DDS on SPARTA

Stefano Zampieri
szampier@eso.org

Agenda

- DDS Overview
- DDS on SPARTA
- Network considerations
- Conclusions

DDS in a nutshell

- DDS is a Real-Time Data-Centric Networking Middleware
- DDS focuses on
  - Performance
    - High-performance data-access APIs (zero copy access)
  - Configurability
    - Quality of Service
  - Scalability
    - UDP, multicast, reliable multicast
  - DDS does not require the presence of intermediate brokers
    - Applications can communicate directly peer-to-peer
  - DDS supports advanced features
    - E.g. source filtering (via Content-based and Time-based filters)
  - Integration
    - E.g. with Database Management Systems

DDS Standards

- Data Distribution Service for Real-Time Systems (DDS)
  - API specification for Data-Centric Publish-Subscribe communication for distributed real-time systems.
  - Current version 1.2
- DDS Interoperability wire Protocol (DDSI/RTPS)
  - Ensure that applications based on different vendors' implementations of DDS can interoperate.
  - Current version 2.1
- Related Standards
  - UML Profile for DDS adopted June 2008
  - DDS for light weight CCM adopted 2008
  - Extensible and Dynamic Topic Types for DDS adopted 2010
- Standards under Development
  - Native Language C++ API for DDS
  - DDS-Java

DDS Vendors

- Real-Time innovations, Inc. (Commercial, Open Community Source)
- PrismTech (Commercial & Open Source)
- Object Computing, Inc. (OpenDDS, Open Source)
- Twin Oaks Computing, Inc. (CoreDX, Commercial)
- Etc.

DDS Model

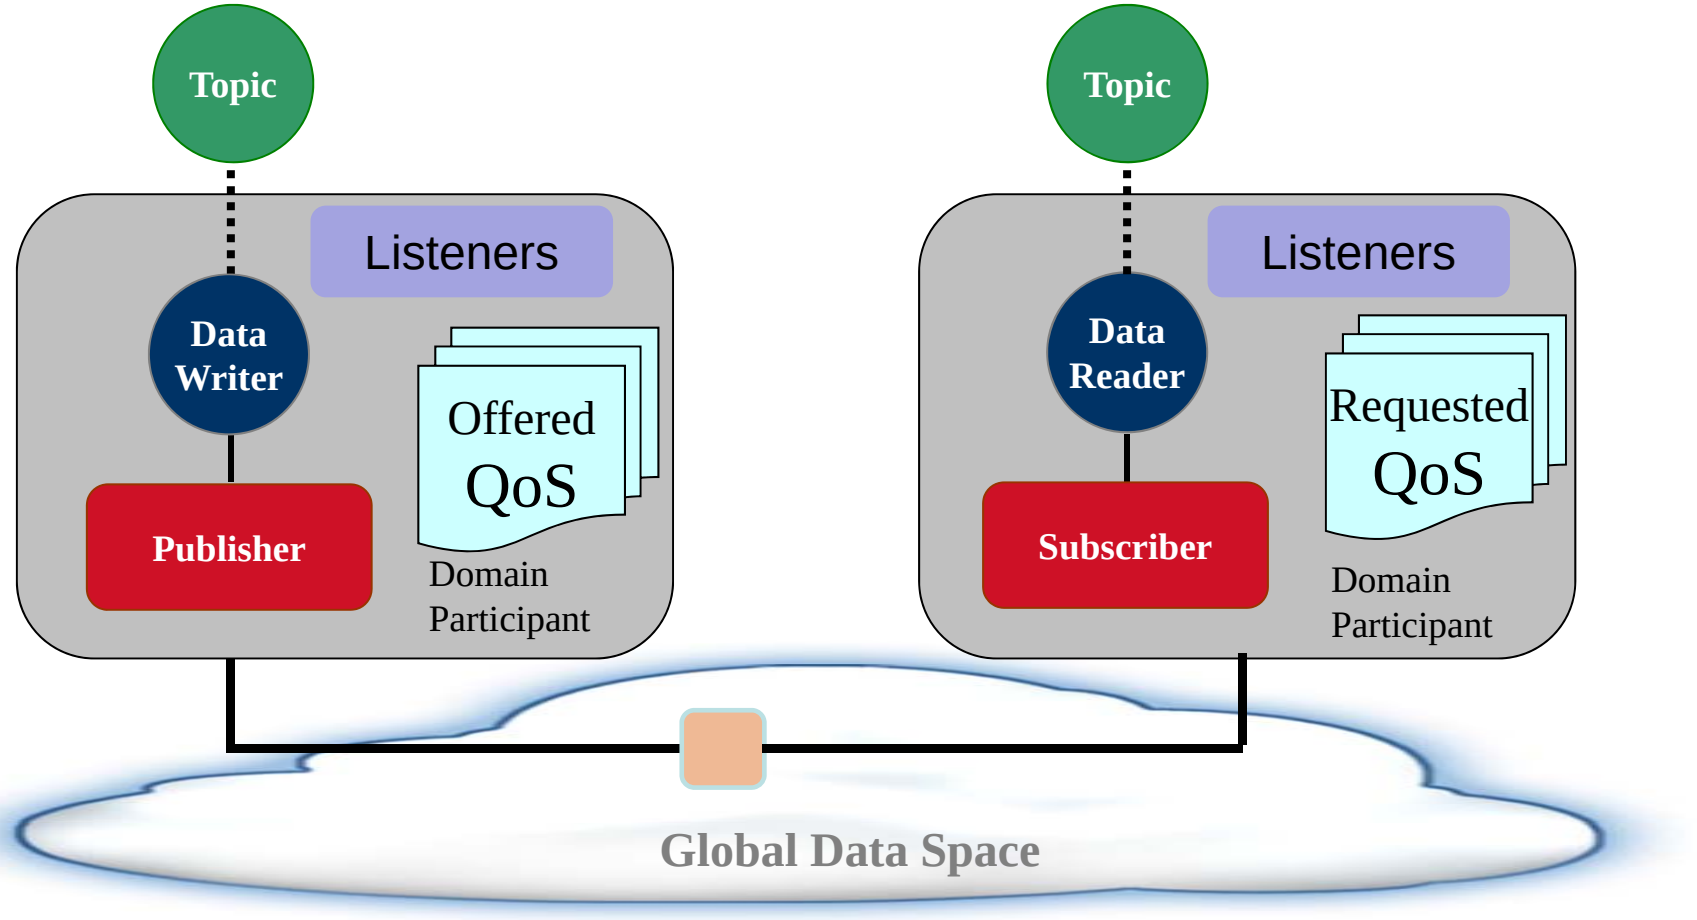
Topic
Listeners
Data Writer
Offered QoS
Publisher
Domain Participant
Topic
Listeners
Data Reader
Requested QoS
Subscriber
Domain Participant
Global Data Space

QoS: Quality of Service

QoS Policy	QoS Policy
DURABILITY	USER DATA
HISTORY	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

Volatility

Infrastructure

Delivery

User QoS

Presentation

Redundancy

Transport

Example QoS

```
<durability>
    <kind>DDS_TRANSIENT_LOCAL_DURABILITY_QOS</kind>
</durability>
```

```
<time_based_filter>
        <minimum_separation>
            <sec>1</sec>
            <nanosec>0</nanosec>
        </minimum_separation>
</time_based_filter>
```

```
<history>
        <kind>DDS_KEEP_ALL_HISTORY_QOS</kind>
</history>
```

```
<reliability>
        <kind>DDS_RELIABLE_RELIABILITY_QOS</kind>
            <max_blocking_time>
                <sec>0</sec>
                <nanosec>0</nanosec>
            </max_blocking_time>
</reliability>
```

Agenda

- DDS Overview
- DDS on SPARTA
- Network considerations
- Conclusions

SPARTA Real Time Data Flows

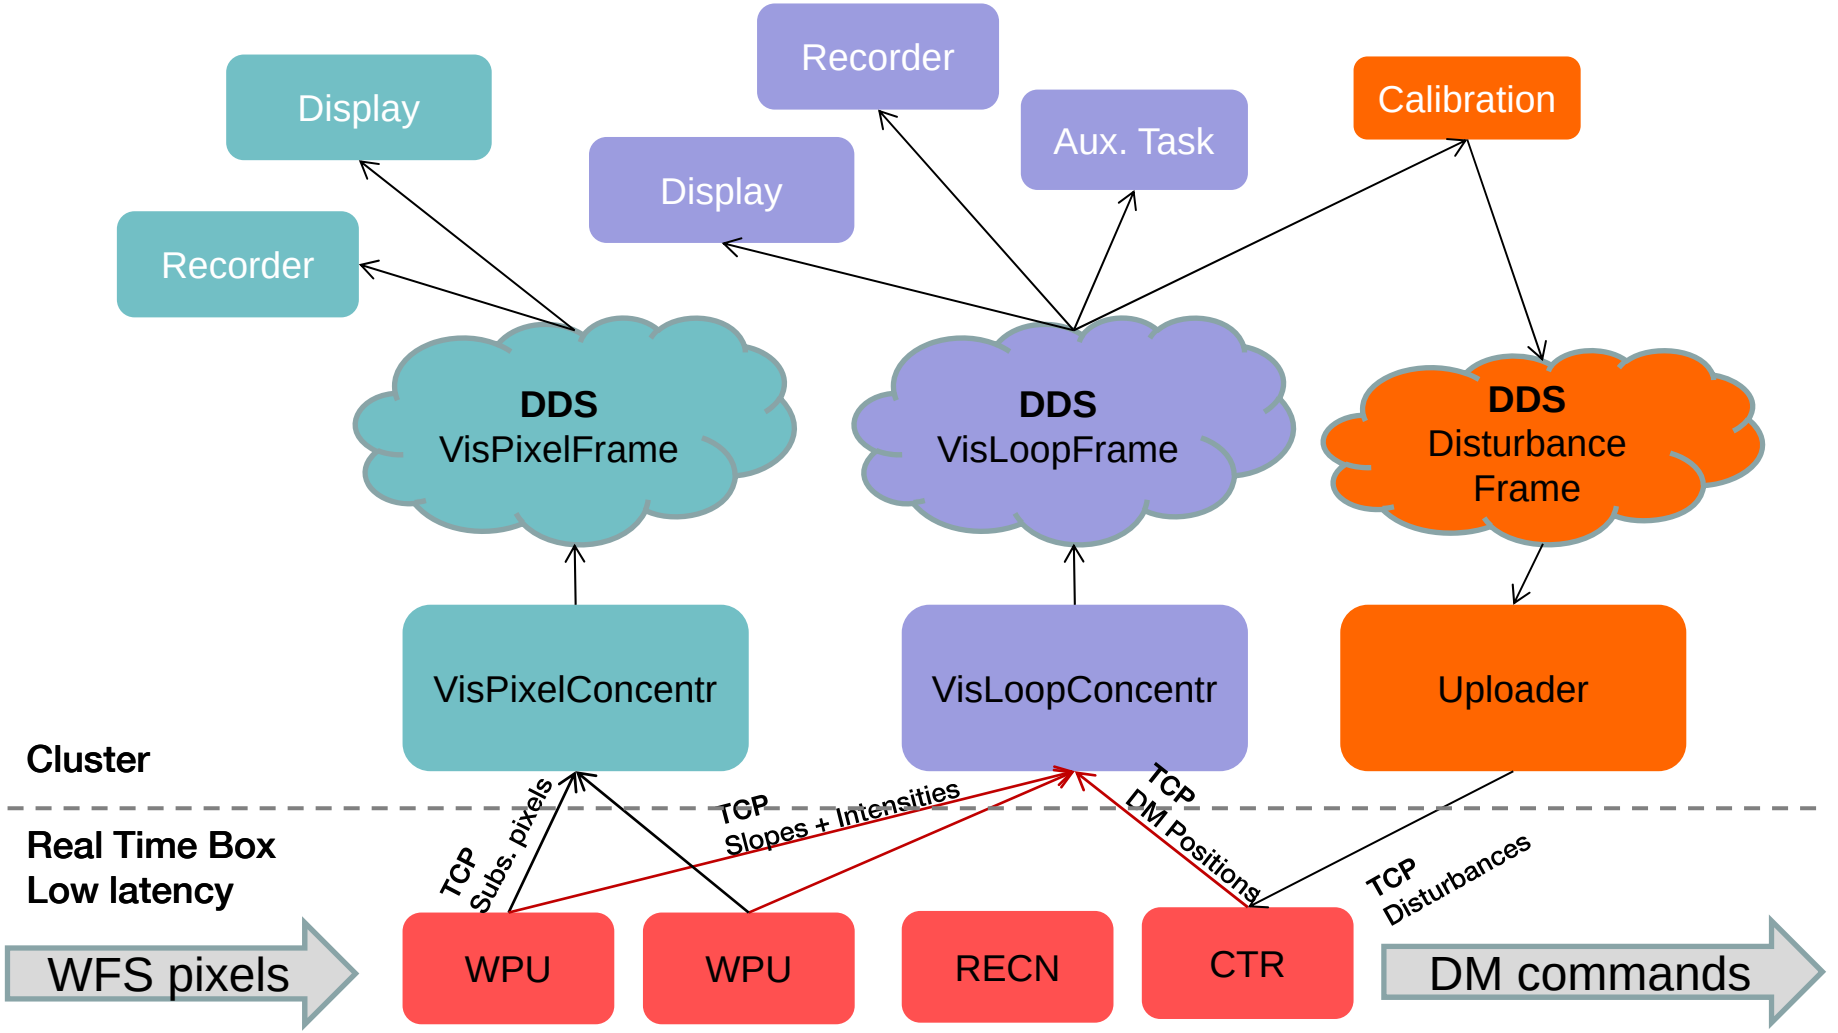

Additional Data Flows

- DB Events (n:m)
  - Database updates sent to DBGateway and to Main
  - Including Alarms
  - ~300 events/s (measured on SPHERE)
- CDMS Events (1:n)
  - Upon object updates in the SPARTA Cfg. DB
  - Trigger chain of events
- Log Events (n:1)
  - Log messages sent to LogGateway

Throughput requirements

- SAXO
  - VisLoop: 20KB @ 1.2KHz
  - VisPixel: 112KB @ 10Hz
  - Tot: ~25MB/s
- AOF
  - LGSLoop: 67KB @ 1KHz
  - LGSPixel: 450KB @ 10 Hz
  - Tot: ~72MB/s
- Multicast !

SPARTA DDS Model

- SPARTA DDS Wrapper (spadds)
  - Simplified API (DDS-like)
  - **Publisher** (*write*)
  - **Subscriber** + **DataListener** (*onDataAvailable*)
  - **Topic**: template parameter + string
  - QoS defined in XML configuration file
  - QoS Profiles, referenced by name when creating Publishers and Subscribers
    - *HighThroughputReliableProfile* (reliable, large send queue)
    - *LargePacketsReliableProfile* (>64KB, asynchronous publisher)
    - *ReliableEventProfile* (durability)
    - *PixelDisplayProfile* (time based filter)

SPARTA Data Task

- Simplifies development of data tasks
- Simple model: receive N samples then process them in a separate thread
- Developer must implement virtual methods *received_*, *process_*, and *deadlineMissed_*
- Examples: Garbage Collector, Loop Optimiser, Atmospheric Monitor, etc

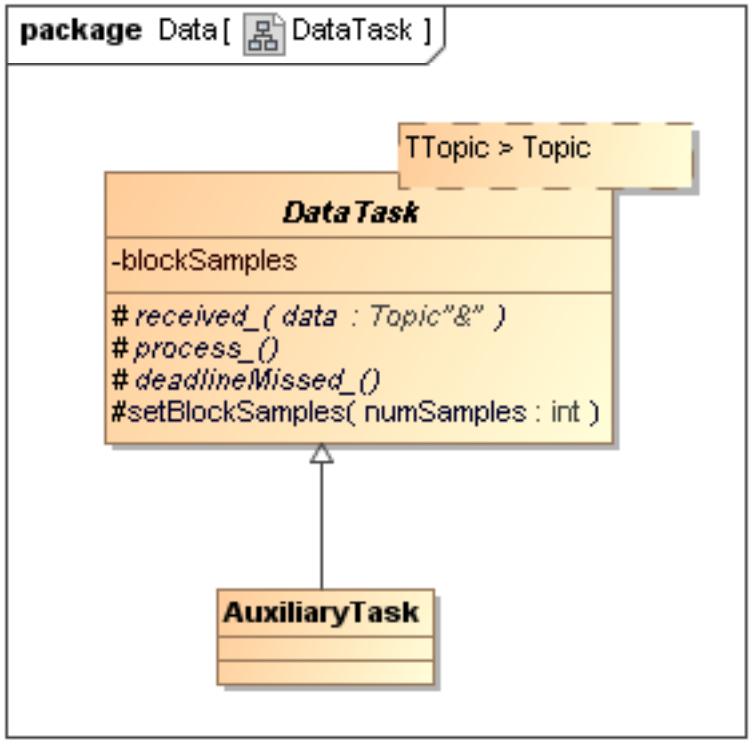

Agenda

- DDS Overview
- DDS on SPARTA
- Network considerations
- Conclusions

Multicast & IGMP

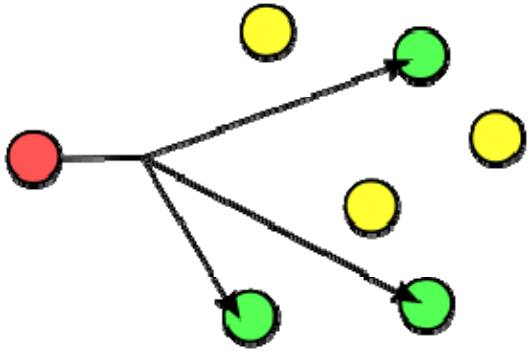

Multicast vs Broadcast

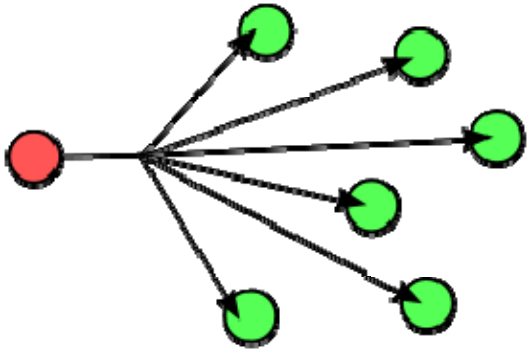

- IP Multicast
  - Take advantage of multicast efficiency in network
  - IP address range: 224.0.0.0 to 239.255.255.255.

- IGMP snooping switches
  - No IGMP snooping
    - Multicast traffic broadcasted to each port
  - IGM Snooping
    - Switch forwards multicast packets to correct ports
      - Monitors IGMP join messages
    - Multicast addresses configured by subscriber

Wireshark & RTPS2

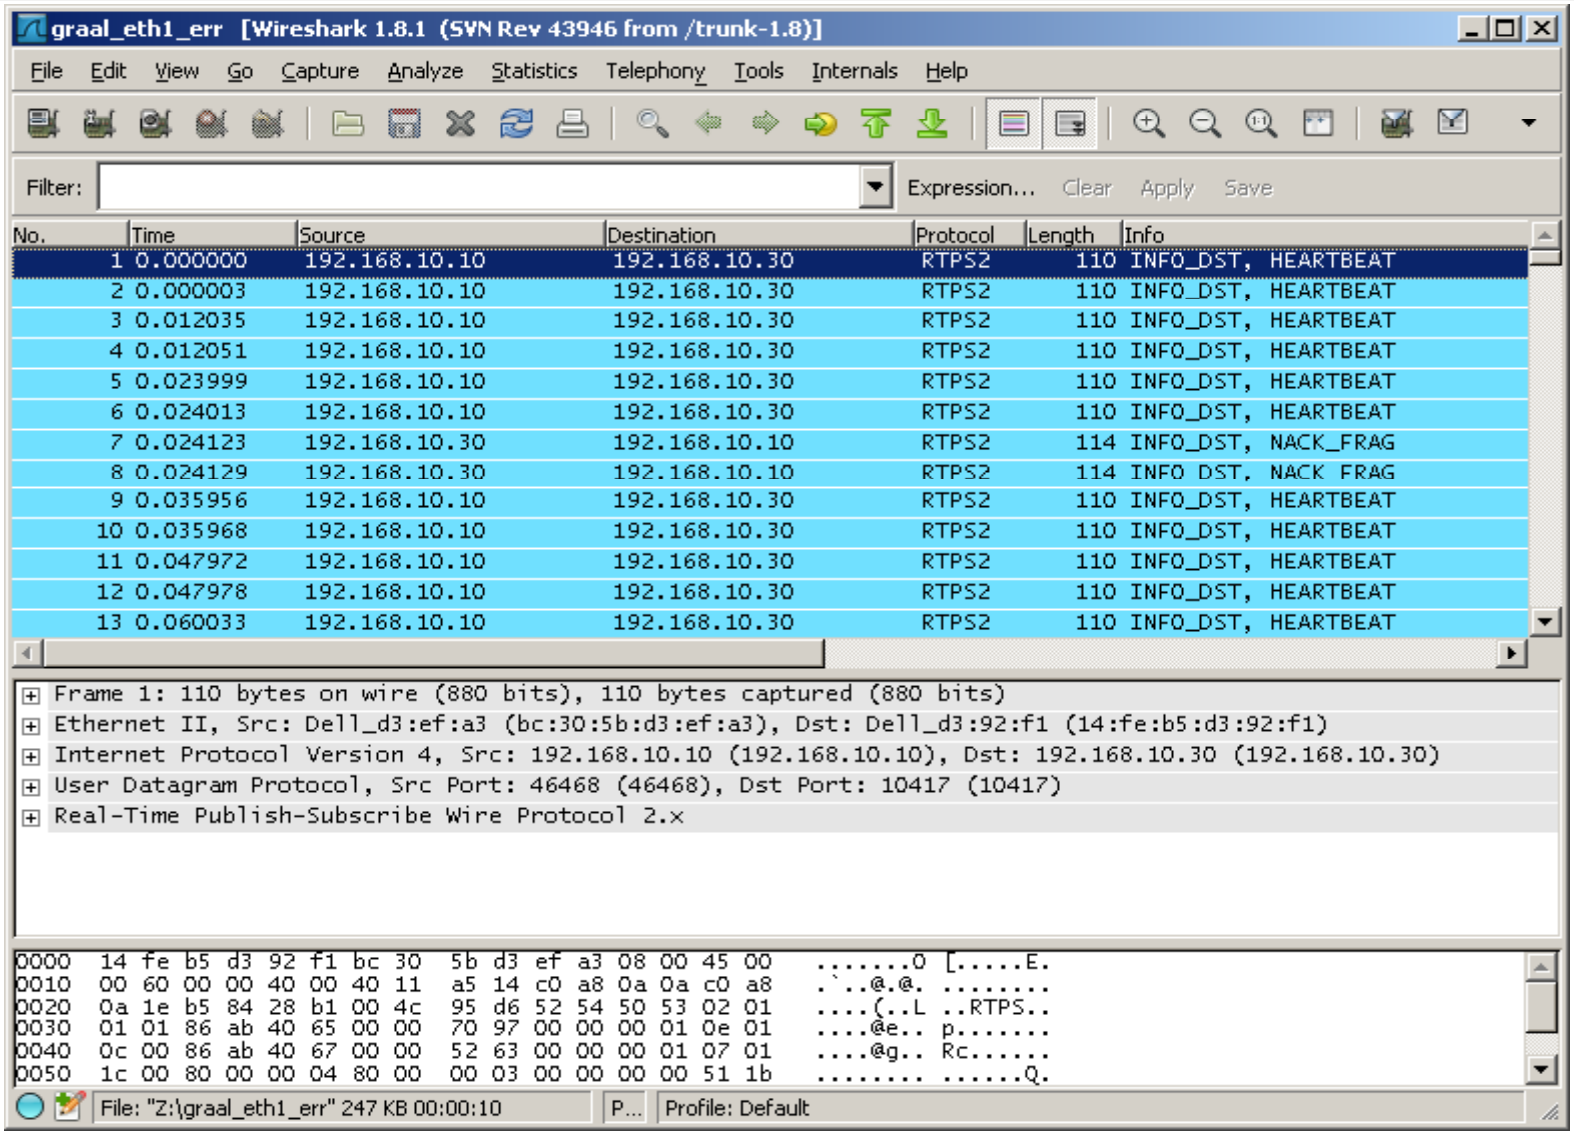

Scaling up

- Initial tests on 10 GigE using rtiperftest (no tuning)

	One way Latency (us)	Packets/s	Mb/s	Packet loss
Best effort Small packets	**105**	**130000**	104	Very low
Reliable Small packets	510?	35000	28	0
Best effort Large packets	357	11160	**5620**	About 1%
Reliable Large packets	372	10400	**5000**	0

- Jumbo frames ?

Agenda

- DDS Overview
- DDS on SPARTA
- Network stuff
- Conclusions

Conclusions

- DDS works (and saves development)
  - Only 1 serious issue up to now, solved by upgrading
  - Reliable (no intermediate brokers), efficient
  - Simple programming model, also thank to wrapper API
  - Highly configurable, through external QoS

- Future perspectives
  - DDS/RTPS on Real-Time Box ?

Questions

Thank You !

References

http://portals.omg.org/dds/sites/default/files/DDS_Tutorial_RT_Worskshop_2010.pdf