participants /* remote_participants_
↪ */ \
    OUL, /* custom_flow_controller_count */ \
    NULL, /* custom_flow_controllers */ \
}

#define RTI_APP_GEN__RPD_HelloWorldAppLibrary_HelloWorldDPSESubDP_HelloWorldAppLibrary_
↪HelloWorldDPSEPubDP_HelloWorldDPSEPub_HelloWorldDPSEDW \
{ \
    { /* publication_data */ \
        { \
            { 0, 0, 0, 1 } /* key */ \
        }, \
        { \
            { 0, 0, 0, 0 } /* participant_key */ \
        }, \
        "HelloWorldTopic", /* topic_name */ \
        "HelloWorldType", /* type_name */ \
        DDS_DEADLINE_QOS_POLICY_DEFAULT, \
        DDS_OWNERSHIP_QOS_POLICY_DEFAULT, \
        DDS_OWNERSHIP_STRENGTH_QOS_POLICY_DEFAULT, \
        DDS_LATENCY_BUDGET_QOS_POLICY_DEFAULT, \
        { /* reliability */ \
            DDS_RELIABLE_RELIABILITY_QOS, /* kind */ \
            { /* max_blocking_time */ \
                OL, /* sec */ \
                100000000L /* nanosec */ \
            } \
        }, \
        DDS_LIVELINESS_QOS_POLICY_DEFAULT, \
        DDS_DURABILITY_QOS_POLICY_DEFAULT, \
        DDS_DESTINATION_ORDER_QOS_POLICY_DEFAULT, \
        DDS_SEQUENCE_INITIALIZER, \
        DDS_DATA_REPRESENTATION_QOS_POLICY_DEFAULT \
    } \
}

```

(continues on next page)



(continued from previous page)

```

        DDS_TRUST_PUBLICATION_DATA_INITIALIZER \
    }, \
    HelloWorldTypePlugin_get /* get_type_plugin */ \
}

extern const struct APPGEN_RemotePublicationModel HelloWorldAppLibrary_
↳HelloWorldDPSESubDP_remote_publishers[1];
extern const struct APPGEN_RemoteParticipantModel HelloWorldAppLibrary_
↳HelloWorldDPSESubDP_remote_participants[1];

#define RTI_APP_GEN__DP_HelloWorldAppLibrary_HelloWorldDPSESubDP \
{ \
    "HelloWorldDPSESubDP", /* name */ \
    RTI_APP_GEN__DPF_HelloWorldAppLibrary_HelloWorldDPSESubDP, /* domain_participant_
↳factory */ \
    RTI_APP_GEN__DP_QOS_HelloWorldAppLibrary_HelloWorldDPSESubDP, /* domain_participant_
↳qos */ \
    0L, /* domain_id */ \
    1UL, /* type_registration_count */ \
    HelloWorldAppLibrary_HelloWorldDPSESubDP_type_registrations, /* type_registrations */
↳ \
    1UL, /* topic_count */ \
    HelloWorldAppLibrary_HelloWorldDPSESubDP_topics, /* topics */ \
    0UL, /* publisher_count */ \
    NULL, /* publishers */ \
    1UL, /* subscriber_count */ \
    HelloWorldAppLibrary_HelloWorldDPSESubDP_subscribers, /* subscribers */ \
    1UL, /* remote_participant_count */ \
    HelloWorldAppLibrary_HelloWorldDPSESubDP_remote_participants /* remote_participants_
↳*/ \
    0UL, /* custom_flow_controller_count */ \
    NULL /* custom_flow_controllers */ \
}

#define RTI_APP_GEN__RSD_HelloWorldAppLibrary_HelloWorldDPSEPubDP_HelloWorldAppLibrary_
↳HelloWorldDPSESubDP_HelloWorldDPSESub_HelloWorldDPSEDR \
{ \
    { /* subscription_data */ \
        { \
            { 0, 0, 0, 2 } /* key */ \
        }, \
        { \
            { 0, 0, 0, 0 } /* participant_key */ \
        }, \
        "HelloWorldTopic", /* topic_name */ \
        "HelloWorldType", /* type_name */ \
        DDS_DEADLINE_QOS_POLICY_DEFAULT, \
        DDS_OWNERSHIP_QOS_POLICY_DEFAULT, \
        DDS_LATENCY_BUDGET_QOS_POLICY_DEFAULT, \
        { /* reliability */ \
            DDS_RELIABLE_RELIABILITY_QOS, /* kind */ \
            { /* max_blocking_time */ \

```

(continues on next page)

(continued from previous page)

```

        OL, /* sec */ \
        OL /* nanosec */ \
    } \
}, \
DDS_LIVELINESS_QOS_POLICY_DEFAULT, \
DDS_DURABILITY_QOS_POLICY_DEFAULT, \
DDS_DESTINATION_ORDER_QOS_POLICY_DEFAULT, \
DDS_SEQUENCE_INITIALIZER, \
DDS_SEQUENCE_INITIALIZER, \
DDS_DATA_REPRESENTATION_QOS_POLICY_DEFAULT \
DDS_TRUST_SUBSCRIPTION_DATA_INITIALIZER \
}, \
HelloWorldTypePlugin_get /* get_type_plugin */ \
}
extern const struct APPGEN_RemoteSubscriptionModel HelloWorldAppLibrary_
↪HelloWorldDPSEPubDP_remote_subscribers[1];
extern const struct APPGEN_RemoteParticipantModel HelloWorldAppLibrary_
↪HelloWorldDPSEPubDP_remote_participants[1];

#define RTI_APP_GEN__DP_HelloWorldAppLibrary_HelloWorldDPSEPubDP \
{ \
    "HelloWorldDPSEPubDP", /* name */ \
    RTI_APP_GEN__DPF_HelloWorldAppLibrary_HelloWorldDPSEPubDP, /* domain_participant_
↪factory */ \
    RTI_APP_GEN__DP_QOS_HelloWorldAppLibrary_HelloWorldDPSEPubDP, /* domain_participant_
↪qos */ \
    OL, /* domain_id */ \
    1UL, /* type_registration_count */ \
    HelloWorldAppLibrary_HelloWorldDPSEPubDP_type_registrations, /* type_registrations */
↪ \
    1UL, /* topic_count */ \
    HelloWorldAppLibrary_HelloWorldDPSEPubDP_topics, /* topics */ \
    1UL, /* publisher_count */ \
    HelloWorldAppLibrary_HelloWorldDPSEPubDP_publishers, /* publishers */ \
    OUL, /* subscriber_count */ \
    NULL, /* subscribers */ \
    1UL, /* remote_participant_count */ \
    HelloWorldAppLibrary_HelloWorldDPSEPubDP_remote_participants /* remote_participants_
↪ */ \
    OUL, /* custom_flow_controller_count */ \
    NULL, /* custom_flow_controllers */ \
}

#define RTI_APP_GEN__RPD_HelloWorldAppLibrary_HelloWorldDPSESubDP_HelloWorldAppLibrary_
↪HelloWorldDPSEPubDP_HelloWorldDPSEPubDP_HelloWorldDPSEDW \
{ \
    { /* publication_data */ \
        { \
            { 0, 0, 0, 1 } /* key */ \
        }, \
        { \

```

(continues on next page)

(continued from previous page)

```

        { 0, 0, 0, 0 } /* participant_key */ \
    }, \
    "HelloWorldTopic", /* topic_name */ \
    "HelloWorldType", /* type_name */ \
    DDS_DEADLINE_QOS_POLICY_DEFAULT, \
    DDS_OWNERSHIP_QOS_POLICY_DEFAULT, \
    DDS_OWNERSHIP_STRENGTH_QOS_POLICY_DEFAULT, \
    DDS_LATENCY_BUDGET_QOS_POLICY_DEFAULT, \
    { /* reliability */ \
        DDS_RELIABLE_RELIABILITY_QOS, /* kind */ \
        { /* max_blocking_time */ \
            0L, /* sec */ \
            100000000L /* nanosec */ \
        } \
    }, \
    DDS_LIVELINESS_QOS_POLICY_DEFAULT, \
    DDS_DURABILITY_QOS_POLICY_DEFAULT, \
    DDS_DESTINATION_ORDER_QOS_POLICY_DEFAULT, \
    DDS_SEQUENCE_INITIALIZER, \
    DDS_DATA_REPRESENTATION_QOS_POLICY_DEFAULT \
    DDS_TRUST_PUBLICATION_DATA_INITIALIZER \
}, \
HelloWorldTypePlugin_get /* get_type_plugin */ \
}

extern const struct APPGEN_RemotePublicationModel HelloWorldAppLibrary_
↳HelloWorldDPSESubDP_remote_publishers[1];
extern const struct APPGEN_RemoteParticipantModel HelloWorldAppLibrary_
↳HelloWorldDPSESubDP_remote_participants[1];

#define RTI_APP_GEN__DP_HelloWorldAppLibrary_HelloWorldDPSESubDP \
{ \
    "HelloWorldDPSESubDP", /* name */ \
    RTI_APP_GEN__DPF_HelloWorldAppLibrary_HelloWorldDPSESubDP, /* domain_participant_
↳factory */ \
    RTI_APP_GEN__DP_QOS_HelloWorldAppLibrary_HelloWorldDPSESubDP, /* domain_participant_
↳qos */ \
    0L, /* domain_id */ \
    1UL, /* type_registration_count */ \
    HelloWorldAppLibrary_HelloWorldDPSESubDP_type_registrations, /* type_registrations */
↳ \
    1UL, /* topic_count */ \
    HelloWorldAppLibrary_HelloWorldDPSESubDP_topics, /* topics */ \
    0UL, /* publisher_count */ \
    NULL, /* publishers */ \
    1UL, /* subscriber_count */ \
    HelloWorldAppLibrary_HelloWorldDPSESubDP_subscribers, /* subscribers */ \
    1UL, /* remote_participant_count */ \
    HelloWorldAppLibrary_HelloWorldDPSESubDP_remote_participants /* remote_participants_
↳*/ \
    0UL, /* custom_flow_controller_count */ \

```

(continues on next page)

(continued from previous page)

```

    NULL /* custom_flow_controllers */ \
}

```

## Lightweight Security Plugin

*rtiddsmag* can generate code to configure the *Lightweight Security Plugin* for *Connext Micro*. You can configure the plugin through properties in the XML file.

*rtiddsmag* requires the following properties to configure the plugin:

- `dds.sec.crypto.enable`
- `dds.sec.crypto.rtps_psk_secret_passphrase`
- `dds.sec.crypto.rtps_psk_symmetric_cipher_algorithm`

The following properties are optional:

- `dds.sec.access.rtps_psk_protection_kind`
- `com.rti.serv.secure.cryptography.max_blocks_per_session`

*rtiddsmag* automatically generates the name of the pre-shared key (PSK) component. *rtiddsmag* also appends a *DomainParticipant* number to the name, as long as the particular configuration is not being used by any other *DomainParticipant*:

```
psk + participant_number
```

For example, the first unique pre-shared key configuration would be named `psk1`; the second, `psk2`; and so on.

This name can be changed with the `-pskName` option described in *MAG Command-Line Options*.

---

**Note:** If the length of the name exceeds the maximum length, *rtiddsmag* will throw an error.

---

Refer to the following code snippet as an example of how to configure the *Lightweight Security Plugin* in XML and the resulting code that *rtiddsmag* generates:

```

<property>
  <value>
    <element>
      <name>dds.micro.crypto.enable</name>
      <value>true</value>
    </element>
    <element>
      <name>dds.sec.crypto.rtps_psk_secret_passphrase</name>
      <value>data:,404166165:castle super radar denial swing lunar kind swarm wet_
↪toilet output harbor basic begin margin huge year visit</value>
    </element>
  </element>

```

(continues on next page)

(continued from previous page)

```

        <name>dds.sec.crypto.rtps_psk_symmetric_cipher_algorithm</name>
        <value>AUTO</value>
    </element>
    <element>
        <name>dds.sec.access.rtps_psk_protection_kind</name>
        <value>ENCRYPT</value>
    </element>
    <element>
        <name>com.rti.serv.secure.cryptography.max_blocks_per_session</name>
        <value>12</value>
    </element>
</value>
</property>

```

### HelloWorldAppgen.c

```

struct DDS_PskServiceFactoryProperty HelloWorldAppLibrary_HelloWorldDPDEPubDP_psk[1] =
{
    RTI_APP_GEN___psk__HelloWorldAppLibrary_HelloWorldDPDEPubDP_psk1
};
const struct ComponentFactoryRegisterModel HelloWorldAppLibrary_HelloWorldDPDEPubDP_
↪register_components[2] =
{
    /* ... */
    {
        "psk1", /* register_name */
        DDS_PskServiceFactory_get_interface, /* register_intf */
        &HelloWorldAppLibrary_HelloWorldDPDEPubDP_psk[0]._parent, /* register_property */
        NULL /* register_listener */
    }
};

```

### HelloWorldAppgen.h

```

#define RTI_APP_GEN___psk__HelloWorldAppLibrary_HelloWorldDPDEPubDP_psk1 \
{ \
    RT_ComponentFactoryProperty_INITIALIZER, \
    "psk1", /* service_name */ \
    PSK_OSSL_get_interface, /* psl_get_interface_func */ \
    NULL /* psl_config */ \
}

extern struct DDS_PskServiceFactoryProperty HelloWorldAppLibrary_HelloWorldDPDEPubDP_
↪psk[1];

```

## Content filtering

**Attention:** Content filtering is an experimental feature in *Connext Micro* <version>. It is not guaranteed to be consistent or supported and should not be used in production.

Refer to *Experimental Features* for more information.

*rtiddsmag* can configure *Content Filtering* with syntax that is compatible with *Connext Professional* application configuration.

## Configuring DataReaders

Unlike in manual configuration, we don't configure content filters in the *DataReader*'s QoS. Instead, we configure it on the *DataReader* itself. Refer to the following code snippet for an example *DataReader* configuration and the code that *rtiddsmag* generates:

```
<subscriber name="HelloWorldDPDESub">
  <data_reader topic_ref="HelloWorldTopic" name="HelloWorldDPDEDR">
    <datareader_qos base_name="QosLibrary::DPDEProfile"/>
    <content_filter name="myFilter1" kind="builtin.sql">
      <expression>id == 1</expression>
    </content_filter>
  </data_reader>
</subscriber>
```

**Note:** When using *rtiddsmag*, the only supported filter kind is `builtin.sql`. The syntax is the same as when configuring a `DDS_ContentFilterQos`.

## HelloWorldAppgen.h

```
#define RTI_APP_GEN__DR_QOS_HelloWorldAppLibrary_HelloWorldDPDESubDP_HelloWorldDPDESub_
↪HelloWorldDPDEDR \
{ \
  DDS_DEADLINE_QOS_POLICY_DEFAULT, \
  DDS_LIVELINESS_QOS_POLICY_DEFAULT, \
  { /* history */ \
    DDS_KEEP_LAST_HISTORY_QOS, /* kind */ \
    32L /* depth */ \
  }, \
  { /* resource_limits */ \
    64L, /* max_samples */ \
    2L, /* max_instances */ \
    32L /* max_samples_per_instance */ \
  }, \
  DDS_OWNERSHIP_QOS_POLICY_DEFAULT, \
  DDS_LATENCY_BUDGET_QOS_POLICY_DEFAULT, \
```

(continues on next page)

(continued from previous page)

```

{ /* reliability */ \
  DDS_RELIABLE_RELIABILITY_QOS, /* kind */ \
  { /* max_blocking_time */ \
    0L, /* sec */ \
    0L /* nanosec */ \
  } \
}, \
DDS_DURABILITY_QOS_POLICY_DEFAULT, \
DDS_DESTINATION_ORDER_QOS_POLICY_DEFAULT, \
DDS_TRANSPORT_ENCAPSULATION_QOS_POLICY_DEFAULT, \
DDS_DATA_REPRESENTATION_QOS_POLICY_DEFAULT, \
DDS_TYPESUPPORT_QOS_POLICY_DEFAULT, \
DDS_DATA_READER_PROTOCOL_QOS_POLICY_DEFAULT, \
{ /* transports */ \
  REDA_StringSeq_INITIALIZER_W_LOAN>HelloWorldAppLibrary>HelloWorldDPDESubDP_
↳HelloWorldDPDESub>HelloWorldDPDEDR_transport_enabled_transports, 2, 2) /* enabled_
↳transports */ \
}, \
{ /* reader_resource_limits */ \
  10L, /* max_remote_writers */ \
  10L, /* max_remote_writers_per_instance */ \
  1L, /* max_samples_per_remote_writer */ \
  1L, /* max_outstanding_reads */ \
  DDS_NO_INSTANCE_REPLACEMENT_QOS, /* instance_replacement */ \
  4L, /* max_routes_per_writer */ \
  DDS_MAX_AUTO, /* max_fragmented_samples */ \
  DDS_MAX_AUTO, /* max_fragmented_samples_per_remote_writer */ \
  DDS_SIZE_AUTO /* shm_ref_transfer_mode_attached_segment_allocation */ \
}, \
RTI_MANAGEMENT_QOS_POLICY_DEFAULT, \
DDS_USER_DATA_QOS_POLICY_DEFAULT, \
DDS_PROPERTY_QOS_POLICY_DEFAULT, \
{ /* content_filter */ \
  "myFilter1", /* filter_name */ \
  "DDSSQL", /* filter_class_name */ \
  "id == 1", /* filter_expression */ \
  DDS_SEQUENCE_INITIALIZER /* expression_parameters */ \
}, \
DDS_DATAREADERQOS_TRUST_INITIALIZER \
DDS_DATAREADERQOS_APPGEN_INITIALIZER \
NULL \
}

```

### HelloWorldAppgen.c

```

const struct ComponentFactoryRegisterModel
HelloWorldAppLibrary>HelloWorldDPDEPubDP_register_components[3] =
{
  {
    "dpde1", /* register_name */
    DPDE_DiscoveryFactory_get_interface, /* register_intf */

```

(continues on next page)

(continued from previous page)

```

    &HelloWorldAppLibrary_HelloWorldDPDEPubDP_dpde[0]._parent, /* register_property */
    ↪ */
    NULL /* register_listener */
},
{
    "udp1", /* register_name */
    UDP_InterfaceFactory_get_interface, /* register_intf */
    &HelloWorldAppLibrary_HelloWorldDPDEPubDP_udp4[0]._parent._parent, /* register_
    ↪ property */
    NULL /* register_listener */
},
{
    DDS_FILTER_PLUGIN_FACTORY_DEFAULT_NAME, /* register_name */
    DDS_FilterPluginFactory_get_interface, /* register_intf */
    NULL, /* register_property */
    NULL /* register_listener */
},
};

```

## Configuring DomainParticipants

To enable filtering support on a specific *DomainParticipant*, set the following property in the *DomainParticipant*'s QoS:

```

<participant_qos>
  <property>
    <value>
      <element>
        <name>dds.filter.micro.enable</name>
        <value>true</value>
      </element>
    </value>
  </property>
</participant_qos>

```

When filtering support is enabled on any *DomainParticipant*, the filter plugin factory will be registered automatically. Unlike manual configuration, enabling filter support for a single *DomainParticipant* does not automatically enable it for all *DomainParticipants*; the property must be set to **true** for each *DomainParticipant* that wishes to use the feature.

The following optional properties are supported on the *DomainParticipant*. These options are all equivalent to the fields of the `DDS_FilterQoS`.



Property	Type
<code>dds.filter.micro.enable<sup>3</sup></code>	Boolean
<code>dds.filter.micro.disable_writer_filtering</code>	Boolean
<code>dds.filter.micro.disable_builtin_sql_filter</code>	Boolean
<code>dds.filter.micro.resource_limits.filter_class_max_length</code>	DDS_Long, (0, 16)
<code>dds.filter.micro.resource_limits.filter_class_max_count</code>	DDS_Long (0, LONG_MAX]
<code>dds.filter.micro.resource_limits.filter_expression_max_length</code>	DDS_Long (0, 256)
<code>dds.filter.micro.resource_limits.filter_expression_max_count</code>	DDS_Long (0, LONG_MAX] or DDS_SIZE_AUTO
<code>dds.filter.micro.resource_limits.filter_parameter_max_count_per_expression</code>	DDS_Long [0, 16]
<code>dds.filter.micro.resource_limits.filter_parameter_max_length</code>	DDS_Long [0, LONG_MAX]
<code>dds.filter.micro.resource_limits.sql_predicate_max_count_per_expression</code>	DDS_Long (0, LONG_MAX]

Refer to the following code snippet for an example *DomainParticipant* configuration and the code that *rtiddsmag* generates:

```

<!-- CFT properties -->
<element>
  <name>dds.filter.micro.disable_writer_filtering</name>
  <value>true</value>
  <propagate>>false</propagate>
</element>
<element>
  <name>dds.filter.micro.disable_builtin_sql_filter</name>
  <value>true</value>
  <propagate>>false</propagate>
</element>
<element>
  <name>dds.filter.micro.resource_limits.filter_class_max_length</name>
  <value>7</value>
  <propagate>>false</propagate>
</element>
<element>
  <name>dds.filter.micro.resource_limits.filter_class_max_count</name>
  <value>1</value>
  <propagate>>false</propagate>
</element>
<element>
  <name>dds.filter.micro.resource_limits.filter_expression_max_length</name>
  <value>63</value>

```

(continues on next page)

<sup>3</sup> This is an optional property, because by default *rtiddsmag* will automatically generate code to register the component in the *DomainParticipantFactory* if there is at least one content filter used by one *DataReader*.

(continued from previous page)

```

        <propagate>false</propagate>
    </element>
    <element>
        <name>dds.filter.micro.resource_limits.filter_expression_max_count</name>
        <value>LENGTH_AUTO</value>
        <propagate>false</propagate>
    </element>
    <element>
        <name>dds.filter.micro.resource_limits.filter_parameter_max_count_per_expression
↪</name>
        <value>4</value>
        <propagate>false</propagate>
    </element>
    <element>
        <name>dds.filter.micro.resource_limits.filter_parameter_max_length</name>
        <value>15</value>
        <propagate>false</propagate>
    </element>
    <element>
        <name>dds.filter.micro.resource_limits.sql_predicate_max_count_per_expression</
↪name>
        <value>4</value>
        <propagate>false</propagate>
    </element>

```

## HelloWorldAppgen.c

```

#define RTI_APP_GEN__DP_QOS_HelloWorldAppLibrary_HelloWorldDPDEPubDP \
{ \
    { /* entity_factory */ \
        DDS_BOOLEAN_TRUE /* autoenable_created_entities */ \
    }, \
    { /* discovery */ \
        REDA_StringSeq_INITIALIZER_W_LOAN(HelloWorldAppLibrary_HelloWorldDPDEPubDP_
↪initial_peers, 2, 2), /* initial_peers */ \
        REDA_StringSeq_INITIALIZER_W_LOAN(HelloWorldAppLibrary_HelloWorldDPDEPubDP_
↪discovery_enabled_transports, 3, 3), /* enabled_transports */ \
        { \
            { { "dpde1" } }, /* RT_ComponentFactoryId_INITIALIZER */ \
            NDDS_Discovery_Property_INITIALIZER \
        }, /* discovery_component */ \
        DDS_BOOLEAN_FALSE /* accept_unknown_peers */ \
    }, \
    { /* resource_limits */ \
        1L, /* local_writer_allocation */ \
        1L, /* local_reader_allocation */ \
        1L, /* local_publisher_allocation */ \
        1L, /* local_subscriber_allocation */ \
        1L, /* local_topic_allocation */ \
        1L, /* local_type_allocation */ \
        8L, /* remote_participant_allocation */ \
    } \
}

```

(continues on next page)

(continued from previous page)

```

8L, /* remote_writer_allocation */ \
8L, /* remote_reader_allocation */ \
32L, /* matching_writer_reader_pair_allocation */ \
32L, /* matching_reader_writer_pair_allocation */ \
32L, /* max_receive_ports */ \
32L, /* max_destination_ports */ \
65536, /* unbound_data_buffer_size */ \
500UL, /* shm_ref_transfer_mode_max_segments */ \
0L, /* participant_user_data_max_length */ \
DDS_SIZE_AUTO, /* participant_user_data_max_count */ \
0L, /* topic_data_max_length */ \
DDS_SIZE_AUTO, /* topic_data_max_count */ \
0L, /* publisher_group_data_max_length */ \
DDS_SIZE_AUTO, /* publisher_group_data_max_count */ \
0L, /* subscriber_group_data_max_length */ \
DDS_SIZE_AUTO, /* subscriber_group_data_max_count */ \
0L, /* writer_user_data_max_length */ \
DDS_SIZE_AUTO, /* writer_user_data_max_count */ \
0L, /* reader_user_data_max_length */ \
DDS_SIZE_AUTO, /* reader_user_data_max_count */ \
64L, /* max_partitions */ \
256L, /* max_partition_cumulative_characters */ \
DDS_LENGTH_UNLIMITED, /* max_partition_string_size */ \
DDS_LENGTH_UNLIMITED /* max_partition_string_allocation */ \
}, \
DDS_ENTITY_NAME_QOS_POLICY_DEFAULT, \
DDS_WIRE_PROTOCOL_QOS_POLICY_DEFAULT, \
{ /* transports */ \
    REDA_StringSeq_INITIALIZER_W_LOAN>HelloWorldAppLibrary>HelloWorldDPDEPubDP_
↪transport_enabled_transports, 1, 1) /* enabled_transports */ \
}, \
{ /* user_traffic */ \
    REDA_StringSeq_INITIALIZER_W_LOAN>HelloWorldAppLibrary>HelloWorldDPDEPubDP_user_
↪traffic_enabled_transports, 1, 1) /* enabled_transports */ \
}, \
DDS_TRUST_QOS_POLICY_DEFAULT, \
DDS_PROPERTY_QOS_POLICY_DEFAULT, \
DDS_USER_DATA_QOS_POLICY_DEFAULT \
,{ /* filter */ \
    DDS_FILTER_PLUGIN_FACTORY_DEFAULT_ID, /* name */ \
    { /* resource_limits */ \
        7L, /* filter_class_max_length */ \
        2L, /* filter_class_max_count */ \
        63L, /* filter_expression_max_length */ \
        DDS_LENGTH_AUTO, /* filter_expression_max_count */ \
        4L, /* filter_parameter_max_count_per_expression */ \
        15L, /* filter_parameter_max_length */ \
        4L /* sql_predicate_max_count_per_expression */ \
    }, \
    DDS_BOOLEAN_TRUE, /* disable_writer_filtering */ \
    DDS_BOOLEAN_TRUE /* disable_builtin_sql_filter */ \
}

```

(continues on next page)

(continued from previous page)

```

} \
}

```

### 3.18.5 Errors Caused by Invalid Configurations and QoS

This section explains the different results thrown by MAG if it receives invalid configuration files.

- Invalid XML content

MAG will fail to validate the configuration file if it contains invalid content, such as elements/attributes that don't exist in the schema or values that aren't supported by any of the existing types. For example:

```

<dds>
  ...
  <!-- Participant Library -->
  <domain_participant_library name="FeatureTestLibrary">
    <domain_participant name="01_EmptyDomainParticipant"
      domain_ref="HelloWorldLibrary::HelloWorldDomain">
      <invalid_tag></invalid_tag>
    </domain_participant>
  </domain_participant_library>
  ...
</dds>

```

```

07:41:48.334 [main] INFO com.rti.micro.appgen.MicroAppGen - Processing file : /
home/test/Error.xml
07:41:49.827 [main] ERROR com.rti.micro.appgen.MicroAppGen - Failed to parse inp
ut file : /home/test/Error.xml
07:41:49.837 [main] ERROR com.rti.micro.appgen.MicroAppGen - cvc-complex-type.2.
4.a: Invalid content was found starting with element 'invalid_tag'. One of '{dat
a_writer, publisher_qos}' is expected.
07:41:49.837 [main] INFO com.rti.micro.appgen.MicroAppGen - Exiting.

```

- Unsupported elements

MAG will throw a warning for any elements that are not supported by *Connext Micro*. Unsupported elements will be ignored, such as the `user_data` in the following:

```

<dds>
  ...
  <!-- Participant Library -->
  <domain_participant_library name="FeatureTestLibrary">
    <domain_participant name="01_EmptyDomainParticipant"
      domain_ref="HelloWorldLibrary::HelloWorldDomain">
      <domain_participant_qos>
        <!-- user_data is not supported by Micro -->
        <user_data/>
      </domain_participant_qos>
    </domain_participant>
  </domain_participant_library>

```

(continues on next page)

(continued from previous page)

```

    </domain_participant_library>
</dds>

```

```

07:39:52.643 [main] INFO com.rti.micro.appgen.MicroAppGen - Processing file : /
home/test/Warning.xml
07:39:53.439 [main] WARN com.rti.micro.appgen.utils.ConverterUtils - userData i
s not supported by Micro, the tool will ignore its value.
file=/home/test/Warning.xml, lineNumber=90, columnNumber=38

```

- **Unsupported values**

MAG will throw an error if it finds a value that is not supported by *Connext Micro*.

```

<dds>
...
<!-- Participant Library -->
<domain_participant_library name="FeatureTestLibrary">
  <domain_participant name="01_EmptyDomainParticipant"
    domain_ref="HelloWorldLibrary::HelloWorldDomain">
    <publisher name="test">
      <data_writer topic_ref="HelloWorldTopic1" name="testW">
        <datawriter_qos>
          <durability>
            <!-- transient is not supported by Micro -->
            <kind>TRANSIENT_DURABILITY_QOS</kind>
          </durability>
        </datawriter_qos>
      </data_writer>
    </publisher>
  </domain_participant>
</domain_participant_library>
</dds>

```

```

07:39:01.248 [main] INFO com.rti.micro.appgen.MicroAppGen - Processing file : /
home/test/Error.xml
07:39:02.069 [main] ERROR com.rti.micro.appgen.utils.ConverterUtils - TRANSIENT_
DURABILITY_QOS is not supported by Micro, only VOLATILE and TRANSIENT_LOCAL are
valid values for the durability kind field.
file=/home/test/Error.xml, lineNumber=35, columnNumber=37
07:39:02.072 [main] ERROR com.rti.micro.appgen.MicroAppGen - Failed to add input
file information into the model.
07:39:02.074 [main] INFO com.rti.micro.appgen.MicroAppGen - Exiting.

```

MAG will throw an error if the QoS values are not consistent with values supported in *Connext Micro*. For example, the following XML contains a deadline period that is too large.

```

<dds>
...
<!-- Participant Library -->
<domain_participant_library name="FeatureTestLibrary">
  <domain_participant name="01_EmptyDomainParticipant"
    domain_ref="HelloWorldLibrary::HelloWorldDomain">
    <publisher name="test">

```

(continues on next page)

(continued from previous page)

```

<data_writer topic_ref="HelloWorldTopic1" name="testW
→">
    <datawriter_qos>
        <deadline>
            <!-- this deadline exceeds the maximum --
→>
        <period>
            <sec>123213123</sec>
            <nanosec>12</nanosec>
        </period>
    </deadline>
</datawriter_qos>
</data_writer>
</publisher>
</domain_participant>
</domain_participant_library>
</dds>

```

```

07:43:26.805 [main] INFO com.rti.micro.appgen.MicroAppGen - Processing file : /
home/test/Error.xml
07:43:27.619 [main] ERROR com.rti.micro.appgen.utils.ConverterUtils - The durati
on of deadline.period=3.90706250000000038052 y exceeded the maximum range [1 ns,
1 year]
file=/home/test/Error.xml, lineNumber=35, columnNumber=11
07:43:27.620 [main] ERROR com.rti.micro.appgen.MicroAppGen - Failed to add input
file information into the model.
07:43:27.620 [main] INFO com.rti.micro.appgen.MicroAppGen - Exiting.

```

MAG will throw an error if the `dds.xtypes.compliance_mask` property uses a different value than `0x00000008`.

```

file=/tmp/Error.xml, lineNumber=104, columnNumber=38
09:29:36.451 [main] ERROR com.rti.micro.appgen.utils.ConverterUtils - 0x00000009 is not a valid value for dds.xtypes.compliance_mask, only 0x00000008 is supported.
file=/tmp/Error.xml, lineNumber=305, columnNumber=27
09:29:36.451 [main] ERROR com.rti.micro.appgen.MicroAppGen - Failed to calculate the system model.
java.lang.Exception: Invalid values found while creating the Model.

```

### • Unsupported QoS

Not all the QoS policies supported by *Connext Micro* can be configured in XML.

- QoS settings related to UDP transformation cannot be configured in XML. See the *UDP Transport* section for more information on UDP transformation.
- MAG does not support any PROPERTY QoS policy properties except the `dds.xtypes.compliance_mask` property.

## 3.19 Building Against FACE Conformance Libraries

This section describes how to build *Connext Micro* using the FACE™ conformance test tools.

### 3.19.1 Requirements

#### Connex Micro Source Code

The *Connex Micro* source code is available from [RTI's Support portal](#).

#### FACE Conformance Tools

RTI does not distribute the FACE conformance tools.

#### CMake

The *Connex Micro* source is distributed with a **CMakeList.txt** project file. CMake is an easy to use tool that generates makefiles or project files for various build-tools, such as Linux/macOS makefiles, Microsoft® Visual Studio® project files, and Xcode.

CMake can be downloaded from <https://www.cmake.org>.

### 3.19.2 FACE Golden Libraries

The FACE conformance tools use a set of golden libraries. There are different golden libraries for different FACE services, languages and profiles. *Connex Micro only* conforms to the safetyExt and safety profile of OSS using the C language.

#### Building the FACE Golden Libraries

The FACE conformance tools ship with their own set of tools to build the golden libraries. Please follow the instructions provided by FACE. In order to build the FACE golden libraries, it is necessary to port to the required platform. RTI has only tested *Connex Micro* on Linux 2.6 systems with GCC 4.4.5. The complete list of files modified by RTI are included below in source form.

### 3.19.3 Building the Connex Micro Source

The following instructions show how to build the *Connex Micro* source:

- Extract the source-code. Please note that the remaining instructions assume that only a single platform is built from the source.
- In the top-level source directory, enter the following:

```
shell> cmake-gui .
```

This will start the CMake GUI where all build configuration takes place.

- Click the “Configure” button.
- Select “Unix Makefiles” from the drop-down list.

- Select “Use default compilers” or “Specify native compilers” as required. Press “Done.”
- Click “Configure” again. There should not be any red lines. If there are, click “Configure” again.

NOTE: A red line means that a variable has not been configured. Some options could add new variables. Thus, if you change an option a new red lines may appear. In this case configure the variable and press “Configure.”

- Expand the CMAKE and RTIMICRO options and configure how to build *Connext Micro*:

```

CMAKE_BUILD_TYPE: Debug or blank. If Debug is used, the |me| debug
                    libraries are built.

RTIMICRO_BUILD_API: C or C++
  C    - Include the C API. For FACE, only C is supported.
  C++  - Include the C++ API.

RTIMICRO_BUILD_DISCOVERY_MODULE: Dynamic | Static | Both
  Dynamic - Include the dynamic discovery module.
  Static   - Include the static discovery module.
  Both     - Include both discovery modules.

RTIMICRO_BUILD_LIBRARY_BUILD:
  Single    - Build a single library.
  RTI style - Build the same libraries RTI normally ships. This is useful
               if RTI libraries are already being used and you want to use
               the libraries built from source.

RTIMICRO_BUILD_LIBRARY_TYPE:
  Static    - Build static libraries.
  Shared    - Build shared libraries.

RTIMICRO_BUILD_LIBRARY_PLATFORM_MODULE: POSIX

RTIMICRO_BUILD_LIBRARY_TARGET_NAME: <target name>
  Enter a string as the name of the target. This is also used as the
  name of the directory where the built libraries are placed.
  If you are building libraries to replace the libraries shipped by RTI,
  you can use the RTI target name here. It is then possible to set
  RTIMEHOME to the source tree (if RTI style is selected for
  RTIMICRO_BUILD_LIBRARY_BUILD).

RTIMICRO_BUILD_ENABLE_FACE_COMPLIANCE: Select level of FACE compliance
  None          - No compliance required
  General        - Build for compliance with the FACE general profile
  Safety Extended - Build for compliance with the FACE safety extended profile
  Safety         - Build for compliance with the FACE safety profile

RTIMICRO_BUILD_LINK_FACE_GOLDEBLIBS:
  Check if linking against the static FACE conformance test libraries.
  NOTE: This check-box is only available if FACE compliance is different
  from "None".

```

(continues on next page)



(continued from previous page)

**RTIMICRO\_BUILD\_LINK\_FACE\_GOLDEBLIBS:**

If the `RTIMICRO_BUILD_LINK_FACE_GOLDEBLIBS` is checked the path to the top-level FACE root must be specified here.

- Click “Configure”.
- Click “Generate”.
- Build the generated project.
- Libraries are placed in `lib/<RTIMICRO_BUILD_LIBRARY_TARGET_NAME>`.

## 3.20 Working With Sequences

### 3.20.1 Introduction

*RTI Connex Micro* uses IDL as the language to define data-types. One of the constructs in IDL is the *sequence*: a variable-length vector where each element is of the same type. This section describes how to work with sequences; in particular, the string sequence since it has special properties.

**Note:** This section references several sequence APIs supported by *Connex Micro*. However, *Connex Cert* only supports a subsection of these APIs. Please refer to [Sequence Support in the C API Reference](#) for a full list; the APIs supported by *Connex Cert* will have a «**cert**» annotation in their description.

### 3.20.2 Working with Sequences

#### Overview

Logically a sequence can be viewed as a variable-length vector with N elements, as illustrated below. Note that sequences indices are 0 based.

```

+----+
0 | T |
+----+
1 | T |
+----+
2 | T |
+----+
  |
  |
+----+
N-1 | T |
+----+
```

There are three types of sequences in *Connex Micro*:

- Builtin sequences of primitive IDL types.
- Sequences defined in IDL using the sequence keyword.
- Sequences defined by the application.

The following builtin sequences exist (please refer to [C API Reference](#) and [C++ API Reference](#) for the complete API).

IDL Type	<i>Connex Micro</i> Type	<i>Connex Micro</i> Sequence
octet	DDS__Octet	DDS__OctetSeq
char	DDS__Char	DDS__CharSeq
boolean	DDS__Boolean	DDS__BooleanSeq
short	DDS__Short	DDS__ShortSeq
unsigned short	DDS__UnsignedShort	DDS__UnsignedShortSeq
long	DDS__Long	DDS__LongSeq
unsigned long	DDS__UnsignedLong	DDS__UnsignedLongSeq
enum	DDS__Enum	DDS__EnumSeq
wchar	DDS__Wchar	DDS__WcharSeq
long long	DDS__LongLong	DDS__LongLongSeq
unsigned long long	DDS__UnsignedLongLong	DDS__UnsignedLongLongSeq
float	DDS__Float	DDS__FloatSeq
double	DDS__Double	DDS__DoubleSeq
long double	DDS__LongDouble	DDS__LongDoubleSeq
string	DDS__String	DDS__StringSeq
wstring	DDS__Wstring	DDS__WstringSeq

The following are important properties of sequences to remember:

- All sequences in *Connex Micro* must be finite.
- All sequences defined in IDL are sized based on IDL properties and *must* not be resized. That is, *never* call [set\\_maximum\(\)](#) on a sequence defined in IDL. This is particularly important for string sequences.
- Application defined sequences can be resized using [set\\_maximum\(\)](#).
- There are two ways to use a **DDS\_\_StringSeq** (they are type-compatible):
  - A **DDS\_\_StringSeq** originating from IDL. This sequence is sized based on maximum sequence length *and* maximum string length.
  - A **DDS\_\_StringSeq** originating from an application. In this case the sequence element memory is unmanaged.
- All sequences have an initial length of 0.

## Working with IDL Sequences

Sequences that originate from IDL are created when the IDL type they belong to is created. IDL sequences are always initialized with the maximum size specified in the IDL file. The maximum size of a type, and hence the sequence size, is used to calculate memory needs for serialization and deserialization buffers. Thus, changing the size of an IDL sequence can lead to hard to find memory corruption.

The string and wstring sequences are special in that not only is the maximum sequence size allocated, but because strings are also always of a finite maximum length, the maximum space needed for each string element is also allocated. This ensure that *Connext Micro* can prevent memory overruns and validate input.

Some typical scenarios with a long sequence and a string sequence defined in IDL is shown below:

```
/* In IDL */
struct SomeIdlType
{
    // A sequence of 20 longs
    sequence<long,20> long_seq;

    // A sequence of 10 strings, each string has a maximum length of 255 bytes
    // (excluding NUL)
    sequence<string<255>,10> string_seq;
}

/* In C source */
SomeIdlType *my_sample = SomeIdlTypeTypeSupport_create_data()

DDS_LongSeq_set_length(&my_sample->long_seq,5);
DDS_StringSeq_set_length(&my_sample->string_seq,5);

/* Assign the first 5 longs in long_seq */
for (i = 0; i < 5; ++i)
{
    *DDS_LongSeq_get_reference(&my_sample->long_seq,i) = i;
    snprintf(*DDS_StringSeq_get_reference(&my_sample->string_seq,0),255,"SomeString %d",
    ↪i);
}

SomeIdlTypeTypeSupport_delete_data(my_sample);

/* In C++ source */
SomeIdlType *my_sample = SomeIdlTypeTypeSupport::create_data()

/* Assign the first 5 longs in long_seq */

my_sample->long_seq.length(5);
my_sample->string_seq.length(5);

for (i = 0; i < 5; ++i)
{
```

(continues on next page)

(continued from previous page)

```

    /* use method */
    *DDSLongSeq_get_reference(&my_sample->long_seq,i) = i;
    snprintf(*DDSStringSeq_get_reference(&my_sample->string_seq,i),255,"SomeString %d",
    ↪i);

    /* or assignment */
    my_sample->long_seq[i] = i;
    snprintf(my_sample->string_seq[i],255,"SomeString %d",i);
}

SomeIdlTypeTypeSupport::delete_data(my_sample);

```

Note that in the example above the sequence length is set. The maximum size for each sequence is set when `my_sample` is allocated.

A special case is to copy a string sequence from a sample to a string sequence defined outside of the sample. This is possible, but care *must* be taken to ensure that the memory is allocated properly:

Consider the IDL type from the previous example. A string sequence of equal size can be allocated as follows:

```

struct DDS_StringSeq app_seq = DDS_SEQUENCE_INITIALIZER;

/* This ensures that memory for the strings are allocated upfront */
DDS_StringSeq_set_maximum_w_max(&app_seq,10,255);

DDS_StringSeq_copy(&app_seq,&my_sample->string_seq);

```

If instead the following code was used, memory for the string in `app_seq` would be allocated as needed.

```

struct DDS_StringSeq app_seq = DDS_SEQUENCE_INITIALIZER;

/* This ensures that memory for the strings are allocated upfront */
DDS_StringSeq_set_maximum(&app_seq,10);

DDS_StringSeq_copy(&app_seq,&my_sample->string_seq);

```

## Working with Application Defined Sequences

Application defined sequences work in the same way as sequences defined in IDL with two exceptions:

- The maximum size is 0 by default. It is necessary to call `set_maximum()` or `ensure_length` to allocate space.
- `DDS_StringSeq_set_maximum` does not allocate space for the string pointers. The memory must be allocated on a per needed basis and calls to `__copy` may reallocate memory as needed. Use `DDS_StringSeq_set_maximum_w_max` or `DDS_StringSeq_ensure_length_w_max` to also allocate pointers. In this case `__copy` will *not* reallocate

memory.

Note that it is not allowed to mix the use of calls that pass the max (ends in `__w__max`) and calls that do not. Doing so may cause memory leaks and/or memory corruption.

```
struct DDS_StringSeq my_seq = DDS_SEQUENCE_INITIALIZER;

DDS_StringSeq_ensure_length(&my_seq,10,20);

for (i = 0; i < 10; i++)
{
    *DDS_StringSeq_get_reference(&my_seq,i) = DDS_String_dup("test");
}

DDS_StringSeq_finalize(&my_seq);
```

`DDS_StringSeq_finalize` automatically frees memory pointed to by each element using `DDS_String_free`. All memory allocated to a string element should be allocated using a `DDS_String` function.

It is possible to assign any memory to a string sequence element if all elements are released manually first:

```
struct DDS_StringSeq my_seq = DDS_SEQUENCE_INITIALIZER;

DDS_StringSeq_ensure_length(&my_seq,10,20);

for (i = 0; i < 10; i++)
{
    *DDS_StringSeq_get_reference(&my_seq,i) = static_string[i];
}

/* Work with the sequence */

for (i = 0; i < 10; i++)
{
    *DDS_StringSeq_get_reference(&my_seq,i) = NULL;
}

DDS_StringSeq_finalize(&my_seq);
```

## 3.21 Debugging

### 3.21.1 Overview

*Connext Micro* maintains a log of events occurring in a *Connext Micro* application. Information on each event is formatted into a log entry. Each entry can be stored in a buffer, stringified into a displayable log message, and/or redirected to a user-defined log handler.

For a list of error codes, please refer to [Logging Reference](#).

### 3.21.2 Configuring Logging

By default, *Connex Micro* sets the log verbosity to *Error*. It can be changed at any time by calling **OSAPI\_Log\_set\_verbosity()** using the desired verbosity as a parameter.

Note that when compiling with `RTI_CERT` defined, logging is completely removed.

The *Connex Micro* log stores new log entries in a log buffer.

The default buffer size is different for Debug and Release libraries. The Debug libraries are configured to use a much larger buffer than the Release ones. A custom buffer size can be configured using the **OSAPI\_Log\_set\_property()** function. For example, to set a buffer size of 128 bytes:

```
struct OSAPI_LogProperty prop = OSAPI_LogProperty_INIITALIZER;

OSAPI_Log_get_property(&prop);
prop.max_buffer_size = 128;
OSAPI_Log_set_property(&prop);
```

Note that if the buffer size is too small, log entries will be truncated in order to fit in the available buffer.

The function used to write the logs can be set during compilation by defining the macro `OSAPI_LOG_WRITE_BUFFER`. This macro shall have the same parameters as the function prototype **OSAPI\_Log\_write\_buffer\_T**.

It is also possible to set this function during runtime by using the function **OSAPI\_Log\_set\_property()**:

```
struct OSAPI_LogProperty prop = OSAPI_LogProperty_INIITALIZER;

OSAPI_Log_get_property(&prop);
prop.write_buffer = <pointer to user defined write function>;
OSAPI_Log_set_property(&prop);
```

A user can install a log handler function to process each new log entry. The handler must conform to the definition `OSAPI_LogHandler_T`, and it is set by **OSAPI\_Log\_set\_log\_handler()**.

When called, the handler has parameters containing the raw log entry and detailed log information (e.g., error code, module, file and function names, line number).

The log handler is called for every new log entry, even when the log buffer is full. An expected use case is redirecting log entries to another logger, such as one native to a particular platform.

### 3.21.3 Log Message Kinds

Each log entry is classified as one of the following kinds:

- *Error*. An unexpected event with negative functional impact.
- *Warning*. An event that may not have negative functional impact but could indicate an unexpected situation.
- *Information*. An event logged for informative purposes.

By default, the log verbosity is set to *Error*, so only error logs will be visible. To change the log verbosity, simply call the function `OSAPI_Log_set_verbosity()` with the desired verbosity level.

### 3.21.4 Interpreting Log Messages and Error Codes

A log entry in *Connext Micro* has a defined format.

Each entry contains a header with the following information:

- *Length*. The length of the log message, in bytes.
- *Module ID*. A numerical ID of the module from which the message was logged.
- *Error Code*. A numerical ID for the log message. It is unique within a module.

Though referred to as an “error” code, it exists for all log kinds (error, warning, info).

The module ID and error code together uniquely identify a log message within *Connext Micro*.

*Connext Micro* can be configured to provide additional details per log message:

- *Line Number*. The line number of the source file from which the message is logged.
- *Module Name*. The name of the module from which the message is logged.
- *Function Name*. The name of the function from which the message is logged.

When an event is logged, by default it is printed as a message to standard output. An example error entry looks like this:

```
[943921909.645099999]ERROR: ModuleID=7 Errcode=200 X=1 E=0 T=1
dds_c/DomainFactory.c:163/DDS_DomainParticipantFactory_get_instance: kind=19
```

- *X* Extended debug information is present, such as file and line number.
- *E* Exception, the log message has been truncated.
- *T* The log message has a valid timestamp (successful call to `OSAPI_System_get_time()`).

A log message will need to be interpreted by the user when an error or warning has occurred and its cause needs to be determined, or the user has set a log handler and is processing each log message based on its contents.

A description of an error code printed in a log message can be determined by following these steps:

- Navigate to the module that corresponds to the Module ID, or the printed module name in the second line. In the above example, “ModuleID=7” corresponds to DDS.
- Search for the error code to find it in the list of the module’s error codes. In the example above, with “Errcode=200,” search for “200” to find the log message that has the value “(DDSC\_LOG\_BASE + 200)”.

## 3.22 Example Applications

The RTI *Code Generator* (*rtiddsgen*) can generate example applications from an IDL file, as shown in *Hello World with Publish/Subscribe*. The last type in the IDL is used as the data-type to publish. Refer to *Generating examples* for instructions on how to build these examples.

*Connext Micro* also provides some buildable applications in the installation directory, as described in *Provided examples*.

### 3.22.1 Generating examples

---

**Note:** Before running *rtiddsgen*, you might need to add `rti_connext_dds-<version>/rtiddsgen/scripts` to your path environment variable folder.

---

#### Default example

To generate an example, run the following command:

```
rtiddsgen -example -language <C|C++> [-namespace] <file with type definition>
```

This generates an example using the default example template, which uses the Dynamic Participant Dynamic Endpoint (DPDE) discovery plugin.

*rtiddsgen* accepts the following options:

- **-example:** Generates type files, example files, and CMakeLists files.
- **-language <C|C++>:** Generates C or C++ code.
- **-namespace:** Enables C++ namespaces when the language option is C++.

The generated example can then be compiled using [CMake](#) and the `CMakeLists.txt` file generated by *Code Generator*. *Code Generator* also creates a `README.txt` file with a description of the example and instructions for how to compile and run it.



## Custom example

*Code Generator* can also generate examples using custom templates with the option `-exampleTemplate <templateName>`.

To generate an example using a custom template instead of the default one, run the following command:

```
rtiddsgen -example -exampleTemplate <template name> -language <C|C++> [-namespace] <file_
↳with type definition>
```

To see the list of the available templates for each language, run the following command:

```
rtiddsgen -showTemplates
```

As an example, the following command will generate an example in the C language, using the **waitsets** custom template instead of the default template:

List of example templates per language:

- C:
  - cert
  - dpse
  - shared\_memory
  - static\_udp
  - waitsets
  - crc
  - zcv2
  - mag/shared\_memory
  - mag/static\_udp
  - mag/dpde
  - mag/dpse
  - psk
- C++:
  - dpse
  - waitsets
  - mag/dpde
- C++ Namespace:
  - dpse
  - waitsets

The following command will generate an example in the C language, using the ‘waitsets’ custom template instead of the default template:

```
rtiddsgen -example -exampleTemplate waitsets -language C <file with type definition>
```

## Descriptions of generated examples

Each example consists of a publication and subscription pair to send and receive the type specified by the user. When compiled, the example creates two applications: one to send samples (a publisher) and another to receive samples (a subscriber).

- **default example (no template specified)**

Discovery of endpoints is done with the *Dynamic Discovery Plugin* (DPDE). Only the *UDP* and *INTRA* transports are enabled. The subscriber application creates a *DataReader*, which uses a listener to receive notifications about new samples and matched publishers. These notifications are received in the middleware thread (instead of the application thread).

- **cert**

An example that only uses APIs that are compatible with *Connext Cert*.

- **dpse**

Discovery of endpoints is done with the *Static Discovery Plugin* (DPSE). Static-endpoint discovery uses function calls to statically assert information about remote endpoints belonging to remote *DomainParticipants*.

- **shared\_memory**

The only transport used is shared memory. Because the *UDP Transport* is disabled and only the *Shared Memory Transport* (*SHMEM*) is enabled, both the publisher and subscriber applications need to run in the same operating system.

- **static\_udp**

This example uses a static *UDP Transport* interface configuration. Using this API, the UDP transport is statically configured. This is useful in systems that are not able to return the installed UDP interfaces (name, IP address, mask, etc.).

- **waitsets**

In this example, the Subscriber application creates a *DataReader* that uses a [Waitset](#) (instead of a listener) to receive notifications about new samples and matched publishers. These notifications are received in the middleware thread (instead of the application thread).

- **crc**

This example includes configuration of the *Message Integrity Checking* settings. The CRC fields in the [WireProtocolQosPolicy](#) enable generating a checksum value based on the data being transmitted. They also determine which checksums are allowed and whether they should be sent.

- **zcv2**

This example uses the *Zero Copy v2 Transport* and handles discovery over the *UDP Transport*. User traffic is handled using the Zero Copy v2 interface.

- **psk**

This example uses pre-shared key (PSK) encryption to secure communications. You must have the *Lightweight Security Plugin* installed in order to compile this example.

### How to compile the generated examples

Before compiling, set the environment variable `RTIMEHOME` to the *Connex Micro* installation directory.

The *Connex Micro* source bundle includes `rtime-make` (on Linux® and macOS® systems) or `rtime-make.bat` (on Windows® systems) to simplify invocation of CMake. This script is a convenient way to invoke CMake with the correct options. For example:

Linux

```
cd <directory with generated example>

rtime-make --config <Debug|Release> --build --target armv8leElfgcc7.3.0-Linux4 --source-
↪dir . \
    -G "Unix Makefiles" --delete [-DRTIME_IDL_ADD_REGENERATE_TYPESUPPORT_RULE=true]
```

macOS

```
cd <directory with generated example>

rtime-make --config <Debug|Release> --build --target x86_64leMachOclang15.0-Darwin23 --
↪source-dir . \
    -G "Unix Makefiles" --delete [-DRTIME_IDL_ADD_REGENERATE_TYPESUPPORT_RULE=true]
```

Windows

```
cd <directory with generated example>

rtime-make.bat --config <Debug|Release> --build --target x86_64lePEvs2017-Win10 --source-
↪dir . \
    -G "Visual Studio 15 2017" --delete [-DRTIME_IDL_ADD_REGENERATE_TYPESUPPORT_RULE_eq_
↪true]
```

**Warning:** RTI recommends using the toolchain file that matches the target architecture to compile the generated examples.

For example, if the target architecture is `--target armv8leElfgcc7.3.0-Linux4`, then the example applications should be compiled with the `armv8leElfgcc7.3.0-Linux4` toolchain file. Failing to do so may cause warnings.

The executable can be found in the directory ‘`objs`’.

It is also possible to compile using CMake, e.g., when the *Connex Micro* source bundle is not installed.

Linux

```
cmake [-DRTIME_IDL_ADD_REGENERATE_TYPESUPPORT_RULE=true] \ [-DCMAKE_BUILD_TYPE=
↪<Debug|Release>] \
    -G "Unix Makefiles" -B./<your build directory> -H. -DRTIME_TARGET_
↪NAME=armv8leElfgcc7.3.0-Linux4

cmake --build ./<your build directory> [--config <Debug|Release>]
```

## macOS

```
cmake [-DRTIME_IDL_ADD_REGENERATE_TYPESUPPORT_RULE=true] \ [-DCMAKE_BUILD_TYPE=
↪<Debug|Release>] \
    -G "Unix Makefiles" -B./<your build directory> -H. -DRTIME_TARGET_NAME=x86_
↪64leMachOclang15.0-Darwin23

cmake --build ./<your build directory> [--config <Debug|Release>]
```

## Windows

```
cmake [-DRTIME_IDL_ADD_REGENERATE_TYPESUPPORT_RULE=true] \ [-DCMAKE_BUILD_TYPE=
↪<Debug|Release>] \
    -G "Visual Studio 15 2017" -B./<your build directory> -H. -DRTIME_TARGET_NAME=x86_
↪64lePEvs2017-Win10

cmake --build .\<your build directory> [--config <Debug|Release>]
```

The executable can be found in the directory ‘objs’.

The following options are accepted:

- `-DRTIME_IDL_ADD_REGENERATE_TYPESUPPORT_RULE=true` adds a rule to regenerate type support plugin source files if the input file with the type definition changes. Default value is ‘false’.

## How to run the generated examples

By default, the example uses all available interfaces to receive samples. This can cause communication problems if the number of available interfaces is greater than the maximum number of interfaces supported by *Connext Micro*. For this reason, it is recommended to restrict the number of interfaces used by the application. Use the option `-udp_intf <interface name>` when running the example.

For example, if the example has been compiled for Linux i86Linux2.6gcc4.4.5, run the subscriber with this command:

```
objs/armv8leElfgcc7.3.0-Linux4/<Type definition file name>_subscriber [-domain
↪<Domain_ID>] [-peer <address>] \
    [-sleep <sleep_time>] [-count <seconds_to_run>] [-udp_intf
↪<interface name>]
```

and run the publisher with this command:

```

objs/armv8leElfgcc7.3.0-Linux4/<Type definition file name>_publisher [-domain
↪<Domain_ID> -peer <address>] \
    [-sleep <sleep_time>] [-count <seconds_to_run>] [-udp_intf
↪<interface name>]

```

---

**Note:** Shared memory examples only accept the following options:

- [-domain <Domain\_ID>]
  - [-sleep <sleep\_time>]
  - [-count <seconds\_to\_run>]
- 

### 3.22.2 Provided examples

*Connext Micro* provides buildable example applications in the `example/` directory of your installation. Each example demonstrates different features, as described below:

- **HelloWorld\_transformations:** A HelloWorld example that uses *UDP transformations* to send encrypted packets using OpenSSL.
- **HelloWorld\_dgram:** A HelloWorld example that uses an example *DGRAM* implementation.

Consult the `README.txt` file included with each example for instructions on how to build and run the application.

## 3.23 Lightweight Security Plugin

The *Lightweight Security Plugin* is an addon for *Connext Micro* that offers an efficient solution for securing RTPS communication at the domain level. This plugin protects your domain from unauthorized participants and ensures the integrity and confidentiality of RTPS messages.

### 3.23.1 Overview

The *Lightweight Security Plugin* uses symmetric encryption, deriving keys from a combination of publicly-available data and a user-provided pre-shared key (PSK) seed. Pre-shared key protection offers metadata and data protection to RTPS messages and restricts communication only to *DomainParticipants* holding the correct PSK. This architecture offers a streamlined path to deploy security without the overhead of full public key infrastructure.

The *Lightweight Security Plugin* consists of a Platform Independent Library (PIL) and a Transform Platform Support Library (PSL). The Transform PSL is an abstraction layer that allows seamless

integration with a cryptographic backend. The plugin includes a reference implementation that integrates with [OpenSSL 3.5](#)<sup>1</sup>.

This modular design makes the *Lightweight Security Plugin* an ideal choice for applications that demand robust protection without compromising on portability or efficiency.

### 3.23.2 Downloading and Installing the Lightweight Security Plugin

Once you have purchased the *Lightweight Security Plugin*, you can download the following package files from [support.rti.com](https://support.rti.com):

- `rti_connext_dds_micro-<version>-lw-security-host.rtipkg`
- `rti_connext_dds_micro-<version>-lw-security-target-openssl-3.5-<architecture>.rtipkg`
- `openssl-3.5-micro-<version>-target-<architecture>.rtipkg` (*Optional*)

Download the host package with `<version>` 4.2.0, and a target package with `<version>` 4.2.0 and `<architecture>` matching your CPU and compiler.

To install the *Lightweight Security Plugins* packages, follow the same instructions for installing *Connext Micro* packages in *Installing Connext Micro*.

#### Package contents

The host package contains the necessary header files for lightweight security development.

The target package provides the Platform Independent Library (PIL) and Transform Platform Support Library (PSL) specific to your target architecture. The Transform PSL is dependent on OpenSSL; RTI builds and tests the PSL against [OpenSSL 3.5](#).

The OpenSSL target package simply contains the OpenSSL 3.5 release. RTI provides this for convenience, but you may use any OpenSSL distribution ABI compatible with 3.5.0. The OpenSSL source can be found at <https://openssl-library.org/source/>.

### 3.23.3 Getting Started

This section will walk you through the steps needed to enable the *Lightweight Security Plugin* in your application and run it in your domain. We will use the provided reference implementation of the Transform PSL with [OpenSSL 3.5](#).

---

**Note:** You can also generate an example application with *rtiddsgen* that uses lightweight security. Refer to *Example Applications* for more information.

---

---

<sup>1</sup> You can substitute in your own Transform PSL to meet performance, certification, or compliance requirements. For details, please contact RTI through [support.rti.com](https://support.rti.com).

## Link application and libraries

To link an application with the lightweight security feature, the following libraries are required in the listed order:

- RTIMEHOME/lib/<arch>/
  1. rti\_me\_ddspsk (must be linked before rti\_me)
- RTIMEHOME/lib/<arch>-<PSL>/
  2. rti\_me\_pskpsl (must be linked before rti\_me)
- RTIMEHOME/thirdparty/openssl-3.5.0-<arch>/<debug|release>-<shared|static>/lib
  3. libcrypto

## Enable lightweight security in your application

To use lightweight security, you must first register the *Lightweight Security Plugin* library with *Connext Micro* via the `DDS_PskLibrary_register()` function, as shown below:

```
{
    struct DDS_PskServiceFactoryProperty psk_svc_property =
        DDS_PskServiceFactoryProperty_INITIALIZER;
    psk_svc_property.psl_get_interface_func = PSK_OSSL_get_interface;
    if(!DDS_PskLibrary_register(
        registry,
        &psk_svc_property))
    {
        printf("Lightweight Security Plugin registration error\n");
    }
}
```

**Note:** When using the C++ API, register the *Lightweight Security Plugin* using the same method as the C API, shown above.

This function takes a property structure with the following fields:

- `service_name`: a unique identifier for this registration within the *DomainParticipantFactory*. The *DomainParticipant* QoS references this name to select which library registration to use. If set to NULL, the default name `DDS_PSK_DEFAULT_SUITE_NAME` will be used.
- `psl_get_interface_func`: a function pointer used to retrieve the Transform PSL interface. When using the reference Transform PSL provided by RTI, set this to `PSK_OSSL_get_interface`<sup>2</sup>.

<sup>2</sup> This function is provided by `rti_me_psl/pskpsl/psk_ssl_transform.h`.

- `psl_config`: An opaque pointer passed to the Transform PSL creation function. When using the RTI-provided Transform PSL, this field resolves to `DDS_SecTransform_Configuration`. Its type is `struct DDS_SecTransform_Configuration *`, as shown in the example below:

```
struct DDS_SecTransform_Configuration config = *PSK_OSSL_get_default_
↪ configuration();
```

### Specify the plugin in the DomainParticipant QoS

Assign the `service_name` that you specified during registration to the `DomainParticipantQos.trust.suite` field with the following function:

```
RT_ComponentFactoryId_set_name(&dp_qos.trust.suite, DDS_PSK_DEFAULT_SUITE_NAME)
```

This function tells the *DomainParticipant* to enable the *Lightweight Security Plugin* and identifies which registered instance to use.

### Configure the plugin

The *Lightweight Security Plugin* accepts the following properties (in the [PROPERTY QoS](#)):

- `dds.sec.crypto.rtps_psk_secret_passphrase` (*Required*)
- `dds.sec.access.rtps_psk_protection_kind` (*Optional*)
- `dds.sec.crypto.rtps_psk_symmetric_cipher_algorithm` (*Optional*)
- `com.rti.serv.secure.cryptography.max_blocks_per_session` (*Optional*)

For more information about these properties, refer to *Configuring the Lightweight Security Plugin*. For now, set these properties as shown in the example snippet below:

```
{
    struct DDS_DomainParticipantQos dp_qos =
        DDS_DomainParticipantQos_INITIALIZER;

    /* Required property */
    DDS_PropertyQosPolicyHelper_add_property(
        &dp_qos.property,
        "dds.sec.crypto.rtps_psk_secret_passphrase",
        "data:,404166165:castle super radar denial swing lunar kind swarm wet toilet_
↪ output harbor basic begin margin huge year visit",
        DDS_BOOLEAN_FALSE);

    /* Optional Properties */
    DDS_PropertyQosPolicyHelper_add_property(
        &dp_qos.property,
        "dds.sec.access.rtps_psk_protection_kind",
        "ENCRYPT",
        DDS_BOOLEAN_FALSE);
}
```

(continues on next page)



(continued from previous page)

```
DDS_PropertyQosPolicyHelper_add_property(  
    &dp_qos.property,  
    "dds.sec.crypto.rtps_psk_symmetric_cipher_algorithm",  
    "AUTO",  
    DDS_BOOLEAN_FALSE);  
  
DDS_PropertyQosPolicyHelper_add_property(  
    &dp_qos.property,  
    "com.rti.serv.secure.cryptography.max_blocks_per_session",  
    "256",  
    DDS_BOOLEAN_FALSE);  
}
```

### Create a protected DomainParticipant

Invoke the function [DDS\\_DomainParticipantFactory\\_create\\_participant](#) with the QoS `dp_qos`.

At this point, your application is protected by the *Lightweight Security Plugin* and ready to run!

#### 3.23.4 Configuring the Lightweight Security Plugin

The following table lists the properties used to configure the *Lightweight Security Plugin* and detailed descriptions of their use. Refer to the OMD [DDS Security Specification 1.1](#) for additional information.

All of the properties below have string values.

Property Name	Property Description
dds.sec.crypto.rtps_psk_secret_passphrase	<p>Passphrase used to derive the per-participant key used for encoding RTPS messages (in combination with other publicly available data).</p> <p>This property is mutable through the <i>DomainParticipant</i> <code>set_qos</code> API. The <i>DomainParticipant</i> will use the updated passphrase to derive the new local and remote pre-shared keys. From that point on, the <i>DomainParticipant</i> will use the new keys to encrypt and decrypt RTPS messages with PSK protection. You should expect warning messages (about receiving messages encoded with a different pre-shared key identifier) as you change the property value throughout your system. Once all your <i>DomainParticipants</i> have the same passphrase, all PSK protected traffic will be encoded and decoded with consistent keys. At that point, warning messages will stop.</p> <p>Updating the value of this property implies a change in the secret key AND in the key identifier. Each secret key should be associated with a unique key identifier. Updating the secret key without changing the key identifier is not allowed and will result in <code>set_qos</code> returning <code>DDS_RETCODE_NOT_ALLOWED_BY_SECURITY</code>.</p> <p>The pre-shared secret identifier should not be reused; you may set a new secret key with an old key identifier, but doing so is not advised. In that case, <i>Lightweight Security Plugin</i> may end up using the old pre-shared key to try (unsuccessfully) to decode an RTPS message that was encoded using the new key.</p> <p>The passphrase must follow the format <code>data:,&lt;passphrase_id&gt;:&lt;passphrase&gt;</code>.</p> <p><code>&lt;passphrase_id&gt;</code> is a unique four-byte integer number that identifies the key. It must be in the <math>[0, 2^{32} - 1]</math> range, excluding the values that have <code>0xFF</code> as the least significant byte (e.g., 255, 1023, 2047, etc.). <code>&lt;passphrase_id&gt;</code> must be provided to support proper handling of different passphrase revisions, and it must be consistent across all the <i>DomainParticipants</i> within a domain. This value must be different with every passphrase change, otherwise the update will be rejected. RTI recommends using incremental subsequent values for every new key provided. Rolling back to a lower number (e.g., 0) after the indexes have been exhausted is acceptable.</p> <p><code>&lt;passphrase&gt;</code> is the secret passphrase string that the <i>Lightweight Security Plugin</i> will use to derive the per-participant pre-shared key. <code>&lt;passphrase&gt;</code> must contain up to 512 ASCII printable characters (character codes 32 to 126, both included). The first and last characters of the <code>&lt;passphrase&gt;</code> shall not be the space character (character code 32).</p> <p><b>Default:</b> not set</p> <div style="border: 1px solid black; padding: 5px; margin-top: 10px;"> <p><b>Attention:</b> It is your responsibility to provide a key seed with good entropy and length. Ideally, the pre-shared key seed should be a true random 256-bit seed. There are a</p> </div>
3.23. Lightweight Security Plugin	245

### 3.23.5 Technical Considerations

Consider the following when using the *Lightweight Security Plugin*:

#### **RTPS transport only**

The *Lightweight Security Plugin* only protects transports that use the RTPS protocol. That means that INTRA and Zero Copy V2 transports cannot be protected with the *Lightweight Security Plugin*.

#### **Maximum Transmission Unit (MTU)**

RTI limits the maximum size of a single packet to 65535 bytes. If using a transport that can support an MTU greater than this, the size will be truncated to 65535 bytes.

---

**Note:** Types that exceed this limit can still be fragmented and sent by *Connex Micro*. This only limits the size of individual packets.

---

### 3.23.6 Interoperability

The *Lightweight Security Plugin* interoperates with the *Connex Professional Lightweight Builtin Security Plugins* security solution out of the box, with some considerations explained in the following sections.

The *Lightweight Security Plugin* is also compatible with *Connex Professional's Builtin Security Plugins* when the *Builtin Security Plugins* are configured for compatibility; refer to [Lightweight Builtin Security Plugins and Builtin Security Plugins Interoperability](#) for more information.

#### **Additional Authenticated Data (AAD)**

The *Lightweight Security Plugin* always uses AAD, which adds a header extension to RTPS messages. *Connex Professional* allows you to disable this. For successful interoperability between *Connex Micro* and *Connex Professional*, you must configure *Connex Professional* to always enable AAD by setting the `cryptography.enable_additional_authenticated_data` property to TRUE.

## Passphrase ID limitations

*Connext Micro* supports a `passphrase_id` (which is part of the passphrase) up to `MAX_UINT64`. However, *Connext 7.3.0* and above only supports a `passphrase_id` up to 254. To ensure interoperability between *Connext Micro* and *Connext Professional*, choose a `passphrase_id` within the range of 0 to 254.

## 3.24 Memory Management

*Connext Micro* is designed for use in real-time systems and uses a predictable, deterministic memory manager to ensure that memory growth is not unbounded, OS memory fragmentation is eliminated, and memory usage can be determined before runtime. This design ensures proper operation after all entities have been created, but requires the system designer to properly configure *Connext Micro*, as described in this section.

### 3.24.1 Resource limits

*Connext Micro* uses resource limits to determine how much to allocate of a particular resource; e.g., the maximum number of samples that can be cached by a *DataReader*, or how many *DataWriters* can be created locally. These limits may incur additional memory allocations. This section provides guidance on how much overhead these resource limits incur.

---

**Note:** Simple resource limits, such as those that specify a number of bytes, are not discussed in this section.

---

Please refer to [Resource Limits](#) for the full list of resource limits, and to [RTI Connext Performance Benchmarks](#) on RTI Community for memory usage benchmarks.

### 3.24.2 Dynamic memory allocation

*Connext Micro* allocates heap memory to create internal data structures dynamically as DDS entities are created. Keep in mind that *Connext Micro* manages allocated memory using its own internal memory management, and *only* returns freed allocated memory when something is deleted. So, if an application never deletes anything, no memory is ever freed.

---

**Note:** *Connext Micro* uses the term “heap” to refer to any memory that is allocated dynamically, typically using the `malloc` system call. However, the Platform Support Library (PSL) is responsible for implementing the memory allocation and freeing functions required by *Connext Micro*; the only requirement is that heap memory is valid from when it is allocated to when it is freed.

---

**Attention:** *Connext Micro* does not support dynamically allocating resources beyond the initial configuration. Therefore, all resource limits must be finite.

### 3.24.3 DDS resource limits

The following table shows which resource limits are applicable to DDS APIs that allocate memory:

Table 3.13: DDS APIs that allocates memory

API	Creates	Applicable Resource Limits
DDS_DomainParticipantFactory_get_instance	DDS_DomainParticipantFactory	<a href="#">DDS_DomainParticipantFactoryQos</a> <a href="#">OSAPI_LogProperty</a> <a href="#">OSAPI_SystemProperty</a>
DDS_DomainParticipantFactory_create_participant	DDS_DomainParticipant	<a href="#">DDS_DomainParticipantQos</a>
DDS_DomainParticipant_create_topic	DDS_Topic	Topic name length
DDS_Publisher_create_datawriter	DDS_DataWriter	<a href="#">DDS_DataWriterQos</a>
DDS_Subscriber_create_datareader	DDS_DataReader	<a href="#">DDS_DataReaderQos</a>
DDS_WaitSet_new	DDS_WaitSet	None
FooTypeSupport_register_type	TypeSupport	Type name length <i>Discovery plugin resource limits</i>

### 3.24.4 Discovery plugin resource limits

*Connext Micro* preallocates memory to store discovery information. If these limits are exceeded, *Connext Micro* discards the discovery information. *Connext Micro* includes two different discovery plugins:

1. Dynamic Participant Dynamic Endpoint (DPDE). Please refer to [DPDE](#) for details. DPDE is the simplest discovery plugin to use; however, its lack of ability to ignore DDS entities may cause *Connext Micro* to run out of resources if running in a non-deterministic environment. For this reason, DPDE is recommended when running in a deterministic environment OR when resource limits are set sufficiently high to accomodate all possible scenarios.
2. Dynamic Participant Static Endpoint (DPSE). Please refer to [DPSE](#) for details. DPSE trades off ease of use, higher memory usage, and non-deterministic behavior with more upfront configuraton, lower memory usage, and deterministic behavior.

The discovery process does not *only* store discovery information; it also *matches* local endpoints with discovered endpoints. While there may be sufficient resources needed to store a discovery message, the resource needed to perform matching may be exhausted. The DDS discovery process can easily exhaust the resource limits in *Connext Micro*.

The following resource limits are important to take into account:

Table 3.14: Resource Limits and Discovery

Resource Limit	Description
<code>DDS_DomainParticipantQos.resource_limits.remote_participant_allocation</code>	The maximum number of <i>DomainParticipants</i> that can be discovered. The <i>DomainParticipant</i> itself is never discovered.
<code>DDS_DomainParticipantQos.resource_limits.remote_reader_allocation</code>	The maximum number of remote <i>DataReaders</i> that can be discovered. Locally created <i>DataReaders</i> are not discovered via discovery plugin and should <i>not</i> count toward this limit.
<code>DDS_DomainParticipantQos.resource_limits.remote_writer_allocation</code>	The maximum number of remote <i>DataWriters</i> that can be discovered. Locally created <i>DataWriters</i> are not discovered via discovery plugin and should <i>not</i> count toward this limit.
<code>DDS_DomainParticipantQos.resource_limits.matching_writer_reader_pair_allocation</code>	The maximum number of RTPS sessions that can be created. An RTPS session is created between a <i>DataReader</i> and a matched <i>DataWriter</i> and between a <i>DataWriter</i> and a matched <i>DataReader</i> .
<code>DDS_DataWriterQos.writer_resource_limits.max_remote_readers</code>	The maximum number of <i>DataReaders</i> the <i>DataWriter</i> can match with. Note that both discovered <i>and</i> locally created <i>DataReaders</i> should count toward this resource limit.
<code>DDS_DataReaderQos.writer_resource_limits.max_remote_writers</code>	The maximum number of <i>DataWriters</i> the <i>DataReader</i> can match with. Note that both discovered <i>and</i> locally created <i>DataWriters</i> should count toward this resource limit.

### 3.24.5 Type support resource limits

*Connext Micro* supports two different types of code generation from IDL:

1. **Interpreted:** This is the default and requires the use of the `rti_me_ddsxtypes` library. However, it also allocates additional memory. You can enable interpreted type support by passing `interpreted 1` to `rtiddsgen`.
2. **Non-Interpreted:** The non-interpreted type support does not support the DDS X-Types specification. However, it uses significantly less memory. This is also the only type support supported by the CERT profile. You can enable the non-interpreted type support by passing `interpreted 0` to `rtiddsgen`.

# Chapter 4

## Platform Notes

### 4.1 Introduction

This section provides platform-specific instructions that you will need to build and run *RTI Connex Micro* applications.

For each supported operating system (OS), this section describes:

- Supported combinations of OS versions, CPUs, and compilers
- How to build your application, including:
  - Required *Connex Micro* and system libraries
  - Required compiler and linker flags
  - Details on how the *Connex Micro* libraries were built

To see a list of all supported platforms, refer to *Supported Platforms and Programming Languages*.

#### 4.1.1 Library types

*Connex Micro* provides precompiled binaries for supported architectures. This section explains the different library types and gives a general description of the binaries shipped by RTI.

In this section, the following terms are used:

- *toolchain* refers to the compiler, linker, and archiver for a specific CPU architecture, excluding dependencies in standard header files and libraries.
- *platform* refers to the hardware, BSPs, OS kernel, and C/C++ libraries that are not included in the toolchain (such as `libc`, `libc++`, and the network stack).

RTI builds two types of binaries for *Connex Micro*: integrated libraries and split libraries. RTI may include either or both types of binaries for a given target architecture.

## Integrated libraries

All *Connex Micro* target packages include a core library called `rti_me`. The `rti_me` library includes all the required basic functionality for *Connex Micro*.

The term “integrated library” refers to an `rti_me` library where all the OS integration and network stack integration is compiled directly into `rti_me`. This means that it is not possible to change how the OS and network integration has been written without recompiling the entire library. This is illustrated below:

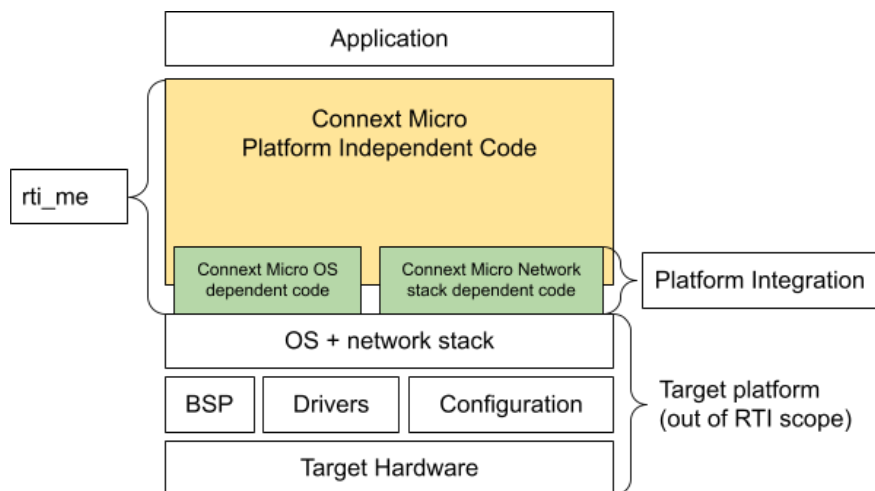


Figure 4.1: An overview of `rti_me` as an integrated library

---

**Note:** All binaries provided for *Connex Micro* version 4.0.1 and below are integrated libraries.

---

## Split libraries

In contrast to an integrated library, split libraries consist of a Platform Independent Library (PIL) and a Platform Support Library (PSL).

The PIL is an `rti_me` library that includes all functionality for *Connex Micro* except for platform integration code.

The PSL consists of two libraries that support OS integration and network stack integration:

- The OS Platform Support Library (`ospsl`): Contains the required OS support, such as mutex and semaphore support. This library is very limited in functionality.
- The Network Support Library (`netiops1`): Includes support for transports, such as UDPv4.

The `ospsl` and `netiops1` libraries are collectively referred to as the PSL (even though it is more than one library).

This is illustrated below:



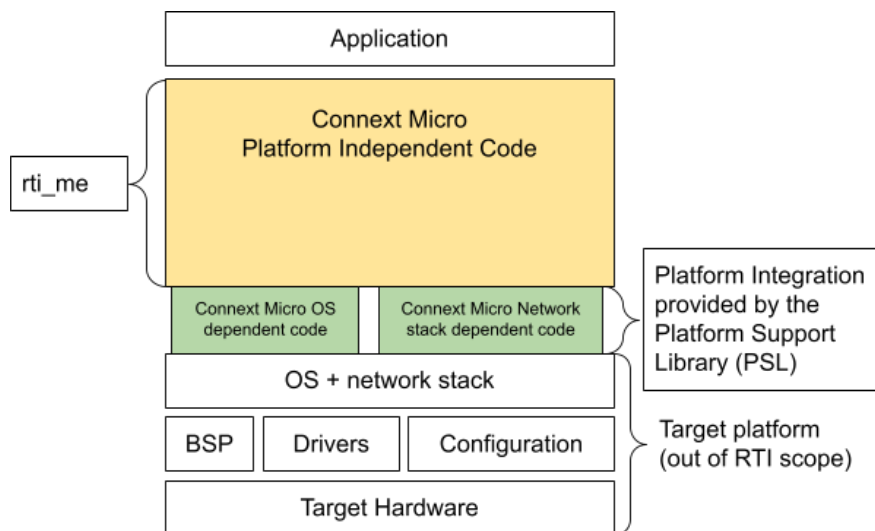


Figure 4.2: An overview of split libraries

The main benefit of split libraries is that different PSLs can be written for the same PIL without having to recompile the platform-independent code.

---

**Note:** The PIL is compiled without standard C header files and is only dependent on the toolchain. This is different from the integrated libraries, which are compiled with standard C header files.

The PSL is always compiled against the standard C header files, as well as other platform-dependent header files.

---

### 4.1.2 Library descriptions

The following libraries are included in the target bundle. Note that the names listed below do not include platform-specific prefixes or suffixes.

Depending on the target architecture, the library name is prefixed with lib and the library suffix also varies between target architectures, such as .so and .dylib.

The following naming conventions are also used:

- Static libraries have a z suffix.
- Shared libraries do not have an additional suffix.
- Debug libraries have a d suffix.
- Release libraries do not have an additional suffix.

For example:

- `rti_mezd` indicates a static debug library.
- `rti_me` indicates a dynamically linked release library.

Table 4.1: Target Bundle Libraries

Library Name	Description
<code>rti_me</code>	The core library, including the DDS C API.
<code>rti_me_discdpse</code>	The Dynamic Participant Static Endpoint (DPSE) plugin.
<code>rti_me_discdpde</code>	The Dynamic Participant Dynamic Endpoint (DPDE) plugin.
<code>rti_me_rhsm</code>	The Reader History plugin.
<code>rti_me_whsm</code>	The Writer History plugin.
<code>rti_me_netiosdm</code>	The Zero Copy v1 over shared memory transport library plugin.
<code>rti_me_netioshmem</code>	The Shared Memory Transport plugin.
<code>rti_me_appgen</code>	The Application Generation plugin.
<code>rti_me_cpp</code>	The C++ API.
<code>rti_me_ospsl</code>	The OS PSL.
<code>rti_me_netioysl</code>	The C NETIO PSL.
<code>rti_me_netioysl_cpp</code>	The C++ NETIO PSL library.
<code>rti_me_netiozcopy</code>	The Zero Copy v2 transport library plugin (not supported on all platforms).
<code>rti_me_ddsfilter</code>	The content filtering library.
<code>rti_me_ddsxtypes</code>	The X-Types library.
<code>rti_me_ddspsk</code>	The Lightweight Security Plugin PIL (installed separately, not supported on all platforms).
<code>rti_me_pskpsl</code>	The Lightweight Security Plugin Transform PSL (installed separately, not supported on all platforms).

### 4.1.3 Build profiles

You can optionally build *Connext Micro* with a CERT profile. This restricts *Connext Micro* to only include features that are available or planned for *Connext Cert*; for more information, see *Building Connext Micro with Compatibility for Connext Cert*.

Any architecture ending with CERT is built with the CERT profile enabled.

Some features are only available on specific platforms; see the footnotes in the table below.

Table 4.2: Features by Profile

Feature/Capability	Non-CERT Profile	CERT Profile
Dynamic Participant Discovery	✓	✓
Static Endpoint Discovery	✓	✓
Dynamic Endpoint Discovery	✓	
C++ API	✓	
Shared Memory Transport (SHMEM)	✓ <sup>1</sup>	
Zero Copy v1	✓ <sup>1</sup>	
Zero Copy v2	✓ <sup>12</sup>	✓ <sup>12</sup>
Micro Application Generator (MAG)	✓	
Content Filtering	✓	
Lightweight Security Plugin	✓ <sup>3</sup>	✓ <sup>3</sup>
X-Types	✓	

#### 4.1.4 Supported libraries by platform

The following table shows which *Connex Micro* libraries are supported on each platform (RTI architecture).

Table 4.3: Supported Libraries by Platform

Platform	RTI Architecture	Supported Libraries
Windows 10 x64	x86_64lePEvs2017	rti_me rti_me_whsm rti_me_rhsm rti_me_discdpse rti_me_discdpde rti_me_netiosdm rti_me_netioshmem rti_me_appgen rti_me_cpp rti_me_ddsxtypes rti_me_ddsfilter rti_me_ddspsk
	x86_64lePEvs2017CERT	rti_me rti_me_ddspsk
macOS 14 x64	x86_64leMachOclang15.0	rti_me rti_me_whsm rti_me_rhsm rti_me_discdpse rti_me_discdpde rti_me_netiosdm rti_me_netioshmem rti_me_appgen rti_me_cpp rti_me_ddsxtypes rti_me_ddsfilter
	x86_64leMa- chOclang15.0CERT	rti_me

continues on next page

<sup>1</sup> Not supported on FreeRTOS platforms.

<sup>2</sup> Not supported on Windows and macOS platforms.

<sup>3</sup> Not supported on macOS and FreeRTOS platforms.

Table 4.3 – continued from previous page

Platform	RTI Architecture	Supported Libraries
macOS 14 arm64	armv8leMachOclang15.0	rti_me rti_me_whsm rti_me_rhsm rti_me_discdpse rti_me_discdpde rti_me_netiosdm rti_me_netioshmem rti_me_appgen rti_me_cpp rti_me_ddsxtypes rti_me_ddsfilter
	armv8leMa- chOclang15.0CERT	rti_me
Ubuntu 22.04 x64	x86_64leElfgcc12.3.0	rti_me rti_me_whsm rti_me_rhsm rti_me_discdpse rti_me_discdpde rti_me_netiosdm rti_me_netioshmem rti_me_netiozcopy rti_me_appgen rti_me_cpp rti_me_ddsxtypes rti_me_ddsfilter rti_me_ddspsk
	x86_64leElfgcc12.3.0CERT	rti_me rti_me_netiozcopy rti_me_ddspsk
Ubuntu 18.04 ARMv8	armv8leElfgcc7.3.0	rti_me rti_me_whsm rti_me_rhsm rti_me_discdpse rti_me_discdpde rti_me_netiosdm rti_me_netioshmem rti_me_netiozcopy rti_me_appgen rti_me_cpp rti_me_ddsxtypes rti_me_ddsfilter rti_me_ddspsk

continues on next page

Table 4.3 – continued from previous page

Platform	RTI Architecture	Supported Libraries
	armv8leElfgcc7.3.0CERT	rti_me rti_me_netiozcopy rti_me_ddspsk
Yocto 5.15.96	armv8leElfgcc11.3.1	rti_me rti_me_whsm rti_me_rhsm rti_me_discdpse rti_me_discdpde rti_me_netiosdm rti_me_netioshmem rti_me_netiozcopy rti_me_appgen rti_me_cpp rti_me_ddsxtypes rti_me_ddsfilter rti_me_ddspsk
	armv8leElfgcc11.3.1CERT	rti_me rti_me_netiozcopy rti_me_ddspsk
QNX 7.1 ARMv8	armv8leElfqnx_qcc8.3.0	rti_me rti_me_whsm rti_me_rhsm rti_me_discdpse rti_me_discdpde rti_me_netiosdm rti_me_netioshmem rti_me_netiozcopy rti_me_appgen rti_me_cpp rti_me_ddsxtypes rti_me_ddsfilter rti_me_ddspsk
	armv8leElfqnx_qcc8.3.0CERT	rti_me rti_me_netiozcopy rti_me_ddspsk

continues on next page

Table 4.3 – continued from previous page

Platform	RTI Architecture	Supported Libraries
QOS 2.2.1 (QNX OS for Safety)	armv8leElfgcc8.3.0	rti_me rti_me_whsm rti_me_rhsm rti_me_discdpse rti_me_discdpde rti_me_netiosdm rti_me_netioshmem rti_me_netiozcopy rti_me_appgen rti_me_cpp rti_me_ddsxtypes rti_me_ddsfilter rti_me_ddspsk
	armv8leElfgcc8.3.0CERT	rti_me rti_me_netiozcopy rti_me_ddspsk
FreeRTOS 10.0.0 ARMv7E-M	armv7emleElfgcc10.3.1	rti_me rti_me_whsm rti_me_rhsm rti_me_discdpse rti_me_discdpde rti_me_appgen rti_me_cpp rti_me_ddsxtypes rti_me_ddsfilter
	armv7em-leElfgcc10.3.1CERT	rti_me

#### 4.1.5 Supported transports by platform

The following table shows which transports are supported on each architecture.

Table 4.4: Supported Transports by Platform

Platform	RTI Architecture	Intra	UDIPv4	SHMEM	Zero Copy v1	Zero Copy v2
Windows 10 x64	x86_64lePEvs2017	✓	✓	✓	✓	
	x86_64lePEvs2017CERT	✓	✓			
macOS 14 x64	x86_64leMachOclang15.0	✓	✓	✓	✓	
	x86_64leMachOclang15.0CERT	✓	✓			
macOS 14 arm64	armv8leMachOclang15.0	✓	✓	✓	✓	
	armv8leMachOclang15.0CERT	✓	✓			
Ubuntu 22.04 x64	x86_64leElfgcc12.3.0	✓	✓	✓	✓	✓
	x86_64leElfgcc12.3.0CERT	✓	✓			✓
Ubuntu 18.04 ARMv8	armv8leElfgcc7.3.0	✓	✓	✓	✓	✓
	armv8leElfgcc7.3.0CERT	✓	✓			✓
Auto Yocto Linux BSP 37.0	armv8leElfgcc11.3.1	✓	✓	✓	✓	✓
	armv8leElfgcc11.3.1CERT	✓	✓			✓
QNX 7.1 ARMv8	armv8leElfqnx_qcc8.3.0	✓	✓	✓	✓	✓
	armv8leElfqnx_qcc8.3.0CERT	✓	✓			✓
QOS 2.2.1 (QNX OS for Safety)	armv8leElfqcc8.3.0	✓	✓	✓	✓	✓
	armv8leElfqcc8.3.0CERT	✓	✓			✓
FreeR-TOS 10.4.6 ARMv7E-M	armv7emleElfgcc10.3.1	✓	✓			
	armv7emleElfgcc10.3.1CERT	✓	✓			

## 4.2 FreeRTOS Platforms

The following table shows the currently supported FreeRTOS platforms.

Table 4.5: Supported FreeRTOS Platforms

OS	Version	CPU	Network Stack	Toolchain	Architecture PIL	Architecture PSL
FreeRTOS	10.0.0	ARMv7E-M	LWiP 2.0.0	GCC 10.3.1	armv7emleElfgcc10.3.1 armv7em- leElfgcc10.3.1CERT	armv7em- leElfgcc10.3.1-FreeRTOS10.0 armv7em- leElfgcc10.3.1CERT-FreeRTOS10.0

### 4.2.1 Port overview

RTI ported *Connex Micro* to run on the FreeRTOS operating system with the lwIP protocol stack. This section contains some additional information on the hardware and software used.

RTI used S32G-VNP-RDB2 as the reference hardware. This reference kit has a S32G274A processor with 3x Arm Cortex-M7 LS pairs and 8MB system RAM. For a full description, please refer to the microcontroller documentation [here](#).

NXP provides a comprehensive software ecosystem for its S32G platform. For development on the M7 core, NXP offers the S32 Design Studio (S32DS)—a free, Eclipse-based integrated development environment (IDE) that supports the full S32G family. S32DS includes a GCC-based C/C++ compiler, GDB debugger integration, and tools for flashing and debugging the M7 core.

To simplify embedded software development, NXP provides the Real-Time Drivers (RTD) 4.0 package. RTD includes all necessary low-level drivers and middleware components to build applications for the M7 core on S32G platforms. It is fully compliant with AUTOSAR 4.x and integrates with FreeRTOS and other lightweight RTOS environments. RTI used the following versions of the different components:

- RTD version V4.0.0
- S32 Design Studio V3.6.0
- FreeRTOS Kernel V10.4.6
- lwIP version V2.0.0



## 4.2.2 How to configure lwIP and FreeRTOS

RTI ported *Connex Micro* to FreeRTOS using the following lwIP and FreeRTOS configurations. These can be tuned according to your needs by modifying the examples below. Details about how to configure these third-party components can also be found in the [FreeRTOS documentation](#) and [lwIP documentation](#).

lwIP

```
#ifndef LWIP_OPTS_H
#define LWIP_OPTS_H

/**
 * @page misra_violations MISRA-C:2012 violations
 *
 * @section [global]
 * Violates MISRA 2012 Advisory Rule 2.5, Global macro not referenced.
 * The global macro will be used in function call of the module.
 */

/* ----- Application APIs ----- */
#define NO_SYS                0
#define LWIP_RAW               1
#define LWIP_SOCKET           (NO_SYS==0)
#define LWIP_NETCONN          (NO_SYS==0)

/* ----- ARP options ----- */
#define LWIP_ARP               1
#define ARP_TABLE_SIZE        10
#define ARP_QUEUEING           1

/* ----- IP options ----- */
#define LWIP_IPV6               0

/* ----- DHCP options ----- */
/* Define LWIP_DHCP to 1 if you want DHCP configuration of interfaces */
#define LWIP_DHCP               0
/* 1 if you want to do an ARP check on the offered address (recommended) */
#define DHCP_DOES_ARP_CHECK    (LWIP_DHCP)
```

(continues on next page)

(continued from previous page)

```

/* ----- AUTOIP options ----- */
#define LWIP_AUTOIP                0
#define LWIP_DHCP_AUTOIP_COOP      (LWIP_DHCP && LWIP_AUTOIP)
#define LWIP_DHCP_AUTOIP_COOP_TRIES 3

/* Define IP_FORWARD to 1 if you wish to have the ability to forward
   IP packets across network interfaces. If you are going to run lwIP
   on a device with only one network interface, define this to 0. */
#define IP_FORWARD                  0

/* IP reassembly and segmentation. These are orthogonal even
   * if they both deal with IP fragments */
#define IP_REASSEMBLY              1
#define IP_REASS_MAX_PBUFS         10
#define MEMP_NUM_REASSDATA         10
#define IP_FRAG                    1

/* ----- IP options ----- */
#define LWIP_IPV4                   1
#define ICMP_TTL                   255
#define LWIP_ICMP                  1
#define LWIP_IGMP                  1

/* ----- TCP options ----- */
#define LWIP_TCP                    1
#define TCP_TTL                    255

/* Controls if TCP should queue segments that arrive out of
   order. Define to 0 if your device is low on memory. */
#define TCP_QUEUE_OOSEQ            0

/* TCP Maximum segment size. */
#define TCP_MSS                    1536

/* TCP sender buffer space (bytes). */

```

(continues on next page)

(continued from previous page)

```

#define TCP_SND_BUF          TCP_MSS * 4

/* TCP sender buffer space (pbufs). This must be at least = 2 *
   TCP_SND_BUF/TCP_MSS for things to work. */
#define TCP_SND_QUEUELEN    (4 * TCP_SND_BUF/TCP_MSS)

/* TCP writable space (bytes). This must be less than or equal
   to TCP_SND_BUF. It is the amount of space which must be
   available in the tcp snd_buf for select to return writable */
#define TCP_SNDLOWAT        (TCP_SND_BUF/2)

/* TCP receive window. */
#define TCP_WND              4096

/* Maximum number of retransmissions of data segments. */
#define TCP_MAXRTX           2

/* Maximum number of retransmissions of SYN segments. */
#define TCP_SYNMAXRTX        2

/* ----- UDP options ----- */
#define LWIP_UDP              1
#define LWIP_UDPLITE          LWIP_UDP
#define UDP_TTL               255

/* ----- Memory options ----- */
/* MEM_ALIGNMENT: should be set to the alignment of the CPU for which
   lwIP is compiled. 4 byte alignment -> define MEM_ALIGNMENT to 4, 2
   byte alignment -> define MEM_ALIGNMENT to 2. */
/* MSVC port: intel processors do not need 4-byte alignment,
   but are faster that way! */

/*!
 * @brief Custom memcpy function used to optimize and improve the lwip copy speed process.
 *
 * This function copies len characters from memory area src to memory area dest using 64 bit alignment.

```

(continues on next page)

(continued from previous page)

```

*/
void memcpy_64(uint64_t *dst, uint64_t * src, unsigned int len);

/*!
 * @brief Custom memcpy function used to optimize and improve the lwip copy speed process.
 *
 * This function copies len characters from memory area src to memory area dest.
 */
void memcpy_custom(void *dst, const void * src, unsigned int len);

#define MEMCPY(dst,src,len)      memcpy(dst,src,len)

/* Memory alignment must match your architecture - for ARM Cortex M7, 8-byte alignment is required */
#define MEM_ALIGNMENT            8

/* MEM_SIZE: the size of the heap memory. If the application will send
a lot of data that needs to be copied, this should be set high. */
#define MEM_SIZE                 128*1024

/* MEMP_NUM_PBUF: the number of memp struct pbufs. If the application
sends a lot of data out of ROM (or other static memory), this
should be set high. */
#define MEMP_NUM_PBUF           8

/* MEMP_NUM_RAW_PCB: the number of UDP protocol control blocks. One
per active RAW connection. */
#define MEMP_NUM_RAW_PCB        4

/* MEMP_NUM_UDP_PCB: the number of UDP protocol control blocks. One
per active UDP connection. */
#define MEMP_NUM_UDP_PCB        4

/* MEMP_NUM_TCP_PCB: the number of simultaneously active TCP connections. */
#define MEMP_NUM_TCP_PCB        4

/* MEMP_NUM_TCP_PCB_LISTEN: the number of listening TCP connections. */

```

(continues on next page)

(continued from previous page)

```

#define MEMP_NUM_TCP_PCB_LISTEN      4

/* MEMP_NUM_TCP_SEG: the number of simultaneously queued TCP segments. */
#define MEMP_NUM_TCP_SEG              64

/* MEMP_NUM_SYS_TIMEOUT: the number of simultaneously active timeouts. */
#define MEMP_NUM_SYS_TIMEOUT          15

/* The following four are used only with the sequential API and can be
   set to 0 if the application only will use the raw API. */
/* MEMP_NUM_NETBUF: the number of struct netbufs. */
#define MEMP_NUM_NETBUF               16

/* MEMP_NUM_NETCONN: the number of struct netconns. */
#define MEMP_NUM_NETCONN              32

/* MEMP_NUM_TCPIP_MSG_*: the number of struct tcpip_msg, which is used
   for sequential API communication and incoming packets. Used in
   src/api/tcpip.c. */
#define MEMP_NUM_TCPIP_MSG_API        40 /* Increased from 20 */
#define MEMP_NUM_TCPIP_MSG_INPKT      30 /* Increased from 20 */

/* MEM_USE_POOLS==1: Use an alternative to malloc()
   To use this, MEMP_USE_CUSTOM_POOLS also has to be enabled. */
#define MEM_USE_POOLS                 0
#define MEMP_USE_CUSTOM_POOLS         0

/* ----- Pbuf options ----- */
/* PBUF_POOL_SIZE: the number of buffers in the pbuf pool. */
#define PBUF_POOL_SIZE                32
#define PBUF_POOL_BUFSIZE             1536

/** SYS_LIGHTWEIGHT_PROT
 * define SYS_LIGHTWEIGHT_PROT in lwipopts.h if you want inter-task protection
 * for certain critical regions during buffer allocation, deallocation and memory
 * allocation and deallocation. */

```

(continues on next page)

(continued from previous page)

```

#define SYS_LIGHTWEIGHT_PROT      (NO_SYS==0)

/* ----- Statistics options ----- */
#define LWIP_STATS                  1
#define LWIP_STATS_DISPLAY         1

#if LWIP_STATS
#define LINK_STATS                  1
#define IP_STATS                    1
#define ICMP_STATS                  1
#define IGMP_STATS                  1
#define IPFRAG_STATS               1
#define UDP_STATS                   1
#define TCP_STATS                   1
#define MEM_STATS                   1
#define MEMP_STATS                  1
#define PBUF_STATS                  1
#define SYS_STATS                   1
#endif /* LWIP_STATS */

/* ----- PPP options ----- */
#define PPP_SUPPORT                  0

/* ----- Other options ----- */

/* (LWIP_ETHERNET == 1) Enable ETHERNET support even though ARP might be disabled */
#define LWIP_ETHERNET               1

#define TCPIP_MBOX_SIZE             24

#define DEFAULT_UDP_RECVMBOX_SIZE   1500
#define DEFAULT_TCP_RECVMBOX_SIZE   1500
#define DEFAULT_RAW_RECVMBOX_SIZE   1500
#define DEFAULT_ACCEPTMBOX_SIZE     1500
#define LWIP_NETIF_TX_SINGLE_PBUF   1

```

(continues on next page)

(continued from previous page)

```

#define LWIP_SUPPORT_CUSTOM_PBUF      1

#define LWIP_SNMP                      LWIP_UDP
#define MIB2_STATS                     LWIP_SNMP

#define LWIP_DNS                       LWIP_UDP
#define LWIP_MDNS_RESPONDER           LWIP_UDP

#define LWIP_NUM_NETIF_CLIENT_DATA     (LWIP_MDNS_RESPONDER)

#define LWIP_HAVE_LOOPIF               0
#define LWIP_NETIF_LOOPBACK            0

#define TCP_LISTEN_BACKLOG              1

#define LWIP_COMPAT_SOCKETS            1
#define LWIP_SO_RCVTIMEO               1
#define LWIP_SO_RCVBUF                 1

#define LWIP_TCPIP_CORE_LOCKING         1

#define LWIP_NETIF_LINK_CALLBACK        1
#define LWIP_NETIF_STATUS_CALLBACK      1

#define LWIP_NETCONN_SEM_PER_THREAD    0

#define LWIP_SOCKET_SET_ERRNO          1

#define LWIP_NETIF_HOSTNAME             1

/* ----- TIMEVAL options ----- */
#define LWIP_TIMEVAL_PRIVATE            0

/* ----- Thread options ----- */
#define TCPIP_THREAD_NAME               "tcpip_thread"
#define TCPIP_THREAD_STACKSIZE          1024*64

```

(continues on next page)

(continued from previous page)

```

#define TCPIP_THREAD_PRIO          4
#define DEFAULT_THREAD_STACKSIZE  4096
#define DEFAULT_THREAD_PRIO       3

/**
 * LWIP_SO_RCVBUF==1: Enable SO_RCVBUF processing.
 */
#define LWIP_SO_RCVBUF             1

#define LWIP_SOCKET_SELECT         0

#endif /* LWIP_OPTS_H */

```

## FreeRTOS

```

#ifndef FREERTOS_CONFIG_H
#define FREERTOS_CONFIG_H

/*-----
 * Application specific definitions.
 *
 * These definitions should be adjusted for your particular hardware and
 * application requirements.
 *
 * THESE PARAMETERS ARE DESCRIBED WITHIN THE 'CONFIGURATION' SECTION OF THE
 * FreeRTOS API DOCUMENTATION AVAILABLE ON THE FreeRTOS.org WEB SITE.
 *
 * See http://www.freertos.org/a00110.html.
 *-----*/

#define configUSE_PREEMPTION          1
#define configCPU_CLOCK_HZ           ( 400000000UL )
#define configTICK_RATE_HZ           ((TickType_t) 1000 )
#define configMAX_PRIORITIES         10
#define configMINIMAL_STACK_SIZE     RTI_configMINIMAL_STACK_SIZE
#define configMAX_TASK_NAME_LEN      16

```

(continues on next page)



(continued from previous page)

```

#define configUSE_16_BIT_TICKS                0
#define configIDLE_SHOULD_YIELD               1
#define configUSE_PREEMPTION                  1
#define configNUM_THREAD_LOCAL_STORAGE_POINTERS 2
#define configUSE_PORT_OPTIMISED_TASK_SELECTION 0
#define configUSE_TASK_NOTIFICATIONS          1
#define configUSE_TIME_SLICING                0
#define configUSE_NEWLIB_REENTRANT            0
#define configENABLE_BACKWARD_COMPATIBILITY    1
#define configUSE_POSIX_ERRNO                 0
#define configUSE_APPLICATION_TASK_TAG        0
#define configRECORD_STACK_HIGH_ADDRESS        1

/* Definition assert() function. */
#define configASSERT(x)                        if((x)==0) { taskDISABLE_INTERRUPTS(); __asm__ volatile ("bkpt #0");  
↪for(;;); }

/* The highest interrupt priority that can be used by any interrupt service  
routine that makes calls to interrupt safe FreeRTOS API functions. DO NOT CALL  
INTERRUPT SAFE FREERTOS API FUNCTIONS FROM ANY INTERRUPT THAT HAS A HIGHER  
PRIORITY THAN THIS! (higher priorities are lower numeric values. */
#define configLIBRARY_MAX_SYSCALL_INTERRUPT_PRIORITY ( 2 )

/* Memory allocation related definitions. */
#define configSUPPORT_STATIC_ALLOCATION         0
#define configSUPPORT_DYNAMIC_ALLOCATION        1
#define configTOTAL_HEAP_SIZE                  (( size_t ) ( __heap_size__ ) )
#define configAPPLICATION_ALLOCATED_HEAP        4

/* Hook function related definitions. */
#define configUSE_IDLE_HOOK                     0
#define configUSE_TICK_HOOK                     0
#define configUSE_MALLOC_FAILED_HOOK            1
#define configCHECK_FOR_STACK_OVERFLOW          1
#define configUSE_DAEMON_TASK_STARTUP_HOOK      0

```

(continues on next page)

(continued from previous page)

```

/* Run time and task stats gathering related definitions. */
#define configGENERATE_RUN_TIME_STATS      1
#define portCONFIGURE_TIMER_FOR_RUN_TIME_STATS()
#define portGET_RUN_TIME_COUNTER_VALUE()   xTaskGetTickCount()
#define configUSE_TRACE_FACILITY          1
#define configUSE_STATS_FORMATTING_FUNCTIONS 1

/* Co-routine related definitions. */
#define configUSE_CO_ROUTINES              0
#define configMAX_CO_ROUTINE_PRIORITIES    2

/* SEMAPHORES and MUTEXS */
#define configUSE_MUTEXES                  1
#define configUSE_RECURSIVE_MUTEXES        1
#define configUSE_COUNTING_SEMAPHORES      1

/* Software timer related definitions. */
/* Software timer related definitions. */
#define configUSE_TIMERS                    1
#define configTIMER_TASK_PRIORITY          6
#define configTIMER_QUEUE_LENGTH          16
#define configTIMER_TASK_STACK_DEPTH       (1024*2)

/* Tickless Idle Mode */
#define configUSE_TICKLESS_IDLE            0

/* QUEUE */
#define configQUEUE_REGISTRY_SIZE          8
#define configUSE_QUEUE_SETS              0

/* Optional functions - most linkers will remove unused functions anyway. */
#define INCLUDE_vTaskPrioritySet            1
#define INCLUDE_uTaskPriorityGet            1
#define INCLUDE_vTaskDelete                1
#define INCLUDE_vTaskSuspend               1
#define INCLUDE_vTaskDelayUntil            1

```

(continues on next page)

(continued from previous page)

```

#define INCLUDE_vTaskDelay 1
#define INCLUDE_xTaskGetSchedulerState 1
#define INCLUDE_xTaskGetCurrentTaskHandle 1
#define INCLUDE_uxTaskGetStackHighWaterMark 1
#define INCLUDE_uxTaskGetStackHighWaterMark2 0
#define INCLUDE_xTaskGetIdleTaskHandle 0
#define INCLUDE_eTaskGetState 1
#define INCLUDE_xEventGroupSetBitFromISR 1
#define INCLUDE_xTimerPendFunctionCall 1
#define INCLUDE_xTaskAbortDelay 1
#define INCLUDE_xTaskGetHandle 1
#define INCLUDE_xTaskResumeFromISR 1
#define INCLUDE_xQueueGetMutexHolder 1

/* Run time stats gathering definitions. */
#ifdef __GNUC__
    /* The #ifdef just prevents this C specific syntax from being included in
    assembly files. */
    void vMainConfigureTimerForRunTimeStats( void );
    unsigned long ulMainGetRunTimeCounterValue( void );
#endif

/* Cortex-M specific definitions. */
#ifdef __NVIC_PRIO_BITS
    #define configPRIO_BITS __NVIC_PRIO_BITS
#else
    #define configPRIO_BITS 4
#endif

/* The lowest interrupt priority that can be used in a call to a "set priority"
function. */
#define configLIBRARY_LOWEST_INTERRUPT_PRIORITY (0x0F)

/* Interrupt priorities used by the kernel port layer itself. These are generic
to all Cortex-M ports, and do not rely on any particular library functions. */
#ifndef configKERNEL_INTERRUPT_PRIORITY

```

(continues on next page)

(continued from previous page)

```

#define configKERNEL_INTERRUPT_PRIORITY      (configLIBRARY_LOWEST_INTERRUPT_PRIORITY << (8-configPRIO_BITS))
#endif

/* !!!! configMAX_SYSCALL_INTERRUPT_PRIORITY must not be set to zero !!!!
See http://www.FreeRTOS.org/RTOS-Cortex-M3-M4.html. */
#ifndef configMAX_SYSCALL_INTERRUPT_PRIORITY
#define configMAX_SYSCALL_INTERRUPT_PRIORITY      (configLIBRARY_MAX_SYSCALL_INTERRUPT_PRIORITY << (8-configPRIO_BITS))
#endif

/* Definitions that map the FreeRTOS port interrupt handlers to their CMSIS
standard names. */
#define vPortSVCHandler                      SVC_Handler
#define xPortPendSVHandler                  PendSV_Handler
#define xPortSysTickHandler                 SysTick_Handler

#endif /* FREERTOS_CONFIG_H */

```

### 4.2.3 How the PIL was built for FreeRTOS

This section describes how RTI built the Platform Independent Library (PIL) for FreeRTOS.

The following table shows the compiler flags RTI used to create the PIL for FreeRTOS platforms:

Table 4.6: PIL Compiler Flags for FreeRTOS Platforms

Architecture PIL	Library Format	Compiler Flags Used by RTI
armv7emleElfgcc10.3.1	Static Release	-specs=nosys.specs -specs=nano.specs -DOSAPI_CC_DEF_H=osapi/osapi_cc_gcc.h -DOSPSL_OS_DEF_H=rti_me_psl/ospsl/ospsl_os_stubs.h -DRTIME_TARGET_NAME="armv7emleElfgcc10.3.1" -DRTI_PIL=1 -D__freertos__ -std=c99 -nostdinc -fstrict-aliasing -fsigned-char -O -mcpu=cortex-m7 -mthumb -mfpv5-sp-d16 -mfloat-abi=hard -ffunction-sections -fdata-sections -fno-short-enums -DNDEBUG
	Static Debug	-specs=nosys.specs -specs=nano.specs -DOSAPI_CC_DEF_H=osapi/osapi_cc_gcc.h -DOSPSL_OS_DEF_H=rti_me_psl/ospsl/ospsl_os_stubs.h -DRTIME_TARGET_NAME="armv7emleElfgcc10.3.1" -DRTI_PIL=1 -D__freertos__ -std=c99 -nostdinc -fstrict-aliasing -fsigned-char -O0 -g -mcpu=cortex-m7 -mthumb -mfpv5-sp-d16 -mfloat-abi=hard -ffunction-sections -fdata-sections -fno-short-enums
armv7emleElfgcc10.3.1CERT	Static Release	-specs=nosys.specs -specs=nano.specs -DOSAPI_CC_DEF_H=osapi/osapi_cc_gcc.h -DOSPSL_OS_DEF_H=rti_me_psl/ospsl/ospsl_os_stubs.h -DRTIME_TARGET_NAME="armv7emleElfgcc10.3.1CERT" -DRTI_CERT -DRTI_PIL=1 -D__freertos__ -std=c99 -nostdinc -fstrict-aliasing -fsigned-char -O -mcpu=cortex-m7 -mthumb -mfpv5-sp-d16 -mfloat-abi=hard -ffunction-sections -fdata-sections -fno-short-enums -DNDEBUG

continues on next page

Table 4.6 – continued from previous page

Architecture PIL	Library Format	Compiler Flags Used by RTI
	Static Debug	<code>-specs=nosys.specs</code> <code>-specs=nano.specs</code> <code>-DOSAPI_CC_DEF_H=osapi/osapi_cc_gcc.h</code> <code>-DOSPSL_OS_DEF_H=rti_me_psl/osppl/osppl_os_stubs.h</code> <code>-DRTIME_TARGET_NAME="armv7emleElfgcc10.3.1CERT"</code> <code>-DRTI_CERT</code> <code>-DRTI_PIL=1</code> <code>-D__freertos__</code> <code>-std=c99</code> <code>-nostdinc</code> <code>-fstrict-aliasing</code> <code>-fsigned-char</code> <code>-O0</code> <code>-g</code> <code>-mcpu=cortex-m7</code> <code>-mthumb</code> <code>-mfpv5-sp-d16</code> <code>-mfloat-abi=hard</code> <code>-ffunction-sections</code> <code>-fdata-sections</code> <code>-fno-short-enums</code>

4.2.4 Building the PSL from source for FreeRTOS platforms

This section describes how to build your own PSL for FreeRTOS.

*Connex Micro* includes support to compile Platform Support Libraries (PSL) for FreeRTOS using [CMake](#). Refer to *Setting up the build environment* before continuing with the following instructions.

1. Make sure CMake is in the path.
2. Define the following environment variables:
  - `CONFIG_PATH`: Path to where the `FreeRTOSConfig.h` and `lwipopts.h` files are located.
  - `FREERTOS_PATH`: Path to the FreeRTOS source code and header files.
  - `LWIP_PATH`: Path to the lwIP source code and header files.
  - `LWIP_PORTS_PATH`: Path to the lwIP ports folder that contains `cc.h`.
  - `RTD_PATH`: Path to the folder that contains the eclipse plugins of RTD.
  - `PATH`: Update your path with the location of the C and C++ compiler. By default, `arm-none-eabi-gcc` and `arm-none-eabi-g++` are used as C and C++ compilers.
3. Enter the following command:

```
RTIMEHOME/resource/scripts/rtime-make --target armv7emleElfgcc10.3.1-FreeRTOS10.0 \
-G "Unix Makefiles" --build
```

**Note:** `rtime-make` uses the architecture specified with `--target` to determine a few settings needed by *Connex Micro*. Please refer to *Preparing to build* for details.

4. The *Connex Micro* PSL is available in:

```
RTIMEHOME/lib/armv7emleElfgcc10.3.1-FreeRTOS10.0
```

#### 4.2.5 Building FreeRTOS applications with Connex Micro

This section describes how RTI built the Platform Support Library (PSL) for FreeRTOS platforms. You must build applications with compatible flags to the PIL and PSL in order to operate with *Connex Micro*. The PSL must also be binary compatible with the PIL. Applications must not specify the `RTI_PSL` or `RTI_PIL` preprocessor definitions.

The following table shows the compiler flags and required options that RTI used to build the PSL for FreeRTOS platforms. When you build the PSL with `rtime-make`, the `--target` argument automatically adds all the necessary flags for the specified architecture.

Table 4.7: PSL Compiler Flags for FreeRTOS Platforms

Architecture PSL	Library Format	Compiler Flags Used by RTI
armv7emleElfgcc10.3.1-FreeRTOS10.0	Static Release	<pre>-specs=nosys.specs      -specs=nano.specs      -DNDEBUG -DOSAPI_CC_DEF_H=osapi/osapi_cc_gcc.h -DOSPSL_OS_DEF_H=rti_me_psl/ospsl/ospsl_os_freertos.h -DRTIME_TARGET_NAME="armv7emleElfgcc10.3.1-FreeRTOS10.0" -DRTI_PSL=1              -DUSING_RTD=1 -D__TIMEVAL_DEFINED=0    -D__freertos__      -std=c99 -fstrict-aliasing -fsigned-char -O -mcpu=cortex-m7 -mthumb -mfpv5-sp-d16 -mfloat-abi=hard -ffunction-sections -fdata-sections -fno-short-enums -DNDEBUG</pre>

continues on next page

Table 4.7 – continued from previous page

Architecture PSL	Library Format	Compiler Flags Used by RTI
	Static Debug	-specs=nosys.specs -specs=nano.specs -DOSAPI_CC_DEF_H=osapi/osapi_cc_gcc.h -DOSPSL_OS_DEF_H=rti_me_psl/osppl/osppl_os_freertos.h -DRTIME_TARGET_NAME="armv7emleElfgcc10.3.1-FreeRTOS10.0" -DRTI_PSL=1 -DUSING_RTD=1 -D_TIMEVAL_DEFINED=0 -D__freertos__ -std=c99 -fstrict-aliasing -fsigned-char -O0 -g -mcpu=cortex-m7 -mthumb -mfpv5-sp-d16 -mfloat-abi=hard -ffunction-sections -fdata-sections -fno-short-enums
armv7em-leElfgcc10.3.1CERT-FreeRTOS10.0	Static Release	-specs=nosys.specs -specs=nano.specs -DNDEBUG -DOSAPI_CC_DEF_H=osapi/osapi_cc_gcc.h -DOSPSL_OS_DEF_H=rti_me_psl/osppl/osppl_os_freertos.h -DRTIME_TARGET_NAME="armv7emleElfgcc10.3.1CERT-FreeRTOS10.0" -DRTI_CERT -DRTI_PSL=1 -DUSING_RTD=1 -D_TIMEVAL_DEFINED=0 -D__freertos__ -std=c99 -fstrict-aliasing -fsigned-char -O -g -mcpu=cortex-m7 -mthumb -mfpv5-sp-d16 -mfloat-abi=hard -ffunction-sections -fdata-sections -fno-short-enums -DNDEBUG
	Static Debug	-specs=nosys.specs -specs=nano.specs -DOSAPI_CC_DEF_H=osapi/osapi_cc_gcc.h -DOSPSL_OS_DEF_H=rti_me_psl/osppl/osppl_os_freertos.h -DRTIME_TARGET_NAME="armv7emleElfgcc10.3.1CERT-FreeRTOS10.0" -DRTI_CERT -DRTI_PSL=1 -DUSING_RTD=1 -D_TIMEVAL_DEFINED=0 -D__freertos__ -std=c99 -fstrict-aliasing -fsigned-char -O0 -g -mcpu=cortex-m7 -mthumb -mfpv5-sp-d16 -mfloat-abi=hard -ffunction-sections -fdata-sections -fno-short-enums



### 4.2.6 System tick rollovers

The [OMG standard](#) does not specify how an implementation of DDS should handle counter rollover. By default, *Connex Micro* does not check for rollover of the system tick count on FreeRTOS platforms.

The [OSAPI\\_SystemFreeRTOS\\_get\\_time](#) function uses the system tick count to calculate the time in milliseconds since the system started. It does so by multiplying the number of ticks since the system started by the FreeRTOS-defined constant `portTICK_RATE_MS`. *Connex Micro* stores the result in an unsigned 32-bit variable.

Two conditions can result from this calculation: the tick count can rollover, and the resulting calculation can be greater than the size of an `unsigned int` ( $2^{32}$ ). This can cause an overflow, which would result in time appearing to go backwards.

If you need to change or mitigate this behavior, you can alter the source code for the reference implementation of [OSAPI\\_SystemFreeRTOS\\_get\\_time](#) (or provide your own implementation) and rebuild the PSL.

## 4.3 Linux Platforms

The following table shows the currently supported Linux platforms.

Table 4.8: Supported Linux Platforms

OS	Version	CPU	Network Stack	Toolchain	Architecture PIL	Architecture PSL
Ubuntu	22.04 LTS	x64	OS Default	GCC 12.3.0	x86_64leElfgcc12.3.0 x86_64leElfgcc12.3.0CERT	x86_64leElfgcc12.3.0-Linux5 x86_64leElfgcc12.3.0CERT-Linux5
Ubuntu	18.04 LTS	ARMv8 (64-bit)	OS Default	GCC 7.3.0	armv8leElfgcc7.3.0 armv8leElfgcc7.3.0CERT	armv8leElfgcc7.3.0-Linux4 armv8leElfgcc7.3.0CERT-Linux4
Yocto	5.15.96	ARMv8-A (64-bit)	OS Default	GCC 11.3.1	armv8aleElfgcc11.3.1 armv8aleElfgcc11.3.1CERT	armv8aleElfgcc11.3.1-Linux5 armv8aleElfgcc11.3.1CERT-Linux5

### 4.3.1 How the PIL was built for Linux platforms

This section describes how RTI built the Platform Independent Library (PIL) for Linux.

The following table shows the compiler flags RTI used to create the PIL for Linux platforms:

Table 4.9: PIL Compiler Flags for Linux Platforms

Architecture	Library Format	Compiler Flags Used by RTI
x86_64leElfgcc12.3.0	Static Release	<b>C Flags:</b> -fsigned-char -O2 -nostdinc -fstrict-aliasing -fPIC -DNDEBUG <b>C++ Flags:</b> -std=c++11 -fsigned-char -O2 -nostdinc -fstrict-aliasing -fPIC -DNDEBUG
	Shared Release	<b>C Flags:</b> -fsigned-char -O2 -nostdinc -fstrict-aliasing -fPIC -DNDEBUG <b>C++ Flags:</b> -std=c++11 -fsigned-char -O2 -nostdinc -fstrict-aliasing -fPIC -DNDEBUG
	Static Debug	<b>C Flags:</b> -fsigned-char -nostdinc -fstrict-aliasing -g -fPIC <b>C++ Flags:</b> -std=c++11 -fsigned-char -nostdinc -fstrict-aliasing -g -fPIC

continues on next page

Table 4.9 – continued from previous page

Architecture	Library Format	Compiler Flags Used by RTI
	Shared Debug	<b>C Flags:</b> -fsigned-char -nostdinc -fstrict-aliasing -g -fPIC <b>C++ Flags:</b> -std=c++11 -fsigned-char -nostdinc -fstrict-aliasing -g -fPIC
x86_64leElfgcc12.3.0CERT	Static Release	<b>C Flags:</b> -fsigned-char -O2 -nostdinc -fstrict-aliasing -fPIC -DNDEBUG
	Static Debug	<b>C Flags:</b> -fsigned-char -nostdinc -fstrict-aliasing -g -fPIC
armv8leElfgcc7.3.0	Static Release	<b>C Flags:</b> -fsigned-char -O2 -nostdinc -fstrict-aliasing -fPIC -DNDEBUG <b>C++ Flags:</b> -std=c++11 -fsigned-char -nostdinc -O2 -fstrict-aliasing -fPIC -DNDEBUG
	Shared Release	<b>C Flags:</b> -fsigned-char -O2 -nostdinc -fstrict-aliasing -fPIC -DNDEBUG <b>C++ Flags:</b> -std=c++11 -fsigned-char -nostdinc -O2 -fstrict-aliasing -fPIC -DNDEBUG
	Static Debug	<b>C Flags:</b> -fsigned-char -nostdinc -fstrict-aliasing -g -fPIC <b>C++ Flags:</b> -std=c++11 -fsigned-char -nostdinc -fstrict-aliasing -g -fPIC
	Shared Debug	<b>C Flags:</b> -fsigned-char -nostdinc -fstrict-aliasing -g -fPIC <b>C++ Flags:</b> -std=c++11 -fsigned-char -nostdinc -fstrict-aliasing -g -fPIC
armv8leElfgcc7.3.0CERT	Static Release	<b>C Flags:</b> -fsigned-char -O2 -nostdinc -fstrict-aliasing -fPIC -DNDEBUG
	Static Debug	<b>C Flags:</b> -fsigned-char -nostdinc -fstrict-aliasing -g -fPIC
armv8leElfgcc11.3.1	Static Release	<b>C Flags:</b> -nostdinc -fstrict-aliasing -O -DNDEBUG -fPIC -DNDEBUG <b>C++ Flags:</b> -O -DNDEBUG -fPIC -DNDEBUG
	Shared Release	<b>C Flags:</b> -nostdinc -fstrict-aliasing -O -DNDEBUG -fPIC -DNDEBUG <b>C++ Flags:</b> -O -DNDEBUG -fPIC -DNDEBUG
	Static Debug	<b>C Flags:</b> -nostdinc -fstrict-aliasing -fsigned-char -O0 -g -fPIC <b>C++ Flags:</b> -O0 -g -fPIC
	Shared Debug	<b>C Flags:</b> -nostdinc -fstrict-aliasing -fsigned-char -O0 -g -fPIC <b>C++ Flags:</b> -O0 -g -fPIC
armv8leElfgcc11.3.1CERT	Static Release	<b>C Flags:</b> -nostdinc -fstrict-aliasing -fsigned-char -O -DNDEBUG

continues on next page

Table 4.9 – continued from previous page

Architecture	Library Format	Compiler Flags Used by RTI
	Static Debug	<b>C Flags:</b> -fsigned-char -nostdinc -fstrict-aliasing -O0 -g
armv8aleElfgcc11.3.1	Static Release	<b>C Flags:</b> -nostdinc -fstrict-aliasing -fsigned-char -O -DNDEBUG -fPIC -DNDEBUG <b>C++ Flags:</b> -O -DNDEBUG -fPIC
	Shared Release	<b>C Flags:</b> -nostdinc -fstrict-aliasing -fsigned-char -O -DNDEBUG -fPIC -DNDEBUG <b>C++ Flags:</b> -O -DNDEBUG -fPIC
	Static Debug	<b>C Flags:</b> -std=c99 -nostdinc -fstrict-aliasing -fsigned-char -O0 -g -fPIC <b>C++ Flags:</b> -O0 -g -fPIC
	Shared Debug	<b>C Flags:</b> -std=c99 -nostdinc -fstrict-aliasing -fsigned-char -O0 -g -fPIC <b>C++ Flags:</b> -O0 -g -fPIC
armv8aleElfgcc11.3.1CERT	Static Release	<b>C Flags:</b> -nostdinc -fstrict-aliasing -fsigned-char -O -DNDEBUG
	Static Debug	<b>C Flags:</b> -fsigned-char -nostdinc -fstrict-aliasing -O0 -g

**Warning:** The *RTI Connex Micro* platform independent libraries are built without the standard C header-files. However, *Connex Micro* makes one direct call to the C library API `qsort`. In addition, GCC may insert direct calls to GLIBC functions and other required functions, such as default C++ constructors and destructors. For this reason, it is necessary to use a GCC version that is compatible with the GCC version used to build the platform independent libraries, or provide a C library with an implementation of the required functions. Future versions of *RTI Connex Micro* will remove these dependencies.

### 4.3.2 Building the PSL from source for Linux platforms

Refer to *Building the PSL* for instructions on how to build your own Platform Support Library (PSL) for Linux platforms.

The following table shows the compiler flags and required options that RTI used to build the PSL for Linux platforms. When you build the PSL with `rtime-make`, the `--target` argument automatically adds all the necessary flags for the specified architecture.

Table 4.10: PSL Compiler Flags for Linux Platforms

Architecture PSL	Library Format	Compiler Flags Used by RTI
x86_64leElfgcc12.3.0-Linux5	Static Release	<b>C Flags:</b> -O2 -fstrict-aliasing -fsigned-char -fPIC -DNDEBUG <b>C++ Flags:</b> -O2 -fstrict-aliasing -fsigned-char -fPIC -DNDEBUG
	Shared Release	<b>C Flags:</b> -O2 -fstrict-aliasing -fsigned-char -fPIC -DNDEBUG <b>C++ Flags:</b> -O2 -fstrict-aliasing -fsigned-char -fPIC -DNDEBUG
	Static Debug	<b>C Flags:</b> -std=c99 -fstrict-aliasing -fsigned-char -g -fPIC <b>C++ Flags:</b> -fstrict-aliasing -fsigned-char -g -fPIC
	Shared Debug	<b>C Flags:</b> -std=c99 -fstrict-aliasing -fsigned-char -g -fPIC <b>C++ Flags:</b> -fstrict-aliasing -fsigned-char -g -fPIC
x86_64leElfgcc12.3.0CERT-Linux5	Static Release	<b>C Flags:</b> -std=c99 -O2 -fstrict-aliasing -fsigned-char -fPIC -DNDEBUG
	Static Debug	<b>C Flags:</b> -std=c99 -fstrict-aliasing -fsigned-char -g -fPIC
armv8leElfgcc7.3.0-Linux4	Static Release	<b>C Flags:</b> -std=c99 -O2 -fstrict-aliasing -fsigned-char -fPIC -DNDEBUG <b>C++ Flags:</b> -O2 -fstrict-aliasing -fsigned-char -fPIC -DNDEBUG
	Shared Release	<b>C Flags:</b> -std=c99 -O2 -fstrict-aliasing -fsigned-char -fPIC -DNDEBUG <b>C++ Flags:</b> -O2 -fstrict-aliasing -fsigned-char -fPIC -DNDEBUG
	Static Debug	<b>C Flags:</b> -std=c99 -fstrict-aliasing -fsigned-char -g -fPIC <b>C++ Flags:</b> -fstrict-aliasing -fsigned-char -g -fPIC
	Shared Debug	<b>C Flags:</b> -std=c99 -fstrict-aliasing -fsigned-char -g -fPIC <b>C++ Flags:</b> -fstrict-aliasing -fsigned-char -g -fPIC
armv8leElfgcc7.3.0CERT-Linux4	Static Release	<b>C Flags:</b> -std=c99 -O2 -fstrict-aliasing -fsigned-char -fPIC -DNDEBUG
	Static Debug	<b>C Flags:</b> -std=c99 -fstrict-aliasing -fsigned-char -g -fPIC
armv8leElfgcc11.3.1-Linux5	Static Release	<b>C Flags:</b> -std=c99 -fstrict-aliasing -fsigned-char -O -DNDEBUG <b>C++ Flags:</b> -O -fPIC -DNDEBUG
	Shared Release	<b>C Flags:</b> -std=c99 -fstrict-aliasing -fsigned-char -O -DNDEBUG <b>C++ Flags:</b> -O -fPIC -DNDEBUG

continues on next page

Table 4.10 – continued from previous page

Architecture PSL	Library Format	Compiler Flags Used by RTI
	Static Debug	<b>C Flags:</b> -std=c99 -fstrict-aliasing -fsigned-char -O0 -g -fPIC <b>C++ Flags:</b> -O0 -g
	Shared Debug	<b>C Flags:</b> -std=c99 -fstrict-aliasing -fsigned-char -O0 -g -fPIC <b>C++ Flags:</b> -O0 -g
armv8aleElfgcc11.3.1CERT-Linux5	Static Release	<b>C Flags:</b> -std=c99 -fstrict-aliasing -fsigned-char -O -DNDEBUG
	Static Debug	<b>C Flags:</b> -std=c99 -fstrict-aliasing -fsigned-char -O0 -g

## 4.4 macOS Platforms

The following table shows the currently supported macOS platforms.

Table 4.11: Supported macOS Platforms

OS	Version	CPU	Network Stack	Toolchain	Architecture PIL	Architecture PSL
macOS	14	x64	OS Default	clang 15.0	x86_64leMachOclang15.0 x86_64leMa- chOclang15.0CERT	x86_64leMa- chOclang15.0-Darwin23 x86_64leMa- chOclang15.0CERT-Darwin23
macOS	14	arm64	OS Default	clang 15.0	armv8leMachOclang15.0 armv8leMachOclang15.0CERT	armv8leMa- chOclang15.0-Darwin23 armv8leMa- chOclang15.0CERT-Darwin23

### 4.4.1 How the PIL was built for macOS platforms

This section describes how RTI built the Platform Independent Library (PIL) for macOS.

The following table shows the compiler flags RTI used to create the PIL for macOS platforms:

Table 4.12: PIL Compiler Flags for macOS Platforms

Architecture PIL	Library Format	Compiler Flags Used by RTI
x86_64leMachOclang15.0	Static Release	<b>C Flags:</b> -std=c99 -nostdinc -fstrict-aliasing -fsigned-char -O2 -fPIC -DNDEBUG -DNDEBUG -DOSAPI_CC_DEF_H=osapi/osapi_cc_clang.h -DOSPSL_OS_DEF_H=rti_me_psl/ospsl/ospsl_os_stubs.h -DRTIME_TARGET_NAME="x86_64leMachOclang15.0" -DRTI_PIL=1 <b>C++ Flags:</b> -fstrict-aliasing -fsigned-char -O2 -fPIC -DNDEBUG -DOSAPI_CC_DEF_H=osapi/osapi_cc_clang.h -DOSPSL_OS_DEF_H=rti_me_psl/ospsl/ospsl_os_stubs.h -DRTIME_TARGET_NAME="x86_64leMachOclang15.0" -DRTI_PIL=1
	Shared Release	<b>C Flags:</b> -std=c99 -nostdinc -fstrict-aliasing -fsigned-char -O2 -fPIC -DNDEBUG <b>C++ Flags:</b> -fstrict-aliasing -fsigned-char -O2 -fPIC -DNDEBUG
	Static Debug	<b>C Flags:</b> -std=c99 -nostdinc -fstrict-aliasing -fsigned-char -g -fPIC <b>C++ Flags:</b> -fstrict-aliasing -fsigned-char -g -fPIC
	Shared Debug	<b>C Flags:</b> -std=c99 -nostdinc -fstrict-aliasing -fsigned-char -g -fPIC <b>C++ Flags:</b> -fstrict-aliasing -fsigned-char -g -fPIC
x86_64leMachOclang15.0CERT	Static Release	<b>C Flags:</b> -std=c99 -nostdinc -O2 -fstrict-aliasing -fsigned-char -fPIC -DNDEBUG
	Static Debug	<b>C Flags:</b> -std=c99 -nostdinc -fstrict-aliasing -fsigned-char -g -fPIC
armv8leMachOclang15.0	Static Release	<b>C Flags:</b> -std=c99 -nostdinc -O2 -fstrict-aliasing -fsigned-char -fPIC -DNDEBUG <b>C++ Flags:</b> -fstrict-aliasing -fsigned-char -O2 -fPIC -DNDEBUG
	Shared Release	<b>C Flags:</b> -std=c99 -nostdinc -O2 -fstrict-aliasing -fsigned-char -fPIC -DNDEBUG <b>C++ Flags:</b> -fstrict-aliasing -fsigned-char -O2 -fPIC -DNDEBUG
	Static Debug	<b>C Flags:</b> -std=c99 -fsigned-char -nostdinc -fstrict-aliasing -g -fPIC <b>C++ Flags:</b> -fsigned-char -fstrict-aliasing -g -fPIC

continues on next page



Table 4.12 – continued from previous page

Architecture PIL	Library Format	Compiler Flags Used by RTI
	Shared Debug	<b>C Flags:</b> -std=c99 -fsigned-char -nostdinc -fstrict-aliasing -g -fPIC <b>C++ Flags:</b> -fsigned-char -fstrict-aliasing -g -fPIC
armv8leMachOclang15.0CERT	Static Release	<b>C Flags:</b> -std=c99 -O2 -fsigned-char -nostdinc -fstrict-aliasing -fPIC -DNDEBUG
	Static Debug	<b>C Flags:</b> -std=c99 -fsigned-char -nostdinc -fstrict-aliasing -g -fPIC

**Warning:** The *RTI Connex Micro* platform independent libraries are built without the standard C header-files. However, *Connex Micro* makes one direct call to the C library API `qsort`. In addition, CLANG may insert direct calls to GLIBC functions and other required functions, such as default C++ constructors and destructors. For this reason, it is necessary to use a CLANG version that is compatible with the CLANG version used to build the platform independent libraries, or provide a C library with an implementation of the required functions. Future versions of *RTI Connex Micro* will remove these dependencies.

#### 4.4.2 Building the PSL from source for macOS platforms

Refer to *Building the PSL* for instructions on how to build your own Platform Support Library (PSL) for macOS platforms.

The following table shows the compiler flags and required options that RTI used to build the PSL for macOS platforms. When you build the PSL with `rtime-make`, the `--target` argument automatically adds all the necessary flags for the specified architecture.

Table 4.13: PSL Compiler Flags for macOS Platforms

Architecture PSL	Library Format	Compiler Flags Used by RTI
x86_64leMachOclang15.0-Darwin23	Static Release	<b>C Flags:</b> -std=c99 -fsigned-char -O2 -fstrict-aliasing -O2 -isysroot /Applications/Xcode.app/Contents/Developer/Platforms/MacOSX.platform/Developer/SDKs/MacOSX14.5.sdk -mmacosx-version-min=14.3 -fPIC -DNDEBUG <b>C++ Flags:</b> -fstrict-aliasing -fsigned-char -O2 -isysroot /Applications/Xcode.app/Contents/Developer/Platforms/MacOSX.platform/Developer/SDKs/MacOSX14.5.sdk -mmacosx-version-min=14.3 -fPIC -DNDEBUG
	Shared Release	<b>C Flags:</b> -std=c99 -fsigned-char -O2 -fstrict-aliasing -O2 -isysroot /Applications/Xcode.app/Contents/Developer/Platforms/MacOSX.platform/Developer/SDKs/MacOSX14.5.sdk -mmacosx-version-min=14.3 -fPIC -DNDEBUG <b>C++ Flags:</b> -fstrict-aliasing -fsigned-char -O2 -isysroot /Applications/Xcode.app/Contents/Developer/Platforms/MacOSX.platform/Developer/SDKs/MacOSX14.5.sdk -mmacosx-version-min=14.3 -fPIC -DNDEBUG
	Static Debug	<b>C Flags:</b> -std=c99 -fsigned-char -fstrict-aliasing -g -isysroot /Applications/Xcode.app/Contents/Developer/Platforms/MacOSX.platform/Developer/SDKs/MacOSX14.5.sdk -mmacosx-version-min=14.3 -fPIC <b>C++ Flags:</b> -fstrict-aliasing -fsigned-char -g -isysroot /Applications/Xcode.app/Contents/Developer/Platforms/MacOSX.platform/Developer/SDKs/MacOSX14.5.sdk -mmacosx-version-min=14.3 -fPIC
	Shared Debug	<b>C Flags:</b> -std=c99 -fsigned-char -fstrict-aliasing -g -isysroot /Applications/Xcode.app/Contents/Developer/Platforms/MacOSX.platform/Developer/SDKs/MacOSX14.5.sdk -mmacosx-version-min=14.3 -fPIC <b>C++ Flags:</b> -fstrict-aliasing -fsigned-char -g -isysroot /Applications/Xcode.app/Contents/Developer/Platforms/MacOSX.platform/Developer/SDKs/MacOSX14.5.sdk -mmacosx-version-min=14.3 -fPIC

continues on next page

Table 4.13 – continued from previous page

Architecture PSL	Library Format	Compiler Flags Used by RTI
x86_64leMachOclang15.0CERT-Darwin23	Static Release	<b>C Flags:</b> -std=c99 -fsigned-char -O2 -fstrict-aliasing -isysroot /Applications/Xcode.app/Contents/Developer/Platforms/MacOSX.platform/Developer/SDKs/MacOSX14.5.sdk -mmacosx-version-min=14.3 -fPIC -DNDEBUG
	Static Debug	<b>C Flags:</b> -std=c99 -fsigned-char -fstrict-aliasing -g -isysroot /Applications/Xcode.app/Contents/Developer/Platforms/MacOSX.platform/Developer/SDKs/MacOSX14.5.sdk -mmacosx-version-min=14.3 -fPIC
armv8leMachOclang15.0-Darwin23	Static Release	<b>C Flags:</b> -std=c99 -fsigned-char -O2 -fstrict-aliasing -O2 -isysroot /Applications/Xcode.app/Contents/Developer/Platforms/MacOSX.platform/Developer/SDKs/MacOSX14.5.sdk -mmacosx-version-min=11 -fPIC -DNDEBUG <b>C++ Flags:</b> -fstrict-aliasing -fsigned-char -O2 -isysroot /Applications/Xcode.app/Contents/Developer/Platforms/MacOSX.platform/Developer/SDKs/MacOSX14.5.sdk -mmacosx-version-min=11 -fPIC -DNDEBUG
	Shared Release	<b>C Flags:</b> -std=c99 -fsigned-char -O2 -fstrict-aliasing -O2 -isysroot /Applications/Xcode.app/Contents/Developer/Platforms/MacOSX.platform/Developer/SDKs/MacOSX14.5.sdk -mmacosx-version-min=11 -fPIC -DNDEBUG <b>C++ Flags:</b> -fstrict-aliasing -fsigned-char -O2 -isysroot /Applications/Xcode.app/Contents/Developer/Platforms/MacOSX.platform/Developer/SDKs/MacOSX14.5.sdk -mmacosx-version-min=11 -fPIC -DNDEBUG
	Static Debug	<b>C Flags:</b> -std=c99 -fsigned-char -fstrict-aliasing -g -isysroot /Applications/Xcode.app/Contents/Developer/Platforms/MacOSX.platform/Developer/SDKs/MacOSX14.5.sdk -mmacosx-version-min=11 -fPIC <b>C++ Flags:</b> -fsigned-char -fstrict-aliasing -g -isysroot /Applications/Xcode.app/Contents/Developer/Platforms/MacOSX.platform/Developer/SDKs/MacOSX14.5.sdk -mmacosx-version-min=11 -fPIC

continues on next page

Table 4.13 – continued from previous page

Architecture PSL	Library Format	Compiler Flags Used by RTI
	Shared Debug	<b>C Flags:</b> -std=c99 -fsigned-char -fstrict-aliasing -g -isysroot /Applications/Xcode.app/Contents/Developer/Platforms/MacOSX.platform/Developer/SDKs/MacOSX14.5.sdk -mmacosx-version-min=11 -fPIC <b>C++ Flags:</b> -fsigned-char -fstrict-aliasing -g -isysroot /Applications/Xcode.app/Contents/Developer/Platforms/MacOSX.platform/Developer/SDKs/MacOSX14.5.sdk -mmacosx-version-min=11 -fPIC
armv8leMa- chOclang15.0CERT-Darwin23	Static Release	<b>C Flags:</b> -std=c99 -fsigned-char -O2 -fstrict-aliasing -isysroot /Applications/Xcode.app/Contents/Developer/Platforms/MacOSX.platform/Developer/SDKs/MacOSX14.5.sdk -mmacosx-version-min=11 -fPIC -DNDEBUG
	Static Debug	<b>C Flags:</b> -std=c99 -fsigned-char -fstrict-aliasing -g -isysroot /Applications/Xcode.app/Contents/Developer/Platforms/MacOSX.platform/Developer/SDKs/MacOSX14.5.sdk -mmacosx-version-min=11 -fPIC

## 4.5 QNX Platforms

The following table shows the currently supported QNX platforms.

Table 4.14: Supported QNX Platforms

OS	Version	CPU	Network Stack	Toolchain	Architecture PIL	Architecture PSL
QNX	7.1	ARMv8 (64-bit)	OS De-fault: io-pkt	qcc_gpp8.3.0	armv8leElfqnx_qcc8.3.0 armv8leElfqnx_qcc8.3.0CERT	armv8leElfqnx_qcc8.3.0-QNX7.1 armv8leElfqnx_qcc8.3.0CERT-QNX7.1
QOS (QNX OS for Safety)	2.2.1	ARMv8 (64-bit)	OS De-fault: io-pkt	qcc_gpp8.3.0	armv8leElfqcc8.3.0 armv8leElfqcc8.3.0CERT	armv8leElfqcc8.3.0-QOS2.2.1 armv8leElfqcc8.3.0CERT-QOS2.2.1

### 4.5.1 How the PIL was built for QNX platforms

This section describes how RTI built the Platform Independent Library (PIL) for QNX platforms.

The following table shows the compiler flags RTI used to create the PIL for QNX platforms:

Table 4.15: PIL Compiler Flags for QNX Platforms

Architecture PIL	Library Format	Compiler Flags Used by RTI
armv8leElfqnx_qcc8.3.0	Static Release	<b>C</b> <b>Flags:</b> -nostdinc -fsigned-char -Vgcc/{version},gcc_ntoaarch64le -O2 -fPIC -DNDEBUG <b>C++</b> <b>Flags:</b> -nostdinc -fsigned-char -O2 -Vgcc/{version},gcc_ntoaarch64le -fPIC -lang-c++ -DNDEBUG
	Shared Release	<b>C</b> <b>Flags:</b> -nostdinc -fsigned-char -Vgcc/{version},gcc_ntoaarch64le -O2 -fPIC -DNDEBUG <b>C++</b> <b>Flags:</b> -nostdinc -fsigned-char -O2 -Vgcc/{version},gcc_ntoaarch64le -fPIC -lang-c++ -DNDEBUG

continues on next page

Table 4.15 – continued from previous page

Architecture PIL	Library Format	Compiler Flags Used by RTI
	Static Debug	<b>C Flags:</b> -nostdinc -fsigned-char -Vgcc/\${version},gcc_ntoaarch64le -g -fPIC <b>C++ Flags:</b> -nostdinc -fsigned-char -Vgcc/\${version},gcc_ntoaarch64le -g -fPIC -lang-c++
	Shared Debug	<b>C Flags:</b> -nostdinc -fsigned-char -Vgcc/\${version},gcc_ntoaarch64le -g -fPIC <b>C++ Flags:</b> -nostdinc -fsigned-char -Vgcc/\${version},gcc_ntoaarch64le -g -fPIC -lang-c++
armv8leElfqnx_qcc8.3.0CERT	Static Release	<b>C Flags:</b> -nostdinc -fsigned-char -Vgcc/8.3.0,gcc_ntoaarch64le -O2 -fPIC -DNDEBUG
	Static Debug	<b>C Flags:</b> -nostdinc -fsigned-char -Vgcc/8.3.0,gcc_ntoaarch64le -g -fPIC
armv8leElfqcc8.3.0	Static Release	<b>C Flags:</b> -nostdinc -fsigned-char -Vgcc/8.3.0,gcc_ntoaarch64le -O2 -fPIC -DNDEBUG <b>C++ Flags:</b> -nostdinc -fsigned-char -Vgcc/8.3.0,gcc_ntoaarch64le -O2 -fPIC -lang-c++ -DNDEBUG
	Shared Release	<b>C Flags:</b> -nostdinc -fsigned-char -Vgcc/8.3.0,gcc_ntoaarch64le -O2 -fPIC -DNDEBUG <b>C++ Flags:</b> -nostdinc -fsigned-char -Vgcc/8.3.0,gcc_ntoaarch64le -O2 -fPIC -lang-c++ -DNDEBUG
	Static Debug	<b>C Flags:</b> -nostdinc -fsigned-char -Vgcc/8.3.0,gcc_ntoaarch64le -g -fPIC <b>C++ Flags:</b> -nostdinc -fsigned-char -Vgcc/8.3.0,gcc_ntoaarch64le -g -fPIC -lang-c++
	Shared Debug	<b>C Flags:</b> -nostdinc -fsigned-char -Vgcc/8.3.0,gcc_ntoaarch64le -g -fPIC <b>C++ Flags:</b> -nostdinc -fsigned-char -Vgcc/8.3.0,gcc_ntoaarch64le -g -fPIC -lang-c++
armv8leElfqcc8.3.0CERT	Static Release	<b>C Flags:</b> -nostdinc -fsigned-char -O2 -Vgcc/\${version},gcc_ntoaarch64le -O2 -fPIC -DNDEBUG
	Static Debug	<b>C Flags:</b> -nostdinc -fsigned-char -Vgcc/\${version},gcc_ntoaarch64le -g -fPIC

**Warning:** The *RTI Connex Micro* platform independent libraries are built without the standard C header-files. However, *Connex Micro* makes one direct call to the C library API `qsort`. In addition, QCC may insert direct calls to GLIBC functions and other required functions, such as default C++ constructors and destructors. For this reason, it is necessary to use a QCC version that is compatible with the QCC version used to build the platform independent libraries, or provide a C library with an implementation of the required functions. Future versions of *RTI Connex Micro* will remove these dependencies.

### 4.5.2 Building the PSL from source for QNX platforms

Refer to *Building the PSL* for instructions on how to build your own Platform Support Library (PSL) for QNX platforms.

The following table shows the compiler flags and required options that RTI used to build the PSL for QNX platforms. When you build the PSL with `rtime-make`, the `--target` argument automatically adds all the necessary flags for the specified architecture.

Table 4.16: PSL Compiler Flags for QNX Platforms

Architecture PSL	Library Format	Compiler Flags Used by RTI
armv8leElfqnx_qcc8.3.0-QNX7.1	Static Release	<b>C Flags:</b> -Vgcc/\${version},gcc_ntoaarch64le -fsigned-char -O2 -Y_gpp -DNDEBUG <b>C++ Flags:</b> -Vgcc/\${version},gcc_ntoaarch64le -Y_gpp -fsigned-char -O2 -lang-c++ -DNDEBUG
	Static Debug	<b>C Flags:</b> -Vgcc/\${version},gcc_ntoaarch64le -fsigned-char -O2 -Y_gpp -DNDEBUG <b>C++ Flags:</b> -Vgcc/\${version},gcc_ntoaarch64le -Y_gpp -fsigned-char -O2 -lang-c++ -DNDEBUG
	Shared Release	<b>C Flags:</b> -Vgcc/\${version},gcc_ntoaarch64le -fsigned-char -O2 -Y_gpp -DNDEBUG <b>C++ Flags:</b> -Vgcc/\${version},gcc_ntoaarch64le -Y_gpp -fsigned-char -O2 -lang-c++ -DNDEBUG
	Shared Debug	<b>C Flags:</b> -Vgcc/\${version},gcc_ntoaarch64le -fsigned-char -O2 -Y_gpp -DNDEBUG <b>C++ Flags:</b> -Vgcc/\${version},gcc_ntoaarch64le -Y_gpp -fsigned-char -O2 -lang-c++ -DNDEBUG

continues on next page

Table 4.16 – continued from previous page

Architecture PSL	Library Format	Compiler Flags Used by RTI
armv8leElfqnx_qcc8.3.0CERT-QNX7.1	Static Release	<b>C Flags:</b> -Vgcc/\${version},gcc_ntoaarch64le -fsigned-char -Y_gpp -O2 -DNDEBUG
	Static Debug	<b>C Flags:</b> -Vgcc/\${version},gcc_ntoaarch64le -fsigned-char -Y_gpp -g
armv8leElfqcc8.3.0-QOS2.2.1	Static Release	<b>C Flags:</b> -Vgcc/\${version},gcc_ntoaarch64le -fsigned-char -Y_gpp -O2 -DNDEBUG <b>C++ Flags:</b> -Vgcc/\${version},gcc_ntoaarch64le -Y_gpp -fsigned-char -O2 -lang-c++ -DNDEBUG
	Static Debug	<b>C Flags:</b> -Vgcc/\${version},gcc_ntoaarch64le -fsigned-char -Y_gpp -g <b>C++ Flags:</b> -Vgcc/\${version},gcc_ntoaarch64le -Y_gpp -fsigned-char -g -lang-c++
	Shared Release	<b>C Flags:</b> -Vgcc/\${version},gcc_ntoaarch64le -fsigned-char -Y_gpp -O2 -DNDEBUG <b>C++ Flags:</b> -Vgcc/\${version},gcc_ntoaarch64le -Y_gpp -fsigned-char -O2 -lang-c++ -DNDEBUG
	Shared Debug	<b>C Flags:</b> -Vgcc/\${version},gcc_ntoaarch64le -fsigned-char -Y_gpp -g <b>C++ Flags:</b> -Vgcc/\${version},gcc_ntoaarch64le -Y_gpp -fsigned-char -g -lang-c++
armv8leElfqcc8.3.0CERT-QOS2.2.1	Static Release	<b>C Flags:</b> -Vgcc/\${version},gcc_ntoaarch64le -fsigned-char -O2 -Y_gpp -DNDEBUG
	Static Debug	<b>C Flags:</b> -Vgcc/\${version},gcc_ntoaarch64le -fsigned-char -Y_gpp -g



## 4.6 Windows Platforms

The following table shows the currently supported Windows platforms.

Table 4.17: Supported Windows Platforms

OS	Version	CPU	Network Stack	Toolchain	Architecture PIL	Architecture PSL
Windows	10	x64	OS Default	Visual Studio 2017	x86_64lePEvs2017 x86_64lePEvs2017CERT	x86_64lePEvs2017-Win10 x86_64leP- Evs2017CERT-Win10

### 4.6.1 How the PIL was built for Windows platforms

This section describes how RTI built the Platform Independent Library (PIL) for Windows platforms.

The following table shows the compiler flags RTI used to create the PIL for Windows platforms:

Table 4.18: PIL Compiler Flags for Windows Platforms

Architecture PIL	Library Format	Compiler Flags Used by RTI
x86_64lePEvs2017	Static/Shared Release	<b>C Flags:</b> _WINDOWS;WIN32;_WINDOWS;WIN32_LEAN_AND_MEAN;_WIN32_WINNT=0x0600;_CRT_SECURE_NO_WARNINGS; <b>C++ Flags:</b> _WINDOWS;WIN32;_WINDOWS;WIN32_LEAN_AND_MEAN;_WIN32_WINNT=0x0600;_CRT_SECURE_NO_WARNINGS;
	Static/Shared Debug	<b>C Flags:</b> _DEBUG;_WINDOWS;WIN32;_WINDOWS;WIN32_LEAN_AND_MEAN;_WIN32_WINNT=0x0600;_CRT_SECURE_NO_WARNINGS; <b>C++ Flags:</b> _DEBUG;_WINDOWS;WIN32;_WINDOWS;WIN32_LEAN_AND_MEAN;_WIN32_WINNT=0x0600;_CRT_SECURE_NO_WARNINGS;

continues on next page

Table 4.18 – continued from previous page

Architecture PIL	Library Format	Compiler Flags Used by RTI
x86_64lePEvs2017CERT	Static Release	<b>C Flags:</b> _WINDOWS;WIN32;_WINDOWS;WIN32_LEAN_AND_MEAN;_WIN32_WINNT=0x0600;_CRT_SECURE_NO_WARNINGS;
	Static Debug	<b>C Flags:</b> _DEBUG;_WINDOWS;WIN32;_WINDOWS;WIN32_LEAN_AND_MEAN;_WIN32_WINNT=0x0600;_CRT_SECURE_NO_WARNINGS;

**Warning:** The *RTI Connex Micro* platform independent libraries are built without the standard C header-files. However, *Connex Micro* makes one direct call to the C library API `qsort`. In addition, MSVS may insert direct calls to C library functions and other required functions, such as default C++ constructors and destructors. For this reason it is necessary to use an MSVS version that is compatible with the MSVS version used to build the platform independent libraries, or provide a C library with an implementation of the required functions. Future versions of *RTI Connex Micro* will remove these dependencies.

#### 4.6.2 Building the PSL from source for Windows platforms

Refer to *Building the PSL* for instructions on how to build your own Platform Support Library (PSL) for Windows platforms.

The following table shows the compiler flags and required options that RTI used to build the PSL for Windows platforms. When you build the PSL with `rtime-make`, the `--target` argument automatically adds all the necessary flags for the specified architecture.

Table 4.19: PSL Compiler Flags for Windows Platforms

Architecture PSL	Library Format	Compiler Flags Used by RTI
x86_64lePEvs2017-Win10	Static/Shared Release	<b>C Flags:</b> _WINDOWS;WIN32;_WINDOWS;WIN32_LEAN_AND_MEAN;_WIN32_WINNT=0x0600;_CRT_SECURE_NO_WARNINGS; <b>C++ Flags:</b> _WINDOWS;WIN32;_WINDOWS;WIN32_LEAN_AND_MEAN;_WIN32_WINNT=0x0600;_CRT_SECURE_NO_WARNINGS;

continues on next page

Table 4.19 – continued from previous page

Architecture PSL	Library Format	Compiler Flags Used by RTI
	Static/Shared Debug	<b>C</b> <b>Flags:</b> _DEBUG;_WINDOWS;WIN32;_WIN- DOWS;WIN32_LEAN_AND_MEAN;_WIN32_WINNT=0x0600;_CRT_SE- CURE_NO_WARNINGS; <b>C++</b> <b>Flags:</b> _DEBUG;_WIN- DOWS;WIN32;_WINDOWS;WIN32_LEAN_AND_MEAN;_WIN32_WINNT=0x0600 CURE_NO_WARNINGS;
x86_64lePEvs2017CERT-Win10	Static Release	<b>C</b> <b>Flags:</b> _WINDOWS;WIN32;_WIN- DOWS;WIN32_LEAN_AND_MEAN;_WIN32_WINNT=0x0600;_CRT_SE- CURE_NO_WARNINGS;
	Static Debug	<b>C</b> <b>Flags:</b> _DEBUG;_WINDOWS;WIN32;_WIN- DOWS;WIN32_LEAN_AND_MEAN;_WIN32_WINNT=0x0600;_CRT_SE- CURE_NO_WARNINGS;

## Chapter 5

# Building Connex Micro

### 5.1 Connex Micro Platforms

*Connex Micro* includes reference Platform Support Libraries (PSL) as both binaries and source code for the target architectures below. The PSL binaries have been tested and validated by RTI; the included source code is identical to the source code used to build the PSL binaries.

- Windows®
- Linux®
- macOS® (Darwin)
- QNX®
- FreeRTOS®

Refer to the *Supported Platforms and Programming Languages* section for a complete list of all available target architectures.

### 5.2 Building Connex Micro for Common Platforms

This section describes how to compile *Connex Micro*, either the Platform Support Libraries (PSL) or the full source if available, for an architecture supported by RTI (see *Connex Micro Platforms* for more information).

For information about how to compile and link *Connex Micro* applications, please refer to *Prepare Your Development Environment*.

This section is written for developers and engineers with a background in software development. RTI recommends reading this section in order, as one subsection may refer to or assume knowledge about concepts described in a preceding subsection.

### 5.2.1 Setting up the build environment

The following terminology is used to refer to the environment in which *Connex Micro* is built and run:

- The *host* is the machine that runs the software to compile and link *Connex Micro*.
- The *target* is the machine that will run *Connex Micro*.
- In many cases *Connex Micro* is built *and* run on the same machine. This is referred to as a *self-hosted environment*.

The *environment* is the collection of tools, OS, compiler, linker, hardware etc. needed to build and run applications.

The word *must* describes a requirement that must be met. Failure to meet a *must* requirement may result in failure to compile, use, or run *Connex Micro*.

The word *should* describes a requirement that is strongly recommended to be met. A failure to meet a *should* recommendation may require modification to how *Connex Micro* is built, used, or run.

The word *may* is used to describe an optional feature.

#### The host environment

*Connex Micro* has been designed to be easy to build and to require few tools on the host.

The host machine **must**:

- support long filenames (8.3 will not work). *Connex Micro* does not require a case-sensitive file-system.
- have the necessary compiler, linkers, and build-tools installed.

The host machine **should**:

- have [CMake](http://www.cmake.org) ([www.cmake.org](http://www.cmake.org)) installed. Note that it is not required to use [CMake](http://www.cmake.org) to build *Connex Micro*, and in some cases it may also not be recommended. As a rule of thumb, if *Connex Micro* can be built from the command-line, [CMake](http://www.cmake.org) is recommended.
- be able to run bash shell scripts (Linux and macOS systems) or BAT scripts (Windows systems).

Supported host environments are Windows (cygwin and mingw are not tested), Linux, and macOS systems.

Typical examples of host machines are:

- a Linux PC with the GNU tools installed (make, gcc, g++, etc.).
- a Mac computer with Xcode and the command-line tools installed.
- a Windows computer with Microsoft Visual Studio Express edition.
- a Linux, Mac or Windows computer with an embedded development tool-suite.

## The target environment

The target compiler should:

- be C99 compliant. Note that many non-standard compilers work, but may require additional configuration.
- be C++98 compliant.

The remainder of this manual assumes that the target environment is one supported by RTI:

- POSIX (Linux, macOS, QNX®)
- Windows

### 5.2.2 Building the PSL

As described in *Library descriptions*, *Connex Micro* comes with Platform Support Libraries (PSL) that are compatible with specific Platform Independent Libraries (PIL). The source code for the PSL is also available in the host package. This section describes how to build the PSL from the provided source code.

RTI provides the PSL source code because it allows the PSL to be recompiled for a specific platform configuration. This may be important in some use cases if the header files include platform-specific information that is different from the binaries provided by RTI.

There are two recommended methods to compile the PSL: by running the `rtime-make` script (which invokes CMake), or by invoking CMake manually. Both are described in more detail below.

[CMake](#) is the preferred tool to build *Connex Micro* because it simplifies configuring the *Connex Micro* build options and generates build files for a variety of environments. Note that CMake itself does not compile anything. CMake is used to *generate* build files for a number of environments, such as `make`, Eclipse® CDT, Xcode® and Visual Studio. Once the build-files have been generated, any of the tools mentioned can be used to build *Connex Micro*. This system makes it easier to support building *Connex Micro* in different build environments. CMake is easy to install with pre-built binaries for common environments and does not depend on external tools to build *Connex Micro*.

#### Building the PSL with `rtime-make`

The *Connex Micro* source bundle includes a `bash` (on Linux and macOS systems) and `BAT` (on Windows systems) script to simplify the invocation of [CMake](#). These scripts are a convenient way to invoke [CMake](#) with the correct options.

---

**Note:** `rtime-make` must be invoked from the `RTIMEHOME` directory.

---

Run the `rtime-make` script with the following command:

Linux

```
RTIMEHOME/resource/scripts/rtime-make --config Debug --target x86_64leElfgcc7.3.0-Linux4_
↪ \
    -G "Unix Makefiles" --build
```

## Windows

```
RTIMEHOME\resource\scripts\rtime-make --config Debug --target x86_64lePEvs2017-Win10 \
    -G "Visual Studio 15 2017" --build
```

**Note:** When building for Windows systems on any architecture other than x86, you must add the `-A` option for your architecture. For example:

```
-A x64
```

When the compilation has finished, the PSL is copied to the directory `RTIMEHOME/lib/<target>` where `<target>` is the argument passed to the `--target <target>`.

**Warning:** The above command will overwrite the PSL installed by RTI. To use a different output directory refer to *Specifying a different output directory*.

Here is an explanation of each argument in the above command:

- `--config Debug`: Create Debug build. Use `--config Release` to create a release build.
- `--target <target>`: The target for the sources to be built. Refer to *Supported Platforms and Programming Languages* for the architecture abbreviations of supported platforms.
- `--build Build`: Build the generated project files.

To get a list of all the options, run:

```
rtime-make -h
```

To get help for a specific target, run:

```
rtime-make --target <target> --help
```

## Specifying a different output directory

By default, `rtime-make` copies the compiled PSL into a directory named `RTIMEHOME/lib/<target>` where `<target>` is the argument that was passed to the `--target <target>` option.

To copy the compiled PSL to a different output directory, the `--name <name>` option can be used together with `--target <target>`. In this case, the PSL will be compiled using the same options as specified for `--target <target>`, but instead the PSLs will be copied to the directory `RTIMEHOME/lib/<name>`.

---

**Note:** You should use the same naming convention for `--name` as for `--target`. *Connex Micro* may use the directory name to determine the appropriate Platform Independent Library (PIL) for the compiled PSL.

---

For example, the following command will compile the PSL using the same target configuration as for `x86_64leElfgcc7.3.0-Linux4`, but copy the compiled PSL to `RTIMEHOME/lib/x86_64leElfgcc7.3.0-mypsl`

```
RTIMEHOME/resource/scripts/rtime-make --config Debug \
--target x86_64leElfgcc7.3.0-Linux4 \
--name x86_64leElfgcc7.3.0-mypsl \
-G "Unix Makefiles" --build
```

## Building the PSL with CMake

### Preparing to build

RTI recommends creating a unique directory for each build configuration. A build configuration can be created to address specific architectures, compiler settings, or different *Connex Micro* build options.

RTI recommends assigning a descriptive *name* to each build configuration, using a common format. While there are no requirements to the format for functional correctness, the toolchain files in *Cross-compiling Connex Micro* use the `<name>` parameter passed to `--target <name>` to determine various compiler options and selections.

RTI uses the following format for the target architecture PSL:

```
{cpu}{compiler}{profile}-{OS}
```

- `{cpu}`: the CPU that the library was compiled for.
- `{compiler}`: the compiler used to build the library.
- `{profile}`: CERT if the library was built to be Cert-compatible; otherwise empty.
- `{OS}`: The operating system that the PSL was compiled for.

Some examples of target names:

- `x86_64leElfgcc7.3.0-Linux4`: PSL for *Connex Micro* for an x64 CPU compiled using GCC 7.3.0 and running a Linux4 kernel.
- `x86_64lePEvs2017-Win10`: PSL for *Connex Micro* for an x64 CPU compiled using VS2017 and running in Windows 10.

Files built by each build configuration will be stored under `RTIMEHOME/build/[Debug | Release]/<name>`. These directories are referred to as build directories or `RTIMEBUILD`. The structure of the `RTIMEBUILD` depends on the generated build files and should be regarded as an intermediate directory.



## Creating build files from the command line

Open a terminal window in the `RTIMEHOME` directory and create the `RTIMEBUILD` directory. Change to the `RTIMEBUILD` directory and invoke CMake with the following arguments:

**Note:** This section assumes that `cmake` is invoked from the `RTIMEHOME` directory.

```
cmake -G <generator> -DCMAKE_BUILD_TYPE=<Debug | Release> \
      -DCMAKE_TOOLCHAIN_FILE=<toolchain file> \
      -DCMAKE_MODULE_PATH=RTIMEHOME/resource/cmake/architectures \
      -DRTIME_TARGET_NAME=<target-name> \
      -DRTIME_TARGET=<target-name> \
      RTIMEHOME
```

Depending on the generator, do one of the following:

- For IDE generators (such as Eclipse, Visual Studio, Xcode), open the generated solution/project files and build the project/solution.
- For command-line tools (such as `make`, `nmake`, `ninja`), run the build-tool.

After a successful build, the output is placed in `RTIMEHOME/lib/<target-name>`.

The generated build files may contain different sub-projects that are specific to the tool. For example, in Xcode and Visual Studio, the following targets are available:

- `ALL_BUILD`: Builds all the projects.
- `\rti_me_<name>`: Builds only the specific library. Note that that dependent libraries are built first.
- `ZERO_CHECK`: Runs [CMake](#) to regenerate project files in case something changed in the build input. This target does not need to be built manually.

For command-line tools, try `<tool> help` for a list of available targets to build. For example, if makefiles were generated:

```
make help
```

### 5.2.3 Building the source

**Warning:** This section only applies to *Connext Micro* source bundles (`rti_connext_dds_micro-<version>-source.zip`). For other bundles, refer to *Building the PSL*.

When you build *Connext Micro* from the source bundle, you have two options:

- Build a Platform Independent Library (PIL) and a compatible Platform Support Library (PSL).

- Build an integrated library.

The source code for the PIL, PSL, and integrated libraries are all included in the source bundle. Refer to *Library descriptions* for more information on the differences between them. The following sections explain how to compile each library type.

There are two recommended methods to compile libraries: by running the `rtime-make` script (which invokes CMake), or by invoking CMake manually. Both are described in more detail below.

### Building with `rtime-make`

The *Connex Micro* source bundle includes a `bash` (on Linux and macOS systems) and `BAT` (on Windows systems) script to simplify the invocation of [CMake](#). These scripts are a convenient way to invoke [CMake](#) with the correct options.

Run the `rtime-make` script with the following commands:

Linux

To build the PIL:

```
RTIMEHOME/resource/scripts/rtime-make --config Debug --target x86_64leElfgcc7.3.0 \
-G "Unix Makefiles" --build
```

To build the PSL using the above PIL:

```
RTIMEHOME/resource/scripts/rtime-make --config Debug --target x86_64leElfgcc7.3.0-Linux4L
↩\
-G "Unix Makefiles" --build
```

To build the integrated library:

```
RTIMEHOME/resource/scripts/rtime-make --config Debug --target x64Linux4gcc7.3.0 \
-G "Unix Makefiles" --build
```

Windows

To build the PIL:

```
RTIMEHOME\resource\scripts\rtime-make --config Debug --target x86_64lePEvs2017 \
-G "Visual Studio 15 2017" --build
```

To build the PSL using the above PIL:

```
RTIMEHOME\resource\scripts\rtime-make --config Debug --target x86_64lePEvs2017-Win10 \
-G "Visual Studio 15 2017" --build
```

To build the integrated library:

```
RTIMEHOME\resource\scripts\rtime-make --config Debug --target x64Win64VS2017 \
-G "Visual Studio 15 2017" --build
```

Here is an explanation of each argument in the above commands:

- `--config Debug`: Create Debug build.
- `--target <target>`: The target for the sources to be built. Refer to *Supported Platforms and Programming Languages* for the architecture abbreviations of supported platforms.
- `--build Build`: The generated project files.

To get a list of all the options, run:

```
rttime-make -h
```

To get help for a specific target, run:

```
rttime-make --target <target> --help
```

## Building with CMake

### Preparing to build

RTI recommends creating a unique directory for each build configuration. A build configuration can be created to address specific architectures, compiler settings, or different *Connex Micro* build options.

RTI recommends assigning a descriptive *name* to each build configuration, using a common format. While there are no requirements to the format for functional correctness, the toolchain files in *Cross-compiling Connex Micro* use the **RTIME\_TARGET\_NAME** variable to determine various compiler options and selections.

RTI uses the following formats for the target architecture libraries:

PIL

```
{cpu}{compiler}{profile}
```

- `{cpu}`: the CPU that the library was compiled for.
- `{compiler}`: the compiler used to build the library.
- `{profile}`: CERT if the library was built to be Cert-compatible; otherwise empty.

PSL

```
{cpu}{compiler}{profile}-{OS}
```

- `{cpu}`: the CPU that the library was compiled for.
- `{compiler}`: the compiler used to build the library.
- `{profile}`: CERT if the library was built to be Cert-compatible; otherwise empty.
- `{OS}`: The operating system that the PSL was compiled for.

Integrated

`{cpu}{OS}{compiler}{profile}`

- `{cpu}`: the CPU that the library was compiled for.
- `{OS}`: The operating system that the integrated library was compiled for.
- `{compiler}`: the compiler used to build the library.
- `{profile}`: CERT if the library was built to be Cert-compatible; otherwise empty.

Some examples of target names:

- `x86_64leElfgcc7.3.0`: PIL for an x64 CPU, compiled with gcc 7.3.0.
- `x86_64leElfgcc7.3.0-Linux4`: PSL for an x64 CPU, running Ubuntu 18.04 LTS, compiled with gcc 7.3.0.
- `x64Linux4gcc7.3.0`: Integrated library for an x64 CPU, running Ubuntu 18.04 LTS, compiled with gcc 7.3.0.
- `x86_64lePEvs2017`: PIL for an x64 CPU, compiled with Visual Studio 2017.
- `x86_64lePEvs2017-Win10`: PSL for an x64 CPU, running Windows 10, compiled with Visual Studio 2017.
- `x64Win64VS2017`: Integrated library for an x64 CPU, running Windows 10, compiled with Visual Studio 2017.

Files built by each build configuration will be stored under `RTIMEHOME/build/[Debug | Release]/<name>`. These directories are referred to as build directories or `RTIMEBUILD`. The structure of the `RTIMEBUILD` depends on the generated build files and should be regarded as an intermediate directory.

## Creating build files from the command line

---

**Note:** This section assumes that CMake is invoked from the `RTIMEHOME` directory. For out-of-source builds using CMake, refer to *Building with CMake outside of source*.

---

Open a terminal window in the `RTIMEHOME` directory and create the `RTIMEBUILD` directory. Change to the `RTIMEBUILD` directory and invoke CMake with the following arguments:

```
cmake -G <generator> -DCMAKE_BUILD_TYPE=<Debug | Release> \
  -DCMAKE_TOOLCHAIN_FILE=<toolchain file> \
  -DCMAKE_MODULE_PATH=RTIMEHOME/resource/cmake/architectures \
  -DRTIME_TARGET_NAME=<target-name> \
  -DRTIME_TARGET=<target-name> \
  RTIMEHOME
```

Depending on the generator, do one of the following:

- For IDE generators (such as Eclipse, Visual Studio, Xcode), open the generated solution/project files and build the project/solution.

- For command-line tools (such as make, nmake, ninja), run the build-tool.

After a successful build, the output is placed in `RTIMEHOME/lib/<name>`.

The generated build-files may contain different sub-projects that are specific to the tool. For example, in Xcode and Visual Studio the following targets are available:

- `ALL_BUILD`: Builds all the projects.
- `\rti_me_<name>`: Builds only the specific library. Note that that dependent libraries are built first.
- `ZERO_CHECK`: Runs [CMake](#) to regenerate project files in case something changed in the build input. This target does not need to be built manually.

For command-line tools, try `<tool> help` for a list of available targets to build. For example, if makefiles were generated:

```
make help
```

## Building with CMake outside of source

---

**Note:** This option is only available with the *Connex Micro* source bundle.

---

You may need to build *Connex Micro* in a directory that is located outside `RTIMEHOME` and to install *Connex Micro* in a separate installation directory. In this case, do the following:

1. Create a build directory and change the current directory to it.
2. Invoke CMake with the following command to create build files for an integrated library:

```
cmake -DCMAKE_TOOLCHAIN_FILE=<RTI toolchain> \
      -DRTIME_TARGET=<target-name> \
      -DRTIME_TARGET_NAME=<target-name> \
      -DCMAKE_BUILD_TYPE=[debug | Release] \
      <path to RTIMEHOME/CMakeLists.txt>
```

Alternatively, you can create build files for split libraries. This requires two commands, one for the PIL and one for the PSL.

To build the PIL:

```
cmake -DCMAKE_TOOLCHAIN_FILE=<RTI toolchain> \
      -DRTIME_TARGET=<target-name> \
      -DRTIME_TARGET_NAME=<target-name> \
      -DCMAKE_BUILD_TYPE=[debug | Release] \
      <path to RTIMEHOME/CMakeLists.txt>
```

To build the PSL:

```
cmake -DCMAKE_TOOLCHAIN_FILE=<RTI toolchain> \
      -DRTIME_TARGET=<target-name> \
      -DRTIME_TARGET_NAME=<target-name> \
      -DCMAKE_BUILD_TYPE=[debug | Release] \
      -DRTIME_PIL_PATH=<path to RTI's PIL if compiling a PSL> \
      <path to RTIMEHOME/CMakeLists.txt>
```

3. Invoke the build tool on the generated build-files:

```
make
```

4. Install the compiled libraries using `cmake --install`:

```
cmake --install . --prefix <install path>
```

The last command will copy the header files and libraries to:

- `<install path>/include` (header files)
- `<install path>/lib` (libraries)

## 5.2.4 Cross-compiling Connex Micro

Cross-compiling the *Connex Micro* source-code uses the exact same process described in *Building the source*, but requires an additional toolchain file. A toolchain file is a [CMake](#) file that describes the compiler, linker, etc., needed to build the source for the target.

To see a list of available targets, use `--list`:

```
rttime-make --list
```

## 5.3 Building Connex Micro with Compatibility for Connex Cert

*RTI Connex Micro* can be compiled to only include features that are available or planned for *RTI Connex Cert*. This is useful to enable the development of a safety-certified project using *Connex Micro* before certification evidence for *Connex Cert* is available. Once *Connex Cert* certification is available, the transition from *Connex Micro* to *Connex Cert* typically requires few changes in the application. The *Connex Micro* Cert-compatibility profile refers to the subset of *Connex Micro* that is feature-comparable to *Connex Cert*.

**Warning:** Please note that this does not mean that certification evidence is provided for *Connex Micro* for any of these features or that using the Cert-compatibility profile constitutes safety. Please contact RTI for further information about *Connex Cert* and certification evidence.

When compiling *Connext Micro* with compatibility for *Connext Cert*, please refer to the *Platform Notes* for information on supported platforms, libraries, and features in CERT profiles. The following additional restrictions apply:

- The CERT profile `rti_me` library **includes** the following libraries:
  - `rti_me_discdpse`
  - `rti_me_rhsm`
  - `rti_me_whsm`
- Memory deallocation is not possible
- Any API that deallocates memory is not supported. In other words, any API whose name includes “finalize”, “free”, or “delete” is not supported (such as `DDS_DomainParticipantFactory_delete_participant()`, `DDS_DomainParticipantQos_finalize()`, or `OSAPI_Heap_free()`)
- Code generated by the *Connext Micro* code generator is compatible with *Connext Cert*, but the code must be generated with the code generator using the `-interpreted 0` option
- The Log module is only available in the debug build
- The UDP transport *must* be configured statically by using the API `UDP_InterfaceTable_add_entry()` and setting `UDP_InterfaceFactoryProperty.disable_auto_interface_config` equal to `RTI_TRUE`
- `OSAPI_Thread_sleep()` is not available
- Batching reception is not supported
- UDP Transformations are not supported
- The Property, User Data, and Partition Qos APIs are available in the *Connext Micro* Cert-compatibility profile, but are not yet available in *Connext Cert*

### 5.3.1 Compiling with compatibility for Connext Cert

To compile *Connext Micro* with the Cert-compatibility profile, you must use one of the available CERT architectures. To get a list of available CERT architectures, please use the following command:

```
cd <rti_me install directory>
resource/scripts/rtime-make --list
```

Architectures ending in CERT (e.g. `x64Linux5gcc12.3.0CERT`) are representative of Cert-compatibility profiles. To compile, use the following command:

```
cd <rti_me install directory>
resource/scripts/rtime-make --target x64Linux5gcc12.3.0CERT <other options>
```

The library is generated in the directory `lib/x64Linux5gcc12.3.0CERT`.

### 5.3.2 Compiling applications with compatibility for Connext Cert

To compile an application with compatibility for Connext Cert, the application must be compiled with the `RTI_CERT=1` preprocessor flag. This can be achieved with one of the following methods:

- If a `CMakeLists.txt` file generated with `rtiddsgen` is used, pass `-DRTIME_CERT=true` to either `rtime-make` or `cmake`.
- Pass `-DRTI_CERT=1` directly to the C preprocessor

With `rtime-make` on Linux or macOS:

```
resource/scripts/rtime-make -target x64Linux5gcc12.3.0CERT -DRTIME_CERT=true --src-dir .  
↪ <other options>
```

With `rtime-make` On Windows:

```
resource/scripts/rtime-make --target x64Win64VS2015CERT -DRTIME_CERT_eq_true --src-dir .  
↪ <other options>
```

Please refer to *Example Applications* for more information about generating examples.



## Chapter 6

# Working with Connex Micro and Connex Professional

In some cases, it may be necessary to write an application that is compiled against both *Connex Micro* and *Connex Professional*. In general this is not easy to do because *Connex Micro* supports a very limited set of features compared to *Connex Professional*.

However, due to the nature of the DDS API and the philosophy of declaring behavior through QoS profiles instead of using different APIs, it may be possible to share common code. In particular, *Connex Professional* supports configuration through QoS profile files, which eases the job of writing portable code.

### 6.1 Supported DDS Features

*Connex Micro* supports a subset of the DDS DCPS standard. A brief overview of the supported features are listed here. For a detailed list, please refer to the [C API Reference](#) and [C++ API Reference](#).

#### 6.1.1 DDS Entity Support

*Connex Micro* supports the following DDS entities. Please refer to the documentation for details.

- [DomainParticipantFactory](#)
- [DomainParticipant](#)
- [Topic](#)
- [Publisher](#)
- [Subscriber](#)
- [DataWriter](#)
- [DataReader](#)

### 6.1.2 DDS QoS Policy Support

*Connext Micro* supports the following DDS QoS Policies. Please refer to the documentation for details.

- [DDS\\_DataReaderProtocolQosPolicy](#)
- [DDS\\_DataReaderResourceLimitsQosPolicy](#)
- [DDS\\_DataWriterProtocolQosPolicy](#)
- [DDS\\_DataWriterResourceLimitsQosPolicy](#)
- [DDS\\_DeadlineQosPolicy](#)
- [DDS\\_DiscoveryQosPolicy](#)
- [DDS\\_DomainParticipantResourceLimitsQosPolicy](#)
- [DDS\\_DurabilityQosPolicy](#)
- [DDS\\_DestinationOrderQosPolicy](#)
- [DDS\\_EntityFactoryQosPolicy](#)
- [DDS\\_HistoryQosPolicy](#)
- [DDS\\_LivelinessQosPolicy](#)
- [DDS\\_OwnershipQosPolicy](#)
- [DDS\\_OwnershipStrengthQosPolicy](#)
- [DDS\\_ReliabilityQosPolicy](#)
- [DDS\\_ResourceLimitsQosPolicy](#)
- [DDS\\_RtpsReliableWriterProtocol\\_t](#)
- [DDS\\_SystemResourceLimitsQosPolicy](#)
- [DDS\\_TopicDataQosPolicy](#)
- [DDS\\_TransportQosPolicy](#)
- [DDS\\_UserDataQosPolicy](#)
- [DDS\\_UserTrafficQosPolicy](#)
- [DDS\\_WireProtocolQosPolicy](#)

## 6.2 Development Environment

There are no conflicts between *Connext Micro* and *Connext Professional* with respect to library names, header files, etc. It is advisable to keep the two installations separate, which is the normal case.

*Connext Micro* uses the environment variable `RTIMEHOME` to locate the root of the *Connext Micro* installation.

*Connext Professional* uses the environment variable `NDDSHOME` to locate the root of the *Connext Professional* installation.

## 6.3 Non-standard APIs

The DDS specification omits many APIs and policies necessary to configure a DDS application, such as transport, discovery, memory, logging, etc. In general, *Connext Micro* and *Connext Professional* do not share APIs for these functions.

It is recommended to configure *Connext Professional* using QoS profiles as much as possible.

## 6.4 QoS Policies

QoS policies defined by the DDS standard behave the same between *Connext Micro* and *Connext Professional*. However, note that *Connext Micro* does not always support all the values for a policy and in particular unlimited resources are not supported.

Unsupported QoS policies are the most likely reason for not being able to switch between *Connext Micro* and *Connext Professional*.

## 6.5 Standard APIs

APIs that are defined by the standard behave the same between *Connext Micro* and *Connext Professional*.

## 6.6 IDL Files

*Connext Micro* and *Connext Professional* use the same IDL compiler (`rtiddsgen`) and *Connext Micro* typically ships with the latest version. However, *Connext Micro* and *Connext Professional* use different templates to generate code and it is not possible to share the generated code. Thus, while the same IDL can be used, the generated output must be saved in different locations.

## 6.7 Interoperability

*Connext Micro* and *Connext Professional* interoperate on the wire unless noted otherwise.

### 6.7.1 Discovery

When trying to establish communication between an *Connext Micro* application that uses the Dynamic Participant Static Endpoint (DPSE) discovery module and an RTI product based on *Connext Professional*, every participant in the DDS system must be configured with a unique participant name. While the static discovery functionality provided by *Connext Professional* allows participants on different hosts to share the same name, *Connext Micro* requires every participant to have a different name to help keep the complexity of its implementation suitable for smaller targets.

Also, *Connext DataWriters* that are configured to send compressed data will not match with *Connext Micro DataReaders*, since *Connext Micro* does not support sending or receiving compressed data. See [DATA\\_REPRESENTATION QosPolicy in the Core Libraries User's Manual](#) for more information on the *Connext* compression feature.

### 6.7.2 Transports

When interoperating with *Connext Professional*, *Connext Micro* must specify at least one unicast transport for each *DataWriter* and *DataReader*, either from [DDS\\_DomainParticipantQos::transports](#) or the endpoint [DDS\\_DataReaderQos::transport](#) and [DDS\\_DataWriterQos::transport](#), as it expects to use the unicast transport's RTPS port mapping to determine automatic participant IDs if needed. This also affects *Connext Micro* itself, where participant IDs must be set manually if only multicast transports are enabled.

Also, when interoperating with *Connext Professional*, only one multicast transport can be specified per *DataReader* of *Connext Micro*.

## 6.8 Connext Tools

In general, *Connext Micro* is compatible with RTI tools and other products. The following sections provide additional information for each product.

### 6.8.1 Admin Console

Admin Console can discover and display *Connext Micro* applications that use full dynamic discovery (DPDE). When using static discovery (DPSE), it is required to use the Limited Bandwidth Endpoint Discovery (LBED) that is available as a separate product for *Connext Professional*. With the library a configuration file with the discovery configuration must be provided (just as in the case for products such as Routing Service, etc.). This is provided through the QoS XML file.

Data can be visualized from *Connext Micro DataWriters*. Keep in mind that *Connext Micro* does not currently distribute type information and the type information has to be provided through an XML file using the “Create Subscription” dialog. Unlike some other products, this information cannot be provided through the QoS XML file. To provide the data types to Admin Console, first run the code generator with the `-convertToXml` option:

```
rtiddsgen -convertToXml <file>
```

Then click on the “Load Data Types from XML file” hyperlink in the “Create Subscription” dialog and add the generated IDL file.

Other Features Supported:

- Match analysis is supported.
- Discovery-based QoS are shown.

The following resource-limits in *Connext Micro* must be incremented as follows when using Admin Console:

- Add 24 to `DDS_DomainParticipantResourceLimitsQosPolicy::remote_reader_allocation`
- Add 24 to `DDS_DomainParticipantResourceLimitsQosPolicy::remote_writer_allocation`
- Add 1 to `DDS_DomainParticipantResourceLimitsQosPolicy::remote_participant_allocation`
- Add 1 to `DDS_DomainParticipantResourceLimitsQosPolicy::remote_participant_allocation` if data-visualization is used

*Connext Micro* does not currently support any administration capabilities or services, and does not match with the Admin Console *DataReaders* and *DataWriters*. However, if matching *DataReaders* and *DataWriters* are created, e.g., by the application, the following resource must be updated:

- Add 48 to `DDS_DomainParticipantResourceLimitsQosPolicy::matching_writer_reader_pair_allocation`

## 6.8.2 Distributed Logger

This product is not supported by *Connext Micro*.

## 6.8.3 LabVIEW

The LabVIEW toolkit uses *Connext Professional*, and it must be configured as any other *Connext Professional* application. A possible option is to use the builtin *Connext Professional* profile: `BuiltinQosLib::Generic.ConnexMicroCompatibility`.

### 6.8.4 Monitor

This product is not supported by *Connex Micro*.

### 6.8.5 Recording Service

#### RTI Recorder

RTI Recorder is compatible with *Connex Micro* in the following ways:

- If static endpoint discovery is used, Recorder is compatible starting with version 5.1.0.3 and onwards.
- If dynamic endpoint discovery is used, Recorder is compatible with *Connex Micro* the same way it is with any other DDS application.
- In both cases, type information has to be provided via XML. Read [Recording Data with RTI Connex Micro](#) for more information.

#### RTI Replay

RTI Replay is compatible with *Connex Micro* in the following ways:

- If static endpoint discovery is used, Replay is compatible starting with version 5.1.0.3 and onwards.
- If dynamic endpoint discovery is used, Replay is compatible with *Connex Micro* the same way it is with any other DDS application.
- In both cases, type information has to be provided via XML. Read [Recording Data with RTI Connex Micro](#) for more information on how to convert from IDL to XML.

#### RTI Converter

Databases recorded with *Connex Micro* contains no type information in the DCPSPublication table, but the type information can be provided via XML. Read [Recording Data with RTI Connex Micro](#) for more information on how to convert from IDL to XML.

### 6.8.6 Wireshark

Wireshark fully supports *Connex Micro*.

### 6.8.7 Persistence Service

*Connex Micro* only supports VOLATILE and TRANSIENT\_LOCAL durability and does not support the use of Persistence Service.

### 6.8.8 Application Generation Using XML

An application defined in XML can be shared between *Connex Micro* and *Connex Professional*, with the limitations documented in *Application Generation Using XML*.

## Chapter 7

# API Reference

*RTI Connex Micro* features API support for C and C++. Select the appropriate language below in order to access the corresponding API Reference HTML documentation.

- [C API Reference](#)
- [C++ API Reference](#)



## Chapter 8

# Release Notes

### 8.1 Compatibility

*Connex Micro* 4.2.0 requires the following minimum versions of *Connex* products for compatibility:

- *Connex Professional* 7.3.0 and above<sup>1</sup>
- *Connex Drive* 3.1.0 and above<sup>1</sup>

Refer to *Working with Connex Micro and Connex Professional* and the *API Reference* for more details about interoperable APIs, QoS policies, and *Connex* tools.

### 8.2 Supported Platforms and Programming Languages

*Connex Micro* supports the C and traditional C++ language bindings.

Note that RTI only tests on a subset of the possible combinations of OSs and CPUs. Please refer to the following table for a list of specific platforms and the specific configurations that are tested by RTI.

For more information on the library types (PIL, PSL, and integrated) that RTI provides, refer to *Library descriptions*.

PIL

RTI provides PILs for the platforms listed below. Architecture abbreviations utilize the following format:

<code>{cpu}{compiler}{profile}</code>
---------------------------------------

- `{cpu}`: the CPU that the library was compiled for.
- `{compiler}`: the compiler used to build the library.
- `{profile}`: CERT if the library was built to be Cert-compatible; otherwise empty.

---

<sup>1</sup> *Connex Micro* 4.2.0 is not guaranteed to be compatible with all future releases of *Connex* products. Consult the **Downloads** page on [support.rti.com](http://support.rti.com) for up-to-date compatibility information.

Table 8.1: Supported Platforms (PIL)

OS	CPU	Compiler	RTI Architecture Abbreviations
Windows® 10	x64	VS 2017	x86_64lePEvs2017 x86_64lePEvs2017CERT
macOS® 14	x64	clang 15.0	x86_64leMachOclang15.0 x86_64leMachOclang15.0CERT
macOS® 14	arm64	clang 15.0	armv8leMachOclang15.0 armv8leMachOclang15.0CERT
Ubuntu® 22.04 LTS	x64	gcc 12.3.0	x86_64leElfgcc12.3.0 x86_64leElfgcc12.3.0CERT
Ubuntu® 18.04 LTS	ARMv8 (64-bit)	gcc 7.3.0	armv8leElfgcc7.3.0 armv8leElfgcc7.3.0CERT
Yocto 5.15.96	ARMv8-A (64-bit)	gcc 11.3.1	armv8aleElfgcc11.3.1 armv8aleElfgcc11.3.1CERT
QNX® 7.1	ARMv8 (64-bit)	qcc_gpp8.3.0	armv8leElfqnx_qcc8.3.0 armv8leElfqnx_qcc8.3.0CERT
QOS 2.2.1 (QNX OS for Safety)	ARMv8 (64-bit)	qcc_gpp8.3.0	armv8leElfqcc8.3.0 armv8leElfqcc8.3.0CERT
FreeRTOS® 10.0.0	Armv7E-M	gcc 10.3.1	armv7emleElfgcc10.3.1 armv7emleElfgcc10.3.1CERT

## PSL

RTI provides PSLs for the platforms listed below. Architecture abbreviations utilize the following format:

```
{cpu}{compiler}{profile}-{OS}
```

- `{cpu}`: the CPU that the library was compiled for.
- `{compiler}`: the compiler used to build the library.
- `{profile}`: CERT if the library was built to be Cert-compatible; otherwise empty.
- `{OS}`: The operating system that the PSL was compiled for.

Table 8.2: Supported Platforms (PSL)

OS	CPU	Compiler	RTI Architecture Abbreviations
Windows® 10	x64	VS 2017	x86_64lePEvs2017-Win10 x86_64lePEvs2017CERT-Win10
macOS® 14	x64	clang 15.0	x86_64leMachOclang15.0-Darwin23 x86_64leMa- chOclang15.0CERT-Darwin23
macOS® 14	arm64	clang 15.0	armv8leMachOclang15.0-Darwin23 armv8leMa- chOclang15.0CERT-Darwin23

continues on next page

Table 8.2 – continued from previous page

OS	CPU	Compiler	RTI Architecture Abbreviations
Ubuntu® 22.04 LTS	x64	gcc 12.3.0	x86_64leElfgcc12.3.0-Linux5 x86_64leElfgcc12.3.0CERT-Linux5
Ubuntu® 18.04 LTS	ARMv8 (64-bit)	gcc 7.3.0	armv8leElfgcc7.3.0-Linux4 armv8leElfgcc7.3.0CERT-Linux4
Yocto 5.15.96	ARMv8-A (64-bit)	gcc 11.3.1	armv8aleElfgcc11.3.1-Linux5 armv8aleElfgcc11.3.1CERT-Linux5
QNX® 7.1	ARMv8 (64-bit)	qcc_gpp8.3.0	armv8leElfqnx_qcc8.3.0-QNX7.1 armv8leElfqnx_qcc8.3.0CERT-QNX7.1
QOS 2.2.1 (QNX OS for Safety)	ARMv8 (64-bit)	qcc_gpp8.3.0	armv8leElfqcc8.3.0-QOS2.2.1 armv8leElfqcc8.3.0CERT-QOS2.2.1
FreeRTOS® 10.0.0	Armv7E-M	gcc 10.3.1	armv7em- leElfgcc10.3.1-FreeRTOS10.0 armv7em- leElfgcc10.3.1CERT-FreeRTOS10.0

Integrated

RTI does not provide integrated libraries for this release.

## 8.3 What's New in 4.2.0

*RTI Connex Micro* 4.2.0 is a feature release. See the [Connex Releases](#) page on the RTI website for more information on RTI's software release model.

The following features are new since *Connex Micro* 4.1.0.

### 8.3.1 Reduced core library size

This release separates the Extensible Types (XTypes) module from the core library, `rti_me`, into a separate library, `rti_me_ddsxtypes`. This change reduces the size of `rti_me`. The XTypes module is now optional and can be omitted if your application does not use its functionality.

Refer to the [Migration Guide](#) (if you have internet access) for more information on how this change affects existing *Connex Micro* applications.

### 8.3.2 Secure RTPS communication with the Lightweight Security Plugin

This release introduces the *Lightweight Security Plugin* for *Connex Micro*, a module that enables you to secure your application's RTPS messages with pre-shared key (PSK) encryption. The *Lightweight Security Plugin* includes a Platform Independent Library (PIL) and a Transform Platform Support Library (PSL) for interfacing with [OpenSSL 3.5](#).

For more information on this feature, refer to the *Lightweight Security Plugin* section.

### 8.3.3 Optimize network bandwidth usage with content filtering

**Attention:** This is an experimental feature in *Connext Micro* 4.2.0. It is not guaranteed to be consistent or supported and should not be used in production.

Refer to *Experimental Features* for more information.

This release introduces content filtering to *Connext Micro*, enabling *DataReaders* to filter incoming sample data. This can reduce the network bandwidth usage in your application, since matching *DataWriters* can be configured to only send samples that pass the *DataReader*'s filter.

For more information on this feature, consult the *Content Filtering* section.

### 8.3.4 Improve system robustness with DomainParticipant rediscovery

This release adds the ability for a *DomainParticipant* with its `DomainParticipantQos.participant_name.name` set to a non-empty value to be rediscovered if it is restarted before its `lease_duration` expires. You can enable this feature by setting the `DomainParticipantQos.discovery.enable_participant_discovery_by_name` property to `DDS_BOOLEAN_TRUE`.

For more information on this feature, refer to the *DomainParticipant Discovery by Name* section.

### 8.3.5 Develop secure applications in XML with MAG

This release allows you to enable and configure the *Lightweight Security Plugin* while defining an application in XML format. Micro Application Generator (MAG) will then create the necessary code for *Connext Micro*.

For more details, refer to the *Lightweight Security Plugin* section of *Application Generation Using XML*.

### 8.3.6 Develop applications with content filtering in XML with MAG

This release allows you to enable and configure content filtering while defining an application in XML format. Micro Application Generator (MAG) will then create the necessary code for *Connext Micro*.

For more details, refer to the *Content filtering* section of *Application Generation Using XML*.

### 8.3.7 Enable DomainParticipant rediscovery in XML with MAG

This release allows you to enable *DomainParticipant* rediscovery while defining an application in XML format. Micro Application Generator (MAG) will then create the necessary code for *Connext Micro*.

For more information on this feature, see *DomainParticipant Discovery by Name*. You can enable *DomainParticipant* rediscovery via the `enable_participant_discovery_by_name` field in the [PROPERTY](#) QoS policy, as described in *Discovery Configuration*.

### 8.3.8 Enable endpoint discovery message queueing in XML with MAG

This release allows you to enable endpoint discovery message queueing while defining an application in XML format. Micro Application Generator (MAG) will then create the necessary code for *Connext Micro*.

This feature is included as a fix for a potential discovery issue; see *[Major] Possible mismatched endpoint messages during discovery* for more information. You can enable discovery message queueing via the `enable_endpoint_discovery_queue` field in the [PROPERTY](#) QoS policy, as described in *Discovery Configuration*.

### 8.3.9 Third-party and open source software changes

This release adds the following open source software as part of the *Lightweight Security Plugin*:

Table 8.3: Open Source Software Additions

Software	Version
OpenSSL	3.5.0

For more information, see the *Third-Party and Open Source Software* section.

## 8.4 What's Fixed in 4.2.0

The following are fixes since *Connext Micro* 4.1.0.

[Critical]: System-stopping issue, such as a crash or data loss. [Major]: Significant issue with no easy workaround. [Minor]: Issue that usually has a workaround. [Trivial]: Small issue, such as a typo in a log.

### 8.4.1 Discovery

#### **[Critical] Lost INFO\_DST messages with GUIDPREFIX\_UNKNOWN**

*Connext Micro* may have silently discarded RTPS INFO\_DST messages containing a *guidPrefix* of 0x0 (GUIDPREFIX\_UNKNOWN).

[RTI Issue ID MICRO-8751]

#### **[Major] DomainParticipants stopped sending announcements to initial peer locator after a discovered DomainParticipant went away**

*This issue only affected patch release 4.1.0.1.*

If *DomainParticipant A* discovered *DomainParticipant B* on an initial peer locator, then lost *DomainParticipant B* and attempted to rediscover it, *DomainParticipant A* would stop sending discovery announcements to that initial peer locator.

[RTI Issue ID MICRO-10899]

#### **[Major] Possible mismatched endpoint messages during discovery**

If a *DomainParticipant* receives an endpoint discovery message and lacks sufficient resources to store the endpoint, *Connext Micro* will (by default) acknowledge the message as received, but discard the discovery information. This can lead to a mismatch in discovery states between *DomainParticipants*.

This release adds a new field, `DomainParticipantQos.discovery.enable_endpoint_discovery_queue`, and if this field is set to `DDS_BOOLEAN_TRUE`, endpoint discovery messages will instead be queued for later processing when resources become available.

---

**Note:** This new field assumes that there are sufficient resources available for discovery information and that the lack of resources is temporary (such as during a system restart). Setting this value to `DDS_BOOLEAN_TRUE` without sufficient resources may cause undefined behavior.

---

[RTI Issue ID MICRO-10044]

#### **[Major] DomainParticipants continued to send DATA(p) messages to nonexistent remote participants**

A *DomainParticipant* might have continued to send DATA(p) discovery messages to remote *DomainParticipants* that were once discovered but no longer exist.

[RTI Issue ID MICRO-8747]

**[Major] DataReader failed to acknowledge samples when it received more instances than it could store**

A *DataReader* failed to add key instances or acknowledge samples when it received more instances than its QoS supported.

[RTI Issue ID MICRO-8872]

**[Minor] Participants using DPDE sent one extra announcement message**

Participants using DPDE sent one more participant announcement message than expected for discovery purposes.

[RTI Issue ID MICRO-8257]

**[Minor] Failed to delete a participant before sending at least one participant announcement**

If `DomainParticipantFactory_delete_participant` was called before sending at least one participant announcement, `DomainParticipantFactory_delete_participant` would return a failure.

[RTI Issue ID MICRO-9828]

## 8.4.2 Reliability Protocol and Wire Representation

**[Major] Asynchronous flow controller may have stopped sending data**

A flow controller configured with with an `added_bits_per_period` of `DDS_LENGTH_UNLIMITED` (such as the default flow controller) may have stopped sending data if the controller lost fragments while using reliable communication.

[RTI Issue ID MICRO-10991]

**[Major] Reliable DataReader may have stopped receiving data**

A reliable *DataReader* may have stopped receiving data if the data was received out of order from a *DataWriter* and the *DataWriter* was deleted before the *DataReader* could request the missing data.

[RTI Issue ID MICRO-10317]

### 8.4.3 APIs (C or Traditional C++)

#### **[Major] DDS\_DomainParticipantFactory\_finalize\_instance failed if INTRA transport had been unregistered**

The [DDS\\_DomainParticipantFactory\\_finalize\\_instance](#) function would fail if the INTRA transport had been unregistered previously in the test.

[RTI Issue ID MICRO-4481]

#### **[Minor] Maximum blocking time is limited to 25 days**

The maximum blocking time for APIs that supported blocking was approximately 25 days, even when the documentation stated a longer max blocking time.

[RTI Issue ID MICRO-9770]

### 8.4.4 Generated Code (C, Traditional C++, and Modern C++)

#### **[Major] Failed to serialize a union with only a default label**

A union with only a default label could not be serialized. For example, the following union would fail to be serialized:

```
union Foo switch(int32) {  
    default: long x;  
};
```

[RTI Issue ID MICRO-9079]

### 8.4.5 Hangs

#### **[Critical] Application may have frozen during discovery of a remote participant**

An application may have frozen during discovery of a remote participant when the built-in participant announcement *DataWriter* used asynchronous publishing. This was due to a race condition caused by the flow controller. Now, the built-in participant announcement *DataWriter* will only use synchronous publishing.

[RTI Issue ID MICRO-8901]



### 8.4.6 Memory Leaks/Growth

#### **[Minor] Memory leak in static discovery examples**

The generated examples for C++ using static discovery leaked memory for the topic and type names.

[RTI Issue ID MICRO-10046]

#### **[Major] Resources not freed when calling unregister\_instance on a reliable Zero Copy DataWriter**

When using a reliable Zero Copy *DataWriter*, calling `FooDataWriter_unregister_instance()` would not return the instance's resources to the appropriate pool for reuse.

[RTI Issue ID MICRO-10039]

### 8.4.7 Platform and Build Changes

#### **[Minor] Linker errors occurred when compiling HelloWorld examples for Windows systems**

Compiling the HelloWorld examples for Windows systems in Release mode or with shared libraries resulted in linker errors. The build system will now correctly include the appropriate C runtime libraries and dynamic link libraries (DLLs) based on the Debug/Release configuration and whether static or dynamic linking is used.

[RTI Issue ID MICRO-8701]

### 8.4.8 Shipped Examples

#### **[Trivial] Failure to compile example generated for MAG**

*This issue was fixed in 4.1.0, but not documented at that time.*

When generating an example for Micro Application Generator (MAG), two files (`<IDL>.xml` and `<IDL_Qos>.xml`) were not generated.

[RTI Issue ID MICRO-6801]

## 8.4.9 Interoperability

### **[Major] Zero Copy communication failed when applications were started from different directories**

When two applications were started from different working directories, they would fail to communicate when using the Zero Copy v2 transport.

---

**Note:** As a result of this fix, applications using Zero Copy transfer are not backwards compatible with previous versions of *Connext Micro*.

Refer to the [Migration Guide](#) for more information.

---

[RTI Issue ID MICRO-10391]

## 8.4.10 Other

### **[Critical] Potential loss of queued samples when using asynchronous communication**

If a sample was removed from a *DataWriter*'s history while the sample was queued for asynchronous publication, *Connext Micro* would not properly reacquire the network lock. This could result in losing queued samples while using BEST\_EFFORT communication, or when resending samples with RELIABLE communication.

[RTI Issue ID MICRO-10947]

### **[Critical] Missed DATA\_FRAG messages with multiple fragments**

Under specific conditions, *Connext Micro* incorrectly processed RTPS messages that contained more than one DATA\_FRAG submessage, resulting in missed messages.

[RTI Issue ID MICRO-9018]

### **[Major] Micro Application Generator (MAG) failed to start on macOS platforms**

The `rtiddsmag` script failed to start on macOS platforms if `java` was not available in the `PATH` or `JAVA_HOME`. `rtiddsmag` will now use the JRE available with the host if `java` is not available in the `PATH` or `JAVA_HOME`.

[RTI Issue ID MAG-188]

**[Major] DataWriter with finite timestamp failed to send data if date exceeded January 2038**

A *DataWriter* configured with a finite source timestamp tolerance and `DDS_BY_SOURCE_TIMESTAMP_DESTINATIONORDER_QOS` would fail to send data after 03:14:07 UTC 19 January 2038.

[RTI Issue ID MICRO-5858]

**[Major] Connext Micro failed to parse malformed DATA\_FRAG submessages**

*Connext Micro* could enter an infinite loop when receiving a `DATA_FRAG` submessage with malformed inline QoS parameters.

[RTI Issue ID MICRO-10288]

**[Major] Connext Micro would not work if year exceeded 2038**

If the date was set beyond the year 2038, *Connext Micro* would not work. This was because the date is reported as a 32-bit unsigned integer; however, *Connext Micro* expects a signed 32-bit value and would therefore interpret the “wrap around” value as a negative number, causing an error.

[RTI Issue ID MICRO-2295]

**[Minor] Multiple endpoints did not match properly when using Zero Copy transfer**

When using the Zero Copy v2 transport, multiple endpoints would not match properly if both *DomainParticipants* were enabled before creating *DataReaders* or *DataWriters*.

[RTI Issue ID MICRO-10162]

**[Minor] SYSTEM\_RESOURCE\_LIMITS.max\_components could not be changed**

*Connext Micro* applications could not change the `SYSTEM_RESOURCE_LIMITS.max_components`.

[RTI Issue ID MICRO-3571]

**[Trivial] rtime-make did not support building multiple targets on Windows systems and failed to report some CMake errors**

The `rtime-make` batch script would fail to build additional targets after the first `--` parameter in the command line on Windows systems. Additionally, the `rtime-make` batch and shell scripts failed to report some CMake errors correctly.

[RTI Issue ID MICRO-10913]

## 8.5 Previous Releases

### 8.5.1 What's New in 4.1.0

*RTI Connext Micro* 4.1.0 is a feature release. See the [Connext Releases](#) page on the RTI website for more information on RTI's software release model.

The following features are new since *Connext Micro* 4.0.1.

#### Platform-independent code is now separate from OS and network stack integration

This release includes precompiled *Connext Micro* binaries in two formats:

- Platform Independent Libraries (PIL): binaries that support the basic features of *Connext Micro*.
- Platform Support Libraries (PSL): binaries that support OS and network stack integration.

This split allows for different platform-specific PSLs to be written for the same PIL without needing to recompile the *Connext Micro* code. Previous releases of *Connext Micro* were delivered as integrated libraries with the OS and network stack code included, which could not be changed without recompiling the entire library. See the *Library descriptions* section for more information on this change.

Since these split libraries are provided precompiled, *Connext Micro* can be used out-of-the-box without building the source code first; see *Getting Started* and *Developing Applications*. However, RTI also provides the source code for the PSL with supported architectures. Refer to *Building the PSL* for instructions on how to build your own PSL from source.

**Warning:** Split libraries **do not** automatically register the UDP transport. This is because *Connext Micro* makes no assumptions about which transports are available when using split libraries. In order to use the UDP transport included in the PSL, add the following code before creating a *DomainParticipant*:

```
/* Register the UDP transport */
struct UDP_InterfaceFactoryProperty *udp_property;

/* allocate and set udp_property */

RT_Registry_register(registry, NETIO_DEFAULT_UDP_NAME,
    UDP_InterfaceFactory_get_interface(),
    (struct RT_ComponentFactoryProperty*)udp_property, NULL),

DDS_StringSeq_set_maximum(&participant_qos.transports.enabled_transports, 1);
DDS_StringSeq_set_length(&participant_qos.transports.enabled_transports, 1);
*DDS_StringSeq_get_reference(&participant_qos.transports.enabled_transports, 0) = DDS_
↪String_dup("_udp");

DDS_StringSeq_set_maximum(&participant_qos.discovery.enabled_transports, 1);
DDS_StringSeq_set_length(&participant_qos.discovery.enabled_transports, 1);
```

```

*DDS_StringSeq_get_reference(&participant_qos.discovery.enabled_transports,0) = DDS_
↳String_dup("_udp://");

DDS_StringSeq_set_maximum(&participant_qos.user_traffic.enabled_transports,1);
DDS_StringSeq_set_length(&participant_qos.user_traffic.enabled_transports,1);
*DDS_StringSeq_get_reference(&participant_qos.user_traffic.enabled_transports,0)
= DDS_String_dup("_udp://");

```

## Transfer large data samples quickly with Zero Copy v2

This release adds a new transport, Zero Copy v2, which can perform Zero Copy data transfer. Zero Copy transfer allows you to move large data samples without copying them, which increases throughput and reduces latency.

You can now set up Zero Copy transfer to use either the Shared Memory Transport (which was available in previous versions of *Connext Micro*) or the new Zero Copy v2 transport. The main difference between the two transports is that the Shared Memory Transport is interoperable with *Connext Professional* and the Zero Copy v2 transport is interoperable with select versions of *Connext Cert*.

For more details, refer to *Zero Copy Transfer* and *Zero Copy v2 Transport*.

## Enable and configure Zero Copy transfer with MAG

This release allows you to enable and configure the *Zero Copy v2 Transport* while defining an application in XML format. Micro Application Generator (MAG) will then create the necessary code to enable and configure the Zero Copy transport in *Connext Micro*.

For details on how to enable Zero Copy v2 with MAG, refer to *MAG Command-Line Options* and *Transport and Discovery Configuration*.

## Enhance data reliability by detecting and discarding corrupted RTPS messages

This release adds support for detecting and discarding corrupted RTPS messages. This improves data reliability and provides basic security by ensuring that the data has not been modified in transit.

A Cyclic Redundancy Check (CRC) is computed over the DDS RTPS message (including the RTPS Header), which is sent as a new RTPS submessage. The subscribing application can detect this new submessage and validate the contained CRC. Optionally, when a corrupted RTPS message is detected, the message can be dropped.

To enable the use of CRC in a *DomainParticipant*, there are three new fields in the [WireProtocolQoSPolicy](#):

- `compute_crc`: when enabled at the sending application, sends the CRC.
- `check_crc`: when enabled, drops corrupted messages.

- `require_crc`: when enabled, ignores participants with `compute_crc` set to false.

Refer to *Message Integrity Checking* for details.

### Develop more reliable applications with MAG

This release adds support to Micro Application Generator (MAG) for the Cyclic Redundancy Check feature (see *Enhance data reliability by detecting and discarding corrupted RTPS messages*). This allows you to develop applications with MAG that have improved data reliability and basic security.

MAG now supports the following fields in the [WireProtocolQosPolicy](#):

- `compute_crc`
- `check_crc`
- `require_crc`
- `computed_crc_kind`
- `allowed_crc_mask`

Refer to *Message Integrity Checking* for details.

### Guarantee compatibility with Connex Professional with MAG when using the Shared Memory Transport

This release adds support to Micro Application Generator (MAG) for the `dds.transport.minimum_compatibility_version` property, which you can set via the [PROPERTY](#) QoS policy for the *DomainParticipant*.

`dds.transport.minimum_compatibility_version` changes the value of the new field [pro\\_minimum\\_compatibility\\_version](#) that has been added to the shared memory interface factory property. This property sets the minimum version of *Connex Professional* to be compatible with when using shared memory.

The default value for this field is `DDS_PRODUCTVERSION_UNKNOWN`.

Refer to *SHMEM Configuration* for details.

### Improve control of data distribution to multicast addresses with new UDP transport options

This release adds the following new options to the UDP transport to further control how *Connex Micro* sends data to multicast addresses:

- [disable\\_multicast\\_bind](#): controls whether *Connex Micro* will bind to a multicast address receive address (if set to 0) or bind to ANY multicast address (if set to 1).
- [multicast\\_loopback\\_disable](#): controls whether *Connex Micro* puts multicast packets onto the loopback interface.

- `disable_multicast_interface_select`: controls whether *Connex Micro* will use `multicast_interface` or `allow_interface/deny_interface` to select the interface for sending to multicast addresses.

Refer to *UDP Transport* for more information on these options.

### Develop applications with new UDP transport options with MAG

This release adds support to Micro Application Generator (MAG) for the following new fields when configuring the UDP transport:

- `disable_multicast_bind`
- `multicast_loopback_disable`
- `disable_multicast_interface_select`

Refer to *UDP Configuration* for more information on these properties.

### Build Connex Micro libraries conveniently with symlinks

This release adds support for using the `cmake --install` command with symbolic links (symlinks) as well as full paths.

Refer to *Building Connex Micro for Common Platforms* for more information on how to use CMake build commands.

## 8.5.2 What's Fixed in 4.1.0

The following are fixes since *Connex Micro* 4.0.1.

[Critical]: System-stopping issue, such as a crash or data loss. [Major]: Significant issue with no easy workaround. [Minor]: Issue that usually has a workaround. [Trivial]: Small issue, such as a typo in a log.

### Discovery

#### [Major] Participants did not perform discovery correctly when `on_data_on_readers` callback was set

When the `on_data_on_readers` callback was enabled on participants, the participants would not perform discovery correctly.

[RTI Issue ID MICRO-7412]

**[Minor] Incorrect data fragmentation of discovery messages for DomainParticipant built-in topic**

*Connext Micro* incorrectly fragmented *DomainParticipant* discovery messages if the Maximum Transmission Unit (MTU) was set to a value lower than the *DomainParticipant* built-in data. Fragmentation of the *DomainParticipant* built-in topic is not supported. The system now ensures that the creation of a *DomainParticipant* will fail if the minimum MTU across all transports used for discovery is less than the size of the *DomainParticipant* discovery message.

[RTI Issue ID MICRO-7545]

**Usability****[Minor] MAG failed to generate warnings for certain unsupported QoS**

*Micro Application Generator* (MAG) did not generate warnings for certain unsupported QoS values in XML configurations. This could lead to runtime errors when creating participants with incorrect configurations.

[RTI Issue ID MAG-176]

**Transports****[Major] Multicast sockets always bound to the multicast address on non-Windows platforms**

*Connext Micro* always bound multicast receive sockets to the multicast address upon creation, even if [enable\\_interface\\_bind](#) had not been set. This only occurred on non-Windows platforms.

Now, [disable\\_multicast\\_bind](#) can optionally disable binding a multicast receive socket to its multicast receive address.

See *Improve control of data distribution to multicast addresses with new UDP transport options* for more information on [disable\\_multicast\\_bind](#).

---

**Note:** On some platforms, you may need to set [disable\\_multicast\\_bind](#) to 1 to interoperate *Connext Micro* with *Connext Professional* because *Connext Professional* does not bind multicast receive sockets to the multicast address.

---

[RTI Issue ID MICRO-8451]



## Reliability Protocol and Wire Representation

### [Critical] Failed to deliver samples when prior samples were lost

*Connext Micro* would sometimes fail to deliver samples to a *DataReader* if the last processed sample was a HEARTBEAT message which indicated lost samples.

[RTI Issue ID MICRO-7317]

### [Major] Delayed or failed to deliver large samples on unreliable networks

The [RELIABILITY](#) protocol may have struggled to repair lost data fragments. In unreliable network conditions, fragmented samples of large data could be delayed or fail to deliver at all if the samples were written faster than they could be repaired. Now, the [RELIABILITY](#) protocol for lost fragments has been improved so that fragmented samples are repaired quickly on an unreliable network.

[RTI Issue ID MICRO-8255]

### [Minor] Piggyback heartbeats were not included when using asynchronous publication

*DataWriters* did not include piggyback heartbeats with non-fragmented samples when using asynchronous publication. They are now included at the rate configured by [DDS\\_RtpsReliableWriter-Protocol\\_t::heartbeats\\_per\\_max\\_samples](#), regardless of the publication mode.

[RTI Issue ID MICRO-8173]

## APIs (C or Traditional C++)

### [Major] DDS\_DomainParticipantFactory\_delete\_participant was not thread-safe

[DDS\\_DomainParticipantFactory\\_delete\\_participant](#) was not thread-safe. A race condition could occur if multiple *DomainParticipants* were deleted at the same time in the same process space.

[RTI Issue ID MICRO-8055]

### [Minor] C++ constructor did not allocate memory for Topic and Type names

The C++ constructor for [DDS\\_PublicationBuiltinTopicData](#) and [DDS\\_SubscriptionBuiltinTopicData](#) did not allocate memory for the **topic\_name** and **type\_name** attributes.

Please refer to the API Reference for more information on these APIs.

[RTI Issue ID MICRO-7110]

**[Trivial] C++ API Reference contained incorrect «cert» references**

The C++ API Reference contained <<cert>> references, but *Connext Cert* does not support the C++ API.

[RTI Issue MICRO-3216]

**[Trivial] Non-descriptive API failure messages**

The following APIs have updated failure messages:

- [FooDataWriter\\_get\\_loan](#) failed with `UNSUPPORTED` if your type was annotated with `@transfer_mode(SHMEM_REF)` and no transports that support Zero Copy transfer were enabled. It now fails with `PRECONDITION_NOT_MET`.
- [FooDataWriter\\_discard\\_loan](#) failed with `UNSUPPORTED` if your type was annotated with `@transfer_mode(SHMEM_REF)` and no transports that support Zero Copy transfer were enabled. It now fails with `PRECONDITION_NOT_MET`.
- [FooDataWriter\\_get\\_loan](#) failed with `ERROR` if all samples had been loaned to your application. It now fails with `OUT_OF_RESOURCES`.

[RTI Issue ID MICRO-7929]

**XML Configuration****[Minor] Invalid code when using a flow controller**

*Micro Application Generator* (MAG) generated some invalid code when using a flow controller in the following scenarios:

- Trying to use one of the built-in flow controllers.
- `SYNCHRONOUS_PUBLISH_MODE_QOS` was configured.

[RTI Issue ID MAG-174]

**Crashes****[Critical] Segmentation fault occurred when an asynchronous publisher tried to send a missing data fragment**

When an asynchronous publisher tried to send a data fragment that was no longer available, a segmentation fault occurred.

[RTI Issue ID MICRO-7982]

**[Critical] Segmentation fault occurred when creating or finalizing a *DataWriter* or *DataReader* while using asynchronous publication**

A segmentation fault could have occurred when creating or finalizing a *DataWriter* or *DataReader* if the built-in writers were using asynchronous publication. The built-in writers would automatically use asynchronous publication if the `max_message_size` property of the UDP transport was set to a low value (~1400). This issue only occurred if the *DataWriter* or *DataReader* was created or finalized while the *DomainParticipant* was matched with another *DomainParticipant*.

[RTI Issue ID MICRO-8111]

**[Critical] Potential race condition and crash occurred when a *DataWriter* waited for resources**

A race condition could occur when a *DataWriter* waited for resources to be freed while sending data asynchronously. This could cause a crash, since the *DataWriter* state was incorrectly updated.

[RTI Issue ID MICRO-8001]

**[Critical] Potential race condition and crash occurred when a *DataWriter* unmatched with a remote entity**

A race condition could occur when a *DataWriter* using asynchronous publishing unmatched with a remote entity while data was being sent to the same entity. This could cause a crash.

[RTI Issue ID MICRO-8170]

**[Minor] Integer overflow when setting MTU lower than 448 bytes**

*Connext Micro* could produce errors if the Maximum Transmission Unit (MTU) for any transport was set lower than 448 bytes. *Connext Micro* now ensures that the MTU across all enabled transports is greater than 448 bytes. If the MTU requirement is not met, entity creation will fail.

[RTI Issue ID MICRO-7530]

**Hangs****[Critical] Ungracefully terminated QNX processes using SHMEM transport prevented startup of new processes due to unclosed POSIX semaphores**

If a QNX application using the shared-memory transport was ungracefully shut down, crashed, or otherwise had an abnormal termination while holding a POSIX semaphore used by the transport (for example, while sending data through the shared-memory transport), *Connext* applications launched after that point on the same domain may have waited forever for that semaphore to be released.

Now on QNX 7.1 and greater, the usage of POSIX semaphores has been replaced with robust pthread mutexes. Abnormal termination of an application while holding a mutex will no longer result in a *Connext* application launched after that point hanging.

As a result, *Connext Micro* 4.1.0 will **not** be backward compatible with previous versions of *Connext Micro* when using the shared memory transport on QNX 7.1 and greater. *Connext Micro* 4.1.0 will now be compatible with *Connext Professional* 7.3.0 on QNX 7.1 and greater.

[RTI Issue ID MICRO-6013]

## Memory Leaks/Growth

### [Minor] Failed to cleanup resources when DomainParticipant creation failed

If the creation of a *DomainParticipant* failed, the allocated resources were not correctly cleaned up.

[RTI Issue ID MICRO-6500]

## Data Corruption

### [Critical] Potential data corruption from race condition when using KEEP\_LAST and publishing asynchronously

Data samples may have been corrupted if both of the following were true:

- A race condition occurred while sending fragmented samples.
- The fragmented samples were sent while a `HISTORY.kind` of `DDS_KEEP_LAST_HISTORY_QOS` caused samples to be removed and reused.

[RTI Issue ID MICRO-8314]

## Interoperability

### [Critical] Incorrect handling of RTPS messages with submessages from different participants

When an RTPS message that contained submessages from multiple participants was received, *Connext Micro* incorrectly treated each submessage as though it was from the participant whose GUID prefix was in the RTPS header. *Connext Micro* does not send RTPS messages with submessages from different participants, but other DDS vendors may do this, which would have led to various communication issues and a lack of interoperability.

[RTI Issue ID MICRO-5984]

## Other

### **[Critical] Lost samples or fragments not repaired when using a flow controller**

When using a flow controller, repair packets may not have been sent by a *DataWriter* to a *DataReader* to replace lost fragments or samples. The writer would only send repair packets when a new sample was written, resulting in repairs never being sent if another sample was not written. Repair packets will now be sent by the flow controller as soon as they are required.

[RTI Issue ID MICRO-7403]

### **[Major] Enabling asynchronous publication on DataWriter caused RTPS messages to fail**

If the `publish_mode.kind` of a *DataWriter*'s QoS was configured to be `ASYNCHRONOUS_PUBLISH_MODE_QOS`, RTPS messages from that *DataWriter* may have failed to send. This condition occurred specifically if asynchronous publication was enabled AND the sample type did not require RTPS fragmentation.

[RTI Issue ID MICRO-7219]

### **[Major] Finalizing a participant might have failed when using DPSE**

`DDS_DomainParticipant_finalize()` may have failed if the participant was using the DPSE discovery plugin. This issue was most likely to occur on macOS® platforms.

[RTI Issue ID MICRO-7870]

### **[Major] Samples with meta-information were not delivered to the user if they arrived when history cache was full**

When a *DataReader*'s history cache was full, samples containing meta-information from matched *DataWriters* were not delivered to the user.

[RTI Issue ID MICRO-8063]

### **[Minor] Flow controllers incorrectly delayed sending packets**

The token-bucket flow controller may have waited to send packets until its next period, even if more tokens were available during the current period.

[RTI Issue ID MICRO-8024]

**[Minor] DataWriters with KEEP\_ALL History may not have sent all historical samples**

*DataWriters* with [History kind](#) set to DDS\_KEEP\_ALL\_HISTORY\_QOS and [Durability kind](#) set to DDS\_TRANSIENT\_LOCAL\_DURABILITY\_QOS would send historical samples only up to the [History depth](#). *DataWriters* now send historical samples up to [max\\_samples\\_per\\_instance](#).

[RTI Issue ID MICRO-8505]

**[Trivial] Illegal reflective access warning when running MAG with OpenJDK™ 11**

*This issue was fixed in 4.0.1, but not documented at that time.*

Running Micro Application Generator (MAG) with OpenJDK 11 generated the following warning:

```
WARNING: An illegal reflective access operation has occurred

WARNING: Illegal reflective access by com.rti.micro.appgen.utils.QosUtils (file:/.../rti_
↪connexrt_dds_micro-4.0.0/rtiddsmag/class/rtiddsmag.jar) to field java.lang.String.value

WARNING: Please consider reporting this to the maintainers of com.rti.micro.appgen.utils.
↪QosUtils

WARNING: Use --illegal-access=warn to enable warnings of further illegal reflective
↪access operations

WARNING: All illegal access operations will be denied in a future release
```

MAG has been updated to use OpenJDK 17, which does not generate this warning.

[RTI Issue ID MAG-172]

**[Trivial] Incorrect example instructions in the User's Manual**

*Examining the application* in the User's Manual stated incorrectly that the files for the example were included with *Connext Micro*. The instructions have been updated to reflect that the files are generated with *rtiddsmag*.

[RTI Issue ID MICRO-7376]

**[Trivial] Micro transformation example failed to compile**

*Connext Micro* failed to compile with the Visual Studio® 2010 solution files included in the transformation example. The build process has been updated to compile with a new CMake file.

[RTI Issue ID MICRO-7171]

### 8.5.3 What's New in 4.0.1

*RTI Connext Micro* 4.0.1 is an Early Access Release, based on release 4.0.0.

The following features are new since *Connext Micro* 4.0.0.

#### Enable or disable padding bits with PROPERTY QoS policy in DomainParticipant and Data Writer

Micro Application Generator (MAG) adds support for enabling and disabling sending padding bits. This feature is part of a fix for a data corruption issue; see *[Critical] DataReader on a Topic with appendable type could receive samples with incorrect value* for more information.

Padding bits can be set with the following property, via the [PROPERTY](#) QoS policy:

- [dds.xtypes.compliance\\_mask](#), to enable or disable padding bits for the *DomainParticipant* or *DataWriter*. The only valid values supported by MAG for this property are 0 and 0x00000008. MAG will report an error if a different value is set.

---

**Note:** In previous releases, MAG ignored [PROPERTY](#) QoS values, but now it parses all [PROPERTY](#) QoS values configured in XML and adds those values when generating the code. However, it ignores every [PROPERTY](#) QoS property that is not [dds.xtypes.compliance\\_mask](#). Future *Connext Micro* releases may add support for additional properties.

---

#### Generate examples with new template options for Code Generator

In this release, some examples that were previously included in a *Connext Micro* installation have been removed. Instead, examples can be generated from templates included with the *RTI Code Generator*.

This release introduces two new *Code Generator* command-line options, `-showTemplates` and `-exampleTemplates`.

The `-showTemplates` option prints and generates an XML file containing a list of available example templates in your *Connext Micro* installation, organized per language.

The `-exampleTemplate` option generates an example you specify, instead of the default one.

When you use the `-exampleTemplate` option, you can specify one of the example templates in `$RTIMEHOME/rtiddsgen/resource/templates/example/<language>/<templateName>/. You may also create your own templates and place them in this directory.`

You must also use one of the following command-line options:

- `-create examplefiles`
- `-update examplefiles`
- `-example`

*Code Generator* will then generate the example you specified. For example:

```
rtiddsgen -language C++ -example -exampleTemplate <exampleTemplateName> foo.idl
```

For more information, please refer to *Example Applications*.

### 8.5.4 What's Fixed in 4.0.1

The following are fixes since *Connex Micro* 4.0.0.

[Critical]: System-stopping issue, such as a crash or data loss. [Major]: Significant issue with no easy workaround. [Minor]: Issue that usually has a workaround. [Trivial]: Small issue, such as a typo in a log.

#### Discovery

##### [Trivial] Possible error message during discovery with Connex Professional

The following log message may have been printed during discovery with a *Connex Professional* application:

```
[1712165262.572102999]ERROR: ModuleID=4 Errcode=27 X=1 E=0 T=1 netio/NETIOPacket.c:87/  
↪NETIO_Packet_set_head: delta=20
```

This message was benign and did not indicate any failures.

[RTI Issue ID MICRO-6594]

#### Usability

##### [Minor] rtiddsgen failed to run if the default shell was not bash compatible

If the default shell on a macOS or Linux system was not bash (e.g., tcsh), *rtiddsgen* would fail to execute.

[RTI Issue ID MICRO-6539]

##### [Trivial] Self toolchain file missing from the Connex Micro bundle

The `self.tc` toolchain file referred to in the documentation was missing from the *Connex Micro* bundle.

[RTI Issue ID MICRO-6536]



**[Trivial] Empty README.txt generated for an example**

When generating an example using the `-example` option for *rtiddsgen*, the generated README.txt file was empty.

[RTI Issue ID MICRO-6642]

**APIs (C or Traditional C++)****[Minor] Unexpected behavior when copying a DDS\_UnsignedShortSeq with 0 length**

When copying a DDS\_UnsignedShortSeq with 0 length, the destination sequence length was not set to 0.

[RTI Issue ID MICRO-2756]

**[Minor] Missing C++ APIs for discovery operations**

The following functions were missing from the C++ API:

- [get\\_discovered\\_participants](#)
- [get\\_discovered\\_participant\\_data](#)
- [get\\_matched\\_subscriptions](#)
- [get\\_matched\\_subscription\\_data](#)
- [get\\_matched\\_publications](#)
- [get\\_matched\\_publication\\_data](#)

For more information on these functions, please refer to the [C++ API Reference](#).

[RTI Issue ID MICRO-6462]

**[Trivial] C++ examples used the undocumented get\_reference API**

The C++ examples used the undocumented `get_reference` API. C++ examples now use the `[]` operator.

[RTI Issue ID MICRO-3104]

## Generated Code (C, Traditional C++, and Modern C++)

### [Major] Incorrect generated code when using IDL whose name starts with a number

The generated code for an IDL whose name started with a number was incorrect and did not compile. The generated code contained some `ifdef` instructions that started with a number, which was not valid because an identifier must start with a letter (or underscore).

Now, invalid identifier characters are converted to `'_'` in the `ifdef` instruction.

[RTI Issue ID MICRO-2066]

### [Major] Code generated for a `FLAT_DATA` type failed to compile when using namespace option

Code generated for a `FLAT_DATA` type failed to compile when using the `-namespace` option to run *rtiddsgen*.

[RTI Issue ID MICRO-6788]

### [Major] Code generated for an aliased sequence of an aliased string failed to compile

The code generated for an aliased sequence of an aliased string failed to compile. For example, the following IDL would fail:

```
typedef string<2> MyString10;
typedef sequence<MyString10,4> MyStringSeq10;

struct SequenceType4 {
    string<64> msg;
    MyString msg2;
    MyStringSeq10 seq;
};
```

[RTI Issue ID MICRO-6824]

## Crashes

### [Minor] Potential segmentation fault while creating entities

A segmentation fault could occur while creating certain entities if *Connext Micro* ran out of memory. *Connext Micro* will now detect this condition and return an error.

This issue only affected non-CERT profiles.

[RTI Issue ID MICRO-3396]

## Data Corruption

### [Critical] `DataReader` on a Topic with appendable type could receive samples with incorrect value

A *DataReader* subscribing to a *Topic* on an appendable type may have received incorrect samples from a matching *DataWriter*.

The problem only occurred when the *DataWriter* published a type with fewer members than the *DataReader* type. For example, consider a *DataWriter* on `FooBase` and a *DataReader* on `FooDerived`:

```
@appendable struct FooBase {
    sequence<uint8,1024>base_value;
};

@appendable struct FooDerived {
    sequence<uint8,1024> base_value;
    @default(12) uint8 derived_value;
};
```

When the *DataWriter* published a sample with type `FooBase`, in some cases the *DataReader* received a sample in which the field `derived_value` was set to 0 instead of 12.

This issue was caused by a bug in which *Connext* did not set the padding bits in the encapsulation header for a serialized sample as required by the [OMG ‘Extensible and Dynamic Topic Types for DDS’ specification, version 1.3](#). As a result, some of the padding bytes were interpreted as data.

---

**Note:** This fix may lead to a compatibility issue causing a *Connext Micro DataWriter* to not match with a *Connext Micro* or *Connext Cert DataReader*.

For more information, see [Extensible Types Compliance Mask](#) in the *Core Libraries Extensible Types Guide* if you have Internet access.

Padding bits can be disabled with the [dds.xtypes.compliance\\_mask](#) property for backwards compatibility with the following releases:

- *Connext Micro* 2.4.12 and earlier
- *Connext Micro* 2.4.13.2-5
- *Connext Micro* 2.4.14 and 2.4.14.1
- *Connext Cert* 2.4.12.1
- *Connext Cert* 2.4.13.1
- *Connext Cert* 2.4.15.1
- *Connext Micro* 3
- *Connext Micro* 4.0.0

[RTI Issue ID MICRO-5930]

### **[Critical] Undefined behavior using XCDR2 with keyed topic types with key union members**

Using XCDR encoding version 2 (XCDR2) with keyed topic types with key union members was not supported. For example:

```
union MyUnion switch(long) {
    case 0:
        long m_long;
    case 1:
        short m_short;
};

struct StructWithUnionKey {
    @key MyUnion m_union;
    long m_long;
};
```

The behavior was undefined if any of your topic types had a union key member. The results varied, from a potential segmentation fault to an incorrect key hash in which two instances were considered equal.

[RTI Issue ID MICRO-5933]

### **[Critical] Incorrect keyhash generated when receiving data without a keyhash from a node with different endianness**

A *DataReader* would generate an incorrect keyhash for a received sample if all of the following were true:

- The *DataReader* did not receive a keyhash in a sample for a keyed type.
- The *DataReader* used the FLAT\_DATA language binding.
- The sample was sent from a node with different endianness than the *DataReader*.

Also, a *DataReader* would generate an incorrect keyhash for a DISPOSE or UNREGISTER sample if:

- The *DataReader* did not receive a keyhash in a DISPOSE or UNREGISTER sample for a keyed type.
- The *DataReader* used IDL compiled with `-interpreted 1` (the default).
- The DISPOSE or UNREGISTER sample was sent from a node with a different endianness than the *DataReader*.

---

**Note:** Both *Connex Micro* and *Connex Professional* send a keyhash by default, but *Connex Professional* can be configured to send without a keyhash.

---

[RTI Issue ID MICRO-6870]

## Interoperability

### [Major] Incorrect deserialization of CDR encapsulation padding bit

The number of padding bytes in a sample were de-serialized incorrectly. This resulted in samples being dropped if the number of padding bytes was not zero.

[RTI Issue ID MICRO-6799]

## 8.5.5 What's New in 4.0.0

*RTI Connex Micro* 4.0.0 is an Engineering Release, based on release 3.0.3.

The following features are new since *Connex Micro* 3.0.3.

### Enhanced performance for asynchronous DataWriters

This release reduces the contention between *DataWriters* and the asynchronous publication thread (used for flow-control of samples and publishing fragmented samples). Previously, *DataWriters* would block while the asynchronous publication thread was sending data. In this release, the asynchronous publication thread uses a separate critical section from the *DataWriter*'s write API, which allows the *DataWriter* to write samples while the asynchronous publication thread is sending data.

Please note the following limitations:

- It is not possible to send and receive data at the same time.
- The asynchronous publication thread and the *DataWriter* will contend for the same critical section when the asynchronous publication thread starts or finishes sending a sample. This is because the *DataWriter* is loaning samples to the asynchronous publication thread instead of copying them, and the ownership transfer of samples from the *DataWriter* to the asynchronous publication thread (and from the asynchronous publication thread to the *DataWriter*) is protected.

In addition, The User's Manual has not been updated for this release; some sections do not reflect the impact of these changes. Specifically, note the following:

- Each *DataWriter* allocates 1 additional mutex.
- Each *DomainParticipant* allocates 2 additional mutexes, plus 1 mutex per flow-controller (3 by default).
- Each *DataWriter* allocates an additional  $(\text{max\_routes\_per\_reader} * \text{max\_fragmented\_samples} * \text{max\_remote\_readers} * 464)$  bytes. Future releases may reduce this.

### Further control which entities communicate with each other using new Partition QoS policy

The PARTITION QoS policy provides a method to prevent *Entities* that have otherwise compatible QoS policies from matching—and thus communicating with—each other. Much in the same way that only applications within the same DDS domain will communicate with each other, only *Entities* that belong to the same partition can talk to each other.

See information on *Partitions* in the User's Manual chapter for more information.

### Store additional entity-related information that is passed between applications during discovery using new User/Topic/Group Data QoS policies

*Connext Micro* now provides areas where your application can store additional information related to DDS *Entities*. How this information is used is up to user code. *Connext Micro* distributes this information to other applications as part of the discovery process; however, *Connext Micro* does not interpret the information. Use cases are usually application-to-application identification, authentication, authorization, and encryption.

There are three User Discovery Data QoS policies:

- USER\_DATA: associated with *DomainParticipants*, *DataWriters*, and *DataReaders*.
- TOPIC\_DATA: associated with *Topics*.
- GROUP\_DATA: associated with *Publishers* and *Subscribers*.

See information on *User Discovery Data* in the User's Manual chapter for more information.

### Verify that locally created participant GUIDs are unique within a DomainParticipantFactory

When a *DomainParticipant* is created, *Connext Micro* now checks that the GUID is not already in use by another *DomainParticipant* created from the same *DomainParticipantFactory*.

## Micro Application Generator (MAG)

### Support for Partition QoS policy in MAG

Micro Application Generator (MAG) now supports the PARTITION QoS policy. Instead of ignoring the Partition QoS values, as it did in previous releases, MAG now parses the values configured in XML and adds those values when generating the code.

The following partition-related *DomainParticipant* QoS resource limits are also now supported:

- **max\_partitions**
- **max\_partition\_cumulative\_characters**
- **max\_partition\_string\_size**
- **max\_partition\_string\_allocation**

See the *Partitions* chapter in the User's Manual for more information on this QoS policy.

### Support for GROUP\_DATA, USER\_DATA, and TOPIC\_DATA QoS policies in MAG

Micro Application Generator (MAG) now supports the GROUP\_DATA, USER\_DATA, and TOPIC\_DATA QoS policies. Instead of ignoring these QoS values, as it did in previous releases, MAG now parses the values configured in XML and adds those values when generating the code.

MAG also supports the group\_data, user\_data, and topic\_data elements:

- **user\_data** in the *DomainParticipant*, *DataWriter*, and *DataReader* QoS
- **topic\_data** in the *Topic* QoS
- **group\_data** in the *Publisher* and *Subscriber* QoS
- The following *DomainParticipant* QoS resource limits:
  - participant\_user\_data\_max\_length
  - participant\_user\_data\_max\_count
  - topic\_data\_max\_length
  - topic\_data\_max\_count
  - publisher\_group\_data\_max\_length
  - publisher\_group\_data\_max\_count
  - subscriber\_group\_data\_max\_length
  - subscriber\_group\_data\_max\_count
  - writer\_user\_data\_max\_length
  - writer\_user\_data\_max\_count
  - reader\_user\_data\_max\_length
  - reader\_user\_data\_max\_count

See the *User Discovery Data* chapter in the User's Manual for more information on these QoS policies.

### Support for environment variable expansion in MAG

Now you can refer to an environment variable set in the command shell within an XML tag. When MAG parses the configuration file, it will expand the environment variable. The way to refer to the environment variable is as follows:

`$(MY_VARIABLE)`

For example:

```
<name>$(MY_VARIABLE)</name>
```

Being able to refer to an environment variable within an XML file increases XML reusability. For example, this will allow you to specify the initial peers, so you do not need to use multiple XML files or XML profiles per application.

### Only check for QoS policies that are used by your system definition

In previous releases, MAG checked whether all of the QoS policies passed to the tool were supported by *Connex Micro*. This has been changed to only check for QoS policies that are used by the system defined in the `<domain_participant_library>`.

### XML fields of type duration have unset tags default to 0 with a warning log message

The duration type tag has two subfields, `<sec>` and `<nanosec>`. Some QoS policies that use these fields, such as the DEADLINE QoS Policy, set the default duration to INFINITE. Therefore, if you had set just one of these fields (such as `<sec>`, but not `<nanosec>`, or vice-versa), the resulting duration value was still INFINITE.

Now if you set only one of these fields (`<sec>` or `<nanosec>`) in the XML file, the other value defaults to 0. (If you set neither one of them, the default duration for that policy would be used.) A warning message will also be logged by the parser specifying the parent tag, the missing subfield, and the line number.

### Support for resource limits in DomainParticipantFactoryQos

This release allows you to configure the resource limits of the `DomainParticipantFactoryQos` (`max_participants`) in XML.

By default, MAG updates the resource limits of the `DomainParticipantFactoryQos` so that MAG can at least support the entities defined in the XML file. However, if your applications communicate with more remote entities than those specified in the XML file, you may need to manually update the resource limits. In that case, you need to use the `-dontUpdateResourceLimits` command-line option. That will prevent MAG from automatically updating the resource limits for the *DomainParticipantFactory*, *DomainParticipants*, *DataReaders*, and *DataWriters*.

### Instance replacement changes affect XML files in MAG

The type used by `<instance_replacement>` in MAG has been changed from a single type to a complex type. Because of this change, XML files used by MAG in previous releases won't work out of the box in this release. For example, the following XML based on MAG in previous releases won't work in the current release:



```
<datareader_qos>
  <reader_resource_limits>
    <instance_replacement>OLDEST_INSTANCE_REPLACEMENT</instance_replacement>
  </reader_resource_limits>
</datareader_qos>
```

You need to update it to the following:

```
<datareader_qos>
  <reader_resource_limits>
    <instance_replacement>
      <alive_instance_removal>ANY_INSTANCE_REMOVAL</alive_instance_removal>
      <disposed_instance_removal>ANY_INSTANCE_REMOVAL</disposed_instance_removal>
      <no_writers_instance_removal>ANY_INSTANCE_REMOVAL</no_writers_instance_
↪removal>
    </instance_replacement>
  </reader_resource_limits>
</datareader_qos>
```

## 8.5.6 What's Fixed in 4.0.0

The following are fixes since *Connext Micro* 3.0.3.

[Critical]: System-stopping issue, such as a crash or data loss. [Major]: Significant issue with no easy workaround. [Minor]: Issue that usually has a workaround. [Trivial]: Small issue, such as a typo in a log.

### Discovery

#### [Critical] Failure to interoperate with other DDS implementations if default multicast locator specified

*Connext Micro* did not interoperate with other DDS implementations when the default multicast locator was specified.

[RTI Issue ID MICRO-5148]

#### [Major] Incorrect `lease_duration` may have been used for a discovered participant.

In previous releases, if the `lease_duration` was not sent by a remote *DomainParticipant*, a previously received value was used instead.

Note that RTI's DDS implementations send the `lease_duration`.

[RTI Issue ID MICRO-3254]

## Serialization and Deserialization

### [Critical] *DataReader* on a *Topic* using an appendable type may receive samples with incorrect value

A *DataReader* subscribing to a *Topic* on an appendable type may have received incorrect samples from a matching *DataWriter*.

The problem only occurred when the *DataWriter* published a type with fewer members than the *DataReader* type. For example, consider a *DataWriter* on *FooBase* and a *DataReader* on *FooDerived*:

```
@appendable struct FooBase {
    sequence<uint8,1024>base_value;
};

@appendable struct FooDerived {
    sequence<uint8,1024> base_value;
    @default(12) uint8 derived_value;
};
```

In this case, the serialized sample stream would be padded with extra bytes to align the stream to 4 bytes as required by the OMG [Extensible and Dynamic Topic Types for DDS](#) specification, version 1.3. However, the additional padding bytes were incorrectly interpreted as part of the data and *derived\_value* may have been set to a random value.

For example, in the case above, when the *DataWriter* published a sample with type *FooBase*, in some cases the *DataReader* received a sample in which the field *derived\_value* was set to 0 instead of 12.

---

**Note:** *Connext Micro* does not support the `@default` annotation.

---

[RTI Issue ID MICRO-6402]

### [Critical] Malformed samples with invalid strings not dropped by *DataReader*

A *DataReader* may have provided the application a malformed sample containing an invalid value (not Null-terminated) for a string member. The string member may not have been Null-terminated, resulting in undefined behavior if the application tried to access it.

Now, the *DataReader* will not deserialize the sample and the sample will not be provided to the application.

[RTI Issue ID MICRO-3039]

**[Major] Float and double ranges may not have been enforced correctly**

Float and double ranges may not have been enforced correctly. Float and double member values that should not have passed the check ended up passing it.

This issue only occurred under any of the following conditions:

For “float”:

- When @min was set to -3.4E38 for a member, a value smaller than @min passed the check when it should not have.
- When @max was set to 3.4E38 for a member, a value greater than @max passed the check when it should not have.

For “double”:

- When @min was set to -1.7E+308 for a member, a value smaller than @min passed the check when it should not have.
- When @max was set to 1.7E+308 for a member, a value greater than @max passed the check when it should not have.

For “float” and “double”:

- When the member value was set to INFINITY, samples passed the range check when they should not have.
- When the member value was set to NaN, samples passed the range check when they should not have.

[RTI Issue ID MICRO-3280]

**[Major] Deserialization of tampered/corrupted samples may have unexpectedly succeeded**

A *DataReader* may not have detected that a truncated sample due to corruption or tampering was invalid. As a result, the application may have received samples with invalid content.

Now, the deserialization of corrupted samples fails, and they are not provided to the application.

[RTI Issue ID MICRO-3057]

**[Major] Invalid serialization of samples with types containing nested structures with primitive members that require padding**

In *Connext DDS* 6.0.1 and earlier, the serialization of samples with a type containing two or more levels of nested complex types, where the nested types have primitive members that require padding, may have failed. This means that a *DataReader* may have received an invalid value for a sample. Example:

```
// Level-2 Nested type
struct Struct1 {
    uint8 m1;
    uint8 m2;
    int32 m3;
};

// Level-1 Nested type
struct Struct2 {
    int32 m1;
    int32 m2;
    uint8 m3;
    uint8 m4;
    Struct1 m5;
};

struct Struct3 {
    Struct2 m1;
};
```

In the above example, Struct2 and Struct1 are nested, and there is padding between Struct1::m2 (1-byte aligned) and Struct1::m3 (4-byte aligned) of 2 bytes.

This issue only applied to nested types that are appendable or final for XCDR1 data representation or final for XCDR2 data representation.

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 or 0 for the -optimization, but this may have had performance implications.

[RTI Issue ID MICRO-2744]

### **[Minor] Serialization of string members did not check for null-terminated strings in C, traditional C++, and modern C++**

The code executed by a *DataWriter* that serializes string members in a Topic type did not check that the strings are null-terminated. This may have led to undefined behavior, because the serialization code calls `strlen`.

This problem has been fixed. The serialization code now checks for null-terminated strings with the maximum allowed length and reports the following error if the string is not well-formed:

```
RTIXcdrInterpreter_serializeString:StrStruct:member2 serialization error. String length
↪(at least 6) is larger than maximum 5
```

[RTI Issue ID MICRO-3040]

## Usability

### [Trivial] Thread names were not set on QNX

In previous releases, the thread names were not set on QNX.

[RTI Issue ID MICRO-5851]

## Transports

### [Critical] Stalled communication when using shared-memory transport

On systems with a weak memory architecture, such as Arm®, the shared-memory (SHMEM) transport may have been corrupted due to a data race in the concurrent queue where the messages are written into the shared-memory segment. This data race may have occurred until **received\_message\_count\_max** messages were sent through the transport. The corrupted transport resulted in parsing errors, which filled up the shared-memory segment, stalling communication.

[RTI Issue ID MICRO-5931]

### [Critical] Undefined behavior when using SHMEM transport in Linux, macOS, QNX, Integrity, and Lynx

There was an issue in the shared-memory (SHMEM) transport implementation that may have led to undefined behavior in your *Connex Micro* application, including data corruption, errors, and hangs. The problem could occur in Linux®, macOS®, QNX®, INTEGRITY®, and LynxOS® systems.

[RTI Issue ID MICRO-5932]

## Reliability Protocol and Wire Representation

### [Critical] Reliable DataWriter may have ignored requests to resend samples

If a *DataWriter* received multiple requests to resend samples before its periodic heartbeat period expired, the *DataWriter* may have ignored the request if the requested sample had been sent and was also the first expected sample by the requesting *DataReader*.

[RTI Issue ID MICRO-5183]

**[Minor] Incorrect heartbeat sent before first sample when first\_write\_sequence\_number is different from 1**

In previous releases, if the `DataWriterQos.protocol.rtps_reliable_writer.first_write_sequence_number` was different from the default value 1, heartbeats sent before the first sample was written would indicate 1 as the first sample available. This caused a *DataReader* to wait for samples with a sequence number less than `DataWriterQos.protocol.rtps_reliable_writer.first_write_sequence_number` until a heartbeat with the correct first sequence number was received.

[RTI Issue ID MICRO-4081]

**Logging****[Major] Race condition and memory corruption in logger**

The following issues have been fixed in the logger:

- Processing log-messages in a log handler was not thread-safe.
- Memory corruption may have occurred.
- Conversion of INT\_MIN was incorrect.

---

**Note:** The `OSAPI_Log_clear` API must not be called outside a log-handler since it is no longer thread-safe.

---

[RTI Issue ID MICRO-5854]

**Performance and Scalability****[Trivial] Asynchronous publication delay**

In previous releases, there was a delay (equal to the OSAPI Task scheduler's clock rate) before sending a fragmented or flow-controlled sample. This delay has been removed.

[RTI Issue ID MICRO-5853]

## APIs (C or Traditional C++)

### [Critical] Segmentation fault when finalizing DataWriter QoS

Finalizing a *DataWriter* QoS could have resulted in a segmentation fault if `publish_mode.name` was set to a builtin Flow Controller name.

[RTI Issue ID MICRO-5966]

### [Major] DDS\_Subscriber\_lookup\_datareader may return a DataReader that was created by a different Subscriber

The `DDS_Subscriber_lookup_datareader` API searches for a *DataReader* for a given *TopicDescription* created by the *Subscriber*. However, in previous releases, it the returned *DataReader* could belong to a different *Subscriber* if multiple *DataReaders* were created for the same *Topic* in different *Subscribers*.

[RTI Issue ID MICRO-4569]

### [Major] DDS\_Publisher\_lookup\_datawriter may return a DataWriter that was created by a different Publisher

The `DDS_Publisher_lookup_datawriter` API searches for a *DataWriter* for a given *Topic* created by the *Publisher*. However, in previous releases, the returned *DataWriter* could belong to a different *Publisher* if multiple *DataWriters* were created for the same *Topic* in different *Publishers*.

[RTI Issue ID MICRO-4570]

### [Major] DDS\_Entity\_enable was not thread-safe for a DomainParticipant

`DDS_Entity_enable` was not thread-safe, which may have led to race conditions.

[RTI Issue ID MICRO-3379]

### [Major] Race condition in DDS enable APIs

A race condition existed if the same DDS entity was enabled from multiple threads at the same time.

[RTI Issue ID MICRO-3311]

**[Minor] DDS\_FLOW\_CONTROLLER\_PROPERTY\_DEFAULT ignored when used as argument to DDS\_DomainParticipant\_create\_flowcontroller**

The following related issues are resolved in this release:

- *Connext Micro* ignored DDS\_FLOW\_CONTROLLER\_PROPERTY\_DEFAULT when passed in to the DDS\_DomainParticipant\_create\_flowcontroller API call.
- The properties used by *Connext Micro* for the builtin Flow Controllers were not aligned with *Connext Professional*.
- The default Flow Controller properties returned were not aligned with *Connext Professional*.

[RTI Issue ID MICRO-6118]

**[Minor] Failure to parse invalid index**

A peer descriptor string consisting of only an invalid range, e.g, “[3” was incorrectly interpreted as the empty peer address string “”.

[RTI Issue ID MICRO-4436]

**Generated Code (C, Traditional C++, and Modern C++)****[Critical] Foo\_create\_data() failed to create samples for data types with long doubles**

Foo\_create\_data() failed to create samples of types that contained arrays or sequences of types that contained long doubles. For example, Foo\_create\_data() failed for the following type Foo:

```
struct S
{
    long double ld;
};

struct Foo
{
    sequence<S> s;
};
```

However, Foo\_create\_data() did not fail for the following type Foo:

```
struct Foo
{
    sequence<long double> s;
};
```

[RTI Issue ID MICRO-3025]



**[Minor] Example code generated from XML or XSD files failed to compile**

Example code generated by *rtiddsgen* from XML or XSD files failed to compile.

[RTI Issue ID MICRO-2505]

**Micro Application Generator****[Major] NullPointerException when using -outputFinalQoS if QoS Profile did not define each internal QoS**

When using MAG with the `-outputFinalQoS` option, if the QoS Profile to check did not contain a definition of each internal QoS (participant\_qos, publisher\_qos, etc.) directly or by inheriting from another QoS Profile, MAG reported this error:

```
Exception in thread "main" java.lang.NullPointerException
  at com.rti.micro.appgen.utils.QoSUtils.removeNullElementsFromList(QoSUtils.java:2332)
  at com.rti.micro.appgen.utils.QoSUtils.removeNullElements(QoSUtils.java:2256)
  at com.rti.micro.appgen.MicroAppGen.main(MicroAppGen.java:328)
```

[RTI Issue ID MAG-121]

**[Minor] MAG failed to generate code when qos\_profile inherited from individual QoS policies**

MAG failed to generate code when a `<qos_profile>` inherited from individual QoS policies. For example, running MAG with the following input file caused an error:

```
<qos_library name="QoSLibrary">
  <qos_profile name="QoSProfile1" is_default_qos="true">
    <participant_qos name="QoSParticipant">
      ...
    </participant_qos>
  </qos_profile>
  <qos_profile name="QoSProfile2" base_name="QoSProfile1::QoSParticipant">
    </qos_profile>
  </qos_profile>
</qos_library>
```

The error was:

```
...
11:31:40.548 [main] ERROR com.rti.micro.appgen.MicroAppGen - Failed to calculate the
↳ system model.
java.lang.Exception: Unable to find QoS library/profile 'QoSProfile1::QoSParticipant'.
...
11:31:40.552 [main] INFO com.rti.micro.appgen.MicroAppGen - Exiting.
```

Now MAG properly handles this case.

[RTI Issue ID MAG-105]

**[Minor] MAG always used default value for disable\_auto\_interface\_config**

MAG always used the default value for `disable_auto_interface_config` in the generated code, regardless of the value specified in the XML.

[RTI Issue ID MAG-110]

**[Minor] MAG failed if arguments contained whitespace on Linux systems**

On Linux systems, MAG failed to run if any arguments contained whitespace. It logged an error similar to the following:

```
12:04:55.205 [main] ERROR com.rti.micro.appgen.MicroAppGen - Only 1 input file
can be processed.
12:04:55.208 [main] INFO  com.rti.micro.appgen.MicroAppGen - Exiting.
```

[RTI Issue ID MAG-118]

**[Trivial] XSD validation failed if flags used a combination of values**

The XSD validation of an XML application file failed if there was a UDPv4 configuration using a combination of values for the `flags` element. For example, this snippet caused an error:

```
<transport_builtin>
  <udp4>
    <interface_table>
      <element>
        <flags>
          UDP_INTERFACE_INTERFACE_UP_FLAG|UDP_INTERFACE_INTERFACE_MULTICAST_
↪FLAG
        </flags>
      </element>
    </interface_table>
  </udp4>
</transport_builtin>
```

The error was:

```
ERROR com.rti.micro.appgen.MicroAppGen - cvc-pattern-valid:
Value 'UDP_INTERFACE_INTERFACE_UP_FLAG|UDP_INTERFACE_INTERFACE_MULTICAST_FLAG'
is not facet-valid with respect to pattern
'(UDP_INTERFACE_INTERFACE_UP_FLAG|UDP_INTERFACE_INTERFACE_MULTICAST_FLAG)'
for type 'udpInterfaceFlagMask'.
```

Now combinations are allowed.

[RTI Issue ID MAG-114]

## OMG Specification Compliance

[Critical]: System-stopping issue, such as a crash or data loss. [Major]: Significant issue with no easy workaround. [Minor]: Issue that usually has a workaround. [Trivial]: Small issue, such as a typo in a log.

### **[Major] DDS\_StatusCondition\_set\_enabled\_statuses did not trigger if an active condition was enabled and had incorrect default value**

In previous releases, if a StatusCondition enabled by a call to `DDS_StatusCondition_set_enabled_statuses` was already active, the StatusCondition did not trigger.

The default enabled status list was incorrectly set to `DDS_STATUS_MASK_NONE`, but is now set to `DDS_STATUS_MASK_ALL` until the first successful call to `DDS_StatusCondition_set_enabled_statuses`.

[RTI Issue ID MICRO-3308]

## Interoperability

### **[Critical] Failure to deserialize fragmented samples sent by Connex Professional 7**

Due to incorrect processing of an inline QoS in a fragmented sample, *Connex Micro* failed to deserialize fragmented samples sent by *Connex Professional 7*, or other implementations that set the length of the `PID_SENTINEL` to a value different than 1.

[RTI Issue ID MICRO-4095]

### **[Critical] Inline QoS offset non-compliant with DDSI-RTPS standard**

The inline QoS offset in DATA was set to 0 instead of the offset to the next field after the inline QoS. This may have caused DDS implementations from other vendors to fail to receive data.

[RTI Issue ID MICRO-4160]

### **[Critical] Connex Micro may have repeated requesting a sample that was no longer available from a DataWriter**

If *Connex Micro* detects a missing sample when using `DDS_RELIABLE_RELIABILITY_QOS` reliability, it will request the sample to be resent, but if the sample is no longer available from the *DataWriter*, the *DataWriter* may send a GAP message to indicate the sample is not longer available.

*Connex Micro* failed to interpret the GAP message correctly if the first sequence number in the GAP message was equal to the bitmap base of the GAP message. In this case, *Connex Micro* failed to ignore the no-longer-available sample and kept sending a request for the sample.

[RTI Issue ID MICRO-4668]

### **[Critical] Failure to deserialize a fragmented sample with multiple fragments in a DATA\_FRAG submessage**

A deserialization error occurred when deserializing a sample that was fragmented into multiple fragments in a single RTPS DATA\_FRAG submessage.

[RTI Issue ID MICRO-2958]

## **Vulnerabilities**

### **[Critical] Vulnerabilities in RTI Micro Application Generator (MAG)**

This release fixes vulnerabilities in Log4j known as “log4shell”. You can find further details in RTI’s **Security Notice 2021-12-log4j** at <https://community.rti.com/kb/apache-log4j-vulnerability-cve-2021-44228cve-2021-45046-impact-rti-connex-products>.

RTI Micro Application Generator uses Apache Log4j version 2.17.1 in this release.

[RTI Issue ID MAG-147]

### **[Critical] Illegal memory access when failing to generate interpreter programs**

Receiving malicious endpoint discovery information might have resulted (very rarely) in an arbitrary read from the thread stack.

User impact with or without security was as follows:

- Remotely exploitable
- Crash application
- Potentially impacting confidentiality of Connex application
- CVSS Base Score: 6.5 MEDIUM
- CVSS v3.1 Vector: [AV:N/AC:H/PR:N/UI:N/S:U/C:L/I:N/A:H](#)

[RTI Issue ID MICRO-3219]

**[Critical] Potential crash when receiving a malformed sample using DDS\_XCDR2\_DATA\_REPRESENTATION**

A *Connext Micro* application could have crashed if a *DataReader* received a malformed serialized sample using DDS\_XCDR2\_DATA\_REPRESENTATION. The issue only affected appendable or mutable types.

User impact with or without security was as follows:

- Remotely exploitable through malicious RTPS messages
- Connext application could crash or potentially leak sensitive information
- CVSS Base Score: 6.5 MEDIUM
- CVSS v3.1 Vector: [AV:N/AC:H/PR:N/UI:N/S:U/C:L/I:N/A:H](#)

[RTI Issue ID MICRO-3118]

**Other****[Minor] Delay in sending data when using a flow-controller**

When using a flow-controller to send data, there was a delay before sending the first sample or fragment (of up to one task period).

[RTI Issue ID MICRO-6494]

**[Minor] Non-default timer resolutions may have caused an incorrect timeout**

Compiling *Connext Micro* with a non-default timer resolution may have caused incorrect timeouts.

[RTI Issue ID MICRO-6476]

## 8.6 Known Issues

### 8.6.1 Samples cannot be recovered if subscribing application fails to return loan

When a subscribing application has taken a loan for the Zero Copy v2 transport (using the API [FooDataReader\\_read\(\)](#) or [FooDataReader\\_take\(\)](#)) and fails to return the loan due to a crash or other circumstances, *Connext Micro* cannot recover those samples. This also affects the matching *DataWriter*, which cannot reclaim the samples and continues to run in a degraded state.

[RTI Issue ID MICRO-5834]

### 8.6.2 Connex Micro does not work with wide-string characters in the network interface name

*Connex Micro* does not work with wide-string characters (such as Japanese or Chinese characters) in the network interface name.

As a workaround, rename all the system interfaces so that none of them contain wide-string characters.

[RTI Issue ID MICRO-2423]

### 8.6.3 64-bit discriminator values greater than $(2^{31}-1)$ or smaller than $(-2^{31})$ not supported

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

- C
- Traditional C++

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

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

[RTI Issue ID MICRO-3056]

### 8.6.4 NaN and INF float and doubles are not detected and will not cause errors

Normally, *Connex Micro* discards samples with values that are out of range during serialization and de-serialization; however, Not a Number (NaN) and Infinite (INF) floating point and doubles are not detected and will not cause serialization or de-serialization errors.

[RTI Issue ID MICRO-5960]

### 8.6.5 Ungracefully terminated QNX processes using SHMEM transport prevents startup of new processes due to unclosed POSIX semaphores

If a QNX 7.0 or earlier application using the shared-memory transport was ungracefully shut down, crashed, or otherwise had an abnormal termination while holding a POSIX semaphore used by the transport (for example, while sending data through the shared-memory transport), *Connex* applications launched after that point on the same domain may wait forever for that semaphore to be released.

Workaround for QNX 7.0 and earlier: to enable new applications to start, RTI recommends stopping all applications, then cleaning up the Inter-Process Communication (IPC) resources before starting new applications.

This problem is resolved for QNX 7.1, as described in the fix for *[Critical] Ungracefully terminated QNX processes using SHMEM transport prevented startup of new processes due to unclosed POSIX semaphores*.

[RTI Issue ID MICRO-6013]

### 8.6.6 Flow Controllers require RTOS

Flow controllers require an RTOS.

[RTI Issue ID MICRO-6648]

### 8.6.7 LatencyBudget is not part of the DataReaderQos or DataWriterQos policy

The LatencyBudgetQos policy is not supported and does not appear as part of the DataReader and DataWriter Qos policy documentation. The default value is 0. When creating earliest deadline first (EDF) flow-controllers, the effective scheduling is round-robin.

[RTI Issue ID MICRO-6649]

### 8.6.8 Porting Guide not included

*RTI Connex Micro* 4.2.0 has many internal changes from previous versions of *RTI Connex Micro*, and the *RTI Connex Micro* 4.2.0 APIs are not considered stable. Therefore, instructions for porting *RTI Connex Micro* 4.2.0 are not included. If you need instructions to port *RTI Connex Micro* 4.2.0, please contact [support@rti.com](mailto:support@rti.com).

[RTI Issue ID MICRO-8618]

### 8.6.9 Platform Independent Library toolchain dependencies

The platform independent libraries (PIL) are not completely independent of the toolchain and standard C library, and thus require a compatible toolchain and standard library to link to. See *Platform Notes* for more information.

[RTI Issue ID MICRO-8154]

## 8.7 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**».

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 features may be only available in a subset of the supported languages and for a subset of the supported platforms.
- Experimental features may change in the future.
- 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/>).



## Chapter 9

# Benchmarks

Performance benchmarks are no longer included with an *RTI Connex Micro* installation. Please refer to the [RTI Connex Performance Benchmarks](#) on RTI Community for more information.

---

**Note:** The *RTI Connex Performance Benchmarks* contain metrics for multiple products and versions, so please ensure that you refer to the appropriate section.

---

# Chapter 10

## Copyrights

© 2017-2025 Real-Time Innovations, Inc.

All rights reserved.

Printed in U.S.A. First printing.

September 2025.

### Trademarks

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

### Copy and Use Restrictions

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form (including electronic, mechanical, photocopy, and facsimile) without the prior written permission of Real-Time Innovations, Inc. The software described in this document is furnished solely under and subject to RTI’s standard terms and conditions available at <https://www.rti.com/terms> and in accordance with your License Acknowledgement Certificate (LAC) and Maintenance and Support Certificate (MSC), except to the extent otherwise agreed to in writing by RTI.

### Third-Party Software

RTI software may contain independent, third-party software or code that are subject to third-party license terms and conditions, including open source license terms and conditions. Copies of applicable third-party licenses and notices are located at [community.rti.com/documentation](https://community.rti.com/documentation). IT IS YOUR RESPONSIBILITY TO ENSURE THAT YOUR USE OF THIRD-PARTY SOFTWARE COMPLIES WITH THE CORRESPONDING THIRD-PARTY LICENSE TERMS AND CONDITIONS.

## Notices

### *Deprecations and Removals*

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

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

### Technical Support

Real-Time Innovations, Inc.

232 E. Java Drive

Sunnyvale, CA 94089

Phone: (408) 990-7444

Email: [support@rti.com](mailto:support@rti.com)

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

## Chapter 11

# Third-Party and Open Source Software

This section outlines Real-Time Innovations (RTI) usage of first-level third-party and open source software in the *RTI Connex Micro* libraries and utilities.

### 11.1 Connex Micro Libraries

#### 11.1.1 fnmatch

- Related to: Content-filtered topics, query conditions, partitions, multicast address management, topic filter in XML QoS profile.
- Software is included in the core middleware libraries.
- Third-Party Software License:

Copyright (c) 1989, 1993, 1994

The Regents of the University of California. All rights reserved.

This code is derived from software contributed to Berkeley by Guido van Rossum.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

(continues on next page)

(continued from previous page)

3. All advertising materials mentioning features or use of this software must display the following acknowledgement:

This product includes software developed by the University of California, Berkeley and its contributors.

4. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

### 11.1.2 crc32c.c

- Related to: RTPS CRC-32 checksum support
- Version 1.1
- Third-Party Software License:

This software is provided 'as-is', without any express or implied warranty. In no event will the author be held liable for any damages arising from the use of the software.

Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions:

1. The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgement in the product documentation would be appreciated but is not required.
2. Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.
3. This notice may not be removed or altered from any source distribution

(continues on next page)

(continued from previous page)

Mark Adler

madler@alumni.caltech.edu

### 11.1.3 MD5

- Related to: DDS keys implementation, content-filtered topics (to sign the filter), persistence service (to generate writer-side unique identification), Integration Toolkit for AUTOSAR (DDS-IDL Service Interface code generation)
- Software is included in core DDS middleware libraries, in the Connex DDS Micro libraries, in the Connex DDS Cert libraries and in the Integration Toolkit for AUTOSAR.
- Third-Party Software License:

Copyright (C) 1999, 2002 Aladdin Enterprises. All rights reserved.

This software is provided 'as-is', without any express or implied warranty. In no event will the authors be held liable for any damages arising from the use of this software.

Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions:

1. The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required.
2. Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.
3. This notice may not be removed or altered from any source distribution.

L. Peter Deutsch  
ghost@aladdin.com

## 11.2 RTI Code Generator (rtiddsgen)

### 11.2.1 ANTLR

- This software is distributed with rtiddsgen (RTI Code Generator) as a jar file. The source code is not modified or shipped. In addition, the output produced by this software from a grammar file is part of the rtiddsgen JAR file.
- Version: Release 3.5.2
- Open Source Software License: <https://www.antlr3.org/license.html>

[The BSD License]

Copyright (c) 2010 Terence Parr

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

Neither the name of the author nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

### 11.2.2 Apache Commons Lang

- Used in rtiddsgen. We only use the class StringUtils.
- Version 2.6
- Open Source Software License: Apache License Version 2.0 (full text found in the *Appendix*).

### 11.2.3 Apache Log4j 2

- This software is distributed with rtiddsmag (Micro Application Generator) as a jar file. The source code is not modified or shipped.
- Source: <https://logging.apache.org/>
- Version: 2.17.1
- Open Source Software License: Apache License Version 2.0 (full text found in the *Appendix*).

### 11.2.4 Apache Velocity

- This software is included in rtiddsgen. The source code is not modified or shipped.
- Source: <http://velocity.apache.org/>
- Version: 2.3
- Open Source Software License: Apache License Version 2.0 (full text found in the *Appendix*).

### 11.2.5 Simple Logging Facade for Java (SLF4J)

- This software is included in rtiddsgen.
- Version 1.7.35
- Open Source Software License: <https://www.slf4j.org/license.html>

Copyright (c) 2004-2025 QOS.ch  
All rights reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

### 11.2.6 Gson

- Portions of rtiddsgen are built using Gson.
- Version 2.9.1
- Open Source Software License: Apache License Version 2.0 (full text found in the *Appendix*).



## 11.3 Micro Application Generator (rtiddsmag)

### 11.3.1 Apache Commons CLI

- Used in Micro Application Generator.
- Version 1.4
- Open Source Software License: Apache License Version 2.0 (full text found in the *Appendix*).

### 11.3.2 Apache Commons Lang

- Used in Micro Application Generator. We only use the class StringUtils.
- Version 3.7
- Open Source Software License: Apache License Version 2.0 (full text found in the *Appendix*).

### 11.3.3 Apache Log4j 2

- This software is distributed with rtiddsmag (Micro Application Generator) as a jar file. The source code is not modified or shipped.
- Source: <https://logging.apache.org/>
- Version: 2.17.1
- Open Source Software License: Apache License Version 2.0 (full text found in the *Appendix*).

### 11.3.4 Apache Velocity

- This software is included in rtiddsmag (Micro Application Generator). The source code is not modified or shipped.
- Source: <http://velocity.apache.org/>
- Version: 2.0
- Open Source Software License: Apache License Version 2.0 (full text found in the *Appendix*).

### 11.3.5 Extended StAX API

- Version 1.8
- Open Source Software License: <https://www.eclipse.org/org/documents/edl-v10.php>

Eclipse Distribution License - v 1.0

Copyright (c) 2007, Eclipse Foundation, Inc. and its licensors.

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- \* Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- \* Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

Neither the name of the Eclipse Foundation, Inc. nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

### 11.3.6 Fast Infoset

- Version 1.2.15
- Open Source Software License: Apache License Version 2.0 (full text found in the *Appendix*).

### 11.3.7 Istack Common Utility Code Runtime

- Version 3.0.7
- Open Source Software License: Common Development and Distribution License (CDDL) Version 1.1 (full text found in the *Appendix*).

### 11.3.8 JavaBeans Activation Framework API

- Version 1.2.0
- Open Source Software License: Common Development and Distribution License (CDDL) Version 1.1 (full text found in the *Appendix*).

### 11.3.9 Javax Annotation API

- Version 1.3.2
- Open Source Software License: Common Development and Distribution License (CDDL) Version 1.1 (full text found in the *Appendix*).

### 11.3.10 JAXB API

- Version 2.3.1
- Open Source Software License: Common Development and Distribution License (CDDL) Version 1.1 (full text found in the *Appendix*).

### 11.3.11 JAXB Runtime

- Version 2.3.1
- Open Source Software License: Common Development and Distribution License (CDDL) Version 1.1 (full text found in the *Appendix*).

### 11.3.12 Simple Logging Facade for Java (SLF4J)

- This software is included in rtiddsmag (Micro Application Generator).
- Version 1.7.25
- Open Source Software License: <https://www.slf4j.org/license.html>

Copyright (c) 2004-2025 QOS.ch  
All rights reserved.

Permission is hereby granted, free of charge, to any person obtaining  
a copy of this software and associated documentation files (the

(continues on next page)

(continued from previous page)

"Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

### 11.3.13 TXW2

- Version 2.3.1
- Open Source Software License: Common Development and Distribution License (CDDL) Version 1.1 (full text found in the *Appendix*).

## 11.4 Lightweight Security Plugin

### 11.4.1 OpenSSL

- Related to: cryptographic functionality
- Software is provided as separate libraries. RTI Security-related plugins use public OpenSSL APIs. We did not copy-paste OpenSSL internal code. RTI Security-related plugins depend on the OpenSSL libraries.
- Release: 3.5.0
- Download link: <https://github.com/openssl/openssl/releases/download/openssl-3.5.0/openssl-3.5.0.tar.gz>
- Open Source Software License: Apache License Version 2.0 (full text found in the *Appendix*).

## 11.5 Appendix

### 11.5.1 Apache License version 2.0, January 2004 (<http://www.apache.org/licenses/>)

Apache License  
Version 2.0, January 2004  
<http://www.apache.org/licenses/>

#### TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

##### 1. Definitions.

"License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files.

"Object" form shall mean any form resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below).

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

(continues on next page)

(continued from previous page)

"Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution."

"Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. Grant of Copyright License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.
3. Grant of Patent License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.
4. Redistribution. You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions:
  - (a) You must give any other recipients of the Work or Derivative Works a copy of this License; and
  - (b) You must cause any modified files to carry prominent notices stating that You changed the files; and

(continues on next page)

(continued from previous page)

- (c) You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and
- (d) If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License.

You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications, or for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.

- 5. Submission of Contributions. Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.
- 6. Trademarks. This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.
- 7. Disclaimer of Warranty. Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the

(continues on next page)

(continued from previous page)

appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.

8. **Limitation of Liability.** In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.
9. **Accepting Warranty or Additional Liability.** While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

APPENDIX: How to apply the Apache License to your work.

To apply the Apache License to your work, attach the following boilerplate notice, with the fields enclosed by brackets "[]" replaced with your own identifying information. (Don't include the brackets!) The text should be enclosed in the appropriate comment syntax for the file format. We also recommend that a file or class name and description of purpose be included on the same "printed page" as the copyright notice for easier identification within third-party archives.

Copyright [yyyy] [name of copyright owner]

Licensed under the Apache License, Version 2.0 (the "License");  
you may not use this file except in compliance with the License.  
You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.



## 11.5.2 Common Development and Distribution License (CDDL) Version 1.1

### 1. Definitions.

1.1. "Contributor" means each individual or entity that creates or contributes to the creation of Modifications.

1.2. "Contributor Version" means the combination of the Original Software, prior Modifications used by a Contributor (if any), and the Modifications made by that particular Contributor.

1.3. "Covered Software" means (a) the Original Software, or (b) Modifications, or (c) the combination of files containing Original Software with files containing Modifications, in each case including portions thereof.

1.4. "Executable" means the Covered Software in any form other than Source Code.

1.5. "Initial Developer" means the individual or entity that first makes Original Software available under this License.

1.6. "Larger Work" means a work which combines Covered Software or portions thereof with code not governed by the terms of this License.

1.7. "License" means this document.

1.8. "Licensable" means having the right to grant, to the maximum extent possible, whether at the time of the initial grant or subsequently acquired, any and all of the rights conveyed herein.

1.9. "Modifications" means the Source Code and Executable form of any of the following:

A. Any file that results from an addition to, deletion from or modification of the contents of a file containing Original Software or previous Modifications;

B. Any new file that contains any part of the Original Software or previous Modification; or

C. Any new file that is contributed or otherwise made available under the terms of this License.

1.10. "Original Software" means the Source Code and Executable form of computer software code that is originally released under this License.

1.11. "Patent Claims" means any patent claim(s), now owned or hereafter acquired, including without limitation, method, process, and apparatus claims, in any patent Licensable by grantor.

(continues on next page)

(continued from previous page)

1.12. "Source Code" means (a) the common form of computer software code in which modifications are made and (b) associated documentation included in or with such code.

1.13. "You" (or "Your") means an individual or a legal entity exercising rights under, and complying with all of the terms of, this License. For legal entities, "You" includes any entity which controls, is controlled by, or is under common control with You. For purposes of this definition, "control" means (a) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (b) ownership of more than fifty percent (50%) of the outstanding shares or beneficial ownership of such entity.

## 2. License Grants.

### 2.1. The Initial Developer Grant.

Conditioned upon Your compliance with Section 3.1 below and subject to third party intellectual property claims, the Initial Developer hereby grants You a world-wide, royalty-free, non-exclusive license:

(a) under intellectual property rights (other than patent or trademark) Licensable by Initial Developer, to use, reproduce, modify, display, perform, sublicense and distribute the Original Software (or portions thereof), with or without Modifications, and/or as part of a Larger Work; and

(b) under Patent Claims infringed by the making, using or selling of Original Software, to make, have made, use, practice, sell, and offer for sale, and/or otherwise dispose of the Original Software (or portions thereof).

(c) The licenses granted in Sections 2.1(a) and (b) are effective on the date Initial Developer first distributes or otherwise makes the Original Software available to a third party under the terms of this License.

(d) Notwithstanding Section 2.1(b) above, no patent license is granted: (1) for code that You delete from the Original Software, or (2) for infringements caused by: (i) the modification of the Original Software, or (ii) the combination of the Original Software with other software or devices.

### 2.2. Contributor Grant.

Conditioned upon Your compliance with Section 3.1 below and subject to third party intellectual property claims, each Contributor hereby grants You a world-wide, royalty-free, non-exclusive license:

(a) under intellectual property rights (other than patent or

(continues on next page)

(continued from previous page)

trademark) Licensable by Contributor to use, reproduce, modify, display, perform, sublicense and distribute the Modifications created by such Contributor (or portions thereof), either on an unmodified basis, with other Modifications, as Covered Software and/or as part of a Larger Work; and

(b) under Patent Claims infringed by the making, using, or selling of Modifications made by that Contributor either alone and/or in combination with its Contributor Version (or portions of such combination), to make, use, sell, offer for sale, have made, and/or otherwise dispose of: (1) Modifications made by that Contributor (or portions thereof); and (2) the combination of Modifications made by that Contributor with its Contributor Version (or portions of such combination).

(c) The licenses granted in Sections 2.2(a) and 2.2(b) are effective on the date Contributor first distributes or otherwise makes the Modifications available to a third party.

(d) Notwithstanding Section 2.2(b) above, no patent license is granted: (1) for any code that Contributor has deleted from the Contributor Version; (2) for infringements caused by: (i) third party modifications of Contributor Version, or (ii) the combination of Modifications made by that Contributor with other software (except as part of the Contributor Version) or other devices; or (3) under Patent Claims infringed by Covered Software in the absence of Modifications made by that Contributor.

### 3. Distribution Obligations.

#### 3.1. Availability of Source Code.

Any Covered Software that You distribute or otherwise make available in Executable form must also be made available in Source Code form and that Source Code form must be distributed only under the terms of this License. You must include a copy of this License with every copy of the Source Code form of the Covered Software You distribute or otherwise make available. You must inform recipients of any such Covered Software in Executable form as to how they can obtain such Covered Software in Source Code form in a reasonable manner on or through a medium customarily used for software exchange.

#### 3.2. Modifications.

The Modifications that You create or to which You contribute are governed by the terms of this License. You represent that You believe Your Modifications are Your original creation(s) and/or You have sufficient rights to grant the rights conveyed by this License.

#### 3.3. Required Notices.

(continues on next page)

(continued from previous page)

You must include a notice in each of Your Modifications that identifies You as the Contributor of the Modification. You may not remove or alter any copyright, patent or trademark notices contained within the Covered Software, or any notices of licensing or any descriptive text giving attribution to any Contributor or the Initial Developer.

#### 3.4. Application of Additional Terms.

You may not offer or impose any terms on any Covered Software in Source Code form that alters or restricts the applicable version of this License or the recipients' rights hereunder. You may choose to offer, and to charge a fee for, warranty, support, indemnity or liability obligations to one or more recipients of Covered Software. However, you may do so only on Your own behalf, and not on behalf of the Initial Developer or any Contributor. You must make it absolutely clear that any such warranty, support, indemnity or liability obligation is offered by You alone, and You hereby agree to indemnify the Initial Developer and every Contributor for any liability incurred by the Initial Developer or such Contributor as a result of warranty, support, indemnity or liability terms You offer.

#### 3.5. Distribution of Executable Versions.

You may distribute the Executable form of the Covered Software under the terms of this License or under the terms of a license of Your choice, which may contain terms different from this License, provided that You are in compliance with the terms of this License and that the license for the Executable form does not attempt to limit or alter the recipient's rights in the Source Code form from the rights set forth in this License. If You distribute the Covered Software in Executable form under a different license, You must make it absolutely clear that any terms which differ from this License are offered by You alone, not by the Initial Developer or Contributor. You hereby agree to indemnify the Initial Developer and every Contributor for any liability incurred by the Initial Developer or such Contributor as a result of any such terms You offer.

#### 3.6. Larger Works.

You may create a Larger Work by combining Covered Software with other code not governed by the terms of this License and distribute the Larger Work as a single product. In such a case, You must make sure the requirements of this License are fulfilled for the Covered Software.

### 4. Versions of the License.

#### 4.1. New Versions.

Oracle is the initial license steward and may publish revised and/or

(continues on next page)

(continued from previous page)

new versions of this License from time to time. Each version will be given a distinguishing version number. Except as provided in Section 4.3, no one other than the license steward has the right to modify this License.

#### 4.2. Effect of New Versions.

You may always continue to use, distribute or otherwise make the Covered Software available under the terms of the version of the License under which You originally received the Covered Software. If the Initial Developer includes a notice in the Original Software prohibiting it from being distributed or otherwise made available under any subsequent version of the License, You must distribute and make the Covered Software available under the terms of the version of the License under which You originally received the Covered Software. Otherwise, You may also choose to use, distribute or otherwise make the Covered Software available under the terms of any subsequent version of the License published by the license steward.

#### 4.3. Modified Versions.

When You are an Initial Developer and You want to create a new license for Your Original Software, You may create and use a modified version of this License if You: (a) rename the license and remove any references to the name of the license steward (except to note that the license differs from this License); and (b) otherwise make it clear that the license contains terms which differ from this License.

### 5. DISCLAIMER OF WARRANTY.

COVERED SOFTWARE IS PROVIDED UNDER THIS LICENSE ON AN "AS IS" BASIS, WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, WARRANTIES THAT THE COVERED SOFTWARE IS FREE OF DEFECTS, MERCHANTABLE, FIT FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE COVERED SOFTWARE IS WITH YOU. SHOULD ANY COVERED SOFTWARE PROVE DEFECTIVE IN ANY RESPECT, YOU (NOT THE INITIAL DEVELOPER OR ANY OTHER CONTRIBUTOR) ASSUME THE COST OF ANY NECESSARY SERVICING, REPAIR OR CORRECTION. THIS DISCLAIMER OF WARRANTY CONSTITUTES AN ESSENTIAL PART OF THIS LICENSE. NO USE OF ANY COVERED SOFTWARE IS AUTHORIZED HEREUNDER EXCEPT UNDER THIS DISCLAIMER.

### 6. TERMINATION.

6.1. This License and the rights granted hereunder will terminate automatically if You fail to comply with terms herein and fail to cure such breach within 30 days of becoming aware of the breach. Provisions which, by their nature, must remain in effect beyond the termination of this License shall survive.

(continues on next page)

(continued from previous page)

6.2. If You assert a patent infringement claim (excluding declaratory judgment actions) against Initial Developer or a Contributor (the Initial Developer or Contributor against whom You assert such claim is referred to as "Participant") alleging that the Participant Software (meaning the Contributor Version where the Participant is a Contributor or the Original Software where the Participant is the Initial Developer) directly or indirectly infringes any patent, then any and all rights granted directly or indirectly to You by such Participant, the Initial Developer (if the Initial Developer is not the Participant) and all Contributors under Sections 2.1 and/or 2.2 of this License shall, upon 60 days notice from Participant terminate prospectively and automatically at the expiration of such 60 day notice period, unless if within such 60 day period You withdraw Your claim with respect to the Participant Software against such Participant either unilaterally or pursuant to a written agreement with Participant.

6.3. If You assert a patent infringement claim against Participant alleging that the Participant Software directly or indirectly infringes any patent where such claim is resolved (such as by license or settlement) prior to the initiation of patent infringement litigation, then the reasonable value of the licenses granted by such Participant under Sections 2.1 or 2.2 shall be taken into account in determining the amount or value of any payment or license.

6.4. In the event of termination under Sections 6.1 or 6.2 above, all end user licenses that have been validly granted by You or any distributor hereunder prior to termination (excluding licenses granted to You by any distributor) shall survive termination.

## 7. LIMITATION OF LIABILITY.

UNDER NO CIRCUMSTANCES AND UNDER NO LEGAL THEORY, WHETHER TORT (INCLUDING NEGLIGENCE), CONTRACT, OR OTHERWISE, SHALL YOU, THE INITIAL DEVELOPER, ANY OTHER CONTRIBUTOR, OR ANY DISTRIBUTOR OF COVERED SOFTWARE, OR ANY SUPPLIER OF ANY OF SUCH PARTIES, BE LIABLE TO ANY PERSON FOR ANY INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES OF ANY CHARACTER INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF GOODWILL, WORK STOPPAGE, COMPUTER FAILURE OR MALFUNCTION, OR ANY AND ALL OTHER COMMERCIAL DAMAGES OR LOSSES, EVEN IF SUCH PARTY SHALL HAVE BEEN INFORMED OF THE POSSIBILITY OF SUCH DAMAGES. THIS LIMITATION OF LIABILITY SHALL NOT APPLY TO LIABILITY FOR DEATH OR PERSONAL INJURY RESULTING FROM SUCH PARTY'S NEGLIGENCE TO THE EXTENT APPLICABLE LAW PROHIBITS SUCH LIMITATION. SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OR LIMITATION OF INCIDENTAL OR CONSEQUENTIAL DAMAGES, SO THIS EXCLUSION AND LIMITATION MAY NOT APPLY TO YOU.

## 8. U.S. GOVERNMENT END USERS.

(continues on next page)

(continued from previous page)

The Covered Software is a "commercial item," as that term is defined in 48 C.F.R. 2.101 (Oct. 1995), consisting of "commercial computer software" (as that term is defined at 48 C.F.R. § 252.227-7014(a)(1)) and "commercial computer software documentation" as such terms are used in 48 C.F.R. 12.212 (Sept. 1995). Consistent with 48 C.F.R. 12.212 and 48 C.F.R. 227.7202-1 through 227.7202-4 (June 1995), all U.S. Government End Users acquire Covered Software with only those rights set forth herein. This U.S. Government Rights clause is in lieu of, and supersedes, any other FAR, DFAR, or other clause or provision that addresses Government rights in computer software under this License.

#### 9. MISCELLANEOUS.

This License represents the complete agreement concerning subject matter hereof. If any provision of this License is held to be unenforceable, such provision shall be reformed only to the extent necessary to make it enforceable. This License shall be governed by the law of the jurisdiction specified in a notice contained within the Original Software (except to the extent applicable law, if any, provides otherwise), excluding such jurisdiction's conflict-of-law provisions. Any litigation relating to this License shall be subject to the jurisdiction of the courts located in the jurisdiction and venue specified in a notice contained within the Original Software, with the losing party responsible for costs, including, without limitation, court costs and reasonable attorneys' fees and expenses. The application of the United Nations Convention on Contracts for the International Sale of Goods is expressly excluded. Any law or regulation which provides that the language of a contract shall be construed against the drafter shall not apply to this License. You agree that You alone are responsible for compliance with the United States export administration regulations (and the export control laws and regulation of any other countries) when You use, distribute or otherwise make available any Covered Software.

#### 10. RESPONSIBILITY FOR CLAIMS.

As between Initial Developer and the Contributors, each party is responsible for claims and damages arising, directly or indirectly, out of its utilization of rights under this License and You agree to work with Initial Developer and Contributors to distribute such responsibility on an equitable basis. Nothing herein is intended or shall be deemed to constitute any admission of liability.

# Chapter 12

## Appendix

### 12.1 Package Contents

This section provides an overview of the installation package contents, as well as the resultant directory structure in the *Connex Micro* installation.

#### 12.1.1 Host bundle

When installed, the host bundle (`rti_connext_dds_micro-<version>-host.rtipkg`) adds the following to the *Connex Micro* directory:

```
--rti_connext_dds-<version>/
|
+--rti_connext_dds_micro-<version>/
    |--bin/
    |--doc/
    |--example/
    |--include/
    |--resource/
    |--rtiddsgen/
    |--rtiddsmag/
    |--CMakeLists.txt
    |--ReadMe.html
    +--src/
        +--rti_me_psl
```

- The `bin/` directory contains scripts for *rtiddsgen*, *rtiddsmag*, and *rtime-make*.
- The `doc/` directory contains this documentation, as well as the C and C++ API References.
- The `example/` directory contains buildable example applications, as well as instructions on how to build and run them. Refer to *Provided examples* for more information.
- The `include/` directory contains the public header files to compile applications.
- The `resource/` directory contains the build system used for the examples and Platform Support Libraries (PSL).



- The `rtiddsgen/` directory contains an IDL compiler for type support code.
- The `rtiddsmag/` directory contains a tool for generating application code from XML descriptions.
- `CMakeLists.txt` is the main input file to [CMake](#) and is used to generate build files.
- `ReadMe.html` opens this documentation.
- `src/` contains the source files for all supported Platform Support Libraries (PSL); refer to *Target bundle* for more information on the PSL. These can be recompiled for specific platform configurations, as described in *Building the PSL*.

### 12.1.2 Target bundle

When installed, the target bundle (`rti_connext_dds_micro-<version>-target-<architecture>.rtipkg`) adds the following to the *Connext Micro* directory:

```
--rti_connext_dds-7.3.0/
  |--rti_connext_dds_micro-4.2.0/
    |--lib/
      |  |--<arch>/
      |  |  +---<arch libraries>
      |  |--<arch>CERT/
      |  |  +---<arch CERT libraries>
      |  |--<arch>-<PSL>/
      |  |  +---<arch PSL libraries>
      |  |--<arch>CERT-<PSL>/
      |  |  +---<arch CERT PSL libraries>
      |
    |--sbom/
```

- The `lib/` directory contains the libraries needed to build *Connext Micro*.
  - `<arch>` contains pre-built Static and Dynamic (where supported) Release and Debug libraries. These may be integrated libraries or Platform Independent Libraries; see *Library descriptions* below.
  - `<arch>CERT` contains pre-built CERT profile Release and Debug libraries, if they exist for the specific architecture. These may be integrated libraries or Platform Independent Libraries; see *Library descriptions* below.
  - `<arch>-<PSL>` contains pre-built Platform Support Libraries for the specific architecture.
  - `<arch>CERT-<PSL>` contains CERT profile Platform Support Libraries, if they exist for the specific architecture.
- The `sbom/` directory contains the Software Bill of Materials (SBOM) for the target libraries.

### 12.1.3 Lightweight Security Plugin host bundle

When installed, the Lightweight Security Plugin host bundle (`rti_connext_dds_micro-<version>-lw-security-host.rtipkg`) adds the following to the *Connext Micro* directory:

```
--rti_connext_dds-7.3.0/
  |--rti_connext_dds_micro-4.2.0/
    |--include/
      |--rti_me/
        |--dds_psk/
          |--rti_me_psl/
            |--pskpsl/
      |--rtiddsgen/
        +---<security-specific example templates>
```

- The `include/` directory contains security-related header files to use the security binaries. These header files do not include the security interfaces.
- The `rtiddsgen/` directory includes security-specific example templates.

### 12.1.4 Lightweight Security Plugin target bundle

When installed, the Lightweight Security Plugin target bundle (`rti_connext_dds_micro-<version>-lw-security-target-openssl-3.5-<architecture>.rtipkg`) adds the following to the *Connext Micro* directory:

```
--rti_connext_dds-7.3.0/
  |--rti_connext_dds_micro-4.2.0/
    |--lib/
      |--<arch>/
        +---<arch PSK Security PIL>
      |--<arch>-<PSL>/
        +---<arch Transform PSL libraries>
      |--<arch>CERT/
        +---<arch CERT PSK Security PIL>
      |--<arch>CERT-<PSL>/
        +---<arch CERT Transform PSL libraries>
    |--sbom/
      +---<ddspsk*>
      +---<pskpsl*>
```

- The `lib/` directory includes the libraries for your architecture and platform, including the pre-built Transform PSL (`pskpsl`).
- The `sbom/` directory includes the Software Bill of Materials (SBOM) for the security target libraries.

## 12.2 Directory Structure

The complete, default *Connext Micro* installation is structured as shown below:

```

+--rti_connext_dds-<version>/
|
+--rti_connext_dds_micro-<version> (rti_connext_dds_micro-<version>-host.rtipkg)
|  |--bin/
|  |--doc/
|  |--example/
|  |--include/
|  |--resource/
|  |--rtiddsgen/
|  |--rtiddsmag/
|  |--CMakeLists.txt
|  |--ReadMe.html
|  |--sbom/
+--src/
|  +-- rti_me_psl
+--lib (rti_connext_dds_micro-<version>-target-<arch>.rtipkg)
    +--<arch>
    |  +---<arch libraries>
    +--<arch>CERT
    |  +---<arch libraries>
    +--<arch>-<PSL>
    |  +-- <arch PSL libraries>
    +--<arch>CERT-<PSL>
        +-- <arch PSL libraries>

```

This directory structure is recommended and *should* be used because:

- The source bundle includes a helper script to run [CMake](#) that expects this directory structure.
- This directory structure supports multiple architectures.

## Chapter 13

# Contact Support

We welcome your input on how to improve *RTI Connex Micro* to suit your needs. If you have questions or comments about this release, please visit the RTI Customer Portal, <https://support.rti.com>. The RTI Customer Portal provides access to RTI software, documentation, and support. It also allows you to log support cases.

To access the software, documentation or log support cases, 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](mailto:license@rti.com). Resetting your login password can be done directly at the RTI Customer Portal.

## Chapter 14

# Join the Community

[RTI Community](#) provides a free public knowledge base containing how-to guides, detailed solutions, and example source code for many use cases. Search it whenever you need help using and developing with RTI products.

[RTI Community](#) also provides forums for all RTI users to connect and interact.