



CoreDX DDS

Data Distribution Service

The leading Small Footprint DDS Middleware

Programmer's Guide

Version 5.0

April 2020



Copyright 2008-2020 Twin Oaks Computing, Inc, 230 Third Street, Ste 260 Castle Rock, Colorado 80104 U.S.A. All rights reserved.

This document describes how to install and use the CoreDX DDS software.

CoreDX, CoreDX DDS, and the CoreDX DDS logo are trademarks of Twin Oaks Computing, Inc. Object Management Group, OMG, and DDS are trademarks of the Object Management Group. All other products or company names mentioned are used for identification purposes only, and may be trademarks of their respective owners.

DISCLAIMER OF WARRANTY. THIS DOCUMENT IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Preface

CoreDX DDS is a small-footprint, high-performance communications middleware compliant with the OMG Data Distribution Service (DDS) standard. CoreDX DDS supports multiple hardware architectures and operating systems, and is intended to facilitate the development of robust, near real-time, highly distributed systems.

This manual describes how to install and use the CoreDX DDS software. It is for developers who want to integrate a high-performance, OMG compliant data distribution middleware service into their application.

How this Guide is Organized

This Programmer’s Guide contains a number of Chapters organized in Parts. The first Part provides an overview of the DDS technology and CoreDX DDS. Part 2 provides guidance on installing CoreDX DDS, and walks the reader through creating a simple first CoreDX DDS application.

The next several chapters in Part 3 make up the majority of this document, and go into detail on different aspects of CoreDX DDS features and functionality. This includes: DDS Entities, developing publishing and subscribing applications, Quality of Service (QoS) settings, communication statuses, data instances and samples, data architecture, built-in topics, and the CoreDX DDS transports.

The last few chapters include a discussion of extensions provided by CoreDX DDS such as the logging facility, CoreDX DDS license management, troubleshooting assistance, and finally background about Twin Oaks Computing.

Related Documentation

CoreDX DDS Type System

CoreDX RPC over DDS

CoreDX DDS Secure

CoreDX DDS Reference Manuals

Intended Audience

This document is intended for software developers who are integrating the CoreDX DDS software into their application(s). The guide assumes that the reader is competent in programming languages and software development concepts. CoreDX DDS supports multiple programming languages, and this guide includes examples in C and C++.

Typographic Conventions

Typeface	Meaning	Examples
Courier	Example code	<pre>struct StringMsg { string msg; };</pre>
Courier	Example Commands	<pre>gunzip -c coredx-2.x.tar.gz</pre>

Figure 0-1: Typographic Conventions

Feedback

Twin Oaks Computing welcomes your comments. We are interested in improving our products and we welcome your comments and suggestions. You can provide email feedback about this document to **documents@twinoakscomputing.com**.

Table of Contents

Preface	iii
How this Guide is Organized	iii
Related Documentation.....	iii
Intended Audience	iv
Typographic Conventions	iv
Feedback.....	iv
Part 1: Introduction.....	2
Chapter 1 An Introduction to CoreDX DDS	4
1.1 Why DDS?	4
1.2 The case for Middleware	4
1.3 The case for Publish SubscribeDDS	5
1.4 The case for CoreDX DDS.....	11
Part 2: Getting Started	16
Chapter 2 Installing CoreDX DDS.....	18
2.1 Installation Requirements	18
2.2 CoreDX DDS Distribution Contents.....	21
2.3 CoreDX DDS Installation Procedure.....	23
2.4 Compiling for a different Target Platform.....	23
Chapter 3 First CoreDX DDS Application.....	24
3.1 Using a License	24
3.2 Writing the Application	24
3.3 Compiling Your Application with CoreDX DDS	29
3.4 Running Your Application with CoreDX DDS	33
Chapter 4 Example CoreDX DDS Applications.....	34
4.1 Environment Setup.....	34
4.2 Example 1: The Basic “Hello World” Applications.....	35
4.3 Example 2: Performance Tests	36
4.4 Example 3: Filtering	36

4.5	Example 4: RPC.....	36
4.6	Example 5: QoS Provider.....	36
4.7	Example 6: Security.....	37
4.8	Example 7: Shapes Demonstration	37
Chapter 5	Advanced Compile Options.....	38
5.1	Linux and other UNIX-variant Operating Systems	38
5.2	Windows Operating System.....	42
Part 3:	CoreDX DDS Programming Concepts	49
Chapter 6	DDS Entities.....	52
6.1	DDS Entity Hierarchy	52
6.2	DDS Entity Common Operations	53
6.3	DDS Entity Quality of Service	54
6.4	DDS Status, Listeners, Conditions and WaitSets	54
Chapter 7	Developing a Publishing Application.....	56
7.1	Summary of Developing a Publishing Application	56
7.2	The Data	56
7.3	The Publishing Application.....	56
7.4	Available QoS Settings	62
7.5	Available Listeners	64
Chapter 8	Developing a Subscribing Application.....	68
8.1	Summary of Developing a Subscribing Application	68
8.2	The Data	68
8.3	The Subscribing Application.....	68
8.4	Sample Status Information (SampleInfo).....	75
8.5	Additional Subscriber / DataReader Features	77
8.6	QoS Policies	79
8.7	Available Listeners	81
Chapter 9	Topics	85
9.1	Overview	85
9.2	Built-In Topics.....	86
9.3	Content Filtered Topics	90

9.4	Multi Topics	94
Chapter 10	Instances and Samples	96
10.1	Overview.....	96
10.2	Publishing Data.....	98
10.3	Subscribing to Data.....	99
10.4	Instance Lifecycles	99
10.5	Data Cache.....	103
Chapter 11	Application Data Types.....	108
11.1	Overview.....	108
11.2	Why Define the Data Types?	108
11.3	Registering Data Types	109
11.4	Data Types and Discovery.....	109
11.5	Data Architecture	109
11.6	Compiled Data Type Definition	110
11.7	Data Definition Syntax.....	110
11.8	IDL and XML Language Mappings.....	112
11.9	Creating Data Types.....	113
11.10	Configuring Data Types.....	116
11.11	Using the CoreDX DDS Data Type Compiler	120
Chapter 12	Configuring Reader Specific Locators.....	125
12.1	Overview.....	125
12.2	Configuration.....	125
12.3	Generated Code.....	126
Chapter 13	Quality of Service Features	128
13.1	QoS Compatibility Between Publishing and Subscribing Entities.....	129
13.2	QoS Mutability.....	130
13.3	Quality of Service Details.....	130
Chapter 14	Communication Status	147
14.1	Communication Status Details	149
14.2	Application Access to Communication Status	161
Part 4:	CoreDX DDS Extensions	177

Chapter 14	CoreDX DDS Logging	179
14.1	Logging Libraries	179
14.2	Logging Control	179
14.3	Logging Flags	180
Chapter 15	CoreDX DDS Transport	183
15.1	Overview	183
Chapter 16	CoreDX DDS Discovery	200
16.1	Overview of CoreDX DDS Discovery	200
16.2	Discovering DomainParticipants	201
16.3	Matching DataReaders and DataWriters	202
16.4	Static Discovery	204
16.5	Centralized Discovery	206
16.6	Wait for Discovery	210
16.7	Access to Discovery Information	211
16.8	Discovery and Deterministic Machines	213
Chapter 17	Configuring Reliability Protocol	218
17.1	Reliability Protocol	218
17.2	Reliability QoS Configuration	224
Chapter 18	Configuring Reader Specific Locators	227
18.1	Overview	227
18.2	Configuration	227
Chapter 19	Dynamic Types	229
19.1	Overview	229
19.2	Subscribe with Dynamic Types	229
19.3	Publish with Dynamic Types	230
Chapter 20	Threading Options	231
20.1	Overview	231
20.2	Configuring Threading Options	231
Chapter 21	Transmit Buffers	235
21.1	Overview	235
21.2	Dynamic Transmit Buffers	235

21.3	Static Transmit Buffers	237
Chapter 22	Receive Buffers.....	239
22.1	Overview.....	239
22.2	Configuration.....	239
Chapter 23	Data Batching	242
Chapter 24	QoS Provider.....	244
24.1	Introduction.....	244
24.2	QoS Provider Usage	244
24.3	QoS Provider XML Syntax	247
24.4	QoS Provider Integration.....	256
24.5	Unsupported.....	256
Chapter 25	Licensing	257
25.1	Development Licenses.....	257
25.2	Run-time Licenses.....	257
Chapter 26	Transition Notes	261
26.1	CoreDX DDS v3.x -> v4.0	261
Chapter 27	Troubleshooting	262
27.1	General Troubleshooting Tools	262
27.2	No Communications between DDS applications.....	262
27.3	Missing or lost samples	263
27.4	TypeSupport version mismatch.....	265
27.5	Can’t find it here?.....	266
Chapter 28	About Twin Oaks Computing.....	268
Chapter 29	Contact Information	270

Table of Figures

Figure 0-1: Typographic Conventions.....	iv
Figure 1-1: Middleware	5
Figure 1-2: Client Server Architecture	5
Figure 1-3: Publish Subscribe Architecture	6
Figure 1-4: Example DDS Usage	7
Figure 1-5: DDS Architecture	10
Figure 2-1: CoreDX DDS Directory Structure	22
Figure 6-1: DDS Entity Hierarchy	52
Figure 10-1: Register Instances Example.....	100
Figure 11-1: Example IDL file	114
Figure 11-2: IDL keys example	120
Figure 13-1: Configuring QoS - Sample C Code	128
Figure 14-1: Inconsistent Topic Status Structure	150
Figure 14-2: Sample Rejected Status Structure.....	152
Figure 14-3: Liveliness Changed Status Structure	153
Figure 14-4: Requested Deadline Missed Status Structure.....	154
Figure 14-5: Requested Incompatible QoS Status Structure.....	155
Figure 14-6: Sample Lost Status Structure	156
Figure 14-7: Subscription Matched Status Structure	157
Figure 14-8: Liveliness Lost Status Structure.....	158
Figure 14-9: Offered Deadline Missed Status Structure.....	159
Figure 14-10: Offered Incompatible QoS Status Structure	160
Figure 14-11: Publication Matched Status Structure	161
Figure 14-12: Listener Hierarchy	163
Figure 14-13: Listener Example C Code.....	167
Figure 14-14: Listener Exampe C++ Code.....	168
Figure 14-15: Condition Example C code	174
Figure 14-16: Condition Example C++ code	175
Figure 17: Standard Discovery (peer-to-peer) architecture	207
Figure 18: Centralized Discovery architecture	208
Figure 19: Example Centralized Discovery deployment	210
Figure 17-1: Example DDS Usage.....	220
Figure 17-2: Example DDS Usage.....	220
Figure 25-1: Example CoreDX DDS license file	257
Figure 27-1: Example CoreDX DDS license file	266

Table of Tables

Table 2-1: CoreDX DDS Architectures and Operating Systems	18
Table 2-2: CoreDX DDS Languages and Compilers	20
Table 3-1: Sample IDL File	25
Table 4-1: Example - Running cdxenv.sh	35
Table 5-1: CoreDX DDS Libraries (UNIX Operating Systems)	38
Table 5-2: CoreDX DDS Libraries (Windows Operating System)	42
Table 5-3: Windows Dynamic Library Dependencies	46
Table 7-1: QoS Policies for Publishing Entities	62
Table 7-2: Listeners for Publishing Entities	65
Table 8-1: QoS Policies for Subscribing Entities	79
Table 8-2: Listeners for Subscribing Entities	81
Table 9-1: Topic Variants	85
Table 9-2: Built-in Topics	86
Table 9-3: Participant Built-in Data Type	88
Table 9-4: Topic Built-in Data Type	88
Table 9-5: Publication Built-in Data Type	89
Table 9-6: Subscription Built-in Data Type	89
Table 9-7: create_contentfilteredtopic() parameters	90
Table 9-8: Valid Condition Operators for Content Filters	91
Table 9-9: Creating a ContentFilteredTopic	92
Table 10-1: Instance Example	98
Table 10-2: Instance Example	103
Table 11-1: Basic User Defined Types	110
Table 11-2: Constructed User Defined Types	111
Table 11-3: Primitive Data Type Mapping	112
Table 11-4: coredx_ddl command line options	120
Table 12-1: Generated source code file names	126
Table 13-1: Standard QoS Policies Summary	130
Table 13-2: CoreDX DDS QoS Policies	132
Table 14-1: Communication Statuses	147
Table 14-2: Listener Method Signatures	164
Table 14-2: Logging QoS Configuration Example	180
Table 14-1: CoreDX DDS Logging Flags	181
Table 15-1: UDP Transport Multicast and Unicast Configuration Parameters	189
Table 15-2: UDP Transport Environment Variables	191
Table 15-3: TCP Transport Environment Variables	196
Table 15-4: LMT Transport Environment Variables	198
Table 16-1: Code Example of peer_participants QoS	205

Table 16-2: API Methods to Access Discovery Information.....	212
Table 17-1: CoreDX DDS RTPS_Protocol QoS Policy	224
Table 21-1: Transmit Buffer Configuration	235

Part 1: Introduction

This section provides an introduction of the Data Distribution Service (DDS) and the CoreDX DDS implementation from Twin Oaks Computing, Inc.

Chapter 1 An Introduction to CoreDX DDS

Welcome to CoreDX DDS, a high-performance, small-footprint implementation of the OMG Data Distribution Service (DDS) standard. The CoreDX DDS Data-Centric, Publish-Subscribe messaging infrastructure provides high-throughput, low-latency data communications.

This chapter provides an overview of the Data Distribution Service (DDS), how applications might use DDS to meet their communication requirements, and features of the CoreDX DDS product.

1.1 Why DDS?

Today's enterprise systems, embedded systems, and all systems in between, need flexible, open information systems. Most systems span multiple technologies, hardware platforms, operating systems, and programming languages. In addition, components of these systems have real-time requirements. CoreDX DDS is an open standards-based, communication middleware solution to meet the needs of these real-time distributed systems.

1.2 The case for Middleware

Middleware is a class of software that exists between an application and the Operating System. In deeply embedded environments, middleware exists between the functional software and a network stack of the device. It provides useful capabilities that are above and beyond those found in standard Operating Systems. In the case of CoreDX DDS, the middleware provides a facility for both publish-subscribe and client-server communications. Figure 1-1 illustrates where middleware components fit in the application, and how they logically bridge across multiple operating systems and hardware architectures.




Figure 1-1: Middleware

Applications that employ a communications middleware like CoreDX DDS realize many benefits. The requirements and complexity of data communications in a distributed system are met by the middleware component - leaving developers more time to focus on the important application logic. CoreDX DDS middleware supports many operating systems and hardware architectures - the task of porting complex communications software is already complete.

1.3 The case for Publish SubscribeDDS

Many communication middleware technologies are available. Most are based on a functional model. For example, RPC (Remote Procedure Call) and CORBA (Object Request Broker) are two examples of middleware that allow function calls to be distributed across the network between a client and a server. However, these architectures lead to tight coupling between the client and the server; this makes these systems difficult to extend.




Figure 1-2: Client Server Architecture

The client-server architecture is appropriate for centralized data processing and works well in some systems and some use cases. In some client-server technologies, the drawbacks are increased integration costs for new capabilities and potential single point of failure.

An alternative to this approach is the Publish-Subscribe architecture embodied in DDS. This architecture promotes a loose coupling between data producers and data consumers. The architecture is flexible and dynamic; it is easy to adapt and extend systems to changing environments and requirements. Figure 1-3 illustrates the DDS Publish Subscribe architecture where multiple Publishers and Subscribers exchange strongly typed data through a common Topic. The communications are controlled by a Quality of Service model.




Figure 1-3: Publish Subscribe Architecture

Figure 1-4 is an example of how DDS might be applied in a system. This example has several sources of “raw data”, a data processor that performs some processing on the raw data to produce “processed data”, several end users working with the processed data, and an administrative user performing analysis, maintenance, or auditing functions.

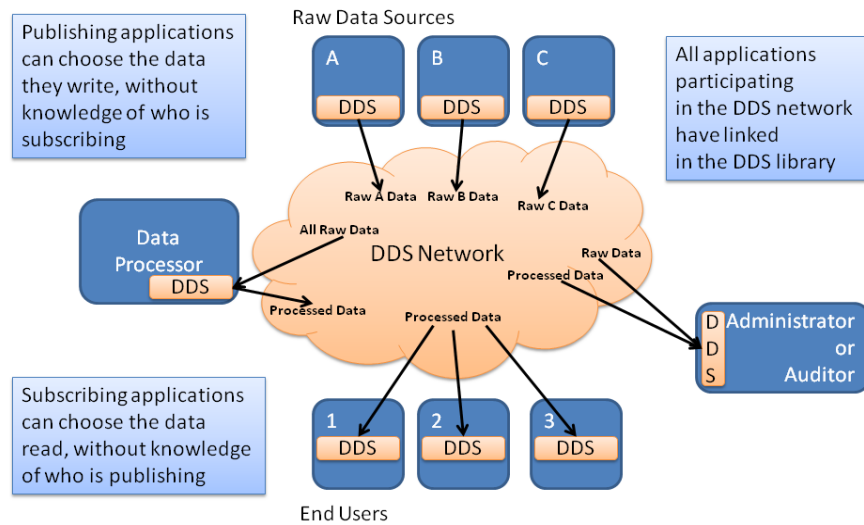


Figure 1-4: Example DDS Usage

In this example, the darker blue boxes represent applications communicating over a DDS network. These applications might be running together on 1 host, or they might be distributed over multiple hosts. A DDS application simply publishes or subscribes to their data, without concern for what, if anything, might be on the other end of its communications. Any of the applications can be dynamically removed (and new applications may be added) without impacting the existing network.

Because many systems include some natural publish-subscribe use cases as well as some natural client-server use cases, the DDS standards include both communication mechanisms. This document focuses on the publish-subscribe interface to DDS. For information on programming with the remote procedure call (RPC) or request reply APIs, please use the CoreDX DDS RPC over DDS Programmer’s Guide.

1.3.1 DDS is an Open Standard

DDS is an open specification (documented by multiple standards) managed by the Object Management Group (OMG). The OMG is an international, open membership, non-profit organization that develops and manages computer industry specifications. Hundreds of organizations, including software end-users and commercial vendors, make up the OMG. Together they develop and manage many of the standards widely used in the computer industry today. The set of Data Distribution Service (DDS) standards is an example of one of the technology standards managed by the

OMG. Other examples include the Unified Modeling Language (UML), Model Driven Architecture (MDA) and the Common Object Request Broker Architecture (CORBA).

There are several advantages to using a technology that conforms to an open standard, and more advantages if that open standard is managed by an open membership organization like the OMG. First, an open standard promotes interoperability. Anyone, even if they are not connected with the managing organization, can pick up an Open Standard and write a conforming application. Second, open standards reduce the dependence on a particular vendor. When an open standard product is available from multiple vendors, the consumer can easily change between them. Finally, anyone can join the managing organization and vote on the direction and advancement of the technology. In the case of DDS, this means vendors and users, both public and private, can influence the future of the technology.

1.3.2 DDS is More than a Communications Middleware

The DDS standards specify the mechanism for moving data – a typical communications middleware technology standard. However, DDS is so much more. In addition to communications, DDS provides advanced data management, storage, organization, filtering, redundancy, extensibility, and security. With a rich set of features, interoperability across languages, operating systems, hardware platforms, and implementations, DDS provides a robust, secure infrastructure foundation for your small-scale, large-scale, enterprise, embedded, and everything in between software system.

1.3.3 Remote Procedure Call (RPC) in addition to Publish-Subscribe

The Data Distribution Service is a publish-subscribe technology, which provides a flexible, loosely coupled architecture suitable for many real-time applications. However, many sophisticated projects are a natural mix of publish-subscribe and client-server (or RPC) requirements.

The DDS Standards include API's for publish-subscribe, request-response, and RPC – all implemented on top of the original DDS publish-subscribe architecture. DDS request-response and RPC have some unique features over other client-server type middleware, including automatic discovery, security, and the ability to use the full set of DDS Quality of Service (QoS) configurations.

The CoreDX DDS RPC API is fully described in the CoreDX DDS RPC Programmer’s Guide.

1.3.4 *DDS is flexible and scalable*

Applications communicating with DDS might be running together on 1 host, or they might be distributed over multiple hosts, each with different architectures and operating systems. Applications using DDS for communications do not need to know the details of where there other applications are residing, or even if they exist.

The *discovery* mechanism built into DDS allows applications to come and go from a DDS network without requiring any changes to the applications or the network. This means a new system can be brought into the network, and start sending or receiving data, without any changes to existing applications.

1.3.5 *DDS is secure*

The DDS Security standard contains a complete state-of-the-art security solution that is completely integrated into the DDS protocols (**not** simply layered on top of SSL). DDS Security includes: Identification, Authentication, Access Control, Integrity, and Confidentiality, allowing the designer full flexibility on a topic-by-topic level.

Security configuration and usage is documented in the CoreDX DDS Security Programmer’s Guide.

1.3.6 *DDS Features*

A DDS application can be a publisher of data, a subscriber of data, or both.

A **Publisher** is responsible for data distribution. It may publish data of different data types. The application uses a typed **DataWriter** attached to the publisher to communicate the data to be published. Both the Publisher and the DataWriter have a Quality of Service (**QoS**) that affects the behavior of the publication.

A **Subscriber** is responsible for receiving published data and making it available to the receiving application. It may receive data of different data types. The application uses a typed **DataReader** attached to the subscriber to access the data. Both the Subscriber and DataReader have a QoS that affects the behavior of the subscription. The subscribing application can

choose to block waiting for data using **WaitSets** or receive data asynchronously, using **Listeners**.

A Topic fits between publications and subscriptions. Subscriptions must be able to refer to specific publications. A topic fulfills this purpose: it associates a name, a data-type, and a QoS related to the data itself.

When an application wants to publish data of a given type, it must use a Publisher and DataWriter with all the characteristics of the desired publication. When an application wants to subscribe to data of a given type, it must use a Subscriber and DataReader with all the characteristics of the desired subscription.

The following figure depicts the common DDS objects used in exchanging data.

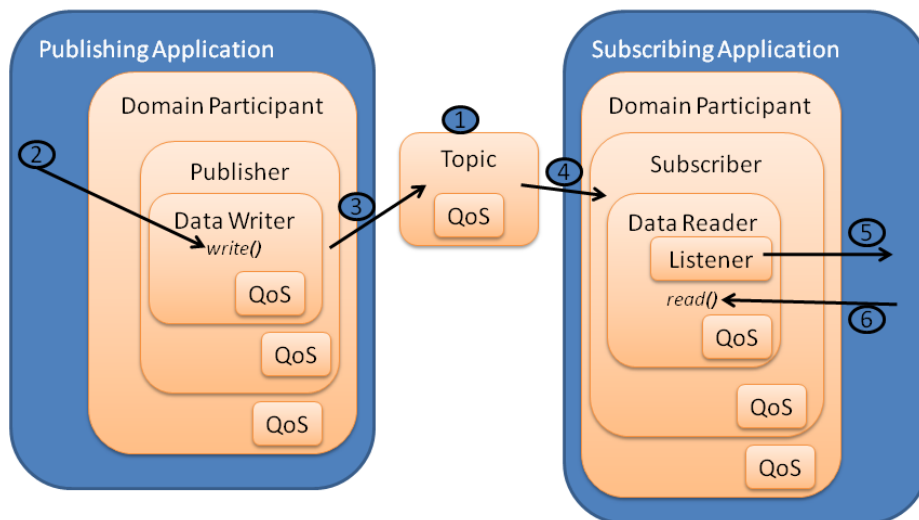


Figure 1-5: DDS Architecture

The following describes the actions depicted in Figure 1-5.

1. *DataReaders* and *DataWriters* are associated with a *Topic*
2. The publishing application calls *DataWriter::write()* to write the data
3. The *Publisher* publishes the data
4. The *Subscriber* receives the data
5. The *Listener* notifies the subscribing application of available data

6. The subscribing application calls *DataReader::read()* to access the data

1.4 The case for CoreDX DDS

The CoreDX DDS provides a quality, high-performance, very small footprint implementation of the DDS standards, including the original publish-subscribe DDS API, RTPS wire protocol, X-Types, RPC over DDS, and DDS Security.

1.4.1 CoreDX DDS is Fast

CoreDX DDS was built from the ground up with performance in mind. The engineering staff at Twin Oaks Computing has a long history of writing and maintaining real-time and near real-time software, and this expertise was used in creating CoreDX DDS. CoreDX DDS is written in ‘C’ (with additional application language bindings available) for low overhead and memory savings. The CoreDX DDS baseline is tested and enhanced for performance at every step of the development process. The result is a quality DDS implementation with extremely low latency and high throughput capacity.

CoreDX DDS data aggregation, multi-core data pipeline, and low latency event notification provide for throughput in the +900Mbps range and latencies below 75 usec over a 1Gbps ETHERNET network. But don’t take our word for it. The CoreDX DDS release includes source code for example benchmarking applications. Use these examples to compile your own benchmark tests and see how CoreDX DDS performs in your environment, with your data.

1.4.2 CoreDX DDS is Small

The CoreDX DDS product is 100% designed and developed by Twin Oaks Computing to meet the OMG’s DDS specification. There is no historical code, no code borrowed from the open source community, no code retrofitted to meet the CoreDX DDS requirements. This allows us to deliver a quality, fully-functional DDS implementation with the smallest footprint. Our entire core library is less than 500 KB, and runs in environments with as little as 100 KB of RAM. The full CoreDX DDS implementation is deployed on FPGA’s, DSP’s, PLC’s, ECU’s and other embedded environments.

This small library size comes with a proportionally small Line of Code Count, perfect for safety critical applications requiring DO-178B certification.

CoreDX DDS is modular and contains additional run-time memory tuning parameters. Space constrained projects can select components of CoreDX DDS to meet their requirements, and tune those components to reduce unnecessary memory utilization.

For those environments that are even smaller: true microcontrollers, CoreDX DDS Micro requires no more than 8K of RAM, allowing the benefit of the interoperability DDS protocols down to the component level of any system.

CoreDX DDS Micro is documented in the CoreDX DDS Micro Programmer's Guide.

1.4.3 CoreDX DDS is Proven & Robust

The small footprint CoreDX DDS software has over 10 years of deployment usage in a wide variety of mission-critical, and business-critical applications.

With over 1 million deployed instances around the world and in space, connecting components in surgical devices, military and commercial vehicles, space exploration platforms, electrical grids, CoreDX DDS has a proven track record of reliability, robustness, and competent technical and business support.

1.4.4 CoreDX DDS is Secure

CoreDX DDS complies with the DDS Security standards, providing integrated and sophisticated security features that are fully configurable. Using state-of-the-art security algorithms, The DDS Security standard was designed to meet the requirements of military and critical national infrastructure systems.

System designers may choose to use the standards compliant Twin Oaks Computing developed security plug-ins for identification, authentication, access control, integrity, and confidentiality, or develop their own with the standardized plug-in API.

1.4.5 CoreDX DDS Uses Multi-Core Technologies

Hardware is moving to multiple core technology. Even embedded processors are shipping with more than one core. This presents a challenge to application developers, because making use of multiple cores requires complex code that is difficult and expensive to develop and maintain. The solution: use a multithreaded communications middleware like CoreDX DDS.

CoreDX DDS was architected from the start to take advantage of multi-core environments. With advanced threading and protections, each CoreDX DDS participant will use a minimum of 3 cores, and typical CoreDX DDS applications will use between 4 and 8 cores. These are single threaded applications, taking advantage of quad-core and higher hardware, just by using CoreDX DDS for data communications.

1.4.6 CoreDX DDS is Self Contained

In order to use CoreDX DDS for communications, the application links in the appropriate CoreDX DDS libraries and that is it. With no daemons and no operating system services that need to be started and maintained, there is no place for data to become “stuck” or for communication states to become corrupted.

1.4.7 CoreDX DDS has Comprehensive Platform Support

With the wide array of language binding, operating system and architecture support, CoreDX DDS runs on a wide variety of platforms, from enterprise servers, to common desktop configurations, to embedded environments and real-time operating systems, to FPGA’s and ‘bare-metal’ configurations. See the *Installation Requirements* section for details on supported platforms.

If you don’t find your specific platform listed, just contact us. We offer extensive engineering services, including (often free!) custom ports to new Operating Systems and Architectures, and well as language ports. And with our low line of code count, custom porting is quick and easy.

1.4.8 CoreDX DDS has a great team behind it

A quality DDS implementation is important. But the organization behind the implementation is critical. When you make a commitment to purchase a software product, you are not only obtaining the rights to run the software contained on the installation disk (or downloaded from the web). You are also obtaining support services, training services, and product enhancements for at least the next year.

The staff at Twin Oaks Computing has been developing and supporting large software systems and global software companies for over 50 years. We have worked beside soldiers in Kuwait, sailors onboard aircraft carriers, and other warfighters around the world. We have supported commercial IoT and IIoT companies with millions of products deployed world-wide. We

understand not only the importance of delivering a software product that works, but also the importance of helping companies and their end users make the most of their investment.

We will do the same for you. Give us a call or send us an email. We promise you will receive prompt, friendly, and helpful service.

Part 2: Getting Started

The Getting Started section includes information to establish a development environment and build a simple DDS publisher and subscriber. This provides a quick overview of the development process and the associated tools.

Chapter 2 Installing CoreDX DDS

This chapter explains how to install CoreDX DDS onto your development system.

2.1 Installation Requirements

2.1.1 Supported Architectures and Operating Systems:

CoreDX DDS is ported and built for additional environments on a regular basis. The following table is intended to provide the reader with an example set of supported platforms – it is not a complete list.

Table 2-1: CoreDX DDS Architectures and Operating Systems

Operating System	Architecture	Build Tools
Linux 2.6+	x86 (32bit), x86_64, atom	gcc-4.3.2+ / glibc 2.8+
	ARMv5, ARMv7, ARMv7-A Raspberry Pi, aarch64	gcc-4.1.2+
	PPC 750, 7400, 440, e500	gcc
	PPC32	gcc
	MIPS 32, MIPS 64	gcc-3.4 / uclibc 0.9.29
	Microblaze, Nios II (soft CPU)	gcc-4.1
OSX, various	x86, x86_64	Llvm
Windows, various	x86 (32bit), x86_64	Visual Studio 2015, 2012, 2010, 2008, 2005
	x86_64 (64bit)	Visual Studio 2015, 2012, 2010, 2008
		MinGW / gcc-3.4.5

Operating System	Architecture	Build Tools
WinCE, Windows Embedded	X86	Visual Studio 2008
	Arm	Visual Studio 2008
QNX 6.5, 6.6, 7	x86 (32bit)	gcc-4.2.4
	ARM v5	gcc
VxWorks 5.5	x86 (32bit)	gcc-4.1.2
	ARM v7	gnu
	PPC405ep	Diab
VxWorks 6.6, 6.8, 6.9 (RTP and DKM)	x86 (32bit)	gcc-4.1.2
	PPC32	gcc
VxWorks-653	PPC	gcc
NexusWare 12.3	x86	gcc
	PowerPC 440gx	gcc
LynxOS SE 6.0	x86	gcc-4.3
LynxOS 178	PPC	gcc
INTEGRITY-178B tuMP	PPC, arm	gcc
INTEGRITY 11	PPC, x86, arm	gcc
DeOS	x86	gcc
uClinux	NIOS2 (softcpu)	gcc
ThreadX	x86	gcc

Operating System	Architecture	Build Tools
	PPCe300	gcc
Android 2.x->8.x	x86	gcc
	ARM v5	gcc
iOS	ARM v7	Apple
Free-RTOS	ARM	gcc
DSP/BIOS	TI DSP	TI

2.1.2 Supported Languages and Compilers

CoreDX DDS is ported and built for additional environments on a regular basis. The following table is intended to provide the reader with an example set of supported platforms – it is not a complete list.

Table 2-2: CoreDX DDS Languages and Compilers

Language	Compiler	Compiler Version
C	gcc	multiple
	MS Visual Studio	VS2012, VS2015+
	MinGW/gcc	3.4.5
	WindRiver Diab	
	Green Hills Multi	

C++	g++	multiple
	MS Visual Studio	VS2012, VS2015+
	Wind River Diab	
Cxx	g++	v11 +
Java	javac	1.5+
C#	Mono (Linux)	2.4
	Visual Studio	VS2008+
Python		

2.2 CoreDX DDS Distribution Contents

The CoreDX DDS distribution includes a top-level directory: **coredx-version**. This top-level directory is referred to throughout this document as **COREDX_TOP**, and contains the following files and subdirectories:

COPYRIGHT	:	File(s) describing the Copyright information for this CoreDX DDS baseline
examples	:	Contains example CoreDX DDS applications
host	:	Contains the files required for the HOST (or build) machines for all installed platforms
README	:	File(s) describing the CoreDX DDS version and build
RELEASE_NOTES	:	Information on changes from previous CoreDX DDS versions
scripts	:	Contains helpful, platform independent, scripts
target	:	Contains the files required for the TARGET (or deployment)

	machines for all installed platforms.
--	---------------------------------------

The host and target subdirectories contain the CoreDX DDS libraries and binaries in platform specific directories. This configuration allows you to install CoreDX DDS for multiple platforms in one COREDX_TOP directory. Figure 2-1 shows the directory structure under COREDX_TOP.




Figure 2-1: CoreDX DDS Directory Structure

The *host* subdirectory contains the HOST tools for CoreDX DDS. The HOST tools include CoreDX DDS DDL compiler and the Twin Oaks Computing license hostid utility for all the installed platforms. The host subdirectory contains two copies of these utilities for each architecture installed. For example, consider the coredx_ddl utility for an x86 Linux system using the gcc4.3 compiler. This utility can be found in two locations:

```

COREDX_TOP/host/bin/Linux_2.6_x86_gcc5_coredx_ddl
COREDX_TOP/host/Linux_2.6_x86_gcc5/bin/coredx_ddl

```

The *target* subdirectory contains the TARGET tools for CoreDX DDS. The TARGET tools include the CoreDX DDS header files, the CoreDX DDS library files, and the Twin Oaks Computing license hostid utility. The CoreDX DDS libraries are located architecture-named subdirectories. For example, consider the DDS library for an x86 Linux system using the gcc5 compiler. This library is located:

```

COREDX_TOP/target/Linux_2.6_x86_gcc5/lib/libdds.a

```

The Compiling Your Application with CoreDX DDS section provides examples for configuring Makefiles to use the CoreDX DDS directory structure.

2.3 CoreDX DDS Installation Procedure

Once you have obtained CoreDX DDS from Twin Oaks Computing (or from an Evaluation CD or USB drive), unpack the distribution somewhere on your system. For example, on a UNIX system this command will extract the distribution into the current directory:

```
gunzip -c coredx-4.x.x-{platform}.tgz | tar xf -
```

Or, for Windows:

```
unzip coredx-4.x.x-{platform}.zip (It's OK to overwrite files if prompted here.)
```

And that's it. There is no additional configuration required.

Customers installing CoreDX DDS for multiple platforms can unpack all the CoreDX DDS releases into the same COREDX_TOP directory. CoreDX DDS uses platform-specific directory names in order to avoid conflicts when working with multiple operating systems and hardware architectures.

2.4 Compiling for a different Target Platform

CoreDX DDS supports cross compiling for a different target platform. For example, if you are developing an INTEGRITY application, you might be developing (compiling) on a Windows/x86 platform and targeting a INTEGRITY/arm platform. Each platform release of CoreDX DDS contains both the HOST and TARGET tools for one platform, so in the above example, you will require two platform versions of CoreDX DDS, one for the HOST (development) platform, and another for the TARGET (run-time) platform. All platform versions of CoreDX DDS may be installed into the same COREDX_TOP directory.

The data type compiler (coredx_ddl) is a HOST tool that generates code to be run on the TARGET platform. When the endian of the HOST is different than the endian of the TARGET, it is important to notify the data type compiler, so that it can generate the correct marshal and un-marshal code for the TARGET platform. This is done using the “-e” option to the data type compiler. See the Using the CoreDX DDS Data Type Compiler section for additional information on the CoreDX DDS data type compiler.

Chapter 3 First CoreDX DDS Application

This chapter describes how to use the CoreDX DDS tools and libraries to integrate basic DDS capabilities into your application. We've provided a sample data type and application that is ported to all CoreDX DDS supported languages and platforms with the distribution (located in `CORED_X_TOP/examples`). You can use these examples or create your own while going through the following steps.

Our example Linux Makefiles were built for `gcc / g++`. The examples contain different Makefiles for additional platforms. All example Makefiles use three environment variables: `CORED_X_TOP`, `CORED_X_HOST`, and `CORED_X_TARGET`. The environment script provided with CoreDX DDS releases can help determine the correct settings for these variables (`CORED_X_TOP/scripts/cdxenv.sh` and `cdxenv.bat`).

3.1 Using a License

Compiling with CoreDX DDS requires a development or evaluation license. The license may be set using environment variables or software API. (Software API is described in the *Licensing* section of this guide.) Use the environment variable: `TWINOAKS_LICENSE_FILE` to set the location of your CoreDX DDS license.

For example:

```
linux% export TWINOAKS_LICENSE_FILE=LICENSE_FILE
```

or

```
windows> set TWINOAKS_LICENSE_FILE=LICENSE_FILE
```

3.2 Writing the Application

While CoreDX DDS provides a consistent looking API across all language bindings, there are slight differences in the code generation and compiling instructions.

3.2.1 The ‘C’ Language Application

Create the data definition for the data type(s) you will use for communications. The syntax used for data definition complies with the OMG IDL v4.0 syntax for describing data types.

Note: There are references to “DDL” and “IDL” throughout this document and the CoreDX DDS examples. Previous versions of CoreDX DDS used the term “Data Definition Language” or “DDL” to describe data types. DDL simply refers to the subset of IDL that defines data types.

Here is the “hello world” example provided with the distribution:

Table 3-1: Sample IDL File

```
hello.ddl

struct StringMsg
{
    string msg;
};
```

Compile the IDL to generate the type specific code using the CoreDX DDSdata type compiler. This requires your TWINOAKS_LICENSE_FILE environment variable be set as described above. Assuming the name of the IDL file is hello.ddl:

```
% COREDX_TOP/host/COREDX_HOST/bin/coredx_ddl -f
hello.ddl
```

The compilation will generate the following files (names are based on the IDL filename):

```
hello.h and .c
helloTypeSupport.h and .c
helloDataReader.h and .c
helloDataWriter.h and .c
```

Create code to publish data of this data type. Our sample Hello World publisher is located in COREDX_TOP/examples/hello_c/hello_pub.c.

Create code to subscribe to data of this data type. Our sample Hello World subscriber is located in COREDX_TOP/examples/hello_c/hello_sub.c.

3.2.2 *The 'C++' Language Application*

Create the data definition for the data type(s) you will use for communications. The syntax used for data definition complies with the OMG IDL v4.0 syntax for describing data types. The C++ application uses the same IDL type definition as the C program in the previous section.

Compile the IDL to generate the type specific code using the CoreDX DDSdata type compiler. This requires your TWINOAKS_LICENSE_FILE environment variable be set as described above. Assuming the name of the IDL file is hello.ddl:

```
% COREDX_TOP/host/COREDX_HOST/bin/coredx_ddl -l cpp -f
hello.ddl
```

The compilation will generate the following files (names are based on the IDL filename):

```
hello.hh and .cc
helloTypeSupport.hh and .cc
helloDataReader.hh and .cc
helloDataWriter.hh and .cc
```

Create code to publish data of this data type. Our sample Hello World publisher is located in COREDX_TOP/examples/hello_cpp/hello_pub.cc.

Create code to subscribe to data of this data type. Our sample Hello World subscriber is located in COREDX_TOP/examples/hello_cpp/hello_sub.cc.

3.2.3 *The 'Cxx' (C++ Modern) Language Application*

Create the data definition for the data type(s) you will use for communications. The syntax used for data definition complies with the OMG IDL v4.0 syntax for describing data types. The Cxx application uses the same IDL type definition as the C program in the previous section.

Compile the IDL to generate the type specific code using the CoreDX DDSdata type compiler. This requires your TWINOAKS_LICENSE_FILE environment variable be set as described above. Assuming the name of the IDL file is hello.ddl:


```
% COREDX_TOP/host/COREDX_HOST/bin/coredx_ddl -l cxx -f  
hello.ddl
```

The compilation will generate the following files (names are based on the IDL filename):

```
hello.hh and .cc  
helloTypeSupport.hh and .cc  
helloDataReader.hh and .cc  
helloDataWriter.hh and .cc
```

Create code to publish data of this data type. Our sample Hello World publisher is located in COREDX_TOP/examples/hello_cxx/hello_pub.cc.

Create code to subscribe to data of this data type. Our sample Hello World subscriber is located in COREDX_TOP/examples/hello_cxx/hello_sub.cc.

3.2.4 The ‘Java’ Language Application

Create the data definition for the data type(s) you will use for communications. The syntax used for data definition complies with the OMG IDL v4.0 syntax for describing data types. The Java application uses the same IDL type definition as the C program in the previous section.

Compile the IDL to generate the type specific code using the CoreDX DDS data type compiler. This requires your TWINOAKS_LICENSE_FILE environment variable be set as described above. Assuming the name of the IDL file is hello.ddl:

```
% COREDX_TOP/host/COREDX_HOST/bin/coredx_ddl -l java -f  
hello.ddl
```

The compilation will generate the following files (names are based on the data type(s) defined in the IDL file):

```
StringMsg.java  
StringMsgTypeSupport.java  
StringMsgDataReader.java  
StringMsgDataWriter.java
```

Create code to publish data of this data type. Our sample Hello World publisher is located in COREDX_TOP/examples/hello_java/HelloPub.java.

Create code to subscribe to data of this data type. Our sample Hello World subscriber is located in COREDX_TOP/examples/hello_java/HelloSub.java.

3.2.5 *The 'C#' Language Application*

Create the data definition for the data type(s) you will use for communications. The syntax used for data definition complies with the OMG IDL v4.0 syntax for describing data types. The C# application uses the same IDL type definition as the C program in the previous section.

Compile the IDL to generate the type specific code using the CoreDX DDS data type compiler. This requires your TWINOAKS_LICENSE_FILE environment variable be set as described above. Assuming the name of the IDL file is hello.idl:

```
% COREDX_TOP/host/COREDX_HOST/bin/coredx_ddl -l csharp -
f hello.idl
```

The compilation will generate the following files (names are based on the data type(s) defined in the IDL file):

```
StringMsg.cs
StringMsgTypeSupport.cs
StringMsgDataReader.cs
StringMsgDataWriter.cs
```

Create code to publish data of this data type. Our sample Hello World publisher is located in COREDX_TOP/examples/hello_csharp/hello_pub.cs.

Create code to subscribe to data of this data type. Our sample Hello World subscriber is located in COREDX_TOP/examples/hello_csharp/hello_sub.cs.

3.2.6 *The 'Python' Language Binding*

Create the data definition for the data type(s) you will use for communications. The syntax used for data definition complies with the OMG IDL v4.0 syntax for describing data types. The Python application uses the same IDL type definition as the C program in the previous section.

Compile the IDL to generate the type specific code using the CoreDX DDS data type compiler. This requires your TWINOAKS_LICENSE_FILE environment variable be set as described above. Assuming the name of the IDL file is hello.idl:

```
% COREDX_TOP/host/COREDX_HOST/bin/coredx_ddl -l python -
f hello.ddl
```

The compilation will generate the following file:

```
__init__.py
```

Create code to publish data of this data type. Our sample Hello World publisher is located in COREDX_TOP/examples/hello_csharp/hello_pub.py.

Create code to subscribe to data of this data type. Our sample Hello World subscriber is located in COREDX_TOP/examples/hello_csharp/hello_sub.py.

3.3 Compiling Your Application with CoreDX DDS

Compile your application(s). Our Hello World example creates two applications, one for the publisher and one for the subscriber. This is not necessary, and is completely dependent on your application architecture. Your application will require the objects from the generated type support code above, as well as your publisher and/or subscriber code.

The examples in this section assume a UNIX-based operating system. Compiling for other operating systems may require additional considerations. Please refer to Part 2:Chapter 5: *Advanced Compile Options* for additional information on compiling CoreDX DDS applications for your specific operating system.

3.3.1 The ‘C’ Language Application

CoreDX DDS will require the following paths and libraries when compiling your application:

```
Include Path: \
               -I${COREDX_TOP}/target/include
Library Path: \
               -L${COREDX_TOP}/target/${COREDX_TARGET}/lib
Static Libraries: -ldds
```

We’ve provided Makefiles and Visual Studio project files for compiling our examples, and you can use these as a reference for compiling your application. Our Makefiles require three environment variables:

1. COREDX_TOP
2. COREDX_HOST
3. COREDX_TARGET

The COREDX_TOP is a path name to the location of your CoreDX DDS distribution(s). COREDX_HOST and COREDX_TARGET are the platform architectures you are compiling on and compiling for (these can be the same). These values can be set manually, or determined by running the script: COREDX_TOP/scripts/cdxenv.sh or cdxenv.bat.

Our Makefiles will run the CoreDX DDS data type compiler to generate the type specific code as well as compile the applications. To compile the Hello World sample application using our Makefile you will need a make program (for example gnu make or Microsoft nmake) and the compiler (for example, gcc or cl.exe) in your path. Then, simply type 'make' (or 'nmake -f NMakefile') in the appropriate directory. This will compile two applications: hello_pub and hello_sub.

3.3.2 The 'C++' Language Application

CoreDX DDS will require the following paths and libraries when compiling your application:

```
Include Path: \
               -I${COREDX_TOP}/target/include
Library Path: \
               -L${COREDX_TOP}/target/${COREDX_TARGET}/lib
Static Libraries: -lddscpp-ldds
```

We've provided Makefiles for compiling our examples, and you can use these as a reference for compiling your application. Our Makefile requires three environment variables:

1. COREDX_TOP
2. COREDX_HOST
3. COREDX_TARGET

The COREDX_TOP is a path name to the location of your CoreDX DDS distribution(s). COREDX_HOST and COREDX_TARGET are the platform architectures you are compiling on and compiling for (these can be the same). These values can be set manually, or determined by running the script: COREDX_TOP/scripts/cdxenv.sh or cdxenv.bat.

Our Makefiles will run the CoreDX DDS data type compiler to generate the type specific code as well as compile the applications. To compile the Hello World sample application using our Linux Makefile you will need a 'make' program and a C++ compiler in your path. Then, simply type 'make' (or 'nmake -f NMakefile') in the appropriate directory.

This will compile two applications: hello_pub and hello_sub.

3.3.3 The ‘Cxx’ Language Application

CoreDX DDS will require the following paths and libraries when compiling your application:

```
Include Path: \  
    -I${COREDX_TOP}/target/include/dds_cxx  
    -I${COREDX_TOP}/target/include  
Library Path: \  
    -L${COREDX_TOP}/target/${COREDX_TARGET}/lib  
Static Libraries: -ldds_cxx -ldds_dyntype -ldds
```

We’ve provided Makefiles for compiling our examples, and you can use these as a reference for compiling your application. Our Makefile requires three environment variables:

4. COREDX_TOP
5. COREDX_HOST
6. COREDX_TARGET

The COREDX_TOP is a path name to the location of your CoreDX DDS distribution(s). COREDX_HOST and COREDX_TARGET are the platform architectures you are compiling on and compiling for (these can be the same). These values can be set manually, or determined by running the script: COREDX_TOP/scripts/cdxenv.sh or cdxenv.bat.

Our Makefiles will run the CoreDX DDS data type compiler to generate the type specific code as well as compile the applications. To compile the Hello World sample application using ourLinux Makefile you will need a ‘make’ program and a C++ compiler in your path. Then, simply type ‘make’ (or ‘nmake -f NMakefile’) in the appropriate directory.

This will compile two applications: hello_pub and hello_sub.

3.3.4 The ‘Java’ Language Application

We’ve provided scripts for compiling our java examples, and you can use these as a reference for compiling your application. Our scripts require three environment variables:

1. COREDX_TOP

2. COREDX_HOST
3. COREDX_TARGET

The COREDX_TOP is a path name to the location of your CoreDX DDS distribution(s). COREDX_HOST and COREDX_TARGET are the platform architectures you are compiling on and compiling for (these can be the same). These values can be set manually, or determined by running the script: COREDX_TOP/scripts/cdxenv.sh or cdxenv.bat.

Our scripts will run the CoreDX DDS data type compiler to generate the type specific code as well as compile the applications. To compile the Hello World sample application using our Linux scripts you will need 'javac' in your path. Then, simply type 'compile.sh' (or 'compile.bat' for Windows) in the appropriate directory.

This will create a jar file with the two hello applications: HelloPub and HelloSub. We also provide scripts to run these java applications: run_pub.sh and run_sub.sh (.bat scripts are provided for Windows).

3.3.5 The 'C#' Language Application

We've provided Makefiles and Visual Studio project files for compiling our examples, and you can use these as a reference for compiling your application. Our Makefiles require three environment variables:

1. COREDX_TOP
2. COREDX_HOST
3. COREDX_TARGET

The COREDX_TOP is a path name to the location of your CoreDX DDS distribution(s). COREDX_HOST and COREDX_TARGET are the platform architectures you are compiling on and compiling for (these can be the same). These values can be set manually, or determined by running the script: COREDX_TOP/scripts/cdxenv.sh or cdxenv.bat.

Our Makefiles will run the CoreDX DDS data type compiler to generate the type specific code as well as compile the applications. To compile the Hello World sample application using our Makefile on Linux you will need the MONO CSharp compiler (gmcs) in your path. Then, simply type 'make' in the appropriate directory. This will compile two applications: hello_pub.exe and hello_sub.exe.

3.3.6 The ‘Python’ Language Application

Python is an interpreted language – no compiling needed.

3.4 Running Your Application with CoreDX DDS

You will need at least one environment variable to run your applications with CoreDX DDS:

`TWINOAKS_LICENSE_FILE:` (refer to *Using a License*, above)

Run your application(s). The sample Hello World has two applications: `hello_pub` (or `hello_pub.exe` or `HelloPub.class`) and `hello_sub` (or `hello_sub.exe` or `HelloSub.class`). The subscriber application (`hello_sub`) will print out the messages it receives from the publishing application (`hello_pub`). The publisher will keep sending out messages until killed. The subscriber will keep listening for messages until killed.

You can run multiple Publishers and multiple Subscribers to immediately see the dynamic nature of the DDS network infrastructure.

Chapter 4 Example CoreDX DDS Applications

CoreDX DDS is bundled with a number of example applications. These applications including working source code and Makefiles to demonstrate how to write, compile, and run CoreDX DDS applications. Most of the example applications are included in the `${COREDX_TOP}/examples` directory (exceptions are noted below).

This section includes a description of each of the example programs included in the CoreDX DDS release.

4.1 Environment Setup

All the Makefiles in the CoreDX DDS example application require a few environment variables.

1. `COREDX_TOP`
2. `COREDX_HOST`
3. `COREDX_TARGET`

The `COREDX_TOP` is a path name to the location of your CoreDX DDS distribution(s). `COREDX_HOST` and `COREDX_TARGET` are the platform architectures you are compiling on and compiling for (these can be the same). These values can be set manually, or determined by running the script: `COREDX_TOP/scripts/cdxenv.sh` or `cdxenv.bat`.

If you have one CoreDX DDS platform architecture installed, the **`cdxenv`** script will print out the appropriate commands to set these environment variables. If you have multiple CoreDX DDS platform architectures installed, the **`cdxenv`** script will list all your platform architectures and prompt you for the correct HOST and PLATFORM architectures before printing the commands to set these environment variables.

Table 4-1: Example - Running `cdxenv.sh` provides an example of running `cdxenv.sh` on a Linux machine where multiple CoreDX DDS platform architectures are installed.

Table 4-1: Example - Running cdxenv.sh

```
cdxenv Example

/home/bob/coredx-5.0.0/scripts> ./cdxenv.sh
1: Linux_2.6_x86_64_gcc5
2: Windows_x86_64_vs2013
3: Windows_x86_64_vs2015
4: Windows_x86_vs2013
5: Windows_x86_vs2015
Please select the HOST platform [1]:
1: Linux_2.6_x86_64_gcc5
2: Windows_x86_64_vs2013
3: Windows_x86_64_vs2015
4: Windows_x86_vs2013
5: Windows_x86_vs2015
Please select the TARGET platform [1]:
export COREDX_TOP=/home/bob/dds_test/coredx-5.0.0;
export COREDX_HOST=Linux_2.6_x86_64_gcc5;
export COREDX_TARGET=Linux_2.6_x86_64_gcc5;
export LD_LIBRARY_PATH=/home/bob/dds_test/coredx-5.0.0/
    target/Linux_2.6_x86_64_gcc5/lib:
    /home/bob/dds_test/coredx-5.0.0/
    target/Linux_2.6_x86_64_gcc5/lib;
export DYLD_LIBRARY_PATH=/home/bob/dds_test/coredx-5.0.0/
    target/Linux_2.6_x86_64_gcc5/lib:
    /home/bob/dds_test/coredx-5.0.0/
    target/Linux_2.6_x86_64_gcc5/lib;
export COREDX_VERSION=5.0.0;
```

4.2 Example 1: The Basic “Hello World” Applications

CoreDX DDS provides several example “Hello World” applications: one for each supported programming language. These are simple applications that show the basic usage of the CoreDX DDS entities for sending and receiving data.

Each of these “Hello World” examples contains two applications: one that will publish a “Hello World” message and one that will subscribe to and receive the “Hello World” message.

The hello world examples are located in the examples directory:

```
hello_c/
hello_cpp/
```

```
hello_cxx/  
hello_csharp/  
hello_java/  
hello_python/
```

Makefiles are provided for all platform architectures supported by CoreDX DDS, so these applications can be run on a variety of operating systems and languages.

4.3 Example 2: Performance Tests

CoreDX DDS also provides sample performance benchmarking source code for latency and bandwidth benchmarking. These applications provide an example of a more sophisticated CoreDX DDS application. In addition, they allow you to determine the performance of CoreDX DDS with your computers and networking hardware.

The performance examples are located in the examples directory:

```
latency_test\  
bwtest\
```

Makefiles are provided for a variety of platform architectures.

4.4 Example 3: Filtering

CoreDX DDS provides an example of using Content Filtered Topics in the “dds_filter” and “dds_filter_cpp” example applications. Makefiles are provided for a variety of platform architectures.

4.5 Example 4: RPC

CoreDX DDS provides an example of using the CoreDX DDS RPC API: the “rpc_fcall” example application. For more information about using the RPC and request-response API, refer to the RPC Programmer's Guide.

4.6 Example 5: QoS Provider

CoreDX DDS provides an example of configuring QoS policies in an external XML file, and applying them dynamically to entities created at run-time. This example application can be found in the “qos_provider” directory.

4.7 Example 6: Security

CoreDX DDS provides an example of using CoreDX DDS Secure. Note, CoreDX DDS Security is provided in it’s own distribution, not part of the standard CoreDX DDS distribution.

4.8 Example 7: Shapes Demonstration

The java source code for the CoreDX DDS Shapes demonstration is freely available from the Twin Oaks Computing website (it is not packaged as part of the CoreDX DDS release). The Shapes demonstration provides examples of medium complexity C and Java applications using CoreDX DDS for communications. A specialized version of the CoreDX DDS Shapes demonstration is available for Android platforms.

For additional information and instructions for downloading and using, visit the Twin Oaks Computing website:

http://www.twinoakscomputing.com/coredx/shapes_demo

Chapter 5 Advanced Compile Options

CoreDX DDS includes several libraries that can be used to develop CoreDX DDS applications. Some libraries are required for any CoreDX DDS application. Some libraries include advanced DDS features that should only be used if necessary. Some libraries provide additional debugging information. For all of these libraries, we include a static and dynamic library option (for operating systems that support this distinction).

This section describes the different options available to compile a CoreDX DDS application for different operating system environments.

5.1 Linux and other UNIX-variant Operating Systems

Table 5-1 lists all the libraries provided with the Linux platform release of CoreDX DDS. Other UNIX-variant operating systems may include a sub-set of these libraries (depending on the CoreDX DDS features supported for each particular operating system).

Table 5-1: CoreDX DDS Libraries (UNIX Operating Systems)

Language	library file name	Description
C libraries		
	libdds	core library (static and dynamic libraries)
	libdds_cf	core with content filter library (static and dynamic libraries)
	libdds_noto	Minimal core library: no type object (static and dynamic libraries)
	libdds_dyntype	Dynamic Types extension (static and dynamic libraries)
	libdds_qos_provider	QoS Provider extension (static and dynamic libraries)

Language	library file name	Description
	libdds_libxml	XML API extension (static and dynamic libraries)
	libcdx_tport_XXX	Various transport extensions
	libxxx_log	Every C library has an equivalent library version with logging enabled.
C++ libraries		
	libdds_cpp	C++ language binding extension (static and dynamic libraries)
	libdds_cpp_cf	C++ language binding for content filters extension, this extension replaces the libdds_cpp extension (static and dynamic libraries)
	Libdds_cpp_noto	Minimal C++ language binding: no type object, this extension replaces the libdds_cpp extension (static and dynamic libraries)
	libdds_cpp_dyntype	C++ Dynamic Types extension (static and dynamic libraries)
	libdds_cpp_qos_provider	QoS Provider extension for the cpp language binding (static and dynamic libraries)
	libdds_rpc_cpp	RPC and request-reply extension (static and dynamic libraries)
	libxxx_cpp_log	Every C++ library has an equivalent library version with logging enabled.
Cxx libraries		
	libdds_cxx, libdds_cxx_fpic	Cxx language binding: contains all features in one library
C# libraries		

Language	library file name	Description
	coredx_csharp.dll	C# language binding: contains all features in one library
	libdds_csharp	C# language binding: native library
	libdds_csharp_log	C# language binding – native library with logging
Java libraries		
	libdds_java	Java language binding – native library
	libdds_java_log	Java language binding – native library with logging
Python libraries		
	python2.7/*	Python language binding – to support Python v2.7
	python3.6/*	Python language binding – to support Python v3.6

CoreDX DDS provides both static (".a") and dynamic (".so") libraries, so the application developer can choose the appropriate type of library to use. All our example applications contain Makefiles that illustrate the use of the different libraries and provide a good reference for CoreDX DDS library usage.

5.1.1 Linking with Static Libraries

A simple CoreDX DDS application written in 'C' requires the core library: libdds.a. The Makefile in the "hello_c" sample application provides an example of using the core library.

A simple CoreDX DDS application written in 'C++' requires the C++ language binding library: libdds_cpp.a **in addition to** the core library: libdds.a. The

Makefile in the “hello_cpp” sample application provides an example of using the C++ library.

There are 4 versions of the base (required) CoreDX DDS library:

```
⇒ libdds
⇒ libdds_cf
⇒ libdds_noto
⇒ libdds_log
```

An application will link only 1 of the above libraries. Other libraries are *extensions* of one the above base libraries, and will be linked *in addition* to the base library. For example: a CoreDX DDS application using content filters must use the content filter library *in place of* the core library. A ‘C’ application will use libdds_cf.a. A ‘C++’ application will use libdds_cf.a *and* libdds_cpp_cf.a (to provide the C++ language binding extension). The Makefile in the “dds_filter” sample application provides an example of using the content filter library.

The C#, Java, and Python languages bindings have one library (or one set of libraries) containing all features available for that language binding.

To enable CoreDX DDS logging (refer to *CoreDX DDS Logging* for information on CoreDX DDS logging), a CoreDX DDS application must use the logging library *in place of* the core library, and to be complete, may also link in the log versions of all the extension libraries. A ‘C’ application will use libdds_log.a. A ‘C++’ application will use libdds_log.a *and* libdds_cpp_log.a.

There is no special configuration required to run a CoreDX DDS application linked with static libraries. Simply set the TWINOAKS_LICENSE_FILE environment variable appropriately as for any CoreDX DDS application.

5.1.2 Linking with Dynamic Libraries

In addition to linking in the correct CoreDX DDS libraries, the LD_LIBRARY_PATH environment variable must be set in order to run dynamically linked CoreDX DDS applications.

To link with dynamic libraries, refer to section 5.1.1: *Linking with Static Libraries* above and replace the “.a” libraries with the “.so” version of the libraries.

To run a CoreDX DDS application compiled with dynamic libraries, the LD_LIBRARY_PATH environment variable must be set, in addition to the TWIN OAKS_LICENSE_FILE environment variable.

5.1.3 Comiling with -fpic (Linux platforms)

For Linux builds, CoreDX DDS also provides versions of the libraries that may be compiled with the -fpic option. These libraries are located in an "fpic" subdirectory: <COREDX_TARGET>/lib/fpic.

5.2 Windows Operating System

Table 5-2 lists all the libraries provided with the Windows platform release of CoreDX DDS.

Table 5-2: CoreDX DDS Libraries (Windows Operating System)

Language	library file name	Description
C libraries		
	dds	core library (static and dynamic libraries)
	dds_cf	core with content filter library (static and dynamic libraries)
	dds_noto	Minimal core library: no type object (static and dynamic libraries)
	dds_dyntype	Dynamic Types extension (static and dynamic library)
	dds_qos_provider	QoS Provider extension (static and dynamic libraries)
	dds_xml	XML API extension (static and dynamic libraries)

Language	library file name	Description
	cdx_tport_XXX	Various transport extension (static and dynamic libraries)
	dds_XXX_log dds_XXX_d dds_XXX_d_log	Every C library has an equivalent library version with logging enabled, with debug enabled, and with both logging and debug enabled.
C++ libraries		
	dds_cpp	C++ language binding extension (static and dynamic libraries)
	dds_cpp_cf	C++ language binding for content filters extension, this extension replaces the libdds_cpp extension (static and dynamic libraries)
	dds_cpp_noto	Minimal C++ language binding: no type object, this extension replaces the libdds_cpp extension (static and dynamic libraries)
	dds_cpp_dyntype	C++ Dynamic Type extension (dynamic library) (debug version)
	dds_cpp_qos_provider	C++ Dynamic Types extension (static library) (debug version)
	dds_rpc_cpp	RPC and request-reply extension (static and dynamic libraries)

Language	library file name	Description
	dds_cpp_xxx_log dds_cpp_xxx_d dds_cpp_xxx_d_log	Every C++ library has an equivalent library version with logging enabled, with debug enabled, and with both logging and debug enabled.
C# libraries		
	coredx_csharp	C# language binding: contains all features in one library
	dds_csharp	C# language binding – native library
	dds_rpc_csharp	RPC and request/reply extension for C# language binding
	dds_csharp_log	C# language binding – native library with logging
Java libraries		
	dds_java	Java language binding – native library
	dds_java_log	Java language binding – native library with logging
Python libraries		
	site-packages/*	Python language binding

Both static (“_static.lib”) and dynamic (“.dll and .lib”) libraries are provided, along with variant libraries with debug and logging enabled. The application developer can choose the appropriate type of library to use. All our example applications contain Makefiles that illustrate the use of the different libraries and provide a good reference for CoreDX DDS library usage.

5.2.1 Linking with Static Libraries

A simple CoreDX DDS application written in ‘C’ requires the core library: `dds_static.lib`. The NMakefile in the “hello_c” sample application provides an example of using the core library.

A simple CoreDX DDS application written in ‘C++’ requires the C++ language binding library: `dds_cpp_static.lib` **in addition to** the core library: `dds_static.lib`. The NMakefile in the “hello_cpp” sample application provides an example of using the C++ library.

There are 4 versions of the base (required) CoreDX DDS library:

- ⇒ `libdds`
- ⇒ `libdds_cf`
- ⇒ `libdds_noto`
- ⇒ `libdds_log`

An application will link only 1 of the above libraries. Other libraries are *extensions* of one the above base libraries, and will be linked *in addition* to the base library. For example: a CoreDX DDS application using content filters must use the special content filter library **in place of** the core library. A ‘C’ application will use `dds_cf_static.lib`. A ‘C++’ application will use `dds_cf_static.lib` **and** `dds_cpp_cf_static.lib`. The NMakefile in the “dds_filter” sample application provides an example of using the content filter library.

A CoreDX DDS application using dynamic types must use the special dynamic type library **in addition to** the core library. A ‘C’ application will use `dds_static.lib` **and** `dds_cf_static.lib`. Dynamic types for C++ are not yet supported. The NMakefile in the “dyntype” sample application provides an example of using the dynamic type library.

The C#, Java, and Python languages bindings have one library (or one set of libraries) containing all features available for that language binding.

To enable CoreDX DDS logging (refer to *CoreDX DDS Logging* for information on CoreDX DDS logging), a CoreDX DDS application must use the logging library **in place of** the core library. A ‘C’ application will use `dds_log_static.lib`. A ‘C++’ application will use `dds_log_static.lib` **and** `dds_cpp_log_static.lib`. There are also content filter and dynamic type libraries with logging enabled.

There is no special configuration required to run a CoreDX DDS application linked with static libraries. Simply set the `TWINOAKS_LICENSE_FILE` environment variable appropriately as for any CoreDX DDS application.

NOTE: To build an application linked to static libraries, **do not include** `/DCOREDX_DLL` in the compile flags.

5.2.2 Linking with Dynamic Libraries

Additional compiler flags are required to build an application linked with dynamic libraries. Include the following in the compile flags:

```
/MD
/DCOREDX_DLL
```

Note: the `"/MD"` specifies dynamically linked applications, while `"/MT"` is for statically linked applications. To link a dynamic application with debug support, replace `"/MD"` with `"/MDd"` and link to the libraries with `"_d"` in their name.

To link with dynamic libraries, refer to section 5.2.1: *Linking with Static Libraries* above but replace the `"_static.lib"` libraries with the `".lib"` version of the libraries.

There are a few additional rules for linking the correct CoreDX DDS Windows dynamic libraries, due to the nature of Windows dynamic libraries. These rules are described below.

Table 5-3: Windows Dynamic Library Dependencies

If your application is using:	It must also link in the following:
dds_cf.dll, .lib	N/A (standalone library)
dds_cf_log.dll, .lib	N/A (standalone library)
dds_cpp_cf.dll, lib	dds_cf.lib
dds_cpp_cf_log.dll, .lib	dds_cf_log.lib
dds_cpp.dll, lib	dds.lib

If your application is using:	It must also link in the following:
dds_cpp_log.dll, lib	dds_log.lib
dds.dll, lib	N/A (standalone library)
dds_dyntype.dll, lib	dds_cf.lib
dds_dyntype._log.dll, lib	dds_cf_log.lib
dds_java.dll, lib	N/A (applications do not link this library)
dds_log.dll, lib	N/A (standalone library)

To run an application linked to CoreDX DDS dynamic libraries, the CoreDX DDS libraries must be in your PATH environment variable. For example:

```
set
PATH=%PATH%;%COREDX_TOP%\target\%COREDX_TARGET%\lib
```

In addition, the TWINOAKS_LICENSE_FILE environment variable must be set correctly.

5.2.3 Dynamic Type Support Library

It may be desirable to create a dynamic library containing the CoreDX DDS generated type support code. This requires some special configuration for Windows.

First, the **generated type support code** must be compiled with these additional flags:

```
/MD (for dynamic applications. Replace with "/MDd" for debug)
/DCOREDX_DLL
/DCOREDX_DLL_TS
/DCOREDX_DLL_TS_EXPORTS
/LD
```

Note the “/MD” is for dynamically linked applications, while “/MT” is for statically linked applications. To link a dynamic application with debug, replace “/MD” with “/MDd”.

The ***application*** linking in the dynamic type support library must be compiled with these additional flags:

```
/MD  
/DCOREDX_DLL  
/DCOREDX_DLL_TS
```

Then include the dynamic type support library on the link line with the appropriate CoreDX DDS libraries.

To run this resulting application, ensure the generated dynamic type support library is available in your path (in addition to the CoreDX DDS libraries). In addition, the TWINOAKS_LICENSE_FILE must be set appropriately.

Part 3: CoreDX DDS Programming Concepts

This section provides a more detailed examination of the considerations for designing and developing applications that make effective use of the CoreDX DDS middleware. This includes data architecture, Quality of Service configuration, data access patterns, and status event handling.

Chapter 6 DDS Entities

The DDS Standard defines an architecture that represents an object-oriented model of entities that compose the DDS API. These entities serve as the interface between the middleware and the application software. In order to develop a DDS enabled application, a developer must create, interact with, and destroy these DDS entities.

This chapter provides an overview of the DDS Entities and related concepts. Subsequent chapters will provide more details and examples to fully illustrate the API.

6.1 DDS Entity Hierarchy

The primary entities that make up the DDS API are structured in a hierarchy. Each entity in the hierarchy exposes a related set of operations from the API. The programmer interacts with the CoreDX DDS middleware through these DDS entities. For example, the common operations on these entities include creation, destruction, get and set QoS, get and set listeners, get status.

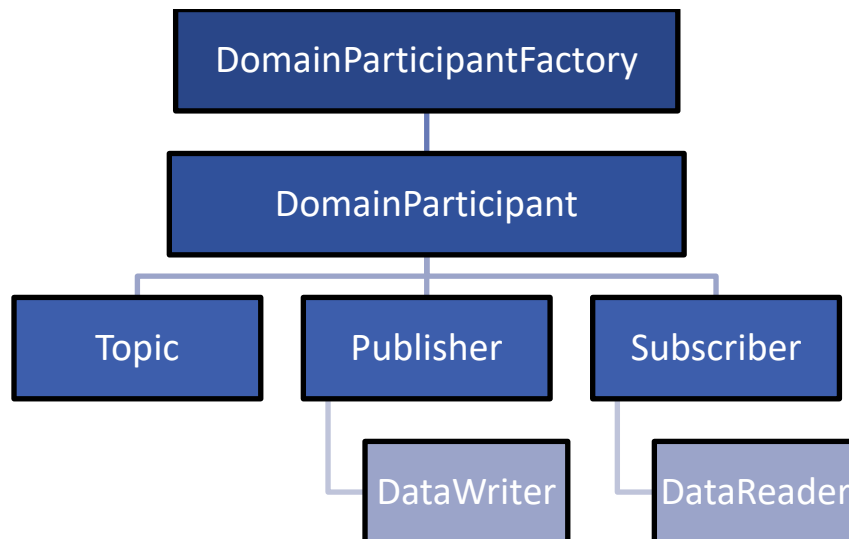


Figure 6-1: DDS Entity Hierarchy

The **DomainParticipantFactory** is the initial object available to applications. It is a singleton, meaning that only one instance of this object exists within

an application. It is used as a factory to create and delete DomainParticipants.

The DomainParticipant object is the factory for Publishers, Subscribers, and Topics. The DomainParticipant is a container for all of the entities that it creates. The DomainParticipant exists within a ‘DOMAIN’. All entities created from a DomainParticipant belong to the same DOMAIN as their parent participant. Entities within a DOMAIN may communicate. Entities in different DOMAINS will not communicate. The DOMAIN specifies a fundamental separation or scope of data.

The Topic entity provides a logical set of homogenous data. The Topic exists within a domain, and has a name and a data type. DataReaders and DataWriters are logically connected by a common topic.

A Publisher entity is a factory for DataWriters. The Publisher is a container for all the DataWriters that it creates. A DataWriter exists within a single Publisher, and is associated with a single Topic. A DataWriter is capable of publishing a single data type that matches its Topic.

A Subscriber entity is a factory for DataReaders. The Subscriber is a container for all the DataReaders that it creates. A DataReader exists within a single Subscriber, and is associated with a single Topic. A DataReader is capable of reading a single data type that matches its Topic.

6.2 DDS Entity Common Operations

Each Entity shares a set of common operations. These operations allow the application to control basic aspects of the entity. For example, an Entity must be enabled before it will participate in communications. If the entity is not enabled automatically by its parent factory, then the enable() method must be called manually.

```
enable()  
get_qos()  
set_qos()  
get_listener()  
set_listener()  
get_statuscondition()  
get_status_changes()
```

6.3 DDS Entity Quality of Service

The behavior of DDS communication is highly configurable. This configuration is performed using Quality of Service policies. Each of these primary DDS Entities accepts Quality of Service parameters to control their behavior. The parent factory maintains a default configuration of QoS for its children. If no QoS is provided at creation time, then these defaults are used. An entity's QoS can be accessed by calling `get_qos()` on the entity. After an entity is created, its QoS can be changed by directly calling `set_qos()` on the entity. If the entity is not yet enabled, then any QoS policy can be changed. However, after the entity is enabled, only a subset of QoS policies can be changed on the fly. See the Chapter on QoS Policies for details.

6.4 DDS Status, Listeners, Conditions and WaitSets

Each entity maintains a set of status information. This information represents the occurrence of significant events within the middleware. For example, the "Data Available" or "Publication Matched" statuses. The CoreDX DDS middleware supports multiple notification methods to communicate status information to the application. Listener callbacks can be installed, supporting asynchronous notification. Alternatively, the application can initialize a WaitSet and block waiting for a specific set of status conditions. This represents synchronous notification. Finally, the application can choose to poll the middleware using the `get_status_changes()` method.

Chapter 7 Developing a Publishing Application

This chapter describes the development process for DDS publishing applications.

7.1 Summary of Developing a Publishing Application

The steps for creating a publishing application are as follows:

1. Create or obtain the data types for the application data.
2. Use the CoreDX DDS data types compiler to compile the data types. The type-specific support and DataWriter are created as a result of compiling the data types.
3. Write the publishing application.
4. Compile the publishing application, linking with the generated code from step 2 and the applicable CoreDX DDS libraries.

7.2 The Data

DDS enabled applications are inherently data-centric. In order for these data-centric applications to perform efficiently, it is necessary to have a well-considered data model, which is usually implemented in either IDL or XML. [Data types may also be created dynamically at run-time, or discovered via DDS Discovery.]

For more information on the data types supported by the CoreDX DDS data type compiler, see the *Application Data Types* section of this guide.

7.3 The Publishing Application

Note: DDS API names and approaches may be different between the different language bindings.

There are two basic approaches to the DDS API for creating DDS Entities. The original approach is used for the CoreDX DDS C, C++, C#, and Java language bindings. The more recent approach is used for the Cxx and Python language bindings. One approach is not 'better' than the other, they are simply two different ways to accomplish the same task.

This section illustrates the original approach used by C, C++, C#, and Java.

The Reference Guides for Cxx and Python include ‘cookbook’ examples illustrating the alternative approach to DDS Entity creation.

Some languages use a DDS namespace, while the C language binding augments names to include the namespace as a prefix. For example, in C++ we might reference `DDS::DomainParticipantFactory::create_participant()`, whereas in C, this would look like `DDS_DomainParticipantFactory_create_participant()`. For purposes of illustration, the examples in this section are written in C++, and assume that the DDS namespace is in use.

A publishing application, if using compiled data types, must include the generated Type, TypeSupport, and DataWriter header files.

7.3.1 Initialize the DomainParticipantFactory

This is the first step for any application communicating using DDS, and initializes the CoreDX DDS middleware for use.

In the ‘C’ language binding, there is initialization the `DomainParticipantFactory`.

In the ‘C++’, ‘C#’, and Java language bindings, the application must get an instance of the `DomainParticipantFactory` in order to use it (for example: to set a license or create a `DomainParticipant`):

```
DomainParticipantFactory * dpf =  
DomainParticipantFactory::get_instance();
```

7.3.2 Create a DomainParticipant

The `DomainParticipant` represents the participation of an application in a virtual DDS network, linking all applications that share the same Domain ID. Several independent DDS “networks” (in DDS terms, these are different Domains) can coexist in the same physical network without interfering (or even being aware) of each other.

A DDS application may contain more than 1 `DomainParticipant`, in order to communicate in multiple DDS Domains, or otherwise organize the application’s communication events.

In addition, the DomainParticipant acts as a factory for creating Topics, Publishers (and Subscribers).

When creating a DomainParticipant, you will be able to specify:

Domain ID	The unique identifier for the domain this application will be publishing in
QoS for the DomainParticipant	Describes the QoS for the Domain Participant
Listener	Allows the application to attach listeners to the domain participant.
Listener Status Mask	Sets which listeners are active
Security Properties	If using CoreDX DDS Secure, the security configuration for this DomainParticipant

[There are additional items to initialize a secure DomainParticipant, this is described in the CoreDX DDS Security Programmer's Guide.] To create a DomainParticipant in the 123 domain, using the default DomainParticipant QoS, and no listeners, use the following:

```
DomainParticipant * dp = dpf->create_participant( 123,
    PARTICIPANT_QOS_DEFAULT, NULL, 0);
```

By default, the create_participant() call will initialize and enable the DomainParticipant for communications.

7.3.3 Create a Publisher

The publisher is responsible for disseminating (publishing) data. It also acts as a factory for creating DataWriters.

When creating a Publisher, you will be able to specify:

QoS for the Publisher	Describes the QoS for the Publisher
Listeners	Allows the application to attach listeners to the publisher

Listener Status Mask	Sets which listeners are active
-----------------------------	---------------------------------

To create a Publisher using the default Publisher QoS and no listeners, use the following:

```
Publisher * pub = dp->create_publisher( PUBLISHER_QOS_DEFAULT,  
    NULL, 0 );
```

By default, the create_publisher() call will initialize and enable the Publisher for communications.

7.3.4 Register a Data Type

In order to publish (or subscribe to) data, the data *type* must be registered in the CoreDX DDS middleware.

The most common mechanism to register a data type is to use the TypeSupport generated from the IDL or XML. (Refer to the DDS Type System Programmer’s Guide for additional information on registering data types.) The example code uses a type name: StringMsg. To register this type using the default type name (StringMsg) use the following:

```
StringMsgTypeSupport ts;  
ts.register_type( dp, NULL );
```

The default type name is the ‘base’ type name, without any namespace prefixes. You can supply an alternate type name by replacing the 2nd argument with a string name.

7.3.5 Create a Topic

The Topic essentially links the publishers of data with the subscribers of data. A Topic is identified by a unique topic name, and one data type.

When creating a Topic, you are able to specify:

Topic Name	Must be unique in the Domain
Type Name	Must already be registered in the Domain

QoS for the Topic	Describes the QoS for the Topic
Listeners	Allows the application to attach listeners to the Topic
Listener Status Mask	Sets which listeners are active

To create a Topic named “HelloTopic” with the “StringMsg” type, default QoS, and no listeners use the following:

```
Topic topic = dp->create_topic( "HelloTopic", "StringMsg",
    TOPIC_QOS_DEFAULT, NULL, 0 );
```

7.3.6 Create a DataWriter

The DataWriter is associated with exactly 1 Topic, and can write data of the specific data type assigned to that Topic. The application typically uses a type-specific DataWriter to publish data. [The alternative to a type-specific DataWriter is a DynamicTypeDataWriter. More information may be found in the Dynamic Types section of this Programmer's Guide.]

When creating a DataWriter, you will be able to specify:

Topic	The Topic to write “on”
QoS for the DataWriter	Describes the QoS for the DataWriter
Listeners	Allows the application to attach listeners to the DataWriter
Listener Status Mask	Sets which listeners are active

To create a DataWriter with the Topic created above, default QoS, and no listeners, use the following:

```
DataWriter * dw = pub->create_datawriter( topic,
    DATAWRITER_QOS_DEFAULT, NULL, 0 );
```

For C++ only: This command creates a “generic” DataWriter (the publisher does not know what type of data will be written). This generic DataWriter

will work, however, there is no type checking on the data passed to this generic `DataWriter` on a `write()`. In order to have that type checking, use the `narrow()` method to obtain a type-specific `DataWriter`:

```
StringMsgDataWriter * sdw = StringMsgDataWriter::narrow(  
dw );
```

7.3.7 Write Data

All the necessary pieces are now in place to start publishing (writing) data. There are two methods that can be used for writing:

```
ReturnCode_t retval = sdw->write( data, HANDLE_NIL );
```

And:

```
ReturnCode_t retval = sdw->write_w_timestamp( data,  
HANDLE_NIL, time );
```

The `write()` method writes the data and uses the current time as the source time stamp. The `write_w_timestamp()` method allows the application to specify the source time stamp. The source time stamp is sent along with the data, and is located in the `SampleInfo` on the subscribing end.

Both the `write()` and `write_w_timestamp()` methods take an **instance handle** argument (the examples above used the empty handle: `HANDLE_NIL`). Each piece of data written is associated in the CoreDX DDS middleware by an instance handle. For more information on instance handles, see the [Instances and Samples](#) chapter. When the data contains a key, calling `write()` with `HANDLE_NIL` results in the infrastructure looking up the key data to find the associated instance handle. If you are doing multiple writes with data contain the same key (or set of keys), you can optimize the code by calling `register()` before `write()`. For example:

```
InstanceHandle_t handle = sdw->register_instance(data);  
ReturnCode_t      retval = sdw->write( data, handle );  
ReturnCode_t      retval = sdw->write( data, handle );
```

The `write()` API is asynchronous: the `write()` call returns once the CoreDX DDS middleware has scheduled the Sample to be written, but the sample is not necessarily written ‘on the wire’ at that time. The return value from `write` indicates CoreDX DDS has successfully accepted this sample, or not. Possible reasons for `write()` returning a non-success return code include:

1. There is no room in the DataWriter cache to hold the written sample, or the written sample's instance.
2. The provided instance_handle (if not HANDLE_NIL) does not refer to a valid instance

7.4 Available QoS Settings

Every created DDS entity has an associated **Quality of Service** (QoS) that can be specified when creating the entity, or later by calling the `set_qos()` method on that entity. The QoS for each entity is a comprehensive set of configuration policies that affect the behavior of that entity. The middleware defines a default QoS for each entity, which was used in the examples above.

Below is a table listing the available QoS for the publishing entities (DomainParticipant, Publisher, Topic, and DataWriter). This table lists only brief descriptions. A more complete list of QoS Policies and their descriptions can be found in the *Quality of Service Features* chapter.

Table 7-1: QoS Policies for Publishing Entities

QoS Policy	Description	Available To
USER_DATA	Data not used by the CoreDX DDS middleware, the application can use this data for its own purposes.	DomainParticipant, DataWriter
TOPIC_DATA	Data not used by the CoreDX DDS middleware, the application can use this data for its own purposes.	Topic
GROUP_DATA	Data not used by the CoreDX DDS middleware, the application can use this data for its own purposes.	Publisher
DURABILITY	Specifies if published data should be saved for later-joining DataReaders to receive.	Topic, DataWriter
DURABILITY_SERVICE	Not Yet Implemented Specifies configuration for the disabilities: TRANSIENT and	Topic, DataWriter

	PERSISTENT.	
PRESENTATION	Affects how data is presented to the subscribing application. For example, published data can be grouped together such that the DataReaders receive all coherent data together.	Publisher
DEADLINE	Establishes an agreement that the DataWriter will update each instance of the data at least once every specified time period. If the writing application fails to meet this contract, the deadline_missed status on the DataWriter is triggered.	Topic, DataWriter
OWNERSHIP	Specifies if multiple DataWriters are allowed to write (or update) the same instance of the data, and how these modifications should be handled.	Topic, DataWriter
OWNERSHIP_STRENGTH	Specifies the strength used to arbitrate among multiple DataWriters writing (or updating) the same instance of the data.	DataWriter
LIVELINESS	Indicates a commitment by the DataWriter to signal it’s liveliness to DataReaders in the specified interval. This may mean the DataWriter updates its samples, or simply asserts it is still alive. If the DataWriter (or application) fails to meet this liveliness contract, the LIVELINESS_LOST status is triggered on the DataWriter.	Topic, DataWriter
PARTITION	A logical partition among Topics visible to Publishers and Subscribers. A publisher will only communicate with a subscriber if their Partitions match (wildcards allowed).	Publisher

RELIABILITY	Indicates the level of reliability offered by the DataWriter.	Topic, DataWriter
DESTINATION_ORDER	Specifies the order in which changes to an instance will be published.	Topic, DataWriter
HISTORY	Specifies if the publishing middleware should keep any (or all) updates to an instance on behalf of existing DataReaders.	Topic, DataWriter
RESOURCE_LIMITS	Specifies the resources the middleware can consume in order to meet the requested QoS.	Topic, DataWriter
RTPS_Writer	Configures CoreDX DDS extensions to DataWriter behavior	DataWriter
DISCOVERY	Configure CoreDX DDS extensions for transport and discovery options	DomainParticipant
PEER_PARTICIPANT	Configure static peers to communicate with	DomainParticipant
THREAD_MODEL	Configuring threading options (single threaded / multithreaded)	DomainParticipant
ENTITY_FACTORY	Specifies if a factory should automatically enable created entities. If the factory does not automatically enable those entities, the application must specifically enable them before they can be used for publishing or writing data.	DomainParticipantFactory, DomainParticipant, Publisher

7.5 Available Listeners

Listeners are one mechanism allowing the application to be made aware of events and changes in the CoreDX DDS middleware communication status.

A listener has a number of methods defined, one for each applicable communication status. The application can define one or more listener methods and attach them to an appropriate DDS entity when the entity is created, or later by using the `set_listener()` method on the entity.

The table below lists the listeners that can be attached to publishing Entities. A more complete list and description of listeners can be found in the *Communication Status* chapter.

Table 7-2: Listeners for Publishing Entities

Listener	Listener Methods	Description
DataWriterListener	<code>on_offered_deadline_missed()</code>	A deadline offered through the DEADLINE QoS setting was missed.
	<code>on_offered_incompatible_qos()</code>	A DataReader was discovered for the same Topic as this DataWriter, but the QoS requested by that DataReader was incompatible with this DataWriter’s offered QoS.
	<code>on_liveliness_lost()</code>	The liveliness specified in the LIVELINESS QoS was not respected, and DataReaders will consider this DataWriter no longer active.
	<code>on_publication_matched()</code>	A DataReader has been found that matches the Topic and QoS of this DataWriter (or a DataReader that was previously matched is no longer matched).
PublisherListener	(none)	(PublisherListener inherits methods from the DataWriterListener)
TopicListener	<code>on_inconsistent_topic()</code>	Another, different, Topic exists with the same name as this

Topic.		
DomainParticipantListener	(none)	(DomainParticipantListener inherits methods from the DataWriterListener, PublisherListener, and TopicListener)

Chapter 8 Developing a Subscribing Application

This chapter describes the development process for DDS subscribing applications.

8.1 Summary of Developing a Subscribing Application

The steps for creating a subscribing application are as follows:

1. Create or obtain the data definitions for the DDS interfaces
2. Use the CoreDX DDS data type compiler to compile the data types. The type-specific support and DataReader are created as a result of compiling the data types.
3. Write the subscribing application
4. Compile the subscribing application, linking with the generated code from step 2 and the applicable CoreDX DDS libraries.

8.2 The Data

DDS enabled applications are inherently data-centric. In order for these data-centric applications to perform efficiently, it is necessary to have a well-considered data model, which is usually implemented in either IDL or XML. [Data types may also be created dynamically at run-time, or discovered via DDS Discovery.]

For more information on the data types supported by the CoreDX DDS data type compiler, see *Application Data Types*.

8.3 The Subscribing Application

Note: API names and approaches may be different between the different language bindings.

There are two basic approaches to the DDS API for creating DDS Entities. The original approach is used for the CoreDX DDS C, C++, C#, and Java language bindings. The more recent approach is used for the Cxx and Python language bindings. One approach is not 'better' than the other, they are simply two different ways to accomplish the same task.

This section illustrates the original approach used by C, C++, C#, and Java.

The Reference Guides for Cxx and Python include ‘cookbook’ examples illustrating the alternative approach to DDS Entity creation.

Some language bindings use a DDS namespace, and the C language binding augments names to include the namespace as a prefix. For example, in C++ we might reference `DDS::DomainParticipantFactory::create_participant()`, whereas in C, this is the syntax:

`DDS_DomainParticipantFactory_create_participant()`. For purposes of illustration, the examples in this section are written in C++, and assume that the DDS namespace is in use.

A subscribing application must include the generated Type, TypeSupport, and DataReader header files.

8.3.1 Initialize the DomainParticipantFactory

This is the first step for any application communicating using DDS, and initializes the CoreDX DDS middleware for use.

In the ‘C’ language binding, there is initialization the `DomainParticipantFactory`.

In the ‘C++’, ‘C#’, and Java language bindings, the application must get an instance of the `DomainParticipantFactory` in order to use it (for example: to set a license or create a `DomainParticipant`):

```
DomainParticipantFactory * dpf =  
DomainParticipantFactory::get_instance();
```

8.3.2 Create a DomainParticipant

The `DomainParticipant` represents the participation of an application in a virtual network, linking all applications that share the same domain ID. Several independent DDS “networks” (in DDS terms, these are different Domains) can coexist in the same physical network without interfering (or even being aware) of each other.

In addition, the `DomainParticipant` acts as a factory for creating Topics and Subscribers (and Publishers), or otherwise organize the application’s communication events.

When creating a `DomainParticipant`, you will be able to specify:

Domain ID	The unique identifier for the domain this application will be publishing in
QoS for the DomainParticipant	Describes the QoS for the Domain Participant
Listener	Allows the application to attach listeners to the domain participant.
Listener status mask	Sets which listeners are active

To create a DomainParticipant in the 123 domain, using the default DomainParticipant QoS, and no listeners use the following:

```
DomainParticipant * dp = dpf->create_participant( 123,
    PARTICIPANT_QOS_DEFAULT, NULL, 0 );
```

By default, the create_publisher() call will initialize and enable the Publisher for communications.

8.3.3 Create a Subscriber

The subscriber is responsible for receiving data. It also acts as a factory for creating DataReaders. When creating a Subscriber, you are able to specify:

QoS for the Subscriber	Describes the QoS for the Subscriber
Listeners	Allows the application to attach listeners to the subscriber
Listener Status Mask	Sets which listeners are active

To create a Subscriber using the default Publisher QoS and no listeners, use the following:

```
Subscriber * sub = dp->create_subscriber(
    SUBSCRIBER_QOS_DEFAULT, NULL, 0 );
```

8.3.4 Register a Data Type

In order to subscribe to data, the data *type* must be registered in the CoreDX DDS middleware.

To register a data type, use the TypeSupport generated from the CoreDX DDS data type compiler. The example code uses a type name: StringMsg. To register this type using the default type name (StringMsg) use the following:

```
StringMsgTypeSupport ts;  
ts.register_type( dp, NULL );
```

You can supply an alternate type name by replacing the 2nd argument with a string name.

8.3.5 Create a Topic

The Topic essentially links the publishers of data with the subscribers of data. A Topic is identified by a unique topic name, and a type.

When creating a Topic, you will be able to specify:

Topic Name	Must be unique in the Domain
Type Name	Must already be registered in the Domain
QoS for the Topic	Describes the QoS for the Topic
Listeners	Allows the application to attach listeners to the Topic
Listener Status Mask	Sets which listeners are active

To create a Topic named “HelloTopic” with the “StringMsg” type, default QoS, and no listeners use the following:

```
Topic topic = dp->create_topic( "HelloTopic",  
"StringMsg", TOPIC_QOS_DEFAULT, NULL, 0 );
```

8.3.6 Create a DataReader

The DataReader can read data of a specific type. The application uses a type-specific DataReader to read data. When creating a DataReader, you will be able to specify:

Topic	The Topic to read "on"
QoS for the DataReader	Describes the QoS for the DataReader
Listeners	Allows the application to attach listeners to the DataReader
Listener Status Mask	Sets which listeners are active

To create a DataReader with the Topic created above, default QoS, and no listeners, use the following:

```
DataReader * dr = sub->create_datareader( topic,
                                         DATAREADER_QOS_DEFAULT, NULL, 0 );
```

For C++ applications: An application should cast the DataReader to a type specific DataReader to use it in a type-safe manner. The narrow() operation can be used to obtain the type specific DataReader. For example:

```
StringMsgDataReader * sdr =
    StringMsgDataReader::narrow( dr );
```

8.3.7 Read (or Take) Data

Ultimately, the application will call one of the read() or take() methods on the DataReader to access data. These operations are non-blocking, and deliver the data that has been received and is currently available to the application. The read() and take() operations return an ordered collection of data samples, and their associated sample information, that match the QoS policies set on the Subscriber and DataReader and the parameters passed to read() and take().

The read() and take() operations have a similar API signature, and a set of variants that provide the application with additional control over the returned data. The basic read() operation provides the application with access to data managed by the DataReader. After the read() operation, the

data is still managed by the DataReader, and is available for access by subsequent read() operations. The take() operation, also provides access to data managed by the DataReader. It differs from read() because data samples accessed by the take() operation are removed from the DataReader, and are not available to subsequent read() or take() operations. As an analogy, read() peeks at the data available in the DataReader while take() actually removes the data from the DataReader.

The basic read() and take() operations both allow the application to specify a filter for view, sample, and instance states. This allows the application to request only those samples that have the requested state. For more information on these states see the *Sample Status Information (SampleInfo)* section.

The variations of read() and take() provided by the API are as follows:

method	Behavior
read_w_condition() take_w_condition()	Applies a ReadCondition filter to the samples before returning. The ReadCondition can be a QueryCondition (specialization) which includes an SQL like select statement. This provides for complex filtering of data samples based on status and data content.
read_next_sample() take_next_sample()	This accesses a single sample in order as dictated by the QoS settings.
read_instance() take_instance()	This accesses the data samples of a particular instance specified as an argument.

<code>read_next_instance()</code> <code>take_next_instance()</code>	This provides a mechanism to iteratively access the data samples of all instances. By providing a NIL HANDLE to the first invocation, and then providing the instance handle of the returned samples to subsequent invocations, it is possible to iterate through all instances contained in the DataReader.
<code>read_next_instance_w_condition()</code> <code>take_next_instance_w_condition()</code>	This combines filtering capabilities with instance iteration.

8.3.8 Notification Options (*Determine When Data is Available*)

CoreDX DDS provides a number of status and notifications that are available for the application to learn about events within the CoreDX DDS middleware. An example is the event that data has been received by the DataReader and is available for the application to read (or take). These notification options may also be used to notify the application of other events that may happen within the CoreDX DDS middleware.

8.3.8.1 Using Listeners

Listeners provide asynchronous notification when data is available. There are two listener operations that indicate data is available: the `on_data_available()` method in the `DataReaderListener` and the `on_data_on_readers()` method in the `SubscriberListener`. The subscribing application attaches a listener to the DataReader or Subscriber, and that listener is invoked when data is available.

Additional information on listeners and example listener code can be found in the *Listeners* section of this manual.

8.3.8.2 Using Conditions

Conditions, when combined with WaitSets, provide synchronous notification when data is available by allowing the subscribing application to block until data is available. There are two types of conditions the

subscribing application can use to be notified of available data. The first is a `StatusCondition`. The `DataReader` and `Subscriber` both have a `StatusCondition`. The `DataReader`’s `StatusCondition` will trigger when the `DATA_AVAILABLE_STATUS` on the `DataReader` changes. The `Subscriber`’s `StatusCondition` will trigger when the subscriber’s `DATA_ON_READERS_STATUS` changes. The second type of condition is a `ReadCondition`, which is triggered when data is available on the `DataReader`. The `ReadCondition` allows the application to specify additional criteria that must be met before the Condition is triggered, including instance state, sample state, and view state.

Additional information on Conditions and WaitSets, along with example code can be found in the *Conditions and WaitSets* section of this manual.

8.3.8.3 Using Polling

The application can choose to poll for data, rather than blocking or using callbacks. When polling for data, the application calls `DataReader::read()` or `take()` operation in a loop. If there is data available, these methods return `DDS::RETCODE_OK`, otherwise they return `DDS::RETCODE_NO_DATA`.

8.4 Sample Status Information (SampleInfo)

Calls to any of the `read()` or `take()` variants described above return one or more samples and corresponding *SampleInfo* structures. The *SampleInfo* structure contains metadata about the received sample and includes the following information:

8.4.1 sample_state

The **sample state** is the state of the data *sample*. Valid states are: `READ` and `NOT_READ`.

The sample state is `READ` if this `DataReader` has read this sample previously, otherwise the state is `NOT_READ`.

8.4.2 view_state

The **view state** indicates this `DataReader`’s view of the data *instance*. Valid states are: `NEW` and `NOT_NEW`.

The view state is `NEW` if this `DataReader` has never read a sample from this instance, otherwise the state is `NOT_NEW`. The view state can also be `NEW` if this is the first sample received since the instance was disposed.

8.4.3 *instance_state*

The **instance state** is the state of the instance. Valid states are: ALIVE, NOT_ALIVE_DISPOSED, and NOT_ALIVE_NO_WRITERS.

The instance state is ALIVE if there is at least one DataWriter actively writing samples on this instance. The instance state is NOT_ALIVE_DISPOSED if the instance was explicitly disposed by a DataWriter. The instance state is NOT_ALIVE_NO_WRITERS if there are no DataWriters actively writing this instance.

8.4.4 *source_timestamp*

The **source timestamp** is the timestamp provided by the DataWriter at the time the sample was produced.

8.4.5 *instance_handle*

The **instance handle** is a unique identifier for this sample's instance.

8.4.6 *publication_handle*

The **publication handle** is a unique identifier for the DataWriter who wrote this sample.

8.4.7 *disposed_generation_count*

The **disposed generation count** is a count of the number of times the instance has come alive *after* being disposed. In other words, any time the instance state changes from NOT_ALIVE_DISPOSED to ALIVE.

This count can be used to determine the number of times an instance has been disposed. Initially, it is 0.

8.4.8 *no_writers_generation_count*

The **no writers generation count** is a count of the number of times a DataWriter has started writing data on the instance *after* being declared NOT_ALIVE_NO_WRITERS. In other words, any time the instance state changes from NOT_ALIVE_NO_WRITERS to ALIVE.

This count can be used to determine the number of times an instance has not been alive due to no active readers. Initially, it is 0.

8.4.9 *sample_rank*

The **sample rank** is the number of samples in this instance that follow this one in the current read (or take) collection.

The sample rank can be used to determine the ‘sample age’ of the current sample, relative to the number samples for the instance in the returned sample set.

8.4.10 *generation_rank*

The **generation rank** is the number of times this instance has transitioned from not-alive to alive in the time between the reception of this sample and the latest sample for this instance in the current read (or take) collection.

The generation rank can be used to determine the ‘generation age’ of the current sample, relative to the number of samples for this instances in the returned sample set.

8.4.11 *valid_data*

The **valid data** flag indicates the sample data associated with this SampleInfo is valid. Valid values are zero (FALSE) and non-zero (TRUE).

The valid data flag is set to true when a data sample is received. The valid data flag is set to false when an unregister or dispose command is received. (There is still a sample returned, however, only the key fields, if any, will be valid in this sample.)

8.5 Additional Subscriber / DataReader Features

8.5.1 *Filtering Data*

There are two basic options for filtering received data.

1. Filter data that is received by the DataReader (filtered data is never available to the application).
2. Filter data as it is read by the application (filtered data is still available in the DataReader for future reads or takes by the application).

A ContentFilteredTopic or the Time Based Filter QoS policy is used to achieve the first option. A DataReader created on a ContentFilteredTopic will not store the filtered data, and so it is never available to the application.

Refer to the *Content Filtered Topics* section for additional information on `ContentFilteredTopics`. Similarly, data filtered by a configured Time Based Filter QoS policy is not added to the `DataReader` cache, and so it is never available to the application.

The second option may be achieved by using the `read_w_condition()` (or `take_w_condition()`) API, or by using a `WaitSet` with a read condition attached. Both the `read_w_condition()` / `take_w_condition()` API and `WaitSets` allow filtering using *Read Conditions* or *Query Conditions*.

Read Conditions allow the application to filter by the state and view of the data in the `DataReader` cache. *Read Condition* filters parameters include:

- Sample State: has the data sample been 'read' or not
- View State: is the data sample newly received since the application last accessed the data cache (via a `read()` or `take()` operation)
- Instance State: is the instance *alive*, *disposed*, or *unregistered* (see xyz for additional information on Instances).

Query Conditions allow the application to filter on the data contents of each sample. These conditions are provided as an SQL-like query string, and only data that matches the specified query is returned to the application. Refer to the *Content Filtered Topics* section for additional information on the query syntax.

8.5.2 Wait for Historical Data

`DataReaders` with a Durability QoS policy configured to *Transient Local*, *Transient*, or *Persistent* may receive historical data published before this `DataReader` was enabled. The `DataReader` provides an API that will block the application until all available historical data has been received:

```
DataReader::wait_for_historical_data( Duration_t
max_wait)
```

When this method is invoked, the application will block until all historical data (all previously published data samples) have been received and are available for the application to read or until `max_wait` has expired, whichever occurs first.

This method is not applicable when the `DataReader`'s Durability QoS policy is configured to *Volatile*; in this case, `wait_for_historical_data()` will return immediately.

8.6 QoS Policies

Every created DDS entity has an associated **Quality of Service** (QoS) that can be specified when creating the entity, or later by calling the `set_qos()` method on that entity. The QoS for each entity is a comprehensive set of configuration policies that affect the behavior of that entity. The middleware defines a default QoS for each entity, which was used in the examples above.

Below is a table listing the available QoS for the subscribing entities (DomainParticipant, Subscriber, Topic, and DataReader). This table lists only brief descriptions. A more complete list of QoS Policies and their descriptions can be found in the *Quality of Service Features* chapter.

Table 8-1: QoS Policies for Subscribing Entities

QoS Policy	Description	Available To
USER_DATA	Data not used by the CoreDX DDS middleware, the application can use this data for its own purposes.	DomainParticipant, DataReader
TOPIC_DATA	Data not used by the CoreDX DDS middleware, the application can use this data for its own purposes.	Topic
GROUP_DATA	Data not used by the CoreDX DDS middleware, the application can use this data for its own purposes.	Subscriber
DURABILITY	Specifies if the DataReader would like to receive older data that was published before the DataReader came online (in other words, the history data).	Topic, DataReader
DURABILITY_SERVICE	Not Yet Implemented Specifies configuration for the disabilities: TRANSIENT and PERSISTENT.	Topic, DataReader
PRESENTATION	Affects how data is presented to the subscribing application. For example, published data can be grouped together	Subscriber

	such that the DataReaders receive all coherent data together.	
DEADLINE	Establishes an expectation that the publisher of data will updated the data at least once every specified time period.	Topic, DataReader
OWNERSHIP	Specifies if multiple DataWriters are allowed to write (or update) the same instance of the data, and how these modifications should be handled.	Topic, DataReader
LIVELINESS	Indicates an expectation of the DataReader that the DataWriter will signal it's liveliness in the specified interval.	Topic, DataReader
PARTITION	A logical partition among Topics visible to Publishers and Subscribers. A publisher will only communicate with a subscriber if their Partitions match.	Subscriber
RELIABILITY	Indicates the level of reliability expected of all matched DataWriters.	Topic, DataReader
DESTINATION_ORDER	Specifies the order in which changes to an instance will be ordered by the Subscriber.	Topic, DataReader
HISTORY	Specifies if the subscribing middleware should keep any (or all) updates to an instance (history).	Topic, DataReader
RESOURCE_LIMITS	Specifies the resources the middleware can consume in order to meet the requested QoS.	Topic, DataReader
RTPS_Reader	Configures CoreDX DDS extensions to DataReader behavior	DataReader

DISCOVERY	Configure CoreDX DDS extensions for transport and discovery options	DomainParticipant
PEER_PARTICIPANT	Configure static peers to communicate with	DomainParticipant
THREAD_MODEL	Configuring threading options (single threaded / multithreaded)	DomainParticipant
ENTITY_FACTORY	Specifies if a factory should automatically enable created entities. If the factory does not automatically enable those entities, the application must specifically enable them before they can be used for receiving data.	DomainParticipantFactory, DomainParticipant, Subscriber

8.7 Available Listeners

Listeners are one mechanism allowing the application to be made aware of events and changes in the CoreDX DDS middleware communication status. We illustrated an example of one kind of listener above, the `DataReaderListener`, used for receiving data. There are additional listeners available to a subscribing application. The application can define one or more listener methods and attach them to an appropriate DDS entity when the entity is created, or later by using the `set_listener()` method on the entity.

The table below lists the listeners that can be attached to a subscribing application. A more complete list and description of listeners can be found in the *Communication Status* chapter.

Table 8-2: Listeners for Subscribing Entities

Listener	Listener Methods	Description
DataReaderListener	<code>on_requested_deadline_missed()</code>	The deadline this DataReader was expecting through its QoS DEADLINE was missed.

Listener	Listener Methods	Description
	<code>on_requested_incompatible_qos()</code>	A DataWriter has been discovered that has a QoS configuration incompatible with this DataReader's QoS
	<code>on_liveliness_changed()</code>	One or more of the DataWriters this DataReader was receiving samples data from has changed liveliness (either became ACTIVE or INACTIVE)
	<code>on_subscription_match()</code>	A DataWriter has been discovered that matches the Topic and has a compatible QoS configuration to this DataReader
	<code>on_sample_rejected()</code>	A received sample has been rejected by this DataReader because a RESOURCE_LIMITS QoS setting has been exceeded. The sample is not available to the application.
	<code>on_data_available()</code>	New information (data sample(s) or sample information) is available
	<code>on_sample_lost()</code>	A sample has been lost (not received by this DataReader). This sample is not available to the application.
SubscriberListener	<code>on_data_on_readers()</code>	<p>Data has been received and is available on one or more DataReaders attached to this Subscriber.</p> <p>(SubscriberListener also inherits methods from the DataReaderListener)</p>
TopicListener	<code>on_inconsistent_topic()</code>	Another, different, Topic exists with the same name as this Topic.

Listener	Listener Methods	Description
DomainParticipantListener	(none)	(DomainParticipantListener inherits methods from the DataReaderListener, SubscriberListener and the TopicListener)

Chapter 9 Topics

9.1 Overview

Topics logically connect publications and subscriptions. Publications must be known in such a way that subscriptions can refer to them unambiguously. A Topic is meant to fulfill that purpose: it associates a name, a data-type, and QoS related to the data itself.

Each topic corresponds to one data type. However, several topics may refer to the same data type.

Topics have a **Quality of Service** (QoS) that conceptually describes the data written to that topic. The topic QoS can be specified when creating the topic, or later by calling the `set_qos()` operation on the topic. The QoS defined for the topic is **not used** by CoreDX DDS, but may be used by the application as a hint for the QoS of the corresponding DataReaders and DataWriters. Additional information on QoS policies can be found in the *Quality of Service Features* chapter.

There are several variations of a topic. The *base class* for all topics is a **TopicDescription**. The TopicDescription contains the topic name and data-type name. There are three (3) variations of a TopicDescription. They are listed in Table 9-1.

Table 9-1: Topic Variants

Topic Variants	Description
Topic	The basic form of a TopicDescription, it contains a description of the data to be published and subscribed to, including QoS and Listeners.
ContentFilteredTopic	This topic allows for content-based subscriptions, that is, a subscription that receives a subset of the data published based on a SQL-like query condition.

MultiTopic	This topic allows for combining and filtering data from several topics.
	CoreDX DDS currently does not support MultiTopics .

Topics are created using one of the `create_topic()` operation variations provided from the `DomainParticipant`.

In general, a publishing application will create a `Topic`, and associate each `DataWriter` to exactly one `Topic`. A subscribing application will create a `Topic` and if using a content filter, create a `ContentFilteredTopic` based on that `Topic`, and associate each `DataReader` to exactly one `TopicDescription` or `ContentFilteredTopicDescription`.

9.2 Built-In Topics

The CoreDX DDS infrastructure manages a set of built-in topics. These topics are created when a `DomainParticipant` is initialized, and keep track of discovery information about other DDS participants, Topics, `DataReaders`, and `DataWriters`. This information is necessary for the DDS discovery to work properly, and may also be useful to applications that want to react to this discovery information.

Table 9-2 lists the built-in topics and their associated data types.

Table 9-2: Built-in Topics

Built-in Topic Name	Data Type Name	Description
DCPSParticipant	<code>DDS::ParticipantBuiltinTopicData</code>	Each sample is a description of a DDS participant that has been discovered by this <code>DomainParticipant</code>
DCPSTopic	<code>DDS::TopicBuiltinTopicData</code>	Not supported by CoreDX DDS. The presence of a <code>Topic</code> is implied by the presence of a <code>DataReader</code> or <code>DataWriter</code> on the <code>Topic</code> .
DCPSPublication	<code>DDS::PublicationBuiltinTopicData</code>	Each samples is a description of a <code>DataWriter</code> discovered by this

DomainParticipant		
DCPSSubscription	DDS::SubscriptionBuiltinTopicData	Each sample is a description of a DataReader discovered by this DomainParticipant

In general, the built-in data types hold information about the discovered entity’s QoS configuration, along with other useful information. For a detailed description of these built-in data types, refer to the `dds_builtin.h` or `dds_builtin.hh` header files in the `COREDX_TOP/target/COREDX_TARGET/include/dds` directory.

Each built-in topic has a type-specific DataReader associated with it (DCPSParticipantDataReader, etc.). The application can use these DataReaders to access the data and statuses from the built-in topics in the same way any user defined DataReader would do this.

To get access to these built-in DataReaders, the application can call

```
DomainParticipant::get_builtin_subscriber()
```

This Subscriber can be used to access specific built-in data readers by calling

```
Subscriber::lookup_datareader(topic_name)
```

and using the appropriate topic name from

Table 9-2. For example:

```
Subscriber::lookup_datareader("DCPSPublication")
```

Once the application has gotten the desired built-in DataReader, it can use any of the normal data access mechanisms: asynchronous notification via a listener, direct polling via `read()`, or blocking with a `WaitSet`.

The built-in topics use data types specified in the DDS standard for communicating discovery data. The following tables illustrate the data type of each of the built-in topics.

Table 9-3: Participant Built-in Data Type

ParticipantBuiltinTopicData

```

struct ParticipantBuiltinTopicData
{
    DDS_KEY    BuiltinTopicKey_t    key;
              UserDataQosPolicy    user_data;
};

```

Table 9-4: Topic Built-in Data Type

TopicBuiltinTopicData

```

struct TopicBuiltinTopicData
{
    DDS_KEY    BuiltinTopicKey_t    key;
              string                name;
              string                type_name;
              DurabilityQosPolicy    durability;
              DurabilityServiceQosPolicy durability_service;
              DeadlineQosPolicy      deadline;
              LatencyBudgetQosPolicy latency_budget;
              LivelinessQosPolicy    liveliness;
              ReliabilityQosPolicy    reliability;
              TransportPriorityQosPolicy transport_priority;
              LifespanQosPolicy       lifespan;
              DestinationOrderQosPolicy destination_order;
              HistoryQosPolicy        history;
              ResourceLimitsQosPolicy resource_limits;
              OwnershipQosPolicy      ownership;
              TopicDataQosPolicy      topic_data;
};

```

Table 9-5: Publication Built-in Data Type

PublicationBuiltinTopicData		
struct PublicationBuiltinTopicData		
{		
DDS_KEY	BuiltinTopicKey_t	key;
	BuiltinTopicKey_t	participant_key;
	string	topic_name;
	string	type_name;
	DurabilityQosPolicy	durability;
	DurabilityServiceQosPolicy	durability_service;
	DeadlineQosPolicy	deadline;
	LatencyBudgetQosPolicy	latency_budget;
	LivelinessQosPolicy	liveliness;
	ReliabilityQosPolicy	reliability;
	LifespanQosPolicy	lifespan;
	UserDataQosPolicy	user_data;
	OwnershipQosPolicy	ownership;
	OwnershipStrengthQosPolicy	ownership_strength;
	DestinationOrderQosPolicy	destination_order;
	PresentationQosPolicy	presentation;
	PartitionQosPolicy	partition;
	TopicDataQosPolicy	topic_data;
	GroupDataQosPolicy	group_data;
};		

Table 9-6: Subscription Built-in Data Type

SubscriptionBuiltinTopicData		
struct SubscriptionBuiltinTopicData		
{		
DDS_KEY	BuiltinTopicKey_t	key;
	BuiltinTopicKey_t	participant_key;
	string	topic_name;
	string	type_name;
	DurabilityQosPolicy	durability;
	DeadlineQosPolicy	deadline;
	LatencyBudgetQosPolicy	latency_budget;

```

    LivelinessQosPolicy      liveliness;
    ReliabilityQosPolicy     reliability;
    OwnershipQosPolicy       ownership;
    DestinationOrderQosPolicy destination_order;
    UserDataQosPolicy        user_data;
    TimeBasedFilterQosPolicy time_based_filter;
    PresentationQosPolicy    presentation;
    PartitionQosPolicy        partition;
    TopicDataQosPolicy        topic_data;
    GroupDataQosPolicy        group_data;
};

```

9.3 Content Filtered Topics

A ContentFilteredTopic is a specialization of TopicDescription that allows the subscribing application to describe a subscription where it will only see a subset of the data published, based on a defined content filter. The filter is an SQL like statement. The ContentFilteredTopic is associated with another known Topic and applies a filter to the data available on that related topic.

ContentFilteredTopics are created by a DomainParticipant, just like normal Topics.

```
DomainParticipant::create_contentfilteredtopic()
```

This method call has additional parameters to specify which topic the ContentFilteredTopic is associated with, the SQL query expression, and parameters (if any) for use in evaluating the filter expression.

Table 9-7: create_contentfilteredtopic() parameters

Parameter	Description
Topic	Related topic. The ContentFilteredTopic presents a filtered subset of data available on the related topic.
filter_expression	SQL like condition expression
filter_parameters	String sequence of parameters used in the filter_expression .

The **filter_expression** must be a valid SQL WHERE clause (without the **WHERE** keyword). For example “x<=4”. The filter expression refers to structure members in the application defined data type associated with the related Topic. For embedded structures, the naming convention uses a dot (‘.’) to separate field names. For example, “time.sec > 10”. Table 9-8 lists all the operators available when constructing the condition.

Table 9-8: Valid Condition Operators for Content Filters

Operator	Description
=	Equals
<>	Not equals
>=	Greater than or equal
>	Greater than
<=	Less than or equal
<	Less than
NOT, not	Not operator
()	Parenthesis are used for nesting conditions
AND, and	And operator
OR, or	Or operator
IN, in	In operator for matching a value to something in a list
BETWEEN, between	For future use: the between operator is not yet supported by CoreDX DDS.
LIKE, like	For future use: the like operator is not yet supported by CoreDX DDS.

The filter expression can also contain references to parameters, present in the **filter_parameters** argument. Parameter references take the form “%<number>”. The number is an index into the filter_parameters sequence, and starts at zero. For example “%2” refers to the third parameter in the filter_parameters sequence.

Once created ContentFilteredTopics can be used by a DataReader just like a normal Topic. The filter expression is static, and cannot be changed after the ContentFilteredTopic is created; however, filter parameters can be changed on the fly with a call to

```
ContentFilteredTopic::set_expression_parameters()
```

Content Filter expressions can contain references to basic data types. For example, data members of type int, short, long, and string are all valid field types for a filter expression; however, sequences, arrays, or unions are not.

9.3.1 Content Filter Example

The full code for a content filter example can be found in the examples directory of the CoreDX DDS release.

Table 9-9: Creating a ContentFilteredTopic

Creating a ContentFilteredTopic (C language)

```
DDS_DomainParticipant    dp;
DDS_Subscriber           sub;
DDS_Topic                top;
DDS_ContentFilteredTopic cftop;
DDS_TopicDescription      td;
DDS_DataReader           dr;
DDS_StringSeq            cf_params;

dp = DDS_DomainParticipantFactory_create_participant( 1,
    DDS_PARTICIPANT_QOS_DEFAULT, NULL, 0);
sub = DDS_DomainParticipant_create_subscriber(dp,
    DDS_SUBSCRIBER_QOS_DEFAULT, NULL, 0);
MyTypeTypeSupport_register_type( dp, "topic_type" );
top = DDS_DomainParticipant_create_topic(dp, "topic_name",
    "topic_type", DDS_TOPIC_QOS_DEFAULT, NULL, 0);

/* BUILD A CONTENT_FILTERED TOPIC */
cftop = DDS_DomainParticipant_create_contentfilteredtopic(dp,
    "cf_topic_name", top, "x<%0", NULL);
```

```
/* parameters can be specified/modified after creation */
INIT_SEQ(cf_params);
seq_set_size(&cf_params, 1);
seq_set_length(&cf_params, 1);
cf_params._buffer[0] = "5";
DDS_ContentFilteredTopic_set_expression_parameters(cftop,
    &cf_params);

td = DDS_Topic_TopicDescription((DDS_Topic)cftop);
dr = DDS_Subscriber_create_datareader(sub, td, &dr_qos,
    NULL, 0);
if (!dr)
    printf("FAILED to create DR!\n");
```

9.3.2 Configuring Content Filters

When a DataReader uses a content filter to filter the received data, the filter expression is communicated to any matching DataWriter(s). This allows CoreDX DDS to filter data either at the DataWriter or the DataReader.

The DataReader's content filter is always enabled. This ensures the specified data is always filtered, even when the DataReader is matched with a DataWriter that does not support writer-side-filtering (for example, this may happen when interoperating with another DDS implementation).

DataWriter content filtering is configurable by the DataWriter QoS policy: RTPSWriterQoSPolicy.apply_filters (true or false value).

DataWriter filtering is enabled by default. This means the DataWriter will write data to a DataReader only when it passes the DataReader's content filter (or if the DataReader does not have a content filter configured). This configuration can reduce the 'work' performed by the DataReader to apply the filter, but also means the DataWriter must unicast all samples that are filtered by at least 1 matched DataReader.

When the DataWriter is configured to NOT apply filters, the DataWriter will always multicast written data samples, allowing the DataReader to apply the filter.

9.3.3 Compiling an application with Content Filters

CoreDX DDS provides an alternate library that contains the Content Filter capability. This allows the basic CoreDX DDS library to stay extremely small.

To compile an application that uses content filters, change the library line in your Makefile, replacing `libdds` with `libdds_cf`. For the C++ language binding, you will *also* need to replace `libdds_cpp` with `libdds_cpp_cf`.

The Java and C# language binding contain all CoreDX DDS functionality in one library, including the content filter API and functionality.

9.4 Multi Topics

A MultiTopic is a specialization of TopicDescription that allows the application to describe a subscription that combines, filters, and orders data from multiple Topics. This is similar to the **JOIN** statement in SQL and Relational Database Systems.

CoreDX DDS currently does not support multi topics.

Chapter 10 Instances and Samples

Data is the core of any communications middleware, and it is especially important to a data-centric, publish-subscribe middleware like CoreDX DDS. This chapter describes how the CoreDX DDS middleware handles and classifies the data, and how the data is packaged and communicated between the application and the CoreDX DDS middleware.

10.1 Overview

Each Topic is attached to a DDS **data type**. Only data of that type may be published on the Topic. The DDS data type is always a structure, which can be made up of virtually any user defined type. The following is an example of a DDS data type.

```
struct HelloMessage {
    long        time_sent;
    long        sender_id;
    string       sender_name;
    string       msg;
    sequence<string> msg_history;
}
```

The type name for the above DDS data type is “HelloMessage”.

Additional information on the generating DDS data types can be found in the *Application Data Types* chapter.

The CoreDX DDS middleware classifies data into **samples** and **instances**. A **sample** is data of the appropriate DDS data type that has been published to a DDS Topic. The following is a possible **sample** of the HelloMessage data type:

```
123456789
42
"Bob The Builder"
"Hello out there!"
<empty sequence>
```

The application developer, when creating a DDS data type, can specify one or more attributes of the DDS data type as a **key**. The CoreDX DDS middleware uses those key attributes to organize the published data. Going back to the HelloMessage example above, an application developer might specify the “sender_id” field as the key for the HelloMessage data type.

A publishing application might write the following four (4) **samples**:

```
123456789
42
"Bob The Builder"
"Hello out there!"
<empty sequence>
```

```
223456789
45
"Wendy"
"Busy day today"
<empty sequence>
```

```
323456789
42
"Bob The Builder"
"Can We Build It?"
<empty sequence>
```

```
423456789
12
"Scoop"
"Yes we can!"
<empty sequence>
```

If the Data Type defines the “sender_id” field as the key for the HelloMessage data type, then in the 4 samples published, there are 3 unique keys. The CoreDX DDS middleware classifies all samples with the same key value to be one (1) **instance**. In this example, 3 instances have been published, with the following keys: 42, 45, and 12. This is depicted in Table 10-1.

Table 10-1: Instance Example

Instance 1, Key=42	Instance 2, Key=45	Instance3, Key=12
2 samples	1 sample	1 sample
123456789 42 "Bob The Builder" "Hello out there!" <empty sequence>	223456789 45 "Wendy" "Busy day today" <empty sequence>	423456789 12 "Scoop" "Yes we can!" <empty sequence>
323456789 42 "Bob The Builder" "Can We Build It?" <empty sequence>		

The CoreDX DDS middleware stores and manages data ***samples*** (an individual structure provided to write() and returned from read()) and ***instances*** (a collection of zero or more samples with the same ***key*** value).

10.2 Publishing Data

On the publishing side of DDS communications, ***samples*** represent data that is sent to DataReaders. Samples are created for every write(), unregister(), and dispose() call made by the application. Each sample written is associated with a particular instance. In general, samples and instances are stored by the DataWriter until they are delivered to all appropriate DataReaders, at which point the samples and instances may be removed. The specific rules for maintaining ***samples*** in the DataWriter are different from the rules for managing ***instances***. For this reason, it is possible for all samples on an instance to be removed from the DataWriter, while the instance remains (with no associated samples). In contrast, it is *not* possible to remove an instance from the DataWriter while any samples associated with it remain.

Data ***instances*** are used to manage several DataWriter QoS policies. ***Instances*** allow the application to set *Deadlines*, keep *History*, manage *Ownership*, and follow *Resource Limits*. For additional information on these QoS policies, see the *Quality of Service Features* Chapter.

10.3 Subscribing to Data

On the subscribing side of CoreDX DDS communications, **samples** represent the data that has been received by the middleware and may be made available to the subscribing application (filters or other QoS policy settings may preclude samples from reaching the application). Each received sample is associated with a particular **instance**.

In general, samples and instances are stored in the DataReader until they are explicitly removed by the subscribing application (see *Read (or Take) Data* for details), or the CoreDX DDS middleware removes them based on various QoS policy settings. Similar to the DataWriter management, the specific rules for maintaining and deleting samples from the DataReader are different from the rules for maintaining and deleting instances. For this reason, it is possible for all samples on an instance to be removed, while the instance remains (with no associated samples). In contrast, it is *not* possible to remove an instance while any samples associated with it remain.

10.4 Instance Lifecycles

Instances are used to manage the data lifecycle. Instances are created (registered), updated (written), and deleted (unregistered or disposed). For example, consider the three instances above. Instance key=12 (Scoop) may shut down for the day while the other two instances (Bob and Wendy) are still working (and updating). The publishing application can call `dispose()` on Scoop (instance key=12) and the remaining instances (Bob and Wendy) will still be alive. When Scoop comes online again, the publishing application can start updating that instance (which will retain the same instance handle). These data lifecycle operations are covered in detail in the following sections.

10.4.1 Registering Instances

Instances must be **registered** with the DataWriter before any samples associated with that instance can be written (or deleted). As a convenience, CoreDX DDS will automatically register an instance when the application calls one of the `write()`, `unregister_instance()`, or `dispose()` operations without first registering the instance.

Publishing applications can also explicitly register an instance by calling `DataWriter::register_instance()`. The `register_instance()` operation returns an instance handle which can be used to improve the performance of subsequent calls to `write()`.

The below example illustrates the use of `DataWriter::register_instance()` and `DataWriter::write()`.

Example (C++)

```

HelloMessageDataWriter  dw;
HelloMessage            bobData, wendyData;
InstanceHandle_t        bobHandle, wendyHandle;
ReturnCode_t            retval;

bobData.sender_id       = 42;
bobData.sender_name     = strdup("Bob the Builder");
bobData.msg             = strdup("Hello");

wendyData.sender_id     = 45;
wendyData.sender_name   = strdup("Wendy");
wendyData.msg           = strdup("Good Morning, Bob!");

/* calling write() without an instance handle forces
 * CoreDX DDS to register this instance (key=42)
 */
retval = dw .write( bobData, HANDLE_NIL );

/* the instance can later be 'looked up'
 */
bobHandle = dw .lookup_instance( bobData );

/* Calling register_instance() first allows for
 * subsequent optimized calls to write()
 */
wendyHandle = dw .register_instance( wendyData );
retval = dw .write( wendyData, wendyHandle );

delete[] wendyData.msg;
wendyData . msg = strdup( "Good night, everyone!" );

/* Changing the 'msg' in wendyData does not change the
 * key value, and therefore does not change the instance
 * or instance handle.
 */
retval = dw .write( wendyData, wendyHandle );

```

Figure 10-1: Register Instances Example

Register instance operations are applicable only to DataWriters, and are not communicated to DataReaders. In fact, DataReaders do not have a concept of *registered* instances. Instead, a DataReader has a concept of an *Alive*

instance state. Instances have an Alive state if they have at least one alive DataWriter actively writing data samples on the instance. Refer to *Section 8.4 Sample Status Information (SampleInfo)* for additional information on *instance states*. In order to manage instance states, DataReaders store a list of alive, actively writing DataWriters for each instance.

CoreDX DDS can support millions of unique instances in each DataWriter and DataReader.

10.4.2 Unregistering Instances

The publishing application can **unregister** previously registered instances by calling the `DataWriter::unregister_instance()` operation. This indicates the application will no longer be writing any samples for this instance. This is not the same as disposing the instance (which indicates the instance no longer exists). The unregister operation simply indicates that *this* DataWriter will no longer be publishing updates on the instance. CoreDX DDS treats an unregister command very much like an application written data sample. The unregister command is stored by the DataWriter to be communicated to matched DataReaders. If the instance for this unregister sample is not already registered, CoreDX DDS will automatically register it before processing the unregister operation.

After an unregister operation, the instance handle associated with the unregistered instance is invalid. This is because the CoreDX DDS middleware may have removed all records of that instance. After an unregister operation, the instance handle may be reused for a different instance. The application may re-register the instance and then continue to publish samples or a dispose on the instance. Unregister operations are communicated to matched DataReaders as a sample with the *valid_data* flag set to false and indicate that the DataWriter is no longer actively writing on this instance. The DataReader will remove this DataWriter from the instance’s list of active DataWriters. When this list of alive, actively writing DataWriters becomes empty, the state of the instance in the DataReader Cache will change to `NOT_ALIVE_NO_WRITERS`. Refer to *Sample Status Information (SampleInfo)* for additional information on *instance states*.

When a DataWriter is deleted by the publishing application, CoreDX DDS will automatically send an unregister message to matched DataReaders for every instance the DataWriter has written and still knows about. If the publishing application exits without deleting its DataWriter Entities, the DataReader will not receive the unregister commands. In this case, the DataReader will eventually determine the DataWriter is not alive (based on

Liveliness QoS configurations), and remove the DataWriter from the list of alive and actively writing DataWriters on each instance the DataReader knows about.

10.4.2.1 Relationship between DataWriter::unregister and other QoS Policies

Unregistering an instance on a DataWriter configured for Exclusive Ownership (via the Ownership QoS Policy) will cause the DataWriter to relinquish ownership of the instance. Other Exclusive Ownership DataWriters may take over ownership for the instance.

DataWriters configured with Volatile Durability will remove the instance (and all associated samples) after all *currently alive and matched* DataReaders acknowledge receiving all samples on the instance.

DataReaders configured with `auto_purge_no_writers` set in their Reader Data Lifecycle may delete the instance (and all associated samples), if the DataWriter that sent this unregister command was the last active DataWriter on this instance.

10.4.3 Disposing Instances

The publishing application can **dispose** previously registered instances by calling the DataWriter::dispose() operation. This indicates that the instance no longer exists. This is different than the unregister operation (which indicates this DataWriter is no longer writing on this instance). CoreDX DDS treats a dispose very much like an application written data sample. The dispose command is stored by the DataWriter to be communicated to matched DataReaders. If the instance for this dispose sample is not already registered, CoreDX DDS will automatically register it before processing the dispose operation.

Dispose operations are communicated to matched DataReaders as a sample with the *valid_data* flag set to false. The DataReader will change the state of the associated instance to NOT_ALIVE_DISPOSED. Refer to *Sample Status Information (SampleInfo)* for additional information on *instance states*. Note that the DataWriter will still be considered an alive, actively writing DataWriter on this instance.

10.4.3.1 Relationship between Dispose and other QoS Policies

DataReaders configured with `auto_purge_disposed` set in their Reader Data Lifecycle will delete the instance (and all associated samples) on receipt of a dispose command.

10.4.4 Instance Handles

An **instance handle** is a value that can be used to uniquely identify a *registered* instance within one DataWriter or DataReader. An instance handle is generated by a DataWriter when an instance is registered (returned from a DataWriter::register_instance() call) and for that DataWriter it will be used to identify the instance until that instance is unregistered. ***The instance handle is valid only for the DataWriter that created it, and only while the instance is registered.*** Once an instance is unregistered, the instance handle can no longer be used to identify that instance. This is a critical detail of the CoreDX DDS middleware. These instance handles are reused, and after an unregister operation, the old instance handle may identify a *different* instance. If the unregistered instance is re-registered, a different handle may be assigned for the next ‘life’ of that instance.

10.5 Data Cache

DataReaders and DataWriters contain a Data Cache for storing data samples and instances. The DataWriter’s Data Cache contains samples and instances that have been written, and the DataReader’s Data Cache contains samples and instances that have been received. These data caches are sized and managed according to the configuration of the several QoS policies as shown in the figure below.

Table 10-2: Instance Example

QoS Policy	Configuring this policy on the:	Will effect Cache Management on the:
RELIABILITY	DataWriter and DataReader	DataWriter
	DataReader	DataReader
DURABILITY	DataWriter	DataWriter and DataReader
HISTORY	DataWriter	DataWriter
	DataReader	DataReader
RESOURCE_LIMITS	DataWriter	DataWriter
	DataReader	DataReader

QoS Policy	Configuring this policy on the:	Will effect Cache Management on the:
READER_DATA_LIFECYCLE	DataReader	DataReader
WRITER_DATA_LIFECYCLE	DataWriter	DataReader
OWNERSHIP	DataWriter	DataReader
LIFESPAN	DataWriter	DataWriter and DataReader
Filters (TIME_BASED_FILTER , content filters)	DataReader	DataReader

In general, data *samples* are added to the Data Cache as they are written (for a DataWriter) or received (for a DataReader). In general, these samples are removed when they are no longer needed by the application or the CoreDX DDS middleware. The specific management of samples in the Data Caches is described in the following sections.

In general, data *instances* are added to the Data Cache as they are registered (for a DataWriter) or received (for a DataReader). In general, these instances are removed when they are no longer needed by the application or the CoreDX DDS middleware. The management and removal of instances is different than samples, and in fact, it is possible for instances to be managed in the DataWriter and DataReader Data Caches even when there are no samples associated with the instance in the DataCache. The specific management of instances in the Data Caches is described in the following sections.

10.5.1 DataWriter Cache

The DataWriter's Data Cache contains samples that have been written by a call to DataWriter::write() variant and instances that have been registered (or written, since calling DataWriter::write() automatically registers an instance).

10.5.1.1 Samples in the DataWriter Cache

Samples are added to the Data Cache as they are written by the application. Samples will be stored in the DataWriter Cache until they are initially written to the wire, and may be stored longer to meet Reliability and Durability QoS settings. Samples can be removed from the cache by the application (calling `DataSetWriter::unregister()`) or by the CoreDX DDS infrastructure, based on QoS policy settings. Samples are removed from the DataWriter cache under the following conditions:

1. The CoreDX DDS middleware in the publishing application has completed writing the sample. This happens when:
 - a. The CoreDX DDS middleware writes the sample to all *Best Effort* DataReaders AND
 - b. (Only if the DataWriter is *Reliable* and *Volatile*) The CoreDX DDS middleware has received an acknowledgement from all matched *Reliable* DataReaders
2. OR: Samples expire based on the Lifespan duration
3. OR: A DataWriter has a History QoS Policy *kind* of *KEEP_LAST* and the cache already holds History *depth* samples and a new sample is created by `write()`, `unregister()` or `dispose()`
4. OR: A Best Effort DataWriter has non-INFINITE *max_samples* or *max_samples_per_instance* Resource Limits and the cache already holds the maximum samples and a new sample is created by `write()`, `unregister()` or `dispose()`.

It is important to note that some combinations of QoS settings will cause the DataWriter cache to grow infinitely, consuming more and more run time memory. This can happen with a Durability *kind* of *TRANSIENT_LOCAL* and Reliability *kind* of *RELIABLE*, a History *kind* of *KEEP_ALL*, and Resource Limits set to infinite. With this combination of QoS settings, the application must manage instances to meet application or machine resource limitations.

10.5.1.2 Instances in the DataWriter Cache

Instances are added to the DataWriter Cache when the publishing application registers an instance that is not already registered. Every sample belongs to an instance, and the instance must be registered before a sample on that instance samples can be created. The application can explicitly register an instance by calling `DataSetWriter::register_instance()`, or CoreDX DDS will automatically register the instance when the application attempts to create a sample without first registering its associated instance.

Instances are removed from the DataWriter Cache when the publishing application unregisters an instance. This must be done explicitly by calling `DataWriter::unregister_instance()`.

10.5.2 DataReader Cache

The DataReader's Data Cache contains samples and instances that have been received (subject to filters). These samples and instances may or may not have been already read (seen) by the application.

10.5.2.1 Samples in the DataReader Cache

Samples are added to the DataReader Cache when they are received by the DataReader **and** the sample passes any filters configured on the DataReader **and** there is room in the DataReader Cache for the new sample. If there is not room in the DataReader Cache, the new sample will be added only if the History QoS policy is configured to `KEEP_LAST`.

Samples are eligible to be removed from the DataReader cache under the following conditions:

1. When the application uses `DataReader::take()`.
2. If the Lifespan QoS is used, samples will be removed from the data cache after they are expired.
3. If the Data Cache is full and the History kind is `KEEP_LAST`, the oldest sample will be removed to make room for a newly received sample.
4. When a Best Effort DataReader has non-INFINITE `max_samples` or `max_samples_per_instance` Resource Limits and the cache already holds the maximum samples and a new sample is received.
5. When a DataReader has non-INFINITE `autopurge_nowriter_samples_delay` or `autopurge_disposed_samples_delay` and an instance state is determined to be `NOT_ALIVE_NO_WRITERS` or `NOT_ALIVE_DISPOSED`; (associated samples will be removed after the specified delay)

10.5.2.2 Instances in the DataReader Cache

Instances are added to the DataReader Cache when a sample belonging to the instance is received by the DataReader and the instance does not already exist in the DataReader Cache. The instance is created even if the associated sample is *not* added to the DataReader Cache. Samples may not be added to the DataReader Cache because of filters applied on the DataReader or because of the configuration of the History and Resource Limits QoS policies.

Instances are eligible to be removed from the DataReader Cache when an instance has a state of NOT_ALIVE_NO_WRITERS and there are no associated samples.

Chapter 11 Application Data Types

11.1 Overview

Every DDS Topic contains one and only one **Data Type**, a user defined data type that is used when communicating on the Topic. In most cases, the application developer defines these DDS Data Type(s) ahead of time so they may be compiled into the application in a language-independent format. The Interface Definition Language (IDL) and eXtensible Markup Language (XML) formats are available for DDS data type definition. A compiler is used to translate these type definitions into the appropriate programming language for inclusion into an application.

CoreDX DDS also supports *dynamic types*, which are Data Types that are not defined at compile time. With dynamic types, the application developer may create data types at run-time using the dynamic types API, or create a data type from a discovered peer DataWriter or DataReader. In this scenario, the application has no knowledge of the data type until the Topic is discovered at run time. A complete discussion of dynamic types can be found in the *Dynamic Types* section of this guide.

11.2 Why Define the Data Types?

CoreDX DDS is data-centric. This means that the structure and contents of application data is known and used by the CoreDX DDS middleware. This allows the CoreDX DDS middleware to perform advanced data management operations that are not available in other message oriented middleware technologies. For example, instance and sample history are enabled by identifying 'key' fields in the data to identify unique data instances. This can be compared to key fields in relational database technologies. Each key uniquely identifies a collection of related records. In DDS, the key is used to identify a data 'instance'. Updates to a data instance are referred to as 'samples'. The CoreDX DDS middleware can maintain historical samples for each instance (see the HISTORY Quality of Service). Furthermore, the CoreDX DDS middleware can apply a Content Filter to data samples. Content filters are comparable to an SQL query that selects a sub-set of available data based on conditions. These are a few examples of the power that a data-centric middleware can offer.

Furthermore, because the data types are well known, the DDS API and compilers can enforce the usage of proper types throughout the application

code. This can avoid potentially difficult-to-find bugs related to subtle mismatches of data throughout the distributed system.

The benefits of data-centric middleware are possible because the middleware has an understanding of the data structures used in your application. Part of the DDS development workflow includes defining the application data types and registering them with the CoreDX DDS middleware.

11.3 Registering Data Types

The `register_type()` function informs the middleware of a `DataType` that will be written (using a `DataWriter::write()` variant) or read (using a `DataReader::read()` or `take()` variant). Along with the data type definition, the `register_type()` method initializes hooks to the type specific code to marshal (serialize) and unmarshal (deserialize) the data, and other useful sorting and filtering routines specific to that data type. These type-specific methods are collectively referred to as the data type’s **TypeSupport**.

CoreDX DDS will allow multiple calls to `register_type()` for the same type name, as long as the associated data type and `TypeSupport` are the same. CoreDX DDS will not allow two different data types to be registered with the same type name. Similarly, `register_type()` may not be called with the same type name if the `TypeSupports` are different. For example, an application may include a compiled “HelloType”, and create a dynamic type with an identical type structure as the compiled “Hello Type”. Both of these types may be registered via `register_type()`, however, the type names **must be different**, because the `TypeSupports` are different.

11.4 Data Types and Discovery

Data Types play an important role in the process of DDS dynamic *discovery*. (Additional information on CoreDX DDS discovery can be found in the *CoreDX DDS Discovery* section of this guide). A `DataWriter` will match with a `DataReader` only if the data type published by the `DataWriter` is compatible the data type subscribed to by the `DataReader`. This provides additional type safety for DDS applications because a `DataReader` cannot receive data in an unexpected format (a commonly occurring programming error when using other communication middleware products).

11.5 Data Architecture

The process of architecting DDS Data Types is very similar to that of designing a Relational Database schema. Because the data structures used

by your application will be conveyed between applications, and possibly over a network, it is important that they be designed with efficiency in mind. This 'data normalization' process is important to effective deployment of DDS technology.

For large or complex systems, data interfaces, connections, and communicated types will become very complex. In these cases, data modeling tools may provide a necessary framework for managing the data types across a system. Many data modeling tools support generating DDS compatible IDL or XML that can be processed by the CoreDX DDS data type compiler.

11.6 Compiled Data Type Definition

Once the appropriate data structures have been designed, they must be written in a format that the CoreDX DDS middleware can understand. CoreDX DDS supports 2 language independent formats: IDL and XML.

IDL is a simple and flexible language for defining data types. It has a rich set of primitive and complex data types, and there are defined mappings from IDL to many common programming languages. This makes IDL a good language for DDS.

11.7 Data Definition Syntax

A DDS Data Type is always a **structure**, which may contain any combination of basic and constructed types, including embedded structures. Table 11-1 is a list of basic types supported in the CoreDX DDS IDL and XML.

Table 11-1: Basic User Defined Types

Basic Type	Description
short	2 bytes
long	4 bytes
long long	8 bytes
unsigned short	2 bytes
unsigned long	4 bytes

unsigned long long	8 bytes
float	4 bytes
double	8 bytes
char	1 byte
wchar	4 bytes
boolean	1 byte
octet	binary data, 1 byte
string	bounded and unbounded
constant	constant type, always a number
typedef	Alias or alternate name

Table 11-2 is a list of constructed types supported by CoreDX DDS.

Table 11-2: Constructed User Defined Types

Constructed Type	Description
struct	Structure type
union	Union type
array	Single or multi dimensional
sequence	bounded and unbounded
enum	Variable sized collection of names with constant values

bitset	Variable sized list of named bit-sized flags
map	Key-value pairs

CoreDX DDS IDL does not support the *any* type.

11.8 IDL and XML Language Mappings

The CoreDX data type compiler generates source code for use in application programs. These files contain, among other things, a translation of the IDL or XML specified data type into a language specific data type. The mapping of IDL or XML to programming language follows the standards of the OMG.

The following subsections provide an overview of the available data types and how they are mapped to each of the supported CoreDX DDS programming languages. A full description of available data types and configuration options as specified in the X-Types and IDL specifications is provided in the CoreDX DDS Type System Programmer's Guide.

11.8.1 Basic Types

Table 11-3: Primitive Data Type Mapping

IDL Type	XML Type	C	C++, Cxx	C#	Java
char	char8	char	char	char	char
wchar	char32	char32	char32	int	int
octet	byte	unsigned char	unsigned char	byte	byte
short	int16	short	short	short	short
unsigned short	uint16	unsigned short	unsigned short	ushort	short
long	int32	int	int	int	int

IDL Type	XML Type	C	C++, Cxx	C#	Java
unsigned long	uint32	unsigned int	unsigned int	uint	int
long long	int64	int64_t	int64_t	long	long
unsigned long long	uint64	unsigned int64_t	unsigned int64_t	ulong	long
float	float32	float	float	float	float
double	float64	double	double	double	double
string	string	char *	char *	String	String

11.8.2 Complex Types

Complex types are described in detail in the CoreDX DDS Type System Programmer’s Guide.

11.9 Creating Data Types

11.9.1 Using IDL

The syntax of IDL is flexible and is similar to other languages such as C. The IDL file can contain any combination of:

- C/C++ style comments
- C compiler directives
 #if, #ifdef, #else, #endif,
 #include, etc.
- namespace / module
- constant definitions
- enumerated type definitions
- typedefs
- structures (these become DDS data types)
- unions
- annotations

Figure 11-1: Example IDL file provides an example of a IDL file.

Example IDL data definition file

```
//=====
// CoreDX DDS IDL example
//=====

struct SenderType
{
    string firstname;
    string lastname;
};

struct StringMsg
{
    SenderType      sender;
    long            time_sent;
    sequence<string> old_msgs;
    string          msg;
};
```

Figure 11-1: Example IDL file

The syntax for these IDL file elements adheres to the OMG's IDL syntax specification as defined in the IDL 4 specification. The IDL type definitions are the primary focus of IDL files. The CoreDX DDS IDL compiler (coredx_ddl) can parse fully compliant IDL files; it will use the DDS type definitions and ignore any non-DDS definitions when generating CoreDX DDS code.

11.9.2 USING XML

The XML syntax allows the same data definitions as IDL, following the format of the TBD schema. The XML file may contain:

- XML comments
- namespace / module
- constant definitions
- enumerated type definitions
- typedefs

- structures (these become DDS data types)
- unions
- annotations

Figure 11-1: Example IDL file provides an example of a XML file.

Example XML dat type definition file

```
//=====
// CoreDX DDS XML example
//=====

<dds:types

xmlns:dds="http://www.omg.org/ptc/2011/01/07/XML Type Representation"

xmlns="http://www.omg.org/ptc/2011/01/07/XML Type Representation">
<struct name="SenderType">
<member name="firstname"
        id="0"
        type="string"/>
<member name="lastname"
        id="1"
        type="string"/>
</struct>
<struct name="StringMsg">
<member name="sender"
        id="0"
        type="nonBasic"
        nonBasicTypeName="SenderType"/>
<member name="time_sent"
        id="1"
        type="int32"/>
<member name="old_msgs"
        id="2"
        type="string"
        sequenceMaxLength="(-1)"/>
<member name="msg"
        id="3"
        type="string"/>
</struct>
</dds:types>
```

11.10 Configuring Data Types

Data types may be configured using *annotations*. Configuration options include specifying key fields, marking fields as optional or mandatory, and specifying the extensibility of the data type.

This document describes the syntax for defining keys. The CoreDX DDS Type System Programmer's Guide describes data type configuration options in detail.

11.10.1 Specifying Keys

The application developer, when creating a DDS data type, can specify one or more attributes of the DDS data type as a **key**. The CoreDX DDS middleware uses those key attributes to organize the data into **instances** (see the *Instances and Samples* chapter for additional information).

A key can be any field in the DDS Data Type, except for *sequences*, *unions*, and *enums*. The application developer may define any number of fields to be a key field. All the fields labeled as a key field are concatenated together to form one key for the DDS Data Type.

There are a few methods to specify a key field in IDL.

1. Annotation before the field

The following example shows how to define keys using the IDL 4 annotation *before* the key field, using IDL.

Example IDL file

```
//=====
// CoreDX DDS IDL example - using keys
//=====

// the 'sender_id' field is defined as a key
struct StringMsg1
{
    struct SenderType {
        string first_name;
```

```
        string last_name;
    }
    sender;
    @key long        sender_id;
    long            time_sent;
    sequence<string> old_msgs;
    string          msg;
};

// the 'sender' field is defined as a key
struct StringMsg1
{
    @key struct SenderType {
        string first_name;
        string last_name;
    }
    sender;
    long        sender_id;
    long        time_sent;
    sequence<string> old_msgs;
    string      msg;
};

// the key is a combination of the sender's last name
// and the sender_id
struct StringMsg1
{
    struct SenderType {
        string first_name;
        @key string last_name;
    }
    sender;
    @key long        sender_id;
    long            time_sent;
    sequence<string> old_msgs;
    string          msg;
};
```

2. Annotation after the field

The following example shows how to define keys using the IDL 4 annotation *after* the key field, using IDL.

Example IDL file

```
//=====
// CoreDX DDS IDL example - using keys
//=====
```

```

// the 'sender_id' field is defined as a key
struct StringMsg1
{
    struct SenderType {
        string first_name;
        string last_name;
    }
    long                sender;
    long                sender_id; //@key
    long                time_sent;
    sequence<string>    old_msgs;
    string              msg;
};

// the 'sender' field is defined as a key
struct StringMsg1
{
    struct SenderType {
        string first_name;
        string last_name;
    }
    long                sender; //@key
    long                sender_id;
    long                time_sent;
    sequence<string>    old_msgs;
    string              msg;
};

// the key is a combination of the sender's last name
// and the sender_id
struct StringMsg1
{
    struct SenderType {
        string first_name;
        string last_name; //@key
    }
    long                sender;
    long                sender_id; //@key
    long                time_sent;
    sequence<string>    old_msgs;
    string              msg;
};

```

3. Original CoreDX DDS syntax

Before the formalization of the IDL 4 specification, the definition of DDS keys was not specified in any of the DDS standards. New CoreDX DDS

versions accept both the IDL 4 key definition syntax, and the original CoreDX DDS key definition syntax.

This section describes the original CoreDX DDS syntax. To specify a field to be a key, add the string “__dds_key” in front of the field definition. The CoreDX DDS data type compiler defines a compiler flag: “DDS_IDL”, and it is common to use this to provide compatible IDL syntax in the IDL files. The following IDL provides several examples of using keys in the IDL.

Example IDL file

```
//=====
// CoreDX DDS IDL example - using keys
//=====

// This #ifdef block allows the following CoreDX DDS IDL to be
// compatible with all other IDL

#ifdef DDS_IDL
#define DDS_KEY __dds_key
#else
#define DDS_KEY
#endif

// the 'sender_id' field is defined as a key
struct StringMsg1
{
    struct SenderType {
        string first_name;
        string last_name;
    } sender;
    DDS_KEY long sender_id;
    long time_sent;
    sequence<string> old_msgs;
    string msg;
};

// the 'sender' field is defined as a key
struct StringMsg1
{
    DDS_KEY struct SenderType {
        string first_name;
        string last_name;
    } sender;
    long sender_id;
    long time_sent;
    sequence<string> old_msgs;
};
```

```

    string                msg;
};

// the key is a combination of the sender's last name
// and the sender_id
struct StringMsg1
{
    struct SenderType {
        string first_name;
        DDS_KEY string last_name;
    } sender;
    DDS_KEY long          sender_id;
    long               time_sent;
    sequence<string>    old_msgs;
    string              msg;
};

```

Figure 11-2: IDL keys example

More information on data type configuration, keys, and annotation syntax is available in the CoreDX DDS Type System Programmer's Guide.

11.11 Using the CoreDX DDS Data Type Compiler

The CoreDX DDS data type compiler (coredx_ddl or coredx_ddl.exe) compiles the DDS types defined in IDL or XML and generates type-specific source code to be compiled and linked with a CoreDX DDS application. The CoreDX DDS data type compiler provides some command line arguments to tailor the behavior.

Table 11-4: coredx_ddl command line options

coredx_ddl option	Default	Description
-h	n/a	Help: print coredx_ddl usage information

coredx_ddl option	Default	Description
-a <i>alignment flag</i>	0	<p>0 == don’t count CDR ENCAP HDR in alignment. This is the default in CoreDX DDS versions 3.5.3 and newer, and is interoperable with other DDS implementations.</p> <p>1 == count CDR_ENCAP_HDR is alignment. This value is interoperable with older CoreDX DDS versions (before v3.5.3).</p>
-b <i>name</i>	IDL or XML filename	<p>Basename: provide an alternate filename prefix to use for the generated files. The default is the basename of the IDL or XML filename. For example, a IDL file named “hello.ddl” will, by default, produce generated files named like “hello.h”, “helloTypeSupport.h”, etc. The user can provide this argument to change the generated filename prefix to a string specified by <i>name</i>.</p> <p>-b option only valid for C, C++, and Cxx languages</p>
-d <i>directory name</i>	Current directory	Output Directory: provide an alternate directory to put the resulting generated code. The default is the current working directory.
-D <i>preprocessor symbol</i>	n/a	This option is used to specify pre-processor defines. Predefined by CoreDX DDS: <i>DDS_IDL</i> , <i>COREDXX_DDS</i>
-e <i>endian</i>	HOST CPU architecture	<p>Endian: one of “b” or “l”, provide the byte order (big endian or little endian) to use when marshalling published data to transmit over the wire. The default is the endian of the HOST platform (the platform where the coredx_ddl compiler is running). This should be used when the TARGET platform has a different endianness than the HOST platform.</p> <p>If the endianness of the TARGET platform is not set correctly, DDS applications that perform a read or take operation will have undefined behavior and may segfault.</p>

coredx_ddl option	Default	Description
-E <i>def extensibility</i>	EXTENSIBLE	Specify the default Extensibility of structures, if the extensibility annotation is not provided. 'e' -> EXTENSIBLE 'f' -> FINAL 'm' -> MUTABLE
-f <i>filename</i>	No default, must be specified	File: provide the IDL or XML file to process
-F		Enable support for the full X-Types type system. Without this flag, types are fully backwards compatible, and errors are generated with types and annotations that are <i>not</i> backwards compatible.
-g		Construct preprocessor guard from full input filename (default is base input filename)
-G <i>guard_variable</i>		Specify preprocessor guard variable used in c/c++ headers (default is base input filename).

coredx_ddl option	Default	Description
-I include_flags		<p>Enable or disable generation of certain code components. If a flag is prefixed with '^', then disable generation of that component. Otherwise, enable generation of that component. Flags include:</p> <p>D: generate extra X-Types defined sample <-> Dynamic Data conversion API's (default is no)</p> <p>e: generate scoped enum types (impacts traditional C++ mapping only) (default is no)</p> <p>g: generate 'get_field()' routine for struct/union types (default is enabled)</p> <p>k: strict keyword case checking (default is no)</p> <p>K: keyword-identifier collision is an error (only effective if strict key word case check is on) (default is off)</p> <p>O: generate TypeObject in TypeSupport (default is enabled)</p> <p>p: generate 'print' routine for struct types (default is disabled)</p> <p>T: generate TypeCode in TypeSupport source (default is enabled). Note '-i ^T' is equivalent to '-T'</p> <p>x: generate extra Type typedefs API's (default is enabled)</p> <p>X: generate extra X-Types defined TypeSupport API's (default is disabled)</p>
-I path	empty	This option provides a path that will be searched to satisfy '#include' directives found in the IDL file(s).
-j java_version	7	Controls some aspects of Java code generation (5 or 7).

coredx_ddl option	Default	Description
-l <i>language</i>	c	Language: one of "c", "cpp", "csharp", "java" provide the language to use for the generated source code.
-L license_file		Specify the full path to the CoreDX DDS license file.
-P preprocessor	cpp (linux) cl.exe /E (win)	Preprocessor: provide the preprocessor to use. The location of the specified preprocessor must be in the PATH.
-s		This option suppresses the generation of code for 'included' IDL files.
-S		This option strips the path from generated #include statements. It is only relevant if '-s' is in effect.
-T		This option suppresses the generation of 'typecode' data inline in TypeSupport.
-X		Specify that the input is in XML format (as opposed to IDL or DDL).
-Wno-<warning>		Disable the warning <warning>. For example: -Wno-1091

The CoreDX DDS data type compiler generates several source code files. All generated files are written to the current working directory (the directory from which coredx_ddl was run), unless the `-d <output_directory>` option is used.

For additional CoreDX data type compiler options for use with Dynamic Types, refer to the [section](#) of this guide.

Chapter 12 Configuring Reader Specific Locators

12.1 Overview

Some network environments can benefit from having a subset of traffic isolated to a specific multicast address (different from other traffic).

DataReaders and DataWriters in CoreDX DDS inherit the 'default' multicast and unicast locators from the DomainParticipant. If no specific locators are specified, then the entities will use the defaults. If a specific multicast locator is specified, it will be used in place of any default multicast locators. If a specific unicast locator is specified, it will be used in place of any default unicast locators.

The default UDP IPv4 multicast address used for user data is 239.255.0.1.

By configuring a DataReader to accept data on a specific address (Locator), then all DataWriters matched with that DataReader will send data to that specific address.

12.2 Configuration

The specific locators are configured in the **DataReaderQos.rtps_reader.locators** QoS policy. This field is a sequence of CoreDX_Locators. A CoreDX_Locators has the following structure:

```
typedef struct CoreDX_Locator_t {  
    int          kind;  
    uint32_t     port;  
    unsigned char addr[COREDX_LOCATOR_ADDR_LEN];  
} CoreDX_Locator;
```

To define a specific multicast locator in 'C' for the UDP transport, use the following:

```
CoreDX_Locator l;  
DDS_DataReaderQos dr_qos;  
  
DDS_Subscriber_get_default_datareader_qos(sub,  
&dr_qos);  
  
memset(&l, 0, sizeof(l));
```

```

l.kind = COREDX_UDPV4_LOCATOR_KIND_QOS;
l.port = 7400;
l.addr[12] = 239;
l.addr[13] = 255;
l.addr[14] = 0;
l.addr[15] = 2;
seq_add( &dr_qos.rtps_reader.locators, &l );

```

To define a specific multicast locator in 'C++' for the UDP transport, use the following:

```

sub->get_default_datareader_qos( dr_qos );

DDS::Locator l;
l.kind = COREDX_UDPV4_LOCATOR_KIND_QOS;
l.port = 7400;
memset( l.addr, 0, COREDX_LOCATOR_ADDR_LEN );
l.addr[12] = 239;
l.addr[13] = 255;
l.addr[14] = 0;
l.addr[15] = 2;
dr_qos.rtps_reader.locators.push_back( l );

```

CoreDX DDS supports setting specific UDP MULTICAST locators on the DataReader. Other combinations (other transports, or unicast locators) are not supported at this time.

Dynamic Types

12.3 Generated Code

The CoreDX DDS data type compiler will generate several source files, listed in Table 12-1. Detailed descriptions of the generated files are listed below. [If RPC interfaces are defined in the data type file, additional source files are generated. Those are described in the CoreDX DDS RPC Programmer's Guide.]

Table 12-1: Generated source code file names

C filenames	C++ filenames	C# filenames	Java filenames
-------------	---------------	--------------	----------------

name.h	name.hh	name.cs	name.java
name.c	name.cc		
nameTypeSupport.h	nameTypeSupport.hh	nameTypeSupport.cs	nameTypeSupport.java
nameTypeSupport.c	nameTypeSupport.cc		
nameDataReader.h	nameDataReader.hh	nameDataReader.cs	nameDataReader.java
nameDataReader.c	nameDataReader.cc		
nameDataWriter.h	nameDataWriter.hh	nameDataWriter.cs	nameDataWriter.java
nameDataWriter.c	nameDataWriter.cc		

12.3.1 Type Definition

The type definition files (name.h, name.c in C) contain the language specific data type declarations for the DDS data types defined in the IDL or XML file. They also include basic initialization, copy, and delete operations for the data types.

12.3.2 Typed TypeSupport Definition

The type support files (nameTypeSupport.h, nameTypeSupport.c in C) contain the language and type specific TypeSupport structures. For languages that support heredity, these are classes that derive from the base TypeSupport class.

12.3.3 Typed DataReader and DataWriter Definitions

The data reader (nameDataReader.h, nameDataReader.c) and data writer files (nameDataWriter.h, nameDataWriter.c) contain the language and type specific DataReader and DataWriter structures. For languages that support heredity, these are classes that derive from the base DataReader and DataWriter classes.

The application should use the type specific operations for DataReader and DataWriter calls.

Chapter 13 Quality of Service Features

One of the powerful features of DDS is the support for various **Quality of Service** (QoS) settings. QoS settings allow the application developer to tailor the behavior of publishers, subscribers, and the communications between them.

Most DDS Entities, from a DomainParticipantFactory down to the DataReader and DataWriter, have a set of QoS settings that apply. The QoS settings are contained in a structure. For example, a DomainParticipantFactory has a DomainParticipantFactoryQos structure containing all the applicable QoS settings.

The QoS settings for an entity can be established when the entity is created, or by using the `get_qos()` and `set_qos()` methods on each entity, as the following C example illustrates.

Sample C Application Code

```
DDS_Subscriber      subscriber;
DDS_DataReader      dr;
DDS_DataReaderQos   drqos;
DDS_DataReaderListener dr_listener;

/* ... code deleted ... */

/* Setup a non-default DataReader QoS structure */

DDS_Subscriber_get_default_datareader_qos(subscriber, &drqos);
drqos.history.kind = DDS_KEEP_LAST_HISTORY_QOS;
drqos.history.depth = 5;

/* EXAMPLE 1: Define the DataReader QoS at creation */

dr = DDS_Subscriber_create_datareader(subscriber,
                                     topic_descr,
                                     &drqos, &dr_listener,
                                     DDS_DATA_AVAILABLE_STATUS);

/* EXAMPLE 2: Set the QoS after creating the DataReader */
DDS_DataReader_set_qos(dr, &drqos);
```

Figure 13-1: Configuring QoS - Sample C Code

13.1 QoS Compatibility Between Publishing and Subscribing Entities

Som QoS policies apply only to the entity on which they are configured. Some QoS policies are shared with peer DDS applications through DDS discovery. And some QoS policies are used to determine matches between publishing and subscribing entities. In order to have effective communications, these QoS settings used to determine matches must be compatible between publishing entities and subscribing entities. Publishers and DataWriters **offer** a QoS configuration. Subscribers and DataReaders **request** a QoS configuration. If the Publisher and DataWriter offer a configuration setting that is **at least** what the Subscriber and DataReader requested, this is considered a match, even if the QoS configurations are not the same.

For example, the DURABILITY QoS setting indicates weather a publishing application will save previously published data and whether a subscribing application expects to receive data that was published before it was created. Consider a subscribing application requesting a DURABILITY QoS to be able to receive history (data published before the DataReader existed), and a publishing application offering a DURABILITY QoS indicating it will not save any data after it has been published. It is impossible for this publishing application to meet the request of this subscribing application, and effective communication will not occur. However, if the subscribing application requests a DURABILITY QoS to not receive any history, and a publishing application offers a DURABILITY QoS to make history available, the QoS settings are compatible and the DataWriter and DataReader will match. In this case, the publishing application is offering more than the subscribing application is requesting.

When attempting to match publishing applications with subscribing applications, the CoreDX DDS middleware will consider the QoS settings on both sides. If all QoS settings are compatible, the publishing application and subscribing application will be “matched”. If any QoS settings are incompatible both the publishing and subscribing applications are notified. The OfferedIncompatibleQos status is updated on the publishing application and the RequestedIncompatibleQos status is updated on the subscribing application. For information on how to retrieve communication statuses, see the *Communication Status* chapter.

Table 13-1 lists the compatibility of each QoS setting (whether or not the QoS setting must be compatible between publishing entities, subscribing entities, and topics).

13.2 QoS Mutability

Many QoS settings must be configured before the DDS Entity is *enabled*. However, there are some QoS settings that can be changed dynamically at any time. These QoS settings are considered *mutable* or changeable. Table 13-1 lists the mutability of each QoS setting (whether or not the QoS setting can be dynamically changed at any time). Attempting to change a QoS setting that is not *mutable* using a `set_qos()` operation will return `DDS::RETCODE_IMMUTABLE_POLICY`.

13.3 Quality of Service Details

Table 13-1 lists all the standard DDS QoS policies, along with compatibility and mutability characteristics. Following the table is a detailed list of all supported QoS features from the OMG DDS specifications and description of their use. CoreDX DDS supports additional policies beyond those defined in the OMG DDS specifications. QoS policies available with CoreDX DDS are listed in the 2 tables below.

Table 13-1: Standard QoS Policies Summary

QoS Setting	Applicable Entities	Must be compatible	Dynamically Changeable
DEADLINE	DataWriter, DataReader	YES	YES
DESTINATION_ORDER	DataWriter, DataReader	YES	no
DURABILITY	DataWriter, DataReader	YES	no
DURABILITY_SERVICE	Not Supported	no	no
ENTITY_FACTORY	DomainParticipantFactory, DomainParticipant, Publisher, Subscriber	no	YES

QoS Setting	Applicable Entities	Must be compatible	Dynamically Changeable
GROUP_DATA	Publisher, Subscriber	no	YES
HISTORY	DataWriter, DataReader	no	no
LATENCY_BUDGET	DataWriter, DataReader	YES	YES
LIFESPAN	DataWriter	(n/a)	YES
LIVELINESS	DataWriter, DataReader	YES	no
OWNERSHIP	DataWriter, DataReader	YES	no
OWNERSHIP_STRENGTH	DataWriter	(n/a)	YES
PARTITION	Publisher, Subscriber	no	YES
PRESENTATION	Publisher, Subscriber	YES	no
PROPERTY	DomainParticipant	no	Depends on use
READER_DATA_LIFECYCLE	DataReader	(n/a)	YES
RELIABILITY	DataWriter, DataReader	YES	no
RESOURCE_LIMITS	DataWriter, DataReader	no	no
TIME_BASED_FILTER	DataReader	(n/a)	YES
TOPIC_DATA	Topic	no	YES
TRANSPORT_PRIORITY	DataWriter (Not Supported)	(n/a)	YES
USER_DATA	DataWriter, DataReader	no	YES
WRITER_DATA_LIFECYCLE	DataWriter	(n/a)	YES

Table 13-2: CoreDX DDS QoS Policies

QoS Setting	Applicable Entites	Must be compatible	Dynamically Changeable
DISCOVERY	DomainParticipant	no	no
ENTITY_NAME	(all)	no	no
LOGGING	(all)	no	YES
PEER_PARTICIPANT	DomainParticipant	no	no
RTPS_READER	DataReader	no	no
RTPS_WRITER	DataWriter	no	no
THREAD_MODEL	DomainParticipant	no	no

13.3.1 DEADLINE

The **Deadline** QoS policy is used when a Topic is expected to have each instance updated periodically. The DataWriter offers a *Deadline* contract where the application guarantees to update each instance every *n* timeperiod. The DataReader requests a *Deadline* contract where the application expects each instance to be updated every *n* time period.

When a writing application does not satisfy the DataWriter's Deadline period (configured in its QoS policy), the *offered_deadline_missed* status is updated. The writing application may choose to be notified of this event through any of the offered notification methods (refer to *Communication Status* for more information).

When a writing application does not satisfy a matched DataReader's Deadline period (configured in its QoS policy), the *requested_deadline_missed* status is updated. The reading application may choose to be notified of this event through any of the available notification methods (refer to *Communication Status* for more information).

This is a QoS policy that must be compatible before DataReaders and DataWriters will match. The DataWriter must have a *Deadline* ≤ the

DataReader’s *Deadline* before the DataWriter and DataReader will communicate. If the *Deadlines* are not compatible, CoreDX DDS will generate an *IncompatibleQos* status (see the *Communication Status* chapter for additional information).

13.3.2 DESTINATION_ORDER

The **Destination Order** policy determines the logical order at reception time of data samples for an instance. This is important when the infrastructure must determine which samples to keep at the DataReader, based on other QoS policies like *HISTORY* and *RESOURCE_LIMITS*. The possible values for **Destination Order** are *by reception timestamp* and *by source timestamp*. When set to *by reception timestamp*, CoreDX DDS will use the reception time to determine the order of samples. When set to *by source timestamp*, CoreDX DDS will use the timestamp set by the publisher to determine the order of samples, regardless of when the data sample was received.

13.3.3 DISCOVERY (CoreDX DDS Extension)

The **Discovery** policy allows the tailoring of DomainParticipant discovery aspects, including configuration of the built-in DataReaders and DataWriters. Details on the various configuration items contained in the Discovery QoS policy are described in CoreDX DDS Discovery chapter.

13.3.4 DURABILITY

The **Durability** policy controls whether or not CoreDX DDS will make already published data available to late joining DataReaders. The publish-subscribe paradigm offered by CoreDX DDS allows applications to write data even when there are no current readers on the network. Further, a DataReader has the option to receive historical data (data published before this DataReader came online) in addition to currently published data. The **Durability** policy allows this configuration.

The possible values for **Durability** are *Volatile*, *Transient Local*, *Transient*, and *Persistent*. When set to *Volatile*, CoreDX DDS will not save previously published data for late joining readers. When set to *Transient Local*, CoreDX DDS will save previously published data for the lifespan of the DataWriter for late joining readers. With a *Transient Local* setting, once the DataWriter is destroyed, its published history data is no longer available. When set to *Transient*, CoreDX DDS will save previously published data for the lifespan of the publishing application for late joining readers. With a *Transient* setting, once the publishing application exits, its published history data is no longer

available. When set to *Persistent*, CoreDX DDS will save previously published data in permanent storage, where it can outlive the publishing application and a system reset. This history data is available to late joining readers.

The 'wait_for_historical_data' API may be used on DataReaders with a **Durability** setting of *Transient Local* or stronger to block the application until all historical data has been received by the DataReader. Refer to the *Wait for Historical Data* chapter for additional information.

This is a QoS policy that must be compatible before DataReaders and DataWriters will match. The DataWriter must have a **Durability** ≤ the DataReader's **Durability** before the DataWriter and DataReader will communicate, where the **Durability** values are ordered so that *Volatile* < *Transient Local* < *Transient* < *Persistent*. If the **Durabilities** are not compatible, CoreDX DDS will generate an *IncompatibleQos* status (see the *Communication Status* chapter for additional information).

The number of data samples and instances stored for any **Durability** setting is determined in part by the configuration of the **History** and **Resources Limits** QoS policies.

CoreDX DDS currently supports only the *volatile* and *transient local* values for **Durability**.

13.3.5 DURABILITY_SERVICE

The **Durability Service** QoS policy is applicable only when the **Durability** QoS policy is set to *Transient* or *Persistent*. This policy configures the duration and amount of data stored.

CoreDX DDS currently does not support the **Durability Service** QoS policy.

13.3.6 ENTITY_FACTORY

The **Entity Factory** QoS policy controls the behavior of `create_xxx()` and `delete_xxx()` operations on a factory entity. Factory entities include: *DomainParticipantFactory*, *DomainPartition*, *Publisher*, and *Subscriber*.

The **Entity Factory** QoS policy contains one configuration item: `autoenable_created_entities`. When set to *TRUE*, all *Entities* returned by `create_xxx()` operations are already enabled. When set to *FALSE*, the application must explicitly call `enable()` on all created *Entities*.

The `autoenabled_created_entities=FALSE` setting is commonly used when configuring one or more Transports for a DomainParticipant. Since the transport must be configured and installed *before* the DomainParticipant is enabled, it is necessary to disable the automatic enable of the DomainParticipant.

The default setting for the **Entity Factory** QoS policy is *TRUE*.

13.3.7 ENTITY_NAME (CoreDX DDS Extension)

The **Entity Name** QoS policy allows the application add a string description to each CoreDX DDS Entity. This information is shared during discovery, and accessible through the applicable built-in Topics.

13.3.8 GROUP_DATA

The **Group Data** QoS policy allows the application to attach additional information to created *Entity* objects. This data is not used by CoreDX DDS, and is made available to the application by the **Built-in Topics**, along with other discovery information. For more information, see the 9.2 *Built-In Topics* section.

13.3.9 HISTORY

The **History** QoS policy (along with the **Resource Limits** QoS policy) controls the size and behavior of the DataReader and DataWriter data caches. The data caches may be used to buffer written data on the DataWriter, and received data on the DataReader. The **History** QoS policy determines the behavior of the data caches when they become full. The **History** QoS policy can provide some amount of buffering on both the publishing and subscribing sides. When combined with the **Durability** QoS policy on a DataWriter, this QoS policy may determine the amount of data history saved for late joining readers. On a DataReader, this policy may determine the number of samples available to return on a `read()` or `take()` operation.

The possible values for the **History** kind are *KEEP ALL* and *KEEP LAST*. When the **History** kind is configured to *KEEP_LAST*, CoreDX DDS may delete older data samples if necessary to make room for newly written (or received) samples. When the **History** kind is configured to *KEEP_LAST*, CoreDX DDS will never ‘bump’ a stored data sample to make room for a newly written (or received) sample. Rather, the `write()` call will block or return an error (depending on how the DataWriter is configured).

When set to *KEEP_LAST*, the application can define a *depth* (number of samples per Instance to keep). [Depth is not used by CoreDX DDS when the **History**.kind is set to *KEEP_ALL*.]

A **History** kind of *KEEP_ALL* may be used in combination with the **Resource Limits** QoS policy in order to bound the number of samples and/or instances stored by CoreDX DDS.

Here is an example of how the **History** kinds can behave. Let us consider a scenario where a Reliable DataWriter is writing samples faster than at least one if it's matched Reliable DataReaders can read or process them. In this case, CoreDX DDS will attempt to buffer the unprocessed samples. Samples will be buffered first at the DataReader. With a **History** kind of *KEEP_LAST*, samples are buffered up to the **History** depth, and then they may be overwritten. With a **History** kind of *KEEP_ALL*, samples are buffered up to the configured **Resource Limits**. If **Resource Limits** are configured to be infinite, samples will be buffered infinitely. If **Resource Limits** are configured to be finite value(s) and the DataReader has buffered that configured number of samples, the DataReader will not process (and not acknowledge) any more received samples. At this point, samples will be buffered at the DataWriter. If the DataWriter is configured with **History** kind of *KEEP_ALL* and finite **Resource Limits**, the write() call will return an error once the DataWriter cache is full.

Note that using a **History** set to *KEEP_ALL* in combination with a **Durability** set to *TRANSIENT_LOCAL* (or higher) can be a dangerous combination. The CoreDX DDS infrastructure may keep every sample ever written by the DataWriter, until the publishing application specifically *disposes* or *unregisters* the Instance, or until **Lifespan** limits are expired (see the *Instances and Samples* chapter). This can quickly utilize all available resources for the publishing application, or the host machine if **Resource Limits** are not specified.

13.3.10 LATENCY_BUDGET

The **Latency Budget** QoS policy specifies a delay is acceptable in the time between when a publishing application writes data and when a subscribing application is notified the data is available. CoreDX DDS uses this policy as a hint – not a contract that must be monitored or enforced. By default, the **Latency Budget** on DataWriters is set to a small (non-zero) value, and on DataReaders is set to zero (0), indicating the delay should be minimized.

It may not be obvious why an application would want to configure a **LatencyBudget** greater than zero. Here are two examples of when it may be appropriate to configure a **Latency Budget**. First, consider a publishing application that is publishing a very high rate of data samples. If the **Latency Budget** is set to zero, CoreDX DDS will attempt to write every data sample onto the network as soon as it is available from the application, which may not be very efficient. In this case, setting a **Latency Budget** greater than zero allows CoreDX DDS to queue multiple data samples to write in a batch, which will reduce the amount of overhead required, and may improve performance. For another example, consider a subscribing application that is receiving a very high rate of data samples. If the **Latency Budget** is set to zero, CoreDX DDS will notify the application for every data sample that arrives at the DataReader. However, if the **Latency Budget** is set to a value greater than zero, CoreDX DDS can queue received data samples, and send 1 notification for multiple available samples. In this case, the subscribing application is issuing fewer calls to read() or take() but receiving all the same data samples.

13.3.11 LIFESPAN

The **Lifespan** QoS policy allows CoreDX DDS to expire old data samples. The application configures an expiration duration time on a DataWriter.

A DataReader receiving data from this DataWriter will periodically check all the samples that have been received, and if any samples have *expired*, they will be removed from the DataReader cache.

A DataWriter will also periodically check all the samples in its DataWriter Cache and may remove any samples that have expired (actual removal may be delayed due to Reliability QoS policy settings).

By default, the expiration duration is 0 (meaning an infinite duration, or no expiration).

13.3.12 LIVELINESS

The **Liveliness** QoS policy controls the mechanism used to ensure *DataWriters* on the network remain “alive” to their matched DataReaders. The possible values for the **Liveliness** QoS policy are *Automatic*, *Manual by Participant*, and *Manual by Topic*.

The *Automatic* setting configures CoreDX DDS ensure all *DataWriters* within a *DomainParticipant* stay alive, without requiring any specific action from the publishing application.

The manual settings: *Manual by Participant* and *Manual by Topic* require the publishing application to periodically assert the liveliness to indicate the corresponding *Entity* is still alive. This can be explicit by calling the *assert_liveliness()* operation or implicit by writing data. The *Manual by Participant* configuration allows any one *DataWriter* to assert liveliness for all *DataWriters* within that *DomainParticipant*. The *Manual by Topic* configuration requires each *DataWriter* to assert its own liveliness.

The **Liveliness** QoS policy includes a *leaseduration*. For a *DataWriter*, the *lease duration* is an offered contract that the writer will assert liveliness at least once every specified duration. For a *DataReader*, the *lease duration* is a request that the writer assert liveliness at least once every specified duration interval.

When a writing application does not satisfy the *DataWriter's Liveliness* lease duration, the *liveliness_lost* status is updated. The writing application may choose to be notified of this event through any of the offered notification methods (refer to *Communication Status* for more information).

When a writing application does not satisfy a matched *DataReader's* liveliness period, the *liveliness_changed* status is updated. The reading application may choose to be notified of this event through any of the available notification methods (refer to *Communication Status* for more information).

This is a QoS policy that must be compatible before *DataReaders* and *DataWriters* will match. When configured to one of the Manual kinds, the *DataWriter* must have a **Liveliness** *lease duration* \geq the *DataReader's Liveliness* *lease duration* before the *DataWriter* and *DataReader* will communicate. If the **Liveliness** is not compatible, CoreDX DDS will generate an *IncompatibleQos* status (see the *Communication Status* chapter for additional information).

13.3.13 LOGGING (CoreDX DDS Extension)

The **Logging** QoS policy provides fine grained control over the levels and categories of log messages produced by CoreDX DDS. Setup and configuration of CoreDX DDS logging is described in the CoreDX DDS Logging section.

13.3.14 OWNERSHIP

The **Ownership** QoS policy controls whether CoreDX DDS will allow multiple *DataWriters* to update the same instance. The possible values for **Ownership** are *Shared* and *Exclusive*. When set to *Shared*, CoreDX DDS does not enforce unique ownership for each instance, and multiple *DataWriters* can update the same instance (*DataReaders* will receive the data written by all matched *DataWriters*). When set to *Exclusive*, each instance can be modified only by one *DataWriter*. In this case, one *DataWriter* “owns” each instance, and while that *DataWriter* is “alive”, matched *DataReaders* will only accept samples on an instance written by the instance owner.

A *DataReader* may automatically change the owner of an instance to a different *DataWriter*. This will happen if the current owner misses a deadline or is otherwise considered to be not actively writing on the instance. The *DataReader* will then assign ownership to the active *DataWriter* with the next highest *strength*.

This is a QoS policy that must be compatible before *DataReaders* and *DataWriters* will match. The *DataReader* and *DataWriter* **Ownerships** must match before the *DataWriter* and *DataReader* will communicate. If the **Ownership** is not compatible, CoreDX DDS will generate an *IncompatibleQos* status (see the *Communication Status* chapter for additional information).

13.3.15 OWNERSHIP_STRENGTH

The **Ownership Strength** QoS policy is applicable only when the **Ownership** QoS policy is set to *Exclusive*. Each *DataWriter* can set its *Strength* with this QoS setting. This strength is used to determine which *DataWriter*’s updates will be received used the subscribing application when more than one *DataWriter* is writing on that instance.

13.3.16 PARTITION

The **Partition** QoS policy allows the application to define logical partitions in a DDS domain. In order for a *DataReader* to see data published by a *DataWriter*, their **Partitions** must match. A **Partition** is a string that may contain a ‘*’ wildcard. Entities may define (and be part of) multiple partitions. The empty string (“”) is a valid partition, and will match only another empty string or ‘*’ wildcard partition.

This is a QoS policy that must be compatible before *DataReaders* and *DataWriters* will match. If the **Partitions** do not match, CoreDX DDS will

generate an `IncompatibleQos` status (see the *Communication Status* chapter for additional information).

13.3.17 **PEER_PARTICIPANT (CoreDX DDS Extension)**

The **PeerParticipant** QoS policy contains a list of peer DomainParticipants. This configuration can be used as an alternative or in addition to the standard automatic and dynamic discovery. More information on configuring the PeerParticipant QoS policy can be found in the Static Discovery section of this guide.

13.3.18 **PRESENTATION**

The **Presentation** QoS policy controls the extent to which published data changes are dependent on each other. The possible values for **Presentation** are *coherent access* and *ordered access*. In addition, there is an additional configuration item: *Access Scope*.

When **Presentation** is configured as *coherent access*, CoreDX DDS will preserve groupings of changes made by the publishing application between calls to the *begin_coherent_change()* and *end_coherent_change()* operations. When *Access Scope* is set to `INSTANCE`, the scope for grouping changes is within each individual instance. In this case, calls to *begin_coherent_change()* and *end_coherent_change()* have no effect. When *Access Scope* is set to `TOPIC`, then coherent changes made by a DataWriter will be preserved and made available as a set to each DataReader. When *Access Scope* is set to `GROUP`, coherent changes made by all DataWriters attached to a Publisher will be preserved and made available to each Subscriber.

When **Presentation** is configured as *ordered access*, CoreDX DDS will preserve the order of changes made by the publishing application, in accordance with the setting of *Access Scope*. When *Access Scope* is set to `INSTANCE`, changes to an instance are unordered relative to other instances. When *Access Scope* is set to `TOPIC`, changes made by a single DataWriter are made available to DataReaders in the same order in which they occurred. When *Access Scope* is set to `GROUP`, changes made by all DataWriters attached to a Publisher are made available to Subscribers in the same order in which they occurred.

Note that while changes are preserved by CoreDX DDS and made available to the DataReaders or Subscribers in order, the application must make the

appropriate calls to the DataReader or Subscriber in order to see the data in the desired order.

This is a QoS policy that must be compatible before DataReaders and DataWriters will match.

13.3.19 **PROPERTY**

The **Property** QoS policy is part of the ParticipantQoSPolicy structure and allows the application to flexibly define name/value pairs. Currently, CoreDX DDS uses the **Property** QoS policy only for configuration of the CoreDX DDS Security Plug-ins.

The Property QoS policy contains a sequence of properties, each one containing a name, value, and propagate_flag.

The propagate_flag determines if the information is shared during participant discovery. Depending on the application use of the Property QoS policy, this information may or may not need to be shared to peers during DDS discovery. One reason *not* to propagate these properties is to keep the size of discovery data small (network and memory utilization). When CoreDX DDS Secure is used, the Property QoS policy contains security configuration specific to this DomainParticipant that has no meaning to peers and could potentially leak sensitive information if shared on the network.

13.3.20 **READER_DATA_LIFECYCLE**

The **Reader Data Lifecycle** QoS policy controls the behavior of the DataReader with regard to the data it has received and is maintaining. A DataReader internally maintains data samples it has received until they have been ‘taken’ by the application according to **History** and **Resource Limits** QoS policy settings. A DataReader will maintain information for an instance, even when the associated DataWriter is no longer alive, until the application has “taken” all samples for that instance.

The **Reader Data Lifecycle** QoS policy offers some memory usage protection to the application, by allowing CoreDX DDS to release resources for instances, even if the application neglects to “take” all samples for these instances. This QoS policy offers two configuration items. The *autopurge nowriter samples delay* configuration item defines the amount of time the DataReader will maintain information for an instance once it becomes NOT_ALIVE_NO_WRITERS (there are no live writers writing this instance).

The *autopurge disposed samples delay* configuration item defines the amount of time the DataReader will maintain information for an instance once it becomes NOT_ALIVE_DISPOSED (a writer has disposed this instance). CoreDX DDS may reclaim an instance on a DataReader when the instance state is NOT_ALIVE_NO_WRITERS and there are no samples associated with the instance.

13.3.21 RELIABILITY

The **Reliability** QoS policy configures the level of reliability CoreDX DDS will guarantee for communications between a DataReader and DataWriter. The possible configuration for **Reliability** are *Best Effort*, *Reliable*, and *ACK-less Reliable*.

With a *Best Effort* configuration, CoreDX DDS will operate in a “fire and forget” mode, making an effort to deliver all published data, with no guarantee all data will be received by all matched DataReaders. It is important to note that on reliable physical networks, with enough bandwidth to support published data, it is rare to have dropped or missed samples in this mode. DDS applications publishing periodic data with high data rates can see better performance with minimal to no data loss using a *Best Effort* configuration.

With a *Reliable* configuration, CoreDX DDS will use an additional reliability protocol to check if written samples are received, and possibly resend them if necessary. The *Reliable* configuration may be used in combination with the **History** and **Resource Limits** QoS policies to guarantee all published data will be received by all matched DataReaders. This configuration requires more resources and overhead to fulfill. The **History** and **Resource Limits** QoS policies will determine the amount of resources CoreDX DDS will maintain in order to meet *Reliable* requirements. If configured resource limits are met, the publishing application may block on a write() operation.

With an *ACK-less Reliable* configuration, CoreDX DDS will use a lighter-weight version of the Reliability protocol that uses negative acknowledgements, but no positive acknowledgments.

Additional information about using and configuring the Reliability QoS policy can be found in the *Configuring Reliability Protocol* chapter of this guide.

This is a QoS policy that must be compatible before DataReaders and DataWriters will match. A *Reliable* DataWriter will match any DataReader,

Best Effort, ACK-less Reliable, or Reliable. A *Best Effort* DataWriter will match only *Best Effort* DataReaders.

13.3.22 **RESOURCE_LIMITS**

The **Resource Limits** QoS policy configures the resources CoreDX DDS can use in order to meet the requirements imposed by the application and other QoS settings (including **History**, **Durability**, and **Reliability**). The **Resources Limits** that can be configured include: the total number of samples (*max_samples*), the total number of instances (*max_instances*), and the number of samples in each instance (*max_samples_per_instance*).

CoreDX DDS DataReaders and DataWriters will not store more samples or instance than is specified by their Resource Limits QoS policies. This provides a convenient mechanism to constrain the amount of memory a CoreDX DDS application will use for application data. It also allows CoreDX DDS to pre-allocate memory resources, which can result in better performance.

The Resource Limits QoS policy provides a hard limit for the number of samples or instances that can be stored by a DataReader or DataWriter. The configuration of other QoS Policies: Reliability, Durability, and History determine the behavior of DataReaders and DataWriters when their Resource Limits are met.

For additional information, refer to the Data Cache section, as well as the sections for these specific QoS policies: **RELIABILITY**, **DURABILITY**, and **HISTORY**.

It is possible for a DataWriter to publish data samples (and instances) faster than a DataReader is consuming them, causing the DataReader to fill up its Data Cache to its configured Resource Limits. This could be with respect to *samples* (*max_samples* or *max_samples_per_instance*) or *instances* (*max_instances*). With the right combination of QoS policies (specifically, Reliable Reliability and Keep All History), published samples will be stored by the DataWriter until all DataReaders can accept them. The DataWriter’s Cache of data samples will continue to grow until it meets the configured **Resources Limits**. When this occurs, CoreDX DDS will return an error on the next call to `write()` (or `dispose()` or `register_instance()`). The application may block first, depending on the configuration of the Reliability blocking time. This is a configuration where one ‘slow’ DataReader can effectively prevent the DataWriter from publishing any new data, affecting all other DataReaders matched to that DataWriter.

The CoreDX DDS constant `DDS::LENGTH_UNLIMITED` is used to indicate the absence of a particular limit.

The *max_samples* and *max_samples_per_instance* settings must be consistent such that *max_samples* \geq *max_samples_per_instance*. In other words, if *max_samples* is set (not `LENGTH_UNLIMITED`), then *max_samples_per_instance* must also be set (to a value \leq *max_samples*).

In addition, the *max_samples_per_instance* setting must be consistent with the **History depth**, such that *depth* \leq *max_samples_per_instance*. If there is an error in these settings, a `create_datareader()` or `create_datawriter()` operation will fail. A `set_qos()` operation will return the error `DDS::RETCODE_INCONSISTENT_POLICY`.

13.3.23 *RTPS READER (CoreDX DDS Extension)*

The **RTPS Reader** QoS policy allows the application to tailor various aspects of the DDS wire protocol for this particular DataReader. These control options include Reliability timing and behavior, data batching, and if typecode should be transmitted for discovery purposes.

Additional information about these control options can be found in the following sections:

- Reliability: Configuring Reliability Protocol
- Data Batching: Data Batching
- Typecode: CoreDX DDS *Discovery* and the CoreDX DDS Type System Programmer's Guide

13.3.24 *RTPS WRITER (CoreDX DDS Extension)*

The **RTPS Writer** QoS policy allows the application to tailor various aspects of the DDS wire protocol for this particular Data Writer. These control options include Reliability timing and behavior parameters, data batching, filtering behavior, transmit buffer size and behavior, and if typecode should be transmitted for discovery purposes.

Additional information about these control options can be found in the following sections:

- Reliability: Configuring Reliability Protocol
- Data Batching: Data Batching
- Content Filters: Content Filtered Topics
- Transmit Buffers: Transmit Buffers
- Typecode: CoreDX DDS *Discovery* and the CoreDX DDS Type System Programmer’s Guide

13.3.25 **THREAD_MODEL (CoreDX DDS Extension)**

The **Thread Model** QoS policy allows the application to control the number of threads created and used by CoreDX DDS. CoreDX DDS can run in single threaded or multithreaded modes. This is described further in the Threading Options section.

13.3.26 **TIME_BASED_FILTER**

The **Time Based Filter** QoS policy allows the application to indicate a particular DataReader does not necessarily want to see all data samples published for a Topic. In fact, the DataReader would to see, for each instance, at most one data sample every n time period. This time period is the *minimum_separation* for the **Time Based Filter**.

Using the **Time Based Filter** QoS policy can reduce the amount of data written by a DataWriter. This is particularly useful in situations where a DataReader cannot keep up with the amount of data published, or where some DataReaders simply do not need all the intermediate data samples published on a Topic.

The **Time Based Filter** *minimum_separation* must be consistent with the **Deadline period**. Setting a *period* < *minimum_separation* is an error. A `create_datareader()` operation will fail. A `set_qos()` operation will return the error `DDS::RETCODE_INCONSISTENT_POLICY`.

13.3.27 **TOPIC_DATA**

The **Topic Data** QoS policy allows the application to attach additional information to created *Entity* objects. This data is not used by CoreDX DDS, and this information is not shared during discovery. For more information, see the *Built-In Topics* section.

13.3.28 **TRANSPORT_PRIORITY**

CoreDX DDS does not currently support this QoS policy.

13.3.29 **USER_DATA**

The **User Data** QoS policy allows the application to attach additional information to created *Entity* objects. This data is not used by CoreDX DDS, and is made available to the application by the **Built-in Topics**, along with other discovery information. For more information, see the *Built-In Topics* section.

13.3.30 **WRITER_DATA_LIFECYCLE**

The **Writer Data Lifecycle** QoS policy controls the behavior of the DataWriter with regard to the data it has published and is maintaining. This QoS policy allows the application to configure CoreDX DDS to automatically dispose instances when they are unregistered (see the *Instance Lifecycles* for additional information).

The **Writer Data Lifecycle** QoS policy contains one configuration item: auto-dispose unregistered instances (*autodispose_unregistered_instances*). Setting this configuration item to *TRUE* causes the DataWriter to dispose the instance each time it is unregistered. Setting this configuration item to *FALSE* will not automatically dispose instances when they are unregistered.

When a DataWriter is deleted, all instances managed by the DataWriter are automatically unregistered. Therefore, only setting the auto-dispose unregistered instances configuration item will ensure the instances managed by a DataWriter are disposed.

Chapter 14 Communication Status

The DDS infrastructure keeps track of a number of statuses and statistics related to data communications. The application may choose to be made aware of some, all, or none of these statuses and statistics.

Each DDS entity has its relevant statuses, as listed in Table 14-1.

Table 14-1: Communication Statuses

Entity	Status Name	Description
Topic	INCONSISTENT_TOPIC	Another, different, Topic exists with the same name as this Topic, but a different data type (if type information is shared for discovery), or different data type name.
Subscriber	DATA_ON_READERS	New data has been received on one or more DataReaders associated with this Subscriber, and is available for the application to read() or take().
DataReader	SAMPLE_REJECTED	A received sample has been rejected because of RESOURCE_LIMITS QoS setting has been reached or exceeded. The sample has been ‘seen’ by the DDS middleware at the subscribing application, but could not be stored. The DataWriter will not re-transmit this sample. The sample is not available to the subscribing application.
	LIVELINESS_CHANGED	One or more DataWriters that were writing data this DataReader was reading has changed its liveliness (becoming active or inactive).

Entity	Status Name	Description
	REQUESTED_DEADLINE_MISSED	A data update for an instance was not received in the expected time interval (configured in the Deadline QoS Policy). Because the offered deadline (DataWriter) must be \leq the requested deadline (DataReader), it is possible for a DataWriter's OFFERED_DEADLINE_MISSED status may be triggered, when the DataReader's OFFERED_DEADLINE_MISSED status is NOT triggered.
	REQUESTED_INCOMPATIBLE_QOS	A DataWriter was discovered whose Topic matches this DataReader, but whose QoS is incompatible with this DataReader.
	DATA_AVAILABLE	New data is now available to be read.
	SAMPLE_LOST	A sample was lost (never received). This may be because the sample was dropped by the network (with BEST_EFFORT reliability) or because the DataWriter no longer has this sample to re-transmit (with RELIABLE reliability).
	SUBSCRIPTION_MATCHED	A DataWriter has been discovered that matches the Topic of this DataReader and has a compatible QoS (or a DataWriter that was previously matched is no longer matched).
DataWriter	LIVELINESS_LOST	The liveliness specified in the LIVELINESS QoS was not respected, and DataReaders will consider this DataWriter no longer active.

Entity	Status Name	Description
	OFFERED_DEADLINE_MISSED	A data update was not received in the expected time interval (configured in the Deadline QoS Policy). Because the offered deadline (DataWriter) must be \leq the requested deadline (DataReader), it is possible for a DataWriter’s OFFERED_DEADLINE_MISSED status may be triggered, when the DataReader’s OFFERED_DEADLINE_MISSED status is NOT triggered.
	OFFERED_INCOMPATIBLE_QOS	A DataReader was discovered for the same Topic as this DataWriter, but the QoS requested by that DataReader was incompatible with this DataWriter’s QoS.
	PUBLICATION_MATCHED	A DataReader has been found that matches the Topic and QoS of this DataWriter (or a DataReader that was previously matched is no longer matched).

Some communication statuses are associated with data being available for the subscribing application. These are referred to as **read communication statuses**, and include: DATA_ON_READERS and DATA_AVAILABLE. Since these two statuses indicate the reception of data (the real purpose for a Data Distribution Service) they treated a little differently from the other communication statuses, referred to as **plain communication statuses**.

14.1 Communication Status Details

Each communication status is described in detail below.

14.1.1 Inconsistent Topic Status

The Inconsistent Topic Status is used to inform the application that another Topic has been registered in the Domain (and Partition, if defined) that has the same name as this Topic, but a different data type (if data types are used for matching) or a different data type name. The CoreDX DDS middleware will allow an application to create multiple Topics of the same name and different types. The Inconsistent Topic Status allows applications to be made aware of these inconsistencies.

Type: Plain Communication Status
 Associated Entity: Topic
 Mask Name: INCONSISTENT_TOPIC_STATUS
 Struct Type Name: InconsistentTopicStatus




Figure 14-1: Inconsistent Topic Status Structure

14.1.2 Data On Readers Status

The Data On Readers Status is one of the read communication statuses, and is used to inform the application that one or more of the DataReaders attached to a Subscriber has new data samples or sample information.

The Data Available Status and Data On Readers Status are communicated to the application together. In other words, if there is Data Available for a DataReader, then there is a DataReader with new data.

Type: Read Communication Status
 Associated Entity: Subscriber
 Mask Name: DATA_ON_READERS_STATUS
 Struct Type Name: N/A

14.1.3 Sample Rejected Status

The Sample Rejected Status is used to inform the application that a data sample was not accepted by the DataReader because of resources limits and history configured in the associated QoS (on the effected DataReader). The HISTORY QoS allows the application to configure the behavior of DDS when the DataReader cache is full. With a configuration of KEEP_ALL, the middleware may not remove existing samples to make room for the new sample, and this has the potential to trigger a Sample Rejected Status. The RESOURCE_LIMITS QoS allows the application to set a limit on the total number of samples, the total number of instances, and the number of samples per each instance kept by the CoreDX DDS middleware. If a DataReader receives a sample that would put any of these numbers over their set limit, and History is configured to KEEP_ALL, the sample is rejected (not added to the reader cache), and the appropriate count on the Sample Reject Status is updated.

Because the sample was ‘seen’ by the DataReader, the sending DataWriter will not retransmit it (even when Reliability = RELIABLE). The application will not be able to access this sample.

Type:	Plain Communication Status
Associated Entity:	DataReader
Mask Name:	SAMPLE_REJECTED_STATUS
Struct Type Name:	SampleRejectedStatus

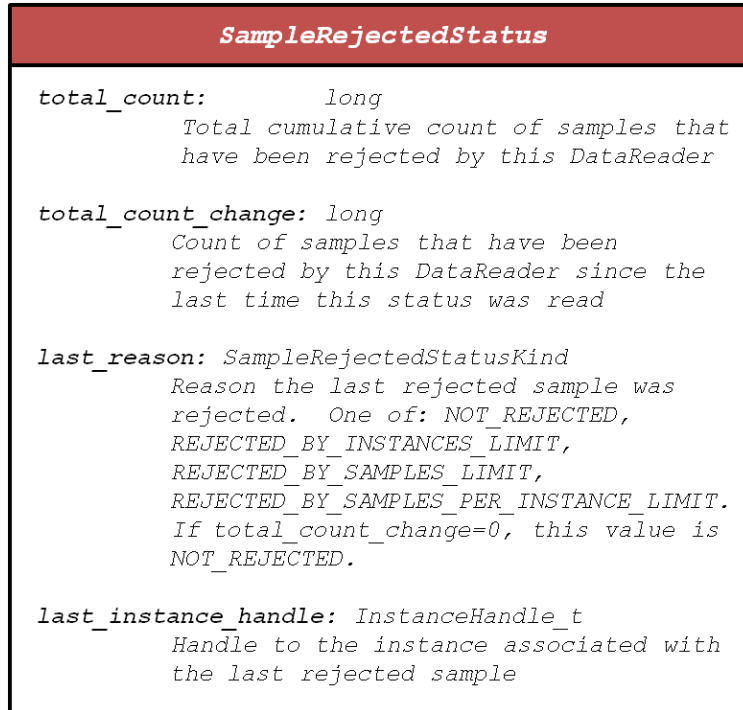


Figure 14-2: Sample Rejected Status Structure

14.1.4 Liveliness Changed Status

The Liveliness Changed Status is used to inform the application of changes to DataWriters known by this DataReader. DataReaders keep track of all DataWriters they are 'matched' with. These matched DataWriters may be ACTIVE (they are either actively writing data or otherwise asserting their liveliness) or INACTIVE (they are not actively writing or asserting their liveliness). Any time one of the DataWriters this DataReader is matched with changes between these states, the Liveliness Changed Status is updated.

Type:	Plain Communication Status
Associated Entity:	DataReader
Mask Name:	LIVELINESS_CHANGED_STATUS
Struct Type Name:	LivelinessChangedStatus

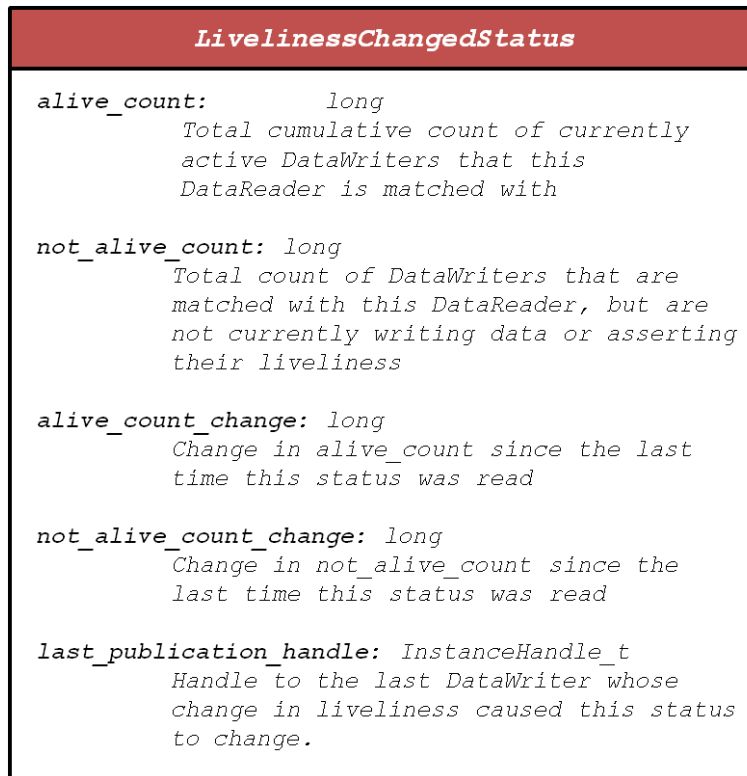


Figure 14-3: Liveliness Changed Status Structure

14.1.5 Requested Deadline Missed Status

The Requested Deadline Missed Status is used to inform the application when a deadline period specified in an associated Deadline QoS (either Topic or DataReader) was missed. The REQUESTED_DEADLINE_QoS allows the application to request that an instance be updated at least once every time interval specified by the QoS. When a DataWriter fails to update an instance within the time interval specified by the DataReader’s Deadline QoS policy, the Requested Deadline Missed Status is updated.

Type:	Plain Communication Status
Associated Entity:	DataReader
Mask Name:	REQUESTED_DEADLINE_MISSED_STATUS
Struct Type Name:	RequestedDeadlineMissedStatus

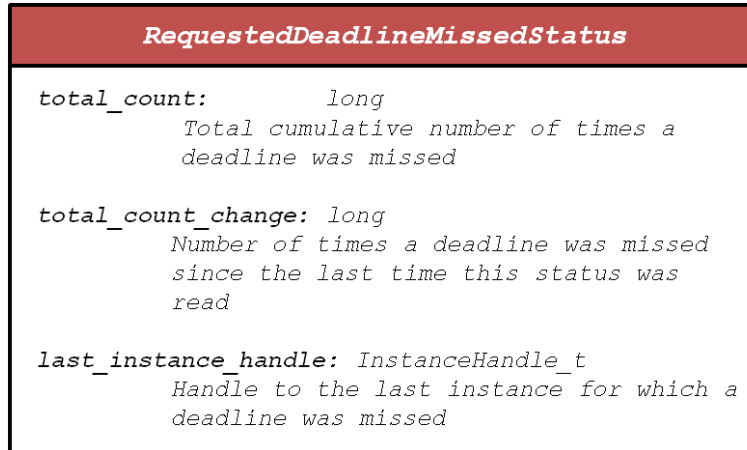


Figure 14-4: Requested Deadline Missed Status Structure

14.1.6 Requested Incompatible QoS Status

The Requested Incompatible QoS Status is used to inform the application that a DataWriter was discovered with a matching Topic and matching data type, but with QoS incompatible to this DataReader's requested QoS. For additional information about compatible and incompatible QoS, see the *Quality of Service Features* chapter.

Type:	Plain Communication Status
Associated Entity:	DataReader
Mask Name:	REQUESTED_INCOMPATIBLE_QOS_STATUS
Struct Type Name:	RequestedIncompatibleQosStatus

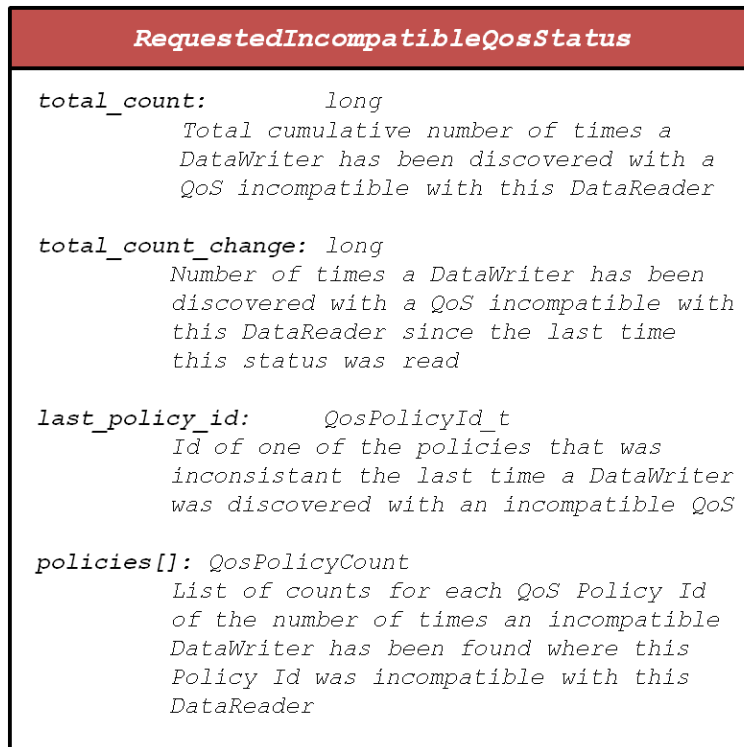


Figure 14-5: Requested Incompatible QoS Status Structure

14.1.7 Data Available Status

The Data Available Status is one of the read communication statuses (along with the Data On Readers Status), and is used to inform the application of new data available to be read on the associated DataReader.

These two read communication statuses are communicated to the application together. In other words, if there is Data Available for a DataReader, then there is a DataReader with new data.

Type:	Read Communication Status
Associated Entity:	DataReader
Mask Name:	DATA_AVAILABLE_STATUS
Struct Type Name:	N/A

14.1.8 Sample Lost Status

The Sample Lost Status is used to inform the application that one or more data samples were not received by a Reader. A sample can be 'lost' for many reasons.

For example, the transport might drop the sample due to network congestion or some other reason. If the Reader is using BEST_EFFORT reliability, then samples dropped by the underlying transport will not be retransmitted, and they are LOST.

Even if reliability is set to RELIABLE, it is still possible to experience lost samples due to other QoS settings. For example, if the Writer is configured to keep very little historical data in its cache (either through HISTORY or RESOURCE_LIMITS QoS), then it is possible that a Reader will fail to get a sample simply because the sample is purged from the Writer's cache before it could be successfully transmitted to the Reader.

Lost Samples are not available to the subscribing application.

Type:	Plain Communication Status
Associated Entity:	DataReader
Mask Name:	SAMPLE_LOST_STATUS
Struct Type Name:	SampleLostStatus

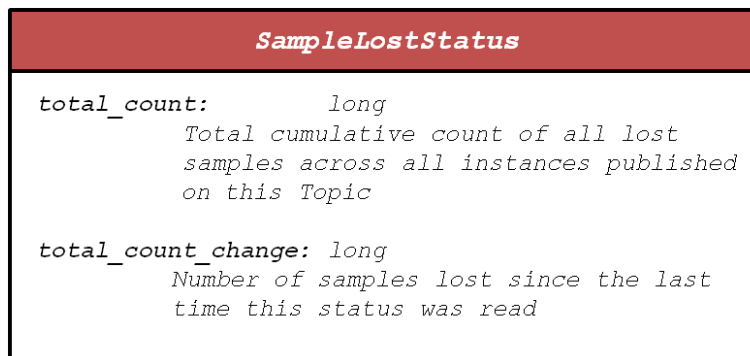


Figure 14-6: Sample Lost Status Structure

14.1.9 Subscription Matched Status

The Subscription Matched Status is used to inform the application that a new DataWriter has been discovered that matches this DataReader’s QoS settings and it is producing data this DataReader is interested in. That is, the DataWriter is writing data on the same Topic this DataReader is reading on and its QoS settings are compatible with this DataReader.

Type: Plain Communication Status
 Associated Entity: DataReader
 Mask Name: SUBSCRIPTION_MATCHED_STATUS
 Struct Type Name: SubscriptionMatchedStatus

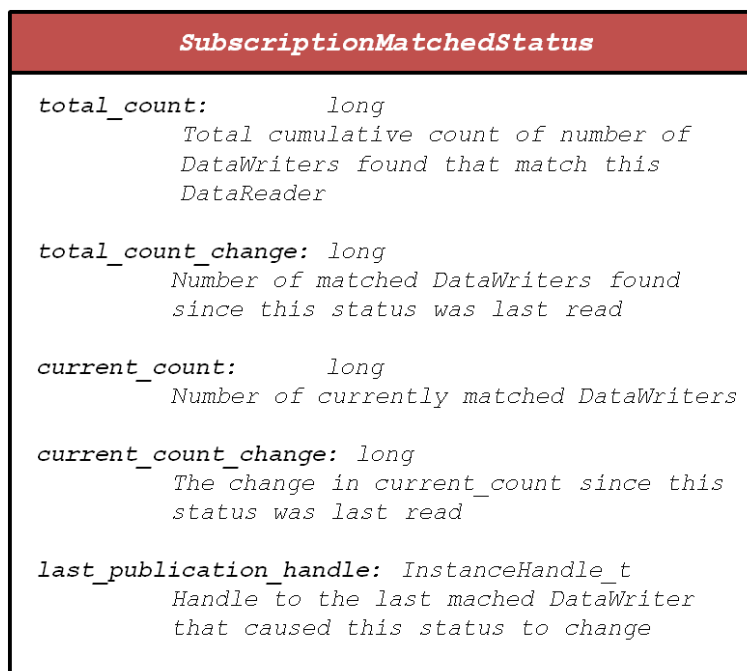


Figure 14-7: Subscription Matched Status Structure

14.1.10 Liveliness Lost Status

The Liveliness Lost Status is used to inform the application that this DataWriter has missed asserting its liveliness in the time period specified by its LIVELINESS QoS policy. The DataWriter’s LIVELINESS QoS policy allows the publishing application to specify the interval in which this DataWriter

will either write a sample or assert its liveliness to all matched DataReaders. If a DataWriter misses this specified window, the Liveliness Lost Status is updated.

Type: Plain Communication Status
 Associated Entity: DataWriter
 Mask Name: LIVELINESS_LOST_STATUS
 Struct Type Name: LivelinessLostStatus

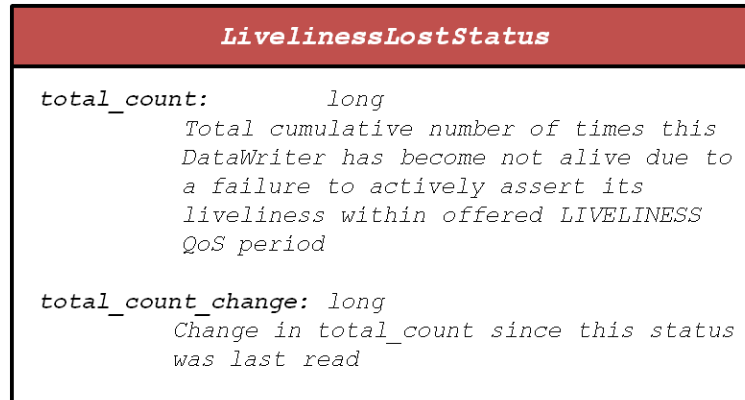


Figure 14-8: Liveliness Lost Status Structure

14.1.11 Offered Deadline Missed Status

The Offered Deadline Missed Status is used to inform the application when a deadline period specified in an associated Deadline QoS (either Topic or DataWriter) was missed. The OFFERED_DEADLINE QoS allows the application to commit to updating each instance at least once every time interval specified by the QoS setting. When the DataWriter fails to update an instance within the specified time interval, the Offered Deadline Missed Status is updated.

Type: Plain Communication Status
 Associated Entity: DataWriter
 Mask Name: OFFERED_DEADLINE_MISSED_STATUS
 Struct Type Name: OfferedDeadlineMissedStatus

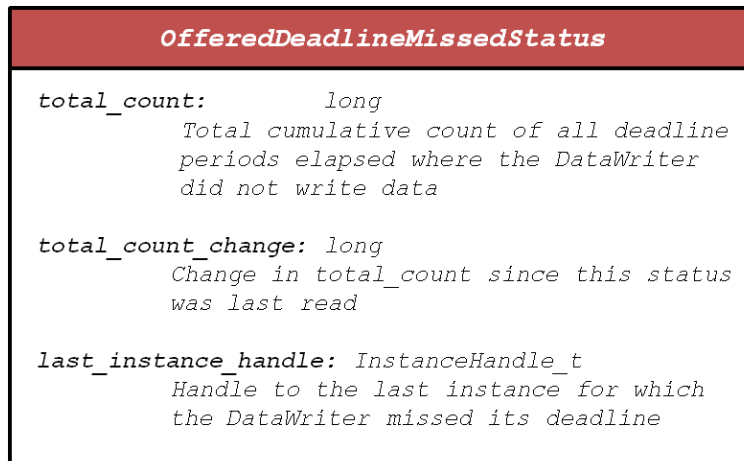


Figure 14-9: Offered Deadline Missed Status Structure

14.1.12 Offered Incompatible QoS Status

The Offered Incompatible QoS Status is used to inform the application that a DataReader was discovered with a matching Topic and matching data type, but with QoS incompatible to this DataWriter’s offered QoS. For additional information about compatible and incompatible QoS, see the *Quality of Service Features* chapter.

Type:	Plain Communication Status
Associated Entity:	DataWriter
Mask Name:	OFFERED_INCOMPATIBLE_QOS_STATUS
Struct Type Name:	OfferedIncompatibleQosStatus

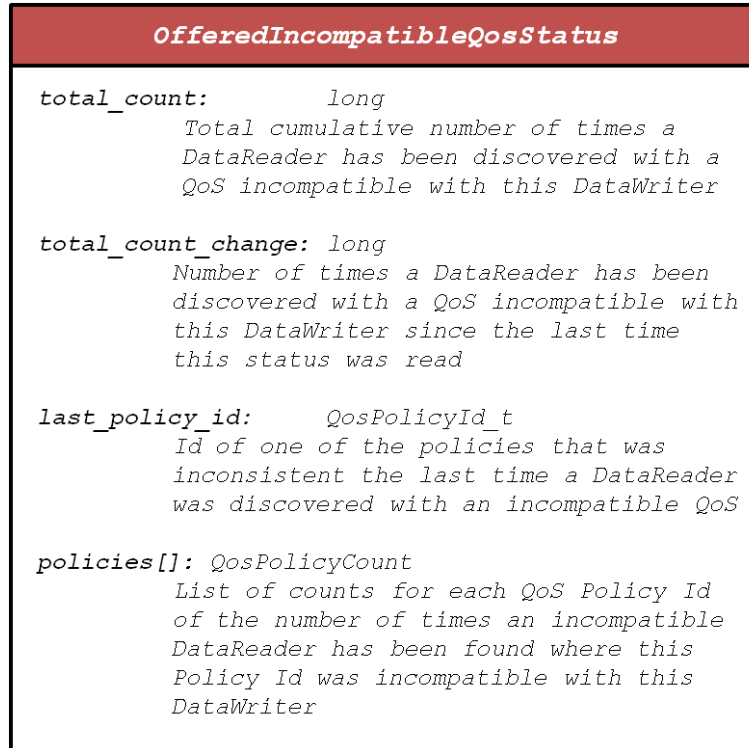


Figure 14-10: Offered Incompatible QoS Status Structure

14.1.13 *Publication Matched Status*

The Publication Matched Status is used to inform the application that a new DataReader has been discovered and matched with this DataWriter.

Type:	Plain Communication Status
Associated Entity:	DataWriter
Mask Name:	PUBLICATION_MATCHED_STATUS
Struct Type Name:	PublicationMatchedStatus

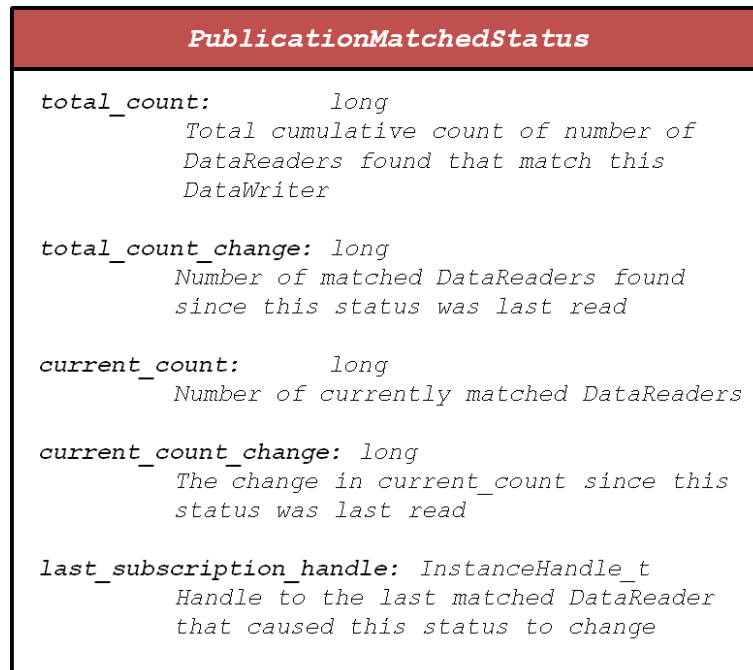


Figure 14-11: Publication Matched Status Structure

14.2 Application Access to Communication Status

An application can access **plain communication statuses** associated with an entity by calling that entity’s `get_<status_name>()` method. For example, a `DataReader` will have a `get_sample_lost_status()` method to obtain a snapshot of its `SampleLostStatus`, and a `Topic` will have a `get_inconsistent_topic_status()` method to obtain a snapshot of its `InconsistentTopicStatus`.

Read communication status are used to indicate the availability of data. The data may then accessed by calling the `DataReader read()` or `take()` methods. The application may call `read()` or `take()` directly at any time, even if there is no data available.

While the application can choose to call `get_xxx_status()`, `read()`, or `take()` at any time, typically the application will wait for notification from the infrastructure that a status has changed (or data is available) before accessing the status information or data.

There are two mechanisms an application may use to learn about changes in communication statuses and statistics. The first is **listeners**, where an application can asynchronously handle a change in communication statuses. The second is **conditions** (using wait sets), where an application can block waiting for a status change.

14.2.1 Listeners

Listeners provide an asynchronous method for the application to be notified of changes in statuses. The application provides a hook for the CoreDX DDS middleware to invoke upon a particular status change. For example, an application interested in the Data Available Status for its DataReader will provide an `on_data_available()` method to the CoreDX DDS middleware, and the CoreDX DDS middleware will call the provided method when new data is available on that DataReader.

All DDS entities support a listener, and all listeners have a type specific to their associated entity. For example, the `DataReaderListener` is associated with a `DataReader`.

Listeners are interfaces. Each listener provides a set of methods that correspond to the relevant communication statuses for that entity.

Listeners are hierarchical. *Figure 14-12: Listener Hierarchy* depicts the hierarchy of all the listeners.




Figure 14-12: Listener Hierarchy

The application must implement an appropriate listener interface in order to receive communication status changes. The following Table 14-2 depicts the listener methods for each entity.

Table 14-2: Listener Method Signatures

Entity	Listener Method Signature
DataReaderListener	void on_requested_deadline_missed (DataReader, RequestedDeadlineMissedStatus); void on_requested_incompatible_qos (DataReader, RequestedIncompatibleQosStatus); void on_sample_rejected (DataReader, SampleRejectedStatus); void on_liveliness_changed (DataReader, SampleRejectedStatus); void on_data_available (DataReader); void on_subscription_matched (DataReader, SubscriptionMatchedStatus); void on_sample_lost (DataReader, SampleLostStatus);
SubscriberListener	void on_data_on_readers (Subscriber); (inherits all DataReaderListener methods)
TopicListener	void on_inconsistent_topic (Topic, InconsistentTopicStatus);
DataWriterListener	void on_liveliness_lost (DataWriter, LivelinessLostStatus); void on_offered_deadline_missed (DataWriter, OfferedDeadlineMissedStatus); void on_offered_incompatible_qos (DataWriter, OfferedIncompatibleQosStatus); void on_publication_matched (DataWriter, PublicationMatchedStatus);
PublisherListener	(inherits all DataWriterListener methods)
DomainParticipantListener	(inherits all SubscriberListener, PublisherListener, and TopicListener methods)

Notice that the listener methods for plain communication statuses follow the same format: they return a void and take the entity and appropriate status as arguments. The listener methods for read communication statuses are a little different. Read communication statuses do not have associated status structures. The only argument is the concerned DataReader (for the `on_data_available()` method) and Subscriber (for the `on_data_on_readres()` method). It is assumed that in handling the read communication status, the application intends to eventually read the available data, and providing the appropriate DataReader or Subscriber allows the application to do this.

The application may use all, some, or none of the listener methods for an Entity. For example, an application may only be interested in the DataReader’s `data_available` status, and implement only the `on_data_available()` listener method. When the application attaches a listener to an entity, it must also set a mask that indicates which listener methods are enabled within this listener.

14.2.1.1 Listener Access to Plain Communication Statuses

Notice in *Figure 14-12* that the listeners form a hierarchy. When a plain communication status changes, the middleware will invoke the most specific relevant listener method that is enabled.

For example, consider the `PublicationMatchedStatus`. The corresponding `on_publication_matched()` listener method comes from the `DataWriterListener`, and is inherited by the `PublisherListener` and the `DomainParticipantListener`. When there is a change to the `PublicationMatchedStatus`, the DDS infrastructure will look for an enabled `on_publication_matched()` listener method to invoke. It will look at the `DataWriterListener` first. If there is not an enabled `on_publication_matched()` listener method, it will then look at the `PublisherListener`. If there is not an enabled `on_publication_matched()` listener method, it will look at the `DomainParticipantListener`. The first enabled listener method will be invoked. If there are no listeners enabled, no listener methods will be invoked. The status is still available and may trigger a configured Condition, or be accessed by calling the associated `get_xxx_status()`.

14.2.1.2 Listener Access to Read Communication Statuses

The read communication status listeners are invoked differently than plain communication status listeners. The two read communication statuses

constitute the real purpose of the Data Distribution Service, and require special consideration.

Each time a read communication status changes the DDS perform the following actions.

1. The infrastructure will attempt to invoke the `on_data_on_readers()` method on the `SusbscriberListener` with a parameter of the related `Subscriber`
2. If this doesn't work (either there was no listener installed or the method was not enabled via the listener mask), the DDS middleware will attempt to trigger the `on_data_available()` method on the related `DataReaderListener` with a parameter of the related `DataReader`.

14.2.1.3 Nil Listeners

The application can choose to install a *nil listener* in place of any listener method. When the infrastructure finds a nil listener, it will perform a NO-OP operation and stop looking for enabled listener methods.

14.2.1.4 Implementing Listeners in C

Listeners in C are implemented as a structure of function pointers. In order to implement a listener method, the application must write a function, and then assign it to the appropriate listener function pointer.

For example, a `DataReader` interested in only the `DATA_AVAILABLE` status might do the following:

on_data_available Listener

```
void my_on_data_available( DDS_DataReader dr )
{
    Printf("Data is available!\n");
    /* process data by calling DDS_DataReader_read()
     * or DDS_DataReader_take()
     */
}

DDS_DataReaderListener drListener =
{
    NULL, NULL, NULL, NULL,
    my_on_data_available,
    NULL, NULL
}
```

```
/* when we create the DataReader */
Dr = DDS_Subscriber_create_datareader( sub,
                                       DDS_Topic_TopicDescription(topic), &drListener,
                                       DDS_DATA_AVAILABLE_STATUS);
```

Figure 14-13: Listener Example C Code

The last argument to the `create_datareader()` method is the listener status mask, telling the DDS middleware which listener methods are enabled.

In C, a nil listener is installed simply by setting the appropriate function pointer to NULL, and then setting that status in the listener mask. Suppose in the above example, the listener mask used in the `create_datareader()` call was:

```
DDS_DATA_AVAILABLE_STATUS |
DDS_LIVELINESS_CHANGED_STATUS
```

Then there would be a nil listener installed for the `on_liveliness_changed()` method, and an actual listener installed for the `on_data_available()` method. All other listener methods would default “down” the listener hierarchy, looking for enabled corresponding listener methods on the Subscriber, and then the DomainParticipant.

The definition of all the C listeners can be found in the CoreDX DDS header file (`DDS_HOME/include/dds/core/dds.h` or `dds.hh`).

14.2.1.5 Implementing Listeners in C++

Listeners in C++ are classes containing virtual `on_<communication_status>()` methods. In order to implement a listener method, the application must create a listener class that derives from the appropriate virtual listener class, and then implement the desired listener method.

For example a DataReader who is only interested in the `DATA_AVAILABLE` status might do the following:

on_data_available Listener

```
class MyDRListener : public DataReaderListener
{
public:
```

```

        void on_data_available( DataReader * dr );
    };

    void MyDRListener::on_data_available(DataReader * dr)
    {
        printf("Data Available!\n");
        /* process data by calling read() or take()
        */
    }

    /* when we create the DataReader */
    MyDRListener drListener;
    Dr = sub->create_datareader( (TopicDescription*)topic,
    DATAREADER_QOS_DEFAULT, &drListener, DATA_AVAILABLE_STATUS );

```

Figure 14-14: Listener Example C++ Code

To install a nil listener in C++, use the `nil_listeners` member on the appropriate listener object. For example, to add a nil listener for the `on_liveliness_changed()` listener method:

```
drListener . nil_listeners = LIVELINESS_CHANGED_STATUS;
```

This indicates to the DDS middleware that the `on_liveliness_changed()` listener method should be treated as a nil listener method.

14.2.1.6 Implementing Listeners in C#

Listeners in C# are classes that may be used as a base class for application data specific Listeners. In order to implement a listener method, the application must create a listener class that derives from the appropriate listener base class, and then implement the desired listener method.

For example a DataReader who is only interested in the `DATA_AVAILABLE` status might do the following:

on_data_available Listener

```

class MyDRListener : DataReaderListener
{
    PublicMyDRListener()
    {
        this.on_requested_deadline_missed = null;
        this.on_requested_incompatible_qos = null;
        this.on_sample_rejected = null;
        this.on_liveliness_changed = null;
        this.on_data_available = data_available;
    }
}

```

```
this.on_subscription_matched      = null;
this.on_sample_lost               = null;
    }

public void data_available( DataReaderdr )
{
    System.Console.WriteLine("Data Available!\n");
    /* process data by calling read() or take()
       */
}

/* when we create the DataReader */
MyDRLListener drListener = new MyDRLListener();
Dr = sub->create_datareader( topic, DDS.DATAREADER_QOS_DEFAULT,
drListener, DDS.DATA_AVAILABLE_STATUS );
```

The last argument to the `create_datareader()` method is the listener status mask, telling the DDS middleware which listener methods are enabled.

In C#, a nil listener is installed simply by setting the appropriate function pointer to NULL, and then setting that status in the listener mask. Suppose in the above example, the listener mask used in the `create_datareader()` call was:

```
DDS.DATA_AVAILABLE_STATUS |
DDS.LIVELINESS_CHANGED_STATUS
```

Then there would be a nil listener installed for the `on_liveliness_changed()` method, and an actual listener installed for the `on_data_available()` method. All other listener methods would default “down” the listener hierarchy, looking for enabled corresponding listener methods on the Subscriber, and then the DomainParticipant.

14.2.1.7 Implementing Listeners in Java

Listeners in Java are interfaces containing empty `on_<communication_status>()` methods. In order to implement a listener method, the application must create a listener class that implements the appropriate listener interface, and then create public versions of all listener methods. Listener methods may be empty.

For example a DataReader who is only interested in the `DATA_AVAILABLE` status might do the following:

on_data_available Listener

```

class MyDRListener implements DataReaderListener
{
    public long get_nil_mask() { return 0; }

    public void on_data_available(DataReader dr)
    {
        System.out.println(" DATA AVAILABLE ");
        /* process data by calling read() or take()*/
    }

    public void on_requested_deadline_missed(DataReader dr,
        RequestedDeadlineMissedStatus status) { };
    public void on_requested_incompatible_qos(DataReader dr,
        RequestedIncompatibleQosStatusstatus) { };
    public void on_sample_rejected (DataReader dr,
        SampleRejectedStatusstatus) { };
    public void on_liveliness_changed (DataReader dr,
        LivelinessChangedStatusstatus) { };
    public void on_subscription_matched (DataReader dr,
        SubscriptionMatchedStatusstatus) { };
    public void on_sample_lost(DataReader dr,
        SampleLostStatusstatus) { };

    /* when we create the DataReader */
    DataReaderListener dr_listener = new MyDRListener();
    dr = sub.create_datareader( topic, DDS.DATAREADER_QOS_DEFAULT,
        drListener, DDS.DATA_AVAILABLE_STATUS );
}

```

To install a nil listener in Java, implement the `get_nil_mask()` method on the appropriate listener object to return the corresponding status. For example, to add a nil listener for the `on_liveliness_changed()` listener method:

```

public long get_nil_mask() { return
    DDS.LIVELINESS_CHANGED_STATUS; }

```

This indicates to the DDS middleware that the `on_liveliness_changed()` listener method should be treated as a nil listener method.

14.2.1.8 Examples for other language bindings

Setting up listeners in other language bindings (including Cxx and Python) are included in the 'cookbook' section of the relevant Reference Manuals.

14.2.2 Conditions and WaitSets

The listener notification method is asynchronous. Conditions and WaitSets provide a wait-based mechanism to be notified of changes in the CoreDX DDS infrastructure. In general, the application using Conditions and WaitSets will use the following pattern.

1. The application indicates which statuses it is interested in by obtaining or creating one or more **Condition** objects and attaching them to a **WaitSet**.
2. The application waits on the *WaitSet* until the trigger value of one or more of the attached *Condition* objects becomes TRUE.
3. The application calls `get_status_changes()` to determine the what changed.
4. The application calls the appropriate `get_<communication_status>()`, `read()`, or `take()` method(s).

Conditions are always used in combination with a WaitSet. The Condition contains a trigger value that is set when there is a change (change in communication status, change in read status, for example). The WaitSet blocks the application until the Condition’s trigger value is set (or until an application-defined timeout is reached). There are different kinds of Conditions available for the application to use. These are described in the sections below.

14.2.2.1 StatusConditions

StatusConditions allow the application to access plain communication statuses. A Condition can be obtained from each entity that contains a communication status, by using the `get_status_condition()` operation.

A StatusCondition contains a mask of enabled statuses. Similar to the listener mask, this mask allows the application to tailor the Condition to trigger only for specific status changes. For example, a `DataWriter` contains four statuses: `liveliness_lost`, `offered_deadline_missed`, `offered_incompatible_qos`, and `publication_matched`. The StatusCondition returned from `DataWriter::get_statuscondition()` will have all four statuses enabled by default. If any one of these statuses changes for the `DataWriter`, the Condition will be set, and the application waiting on the corresponding WaitSet will be signaled. The application can use the enabled statuses mask to enable only the communications statuses that are of interest to the application.

14.2.2.2 ReadConditions and QueryConditions

ReadConditions and QueryConditions allow the application to be notified when data is available. Specifically, these Conditions allow the application to be notified when a specific *kind* of data is available. The ReadCondition allows the application to configure the view_states, instance_states, and sample_states that the ReadCondition should trigger on. The QueryCondition allows the application to define the content of the data samples that should trigger the QueryCondition. This is done using an sql-like query string.

14.2.2.3 Guard Conditions

GuardConditions allow the application to control triggering the condition. GuardConditions are not attached to a CoreDX DDS entity (DataReader, Topic, etc), rather they can be used by the application for synchronization efforts outside the DDS middleware. Because these conditions are not attached to a CoreDX DDS entity, they are created by using the GuardCondition__alloc() operation (C interface), or the GuardCondition constructor (C++, C#, Java interfaces). The application is responsible for releasing the GuardCondition memory.

14.2.2.4 Implementing Conditions in C

Condition Examples

```
/* Variables used in sample code */
DDS_DataReader dr;
DDS_StatusMask sm;
DDS_StatusCondition sc;
DDS_ReadCondition rc;
DDS_WaitSet ws;
DDS_ConditionSeq active_conditions;
DDS_Duration_t timed;
DDS_ReturnCode_t retval;

/* Create DomainParticipant, register data type,
 * create Topic, Subscriber - code not shown here.
 */

/* Create DataReader with no Listeners */
dr = DDS_Subscriber_create_datareader( sub,
                                       DDS_TopicDescription(topic),
```

```
DDS_DATAREADER_QOS_DEFAULT,  
NULL, 0);  
  
/* Get the StatusCondition, configure to only trigger on  
* subscription matched status changes  
*/  
sc      = DDS_DataReader_get_statuscondition( dr );  
sm      = DDS_SUBSCRIPTION_MATCHED_STATUS;  
retval = DDS_StatusCondition_set_enabled_statuses( sc, sm );  
  
/* Create the ReadCondition, configure to trigger  
* only on samples not yet read  
*/  
rc = DDS_DataReader_create_readcondition( dr,  
                                           DDS_NOT_READ_SAMPLE_STATE,  
                                           DDS_ANY_VIEW_STATE,  
                                           DDS_ANY_INSTANCE_STATE );  
  
/* Create a WaitSet and attach all my conditions */  
ws = DDS_WaitSet_init( DDS_WaitSet_alloc() );  
retval = DDS_WaitSet_attach_condition( ws, sc );  
retval = DDS_WaitSet_attach_condition( ws, rc );  
  
/* Wait for something to happen (no timeout) */  
INIT_SEQ(active_conditions);  
timed.sec = 0;      /* infinite */  
timed.nanosec = 0;  
retval = DDS_WaitSet_wait( ws, &active_conditions, &timed );  
  
/* Something woke us up -- check */  
for (i=0 ; i<seq_get_length(&active_conditions) ; i++)  
{  
    if ( active_conditions._buffer[i] == (DDS_Condition)sc )  
    {  
        /* If you don't already know what entity this  
        * StatusCondition is attached to, here's how  
        * to find out.  
        */  
        DDS_Entity e = DDS_StatusCondition_get_entity(sc);  
        DDS_DataReader r = (DDS_DataReader)e;  
  
        DDS_StatusMask s = DDS_DataReader_get_status_changes(r);  
        if ( s& DDS_SUBSCRIPTION_MATCHED_STATUS )  
            /* Handle subscription match */  
    }  
  
    else if ( active_conditions._buffer[i] == (DDS_Condition)rc )  
    {  
        /* handle data available on our DataReader 'dr' */  
    }  
}
```

```

}

/* Cleanup */
retval = DDS_DataReader_delete_readcondition( dr, rc );

```

Figure 14-15: Condition Example C code

14.2.2.5 Implementing Conditions in C++

Condition Examples

```

using namespace DDS;

/* Variables used in sample code */
DataReader *      dr;
StatusMask        sm;
StatusCondition * sc;
ReadCondition *   rc;
WaitSet           ws;
ConditionSeq       active_conditions;
Duration_t         timed;
ReturnCode_t       retval;

/* Create DomainParticipant, register data type,
 * create Topic, Subscriber - code not shown here.
 */

/* Create DataReader with no Listeners */
dr = sub->create_datareader( (TopicDescription)topic,
                           DATAREADER_QOS_DEFAULT,
                           NULL, 0 );

/* Get the StatusCondition, configure to only trigger on
 * subscription matched status changes
 */
sc = dr -> get_statuscondition();
sm = SUBSCRIPTION_MATCHED_STATUS;
retval = sc -> set_enabled_statuses( sm );

/* Create the ReadCondition, configure to trigger
 * only on samples not yet read
 */
rc = dr -> create_readcondition( NOT_READ_SAMPLE_STATE,
                                ANY_VIEW_STATE,
                                ANY_INSTANCE_STATE );

/* Create a WaitSet and attach all my conditions */
retval = ws . attach_condition( sc );
retval = ws . attach_condition( rc );

```

```
/* Wait for something to happen (no timeout) */
timed.sec = 0;          /* infinite */
timed.nanosec = 0;
retval = ws .wait( &active_conditions, &timed );

/* Something woke us up -- check */
for (i=0 ; i<active_conditions . size() ; i++)
{
    if ( active_conditions[i] == (Condition)sc )
    {
        /* If you don't already know what entity this
         * StatusCondition is attached to, here's how
         * to find out.
         */
        Entity e      = sc -> get_entity();
        DataReader r = (DataReader)e;

        StatusMask s = r -> get_status_changes();
        if ( s& SUBSCRIPTION_MATCHED_STATUS )
            /* Handle subscription match */
    }

    else if ( active_conditions[i] == (Condition)rc )
    {
        /* handle data available on our DataReader 'dr' */
    }
}

/* Cleanup */
retval = dr -> delete_readcondition( dr, rc );
```

Figure 14-16: Condition Example C++ code

14.2.3 Using Listeners and Conditions in Combination

The application may choose to use both listeners and conditions in combination. One way to do this is using listeners for some communication statuses and using conditions for other communication status. However, if both listeners and conditions are used for an individual communication status, the listener method is invoked first and then the condition objects are signaled. Since calling a listener has the effect of “resetting” the status, when the Condition is signaled the corresponding status will not be set.

Part 4: CoreDX DDS Extensions

This section introduces some CoreDX DDS specific extensions to the DDS API. These include facilities for logging, transports, licensing, adjusting the discovery mechanisms, naming DDS entities, and other concepts that make using DDS easier.

Chapter 14 CoreDX DDS Logging

This chapter describes the logging facilities provided by CoreDX DDS and how to configure them.

14.1 Logging Libraries

First, it should be noted that there are two versions of CoreDX DDS libraries provided. The first is the optimized, performance-focused library which does not contain the extra logging instrumentation code. This library is faster and smaller than the instrumented library. This library contains very little logging, primarily limited to indicating error or anomalous conditions.

The second version of libraries includes a rich set of logging instrumentation. During application development and debugging, it may be useful to link against the logging version of the CoreDX DDS libraries. Then, deployed applications can be linked with the streamlined libraries for enhanced performance and resource utilization.

For C# and Java users: The `coredx_csharp.dll` (C#) and `coredx_dds.jar` (java) packages refer to the non-logging native library name (`dds_csharp.dll/lib` for C# and `dds_java.so/dll/lib` for java). In order to use the logging versions of these libraries, these libraries must be renamed (or linked to). For example, to use the Java logging library under Linux:

```
% cd ${COREDX_TOP}/target/${TARGET_ARCH}
% ls -l libdds_java*
-rw-r--r-- 1 user grp libdds_java_log.so
-rw-r--r-- 1 user grp libdds_java.so
% mv libdds_java.so libdds_nolog.so
% ln -s libdds_java_log.so libdds_java.so
% ls -l libdds_java*
-rw-r--r-- 1 user grp libdds_java_log.so
-rw-r--r-- 1 user grp libdds_java_nolog.so
-rw-r--r-- 1 user grp libdds_java.so ->
libdds_java_log.so
```

14.2 Logging Control

Once the application has been linked with the instrumented library, the logging output can be controlled either by environment variable or by API via adjusting the `CoreDX_LoggingQosPolicy`.

Controlling logging with an environment variable is quick and easy, and will globally enable the logging flags for the entire application.

Environment variables to enable logging are:

- *COREDX_DEBUG*: turn on debug logging for most DDS Entities
- *COREDX_UDP_DEBUG*: turn on debug logging in the UDP transport code
- *COREDX_TCP_DEBUG*: turn on debug logging in the TCP transport code
- *COREDX_SECURITY_DEBUG*: turn on security debug logging
- *COREDX_TLM_DEBUG*: turn on license related debug logging

The Logging QoS policy offers finer-grained control over which entities are enabled with which logging flags. The following example illustrates enabling the 'DATA' and 'STATUS' logging flags on a DataReader.

Table 14-1: Logging QoS Configuration Example

Setting CoreDX DDS Logging Flags

```
DDS_Subscriber      sub;
DDS_DataReader      dr;
DDS_DataReaderQos   dr_qos;

DDS_Subscriber_get_default_datareader_qos(sub, &dr_qos);
dr_qos.logging.flags |= (COREDX_FACTORY_LOGGING_QOS |
    COREDX_DATA_LOGGING_QOS |
    COREDX_STATUS_LOGGING_QOS);
dr = DDS_Subscriber_create_datareader(sub, td, &dr_qos, NULL, 0);
```

14.3 Logging Flags

Whichever mechanism is used, the level of logging output is controlled by a series of bitmapped flags. Each class of log message is enabled by setting a flag to 1, and is disabled by 0.

The following table lists some of the logging flags available. The full list of logging flags can be found in
 \$COREDX_TOP/target/include/dds/coredx_logging.h

Table 14-2: CoreDX DDS Logging Flags

CoreDX DDS Debug Categories and Flags

COREDX_ERROR_LOGGING_QOS	0x0001
COREDX_DATA_LOGGING_QOS	0x0002
COREDX_DISCOVERY_LOGGING_QOS	0x0004
COREDX_FACTORY_LOGGING_QOS	0x0008
COREDX_LIVELINESS_LOGGING_QOS	0x0010
COREDX_STATUS_LOGGING_QOS	0x0020
COREDX_TRANSPORT_LOGGING_QOS	0x0040
COREDX_SCHEDULE_LOGGING_QOS	0x0080
COREDX_HANDSHAKE_LOGGING_QOS	0x0100
COREDX_CACHE_LOGGING_QOS	0x0200
COREDX_SECURITY_LOGGING_QOS	0x0400
COREDX_TRACE_LOGGING_QOS	0x0800
COREDX_RPC_FACTORY_LOGGING_QOS	0x10000
COREDX_RPC_REQUEST_LOGGING_QOS	0x20000
COREDX_RPC_REPLY_LOGGING_QOS	0x40000

By default, all logging output is directed to **stderr**. There is an alternative API allowing logging messages to be sent to a user-defined function, allowing the user to specify the processing or storing of log messages. The following code illustrates how this API might be used:

Using the CoreDX DDS logio_routine() routine to handle log messages

```

/* In the include/dds/toc_util.h header, this API is provided  *
 *
 * typedef int (*logio_routine)(const char *buf, size_t size); *
 * extern void toc_set_logio_routine(logio_routine logger);    *
 */

/* Here is an example of how to use it */

int my_logger ( const char * buf, size_t len )
{
    /* param: buf is a nul terminated string */
    /* param: len is equivalent to strlen(buf) */
    /* return: the return value from this routine is ignored */

    printf( ">>>  %s", buf );

```

```
    return 0;
}

int main(int argc, char * argv[] )

    /* early in application initialization: */
    toc_set_logio_routine( my_logger );

    /* DDS initialization ... */
    ...
}
```

Chapter 15 CoreDX DDS Transport

The CoreDX DDS transport conforms to the Real-Time Publish-Subscribe (RTPS) Wire Protocol. This transport does not use a stand-alone transport daemon, and does not require configuration of any operating system services.

15.1 Overview

CoreDX DDS includes a modular transport infrastructure that supports configuration and customization. CoreDX DDS ships with support for UDP (on IPv4 and IPv6), TCP (IPv4, IPv6 planned), LMT (Local Machine Transport), and SERIAL. [Some platforms do not support all kinds of transports – check with Twin Oaks to determine the availability for your platform.]

Each transport implementation includes the capability to configure aspects of the transport. The set of configuration options available for each transport are described in detail in subsequent sections.

By default, CoreDX DDS will install and use the UDP transport. Alternatively, the developer can configure and install alternate or additional transports during the initialization of a DomainParticipant.

15.1.1 Transport Common API

Each transport implementation provides a set of methods to facilitate access to configuration parameters and creation of the transport; these are referred to as the **Transport Initialization API**. In addition to this API, this section also describes the DomainParticipant method to add the configured transport(s).

Using the CoreDX DDS Transport Initialization API, the developer can configure and create an instance of a transport. The API provides a mechanism to retrieve the default configuration settings. These settings can be modified by the developer as required before passing them to ‘create’. The ‘create’ operation accepts a structure that contains all of the configuration parameters.

For example, the ‘C’ **Transport Initialization API** for the **UDP** transport consists of the following methods:

```
DDS_ReturnCode_t
CoreDX_UdpTransport_get_default_config(
CoreDX_UdpTransportConfig * config );
```

```
DDS_ReturnCode_t
CoreDX_UdpTransport_get_env_config(
CoreDX_UdpTransportConfig * config );
```

```
DDS_ReturnCode_t
CoreDX_UdpTransport_clear_config(
CoreDX_UdpTransportConfig * config );
```

```
CoreDX_Transport
*CoreDX_UdpTransport_create_transport(
CoreDX_UdpTransportConfig * config );
```

Each transport implementation will provide this set of Transport Initialization API methods. These methods are further described below.

15.1.1.1 Transport::get_default_config()

The `get_default_config()` method will initialize all fields in the provided `TransportConfig` structure with default values. This provides a good starting point for custom configuration modifications. If the user does not call `get_default_config()`, then the user is responsible for manually initializing all fields in the `TransportConfig` structure.

Once the `TransportConfig` structure has been initialized with the desired configuration values, the configuration can be passed to the `create_transport()` method. The transport will utilize the provided configuration values to initialize the characteristics of the transport. Once `create_transport()` returns, the caller can destroy or reuse the `TransportConfig` structure – the transport implementation does not hold any references to the provided structure.

The caller is responsible for clearing any allocated memory within the `TransportConfig` structure. The method `clear_config()` will accomplish this.

15.1.1.2 Transport::get_env_config()

This method will override values in the provided `TransportConfig` structure with values taken from the environment. Some transports allow the user to adjust configuration parameters by setting environment variables. This

routine will query the environment and set configuration parameters based on discovered values. [Not all transports use environment variables for configuration.] In all cases, the environment variable based configuration parameters are provided as a convenience – the same result can be obtained by directly adjusting values within the TransportConfig structure.

The provided TransportConfig structure should be initialized to default values (either manually or by calling `get_default_config()`) prior to calling this routine.

15.1.1.3 Transport::clear_config()

This routine will clear any allocated memory within the provided TransportConfig structure. In some cases, a TransportConfig structure may contain dynamically allocated values. This memory may have been allocated by the `get_default_config()` or `get_env_config()` methods, or by the user. Each transport has a different TransportConfig structure which may or may not include parameters that have dynamic memory allocations. `Clear_config()` provides a mechanism to release any dynamically allocated memory within the TransportConfig structure.

15.1.1.4 Transport::create_transport()

This routine will create an instance of the transport. The returned transport instance can then be registered with a DomainParticipant. The provided TransportConfig structure will be used to configure the transport during creation. The user maintains ownership of the TransportConfig structure, and can clear or reuse it after the `create_transport()` call returns.

On success, the `create_transport()` method will return a pointer to an initialized transport instance. If there is an error during creation, NULL will be returned. Normally, the resulting transport will be passed to `DomainParticipant::add_transport()`. In this case, the DomainParticipant takes ownership of the transport, and the user should no longer access the transport instance. If the user does not register the transport with a DomainParticipant, then the user maintains ownership of the transport instance and is responsible for destroying it. This can be accomplished by invoking the `Transport::destroy` method.

15.1.1.5 DomainParticipant::add_transport()

This routine, while not strictly part of the Transport API, is used to add a configured transport to the DomainParticipant.

```
DDS_ReturnCode_t
DDS_DomainParticipant_add_transport(
    DDS_DomainParticipant dp,
    CoreDX_Transport * transport);
```

If the application does not install any transports using the `add_transport()` method, the `DomainParticipant` will automatically create and register a default UDP transport.

15.1.2 Transport Configuration

This section documents the details of each transport. This includes a discussion of the transport specific configuration parameters.

15.1.2.1 UDP Transport API

The UDP transport is the default CoreDX DDS transport. It provides the fully RTPS compliant, interoperable, DDS transport. There are many configuration options to the UDP transport.

It provides support for UNICAST and MULTICAST transmission and reception. By default, MULTICAST is enabled for both 'built-in' (discovery) data and 'user' data. For discovery configuration options over UDP, refer to Chapter 16 CoreDX DDS Discovery.

The specific configuration parameters available with the UDP transport are described in the following sections.

15.1.2.1.1 Participant Index

```
participant_index : short
```

Participant Index is an integer number that by default is computed automatically by the RTPS implementation. It is used to distinguish `DomainParticipants` on the same host and in the same Domain from one another. In some cases, it is helpful to directly configure this value as it is used to compute unicast port numbers. It is an error to configure 2 or more `DomainParticipants` on the same host, within the same Domain, such that they have the same Participant Index value.

15.1.2.1.2 IPv4 and IPv6

```
use_ipv4 : unsigned char, default = 1
use_ipv6 : unsigned char, default = 0
```


The transport is configured to operate over IPv4 by default; and can be configured to support IPv6. The ‘use_ipv4’ and ‘use_ipv6’ configuration parameters provide control over which version(s) of IP will be used.

15.1.2.1.3 Interfaces and Interface Names

```
interfaces      : CoreDX_IpTransportInterfaceSeq,  
interface_names : StringSequence
```

The transport, by default, will make use of all available active network interfaces. If your machine has multiple network interfaces, this may generate unnecessary network traffic on some of those networks. The transport can be configured to use a subset of available interfaces.

The ‘interfaces’ configuration parameter can be configured with a list of interface IP addresses. These IP addresses are used for sending multicast packets and added to the list of locators advertised to peers (even if the interfaces do not exist).

The ‘interface_names’ configuration parameter can be configured with a list of interface names. These interfaces, if they exist, are used for sending multicast packets and added to the list of locators advertised to peers.

‘Interface_names’ is consulted only if ‘interfaces’ is empty. If ‘interfaces’ is non-empty, then the addresses in ‘interfaces’ define the set of IP addresses used by the UDP transport and ‘interface_names’ is not considered. If both ‘interfaces’ and ‘interface_names’ are empty, then CoreDX DDS will query the operating system to obtain a list of all available active interfaces.

NOTE: When configuring a static list of ‘interfaces’ for CoreDX DDS to use, you may also want to disable dynamic interface detection (below).

15.1.2.1.4 Dynamic Interface Detection

```
dynamic_interfaces : unsigned char
```

On Operating Systems that support it, the CoreDX DDS UDP transport can detect changes to the network interfaces and adjust its configuration in response. For example, if the user brings up a new interface, CoreDX DDS will discover and utilize the new interface on the fly. This dynamic interface detection is configurable. The ‘dynamic_interfaces’ configuration parameter is used to enable or disable this capability.

Interaction with ‘interfaces’ configuration: When ‘dynamic_interfaces’ is enabled and ‘interfaces’ is configured with one or more IP addresses, any

newly detected interfaces will be used in addition to the ones specified in the 'interfaces' list. When 'dynamic_interfaces' is disabled, only the interfaces specified by 'interfaces' will be used.

Interaction with 'interface_names' configuration: When 'dynamic_interfaces' is enabled and 'interface_names' is configured with one or more names, any newly detected interfaces will be used only if they match with an entry in 'interface_names'.

15.1.2.1.5 RX Buffer Sizes

```
rx_init_buffer_size : int,
rx_max_buffer_size  : int
```

The UDP transport maintains buffers to handle incoming data packets. In order to preserve memory, the size and behavior of the receive buffer is configurable. By default, the receive buffer begins small and can grow dynamically as required to handle the incoming data. The initial size of the buffer is determined by the 'rx_init_buffer_size' parameter. The maximum size of the buffer is limited by the 'rx_max_buffer_size' configuration parameter. If these two values are identical, then the buffer will not dynamically resize, and all memory allocation is performed during initialization.

NOTE: If the maximum size of the RX Buffer is limited to some value smaller than that allowed by the underlying transport (in this case, UDP maximum datagram size is 64 KB), then it is possible that the transport will be forced to drop some incoming data. The size of incoming UDP datagrams is determined by the remote writing application. If the remote writer is from a CoreDX DDS DomainParticipant, then the remote peer can be configured to limit the size of transmitted packets. This configuration will enable transmission of large data between two peers without requiring the transport to establish a large RX buffer.

15.1.2.1.6 TX Packet Size Limit

```
tx_max_buffer_size : int
```

The CoreDX DDS UDP Transport transmits information in UDP datagrams. The underlying UDP transport mechanism support datagram sizes up to 64KB. In some cases, it is beneficial to limit the size of datagram put onto the network. For example, some network devices fail to handle large datagrams. The 'tx_max_packet_size' configuration parameter is used to limit the size of UDP datagram produced by the CoreDX DDS UDP Transport.

CoreDX DDS will fragment the data message if it will not fit within the specified maximum size.

15.1.2.1.7 *SNDBUF and RCVBUF*

```
so_rcvbuf : int
so_sndbuf : int
```

The UDP sockets used by the CoreDX DDS UDP Transport have an OS configured send and receive buffer. This is configurable through an Operating System provided API. In general, the OS provided default buffer sizes are appropriate; however, it is possible to override these defaults with the ‘so_rcvbuf’ and ‘so_sndbuf’ configuration parameters. For further information on these buffers, refer to the documentation for your Operating System under the topic of ‘setsockopt’ and SO_RCVBUF or SO_SNDBUF.

15.1.2.1.8 *Multicast and Unicast*

The CoreDX DDS UDP Transport supports the use of Unicast and Multicast datagrams. CoreDX DDS will use Multicast when available to minimize the number of packets written on the network. In general, this is for all communications except for:

- Heartbeat and ACK/NACK messages (RELIABLE Reliability)
- Retransmission of data packets (RELIABLE Reliability)
- Content filtered data with Writer-side filtering enabled

The addresses used for multicast and unicast communications may be configured.

If Multicast is not available or desirable, then CoreDX DDS can be configured to use only Unicast transmissions. There are several configuration parameters available to tailor the use of Unicast and Multicast. In some cases, it may be useful to use Multicast for ‘meta’ (discovery) topics, but not for **user** topics. In some cases, it is useful to transmit multicast, but not receive it.

In order to provide full flexibility, the CoreDX DDS UDP Transport provides the following configuration parameters related to Unicast and Multicast:

Table 15-1: UDP Transport Multicast and Unicast Configuration Parameters

Parameter	Description
-----------	-------------

Parameter	Description
meta_multicast_address_v4 user_multicast_address_v4	Specify the MULTICAST GROUP address used for all multicast communications over IPv4.
meta_multicast_address_v6u ser_multicast_address_v6	Specify the MULTICAST GROUP address used for all multicast communications over IPv6.
multicast_ttl	Specify the MULTICAST TTL value. This defines the number of hops (routers) that multicast packets should traverse.
tx_meta_multicast	Enables the transmission of META (discovery) data over multicast.
tx_meta_unicast	Enables the transmission of META (discovery) data over unicast.
rx_meta_multicast	Enables the reception of META (discovery) data over multicast
rx_multicast	Enables the reception of USER data over multicast.
advertise_meta_multicast	Enables the advertisement of our ability to receive META data via multicast.
advertise_user_multicast	Enables the advertisement of our ability to receive USER data via multicast.
try_to_keep_mcast_local	When the DomainParticipant has matched Entites only on the same host, attempt to keep user multicast packets on the local 'loopback' interface, to avoid unnecessary traffic on external networks.

15.1.2.1.9 Broadcast

```
do_meta_broadcast : unsigned char
broadcast_address : unsigned char[4]
```

In some network environments, Multicast is not available or desirable. In these cases, it may be acceptable to use Broadcast as an alternative to facilitate dynamic discovery. The CoreDX DDS UDP Transport can support

the broadcast of DomainParticipant discovery information. By setting ‘do_meta_broadcast’, the DomainParticipant Data message will be broadcast onto the local network segment.

If operating on a host or network that does not support multicast, but does support broadcast, then the following configuration may be useful:

```
advertise_meta_multicast = 0;
advertise_user_multicast = 0;
do_meta_broadcast = 1;
```

This will prohibit remote peers from attempting multicast communication, but will support dynamic discovery via broadcast.

15.1.2.1.10 Debug

```
debug_flags : unsigned int
```

The ‘debug_flags’ parameter enables debug output from the UDP transport. Useful values are 2 (DATA_LOGGING) and 64 (TRANSPORT_LOGGING). The flags can be combined. See <dds/dds.h> for the full set of LOGGING flags.

15.1.2.2 UDP Transport Environment Variables

CoreDX DDS provides the ability to set many of the UDP transport configuration items through environment variables. All settings available through environment variables are also available through the Transport API.

When only environment variables are used for transport configuration, the default transport is automatically augmented by these environment variables. When a combination of environment variable and transport configuration API are used, the Transport::get_env_config() API must be called in order to apply the environment settings.

Table 15-2: UDP Transport Environment Variables

Environment Variable	Meaning	Example
CORED_X_IP_ADDR	Sets the ‘interfeaces’ parameter This configures the default IP address used by the UDP transport. If this value is defined in the environment, then the transport will use only the interface associated with the	CORED_X_IP_ADDR=192.168.1.5

Environment Variable	Meaning	Example
	specified address. Legacy behavior preserved: Setting this environment variable will disable dynamic interface detection (which can be re-enabled via the Transport API if desired).	
COREDX_IFNAME	Sets the 'interface_names' parameter This configures the default interface names used by the UDP transport. This value is checked only if COREDX_IP_ADDR is not set (or is empty).	
COREDX_UDP_DEBUG	Sets the debug_flags parameter.	COREDX_UDP_DEBUG=64
COREDX_UDP_RX_BUFFER_SIZE	Sets the rx_init_buffer_size parameter.	
COREDX_UDP_MAX_TX_SIZE	Sets the tx_max_packet_size parameter.	
COREDX_UDP_RCVBUF	Sets the so_rcvbuf parameter.	COREDX_UDP_RCVBUF=1024
COREDX_UDP_SNDBUF	Sets the so_sndbuf parameter.	COREDX_UDP_SNDBUF=1024
COREDX_USE_MULTICAST	Sets the 'advertise_user_multicast' and 'advertise_meta_multicast' parameters.	COREDX_USE_MULTICAST=1
COREDX_MULTICAST_TTL	Sets the multicast_ttl parameter.	COREDX_MULTICAST_TTL=2
COREDX_UDP_IPV4	Sets the 'use_ipv4' parameter.	COREDX_UDP_IPV4=1
COREDX_UDP_IPV6	Sets the 'use_ipv6' parameter	COREDX_UDP_IPV6=1

15.1.2.3 TCP Transport API

The TCP transport provides support for CoreDX DDS to communicate using TCP connections. Because TCP is a connection oriented transport, there is no facility for Multicast or Broadcast. Without Multicast or Broadcast, the TCP transport does not provide any facilities for fully Dynamic Discovery.

The current version of the TCP transport supports only IPv4. Support for IPv6 is planned for a subsequent release.

The specific configuration parameters available with the TCP transport are described in the following sections.

15.1.2.3.1 Participant Index

`participant_index : short`

Participant Index is an integer number that by default is computed automatically by the RTPS implementation. It is used to distinguish DomainParticipants on the same host and in the same Domain from one another. In some cases, it is helpful to directly configure this value as it is used to compute unicast port numbers. It is an error to configure 2 or more DomainParticipants on the same host, within the same Domain, such that they have the same Participant Index value.

15.1.2.3.2 Interfaces

`interfaces : CoreDX_IpTransportInterfaceSeq`

The transport, by default, will make use of all available active network interfaces. If your machine has multiple network interfaces, this may generate unnecessary network traffic on some of those networks. The transport can be configured to use a subset of available interfaces. The 'interfaces' configuration parameter can be configured with a list of interface addresses. If the list is empty, then CoreDX DDS will query the operating system to obtain a list of all available active interfaces.

NOTE: When configuring a static list of interfaces for CoreDX DDS to use, you may also want to disable dynamic interface detection (below).

15.1.2.3.3 Dynamic Interface Detection

`dynamic_interfaces : unsigned char`

On Operating Systems that support it, the CoreDX DDS TCP transport can detect changes to the network interfaces and adjust its configuration in

response. For example, if the user brings up a new interface, CoreDX DDS will discover and utilize the new interface on the fly. This dynamic interface detection is configurable. The 'dynamic_interfaces' configuration parameter is used to enable or disable this capability.

Interaction with 'interfaces' configuration: When 'dynamic_interfaces' is enabled and 'interfaces' is configured with one or more IP addresses, any newly detected interfaces will be used in addition to the ones specified in the 'interfaces' list. When 'dynamic_interfaces' is disabled, only the interfaces specified by 'interfaces' will be used.

15.1.2.3.4 TX Packet Size Limit

```
tx_max_buffer_size : int
```

The CoreDX DDS TCP Transport transmits information in RTPS Messages. In some cases, it is beneficial to limit the size of RTPS Messages put onto the network. For example, some network devices fail to handle large packets. The 'tx_max_packet_size' configuration parameter is used to limit the size of RTPS Messages produced by the CoreDX DDS TCP Transport. CoreDX DDS will fragment the data message into multiple smaller messages, if it will not fit within the specified maximum size.

15.1.2.3.5 Checksum

```
add_checksum      : unsigned char
use_checksum      : unsigned char
require_checksum  : unsigned char
```

The TCP transport has the ability to augment each RTPS message with a checksum to address potentially corrupted data on unreliable networks. This is controlled by the add_checksum configuration item. If set to 'true', then the transport will append a checksum submessage to every RTPS Message.

The TCP transport has the ability to parse and verify a received checksum. If use_checksum is 'true', then the transport will look for a checksum in each received RTPS Message. If the checksum is present, it will be validated. If the checksum fails, then the RTPS Message is discarded. If the checksum is not present, and require_checksum is 'true', then the RTPS Message is discarded.

Note, require_checksum is effective only if use_checksum is 'true'.

Note, if either `use_checksum = 'false'`, or `require_checksum = 'false'`, then it is possible for corrupted data to be accepted and passed up to the DDS application. For example, if the sender adds a checksum to a message, but the message is corrupted in transit such that it appears to not contain an checksum; then, if `require_checksum` is not `'true'`, the transport will accept the message [because there is no checksum to validate].

Note, if `require_checksum = 'true'`, then the application will not inter-operate with a CoreDX DDS prior to version 4.4.0 [no backwards compatibility].

15.1.2.3.6 Reconnect Delay

```
reconnect_delay : Duration_t
```

In environments with unreliable networks, or systems where the network interface used for DDS TCP communications is frequently disconnected and reconnected, the transport may be flooded with reconnection attempts. The `'reconnect_delay'` configuration item allows these reconnect attempts to be throttled.

15.1.2.3.7 Debug

```
debug_flags : unsigned int
```

The `'debug_flags'` parameter enables debug output from the UDP transport. Useful values are 2 (`DATA_LOGGING`) and 64 (`TRANSPORT_LOGGING`). The flags can be combined. See `<dds/dds.h>` for the full set of LOGGING flags.

15.1.2.4 TCP Transport Environment Variables

CoreDX DDS provides the ability to set the TCP transport configuration items through environment variables. The `Transport::get_env_config()` API must be called in order to apply those environment settings to a transport.

When only environment variables are used for transport configuration, the default transport is automatically augmented by these environment variables. When a combination of environment variable and transport configuration API are used, the `Transport::get_env_config()` API must be called in order to apply the environment settings.

Table 15-3: TCP Transport Environment Variables

Environment Variable	Meaning	Example
CORED_X_IP_ADDR	<p>This configures the default IP address used by the TCP transport. If this value is defined in the environment, then the transport will use only the interface associated with the specified address.</p> <p>Legacy behavior preserved: Setting this environment variable will disable dynamic interface detection</p>	<code>CORED_X_IP_ADDR=192.168.1.5</code>
CORED_X_TCP_DEBUG	<p>Enables debug output from the TCP transport. Useful values are 2 (DATA_LOGGING) and 64 (TRANSPORT_LOGGING). The flags can be combined. See <dds/dds.h> for the full set of LOGGING flags.</p>	<code>CORED_X_TCP_DEBUG=64</code>

15.1.2.5 LMT Transport API

The LMT (Local Machine Transport) provides support for optimized 'on host' communication. The transport enables DomainParticipants that are co-located on the same host to communicate with lower overhead and latencies than provided by the default UDP transport.

The LMT transport is currently provided on a subset of CoreDX DDS supported operating systems. For a current list of supported operating systems, please contact Twin Oaks Computing at support@twinoakscomputing.com.

NOTE: LMT is not implemented using shared memory. As a result, the safety provided by separate program address spaces is maintained.

The specific configuration parameters available with the UDP transport are described in the following sections.

15.1.2.5.1 *SNDBUF and RCVBUF*

The UDP sockets used by the CoreDX DDS UDP Transport have an OS configured send and receive buffer. This is configurable through an

Operating System provided API. In general, the OS provided default buffer sizes are appropriate; however, it is possible to override these defaults with the ‘so_rcvbuf’ and ‘so_sndbuf’ configuration parameters. For further information on these buffers, refer to the documentation for your Operating System under the topic of ‘setsockopt’ and SO_RCVBUF or SO_SNDBUF.

15.1.2.5.2 TX Packet Size Limit

In some cases, it is beneficial to limit the size of packet written by the transport. The ‘max_tx_size’ configuration parameter is used to limit the size of packet produced by the CoreDX DDS LMT Transport. CoreDX DDS will fragment the data message if it will not fit within the specified maximum size. Refer to the Chapter 21: Transmit Buffers section for additional information on buffer sizing and fragmentation.

15.1.2.5.3 RX Buffer Size

The LMT transport maintains buffers to handle incoming data packets. In order to preserve memory, the size and behavior of the receive buffer is configurable. By default, the receive buffer begins small and can grow dynamically as required to handle the incoming data. The initial size of the buffer is set at initialization to be just large enough to handle an RTPS Message Header. The maximum size of the buffer is limited by the ‘max_rx_buf_size’ configuration parameter.

15.1.2.5.4 Debug

The ‘debug_flags’ parameter enables debug output from the LMT transport. Useful values are 2 (DATA_LOGGING) and 64 (TRANSPORT_LOGGING). The flags can be combined. See <dds/dds.h> for the full set of LOGGING flags.

15.1.2.6 LMT Transport Environment Variables

CoreDX DDS provides the ability to set many of the LMT transport configuration items through environment variables. All settings available through environment variables is also available through the Transport API.

If any transport environment variables are used to configure the LMT transport, the Transport::get_env_config() API must be called in order to apply those environment settings to a transport.

Table 15-4: LMT Transport Environment Variables

Environment Variable	Meaning	Example
COREDX_LMT_DEBUG	Sets the 'debug_flags' parameter.	COREDX_LMT_DEBUG=64
COREDX_LMT_RCVBUF	Sets the 'rcvbuf' parameter.	
COREDX_LMT_SNDBUF	Sets the 'sndbuf' parameter.	
COREDX_LMT_MAX_TX_SIZE	Sets the 'max_tx_size' parameter.	
COREDX_LMT_MAX_RX_BUF_SIZE	Sets the 'max_rx_buf_size' parameter.	

15.1.2.7 UDS Transport

The UDS transport provides the facility for serial or other stream based transports. UDS is supported by a helper program that initializes the serial port, and provides a multi-participant access to the single shared resource. Because of the diverse nature of serial and other related hardware devices, please contact Twin Oaks Computing for assistance in adapting this transport technology to your platform.

Chapter 16 CoreDX DDS Discovery

16.1 Overview of CoreDX DDS Discovery

The automatic discovery process is one of the more powerful and useful features of CoreDX DDS. Automatic discovery of entities allows CoreDX DDS applications to publish and subscribe to data without needing to configure the endpoint(s) they talk to. Whether these endpoints are on the same machine, or across the room, CoreDX DDS applications do not need any knowledge of the other applications they will be communicating with.

The Standard (peer-to-peer, dynamic) Discovery process in CoreDX DDS is encapsulated in every CoreDX DDS application, and does not require any additional daemons or services. Each CoreDX DDS application performs the discovery process, including announcing the presence of its DDS Entities, listening for other DDS Entities, and looking for matches between its own DDS Entities and those discovered. The standard discovery mechanism is interoperable between DDS implementations.

The automatic discovery process includes the following steps:

1. Discovering DomainParticipants
2. Discovering DataReaders and DataWriters within those discovered Participants
3. *Matching* those discovered DataReaders and DataWriters with local DataReaders and DataWriters

Using the CoreDX DDS Secure extensions enhances this automatic discovery process. For more information, refer to the CoreDX DDS Secure Programmer's Guide.

Despite the many benefits of the Standard Discovery mechanism, it does have some drawbacks for certain system architectures. For example, where security requirements prevent dynamic discovery, where deployment environments require alternative multicast or unicast addresses, or where scalability requirements need an alternative solution to the standard peer-to-peer discovery. CoreDX DDS provides several alternatives for configuring the discovery process.

16.2 Discovering DomainParticipants

The first step in the automatic discovery process is to discover remote DomainParticipants. Each DomainParticipant will periodically announce its existence (including how it can be reached directly to learn about containing DataReader and DataWriter Entities) by writing a *SPDPdiscoveredParticipantData* message to the multicast address specified by the DDS standards. Each DomainParticipant will also listen on that same multicast address to learn about other DomainParticipants.

After a DomainParticipant has been discovered, it will be considered ‘alive’ as long as its *SPDPdiscoveredParticipantData* messages continue to be received. If enough time expires without receiving a *SPDPdiscoveredParticipantData* message from a DomainParticipant, that DomainParticipant is no longer considered ‘alive’.

16.2.1 Configuring Participant Discovery

The DDS standards specify default durations for how often *SPDPdiscoveredParticipantData* messages should be sent, and how much time should expire before a DomainParticipant should be considered ‘not alive’.

While these default durations work well for most network environments, they may not work for all environments. For example, networks with very long latencies, or extremely bandwidth constrained networks may need to tailor the timing of discovery messages.

CoreDX DDS allows the application to configure timing of DomainParticipant discovery by using the *CoreDX_DiscoveryQosPolicy*, as described below.

QoS Policy	Default Value	Description
CoreDX_DiscoveryQosPolicy		
dpd_push_period (duration_t)	5 seconds	Configure the amount of time between sending of <i>SPDPdiscoveredParticipantData</i> messages.
dpd_lease_duration (duration_t)	120 seconds	Configure the amount of time the can expire without receiving a

QoS Policy	Default Value	Description
		SPDPdiscoveredParticipantData message before we consider a remote Participant be 'not alive'.

16.3 Matching DataReaders and DataWriters

The matching of DataReaders and DataWriters is a sophisticated process that ensures the subscribers of data are matched appropriately with producers of data. This careful matching helps to reduce programming errors, in addition to reducing unnecessary storage and network communications usage.

The following three conditions are checked between DataReaders and DataWriters, and all must be met in order to create a *match*:

1. The Topic Names must match

The name used when creating a Topic with `DomainParticipant::create_topic()` is communicated to peer DomainParticipants during discovery. The name of the Topic used by a DataReader must match the name of the Topic used by a DataWriter.

2. The Type Names (and type definition, if available) must match

The name used when registering a data type with `TypeSupport::register_type()` is communicated to peer DomainParticipants during discovery. The name of the registered type associated with the Topic used by a DataReader must match the name of the registered type associated with the Topic used by a DataWriter.

Further type checking is performed when the type definition is available. The type definition is encoded for use during discovery in one of two ways: *typecode* or *typeobject*. These type definitions provide more accurate information about the data type actually being published by DataWriters and expected to be received by DataReaders. Using typecodes or typeobjects in Entity matching provides an additional level of type safety.

By default, CoreDX DDS v4 and newer applications exchange typeobject information during discovery. [CoreDX DDS v3 applications exchange typecode information during discovery.] More information on typecode and typeobject and their configuration can be found in the CoreDX DDS Type System Programmer’s Guide.

3. The QoS policies must be compatible

The QoS policy defined with DataReaders and DataWriters is communicated to peer DomainParticipants during discovery. Once the Topics and Types have been verified as matching between a DataReader and DataWriter, the QoS policies are checked for compatibility. Refer to section *QoS Compatibility* for additional information on QoS compatibility matching.

16.3.1 Configuring Entity Matching

Typecode and typeobject are not required for DDS discovery. By default, CoreDX DDS will send and use typeobject for Entity matching, but the application can configure this behavior. The RTPSReader and RTPSWriter QoS policies will control whether typecode and/or typeobject (if they are available) will be used for entity matching. Options to the data type compiler (-i) define if typecode, typeobject, or no type encodings will be generated, and available.

Further, when matching types that are compatible, but not equivalent, CoreDX DDS will, by default, coerce a match. This behavior is configurable using the DataReader QoS policy.

QoS Policy	Default Value	Description
CoreDX_RTPSReaderQosPolicy		
send_typecode (unsigned char)	1 (true)	Configure this DataReader to send (or not send) the typecode for the Topic it is subscribing to.
CoreDX_RTPSWriterQosPolicy		
send_typecode (unsigned char)	1 (true)	Configure this DataWriter to send (or not send) the typecode for the Topic it is publishing.

QoS Policy	Default Value	Description
DataReaderQoSPolicy		
type_consistency.kind	ALLOW_TYPE_ COERCION	Possible values: DISALLOW_TYPE_ COERCION, ALLOW_TYPE_ COERCION.

There are several possible reasons an application may need to disable the sending of typecodes (and therefore, removing the capability to use them for Entity matching during discovery).

Because type encoding was not well specified in the earlier DDS standards, it is possible that DDS implementations are not interoperable in matching all possible typecodes.

In addition, typecodes and typeobjects require system resources: generated code uses additional static or FLASH memory, sending and receiving typecodes uses additional dynamic or RAM memory and network resources. This is especially true of large type definitions. In extreme cases, the typecode may be too large to send over the available transport (this is especially true for low bandwidth transports).

16.4 Static Discovery

The CoreDX DDS middleware supports dynamic discovery by default. This allows DomainParticipants to discover other Participants in the same domain. Once DomainParticipants discover each other, then the contained DataReaders and Writers are dynamically matched.

The standard dynamic discovery protocol is based on UDP MULTICAST, and works effectively in a local area network, and can be configured to work through routers and other networking devices. However, there are some

situations in which UDP MULTICAST is not desirable or possible, or where Dynamic Discovery is not suitable for the application.

In order to address these situations, CoreDX DDS provides QoS support for defining remote peer DomainParticipants. An extension QoS policy is added to the DomainParticipant QoS policies. This policy, `peer_participants`, identifies a list of remote DomainParticipants with which this participant should communicate. This avoids the initial participant discovery process, and initiates direct communication between the identified participants.

This QoS policy can be updated on the fly, after the DomainParticipant is enabled, and can support an application controlled discovery mechanism. Table 16-1 shows a C language example of setting the ‘`peer_participant`’ QoS policy.

Table 16-1: Code Example of `peer_participants` QoS

Example code to pre-define Peer Participants

```
DDS_DomainParticipant dp;
DDS_ReturnCode_t      retval;
DDS_DomainParticipantQos dp_qos;
CoreDX_ParticipantLocator pl;

DDS_DomainParticipantFactory_get_default_participant_qos(&dp_qos);

/* add two 'a-priori' configured peer locators */
/* two different participant id's and two different IP addrs */
pl.participant_id = 0;
pl.participant_id_max = 1; /* try to talk to DP 0 and DP 1 at
                           * the specified IP address */
pl.participant_locator.kind = COREDX_UDPV4_LOCATOR_KIND_QOS;
memset(pl.participant_locator.addr, 0, 16);
pl.participant_locator.addr[12] = 192;
pl.participant_locator.addr[13] = 168;
pl.participant_locator.addr[14] = 1;
pl.participant_locator.addr[15] = 12;
seq_add(&dp_qos.peer_participants.value, &pl);

pl.participant_id = 1;
pl.participant_locator.kind = COREDX_UDPV4_LOCATOR_KIND_QOS;
memset(pl.participant_locator.addr, 0, 16);
pl.participant_locator.addr[12] = 192;
pl.participant_locator.addr[13] = 168;
pl.participant_locator.addr[14] = 1;
pl.participant_locator.addr[15] = 22;
seq_add(&dp_qos.peer_participants.value, &pl);
```

```
dp = DDS_DomainParticipantFactory_create_participant( 0,  
                                                    &dp_qos,  
                                                    NULL, 0 );
```

It is important to note the potential performance impacts to configuring peer_participants. In the majority of environments, using standard dynamic discovery will be more efficient with respect to CPU and network usage than specifying peers. With a configured list of peer_participants, CoreDX DDS must attempt to reach out to each and every potential DomainParticipant (that is, the full range of participant_id to participant_id_max) for each peer IP address configured. With standard multicast (or broadcast) discovery, one message will reach all potential peers.

16.5 Centralized Discovery

16.5.1 Overview

The Standard (peer-to-peer) Discovery process in CoreDX DDS is encapsulated in every CoreDX DDS application, and does not require any additional daemons or services. Each CoreDX DDS application performs the discovery process, including announcing the presence of its DDS Entities, listening for other DDS Entities, and looking for matches between its own DDS Entities and those discovered. The standard discovery mechanism is interoperable between DDS implementations.

This Standard Discovery is depicted in the below diagram.

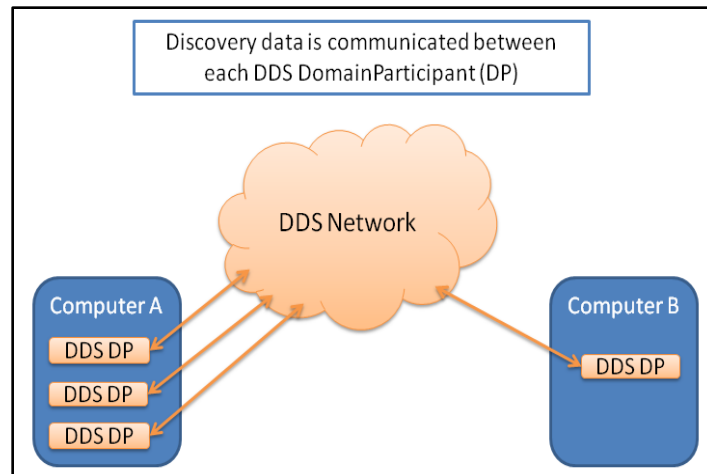


Figure 17: Standard Discovery (peer-to-peer) architecture

Despite the many benefits of the Standard Discovery mechanism, it does have some drawbacks for certain system architectures. In particular, Standard Discovery may not scale well to large DDS domains. In DDS domains with large numbers of DDS entities (Participants, Readers or Writers), the Standard Discovery mechanism can require large amounts of memory as every Participant discovers all other entities in the system. In many cases, this 'world view' of the DDS domain is wasteful. Often, a Participant is required to communicate with only a small sub-set of the entire DDS network.

To address the scalability issues of Standard Discovery, CoreDX DDS supports a specialized discovery mechanism called **Centralized Discovery**. CoreDX Centralized Discovery performs the work of discovering all DDS Entities and appropriately communicating those entities to participants based on 'need to know'. The Centralized Discovery mechanism can scale to very large DDS domains, without the explosion of memory allocation found in Standard Discovery.

Further, Centralized Discovery is designed to be interoperable with Standard Discovery. This means that a DDS domain may combine both discovery mechanisms as necessary: some Domain Participants can use Standard Discovery while others use Centralized Discovery.

This Centralized Discovery is depicted in the below diagram.

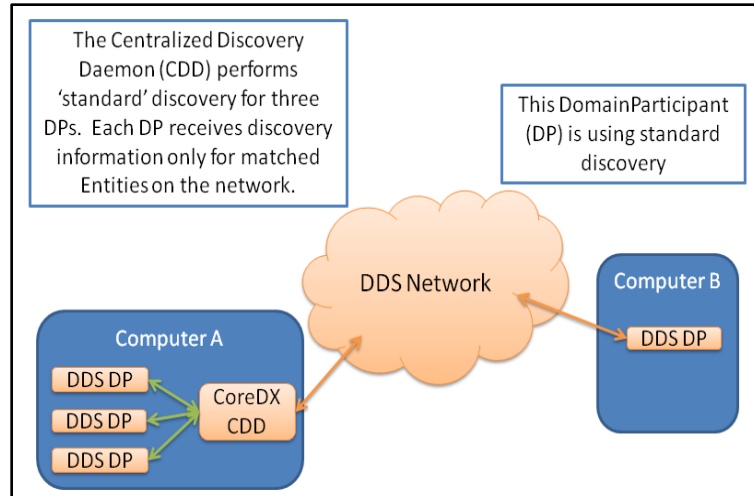


Figure 18: Centralized Discovery architecture

16.5.2 Memory Utilization and Scalability

With Standard Discovery, each DomainParticipant learns and remembers every active DomainParticipant, DataReader, and DataWriter in the DDS Domain. As the number of DDS Entities in the Domain grows, so does the amount of discovery information stored in each DomainParticipant.

For systems that contain many DDS Entities, it may be desirable to reduce the number of copies of this maintained discovery information. This is one benefit of Centralized Discovery. The discovery information about all DDS Entities in a DDS Domain is stored in a centralized location, reducing the overall memory utilization in the system.

The Centralized Discovery Daemon determines the potentialReader/Writer matches for all its connected DomainParticipants. DomainParticipants learn only about potential matches from the Centralized Discovery Daemon.

A CoreDX Centralized Discovery Daemon must reside on the same machine as DomainParticipants that are configured to use Centralized Discovery. Therefore, the greatest benefits for memory reduction are seen when:

1. There are many DDS Entities on one machine that can use Centralized Discovery, and
2. For each DomainParticipant, a small percentage of the DDS Entities in the DDS domain match with its own DDS Entities.

Note that the selection of ‘discovery’ mechanism affects only the exchange of discovery information – application data is not affected. Application data is always exchanged peer-to-peer, even when using Centralized Discovery.

16.5.3 Network Utilization and Discovery

With Standard Discovery, each DomainParticipant communicates with every active DomainParticipant in the DDS Domain. As the number of DDS DomainParticipants in the Domain grows, so does the amount of network traffic generated to communicate with each peer DomainParticipant.

For systems that contain many DomainParticipants, and at least some of these DomainParticipants are co-located on the same computer, it may be desirable to reduce the number messages generated on the network. This is an additional benefit of Centralized Discovery. The DomainParticipants co-located on a computer will communicate with their location Centralized Discovery Daemon, and only the Centralized Discovery Daemon will communicate off-box, reducing the amount of discovery network traffic.

16.5.4 Configuring Centralized Discovery

Configuring CoreDX DDS discovery happens at the DomainParticipant using a QoS policy. A DomainParticipant can be configured to use a certain discovery mechanism at creation time through a QoS policy.

By default, CoreDX DDS uses Standard Discovery (PEER). To use Centralized Discovery, change the DomainParticipant CoreDX_Discovery_Qos_Policy to specify centralized discovery. The DomainParticipant QoS policy for configuring discovery (and built-in entities) is described below.

CoreDX Centralized Discovery is compliant with the OMG’s RTPS Wire Protocol standard, and is therefore interoperable with other DDS implementations. Since DomainParticipants using Centralized Discovery can communicate with DomainParticipants using standard discovery, a mix of discovery types can be configured in the same DDS network.

QoS Policy	Default Value	Description
CoreDX_DiscoveryQosPolicy		

QoS Policy	Default Value	Description
discovery_kind (DiscoveryQoSPolicyDiscoveryKind)	DDS_PEER_DISCOVERY_QOS	Configure this DomainParticipant to use standard (PEER) discovery or centralized discovery. Possible values are: DDS_PEER_DISCOVERY_QOS and DDS_CENTRAL_DISCOVERY_QOS

16.5.5 Deploying Centralized Discovery

A CoreDX Centralized Discovery Daemon must be deployed on each computer that is hosting a DomainParticipant configured to use Centralized Discovery. There should be only one CoreDX Centralized Discovery Daemon running on a computer. Computers that are not hosting DomainParticipants configured to use Centralized Discovery do not need a CoreDX Centralized Discovery Daemon.

An example deployment using Centralized Discovery is shown below.

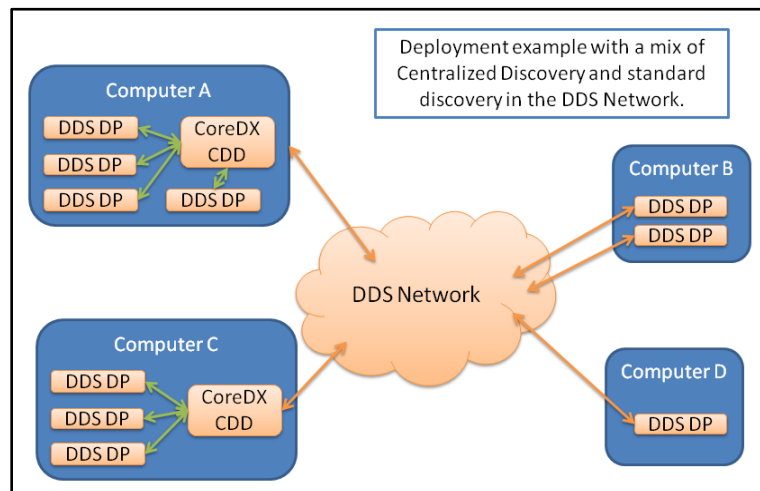


Figure 19: Example Centralized Discovery deployment

16.6 Wait for Discovery

The automatic and dynamic discovery process is defined by the DDS standards and employed by CoreDX DDS. While the discovery algorithms

are efficient, the dynamic nature of discovery means it is impossible to determine the amount of time required for discovery to complete, even when all entities are located within one DomainParticipant.

For example, consider an application with the following execution flow (and no pauses or gaps between steps):

1. Create DomainParticipant
2. Register data types and create Topics
3. Create Subscribers and DataReaders
4. Create Publishers and DataWriters
5. Write data

Discovery and matching of these local DataReaders and DataWriters may not complete before the application writes data. If the discovery and matching is not yet complete, data will not be received by the DataReaders (since they are not yet known to exist by the DataWriter).

To address this problem, CoreDX DDS provides an API on DomainParticipants to wait for built-in acknowledgements:

```
DomainParticipant::builtin_wait_for_acknowledgments(  
    Duration_t *max_wait)
```

While a DomainParticipant configured to use dynamic discovery has no way to know how many, if any remote DDS Entities may be discovered, this API will block the application until all DataReaders and DataWriters within known DomainParticipants have been discovered and matched as appropriate (or until the max_wait duration has expired).

16.7 Access to Discovery Information

Some DDS Entities provide access to a subset of discovery information via their API. Each DomainParticipant provides access to all currently known peer DomainParticipants (including itself) with the get_discovered_participants() method. Each DataReader and DataWriter can provide information about the set of currently matched DataWriters and DataReaders, respectively, with the get_matched_publications() and get_mached_subscriptions() methods.

The relevant methods are summarized in the table below.

Method	Description
--------	-------------

DomainParticipant::get_discovered_participants()	Get the list of currently discovered DomainParticipants (sequence of handles)
DomainParticipant::get_discovered_participant_data()	Get information about a specific discovered DomainParticipant (identified by handle)
DataWriter::get_matched_subscriptions()	Get the list of DataReaders currently matched to this DataWriter (sequence of handles)
DataWriter::get_matched_subscription_data()	Get information about a specific DataReader (identified by handle)
DataReader::get_matched_publications()	Get the list of DataWriters currently matched to this DataReader (sequence of handles)
DataReader::get_matched_publication_data()	Get information about a specific DataReader (identified by handle)

Table 16-2: API Methods to Access Discovery Information

Example usage:

Example C++ Code Snippet

```
#include <dds/dds.hh>
#include <dds/dds_builtin.hh>

...

DDS::InstanceHandleSeq handles;

retval = dp->get_discovered_participants( &handles );

DDS::InstanceHandle_t handle;
for ( i = 0; i < handles.size(); i++ )
{
    handle = handles[i];
    DDS::ParticipantBuiltinTopicData dpdata;
    dp -> get_discovered_participant_data( &dpdata, handle );

    ...
}
```

The discovery information available by these methods is limited to the subset that matches an existing local entity. In some cases an application would like to learn about all discovered entities, regardless of whether it matches a local entity. This can be accomplished by accessing the built-in DataReaders, as described in the Built-In Topics section.

16.8 Discovery and Deterministic Machines

CoreDX DDS discovery adheres to the DDS standards and is interoperable with other DDS implementations. Part of this interoperable discovery process is the assignment of Globally Unique Identifiers (GUIDs) to RTPS Entities that are advertised and may be discovered by other RTPS Entities, including Participants, Readers, and Writers.

The Participant GUID is important for dynamic discovery. Each Participant will generate a unique GUID and include it in the discovery messages it publishes. When a new discovery advertisement message is received, a Participant can use the GUID to determine if this is new Participant, or one that it has already discovered. Newly discovered participants will participate in additional data exchange to share QoS policy settings and information about existing DataReaders and DataWriters.

The discovery process allows an application’s Participant to go away (by normal or abnormal exit, or machine restart), restart, and seamlessly re-join

the existing DDS network as a *new* Participant. This works only if the restarting application's Participants are assigned a *unique* GUID.

According to the standard, the Participant GUID is created using the following data:

- IPv4 address of the computer hosting the Participant
- Process ID of this Participant's application
- One-up counter for each Participant within this application
- Entity Kind (fixed identifier for Participant)

For many computer systems, this algorithm does generate unique GUID's for Participants, even after machine restarts. However, applications running on deterministic Operating Systems, such as VxWorks, may always start with the same process ID, resulting in a Participant always having the same GUID. This can cause a problem when a Participant attempts to re-join DDS communications using the same GUID it had previously. Remote Participants on the DDS network will consider this Participant an already-discovered Participant, and will not participate in the necessary data exchange to share QoS policy settings and existing DataReaders and DataWriters.

To address this problem with deterministic systems, CoreDX DDS provides an additional discovery QoS policy setting for applications to use their own algorithm to set the Process ID portion of the GUID. When used, this should be a number that uniquely identifies an application on a computer, and will not be the same after a machine restart. This might be a one-up counter that is written to persistent storage (disk, writable FLASH memory) or another application defined algorithm.

QoS Policy	Default Value	Description
CoreDX_DiscoveryQosPolicy		
guid_pid (DDS_BUILTIN_TOPIC_KEY_TYP E_NATIVE)	0	A value of '0' will use the default: the application process ID (PID). Values other than '0' will be used in place of the PID in constructing the GUID.

Chapter 17 Configuring Reliability Protocol

17.1 Reliability Protocol

The CoreDX DDS Reliability protocol addresses dropped packets, out of order samples, communication disconnects, and application re-starts to ensure delivery of published data to the intended recipients. This protocol is supported by the standardized RTPS protocol and the `DataReader` and `DataWriter` Data Caches (see Chapter 10.5 for a full discussion on the Data Caches).

This reliability protocol is light-weight and minimizes latency. Dropped packets are quickly detected and repaired. The CoreDX DDS reliability algorithms are carefully designed to provide very good performance across a wide range of operating environments, and will dynamically self-tune and adjust based on current conditions.

In some cases, known network or system characteristics may result in a system that benefits from manual configuration of the CoreDX DDS reliability protocol. CoreDX DDS provides tunable parameters for configuring the reliability protocol to allow the application developer to achieve an optimal balance of overhead and timely data retransmission.

NOTE: Take care when manually adjusting CoreDX DDS Reliability parameters! In most CoreDX DDS networks, using default reliability protocols provides better performance than tuning individual timing configurations, and is much simpler to maintain over time.

17.1.1 Cache Management

The reliable protocol effects more than the handshaking between `DataWriters` and `DataReaders`. The Data Cache also plays an important role in reliable communications.

On the `DataWriter`, the Data Cache contains samples that have not been acknowledged by all reliable subscriptions. [Data may be kept longer based on the Durability QoS configuration.] The data cache size is controlled by the History and the Resource Limits QoS policies. If the data cache becomes full, the History QoS policy *kind* comes into play. With a History kind of `KEEP_ALL`, the `write()` call will block until there is space, or until the `max_blocking_time` has elapsed. With a History kind of `KEEP_LAST`, the

oldest sample will be removed from the cache to make room for a new sample.

On the DataReader, the Data Cache contains samples in order, with respect to the source DataWriter. If samples are lost, the data cache may contain out of order samples (those samples that are received are saved while waiting for the lost sample(s) to be retransmitted and received). These samples are stored in a “forward cache” and are not available to the application until all missing samples are received. This design minimizes retransmissions of data samples, while ensuring the application receives only in order data samples. The DataReader data cache size is controlled by the History and the Resource Limits QoS policies. If the data cache becomes full, the History QoS policy *kind* comes into play. With a History kind of KEEP_ALL, the DataReader will drop new incoming samples, sending a NACK to the DataWriter, and forcing the DataWriter to save and re-transmit the samples. With a History kind of KEEP_LAST, the oldest sample will be removed from the cache to make room for a new sample.

17.1.2 Heartbeats, ACKs, and NACKs

The reliability protocol relies on *Heartbeats* from the DataWriter and *ACK/NACK responses* from the DataReader. *Heartbeat* messages tell the DataReaders the data that is currently available (that has been sent) at the DataWriter. Positive *ACK* and negative *NACK Responses* are sent in response to a *Heartbeat* and confirm the DataReader has received one or more samples, and possibly requests one or more samples to be retransmitted.

Figure 17-1 shows a simple example of the network traffic (including *Heartbeat* and *ACK Responses*) when there are no dropped samples. Notice that *Heartbeats* can be sent in combination with data samples, reducing network overhead.



Figure 17-1: Example DDS Usage

In this example, the DataReader sends an ACK in response to each of the 2 Heartbeats received from the DataWriter. In the first ACK response, the DataReader confirms receipt of all available samples up to sample #3. In the second ACK response, the DataReader confirms receipt of all available samples up to sample #6.

Figure 17-2 shows a similar example, except one data sample has been lost, and must be retransmitted.

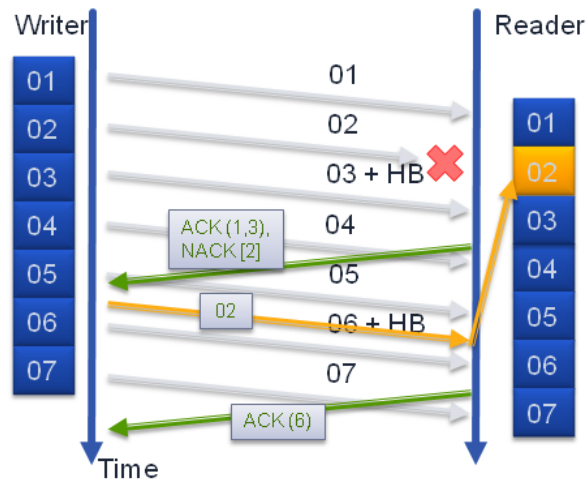


Figure 17-2: Example DDS Usage

In this example, the DataReader sends an ACK/NACK in response to the first Heartbeat from the DataWriter. The DataReader can acknowledge samples #1 and #3, but not sample #2. When the DataWriter receives the NACK, it will retransmit sample #2, sending it directly to this DataReader who needs it. In response to the next Heartbeat from the DataWriter, the DataReader can now acknowledge all data samples through #6.

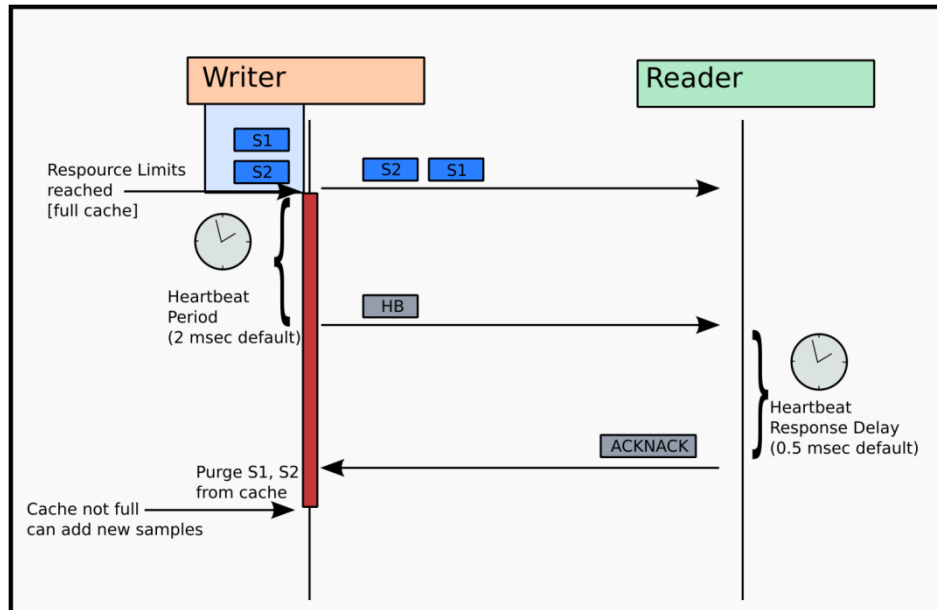
17.1.3 Periodic Data, Reliability Protocol, and Data Cache Configuration

For applications that publish periodic data (data is written at a known frequency) with `Reliability.kind = RELIABLE`, consideration should be given to the DataWriter cache size with respect to the frequency of writes, frequency of Heartbeats, and ACK/NACK response delay. The duration between Heartbeats, along with the ACK/NACK response delay provide the minimum amount of time that will elapse before sent samples may be acknowledged by a matched DataReader, allowing the DataWriter to potentially remove those samples from its cache.

Consider the example, with a DataWriter configured with the QoS:

- `Reliability.kind = DDS_RELIABLE_RELIABILITY_QOS`
- `History.kind = DDS_KEEP_ALL_HISTORY_QOS`
- `Resource_Limits.max_samples = 2`

The following diagram depicts a timeline and the DataWriter cache:



Note the time between when the first samples are written by the writing application and the time those samples are potentially acknowledged by the reading application. The DataWriter cache should be sized to allow for the number of writes() that will occur during that time period, at a minimum.

17.1.4 Unresponsive DataReaders

DataReaders that do not respond to Heartbeats with ACK/NACK messages in a timely manner pose a unique challenge to the standardized DDS Reliability protocol. A DataWriter with certain QoS policy settings will keep each published sample in its data cache until it is acknowledged by all matched Reliable DataReaders. If a DataReader becomes *unresponsive* (that is, the DataWriter is no longer receiving ACK/NACK messages from this DataReader), the DataWriter is unable to purge samples from its cache. After some period of time, the data cache may become full, or available memory will be used to store these un-acknowledged samples. At this time the application may be unable to write additional data (exact behavior will depend on other QoS policy configuration), and the other, responsive

DataReaders will not receive any additional data, all because of 1 *unresponsive* DataReader.

To address this challenge, CoreDX DDS will degrade *unresponsive* DataReaders to a Reliability configuration that not quite Reliable. The DataWriter will continue to send Heartbeats to this DataReader (in the hope that it will eventually respond with ACK/NACK messages), but will no longer expect (or wait for) samples to be acknowledged by this DataReader. In the above example, the DataWriter will remove samples from its data cache once they are acknowledged by all responsive DataReaders, allowing data communications to continue without interruption. [Note that the unresponsive DataReader may miss samples.]

When a DataReader is marked *unresponsive* is configurable by the application through the `ack_deadline` parameter (described below), and this DataReader degrading feature can be disabled by setting the `ack_deadline` to ‘infinite’.

17.1.5 ACK-less Reliable

In addition to the standardized BEST_EFFORT and RELIABLE Reliability QoS modes, CoreDX DDS provides an alternate, modified form of the Reliability QoS. This reliability mode falls between BEST_EFFORT and RELIABLE in terms of the guarantees that it makes for sample delivery. We refer to this as the **ACK-less Reliable** mode. In this mode, the DataWriter behaves very much like a ‘reliable’ writer, but it does not require positive-acknowledgements from DataReaders. Instead, it uses a time period (`reliable_sample_hold_time`) to determine how long to hold a sample before flushing it from the local cache. If a DataReader sends a negative-acknowledgment (asking for retransmission of a sample), the writer will comply if the sample is still in its local cache [this behavior is identical to the standard RELIABLE behavior]. If the Writer has flushed the sample due to expiration of the `sample_hold_time`, then the writer will respond with a ‘gap’ message indicating that the sample is no longer available [again, this is consistent with the standard RELIABLE behavior].

If the `sample_hold_time` is configured to be infinite, then the writer must have some other mechanism that allows it to remove samples from its local cache, for example: limited History depth or non-infinite Lifespan.

In exchange for the possibility of undelivered samples, there are a few subtle benefits of this reliability mode. First, the network traffic is reduced slightly by not requiring positive acknowledgements from matched readers.

Second, an unresponsive reader (one that fails to acknowledge samples) does not adversely impact the writer. In the standardized RELIABLE mode, an unresponsive reader would cause the writer to hold onto all samples until the reader has been determined to be gone. [CoreDX DDS, by default, will handle these 'unresponsive' DataReaders as described earlier in this section.] Without some form of handling of 'unresponsive' DataReaders, this may cause the writer cache to grow excessively, and if resource limits are in effect, then the writing application may find itself unable to add further samples. In this way, a single mis-behaving reader may have a negative impact on the writer and thereby all other matched readers.

17.2 Reliability QoS Configuration

Some applications need the ability to tailor the reliability protocol to achieve an optimal balance of overhead and timely data retransmission.

CoreDX DDS provides additional QoS policies to allow tailoring of the RELIABLE handshaking protocol: the `DataWriter_RTPS_WriterQoSPolicy` and the `DataReader_RTPS_ReaderQoSPolicy`. These QoS policies allow the application to tailor when and how frequently CoreDX DDS sends heartbeats, NACKs, and related responses. Configuring these items can balance latency and overhead in RELIABLE communications, and help avoid packet storms.

These additional QoS policies are described below.

Table 17-1: CoreDX DDS RTPS_Protocol QoS Policy

QoS Policy	Default Value	Description
DataWriter_RTPSWriterQoSPolicy		
heartbeat_period (Duration_t)	2 ms	The duration between Heartbeats sent by the DataWriter.
nack_response_delay (Duration_t)	1 us	The delay between the receipt of a NACK Response and the DataWriter's retransmission of data.

QoS Policy	Default Value	Description
nack_suppress_delay (Duration_t)	0 ns	Not yet used by CoreDX DDS
ack_deadline (Duration_t)	INFINITE	The amount of time the DataWriter should wait for an ACK/NACK message from a DataReader before it is considered <i>unresponsive</i> .
require_acks	1 (true)	Applicable when the DataWriter is configured with Reliability.kind = RELIABLE. When set to ‘TRUE’, the DataWriter uses standard RELIABLE Reliability. When set to ‘FALSE’, the DataWriter uses ACK-less Reliability.
reliable_sample_hold_time	INFINITE	Applicable when require_acks=0 (false). The amount of time an ACK-less DataWriter will hold samples before removing them from its cache (as if they have been ACK’d by matched Reliable DataReaders).
DataReader_RTSPSReaderQoSPolicy		
heartbeat_response_delay (Duration_t)	500 us	The delay between the receipt of a Heartbeat and the DataReader’s ACK/NACK Response.
send_initial_nack	1 (true)	This setting is applicable only to Reliable DataReaders. When set to non-zero (true), the DataReader will send a “NACK” message to each newly discovered DataWriter, jumpstarting the handshaking process to receive any data the DataWriter has to publish. One possible reason to disable this

QoS Policy	Default Value	Description
		'jumpstart', is performance: with lots of DataReaders matching with one DataWriter at the same time.
CoreDX_DiscoveryQosPolicy		
heartbeat_period (Duration_t)	10 ms	For built-in writers, the delay between the receipt of a NACK Response and the DataWriter's retransmission of data.
nack_response_delay (Duration_t)	1 us	For built-in writers, the delay between the receipt of a NACK Response and the DataWriter's retransmission of data.
nack_suppress_delay (Duration_t)	0 sec	No yet used by CoreDX DDS
heartbeat_response_delay (Duration_t)	0 sec	For built-in readers, the delay between the receipt of a Heartbeat and the DataReader's ACK/NACK Response.
send_initial_nack	1 (true)	When set to non-zero (true), the built-in DataReaders will send a "NACK" message to each newly discovered built-in DataWriter, jumpstarting the handshaking process to receive any discovery data the DataWriter has to publish. This may need to be disabled to support interoperability (in our testing to date, one DDS vendor did not support this option).

Chapter 18 Configuring Reader Specific Locators

18.1 Overview

Some network environments can benefit from having a subset of traffic isolated to a specific multicast address (different from other traffic).

DataReaders and DataWriters in CoreDX DDS inherit the 'default' multicast and unicast locators from the DomainParticipant. If no specific locators are specified, then the entities will use the defaults. If a specific multicast locator is specified, it will be used in place of any default multicast locators. If a specific unicast locator is specified, it will be used in place of any default unicast locators.

The default UDP IPv4 multicast address used for user data is 239.255.0.1.

By configuring a DataReader to accept data on a specific address (Locator), then all DataWriters matched with that DataReader will send data to that specific address.

18.2 Configuration

The specific locators are configured in the **DataReaderQos.rtps_reader.locators** QoS policy. This field is a sequence of CoreDX_Locators. A CoreDX_Locators has the following structure:

```
typedef struct CoreDX_Locator_t {  
    int          kind;  
    uint32_t     port;  
    unsigned char addr[COREDX_LOCATOR_ADDR_LEN];  
} CoreDX_Locator;
```

To define a specific multicast locator in 'C' for the UDP transport, use the following:

```
CoreDX_Locator l;  
DDS_DataReaderQos dr_qos;  
  
DDS_Subscriber_get_default_datareader_qos(sub,  
&dr_qos);  
  
memset(&l, 0, sizeof(l));
```

```

l.kind = COREDX_UDPV4_LOCATOR_KIND_QOS;
l.port = 7400;
l.addr[12] = 239;
l.addr[13] = 255;
l.addr[14] = 0;
l.addr[15] = 2;
seq_add( &dr_qos.rtps_reader.locators, &l );

```

To define a specific multicast locator in 'C++' for the UDP transport, use the following:

```

sub->get_default_datareader_qos( dr_qos );

DDS::Locator l;
l.kind = COREDX_UDPV4_LOCATOR_KIND_QOS;
l.port = 7400;
memset( l.addr, 0, COREDX_LOCATOR_ADDR_LEN );
l.addr[12] = 239;
l.addr[13] = 255;
l.addr[14] = 0;
l.addr[15] = 2;
dr_qos.rtps_reader.locators.push_back( l );

```

CoreDX DDS supports setting specific UDP MULTICAST locators on the DataReader. Other combinations (other transports, or unicast locators) are not supported at this time.

Chapter 19 Dynamic Types

19.1 Overview

The original DDS standards were designed around the assumption that types associated with Topics are known at compile time. While this architecture provides better communication performance (throughput and latency) and better type safety, static data typing makes it difficult to dynamically define Topics at run time.

The addition of Dynamic Types to the CoreDX DDS baseline allows greater flexibility to application developers. Developers can define their data types at compile time or discover data types at run time. Once a data type is discovered, the application can dynamically create DataReaders to receive Topic data and use introspection to parse and process received data. DataWriters can also be created to write the discovered data type.

The dynamic type technology can also benefit applications that will be deployed to some very space constrained environments. Using Dynamic Types can help reduce the code size by reducing or eliminating the type specific code generated from IDL or XML.

The Dynamic Type API is fully defined in the CoreDX DDS Type System Programmer's Guide.

19.2 Subscribe with Dynamic Types

A CoreDX DDS application may subscribe to a Topic that is discovered at run time, without any knowledge of the data type associated with the Topic. The basic steps involved in this type of application are:

1. Use the built-in DataReader to Discover a DataWriter
2. Use the Type Object or Type Code information from the DataWriter to register a Data Type
3. Use the topic information from the DataWriter to create a Topic
4. Create a Dynamic DataReader
5. Read data

An example for subscribing with Dynamic Types can be found in the examples directory of the CoreDX DDS package.

19.3 Publish with Dynamic Types

A CoreDX DDS application may publish to a Topic that is discovered at run time, without any knowledge of the data type associated with the Topic. This is possible with the Dynamic Types features and API, although not a common use case for Dynamic Type. The basic steps involved in this type of application are:

1. Create a dynamic data type representing the type of data to be published
2. Register the dynamic data type
3. Create a Topic
4. Create a Dynamic DataWriter
5. Initialize and send data

An example for publishing with Dynamic Types can be found in the examples directory of the CoreDX DDS package.

Chapter 20 Threading Options

20.1 Overview

CoreDX DDS contains advanced multi-threaded technology. This feature allows any application (even non-threaded applications) running on multi-core hardware to make use of multiple cores. Using multiple cores on multi-core hardware provides significant performance benefits, as confirmed using Intel’s Thread Checking benchmarking system.

Because many of our users are not running on multi-core hardware, and in fact are running on significantly reduced-power single core hardware, CoreDX DDS is configurable to provide performance benefits in this type of resource constrained environments.

20.2 Configuring Threading Options

CoreDX DDS runs in an optimized threaded model by default. This mode creates three threads for each DomainParticipant created in the application:

1. The main (application) thread

The main (application) thread contains the main() of the application, and most of the application execution.

2. The CoreDX DDS reading thread

The CoreDX DDS reading thread is responsible for reading data off the transport (UDP socket, or other transport as configured by the application). Data is read off the transport, unmarshaled, and put into the Data Caches of the appropriate DataReaders.

3. The CoreDX DDS work thread

The CoreDX DDS work thread performs all remaining DDS “work”. This includes discovery, writing application data to the transport, maintaining liveliness, performing handshaking and any necessary repairs for Reliable readers and writers, and application notification of events.

CoreDX DDS threading is configurable through the DomainParticipant’s CoreDX_ThreadModelQosPolicy.

20.2.1 Single Threaded Configuration

CoreDX DDS features a *single threaded* model for higher performance on significantly resource constrained single threaded devices. Eliminating threads allows CoreDX DDS to eliminate much of the locking code, and reduce the amount of context switches required by the application, helping to reduce the required memory and CPU resources.

The single threaded model requires the application to periodically “hand over” CPU time to CoreDX DDS to perform its work (discovery, reading data, writing data, maintaining liveliness, etc.). Otherwise, this single threaded model uses the same API as the multi-threaded model. There are no new libraries to link with. This ensures there is not a completely new API to learn, and makes it easy to move applications from multi-threaded to single-threaded modes (and back again).

The CoreDX DDS release packages contain example code illustrating the use of the single threaded mode. This example can be found in the examples directory, “hello_nothr”. If you look at hello_pub.c in this example, you will find that the 'no threads' programming model has a very simple API.

First, use the CoreDX_ThreadModelQoSPolicy on the DomainParticipant to configure CoreDX DDS to use the single threaded model.

QoS Policy	Default Value	Description
CoreDX_ThreadModelQoSPolicy		
use_threads (unsigned char)	1 (true)	Setting use_threads to 0 (false) will configure CoreDX DDS to use the single threaded model.

Next, provide CoreDX DDS with time to perform work. It is important to provide CoreDX DDS with enough opportunities to run so that it can manage its internal tasks. This entails inserting calls to DDS_DomainParticipant_do_work() at strategic points in your 'main' program.

```
DomainParticipant::do_work( duration_t time )
```

The application provides CoreDX DDS a duration in which it can perform its internal work. The `do_work` call will return when this time expires.

20.2.1.1 Single Threaded Mode and DDS Blocking Calls

When a DataWriter method blocks (for example, `DataWriter::write()` when configured with appropriate QoS settings to allow it to block), `do_work()` will be automatically excuted for the duration of the blocking time.

20.2.1.2 Single Threaded Mode and WaitSets

WaitSets in a single threaded mode are not supported.

20.2.2 Listener Thread Configuration

CoreDX DDS can create a 4th thread (only applicable when using the multi threaded model) that will handle application listener callbacks.

In the standard 3-thread threaded model, application listener callbacks are handled by the work thread (along with discovery, writing data, and maintaining liveliness). This means that a long-running application defined listener callback (for example, in a DataReader’s `on_data_available` listener callback) can block CoreDX DDS from performing other internal tasks.

The solution for those applications that cannot reduce their listener callback functions is to create a separate thread for listener callbacks. With the listener thread enabled, an application can block inside a listener callback without effecting other DDS operations.

QoS Policy	Default Value	Description
CoreDX_ThreadModelQosPolicy		
<code>create_listener_thread</code> (unsigned char)	0 (false)	Setting <code>create_listener_thread</code> to 1 (true) will configure the DomainParticipant to create the 4 th thread for application callbacks. This option is applicable only when <code>use_threads</code> is set to non-zero (true).

Chapter 21 Transmit Buffers

21.1 Overview

Each CoreDX DDS DataWriter contains a *transmit buffer* to hold data that is waiting to be published.

Transmit buffers usually hold sample data from `DataWriter::write()` calls, but can also hold instance lifecycle information including unregister or dispose actions. Both built-in DataWriters and application defined DataWriters contain these transmit buffers.

By default, transmit buffers are dynamic, that is, they grow and shrink as necessary to minimize the amount of memory consumed by the CoreDX DDS infrastructure. CoreDX DDS transmit buffers can be configured to be a static size, or configured to be dynamic with specified minimum and maximum sizes.

21.2 Dynamic Transmit Buffers

CoreDX DDS transmit buffers are by default, dynamic. Without any configuration, DataWriter transmit buffers will grow and shrink as necessary to support the size of data written while consuming a minimal amount of memory.

CoreDX DDS transmit buffer sizes can be configured with a QoS policy, or with environment variables. The DataWriter QoS policy to configure the minimum and maximum buffer sizes is described in the following table.

Table 21-1: Transmit Buffer Configuration

QoS Policy	Default Value	Description
DataWriter RTPSWriterQoSPolicy (application-defined DataWriters)		
min_buffer_size (unsigned int)	16	In bytes, the transmit buffer will start at this size, and in dynamic operations, will not shrink smaller

than this size.		
max_buffer_size (unsigned int)	65400	In bytes, the transmit buffer will not grow larger than this size.
DataWriter_RTPSWriterQoSPolicy (built-in DataWriters)		
min_buffer_size (unsigned int)	16	For built-in writers, the transmit buffer will start at this size, and in dynamic operations, will not shrink smaller than this size (in bytes).
max_buffer_size (unsigned int)	32768	For built-in writers, the transmit buffer will not grow larger than this size (in bytes).

The CoreDX DDS environment variables to configure the transmit buffer size are: COREDX_MIN_TX_BUFFER_SIZE and COREDX_MAX_TX_BUFFER_SIZE, and are used in the same manner as the DataWriter QoS policy described above. These environment variables will override the default sizes of all DDS entities (both built-in and application defined).

These are the sizes that bound the dynamic sizing of the buffers. The transmit buffer will not grow beyond max_buffer_size, and the transmit buffer will not shrink below min_buffer_size.

The maximum transmit buffer size will affect how CoreDX DDS aggregates, batches, and/or fragments written data. This is particularly noticeable with applications that perform many, frequent writes, or have bursts of writes. With a small upper bound on the transmit buffer, CoreDX DDS will need to perform many individual writes, rather than aggregating or batching samples together to be sent at one time.

When the application writes a sample that is larger than the configured maximum transmit buffer size for the DataWriter, CoreDX DDS will fragment the data sample as necessary to fit the transmit buffer, and re-assemble the sample on the receiving side before it is available to the receiving application.

The environment variable: COREDX_MAX_PACKET_SIZE (available in earlier CoreDX DDS releases) was equivalent to COREDX_MAX_TX_BUFFER_SIZE. The MAX_PACKET_SIZE environment variable is deprecated, and should no longer be used.

21.3 Static Transmit Buffers

Since allocating and de-allocating memory can be expensive operations, applications interested in very low latencies may benefit from a static transmit buffer that does not grow or shrink through the life of the application. This configuration is possible by setting the min_buffer_size and max_buffer_size to the same value, using either the QoS policies or environment variables described above.

Special care should be taken before setting a static transmit buffer size. Since the transmit buffer must be large enough to hold the complete marshaled data sample, it is important to understand the possible sizes for all possible application data samples written by the application. This is especially true when globally configuring static built-in DataWriter transmit buffers using the environment variables.

Chapter 22 Receive Buffers

22.1 Overview

Each CoreDX DDS DomainParticipant contains a *receive buffer* to hold incoming data that will be passed to its DataReaders and eventually the reading application.

Receive buffers are used to hold data read from the transport until it is processed (parsed) by the CoreDX DDS infrastructure. There is one receive buffer for each DomainParticipant (as opposed to one for each DataReader). Receive buffers are dynamic, that is, they grow and shrink as necessary to minimize the amount of memory consumed by the CoreDX DDS infrastructure.

22.2 Configuration

The application may set the minimum size of the DomainParticipant’s receive buffer. This is commonly used when interoperating with other DDS products, who don’t effectively communicate RTPS message sizes. CoreDX DDS will “learn” how to size the receive buffer, but it may be forced to drop some of the initial messages received during this learning period. Starting with a sufficiently large receive buffer can optimize performance in these mixed DDS product environments.

The application may set the maximum size of the DomainParticipant’s receive buffer, which limits the size the receive buffer may grow to. Special care should be taken when configuring the maximum receive buffer size. If the configured maximum size is smaller than the transport’s maximum message size, CoreDX DDS may be forced to drop received samples.

The minimum and maximum size configuration is located in the transport configuration – see *Transport Configuration* for additional information.

By default, CoreDX DDS includes information about the RTPS message length in each RTPS message. This is a CoreDX DDS proprietary submessage (processed only by CoreDX DDS) that can/should be ignored by other DDS implementations. CoreDX DDS uses this submessage to optimize the memory usage of incoming buffers in the CoreDX DDS Transport[s].

The Discovery QoS Policy provides a configuration item to control whether or not this submessage is used:

QoS Policy	Default Value	Description
Discovery QoS Policy		
send_msg_submsg (unsigned char)	TRUE	When TRUE, The RTPS layer will insert a 'MsgLen' submessage in every RTPS Message.

Chapter 23 Data Batching

Data Batching is the processes of combining data samples into one RTPS message in order to reduce the network overhead and improve throughput, especially with smaller samples.

By default, data batching is disabled in CoreDX DDS DataWriters. DataReaders are configured to accept batch messages by default. Applications can use the following QoS policies to configure *data batching*.

QoS Policy	Default Value	Description
CoreDX_RTPSWriterQoSPolicy		
enable_batch_msg (unsigned char or boolean)	0 (or false)	Configure the DataWriter to use batching. Possible values are: 0 or 1 (false or true)
CoreDX_RTPSReaderQoSPolicy		
accept_batch_msg (unsigned char or boolean)	1 (or true)	Configure the DataReader to use batching. Possible values are: 0 or 1 (false or true)

Chapter 24 QoS Provider

24.1 Introduction

The standardized DDS API provides methods to get and set QoS policies on various DDS entities. Managing QoS policy settings programmatically can prove to be a burden. Some application deployments benefit from the option of specifying QoS settings through an external configuration.

To this end, the **QoS Provider** API provides a mechanism to load QoS policy settings from an XML document. The QoS Provider then provides access to the various QoS policies specified in that document, identified by name.

24.2 QoS Provider Usage

The QoS Provider API includes methods to 'get' a specific set of QoS policy values for a DDS Entity. This QoS can be used in the creation of a DDS Entity or passed to a `set_default_XYZ_qos()` or `set_qos()` call. Once the application has finished with the QoS, it should be returned to the QoS Provider by means of the appropriate 'return' call on the QoSProvider API.

24.2.1 C Language Binding

First the application must include the appropriate header file: "dds/qos_provider.h". Then, the application can instantiate a QoS Provider by means of the function **DDS_QosProvider_newQosProvider()**, passing the URI to the XML document, the name of the profile to be selected from the document, and the XML parser implementation. CoreDX DDS includes a light-weight XML parser that can be used (available in the global name `cdx_xml_impl`), alternatively the application can provide a parser implementation. For example:

```
DDS_QosProvider *qos_provider =
DDS_QosProvider_newQosProvider("file://test_qos.xml",
"profile_1",
&cdx_xml_impl );
```

Then, the `qos_provider` object can be queried to obtain specific QoS policies by name; for example:

```
dp_qos = DDS_QosProvider_get_participant_qos(qos_provider,
"dp_1");
```

```
if (dp_qos) {  
    /* use dp_qos */  
    DDS_QosProvider_return_participant_qos(qos_provider, dp_qos);  
} else { /* error */ }
```

Finally, once the application is finished with the QoS Provider, it must be destroyed with the following:

```
DDS_QosProvider_return_provider( qos_provider );
```

24.2.2 C++ Language Binding

First the application must include the appropriate header file: “dds/qos_provider.hh”. Then, the application can instantiate a QoS Provider by means of its constructor, passing the URI to the XML document, the name of the profile to be selected from the document, and the XML parser implementation. CoreDX DDS includes a light-weight XML parser that can be used (available in the global name **cdx_xml_impl**, alternatively the application can provide a parser implementation.

For example:

```
DDS::QosProvider *  
    qos_provider = new DDS::QosProvider( "file://test_qos.xml",  
    "profile_1", &cdx_xml_impl);
```

Now, the qos_provider object can be queried to obtain specific QoS policies by name; for example:

```
DDS::DomainParticipantQos * dp_qos;  
dp_qos = qos_provider->get_participant_qos("dp_1");  
if ( dp_qos ) {  
    // use the dp_qos policy  
    qos_provider->return_participant_qos(dp_qos);  
} else { /* error */ }
```

Finally, once the application is finished with the QoS Provider, it must be destroyed with the following:

```
delete qos_provider;
```

24.2.3 C# Language Binding

First the application must 'use' the CoreDX DDS package, for example: **use com.toc.coredx.DDS**; Then, the application can instantiate a QoS Provider by means of its constructor, passing the URI to the XML document, the name of the profile to be selected from the document. CoreDX DDS includes a light-weight XML parser that will be used.

For example:

```
QosProvider qosProvider = new QosProvider("file://test_qos.xml",
"profile_1");
```

Now, the qos_provider object can be queried to obtain specific QoS policies by name; for example:

```
dp_qos = qosProvider.get_participant_qos("dp_1");

if ( dp_qos != null ) {
    // use the dp_qos policy
} else { /* error */ }
```

Finally, once the application is finished with the QoS Provider, it can be destroyed with the following:

```
qosProvider.cleanup;
```

24.2.4 Java Language Binding

First the application must import the CoreDX DDS package, for example: **import com.toc.coredx.DDS.***; Then, the application can instantiate a QoS Provider by means of its constructor, passing the URI to the XML document, the name of the profile to be selected from the document. CoreDX DDS includes a light-weight XML parser that will be used.

For example:

```
QosProvider qosProvider = new
QosProvider("file://test_qos.xml", "profile_1");
```

Now, the `qos_provider` object can be queried to obtain specific QoS policies by name; for example:

```
DomainParticipantQos dp_qos =
qosProvider.get_participant_qos("dp_1");

if ( dp_qos != null ) {
    // use the dp_qos policy
} else { /* error */ }
```

Finally, once the application is finished with the QoS Provider, it can be destroyed with the following:

```
qosProvider.cleanup;
```

24.3 QoS Provider XML Syntax

The syntax for the XML QoS Policy file is fairly straight-forward, and is defined in the OMG's **DDS Consolidated XML Syntax Specification**. The syntax is summarized here with examples.

Each individual QOS policy is mapped to an XML element. The contents of which define the value of that policy.

The XML representation of QoS policy values is a mapping of the corresponding IDL types. For example the enumeration **HistoryQosPolicyKind** looks like this in IDL:

```
enum HistoryQosPolicyKind {
    KEEP_LAST_HISTORY_QOS,
    KEEP_ALL_HISTORY_QOS
};
```

These values are represented by a string with the name of the enumeration literals, for example:

```
<kind>KEEP_ALL_HISTORY_QOS</kind>
```

For pre-defined integral values, a string with the name of the IDL constant is used, for example:

```
const long LENGTH_UNLIMITED = -1;
```

is used in XML like this:

```
<depth>LENGTH_UNLIMITED</depth>
```

Putting this together, the History QoS policy is specified like this:

```
<history>
    <kind>KEEP_LAST_HISTORY_QOS</kind>
    <depth>10</depth>
</history>
```

QoS policies are collected together to specify one of the entity QoS structures: ParticipantQos, PublisherQos, SubscriberQos, TopicQos, DataReaderQos, and DataWriterQos. These are encoded respectively in the following XML elements: <domain_participant_qos>, <publisher_qos>, <subscriber_qos>, <topic_qos>, <datareader_qos>, <datawriter_qos>.

These XML elements are contained within a **<qos_profile>** element. Each profile element may be assigned a name, and that name can be used to select a certain profile from the **QoS Provider** API.

A sequence of **qos_profile** elements can be held within a **<qos_library>** element. All of this should be held within a **<dds>** element.

24.3.1 XML Example

The following XML provides example syntax for all standard QoS policies:

```
<dds>
  <qos_library>
    <qos_profile name="default">
      <publisher_qos name="pub_default">
        <entity_factory>
          <autoenable_created_entities>>false</autoenable_created_entities>
        </entity_factory>
      <group_data>
        <value>'A',' ','G','R','O','U','P'</value>
      </group_data>
```

```
<partition>
<name>
<element>partition_1</element>
<element>partition_2</element>
<element>p*</element>
</name>
</partition>
<presentation>
<access_scope>TOPIC_PRESENTATION_QOS</access_scope>
<coherent_access>TRUE</coherent_access>
<ordered_access>FALSE</ordered_access>
</presentation>
</publisher_qos>
</qos_profile>

<qos_profile name="profile_1">
<participantfactory_qos name="dpf_1">
<entity_factory>
<autoenable_created_entities>>false</autoenable_created_entities>
</entity_factory>
</participantfactory_qos>

<participant_qos name="dp_1">
<entity_factory>
<autoenable_created_entities>>false</autoenable_created_entities>
</entity_factory>
<user_data>
<value>'A', '
','P','a','r','t','i','c','i','p','a','n','t'</value>
</user_data>
<peer_participants>
<value>
<participant_locator>
<participant_id>0</participant_id>
<participant_id_max>10</participant_id_max>
<locator>
<kind>UDPV4_LOCATOR_KIND_QOS</kind>
<port>7410</port>
<addr>10,0,0,5</addr>
</locator>
</participant_locator>
</value>
<strict_match>TRUE</strict_match>
</peer_participants>
<logging>
<flags>0x0040</flags>
</logging>
</participant_qos>
```

```

<topic_qos name="top_1">
<deadline>
<period>
<sec>10</sec>
<nanosec>500000000</nanosec>
</period>
</deadline>
<destination_order>
<kind>BY_SOURCE_TIMESTAMP_DESTINATIONORDER_QOS</kind>
</destination_order>
<durability>
<kind>TRANSIENT_LOCAL_DURABILITY_QOS</kind>
</durability>
<durability_service>
<service_cleanup_delay>
<sec>20</sec>
<nanosec>500000000</nanosec>
</service_cleanup_delay>
<history_kind>KEEP_ALL_HISTORY_QOS</history_kind>
<history_depth>-1</history_depth>
<max_samples>LENGTH_UNLIMITED</max_samples>
<max_instances>2</max_instances>
<max_samples_per_instance>3</max_samples_per_instance>
</durability_service>
<history>
<kind>KEEP_ALL_HISTORY_QOS</kind>
<depth>10</depth>
</history>
<latency_budget>
<duration>
<sec>22</sec>
<nanosec>0</nanosec>
</duration>
</latency_budget>
<lifespan>
<duration>
<sec>23</sec>
<nanosec>0</nanosec>
</duration>
</lifespan>
<liveliness>
<kind>MANUAL_BY_PARTICIPANT_LIVELINESS_QOS</kind>
<lease_duration>
<sec>24</sec>
<nanosec>0</nanosec>
</lease_duration>
</liveliness>
<ownership>
<kind>EXCLUSIVE_OWNERSHIP_QOS</kind>

```



```
</ownership>
<reliability>
<kind>RELIABLE_RELIABILITY_QOS</kind>
<max_blocking_time>
<sec>25</sec>
<nanosec>0</nanosec>
</max_blocking_time>
</reliability>
<resource_limits>
<max_samples>100</max_samples>
<max_instances>101</max_instances>
<max_samples_per_instance>102</max_samples_per_instance>
</resource_limits>
<topic_data>
<value>'A',' ','T','o','p','i','c'</value>
</topic_data>
<transport_priority>
<value>501</value>
</transport_priority>
</topic_qos>

<publisher_qos name="pub_1">
<entity_factory>
<autoenable_created_entities>>false</autoenable_created_entities>
</entity_factory>
<group_data>
<value>'A',' ','P','u','b'</value>
</group_data>
<partition>
<name></name>
</partition>
<presentation>
<access_scope>TOPIC_PRESENTATION_QOS</access_scope>
<coherent_access>TRUE</coherent_access>
<ordered_access>TRUE</ordered_access>
</presentation>
</publisher_qos>

<subscriber_qos name="sub_1">
<entity_factory>
<autoenable_created_entities>>false</autoenable_created_entities>
</entity_factory>
<group_data>
<value>'A',' ','S','u','b'</value>
</group_data>
<partition>
<name>
<element>partition_1</element>
```

```

<element>partition_2</element>
<element>p*</element>
</name>
</partition>
<presentation>
<access_scope>TOPIC_PRESENTATION_QOS</access_scope>
<coherent_access>TRUE</coherent_access>
<ordered_access>TRUE</ordered_access>
</presentation>
</subscriber_qos>

<datawriter_qos name="dw_1">
<deadline>
<period>
<sec>10</sec>
<nanosec>500000000</nanosec>
</period>
</deadline>

<destination_order>
<kind>BY_SOURCE_TIMESTAMP_DESTINATIONORDER_QOS</kind>
</destination_order>

<durability>
<kind>TRANSIENT_LOCAL_DURABILITY_QOS</kind>
</durability>

<durability_service>
<service_cleanup_delay>
<sec>20</sec>
<nanosec>500000000</nanosec>
</service_cleanup_delay>
<history_kind>KEEP_ALL_HISTORY_QOS</history_kind>
<history_depth>-1</history_depth>
<max_samples>LENGTH_UNLIMITED</max_samples>
<max_instances>2</max_instances>
<max_samples_per_instance>3</max_samples_per_instance>
</durability_service>

<history>
<kind>KEEP_ALL_HISTORY_QOS</kind>
<depth>10</depth>
</history>

<latency_budget>
<duration>
<sec>22</sec>
<nanosec>0</nanosec>
</duration>
</latency_budget>

```

```
<lifespan>
<duration>
<sec>23</sec>
<nanosec>0</nanosec>
</duration>
</lifespan>

<liveliness>
<kind>MANUAL_BY_PARTICIPANT_LIVELINESS_QOS</kind>
<lease_duration>
<sec>24</sec>
<nanosec>0</nanosec>
</lease_duration>
</liveliness>

<ownership>
<kind>EXCLUSIVE_OWNERSHIP_QOS</kind>
</ownership>

<ownership_strength>
<value>99</value>
</ownership_strength>

<reliability>
<kind>BEST_EFFORT_RELIABILITY_QOS</kind>
<max_blocking_time>
<sec>42</sec>
<nanosec>0</nanosec>
</max_blocking_time>
</reliability>

<resource_limits>
<max_samples>100</max_samples>
<max_instances>101</max_instances>
<max_samples_per_instance>102</max_samples_per_instance>
</resource_limits>

<transport_priority>
<value>501</value>
</transport_priority>

<user_data>
<value>'A',' ','W','r','i','t','e','r'</value>
</user_data>

<writer_data_lifecycle>
<autodispose_unregistered_instances>false
</autodispose_unregistered_instances>
</writer_data_lifecycle>
```

```

</datawriter_qos>

<datareader_qos name="dr_1">
<deadline>
<period>
<sec>10</sec>
<nanosec>500000000</nanosec>
</period>
</deadline>

<destination_order>
<kind>BY_SOURCE_TIMESTAMP_DESTINATIONORDER_QOS</kind>
</destination_order>

<durability>
<kind>TRANSIENT_LOCAL_DURABILITY_QOS</kind>
</durability>

<history>
<kind>KEEP_ALL_HISTORY_QOS</kind>
<depth>21</depth>
</history>

<latency_budget>
<duration>
<sec>22</sec>
<nanosec>0</nanosec>
</duration>
</latency_budget>

<liveliness>
<kind>MANUAL_BY_PARTICIPANT_LIVELINESS_QOS</kind>
<lease_duration>
<sec>23</sec>
<nanosec>0</nanosec>
</lease_duration>
</liveliness>

<ownership>
<kind>EXCLUSIVE_OWNERSHIP_QOS</kind>
</ownership>

<reader_data_lifecycle>
<autopurge_nowriter_samples_delay>
<sec>24</sec>
<nanosec>0</nanosec>
</autopurge_nowriter_samples_delay>
<autopurge_disposed_samples_delay>
<sec>25</sec>

```

```
<nanosec>0</nanosec>
</autopurge_disposed_samples_delay>
</reader_data_lifecycle>

<reliability>
<kind>RELIABLE_RELIABILITY_QOS</kind>
</reliability>

<resource_limits>
<max_samples>100</max_samples>
<max_instances>101</max_instances>
<max_samples_per_instance>102</max_samples_per_instance>
</resource_limits>

<time_based_filter>
<minimum_separation>
<sec>27</sec>
<nanosec>0</nanosec>
</minimum_separation>
</time_based_filter>

<user_data>
<value>'A',' ','R','e','a','d','e','r'</value>
</user_data>

</datareader_qos>

<!-- inherit from previous, override 'deadline' -->
<datareader_qos name="dr_2" base_name="dr_1">
<deadline>
<period>
<sec>20</sec>
<nanosec>500000000</nanosec>
</period>
</deadline>
</datareader_qos>

<publisher_qos name="pub_2" base_name="pub_default">
<group_data><value>'A',' '
','G','R','O','U','P','2'</value></group_data>
</publisher_qos>

<publisher_qos name="pub_3" base_name="default">
</publisher_qos>
</qos_profile>
</qos_library>
</dds>
```

24.4 QoS Provider Integration

The **QoS Provider** implementation is supported in the C, C++, C# and Java language bindings.

24.4.1 C Integration

The “C” QoS Provider API is provided in the header file *include/dds/dds_qos_provider.h*. Under Linux (or other Unix-like operating systems), the binary code is available in *lib/libdds_qos_provider.a* [or *.so*]. For Windows, the binary is available in *lib/dds_qos_provider_static.lib* or *lib/dds_qos_provider.dll*.

24.4.2 C++ Integration

The C++ QoS Provider API is provided in the header file *include/dds/dds_qos_provider.hh*. Under Linux (or other Unix-like operating systems), the binary code is available in *lib/libdds_cpp_qos_provider.a* [or *.so*]. For Windows, the binary is available in *lib/dds_cpp_qos_provider_static.lib* or *lib/dds_cpp_qos_provider.dll*.

24.4.3 C# Integration

The QoS Provider implementation is included in the CoreDX DDS C# library, and is available with no further configuration.

24.4.4 Java Integration

The QoS Provider implementation is included in the CoreDX DDS Java library, and is available with no further configuration.

24.5 Unsupported

The following features are documented in the OMG's **DDS Consolidated XML Syntax Specification**, but are not currently supported by CoreDX DDS:

- **base64** encoding for octet sequences
- **topic_filter** for DataReader and DataWriter QoS policy selection
- QoS Profile shorthand (omitting <qos_profile> element)

Chapter 25 Licensing

CoreDX DDS uses development and run-time licenses. A development license is required for using the CoreDX DDS data type compiler (coredx_ddl). A run-time license is required for running an application built with the CoreDX DDS library. Both run-time and development licenses can be contained in the same license file or in separate files. Here is an example license file containing both development and run-time licenses:

coredx.lic

```
#=====
# CoreDX DDS License file for CompanyX
#
# Created: Jul 22, 2008, by Twin Oaks Computing, Inc.
# Contains: Development and run-time licenses
#
#=====
#
# development LICENSE lines:
#   - Contain your development keys - DO NOT EDIT!
LICENSE PRODUCT=coredx_ddlBUILD=Release OS=linux ARCH=x86 USERID=ntucker
HOSTID=00195b70c3be CUSTOMER=Company_X SIG=abcdefghijklmnopqrstuvwxy
#
# run-time LICENSE lines:
#
#   - Contain your run-time keys - DO NOT EDIT!
LICENSE PRODUCT=coredx_cBUILD=Release HOSTID=00195b70c3be CUSTOMER=Company_X
SIG=abcdefghijklmnopqrstuvwxy
```

Figure 25-1: Example CoreDX DDS license file

25.1 Development Licenses

Development licenses are contained in a license file. These development license keys are required by the CoreDX DDS data type compiler. To develop (compile) with CoreDX DDS, an environment variable `TWINOAKS_LICENSE_FILE` must be set to the full path to the license file.

25.2 Run-time Licenses

There are a few ways to use run-time licenses. Run-time licenses may be contained in a license file, or otherwise coded into the application and provided to CoreDX DDS through the license API.

1. Use an Environment Variable

The environment variable `TWINOAKS_LICENSE_FILE` may be set to one of the following:

- The full path to the license file
- The LICENSE string containing the run-time license

If you have access to the license file from the run-time environment, the simplest way to use the license is to set a `TWINOAKS_LICENSE_FILE` environment variable to be the full path to the license file.

If you do not have access to the license file, you can still use the license by setting the `TWINOAKS_LICENSE_FILE` environment variable to the appropriate run-time LICENSE line. The run-time license line starts with the following:

```
LICENSE PRODUCT=coredx_c
```

Using the license string is a good option for embedded run-time environments. For the run-time license in the above example license file, set your `TWINOAKS_LICENSE_FILE` environment variable like:

```
linux% export TWINOAKS_LICENSE_FILE=<LICENSE
PRODUCT=coredx_c HOSTID=00195b70c3be CUSTOMER=Company_X
SIG=abcdefghijklmnpqrstuvwxyz>
```

2. Use the API

The `DomainParticipantFactory` provides an API to set the license string:

```
DomainParticipantFactory::set_license(const char * lic)
```

The `lic` argument may be set to any of the options that can be used for the `TWINOAKS_LICENSE_FILE` environment variable, described above. That is one of the following:

- The full path to the license file
- The LICENSE string containing the run-time license

The license API is particularly useful for operating systems that do not support environment variables. This allows the application to obtain the license string in any manner acceptable by the environment and system requirements, and then use the API to pass the license string to CoreDX DDS.

Chapter 26 Transition Notes

Every effort is made to make it easy to upgrade between CoreDX DDS versions. Even so, there are occasionally changes that impact some aspects either the CoreDX DDS API or wire protocol. Those changes are documented in this section.

26.1 CoreDX DDS v3.x -> v4.0

26.1.1 C++ Listener Method Signatures

The signature of the listener methods has changed slightly between v3.x and v4.0. Specifically, the ‘status’ parameter passed to most methods (on_data_available() is an exception) is ‘const’ in v4.0. Note: Unfortunately, the c++ compiler doesn’t produce warnings about this, it just quietly doesn’t overload the method.

Example DataReaderListener:

```
class Foo_ReaderListener : public DataReaderListener
{
public:
    void on_data_available( DataReader *dr );
    void on_subscription_matched( DataReader *dr,
    const SubscriptionMatchedStatus status );
    void on_requested_incompatible_qos ( DataReader *dr,
    const RequestedIncompatibleQosStatus status );
    void on_liveliness_changed( DataReader *dr,
    const LivelinessChangedStatus status );
};
```

Chapter 27 Troubleshooting

27.1 General Troubleshooting Tools

Network communication can be complex to troubleshoot. It is recommended that the developer become familiar with standard tools available on the development network. For example, under UNIX, tools such as **ifconfig**, **netstat**, and **route** can be useful to gain an understanding of the network configuration. Further, tools that capture and decode network traffic are very useful. The **wireshark** tool has wide platform support and includes a protocol analyzer for RTPS (the DDS wire protocol). **Wireshark** is an indispensable tool for analyzing DDS network traffic. (See www.wireshark.org).

Twin Oaks Computing offers a tool that is specially designed for analyzing and debugging DDS applications: CoreDX DDS Spy. The CoreDX DDS Spy tool displays, at a glance, all the DDS Entities on the network. This allows the application developer to quickly view of all the DataReaders and DataWriters on the network, what Topic they are communicating on, and which ones are not communicating due to QoS or data type mismatches. In addition, CoreDX DDS Spy to view all the DDS network traffic, including samples written by DDS application for further analysis of DDS applications. The CoreDX DDS Spy tool is also useful in more complex DDS network analysis.

27.2 No Communications between DDS applications

If Readers and Writers are not communicating at all, then there are several items to check. First, it is recommended that Listeners be installed on both the reader and writer to handle all of the events. These events may provide insight into why the entities are not communicating. For example, the `requested_incompatible_qos` and `offered_incompatible_qos` listeners are very useful.

27.2.1 Network Configuration (*if running across a network*)

If all your DDS DomainParticipants are running on one machine, skip this section.

If the DDS DomainParticipants are running across a network, is your network working? Can you use *ping* or another program to talk between your hosts?

27.2.2 Discovery

The first step in DDS communications is the *discovery* process, where DomainParticipants broadcast their existence and look for peer DomainParticipants. This discovery protocol uses *multicast*.

If your DDS DomainParticipants are communication across routers or a Virtual Machine, you may need to increase the reach of your multicast packets by increasing the number of hops (see the *CoreDX DDS Transport* Chapter for more information).

27.2.3 DataReader / DataWriter matching

The next step in DDS communications is matching a DataWriter to a DataReader. Matching requires several compatible attributes:

1. The Topic name must match. Carefully check the `create_topic`, `create_datawriter`, and `create_datareader` calls to ensure the same Topic name string is used for both the DataWriter’s Topic and the DataReader’s Topic.
2. The data types must match. Not only the name of the data type, but also the *types* must match. CoreDX DDS serializes the type of the data into a “typecode”, and compares the typecode of the DataWriter with the typecode of the DataReader. These types must match.
3. Recall that the QoS setting for the DataReader and DataWriter must be compatible for communications to occur (see the *QoS Compatibility* section).

The `SubscriptionMatchedStatus` and `PublicationMatchedStatus` statuses record matching DataReaders and DataWriters. The `OfferedIncompatibleQosStatus` and `RequestedIncompatibleQosStatus` record mis-matching DataWriters and DataReaders due to QoS incompatibility. Use Listeners (see the *Listeners* section) or Conditions (see the *Conditions and WaitSets* section) to check these statuses.

27.3 Missing or lost samples

There are numerous QoS policies that can cause samples to be missing or lost. A few of the more common ones are described below. In addition, QoS settings can interact with each other causing non-intuitive application behavior. While the examples below describe some common problems and solutions, your specific network environment and other QoS settings may result in application behavior different than what is described below.

Twin Oaks Computing is dedicated to the helping customers get the most out of their application communications using CoreDX DDS. Please contact us for additional support with your specific application.

27.3.1 Reliability

If you are communicating over a network, a slow or unreliable network can cause packets to be lost. Similarly, one “slow” subscribing host can have trouble keeping up with publishing hosts. Setting the Reliability QoS policy to RELIABLE can reduce or eliminate lost packets in this scenario.

It is important to note that a RELIABLE Reliability can only happen while both the DataWriter and DataReader are both in existence. Sometimes, a publishing application will exit (killing the DataWriter) before the DataReader has received all the published samples, resulting in lost samples.

27.3.2 Durability

If your DataReader is consistently missing the first one or two samples published by a DataWriter, chances are the discovery process is not completing (matching the DataWriter to the DataReader) before those first samples are written. In general, the solution is to raise the Durability QoS setting to TRANSIENT_LOCAL. This can have other side effects when combined with other QoS settings (see the *HISTORY* section). Other options include having the publishing application wait for a discovered DataReader or simply pause for one or two seconds before starting to write data; allowing the discovery process to complete.

27.3.3 History

By default, the History QoS policy is set to KEEP_LAST, with a depth of 1 (one). Consider a DataWriter writing samples fast enough that the CoreDX DDS infrastructure must queue any before sending, or a DataReader receiving samples fast enough that they must be queued before a read() or take() operation is used. With a History that is only keeping the 1 most recent data sample for each instance, there is a possibility for samples to be dropped. The solution is to increase the History depth to greater than 1, or set the History to KEEP_ALL.

There is only one combination of QoS settings that will guarantee samples are not lost during operations and that is:

Reliability = RELIABLE

History = KEEP_ALL

Resource Limits = Set

This combination of QoS policies forces the publishing application to block on a `DataWriter::write()` operation, if any matched `DataReaders` is unable to accept another sample. The `DataWriter::write` operation will complete (and return) once all matched `DataReaders` have enough room to receive an additional sample.

27.3.4 Putting it all together: Guaranteed Delivery

A common question from DDS users is, “How do I guarantee delivery with DDS?” The goal is to guarantee all data written by a publishing application will be delivered to all matched subscribing applications.

There are three QoS policies that need to be configured to guarantee delivery of data samples:

1. Reliability (`kind = RELIABLE`)
2. History (`kind = KEEP_ALL`)
3. Resource Limits (`set` to something other than `infinite`)

All of these QoS policies must be set to ensure delivery of published data.

The `RELIABLE` Reliability setting allows CoreDX DDS to monitor data reception, and retransmit data if it is not received by any `DataReaders`.

The `KEEP_ALL` History setting instructs CoreDX DDS that it is NOT OK to overwrite any data samples in the `DataWriter`’s cache. This is important if there are any `DataReaders` that are having trouble “keeping up”, and lots of samples must be stored for retransmission.

Setting the Resource Limits allows CoreDX DDS to limit the growth of the `DataWriter`’s cache, even with the `KEEP_ALL` History setting. This is important, not only for resource utilization at run time, but also because it allows the application’s call to `DataWriter::write()` to block if there is no more room in the cache (because at least one `DataWriter` has not acknowledged a number of sent samples).

27.4 TypeSupport version mismatch

CoreDX DDS provides strong data typing. Application developers define the data types that will be used for DDS communications at compile time, and use the CoreDX data type compiler to generate type specific DDS code for each data

type. This generated code interacts very closely with the DDS library to perform type specific operations (for example, serializing data on a write() and de-serializing data on a read()). For this reason, it is important that the data type compiler used to generate the code match the DDS library that is linked into the application. If these versions do not match, CoreDX DDS will print a warning message when register_type() is called:

Sample Warning Message for Version Mismatch

```
WARNING:  MyType TypeSupport version does not match CoreDX
Library version.
This may cause software instability or crashes.
```

Figure 27-1: Example CoreDX DDS license file

To resolve this issue, re-generate your type specific code with the correct version of the CoreDX data type compiler, and check the version of libdds.a that you are linking with your application.

27.5 Can't find it here?

Call us at 720-733-7906, send an email to support@twinoakscomputing.com, or check out our Frequently Asked Questions at <http://www.twinoakscomputing.com/coredx/faq>, or visit our online forums at <http://www.twinoakscomputing.com/forums>.

Chapter 28 About Twin Oaks Computing

Twin Oaks Computing, Inc is a company dedicated to developing and delivering quality software solutions. We leverage our technical experience and abilities to provide innovative and useful services in the domain of intelligence systems.

Twin Oaks Computing specializes in high-performance and embedded communications solutions for commercial and DoD applications. Our CoreDX DDS was first released in 2008. In March 2009, Twin Oaks Computing participated in the first public multi-vendor DDS interoperability demonstration. In addition to CoreDX DDS, we provide a variety of DDS-related products to help develop, integrate, and scale DDS solutions.

Twin Oaks Computing is headquartered in Castle Rock, CO. Our staff has over 30 years of experience developing and supporting complex mission-critical systems. For the past 15 years we have been integrating DDS into a wide variety of commercial and military systems, including:

Medical – connecting devices in operating rooms, monitoring rooms, and remote locations

Robotics and Unmanned vehicles – commercial industrial and military applications

Energy – microgrid and wide distribution control components

Command and Control – land, sea, air, and the ground stations that monitor them

Automotive – infotainment, advanced detection and control systems, telematics

Space – onboard rovers and the ISS, connecting space assets to monitor and control stations on earth

And many more. We would be happy to discuss how we can help you. Please contact us at contact@twinoakscomputing.com.

Chapter 29 Contact Information

Have a question? Don't hesitate to contact us by any means convenient for you:

Web Site: <http://www.twinoakscomputing.com>

Email: support@twinoakscomputing.com

Twitter: @CoreDX_DDS

Phone: 720.733.7906

+33(0)9 62 23 72 20

Address:

Twin Oaks Computing, Inc.
230 Third Street, Suite 260
Castle Rock, CO, 80104