



DDS Tutorial -- Part II Hands On

Gerardo Pardo-Castellote, Ph.D.

**The Real-Time
Middleware Experts**

Gerardo Pardo-Castellote, Ph.D.
Co-chair OMG DDS SIG
CTO, Real-Time Innovations
gerardo.pardo@rti.com



Agenda

- Context: Middleware Technologies
- Overview: DDS Model & Applicability
- Details: DDS in depth

Context: Middleware Technologies




The concept of network middleware:

- Communications Model
- Object Model
- Architecture Model
- Protocol

With increased complexity...



End User Application



With increased complexity...



End User Application

Network Stack




Routing protocols

File System

Kernel

Device Drivers

Operating System



With increased complexity...



End User Application

HTTP Service

Mail Service

FTP Service

Data base

Network Stack




Routing protocols

File System

Kernel

Device Drivers






Operating System



... middleware becomes necessary



End User Application



What is network middleware?

Middleware =

API and service layer above operating system and below
“application” code that abstracts common interaction patterns

Network Middleware =

Most popular class of middleware

Middleware used for developing distributed applications

Distributed Applications =

Those requiring interaction/communication between multiple
computers

Network Middleware Examples

- DDS, RTI DDS, OpenSplice, tao-dds,
- JMS, WebSphere MQ, ActiveMQ, SoniqMQ
- DCE/RPC, DCOM, CORBA, ICE
- TIBCO RV, 29West, GigaSpaces
- Bundles & ESBs:
 - Application Servers (WebSphere, WebLogic, JBOSS) Include network middleware as a component
- NOTE: Middleware “packages” are building blocks, not stand-alone applications like...
 - Skype, gtalk, ...
 - BitTorrent, eMule, ...
- Why aren’t ‘popular’ consumer applications built on top of middleware?

Historical note: From Telephone to Blogs

- Why so many flavors?
- Parallels evolution of general communication patterns:
 - Started with point-to-point connections
 - Then request-reply services
 - Then Message Queue Services
 - Then Publish-Subscribe Services
 - Then Data-Caching services
- Other examples:
 - FTP, email -> WEB -> Blogs, RSS -> Podcasts

Middleware Communication Models




Point-to-Point

Telephone, TCP

Simple, high-bandwidth

Leads to stove-pipe systems




Client-Server

File systems, Database, RPC, CORBA, DCOM

Good if information is naturally centralized

Single point failure, performance bottlenecks




Publish/Subscribe Messaging

Magazines, Newspaper, TV

Excels at *many-to-many* communication

Excels at distributing *time-critical* information



Replicated Data


Libraries, Distributed databases

Excels at data-mining and analysis

RMI vs Pub-Sub/Messaging/Data-Distribution



- RMI (WebServices, CORBA, DCOM) offer a remote method abstraction

- Familiar OO programming model
- Results in a tightly-coupled system
 - Forces synchronous invocations
 - Imposes global object model
 - Limited QoS (appearance of local method call)
 - Lack robustness: cascading points of failure
- Typically built on top of TCP:
 - impacts scalability and time-determinism
- Best-suited to smaller, closely-coupled systems





- Pub-Sub (Messaging Data-Distribution) offer a queue-based and/or replicated-data model

- Subsystems are decoupled in time, space, and synchronization
 - Contracts established by verifying QoS compatibility
- Supports a variety of transports including multicast UDP
- Better suited for high-performance and real-time




Queue versus Pub-Sub

- Queue
 - Message sent to Queue
 - Multiple readers can read from the queue
 - Each message is delivered to ONLY one reader
 - Readers “affect each other”
 - Apps:
 - Job Scheduling
 - Load Balancing
 - Collaboration
- Pub Sub
 - Message Sent to Topic
 - Multiple readers can subscribe to Topic with or without filters
 - Each message delivered to ALL subscribers that pass filter
 - Readers are decoupled
 - Apps:
 - Notifications
 - Information Distribution



Pub-Sub versus Data-Distribution

- Pub-Sub
 - Only messages no concept of data
 - Each message is interpreted without context
 - Messages must be delivered FIFO or according to some “priority” attribute
 - No Caching of data
 - Simple QoS: filters, durability, lifespan
- Data-Distribution
 - Messages represent update to data-objects
 - Data-Objects identify by a key
 - Middleware maintains state of each object
 - Objects are cached. Applications can read at leisure
 - Smart QoS
 - Ownership
 - History (per key)
 - Deadline
 - Subsumes Pub-Sub




DDS Service Model: Communication Model




Provides a “**Global Data Space**” that is accessible to all interested applications.

- Data objects addressed by **Domain**, **Topic** and **Key**
- Subscriptions are **decoupled** from Publications
- Contracts established by means of **QoS**
- Automatic **discovery** and **configuration**



DDS Service Model: Object Model



- **DomainParticipant** – Allows application to join a DDS Domain (Global Data Space)
- **Topic** – A string that addresses a group of objects in the Global Data Space
 - Each Object is identified by a Key (some fields within the object data)
- **Publisher, Subscriber** – Pools resources for DataWriters and DataReaders
- **DataWriter** – Declares intent to publish a Topic and provides type-safe operations to write/send data
- **DataReader** – Declares intent to subscribe to a Topic and provides type-safe operations to read/receive data

Other (non DDS) Commercial Pub-Sub Models



- Older, but widely deployed
 - TIBCO (RendezVous, EMS)
 - IBM WebSphere MQ
- Limited deployment:
 - CORBA Event Service
 - CORBA Notification Service
- Emerging standards – not widely deployed
 - WS-Eventing
 - WS-Notification
- Proprietary Systems
 - 29West
 - IBM LLM

Agenda

- Context: Middleware Technologies
- **Overview: DDS Model & Applicability**
- Details: DDS in depth

History: DDS the Standard(s)

- Data Distribution Service for Real-Time Systems
 - Joint submission (RTI, THALES, OIS)
 - DDS version 1.0 Adopted December 2004
 - DDS version 1.1 Adopted December 2005
 - DDS version 1.2 Adopted October 2006
- Interoperability wire protocol
 - Joint submission (RTI and PrismTech)
 - DDS-RTPS version 1.2 adopted in July 2006
 - DDS-RTPS version 2.0 adopted in June 2007
 - DDS-RTPS version 2.1 approved in June 2008
- Related Standards
 - Joint submission (Sparx, RTI, PrismTech)
 - UML Profile for DDS adopted June 2008
 - DDS for light weight CCM adopted 2008
- Standards under Development
 - Extensible and Dynamic Topic Types for DDS
 - Native Language C++ API for DDS





Navigation

Data Distribution Portal (HOME)

[Data Distribution Intro](#)
[Data Distribution SIG](#)
[DDS User's Forum](#)
[Event Calendar](#)
[Specifications](#)
[Tutorials](#)
[Whitepapers](#)
[Other Materials](#)
[FAQs](#)
[Vendors](#)
[Portal Index](#)
[Portal Search](#)
[Portal Maintenance](#)
[Portal Help](#)
[Recent Changes](#)
[SystemInfoPage](#)

Search

User

[GerardoPardo](#)
[Preferences](#)
[Logout](#)

Recently viewed pages

[DataDistributionSIG](#)
[DataDistributionIntro](#)
[SpecificationsPage](#)
[SpecificationsInProgress](#)
[OMG Data Distribution Portal](#)

OMG Data Distribution Portal

You may discuss this page here.

The Data-Distribution Service for Real-Time Systems (DDS) is a recently-adopted [OMG](#) standard.

DDS is the first open international middleware standard directly addressing *publish-subscribe* communications for *real-time and embedded systems*. DDS introduces a virtual *Global Data Space* where applications can share information by simply reading and writing data-objects addressed by means of an application-defined name (**Topic**) and a **key**. DDS features fine and extensive control of QoS parameters, including *reliability*, *bandwidth*, *delivery deadlines*, and *resource limits*. DDS also supports the construction of local object models on top of the *Global Data Space*.

The DDS portal is maintained by the **OMG Data Distribution SIG (DDSIG)**. For the activities of the DDSIG and other events of interest to the community, please visit the [EventCalendar](#) and the [DDSIG page](#). The portal uses a Wiki to manage the content. Before making edits please visit the [PortalMaintenancePage](#) and read the [PortalUsagePolicies](#). You may also want to look at the [WikiCourse](#) and at [HelpContents](#). The [WikiSandBox](#) is a good place to experiment with editing.

Learning about DDS

A general introduction to DDS can be found in [DataDistributionIntro](#). Howtos, patterns of use of DDS and example code can be found in [DataDistributionExamples](#).

The [TutorialsPage](#) contains presentations on DDS. More material can also be found in the [WhitepapersPage](#) and the [OtherMaterialPage](#). The [TrainingPage](#) lists organizations that hold regular training on DDS or can prepare it on demand.

The [InformationDays](#) contains use-case and vendor presentations OMG DDS information days. These events started in 2006 and are still on-going.

The current DDS specification is [Version 1.2](#) Older and related specifications can be found in the [SpecificationsPage](#).

The current DDS Interoperability Wire Protocol specification is [Version 2.1](#) Older and related specifications can be found in the [SpecificationsPage](#).

Visit the [VendorsPage](#) or the [ProjectsPage](#) learn about vendors, products and projects that support or use DDS.

News and Events

2008.6.26 -- UML Profile for DDS Specification. OMG Technical Meeting held in Ottawa, Canada. The OMG has recommended for adoption the [UML Profile for DDS Specification](#). The specification defines how UML tools can be used to model DDS systems and automatically generate supporting code. For more information visit the [SpecificationsPage](#).

2008.6.26 -- Extensible and Dynamic Topic Types for DDS RFP. OMG Technical Meeting held in Ottawa, Canada. The OMG has issued and RFP with the dual goals of adding Type Extensibility to DDS Topics as well as introducing a dynamic API to allow reading and writing types for which there was no compile-time knowledge. For more information visit [SpecificationsInProgress](#).

2008.5.16 -- Hands-On DDS Workshops. [RTI](#) has scheduled several hands-on DDS training workshops in 2008: May 20-21, August 20-21 and October 22-23. Whether you are evaluating, planning to use, or already using the DDS standard, these two-day workshops provide an excellent opportunity to learn about the capabilities of the standard and how to apply them to your application. For more information and to register visit <http://www.rti.com/services/workshops.html>.

2008.5.16 -- DDS Demonstration Application. [RTI](#) has developed a demonstration application that illustrates DDS concepts such as publish-subscribe, real-time QoS, and data-centric design. Download the demo at https://www.rti.com/mk/shapes_demo.html.

2008.3.14 -- Native C++ Language DDS API RFP. OMG Technical Meeting held in Washington DC. The OMG has issued and RFP with the objective of defining a new C++ API to DDS that takes advantage of the language features present in the ISO C++ Standard. For more information visit [SpecificationsInProgress](#).

DDS mandated for data-distribution

- DISR (formerly JTA)
 - DoD Information Technology Standards Registry
- US Navy Open Architecture
- FCS SOSCOE
 - Future Combat System – System of System Common Operating Environment
- SPAWAR NESI
 - *Net-centric Enterprise Solutions for Interoperability*
 - *Mandates DDS for Pub-Sub SOA*



DDS Adoption – Aerospace & Defense



Boeing
AWAKS

E2C Hawkeye

Northrop?
Qinetiq?
SAAB?



Boeing
Future Combat
Systems

Raytheon
SSDS



Lockheed
AEGIS

Insitu
Unmanned
Air Vehicles



DDS Adoption – Transportation, Industrial



WiTronix
Train and vehicle
Tracking

Schneider Electric
Industrial Automation



Tokyo Japan
Traffic Control

Kuka
Robotics



EU and US
Air Traffic
Management

Varian
Medical Instruments





Many others

GENERAL DYNAMICS
Strength on Your Side™

LOCKHEED MARTIN
We never forget who we're working for™

BOEING®

NORTHROP GRUMMAN

Raytheon

TCG

HARRIS

CAE

Silver Arrow

TCG
Tactical Communications Group, LLC

MITSUBISHI
ELECTRIC

Schneider
Electric

infinera
Solutions For Optical Networks

THALES

BAE SYSTEMS

KUKA
Group

CITI

QinetiQ

WI-TRONIX

OMRON
Sensing tomorrow™

AT&T

SAAB

VARIAN
medical systems

INSITU

gmv
INNOVATING SOLUTIONS

PIMCO

... The Authority on Bonds™

ALSTOM

SAMSUNG

CSC
EXPERIENCE. RESULTS.

Indra

LINCOLN LABORATORY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

SAIC
From Science to Solutions

INFODYNE²
zero tolerance for latency

Rockwell
Collins


Main differences of DDS vs other Pub-Sub

- Flexibility and Power of the data-centric model
- Performance & Scalability
- Rich set of built-in services
- Interoperability across platforms and Languages
- Natural integration with SOA building-blocks

#1 RTI Data-Centric Model

“Global Data Space” generalizes Subject-Based Addressing

- Data objects addressed by **DomainId**, **Topic** and **Key**
- **Domains** provide a level of isolation
- **Topic** groups homogeneous subjects (same data-type & meaning)
- **Key** is a generalization of **subject**
 - **Key** can be any set of fields, not limited to a “x.y.z ...” formatted string




Demo: Concepts

[Start demo](#)




Display Area:
Shows state of objects



Control Area:
Allows selection of objects and QoS

- Topics
 - Square, Circle, Triangle
 - Attributes
- Data types (schemas)
 - Shape (color, x, y, size)
 - Color is instance Key
 - Attributes
 - Shape & color used for key
- QoS
 - Deadline, Liveliness
 - Reliability, Durability
 - History, Partition
 - Ownership

DDS communications model



- **Participants** scope the global data space (domain)
- **Topics** define the data-objects (collections of subjects)
- **Writers** publish data on Topics
- **Readers** subscribe to data on Topics
- **QoS Policies** are used to configure the system
- **Listeners** are used to notify the application of events


QoS: Quality of Service



QoS Policy	QoS Policy
DURABILITY	USER DATA
HISTORY (per subject)	TOPIC DATA
READER DATA LIFECYCLE	GROUP DATA
WRITER DATA LIFECYCLE	PARTITION
LIFESPAN	PRESENTATION
ENTITY FACTORY	DESTINATION ORDER
RESOURCE LIMITS	OWNERSHIP
RELIABILITY	OWNERSHIP STRENGTH
TIME BASED FILTER	LIVELINESS
DEADLINE	LATENCY BUDGET
CONTENT FILTERS	TRANSPORT PRIORITY

Demo: Quality of Service (QoS)

Writers and readers state
Their needs



RTI DDS delivers

[Start demo](#)


- Topics
 - Square, Circle, Triangle
 - Attributes
- Data types (schemas)
 - Shape (color, x, y, size)
 - Color is instance Key
 - Attributes
 - Shape & color used for key
- QoS
 - Deadline, Liveliness
 - Reliability, Durability
 - History, Partition
 - Ownership

#2 Performance & Scalability


DDS was designed to support high performance

- DDS uses high-performance data-access APIs
 - Read data by array (no additional copies)
 - Buffer loaning for zero copy access
- DDS Supports advanced features such as:
 - Message prioritization (via latency budget QoS)
 - Network prioritization (via transport priority QoS)
 - Source filtering (via Content-based and Time-based filters)
- DDS does not require the presence of intermediate brokers
 - Applications can communicate directly peer-to-peer
- DDS Protocol supports UDP, multicast and reliable multicast
 - Multicast can be used for high-performance scalability
 - Use of UDP avoids head-of-line blocking
 - Best efforts can be used for repetitive time-critical data

Data-Distribution and Real-Time




Latency – (Linear Scale)




Adapted from Vanderbilt presentation at July 2006 OMG Workshop on RT Systems

Jitter – (Linear Scale)



Source: Vanderbilt presentation at July 2006 OMG Workshop on RT Systems

DDS compared to JMS






Platform: Linux 2.6 on AMD Athlon, Dual core, 2.2 GHz

Study on impact of WS technologies for future European ATC: XML is not suitable for European Flight Data Distribution



- Data size explodes 10X vs. CDR
 - Flight Plans go from 100KB to 1MB
- Communication speed drops 20X

Not good for Real-Time!!



Source: Christian Esposito and Domenico Cotroneo, Dario Di Crescenzo.





SELEX-SI/Consorzio SESM/University of Naples.

“Flexible Communication Among DDS Publishers and Subscribers”

July 2008, Real-Time Systems Workshop, Washington, DC



Message bus architectures



#3 Powerful Services

Included in the Standard:

- Discovery
- Ownership, Redundancy & Failover
- Persistence (Durable) Data
- Last value and historical caches

Enabled by the DDS protocol:

- Recording service
- Routing Service


Discovery Service

- DDS provides the means for an application to discover other participants on the Domain
 - And the Topics the Publish and Subscribe
 - And the Quality of Service of the remote endpoints
- A participant can determine who it is communicating with...
 - And selectively decide who to communicate with...

[shapes_demo](#)

[discovery_in_excel](#)

Ownership and High Availability



- Owner determined per subject
- Only extant writer with highest strength can publish a subject (or topic for non-keyed topics)
- Automatic failover when highest strength writer:
 - Loses liveness
 - Misses a deadline
 - Stops writing the subject
- Shared Ownership allows any writer to update the subject


[Start demo](#)

Data Persistence

A standalone service that persists data outside of the context of a DataWriter

Can be configured for:

- Redundancy
- Load balancing




Demo:

1. [PersistenceService](#)
2. [ShapesDemo](#)
3. [Application failure](#)
4. [Application \(ShapesDemo\) re-start](#)


Persistent Data Service

- Modes:
 - High-performance “direct” model. Data flows in directly to subscribers and to the persistent service
 - Transactional “relay” mode persists first, then sends to subscribers
- Fault-tolerant: Redundant services are supported




Durable Writers

- Durable Writers keep their history cache in permanent storage
- Upon re-starts they automatically rebuild their history from storage
- New messages will automatically be appended to the history
- Readers will treat the writer re-start as a “temporary” disconnect
- Automatically-assigned sequence numbers ensure no messages are dropped or duplicated



Durable Readers

- Durable Readers Keep information on the messages read on permanent storage
- Information uniquely identifies each message
- Upon re-start they automatically load information so they only request new messages and avoid duplicates
- Combined with the other services it can guarantee exactly-once delivery



Last value cache

- A last-value cache is already built-in into every Writer in the system
 - Can be used in combination with a Durable Writer
- A late joiner will automatically initialize to the last value
- Last value cache can be configured with history depth greater than 1
- The Persistence Service can be used to provide a last value cache for durable data


[Start demo](#)

Historical cache

- A partial historical cache is already built-in into every Writer in the system
 - Can be used in combination with a Durable Writer
- The Persistence Service can be used to provide a historical cache with larger/unlimited depth
- A late joiner will automatically initialize to the desired history
 - Currently amount of history can only be specified as message count.
 - Next release will also allow an age/time based specification.
- Request for historical cache is done by creating a Reader with the desired history depth specified as a QoS

DDS Real-Time Recording Service

- Applications:
 - Future analysis and debugging
 - Post-mortem
 - Compliance checking
 - Replay for testing and simulation purposes
- Record high-rate data arriving in real-time
- Non-intrusive – multicast reception




Demo:


1. Start [RecorderService](#)
2. Start [ShapesDemo](#)
3. See [output files](#)
4. Convert to: [HTML](#) [XML](#) [CSV](#)
5. View Data: [HTML](#) [XML](#) [CSV](#)

#4 Interoperability

- Protocol
 - Interoperability Wire protocol adopted in 2006
- Languages: C/C++, Java, ADA, .NET
- Systems/Platforms/Models
 - Data distribution (publishers and subscribers): **DDS**
 - Data management (storage, retrieval, queries): **SQL**
 - ESB Integration, Business process integration: **WSDL**



#5 Natural with SOA building blocks




Relational Database Integration

Messaging Actions

- Write()
- Read() & Take()
- Dispose()
- Wait() & Listener

Relational Actions


- UPDATE & INSERT
- SELECT
- DELETE




Event driven – The fastest way to observe database changes!

Complex Event Engine Integration

- CEP: programmable engines used to transform “data” into “information”
- CEP engines are programmed using a derivative of SQL
- CEP engines save time: They can implement a lot of the application logic:
 - Classification, Correlation, Aggregation, Filter, Cleansing, Pattern Detection, etc.
- DDS is the perfect ‘data’ and ‘information’ pipe for CEP engines
 - Use high-speed data streams (1,000-1,000,000 msg/sec)
 - Require latency measured in sub-milliseconds
 - Demand access to events from a heterogeneous systems



Web-Enabled DDS (upcoming spec)




Agenda

- Context: Middleware Technologies
- Overview: DDS Model & Applicability
- **Details: DDS in depth**


Do-yourself Message-Centric System

- Model
 - 1-1, FIFO
- Applications coupled in Lifespan & Content
 - Both must be present simultaneously
 - Everything sent is received
- Excellent performance
- Doesn't scale
 - To large-scale systems
 - To loosely-coupled systems




Middleware-based Message-Centric Systems (JMS)

- Model
 - Broker-based, $n \rightarrow n$ communication
 - Independent messages, No state
- Coupling
 - Not coupled in Lifespan
 - Coupled in Order and Content (presentation)
- Worse performance
- Better scalability



Do-yourself Data-Centric System

- Model
 - Shared/replicated structured data/state
 - Asynchronous, Selective sharing
- Coupling:
 - Coupled in Lifespan, Decoupled in Presentation & Content
- Excellent performance & scalability



Middleware-based Data-Centric Systems (DDS)


- Shared data-space, shared state
- Asynchronous communication
- Decoupled in Lifespan, Content, presentation
- Excellent performance & scalability
- Subsumes Message-Centric via QoS



DDS Data-Centric Model


Provides a virtual “**Global Data Space**” that is accessible to all interested applications.

- Data objects addressed by **DomainId**, **Topic** and **Key**
- Subscriptions are **decoupled** from Publications
- Contracts established by means of **QoS**
- Automatic **discovery** and **configuration**




DDS Global Data

- Address in Global Data Space = (DomainId, Topic, Key)
 - Each topic corresponds to a multiple data instances with a common schema
 - A DataWriter can write to any instances of a single topic
 - Multiple DataWriters may write to the same instance
 - A DataReader receives updates from all instances of a single topic
 - Multiple DataReaders may read from the same instances & values



DDS Global Data: Domains

- Address in Global Data Space = (DomainId, Topic, Key)
 - Each Domain is identified by the value of the domainId
 - Each Domain is a separate Global Data Space
 - The same Topic name can mean different things in different domains
 - The same Topic can have different Types on each Domain
 - An application may join multiple Domains
 - Domains can be used for isolation, scalability, modularity



Example: Publication

```
// Entities creation
DomainParticipant participant =
    TheParticipantFactory->create_participant(
        domain_id, participant_qos, participant_listener);

Publisher publisher = domain->create_publisher(
    publisher_qos, publisher_listener);

Topic topic = domain->create_topic(
    "MyTopic", "Text", topic_qos, topic_listener);

DataWriter writer = publisher->create_datawriter(
    topic, writer_qos, writer_listener);

TextDataWriter twriter = TextDataWriter::narrow(writer);

TextStruct my_text;
twriter->write(&my_text);
```

Example: Subscription

```
// Entities creation
Subscriber subscriber = domain->create_subscriber(
    subscriber_qos, subscriber_listener);

Topic topic = domain->create_topic(
    "Track", "TrackStruct",
    topic_qos, topic_listener);

DataReader reader = subscriber->create_datareader(
    topic, reader_qos, reader_listener);

// Use listener-based or wait-based access
```

How to Get Data? (Listener-Based)

```
// Listener creation and attachment
Listener listener = new MyListener();
reader->set_listener(listener);

// Listener code
MyListener::on_data_available( DataReader reader )
{
    TextSeq received_data;
    SampleInfoSeq sample_info;
    TextDataReader reader = TextDataReader::narrow(reader);

    treader->take( &received_data, &sample_info, ... )
    // Use received_data
    printf("Got: %s\n", received_data[0]->contents);
}
```

How to Get Data? (WaitSet-Based)

```
// Creation of condition and attachment
Condition foo_condition =
    treader->create_readcondition(...);
waitset->add_condition(foo_condition);

// Wait
ConditionSeq active_conditions;
waitset->wait(&active_conditions, timeout);

// Wait returns when there is data (or timeout)
FooSeq received_data;
SampleInfoSeq sample_info;

treader->take_w_condition
    (&received_data,
     &sample_info,
     foo_condition);

// Use received_data
printf("Got: %s\n", received_data[0]->contents);
```

Listeners, Conditions & WaitSets

Middleware must notify user application of relevant events:

- Arrival of data
- But also:
 - QoS violations
 - Discovery of relevant entities
- These events may be detected asynchronously by the middleware
 - ... Same issue arises with POSIX signals

DDS allows the application to choice:

- Either to get notified asynchronously using a **Listener**
- Or to wait synchronously using a **WaitSet**

Both approaches are unified using **STATUS** changes

Status Changes

DDS defines

- A set of enumerated STATUS
- The statuses relevant to each kind of DDS Entity

DDS entities maintain a value for each STATUS

STATUS	Entity
INCONSISTENT_TOPIC	Topic
DATA_ON_READERS	Subscriber
LIVELINESS_CHANGED	DataReader
REQUESTED_DEADLINE_MISSED	DataReader
RUQESTED_INCOMPATIBLE_QOS	DataReader
DATA_AVAILABLE	DataReader
SAMPLE_LOST	DataReader
SUBSCRIPTION_MATCH	DataReader
LIVELINESS_LOST	DataWriter
OFFERED_INCOMPATIBLE_QOS	DataWriter
PUBLICATION_MATCH	DataWriter

```
struct LivelinessChangedStatus
{
    long active_count;
    long inactive_count;
    long active_count_change;
    long inactive_count_change;
}
```


Listeners, Conditions and Statuses

- A DDS Entity is associated with:
 - A listener of the proper kind (if attached)
 - A StatusCondition (if activated)
- The Listener for an Entity has a separate operation for each of the relevant statuses


STATUS	Entity	Listener operation
INCONSISTENT_TOPIC	Topic	on_inconsistent_topic
DATA_ON_READERS	Subscriber	on_data_on_readers
LIVELINESS_CHANGED	DataReader	on_liveliness_changed
REQUESTED_DEADLINE_MISSED	DataReader	on_requested_deadline_missed
REQUESTED_INCOMPATIBLE_QOS	DataReader	on_requested_incompatible_qos
DATA_AVAILABLE	DataReader	on_data_available
SAMPLE_LOST	DataReader	on_sample_lost
SUBSCRIPTION_MATCH	DataReader	on_subscription_match
LIVELINESS_LOST	DataWriter	on_liveliness_lost
OFFERED_INCOMPATIBLE_QOS	DataWriter	on_offered_incompatible_qos
PUBLICATION_MATCH	DataWriter	on_publication_match

Listeners & Condition duality

- A StatusCondition can be selectively activated to respond to any subset of the statuses
- An application can wait changes in sets of StatusConditions using a WaitSet
- Each time the value of a STATUS changes DDS
 - Calls the corresponding Listener operation
 - Wakes up any threads waiting on a related status change



Hands-on Example (C++)




Four easy steps:


1. [Define your data](#)
2. [Create your project](#)
3. [Build](#)
4. Run: [publisher](#) [subscriber](#)

Aux:
[File Browser](#)
[Console](#)

Components/Mechanics of the GDS



Components/Mechanics of the GDS



Designing a Data-Centric System


- Define/Model the Global Data Space
- Configure the Cache Management
- Configure Discovery
- Configure the Protocol
- Configure/Use hooks for
 - Fault detection
 - Controlled access

Global Data Space / Global State

- Identify the number of domains
- Domain Information model
 - Topics
 - Types
 - Keys
 - Ownership



Domain and Domain Participants




- Container for applications that want to communicate
- Applications can join or leave a domain in any order
- New Applications are “Auto-Discovered”
- An application that has joined a domain is also called a “Domain Participant”

Single 'Domain' System

Domain and Domain Participants

Using Multiple domains for Scalability, Modularity & Isolation



[demo_domain_0](#)

Multiple Domain System

[demo_domain_1](#)

Topics & Datatypes, Keys & Subjects

Topic “MarketData”

Data-type (name-type-value pairs)

source	type	symbol	Exchange	volume	bid	ask
OPRA		IBM	NYSE	200000	118.30	118.36
OPRA		AAPL	NASDAQ		171.20	171.28
RTFP	EQ					

Key fields → Subject

Additional fields (payload)

Topic “OrderEntry”

Exchange	type	Symbol	Order num	number	limit	stop	expiration
NYSE	BUY	IBM	11956	500	120	-	DAY
NYSE	BUY	IBM	11957	1000	124.5	124	DAY
NASDAQ	SELL	AAPL	11958	400	-	160	DAY

Subject

Key fields


[demo_filters](#)

QoS: Ownership

Specifies whether more than one DataWriter can update the same instance of a data-object


Ownership = EXCLUSIVE

“Only highest-strength data writer can update each data-instance”



Ownership = SHARED


“All data-writers can each update data-instance”



Provides fast, robust, transparent replacement for fail-over and/or take-over.

QoS: Ownership Strength

Specifies which DataWriter is allowed to update the values of data-objects




Note: Only applies to Topics with Ownership = Exclusive

Configure the Cache Management

- Cache State Content

- History
- Lifespan
- Persistence
- Resources



- Reader Cache View

- Partitions
- Content-Based Filter
- Time-Based Filter
- Order

QoS: History – Last x or All


KEEP_ALL:

Publisher: keep all until delivered

Subscriber: keep each sample until the application processes that instance

KEEP_LAST: “depth” integer for the number of samples to keep at any one time

[demo history](#)



QoS: Lifespan


lifespan_pub



lifespan_sub

User can set lifespan duration

Manages samples in
the history queues, attached to each
Sample



Content-Based Filtering

Instance 1

Value = 249

Instance 2

Value = 230

Instance 3

Value = 275

Instance 4

Value = 262

Instance 5

Value = 258

Instance 6

Value = 261

Instance 7

Value = 259

•
•
•

The Filter Expression and Expression Params will determine which instances of the Topic will be received by the subscriber.

Content Filtered Topic

Topic

“Filter Expression ”
Ex. Value > 260

content filter example

Subscriptions: By Topic, Subject, Content

Topic: “Market Data”

Field	Source	Symbol	Type	Exchange	Volume	Payload			...
Value	*	*	*	*		Bid	Ask		*

Topic: “Order Entry”

Field	Symbol	Type	Exchange	OrderNumber	Symbol	OrderKind	Stop	Limit	...
Value	*	*	NYSE						*

Subject Filter (for a Reader)

content filter example

Topic: “Market Data”


Field	Source	Symbol	Type	Exchange	Payload
Value	REUTERS	*	EQ	NYSE	Volume > x, Ask < y

Subject Filter (for a Reader)


Payload Filter (for a Reader)

QoS: TIME_BASED_FILTER

time filter example



Cache Management in Action




- Topics
 - Square, Circle, Triangle
 - Attributes
- Data types (schemas)
 - Shape (color, x, y, size)
 - Color is instance Key
 - Key
 - Color field used for key
- QoS
 - History, Partition
 - Time-Based Filter
 - Content-Based Filter

demo


Configure the Protocol

- Discovery
- Reliability
- Liveliness
- Flow Control
- Asynchronous write
- Network Configuration
 - Enabled Transports + transport properties
 - Multicast addresses
 - Transport Priority
- OS settings
 - Threads
 - Memory



Tunable Reliability Protocol


- Configurable AckNack reply times to eliminate storms
- Fully configurable to bound latency and overhead
 - Heartbeats, delays, buffer sizes
- Performance can be tracked by senders and recipients
 - Configurable high/low watermark, Buffer full
- Flexible handling of slow recipients
 - Dynamically remove slow receivers



Best Efforts




01
02
03
04
05
06
07
08



01
02
X
04
05
06
07
08


Confirmed Reliability I

No packet loss



Confirmed Reliability II


Some packet loss




Configure Notifications, Fault Detection & Management



- Listeners
- Deadline Qos
- Liveliness Qos
- Built-in Readers
- Notification of matching





QoS: Liveliness – Type and Duration

liveliness example




Type: Controls who is responsible for issues of 'liveliness packets'

AUTOMATIC = Infrastructure Managed

MANUAL = Application Managed

kill apps



Summary

- Designing a fault-tolerant high-performance distributed system is not a simple task
- A powerful middleware framework can provide a lot of value and help you focus on the business logic
- The middleware can save you a lot of time and effort.
 - It is worth learning how to use its power!