Angelo Corsaro, Ph.D.
Chief Technology Officer
OMG DDS SIG Co-Chair
angelo.corsaro@prismtech.com

PRISMTECH
Powering Netcentricity

Getting Started with the Community Ed.

OpenSplice | DDS

Delivering Performance, Openness, and Freedom

Background

Addressing Data Distribution Challenges

The OMG DDS Standard

- Introduced in 2004 to **address the Data Distribution challenges** faced by a wide class of **Defense and Aerospace Applications**
- Key requirement for the standard were its ability to **deliver very high performance** while seamlessly **scaling** from **embedded to ultra-large-scale deployments**
- Today **recommended** by **key administration worldwide** and **widely adopted** across several different application domains, such as, Automated Trading, Simulations, SCADA, Telemetry, etc.

DDS is standard designed to address the data-distribution challenges across a wide class of Defense and Aerospace Applications

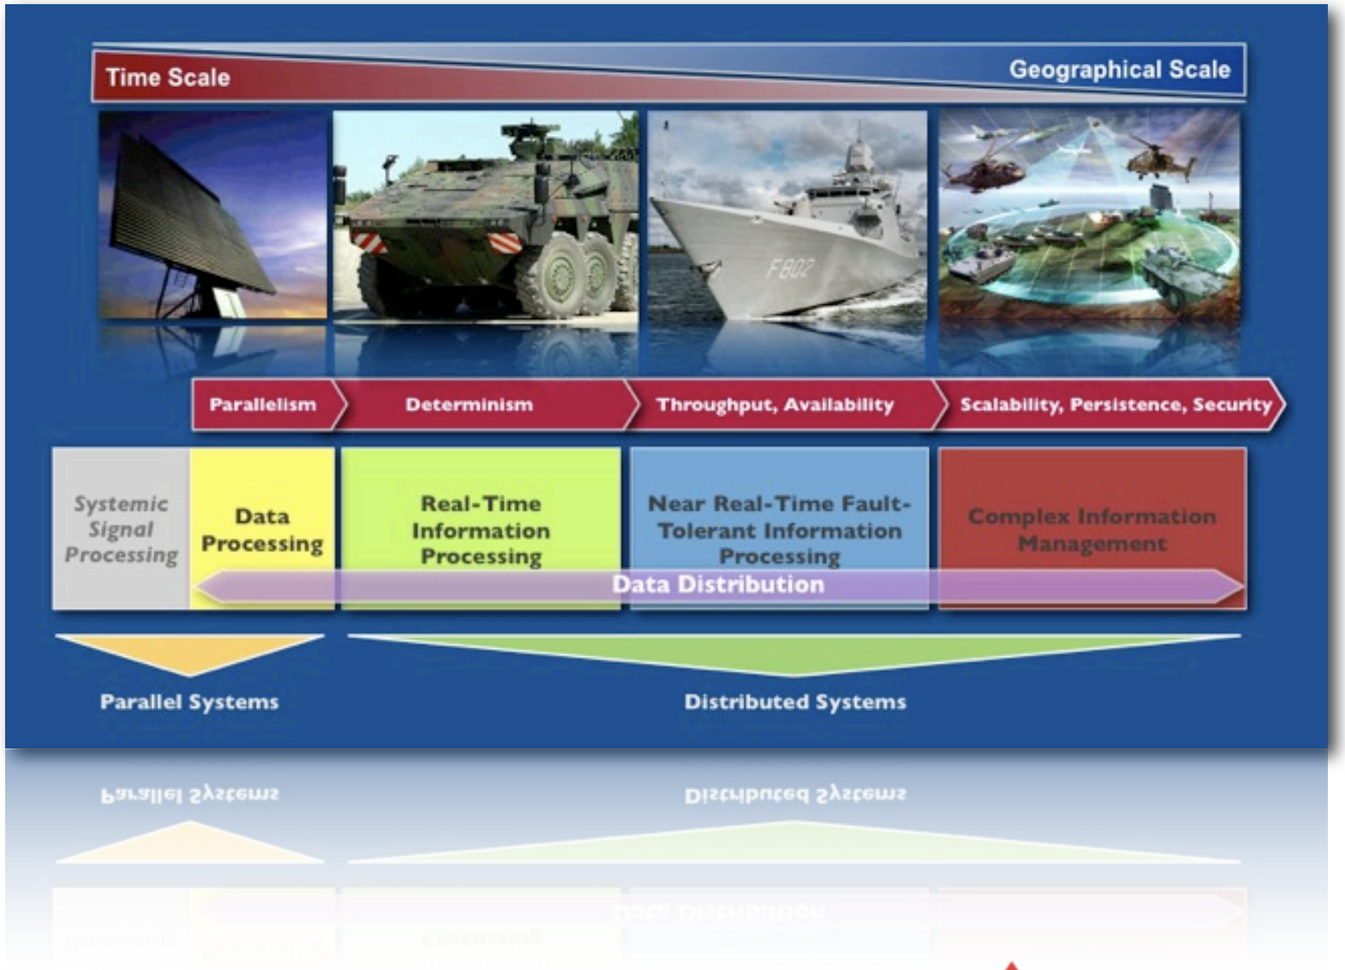

OpenSplice DDS

PRISMTECH

The OMG Data Distribution Service (DDS)

DDS v1.2 API Standard

- Language Independent, OS and HW architecture independent
- **DCPS.** Standard API for Data-Centric, Topic-Based, Real-Time Publish/Subscribe
- **DLRL.** Standard API for creating Object Views out of collection of Topics

DDSI/RTPS v2.1 Wire Protocol Standard

- Standard wire protocol allowing interoperability between different implementations of the DDS standard
- Interoperability demonstrated among key DDS vendors in March 2009

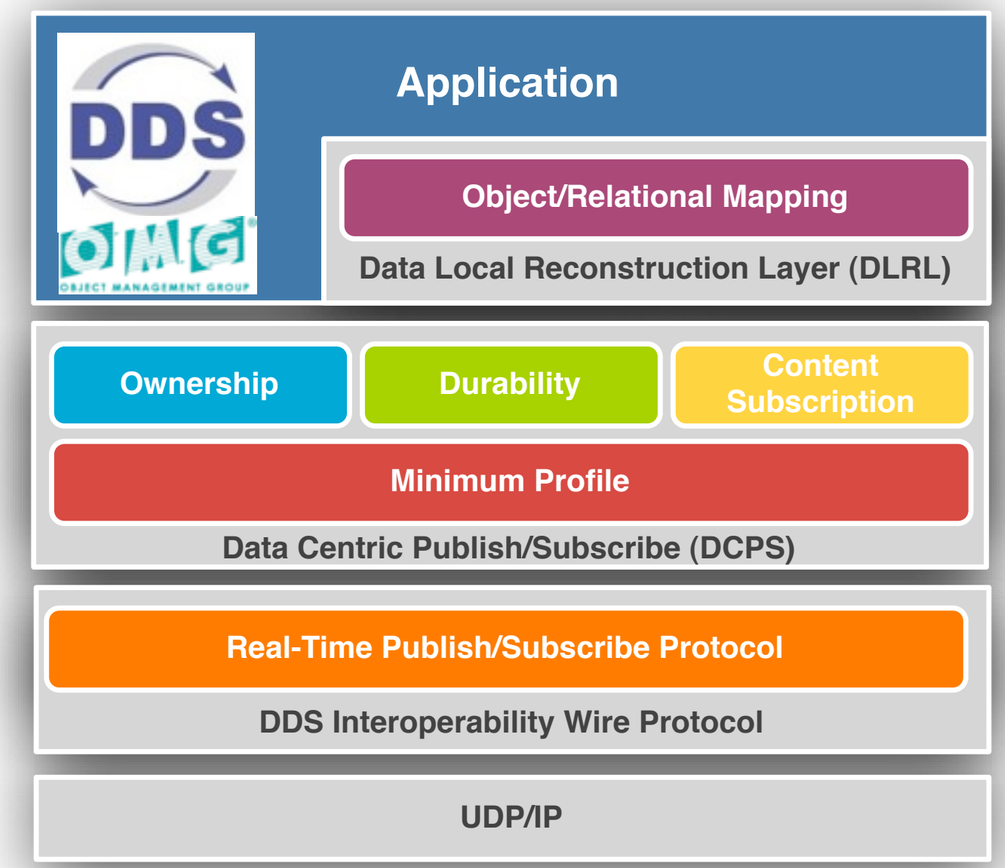

OpenSplice|DDS

PRISMTECH

OpenSplice DDS Community Ed.

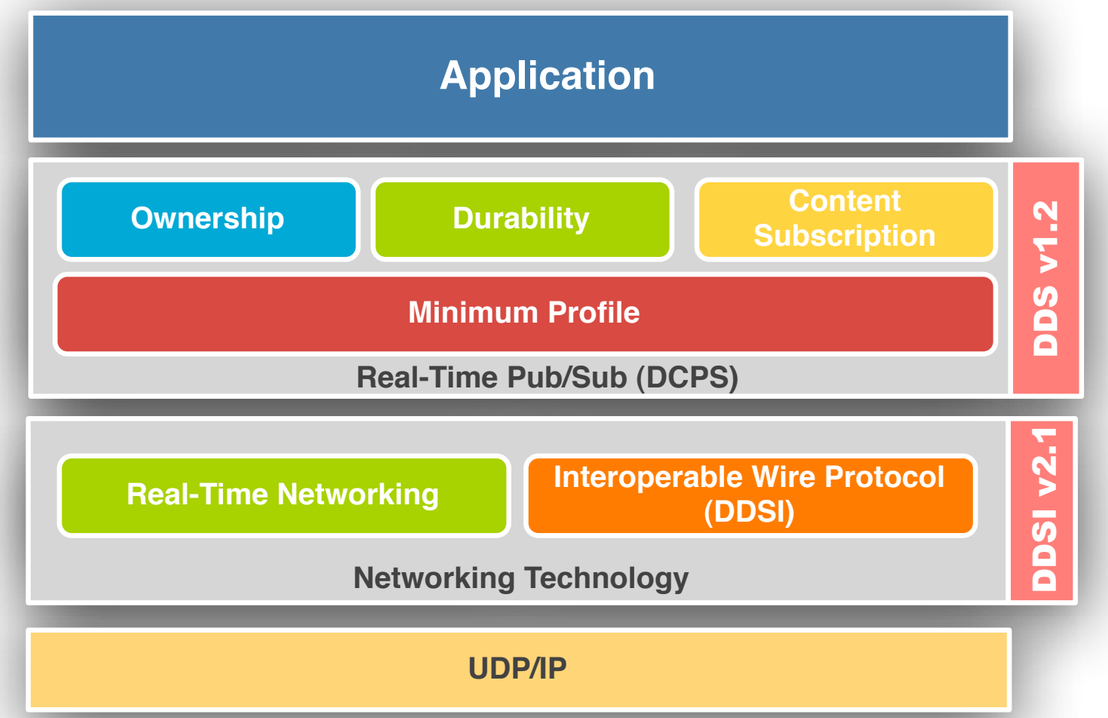

Features

- OMG DDS v1.2 DCPS
  - Minimum Profile
  - Content Subscription Profile
  - Durability Profile
  - Ownership Profile
- Networking
  - DDSI v2.1 Implementation
  - Real-Time Networking Implementation

Licensing

LGPL Free Software

	OMG DDS Standard Compliance					
	DCPS Profiles				DLRL	DDSI/ RTPS
	Minimum	*Content*	*Ownership*	*Durability*		
OpenSplice DDS Community Ed.	Yes	Yes	Yes	Yes	No	Yes
Other DDS (Best Case)	Yes	Partial	Yes	No*	No	Yes

OpenSplice DDS

PRISMTECH

Performance on Commodity HW

Throughput

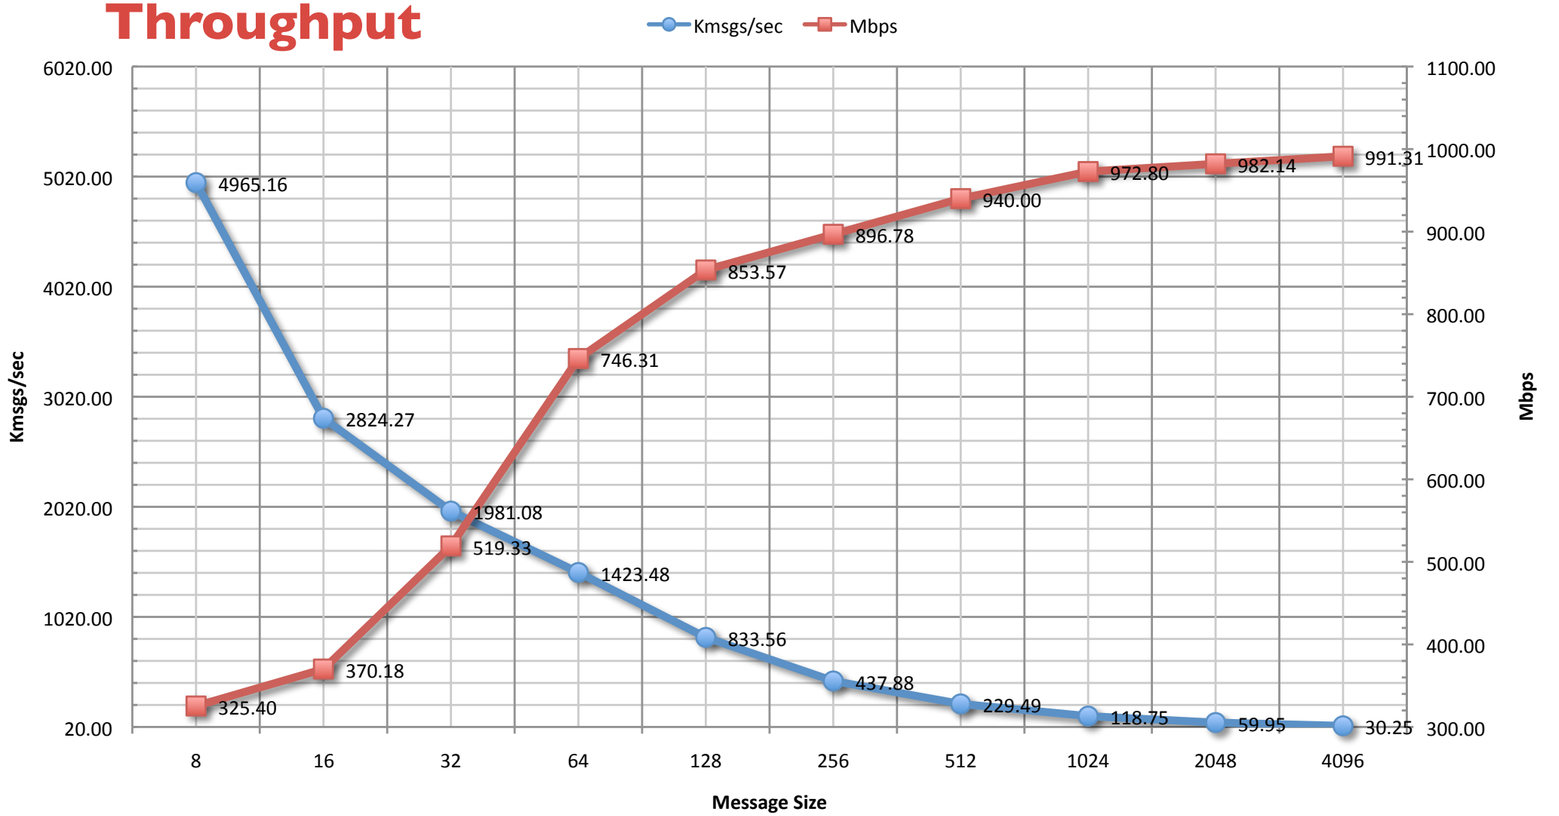

Test Scenario

- Single Threaded Application (multi-threaded networking service)
- 8192 bit message batches

Latency

Inter-Node Latency

- 60 usec

Inter-Core Read-Latency

- 2 usec

Inter-Core Latency

- <10 usec

HW:

- Dell blade-server
- Dual-core, Dual-CPU, AMD Opteron 2.4 Ghz

OS

- Linux 2.6.21-1.3194.fc7

Network

- Gigabit Ethernet cards
- Dell PowerConnect 5324 switch

OpenSplice DDS

PRISMTECH

OpenSplice | DDS

Delivering Performance, Openness, and Freedom

As Simple as it Gets

As Simple as it Gets

- DDS is based around the concept of a **fully distributed Global Data Space (GDS)**

DDS

As Simple as it Gets

- DDS is based around the concept of a **fully distributed Global Data Space (GDS)**
- **Publishers** and **Subscribers** can join and leave the GDS at any time

DDS

As Simple as it Gets

- DDS is based around the concept of a **fully distributed Global Data Space (GDS)**
- **Publishers** and **Subscribers** can join and leave the GDS at any time

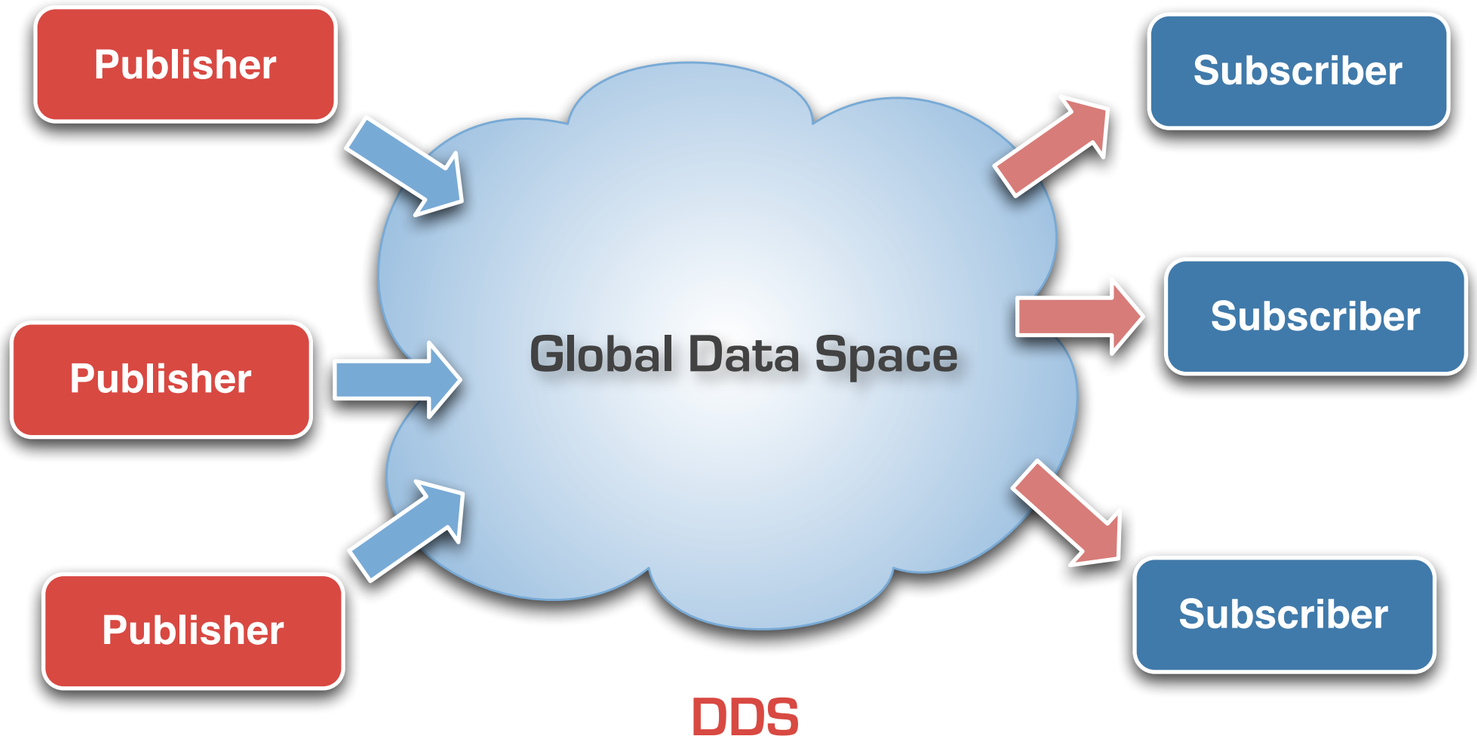

As Simple as it Gets

- DDS is based around the concept of a **fully distributed Global Data Space (GDS)**
- **Publishers** and **Subscribers** can join and leave the GDS at any time
- **Publishers and Subscribers** express their intent to **produce/consume specific type of data**, e.g., **Topics**

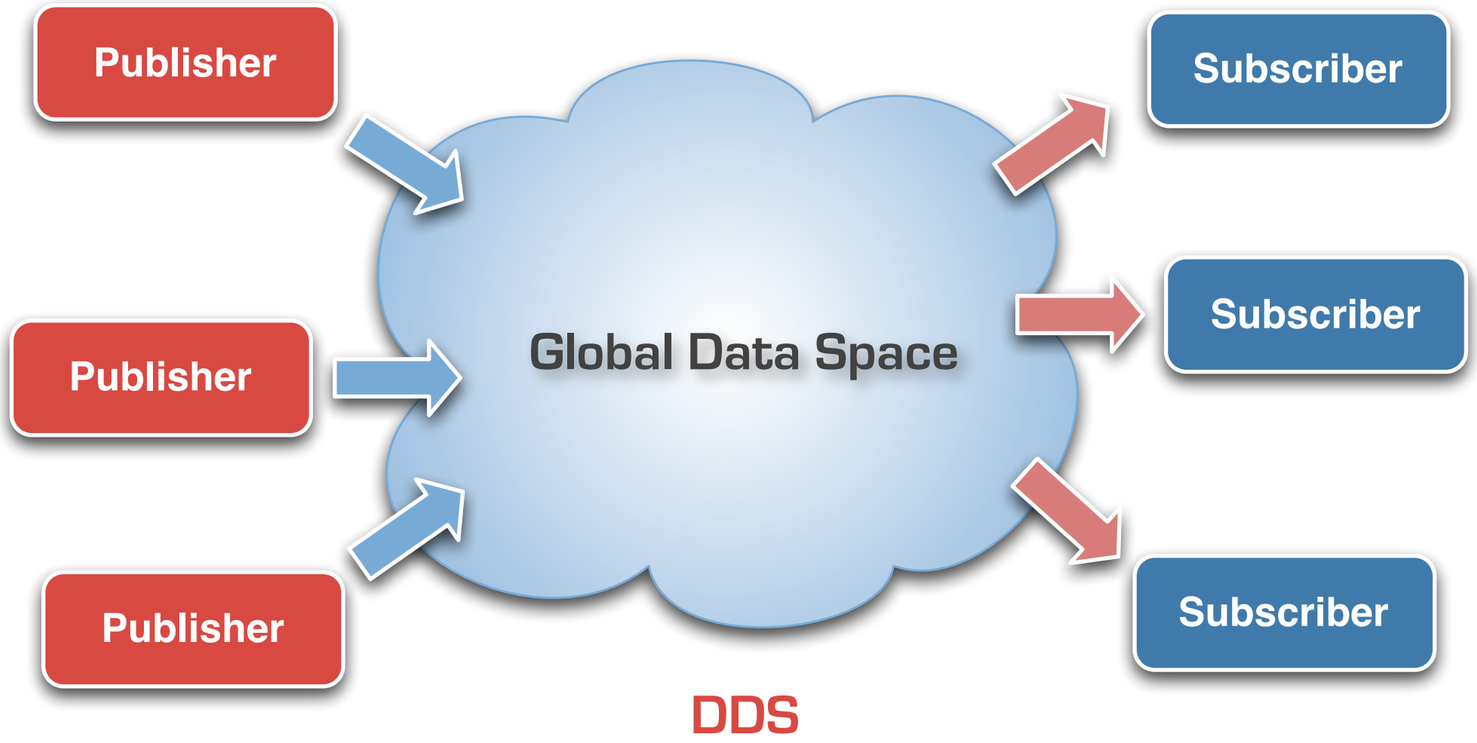

As Simple as it Gets

- DDS is based around the concept of a **fully distributed Global Data Space (GDS)**
- **Publishers** and **Subscribers** can join and leave the GDS at any time
- **Publishers and Subscribers** express their intent to **produce/consume specific type of data**, e.g., **Topics**

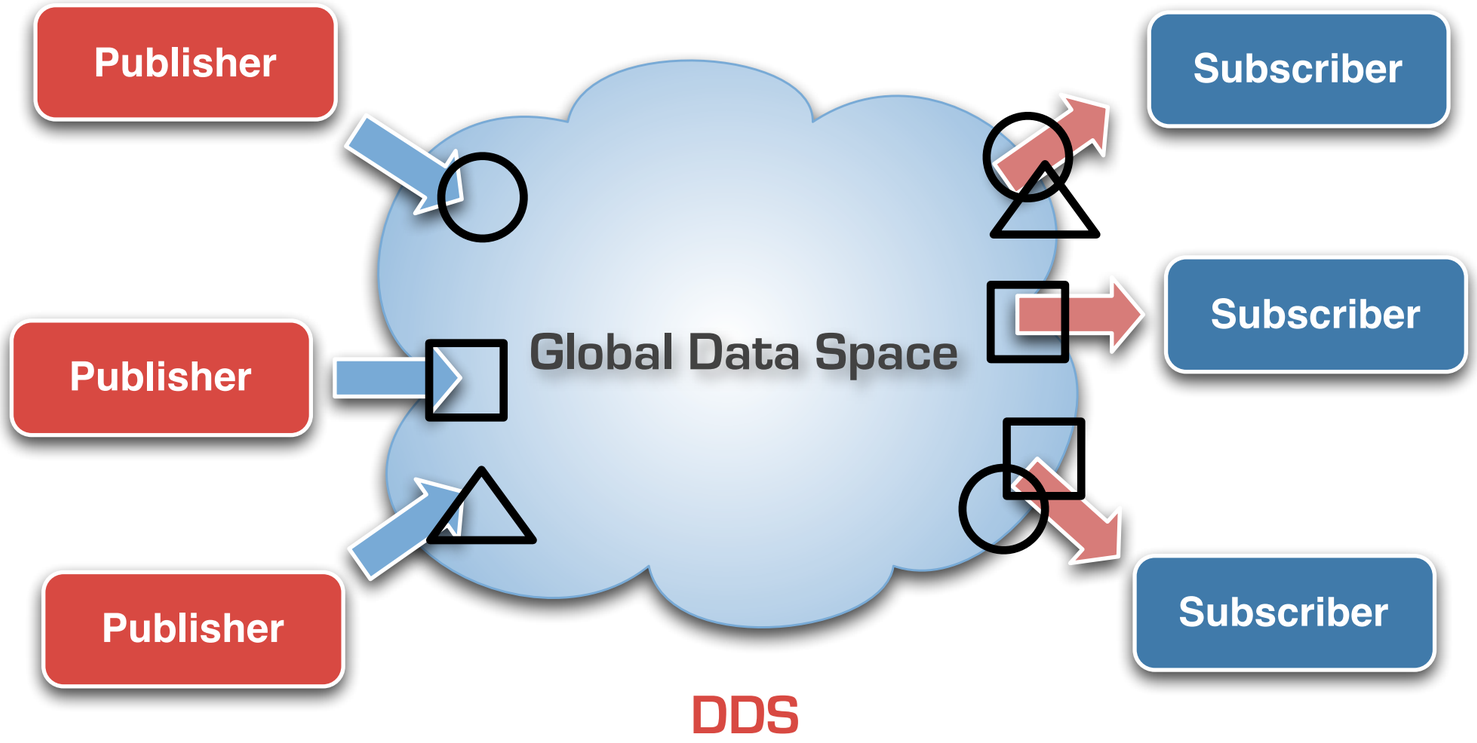

As Simple as it Gets

- DDS is based around the concept of a **fully distributed Global Data Space (GDS)**
- **Publishers** and **Subscribers** can join and leave the GDS at any time
- **Publishers and Subscribers** express their intent to **produce/consume specific type of data**, e.g., **Topics**
- **Data flows from Publisher to Subscribers**

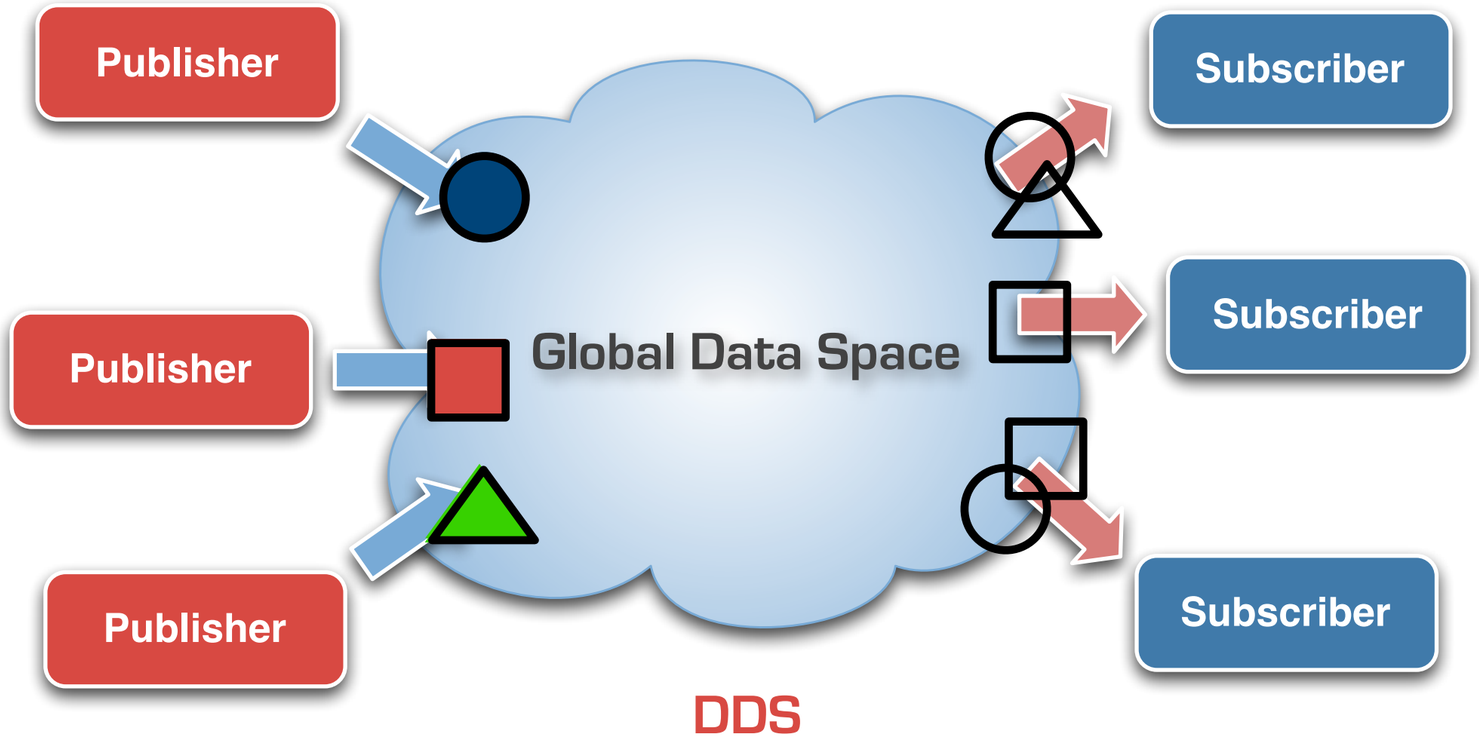

As Simple as it Gets

- DDS is based around the concept of a **fully distributed Global Data Space (GDS)**
- **Publishers** and **Subscribers** can join and leave the GDS at any time
- **Publishers and Subscribers** express their intent to **produce/ consume specific type of data**, e.g., **Topics**
- **Data flows from Publisher to Subscribers**

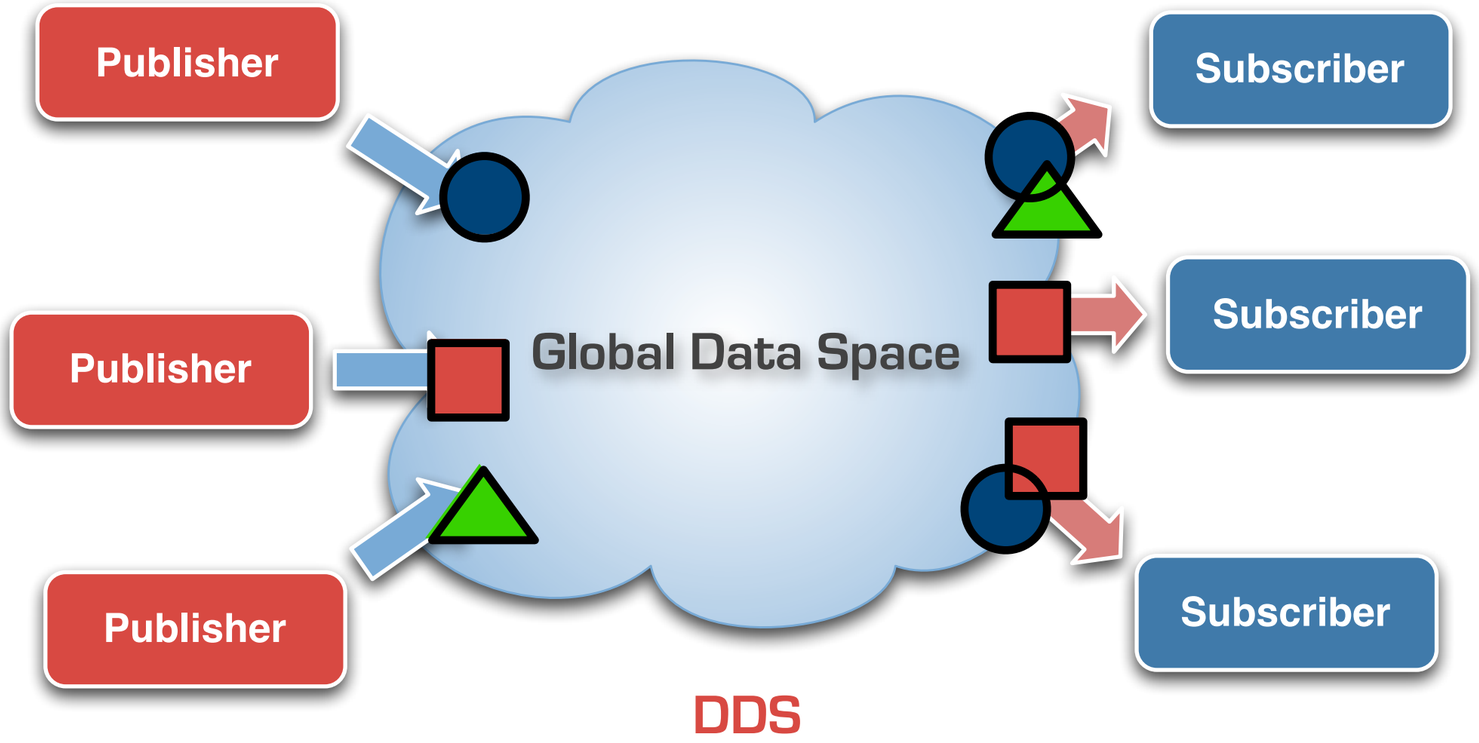

OpenSplice | DDS

Delivering Performance, Openness, and Freedom

Defining Data

DDS Topics

Topic

- Unit of information exchanged between Publisher and Subscribers.
- An association between a unique name, a type and a QoS setting

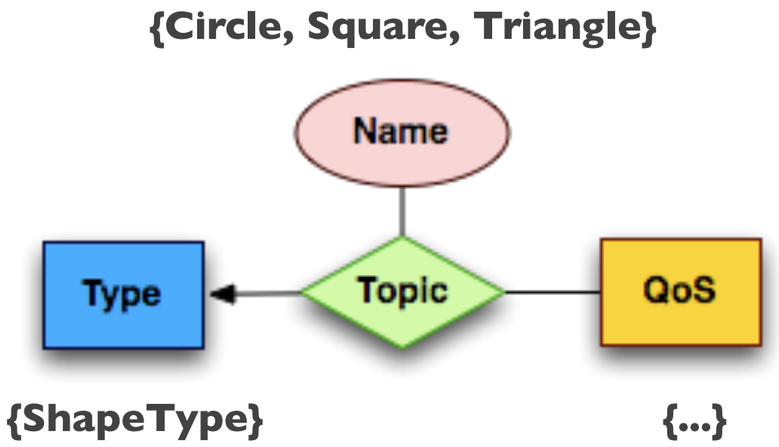

Topic Type.

- Type describing the data associated with one or more Topics
- A Topic type can have a key represented by an arbitrary number of attributes
- Expressed in IDL

```
struct ShapeType {
    long   x;
    long   y;
    long   shapesize;
    string color;
};
#pragma keylist ShapeType color
```

DDS Topics

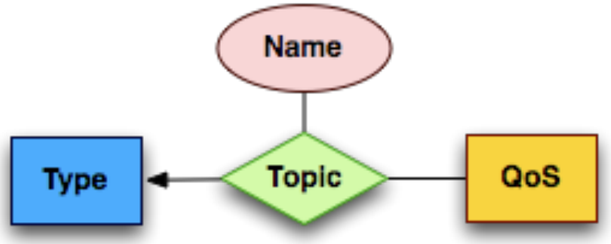

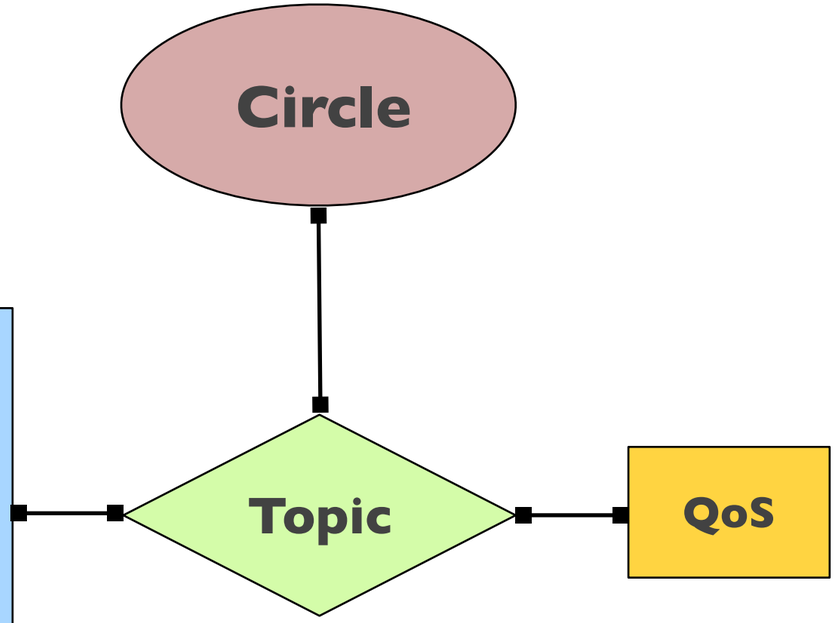


```
struct ShapeType {
    long   x;
    long   y;
    long   shapesize;
    string color;
};
#pragma keylist ShapeType color
```

DDS Topics

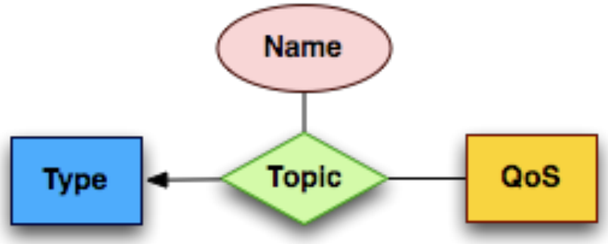

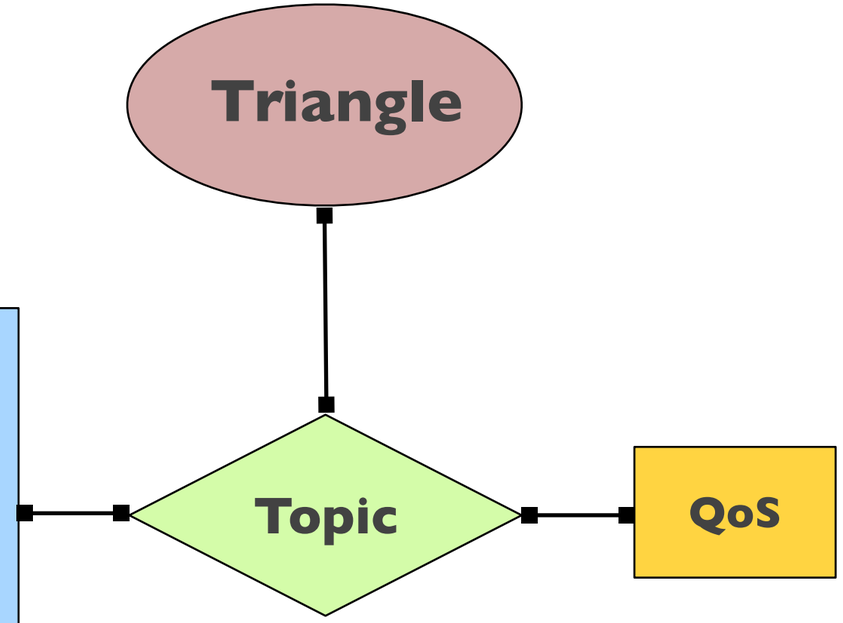


```
struct ShapeType {
    long   x;
    long   y;
    long   shapesize;
    string color;
};
#pragma keylist ShapeType color
```

DDS Topic Instances and Samples

Topic Instances

- Each key value identifies a unique **Topic Instance**,
- Topic's instance lifetime can be explicitly managed in DDS

Topic Samples

- The values assumed by a **Topic Instance** over time are referred as **Instance Sample**

```
struct ShapeType {
    long    x;
    long    y;
    long    shapesize;
    string color;
};
#pragma keylist ShapeType color
```

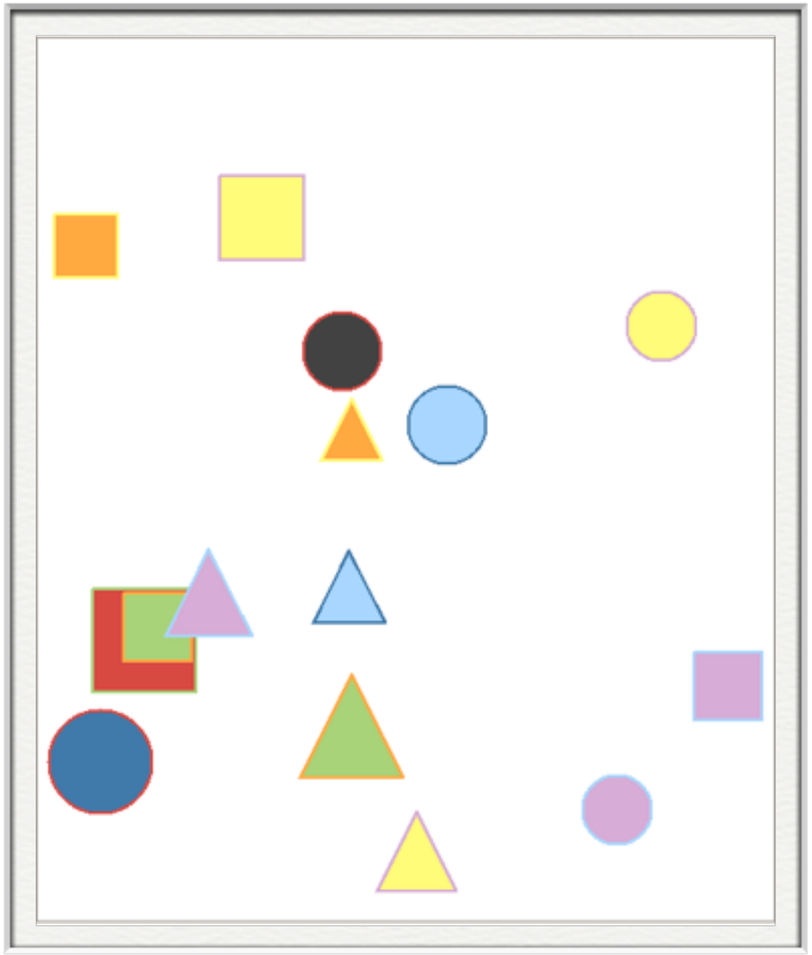

Topic/Instances/Samples Recap.

Topics

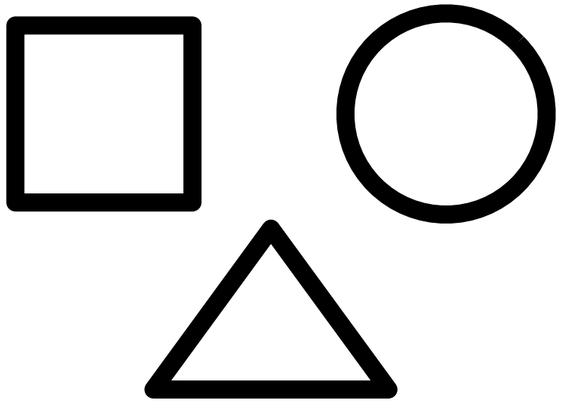

Instances

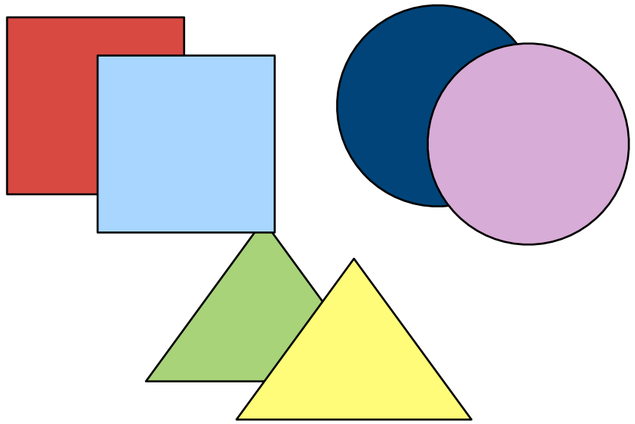

Samples

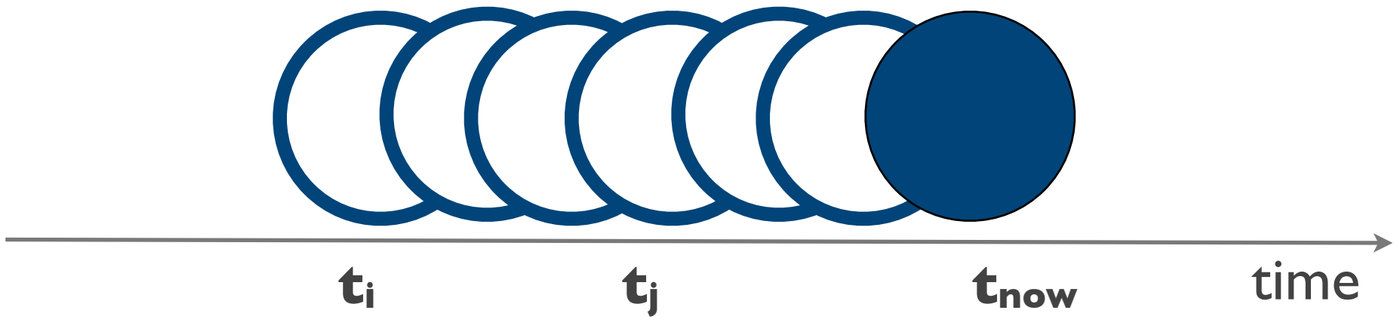

Content Filtering

- DDS allows to specify **content-filtered Topics** for which a subset of SQL92 is used to express the filter condition
- Content filters can be applied on the entire content of the Topic Type
- Content filters are applied by DDS each time a new sample is produced/delivered

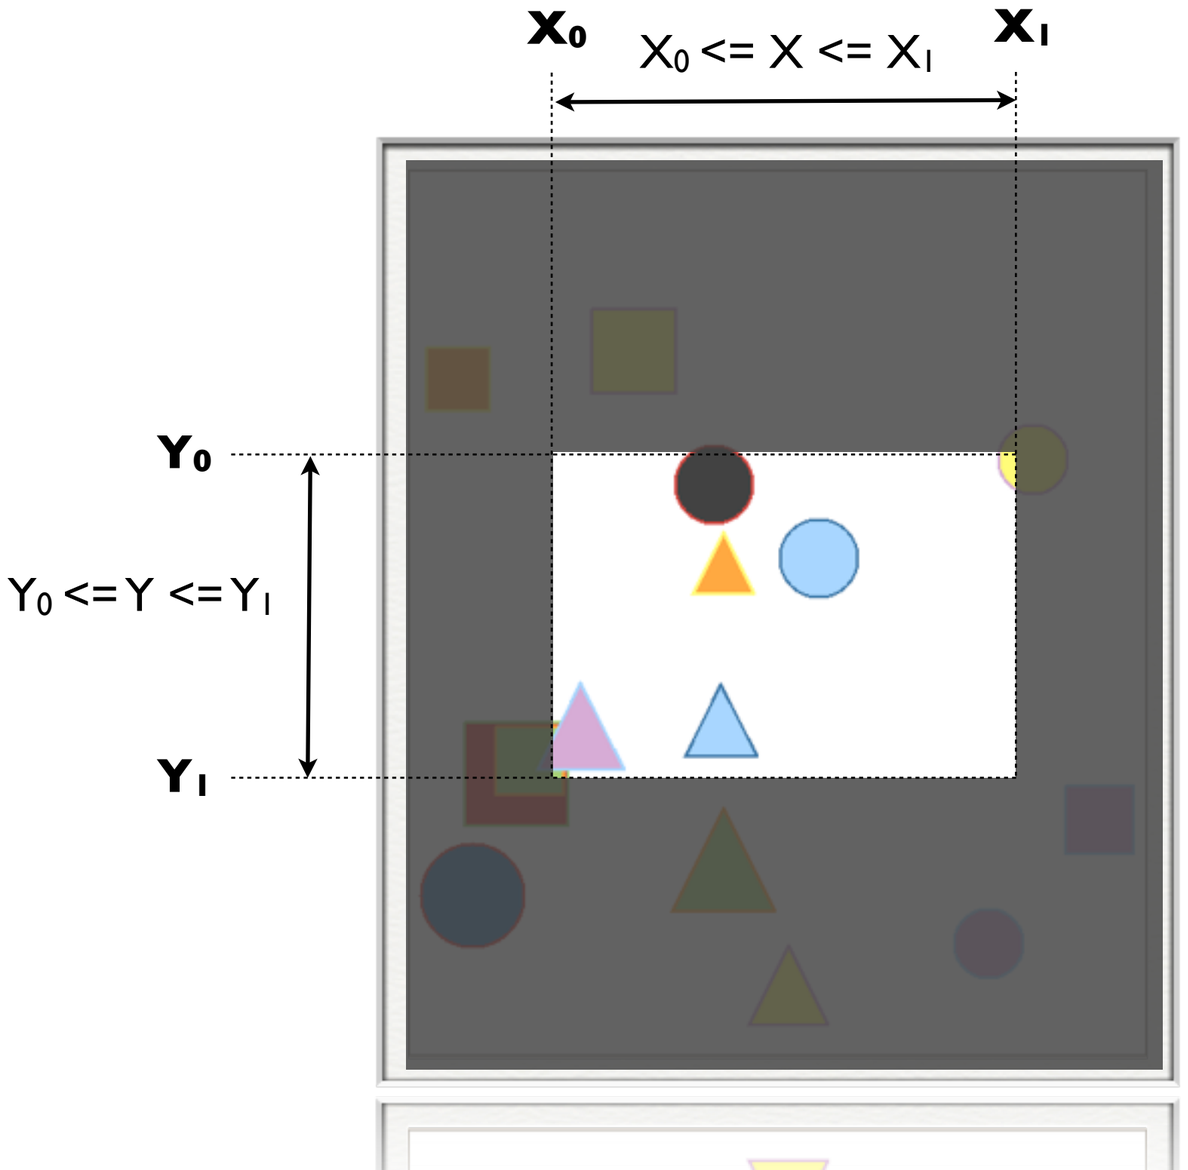

Local Queries

- Subscribed Topics can be seen locally as “Tables”
- A subset of SQL92 can be used for performing queries on multiple topics as well as natural joins
- Queries are performed under user control and provide a result that depends on the current snapshot of the system, e.g., samples currently available

Circle Topic

color	x	y	shapesize
red	57	62	50
blue	90	85	50
yellow	30	25	50

SELECT * FROM ShapeType s
WHERE s.x > 25 AND s.y < 55

color	x	y	shapesize
yellow	30	25	50

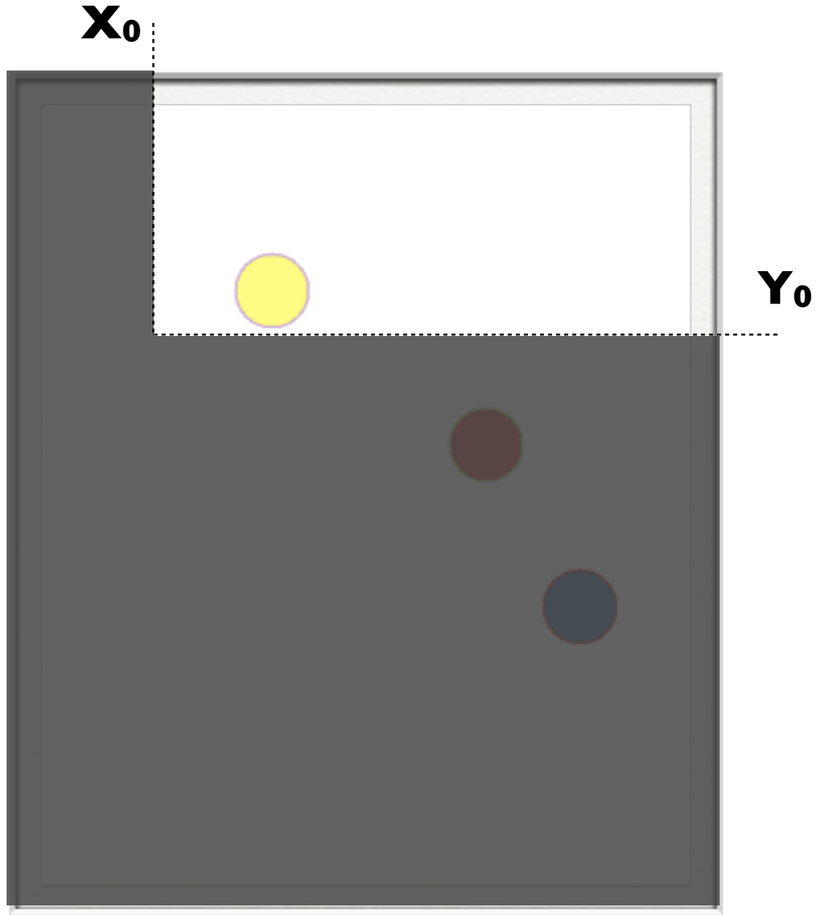

OpenSplice | DDS

Delivering Performance, Openness, and Freedom

Organizing Data

DDS Partitions

- All DDS communication is happens within a **Domain**
- Domain can divided into **Partitions**
- **Topics** are published and subscribed across on or more Partitions

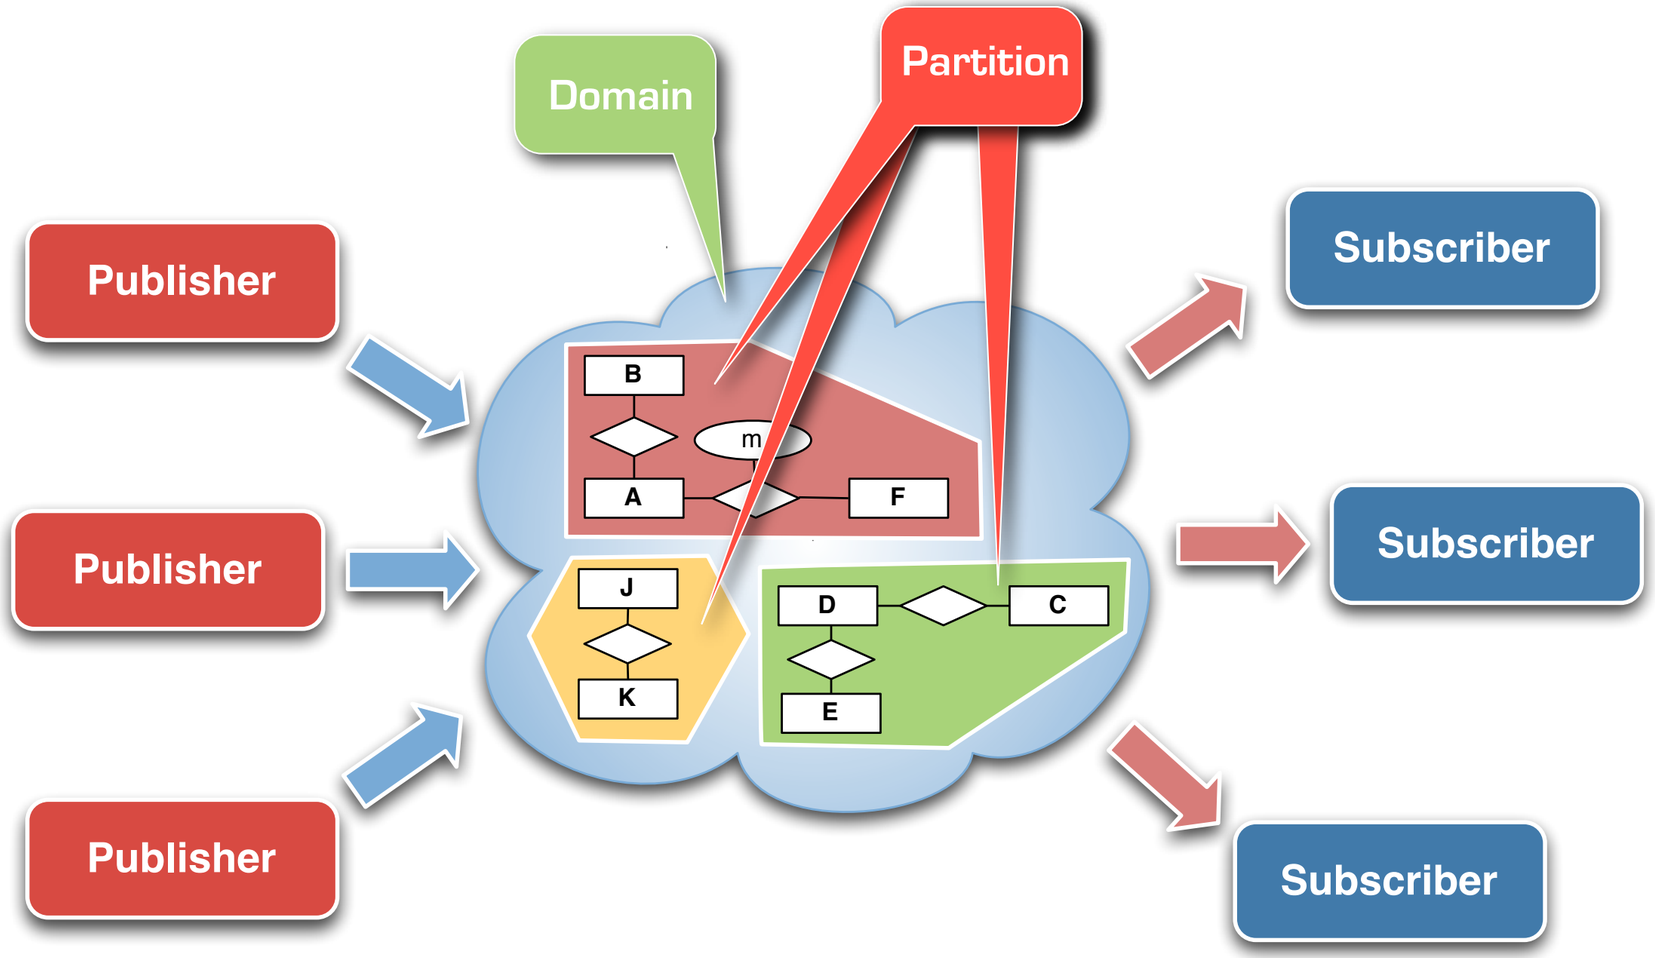

OpenSplice Network Partitions

- OpenSplice DDS allows to define **network partitions** along with **DDS partitions**
- Network partitions are bound to a list of unicast/multicast network addresses
- Partition.Topic combination can be mapped into OpenSplice DDS Network Partitions
- Wildcards can be used when defining the mapping, and in case of multiple matches OpenSplice DDS will always consider the best match

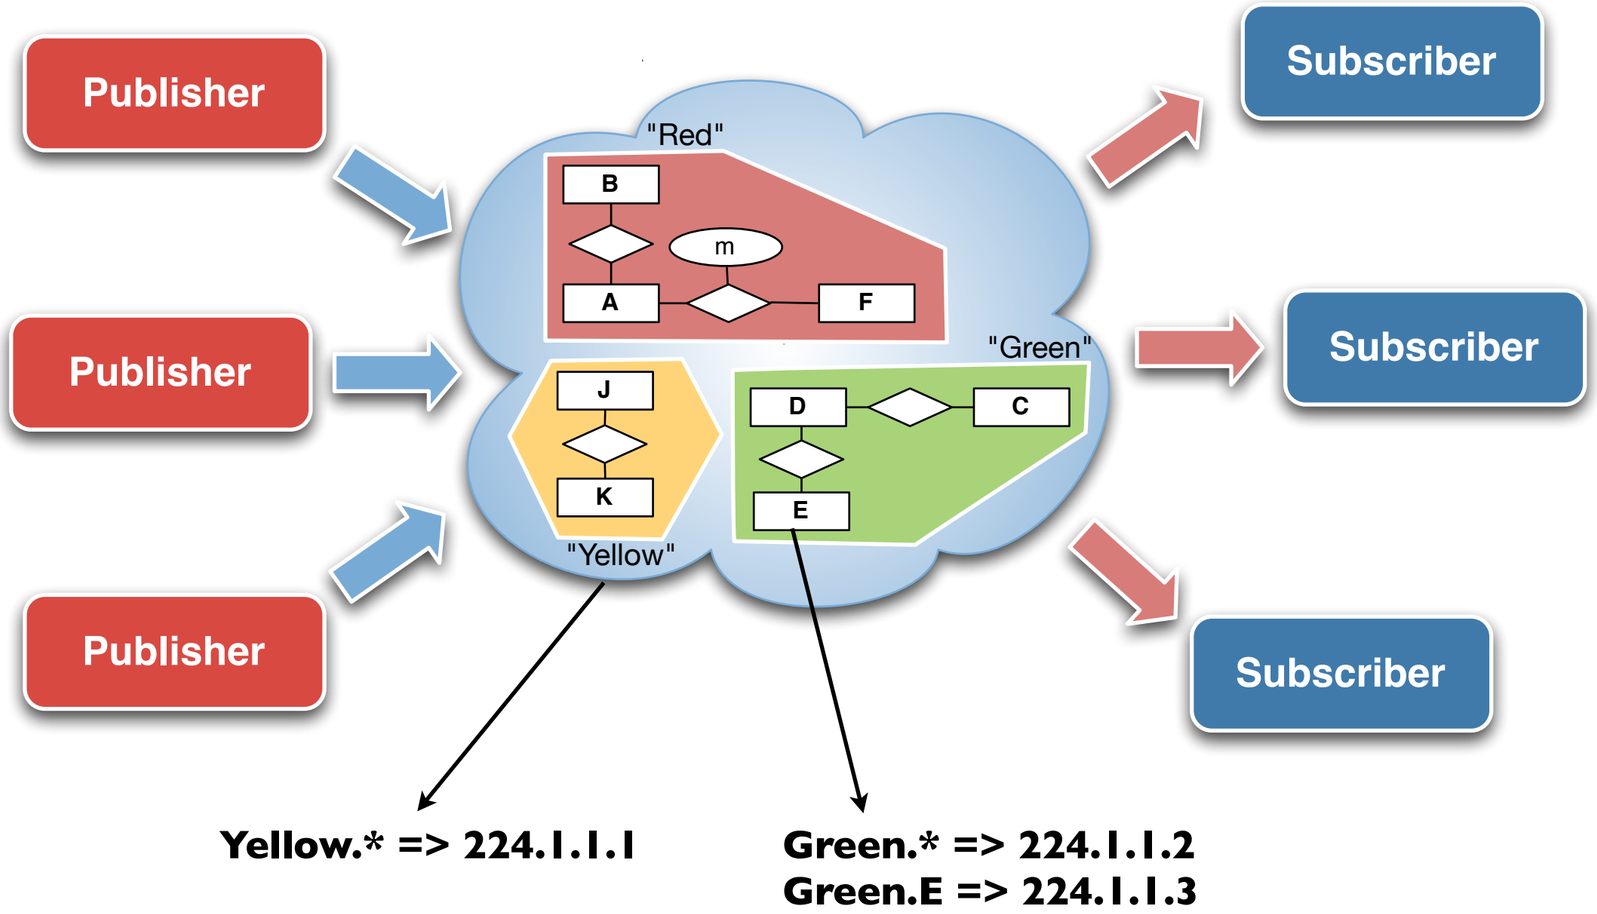

Network Partition (osplconf)

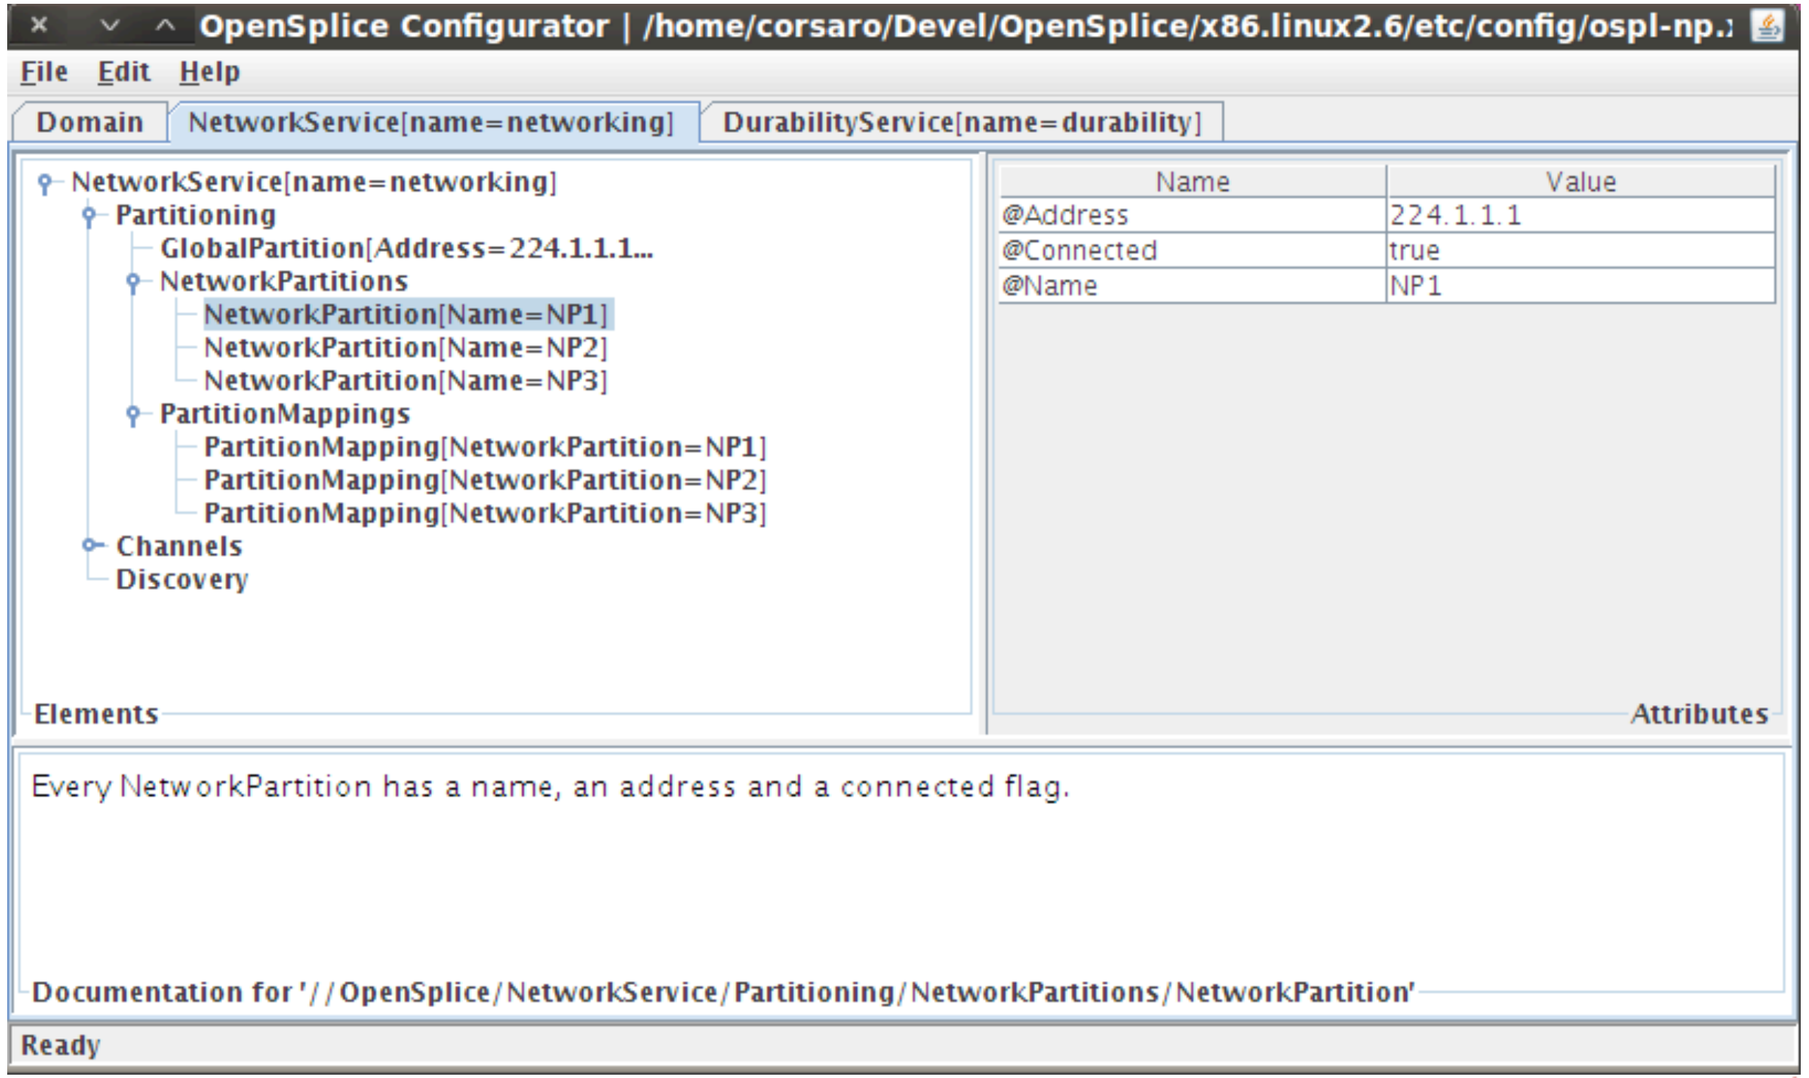

Partition Mapping (osplconf)

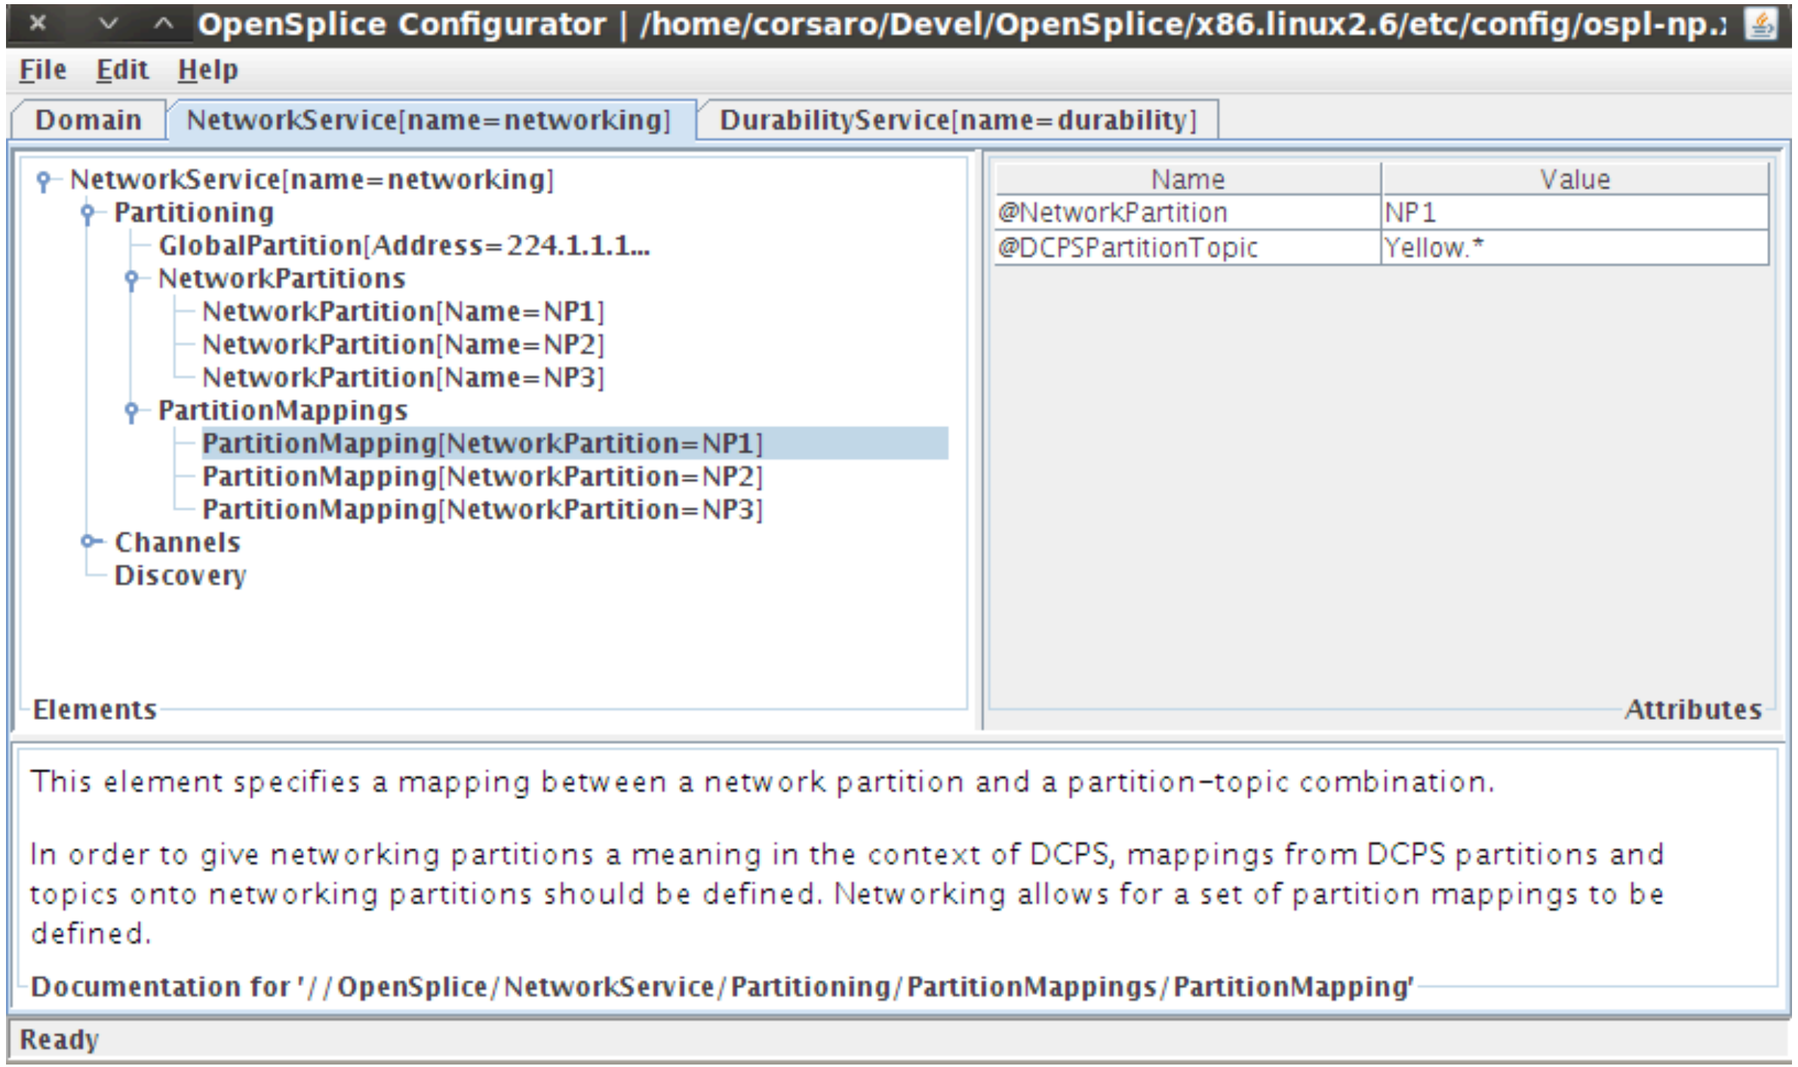

Partition Mapping (osplconf)

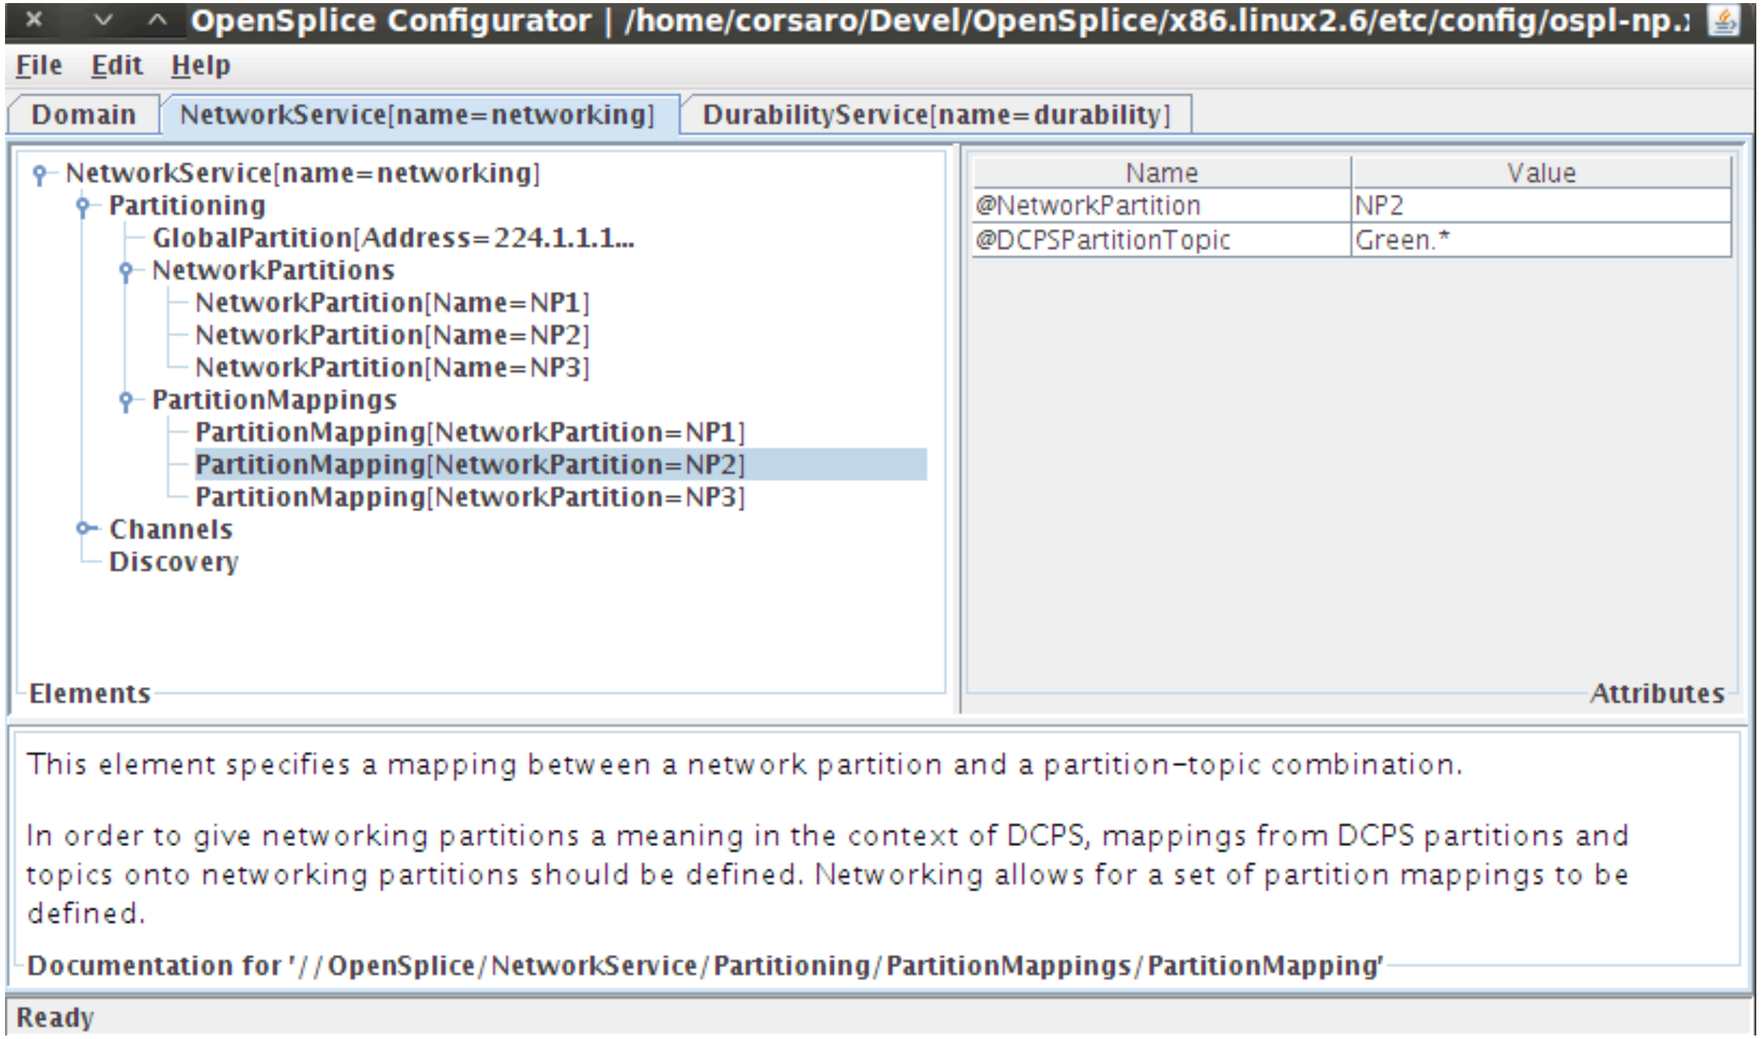

Partition Mapping (osplconf)

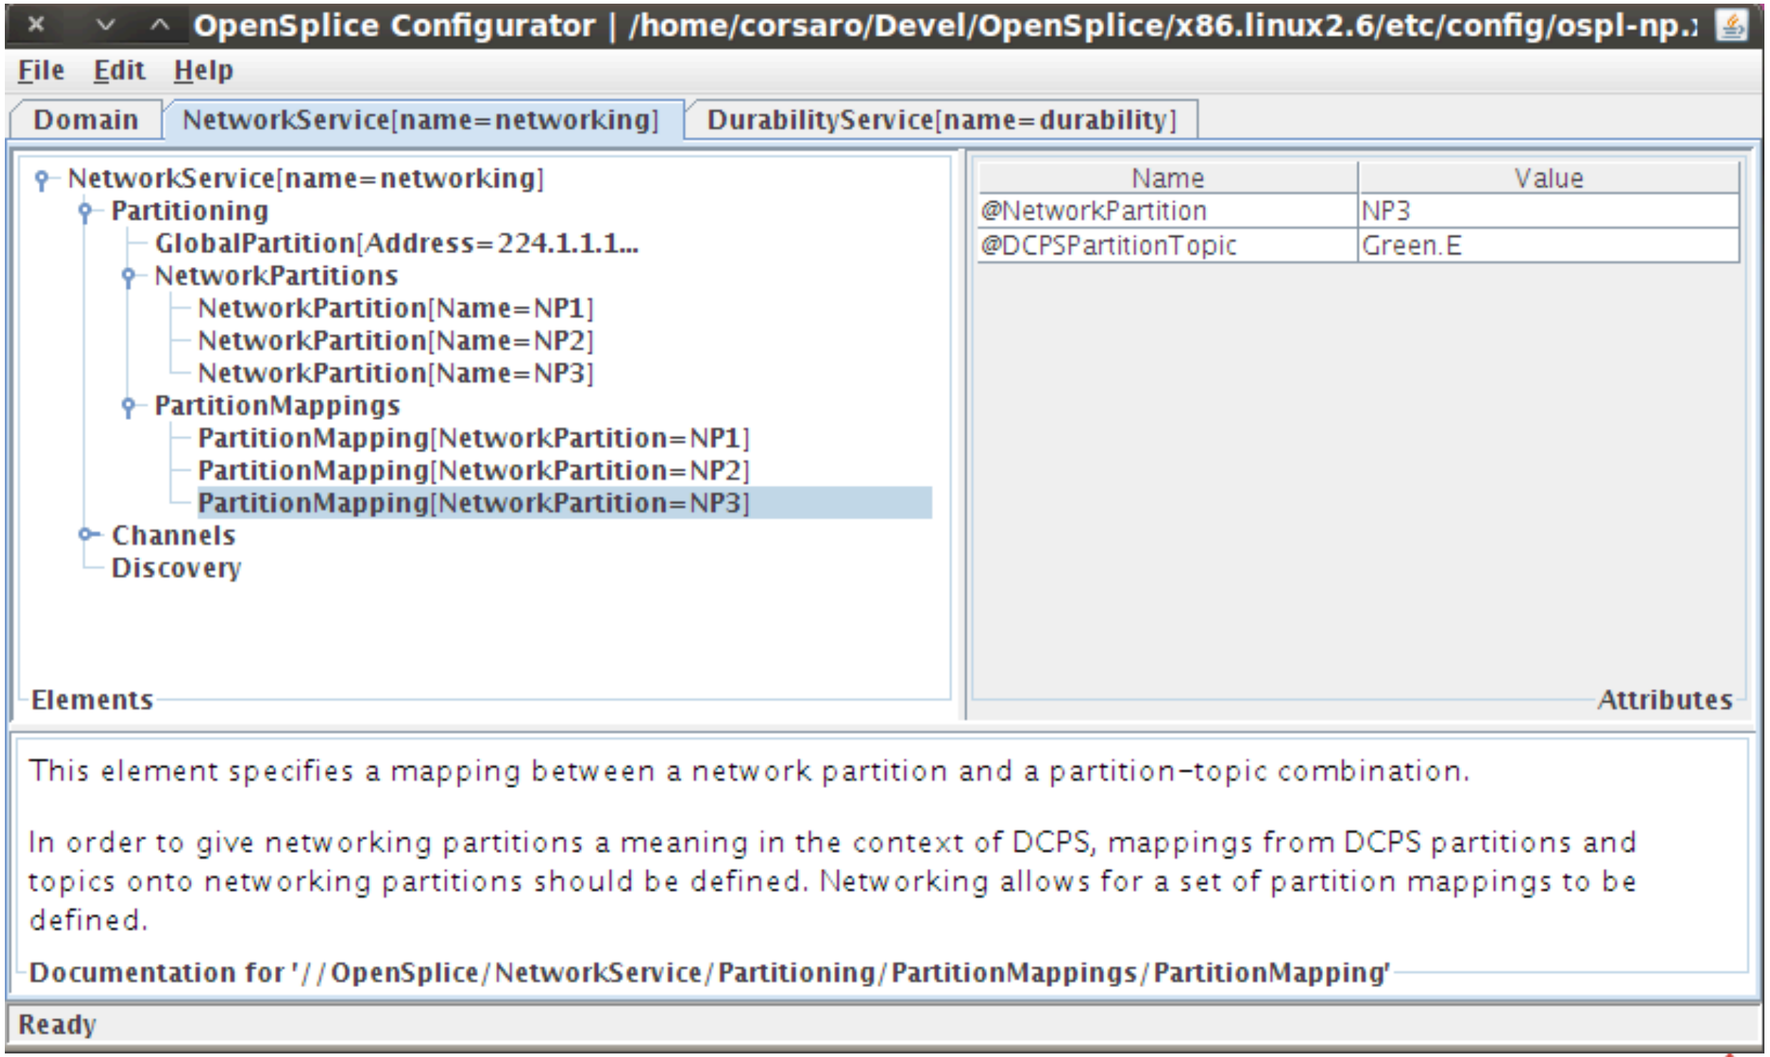

OpenSplice | DDS

Delivering Performance, Openness, and Freedom

Quality of Service

Anatomy of a DDS Application

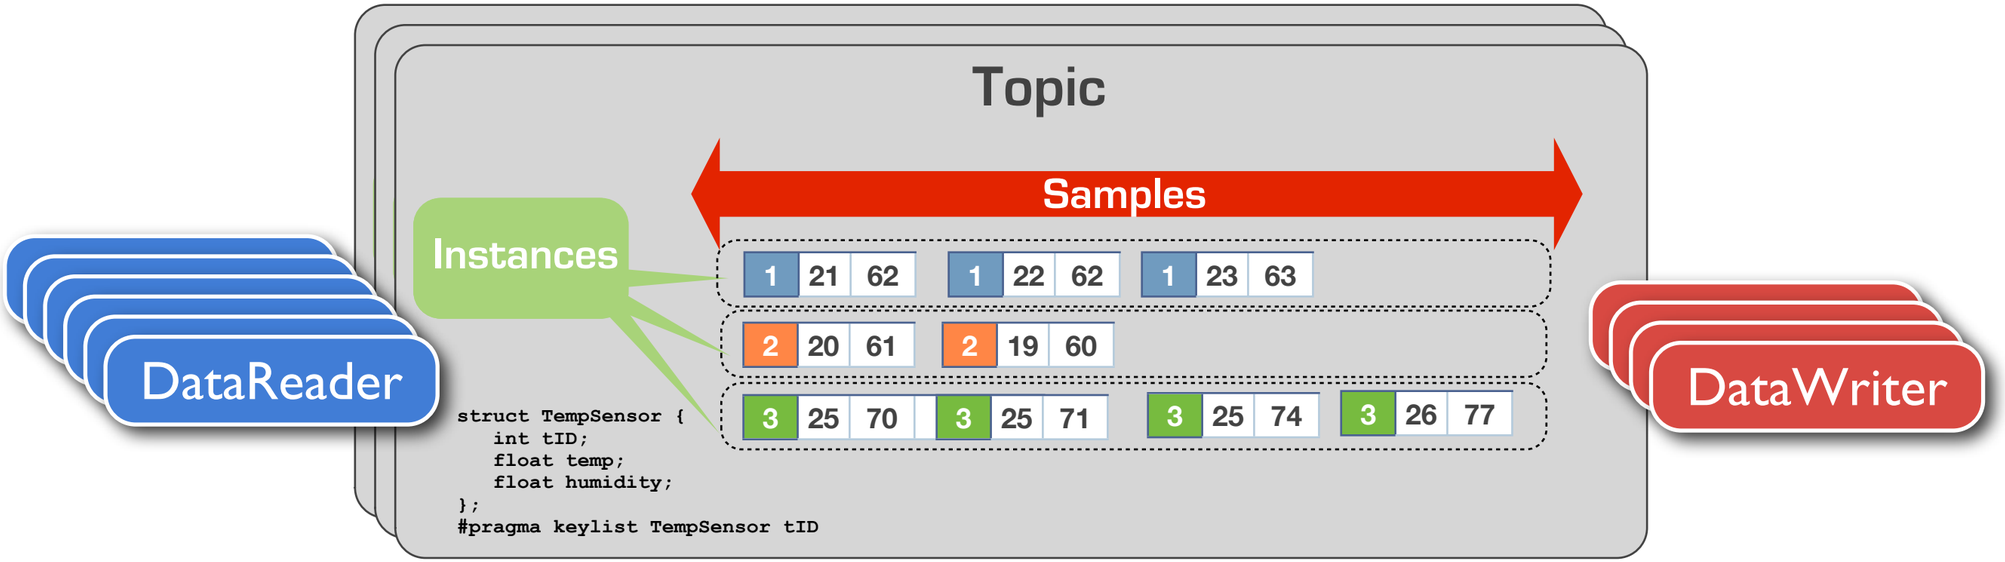

Anatomy of a DDS Application

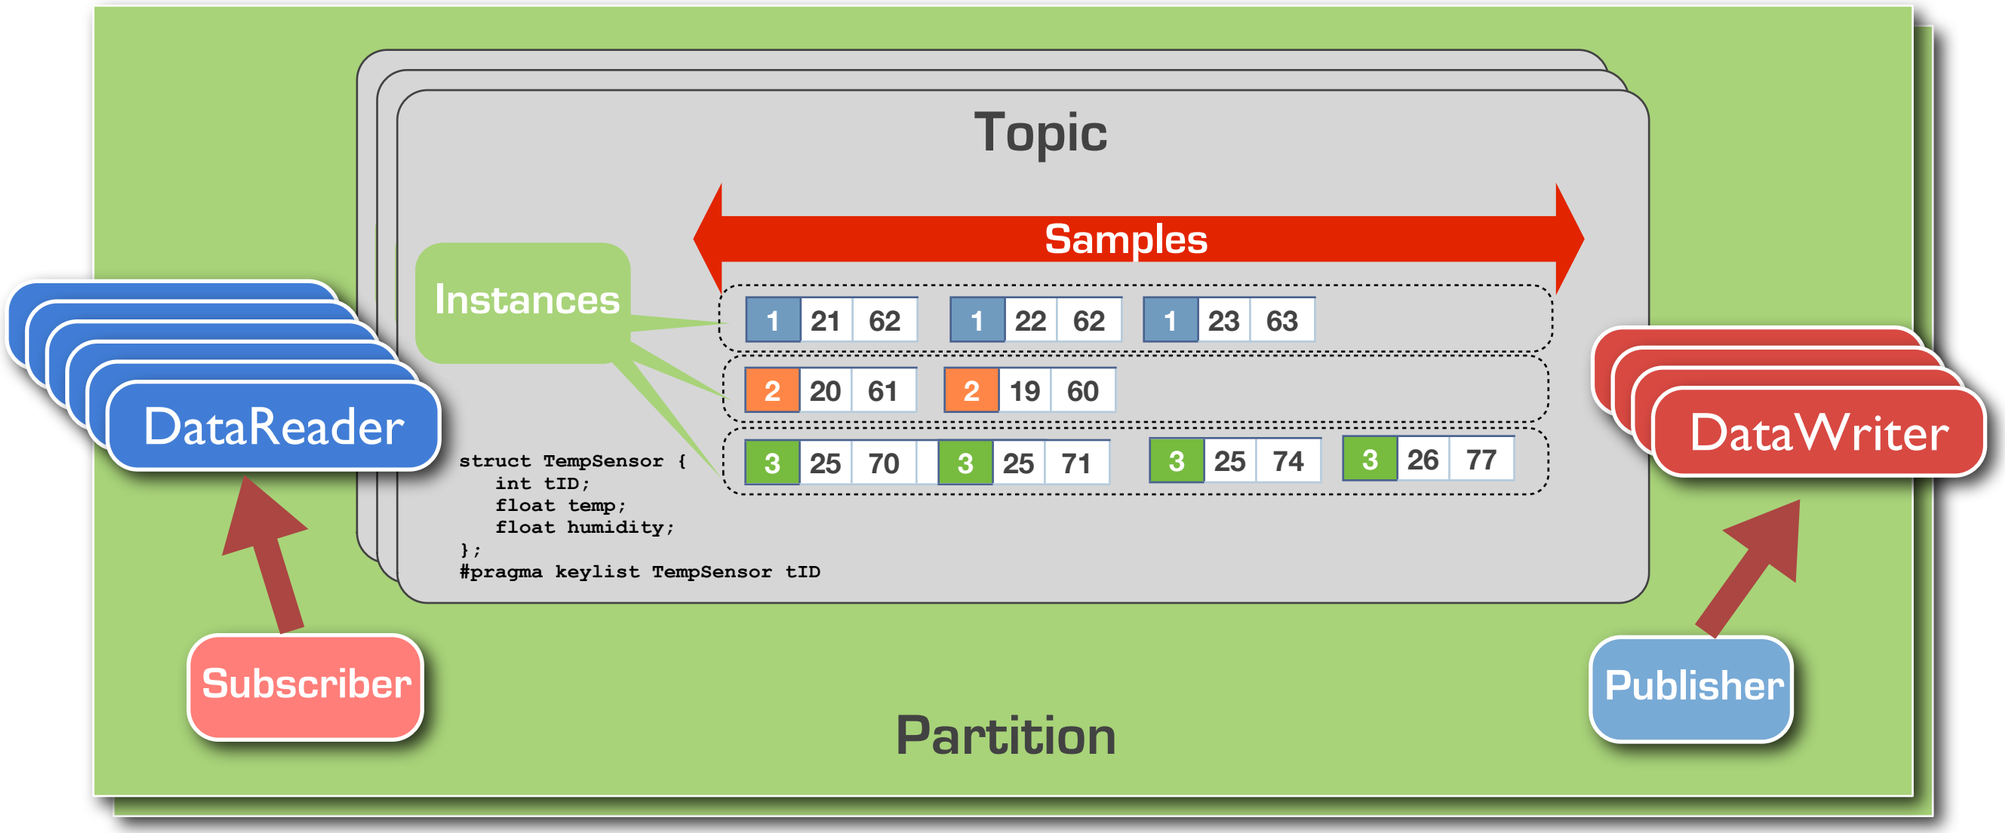

Arrows show structural relationships, not data-flows

Anatomy of a DDS Application

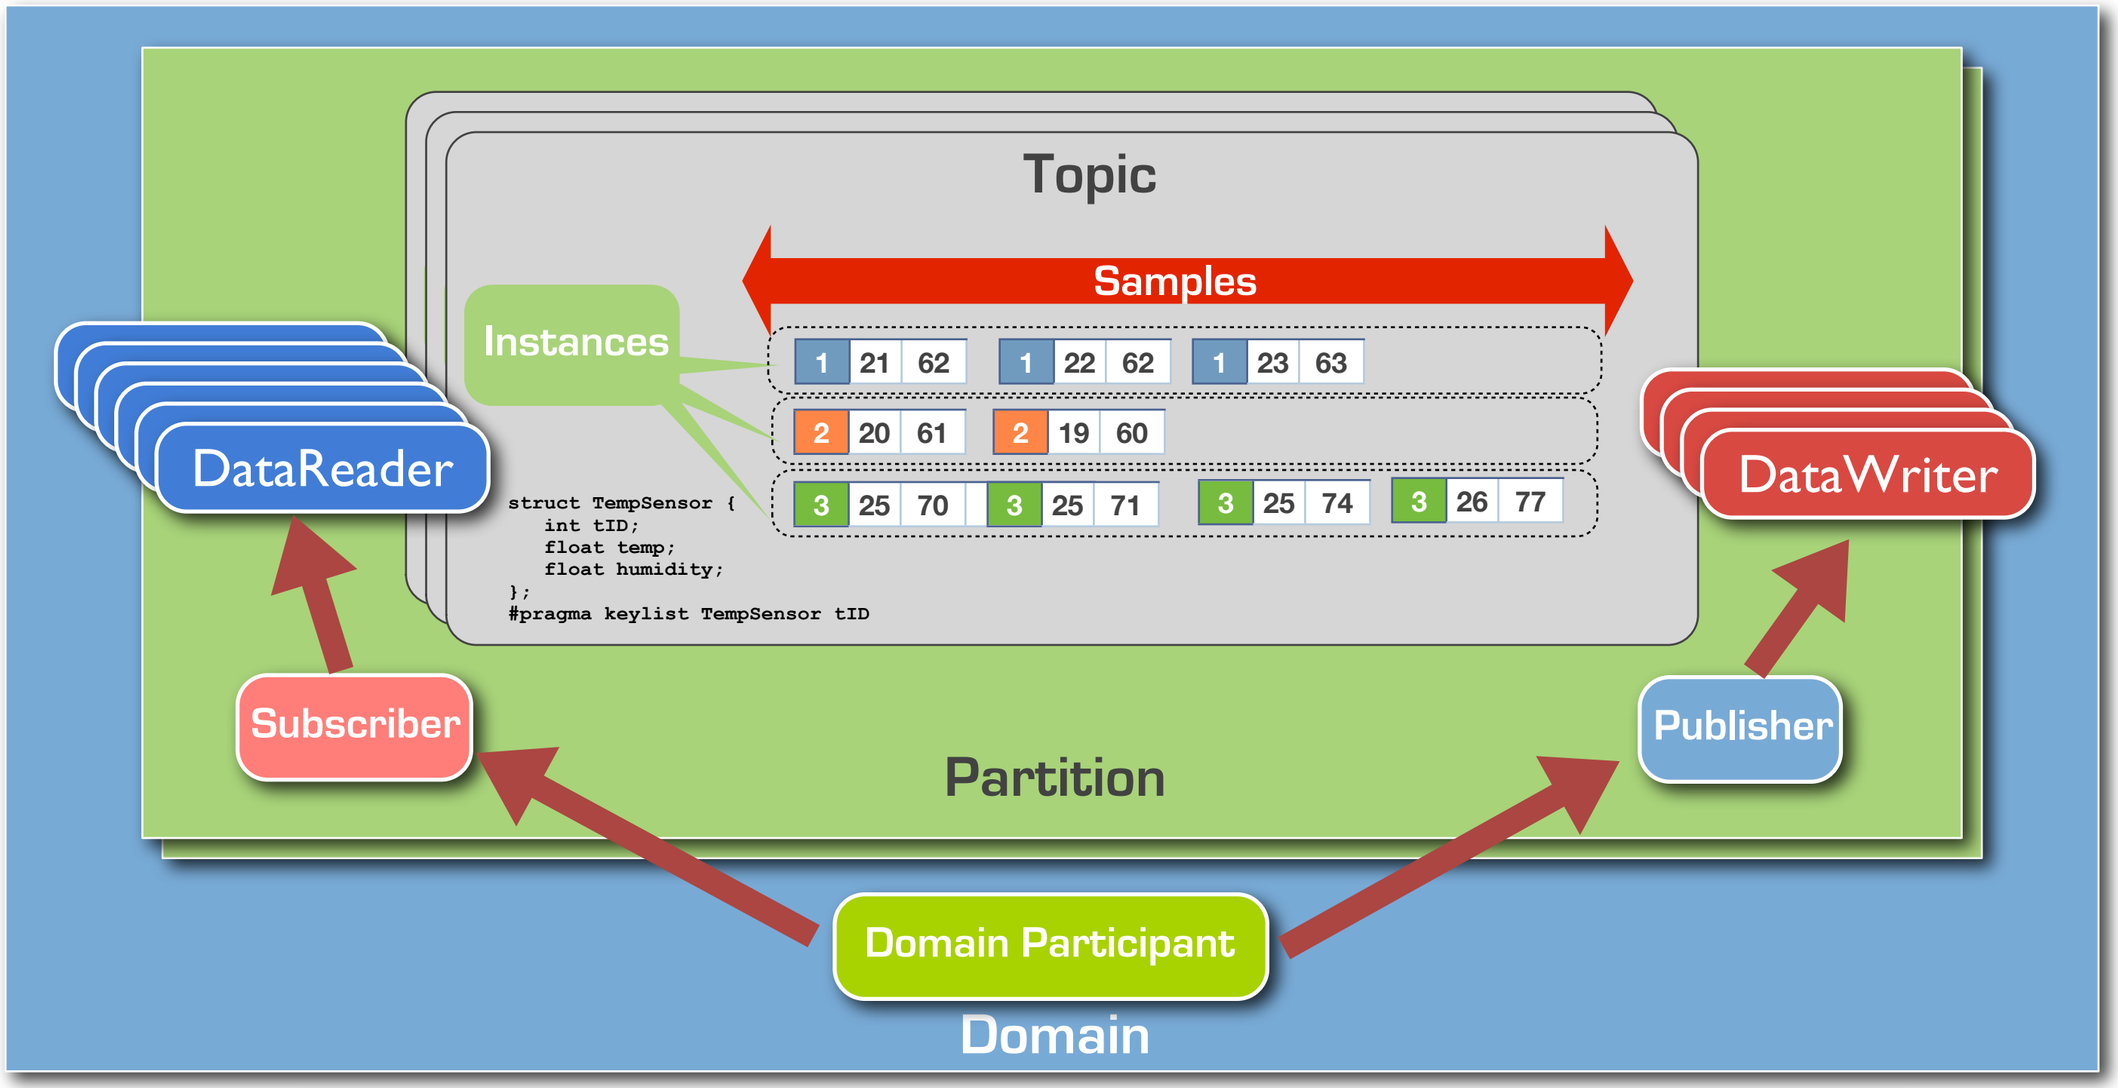

Arrows show structural relationships, not data-flows

QoS Model

- QoS-Policies are used to control relevant properties of OpenSplice DDS entities, such as:
  - Temporal Properties
  - Priority
  - Durability
  - Availability
  - ...
- Some QoS-Policies are matched based on a **Request vs. Offered Model** thus QoS-enforcement

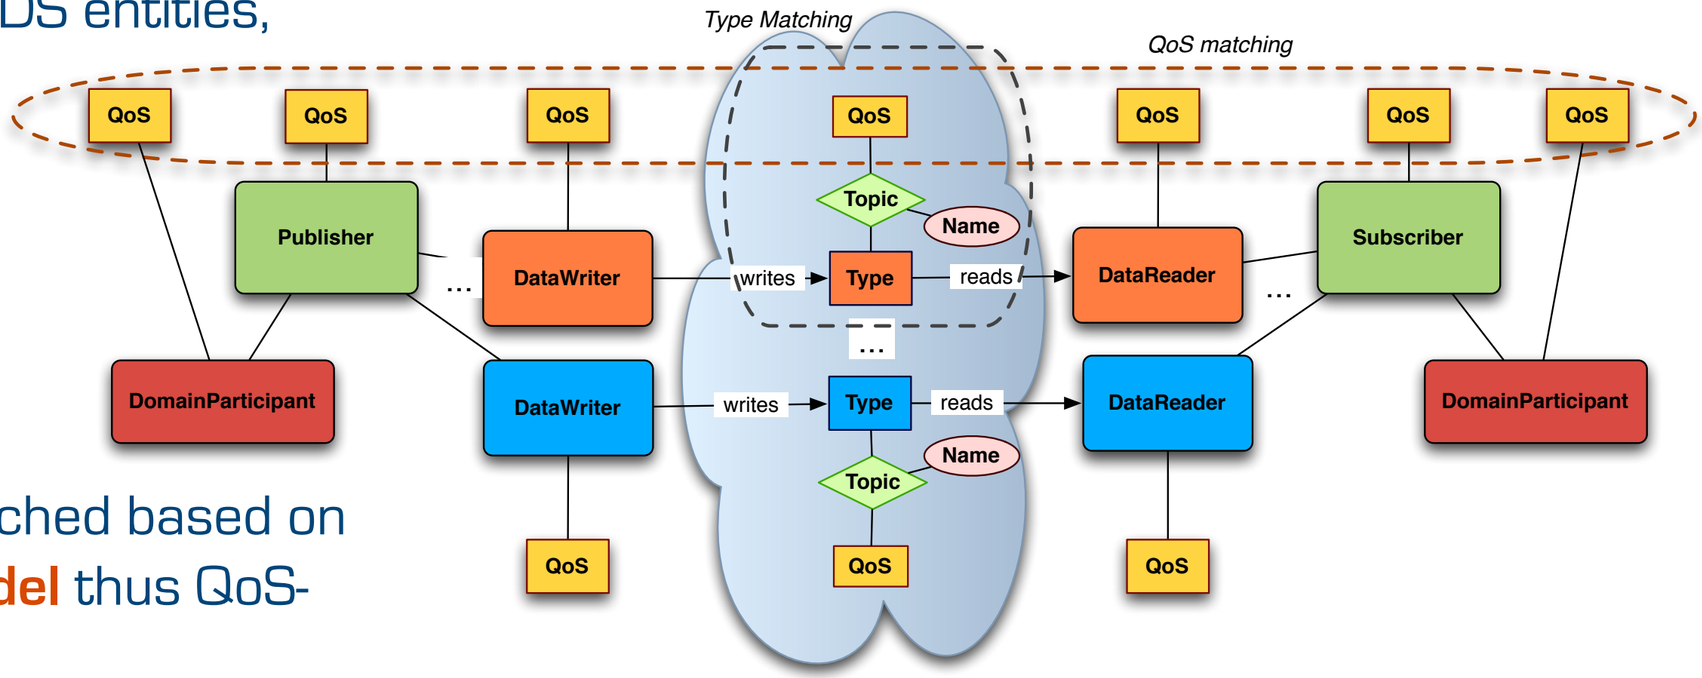

- Publications and Subscriptions match only if the declared vs. requested QoS are compatible
  - e.g., it is not possible to match a publisher which delivers data unreliably with a subscriber which requires reliability

QoS Policies

QoS Policy	Applicability	RxO	Modifiable	
DURABILITY	T, DR, DW	Y	N	Data Availability
DURABILITY SERVICE	T, DW	N	N	
LIFESPAN	T, DW	-	Y	
HISTORY	T, DR, DW	N	N	
PRESENTATION	P, S	Y	N	Data Delivery
RELIABILITY	T, DR, DW	Y	N	
PARTITION	P, S	N	Y	
DESTINATION ORDER	T, DR, DW	Y	N	
OWNERSHIP	T, DR, DW	Y	N	
OWNERSHIP STRENGTH	DW	-	Y	
DEADLINE	T, DR, DW	Y	Y	Data Timeliness
LATENCY BUDGET	T, DR, DW	Y	Y	
TRANSPORT PRIORITY	T, DW	-	Y	
TIME BASED FILTER	DR	-	Y	Resources
RESOURCE LIMITS	T, DR, DW	N	N	
USER_DATA	DP, DR, DW	N	Y	Configuration
TOPIC_DATA	T	N	Y	
GROUP_DATA	P, S	N	Y	

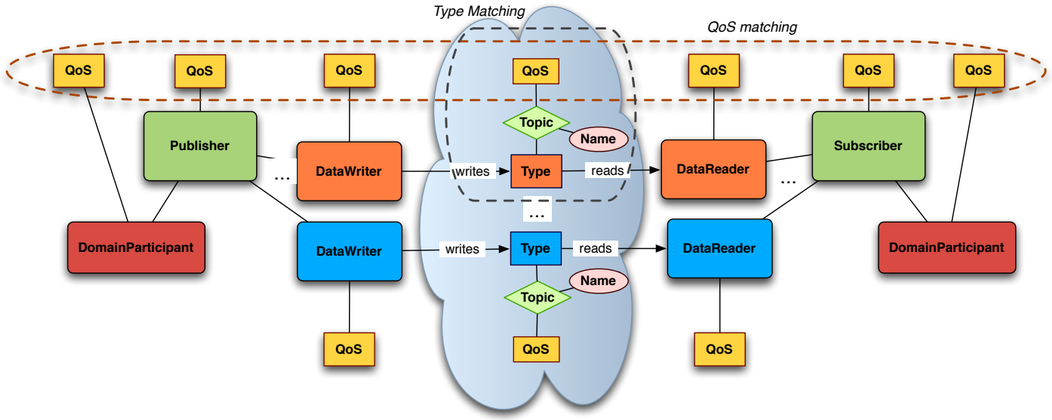

- Rich set of QoS allow to configure several different aspects of data availability, delivery and timeliness
- QoS can be used to control and optimize network as well as computing resource

OpenSplice DDS

PRISMTECH

Reliability

The reliability with which data is delivered to applications is impacted in DDS by the following qualities of service

- RELIABILITY
  - BEST_EFORT
  - RELIABLE
- HISTORY
  - KEEP_LAST (K)
  - KEEP_ALL

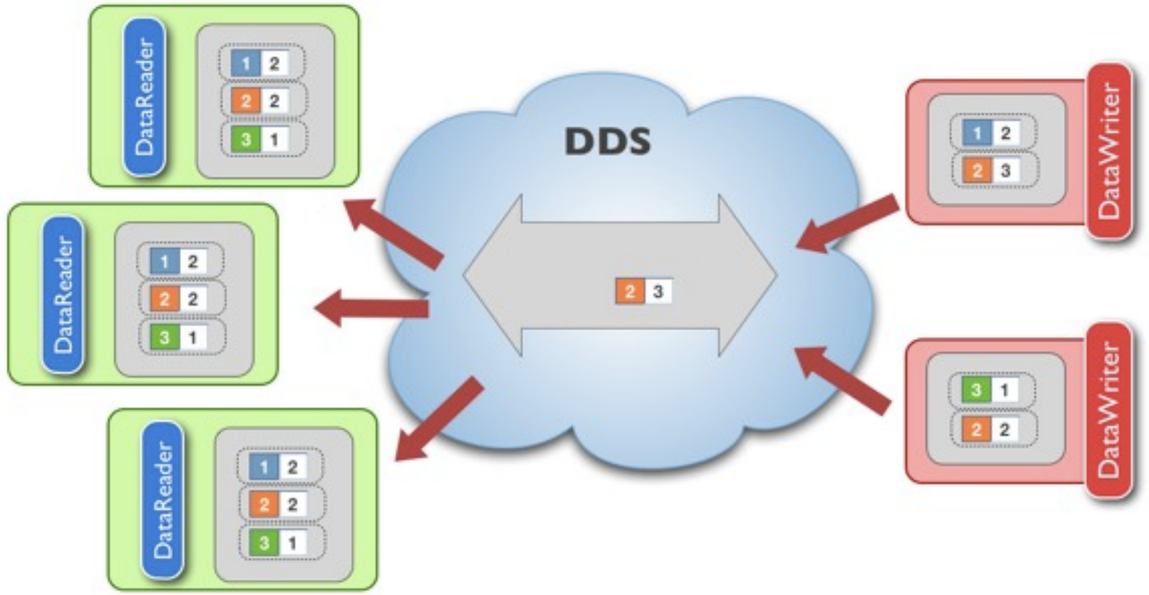

- **Theoretically**, the only way to assure that **an application will see all the samples** produced by a writer is to use **RELIABLE+KEEP_ALL**. Any other combination could induce to samples being discarded on the receiving side because of the HISTORY depth

Real-Time

The real-time properties with which data is delivered to applications is impacted in DDS by the following qualities of service:

- **TRANSPORT_PRIORITY**
- **LATENCY_BUDGET**
- In addition, DDS provides means for detecting performance failure, e.g., Deadline miss, by means of the **DEADLINE** QoS
- Given a periodic task-set {T} with periods Di (with $D_i < D_{i+1}$) and deadline equal to the period, than QoS should be set as follows:
  - Assign to each task T_i a TRANSPORT_PRIORITY Pi such that $P_i > P_{i+1}$
  - Set for each task T_i a DEADLINE QoS of D_i
  - For maximizing throughput and minimizing resource usage set for each Ti a LATENCY_BUDGET QoS between $D_i/2$ and $D_i/3$ (this is a rule of thumb, the upper bound is Di-(RTT/2))

Eventual Consistency & R/W Caches

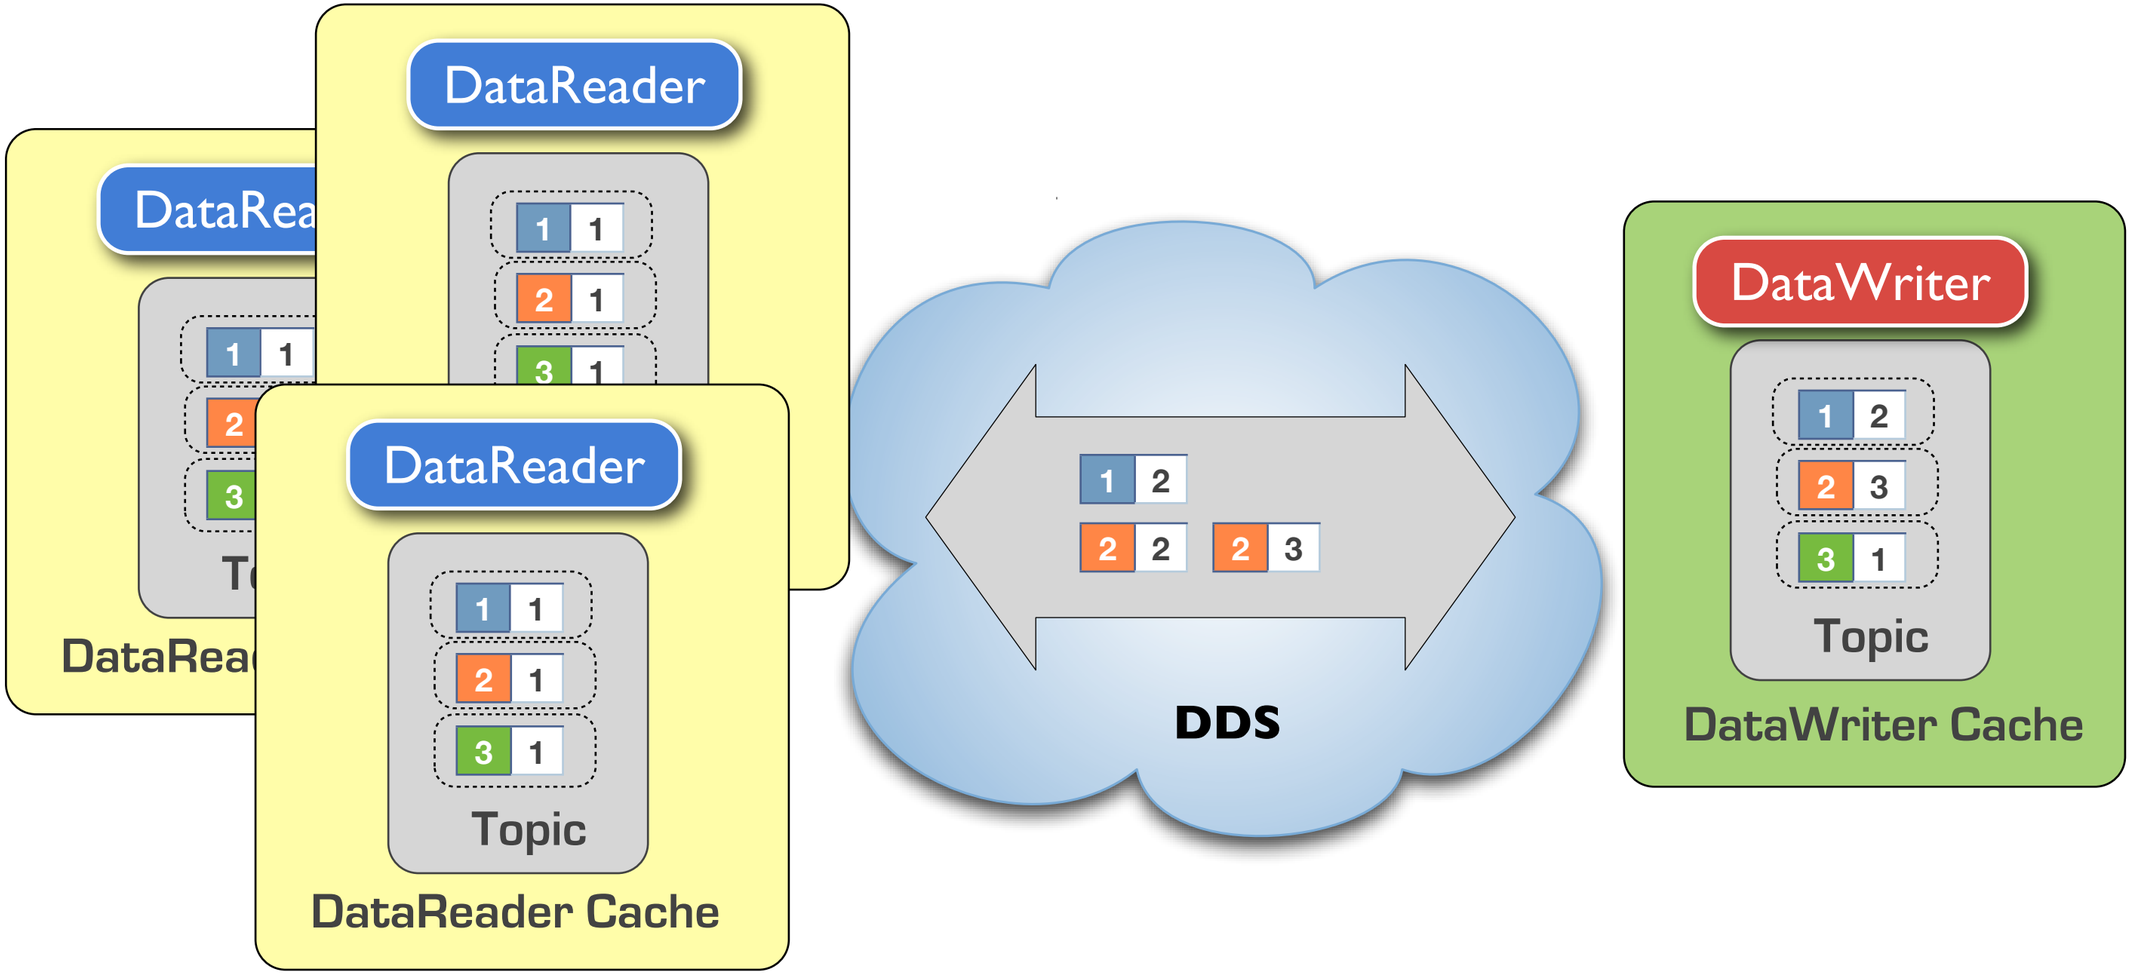

Under an Eventual Consistency Model, DDS guarantees that all matched Reader Caches will eventually be identical of the respective Writer Cache

QoS Impacting the Consistency Model

The DDS Consistency Model is a property that can be associated to Topics or further refined by Reader/Writers. The property is controlled by the following QoS Policies:

- **DURABILITY**
  - VOLATILE | TRANSIENT_LOCAL | TRANSIENT | PERSISTENT
- **LIFESPAN**
- **RELIABILITY**
  - RELIABLE | BEST_EFFORT
- **DESTINATION ORDER**
  - SOURCE_TIMESTAMP | DESTINATION_TIMESTAMP

QoS Policy	Applicability	RxO	Modifiable
DURABILITY	T, DR, DW	Y	N
LIFESPAN	T, DW	-	Y
RELIABILITY	T, DR, DW	Y	N
DESTINATION ORDER	T, DR, DW	Y	N

QoS Impacting the Consistency Model

	DURABILITY	RELIABILITY	DESTINATION_ORDER	LIFESPAN
Eventual Consistency (No Crash / Recovery)	VOLATILE	RELIABLE	SOURCE_TIMESTAMP	INF.
Eventual Consistency (Reader Crash / Recovery)	TRANSIENT_LOCAL	RELIABLE	SOURCE_TIMESTAMP	INF.
Eventual Consistency (Crash/Recovery)	TRANSIENT	RELIABLE	SOURCE_TIMESTAMP	INF.
Eventual Consistency (Crash/Recovery)	PERSISTENT	RELIABLE	SOURCE_TIMESTAMP	INF.
Weak Consistency	ANY	ANY	DESTINATION_TIMESTAMP	ANY
Weak Consistency	ANY	BEST_EFFORT	ANY	ANY
Weak Consistency	ANY	ANY	ANY	N

Eventual Consistency @ Work

	DURABILITY	RELIABILITY	DESTINATION_ORDER	LIFESPAN	
Eventual Consistency (Reader Crash / Recovery)	TRANSIENT_LOCAL	RELIABLE	SOURCE_TIMESTAMP	INF.	{A}
Eventual Consistency (Crash/Recovery)	TRANSIENT	RELIABLE	SOURCE_TIMESTAMP	INF.	{B}
Weak Consistency	ANY	ANY	ANY	N	{J}

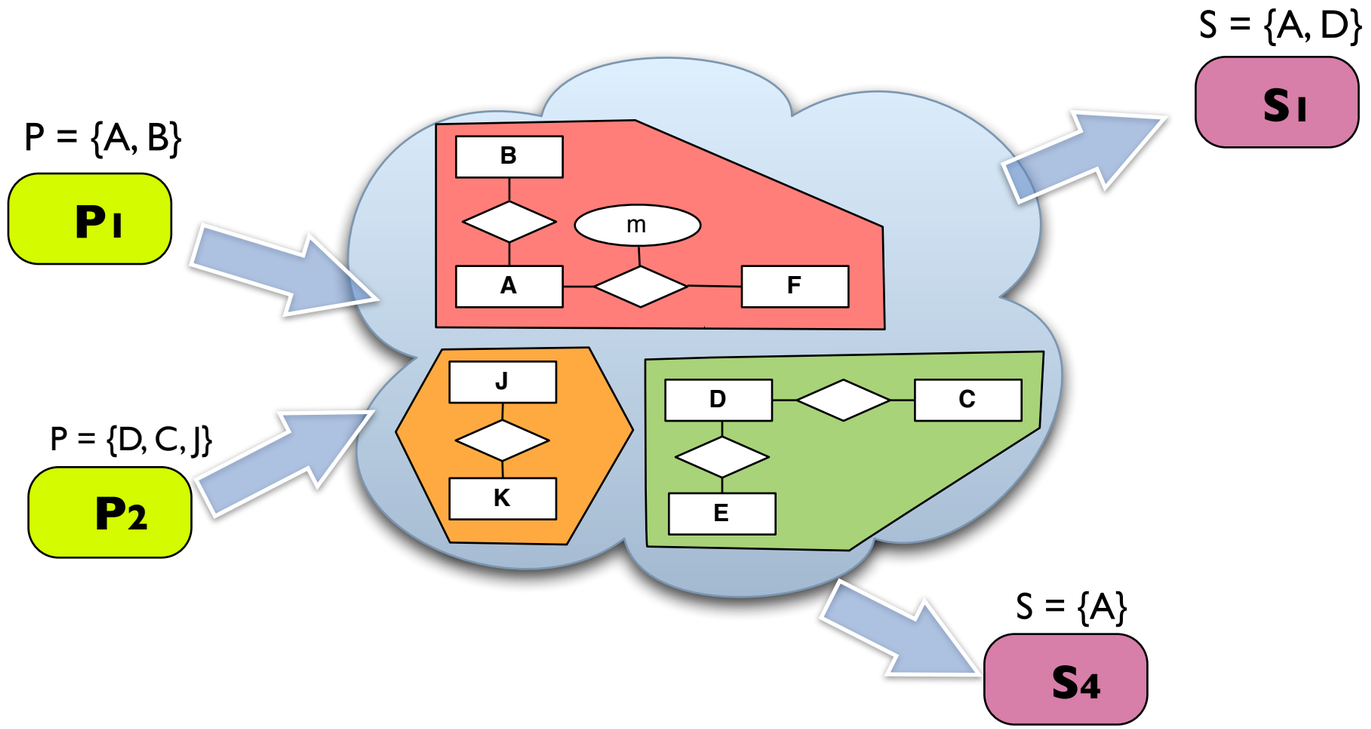

Eventual Consistency @ Work

	DURABILITY	RELIABILITY	DESTINATION_ORDER	LIFESPAN	
Eventual Consistency (Reader Crash / Recovery)	TRANSIENT_LOCAL	RELIABLE	SOURCE_TIMESTAMP	INF.	{A}
Eventual Consistency (Crash/Recovery)	TRANSIENT	RELIABLE	SOURCE_TIMESTAMP	INF.	{B}
Weak Consistency	ANY	ANY	ANY	N	{J}

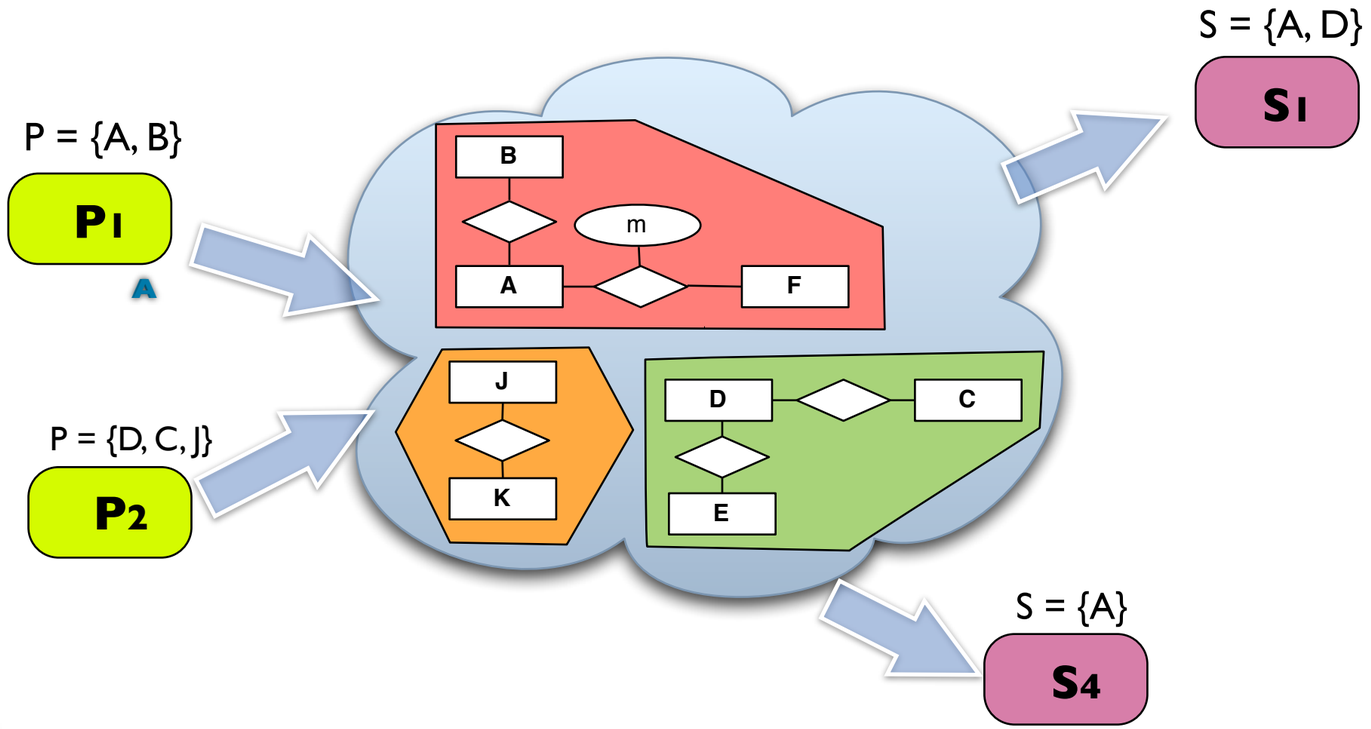

Eventual Consistency @ Work

	DURABILITY	RELIABILITY	DESTINATION_ORDER	LIFESPAN	
Eventual Consistency (Reader Crash / Recovery)	TRANSIENT_LOCAL	RELIABLE	SOURCE_TIMESTAMP	INF.	{A}
Eventual Consistency (Crash/Recovery)	TRANSIENT	RELIABLE	SOURCE_TIMESTAMP	INF.	{B}
Weak Consistency	ANY	ANY	ANY	N	{J}

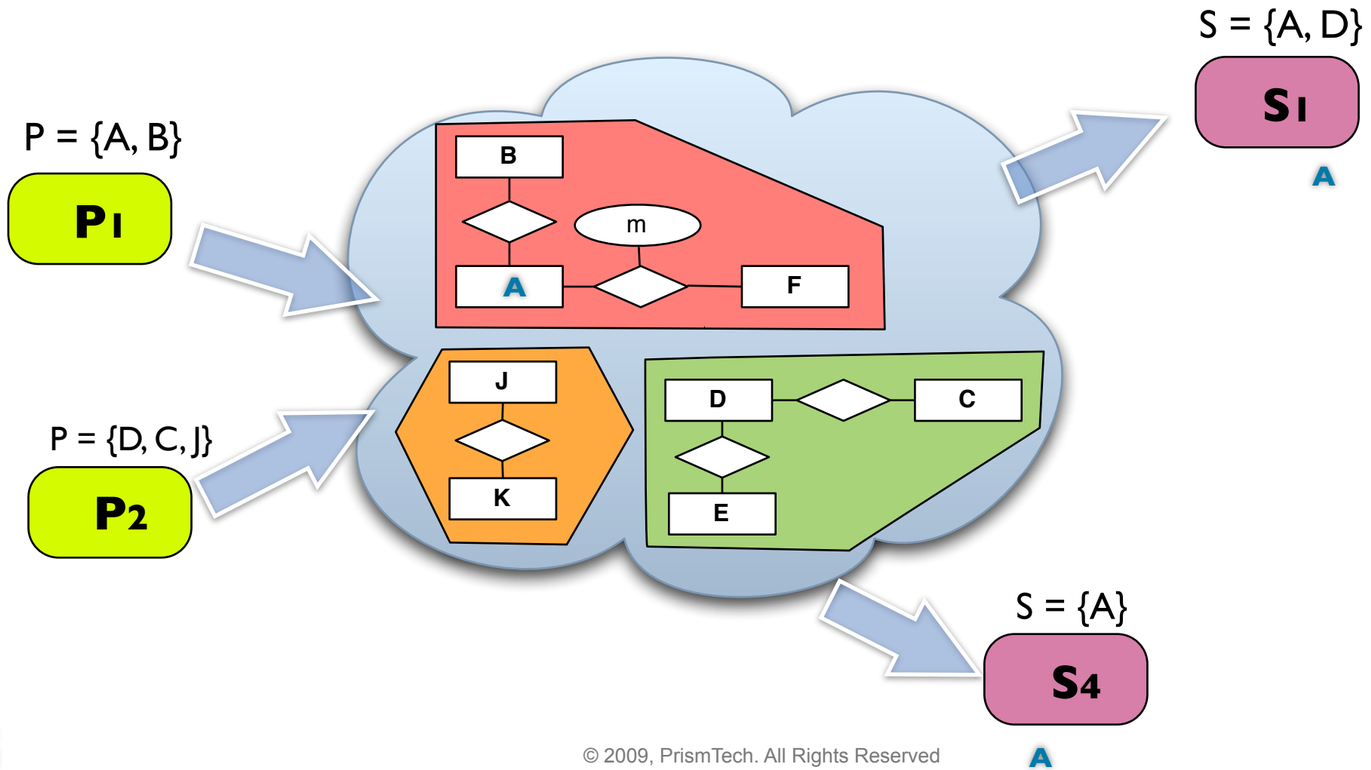

Eventual Consistency @ Work

	DURABILITY	RELIABILITY	DESTINATION_ORDER	LIFESPAN	
Eventual Consistency (Reader Crash / Recovery)	TRANSIENT_LOCAL	RELIABLE	SOURCE_TIMESTAMP	INF.	{A}
Eventual Consistency (Crash/Recovery)	TRANSIENT	RELIABLE	SOURCE_TIMESTAMP	INF.	{B}
Weak Consistency	ANY	ANY	ANY	N	{J}

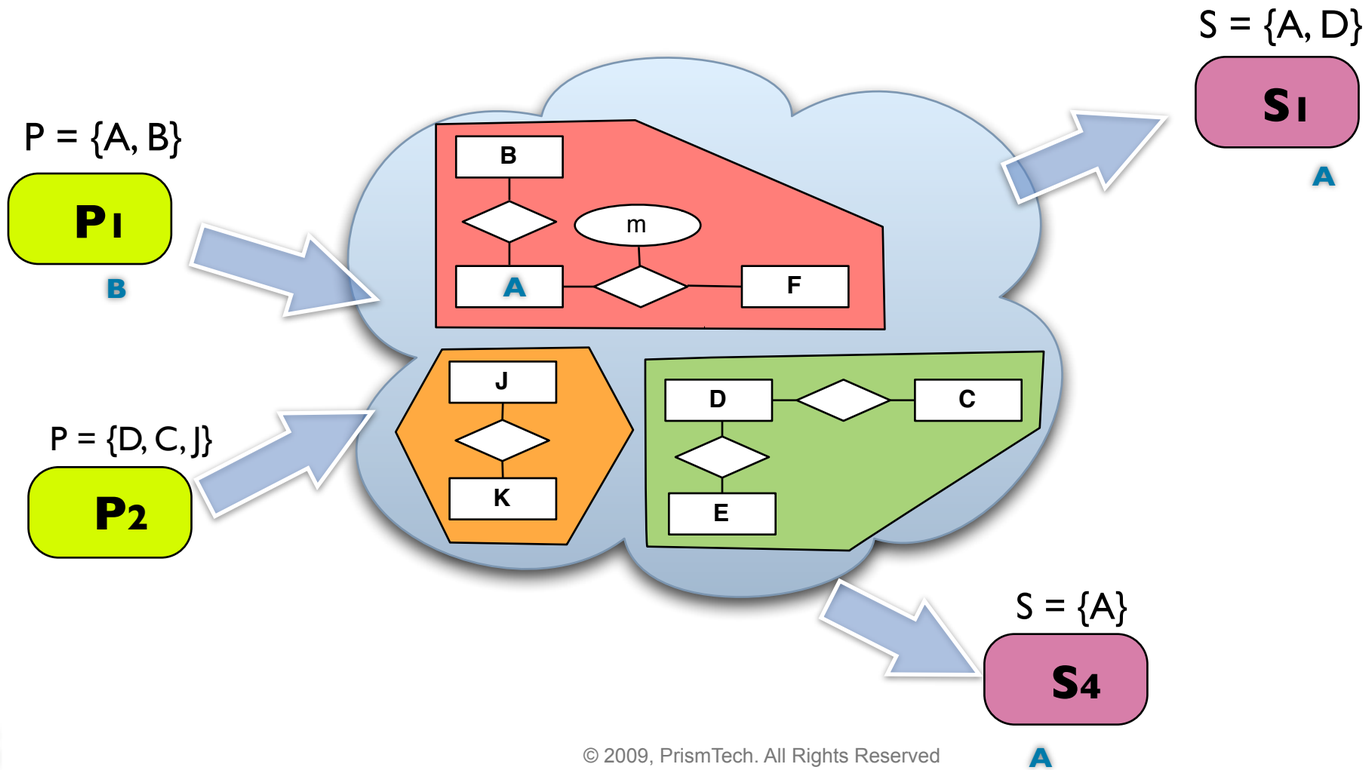

Eventual Consistency @ Work

	DURABILITY	RELIABILITY	DESTINATION_ORDER	LIFESPAN	
Eventual Consistency (Reader Crash / Recovery)	TRANSIENT_LOCAL	RELIABLE	SOURCE_TIMESTAMP	INF.	{A}
Eventual Consistency (Crash/Recovery)	TRANSIENT	RELIABLE	SOURCE_TIMESTAMP	INF.	{B}
Weak Consistency	ANY	ANY	ANY	N	{J}

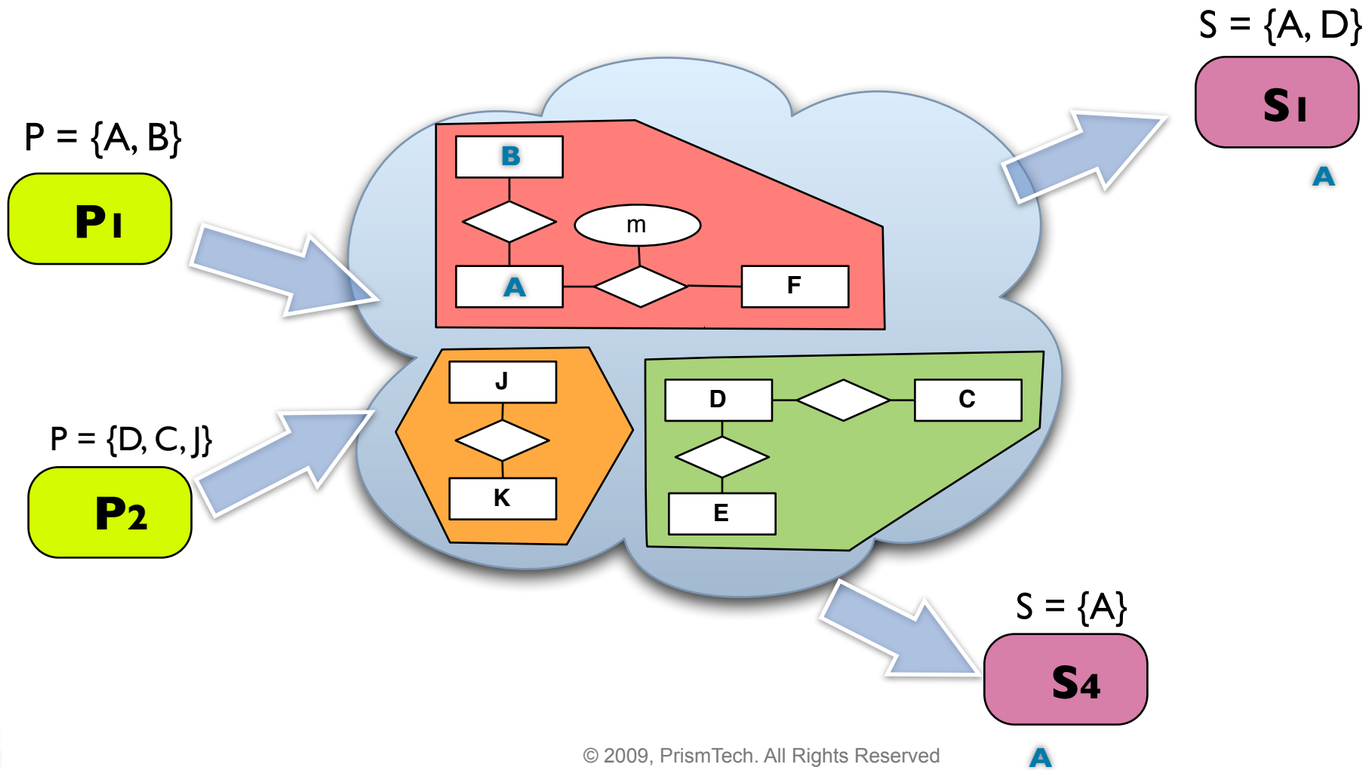

Eventual Consistency @ Work

	DURABILITY	RELIABILITY	DESTINATION_ORDER	LIFESPAN	
Eventual Consistency (Reader Crash / Recovery)	TRANSIENT_LOCAL	RELIABLE	SOURCE_TIMESTAMP	INF.	{A}
Eventual Consistency (Crash/Recovery)	TRANSIENT	RELIABLE	SOURCE_TIMESTAMP	INF.	{B}
Weak Consistency	ANY	ANY	ANY	N	{J}

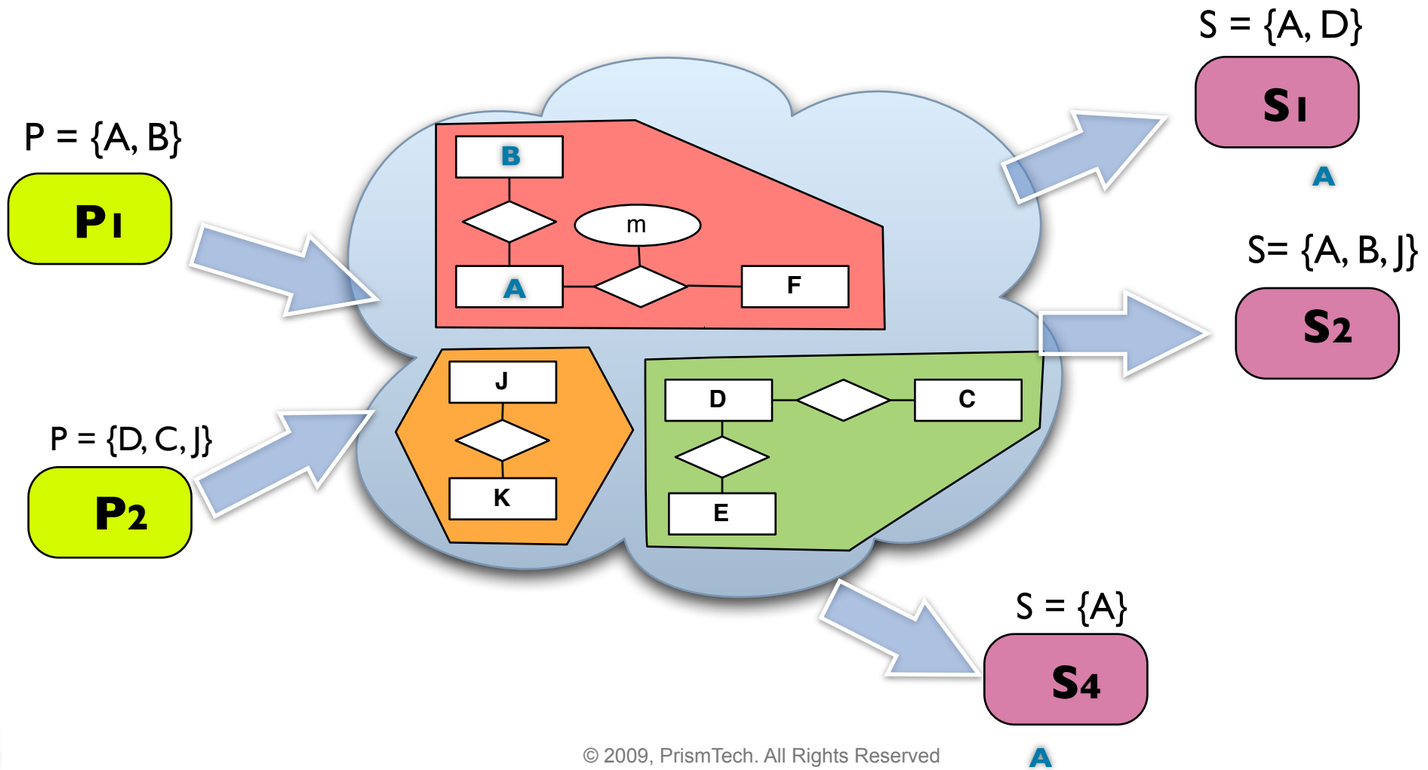

Eventual Consistency @ Work

	DURABILITY	RELIABILITY	DESTINATION_ORDER	LIFESPAN	
Eventual Consistency (Reader Crash / Recovery)	TRANSIENT_LOCAL	RELIABLE	SOURCE_TIMESTAMP	INF.	{A}
Eventual Consistency (Crash/Recovery)	TRANSIENT	RELIABLE	SOURCE_TIMESTAMP	INF.	{B}
Weak Consistency	ANY	ANY	ANY	N	{J}

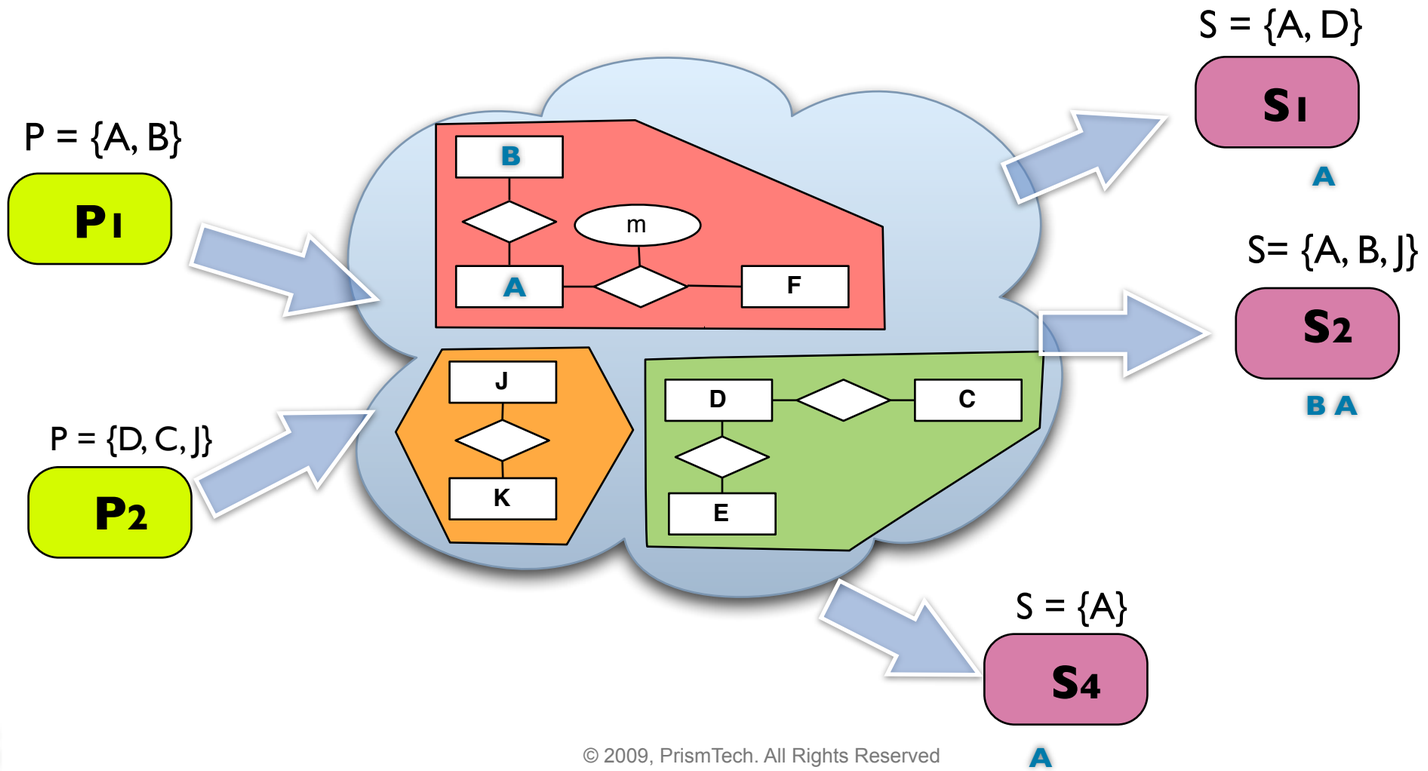

Eventual Consistency @ Work

	DURABILITY	RELIABILITY	DESTINATION_ORDER	LIFESPAN	
Eventual Consistency (Reader Crash / Recovery)	TRANSIENT_LOCAL	RELIABLE	SOURCE_TIMESTAMP	INF.	{A}
Eventual Consistency (Crash/Recovery)	TRANSIENT	RELIABLE	SOURCE_TIMESTAMP	INF.	{B}
Weak Consistency	ANY	ANY	ANY	N	{J}

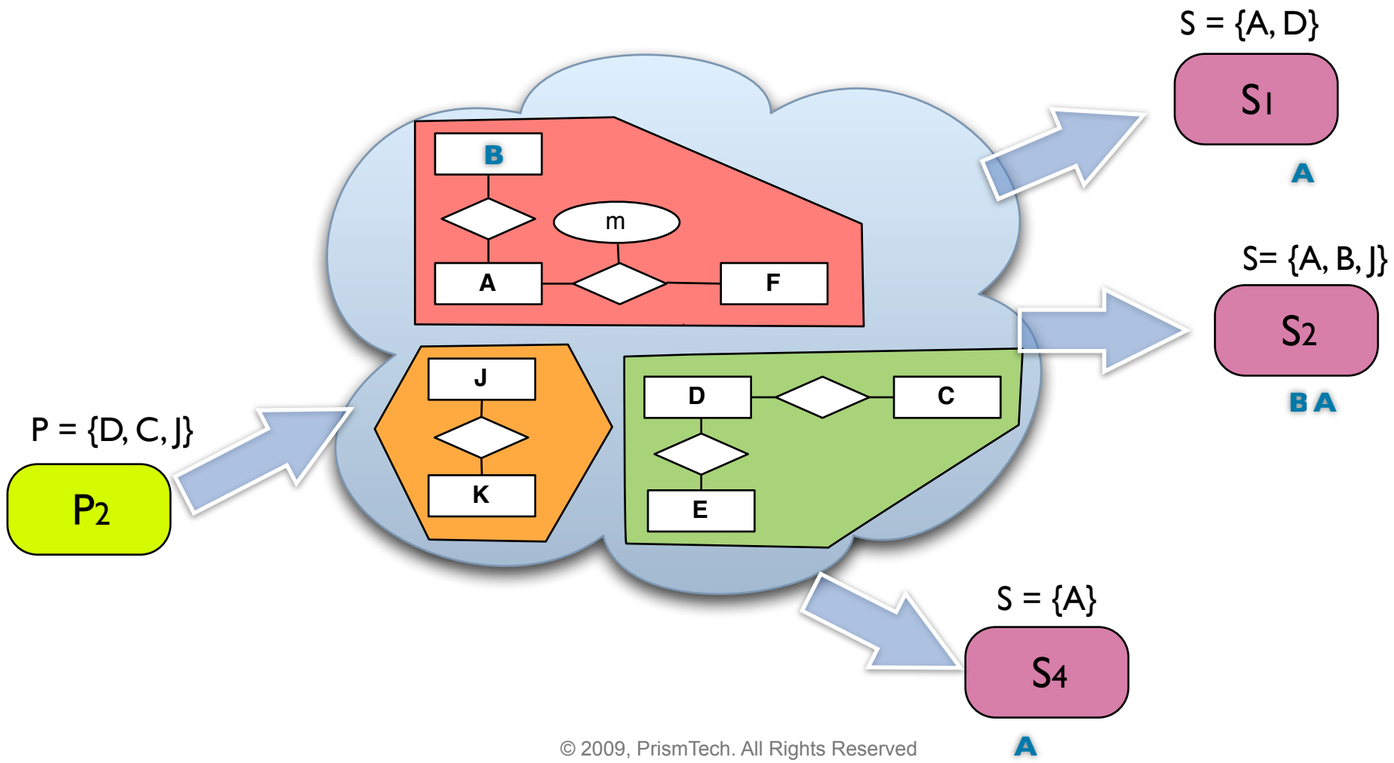

Eventual Consistency @ Work

	DURABILITY	RELIABILITY	DESTINATION_ORDER	LIFESPAN	
Eventual Consistency (Reader Crash / Recovery)	TRANSIENT_LOCAL	RELIABLE	SOURCE_TIMESTAMP	INF.	{A}
Eventual Consistency (Crash/Recovery)	TRANSIENT	RELIABLE	SOURCE_TIMESTAMP	INF.	{B}
Weak Consistency	ANY	ANY	ANY	N	{J}

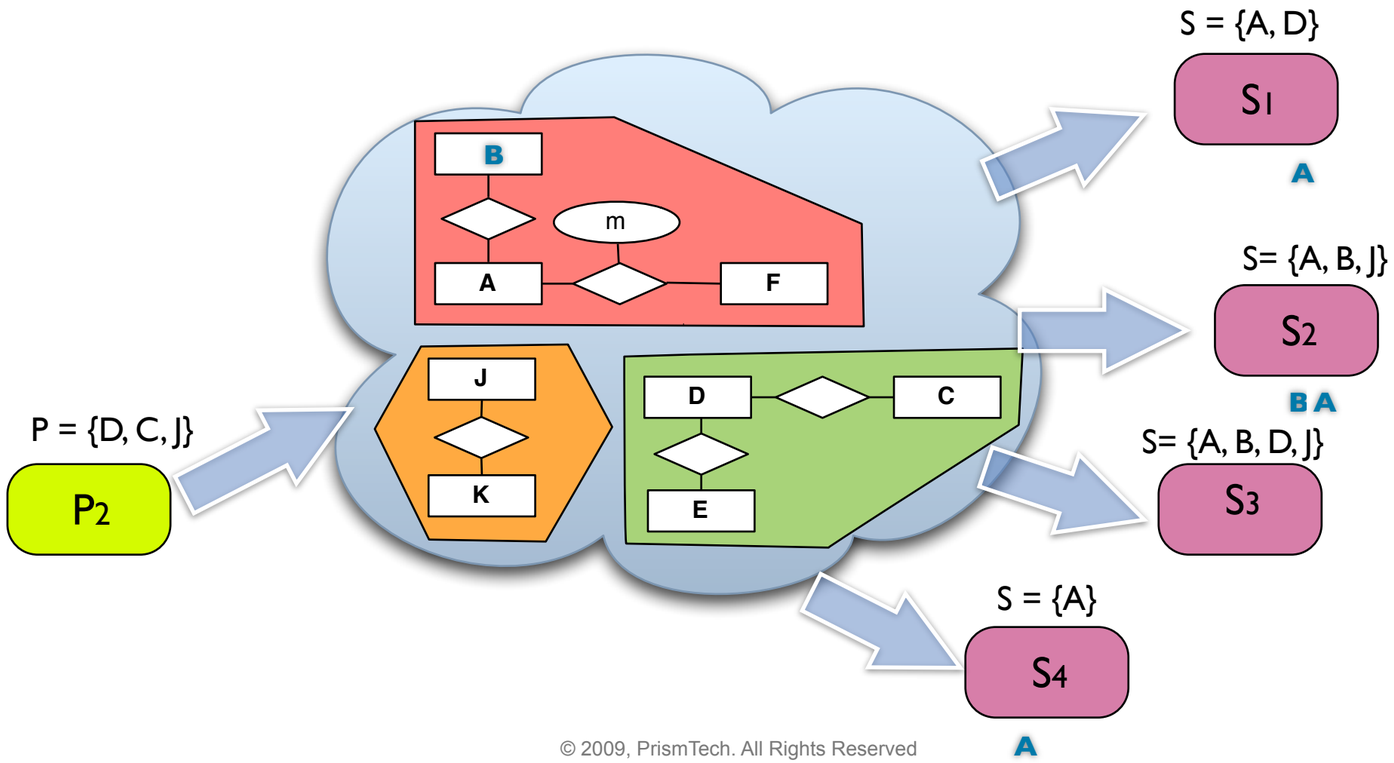

Eventual Consistency @ Work

	DURABILITY	RELIABILITY	DESTINATION_ORDER	LIFESPAN	
Eventual Consistency (Reader Crash / Recovery)	TRANSIENT_LOCAL	RELIABLE	SOURCE_TIMESTAMP	INF.	{A}
Eventual Consistency (Crash/Recovery)	TRANSIENT	RELIABLE	SOURCE_TIMESTAMP	INF.	{B}
Weak Consistency	ANY	ANY	ANY	N	{J}

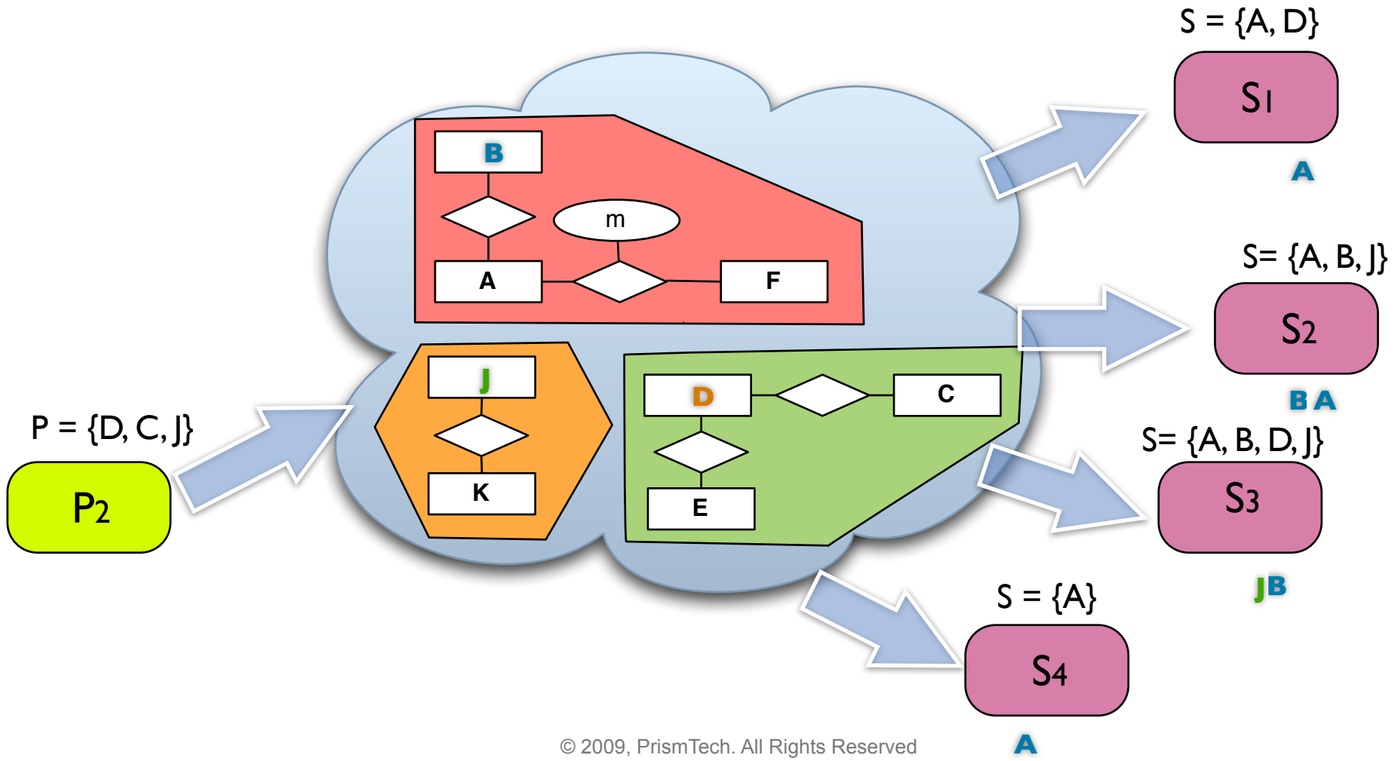

Eventual Consistency @ Work

	DURABILITY	RELIABILITY	DESTINATION_ORDER	LIFESPAN	
Eventual Consistency (Reader Crash / Recovery)	TRANSIENT_LOCAL	RELIABLE	SOURCE_TIMESTAMP	INF.	{A}
Eventual Consistency (Crash/Recovery)	TRANSIENT	RELIABLE	SOURCE_TIMESTAMP	INF.	{B}
Weak Consistency	ANY	ANY	ANY	N	{J}

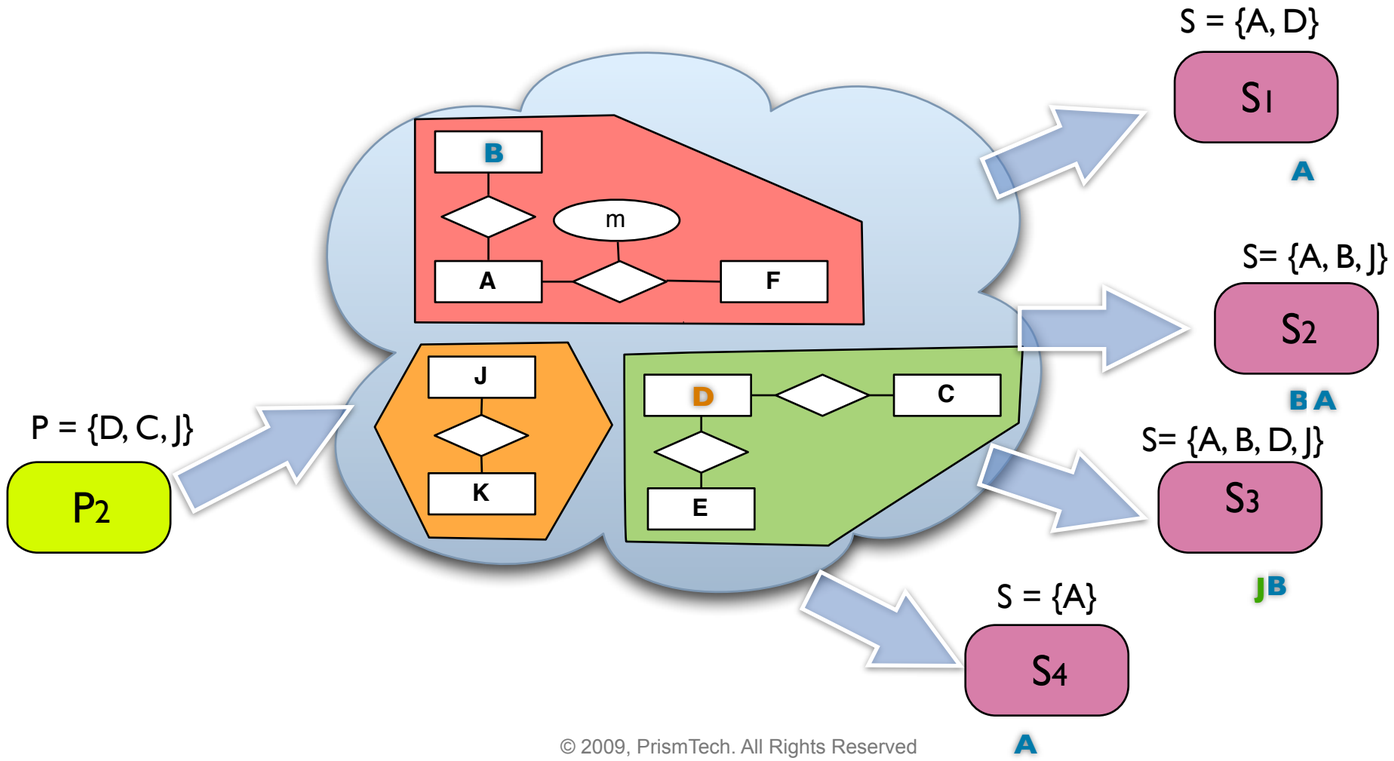

Design Guidelines

- For all (non-periodic) Topics for which an eventually consistent model is required use the following QoS settings:

	DURABILITY	RELIABILITY	DESTINATION_ORDER	LIFESPAN
Eventual Consistency (Crash / Recovery)	TRANSIENT	RELIABLE	SOURCE_TIMESTAMP	INF.

- For information produced periodically, with a period P, where P is small enough to be acceptable as a consistency convergence delay, the following QoS settings will provide an approximation of the eventual consistency:

	DURABILITY	RELIABILITY	DESTINATION_ORDER	LIFESPAN
Eventual Consistency (Crash / Recovery)	VOLATILE	BEST_EFFORT	SOURCE_TIMESTAMP	INF.

OpenSplice | DDS

Delivering Performance, Openness, and Freedom

OpenSplice DDS Architecture

OpenSplice DDS Architectural Outlook

Architectural Highlights

- Shared-Memory based architecture for minimizing intra-nodal latency, as well as maximizing nodal scalability
- Plugglable Service Architecture
- Full control over network scheduling

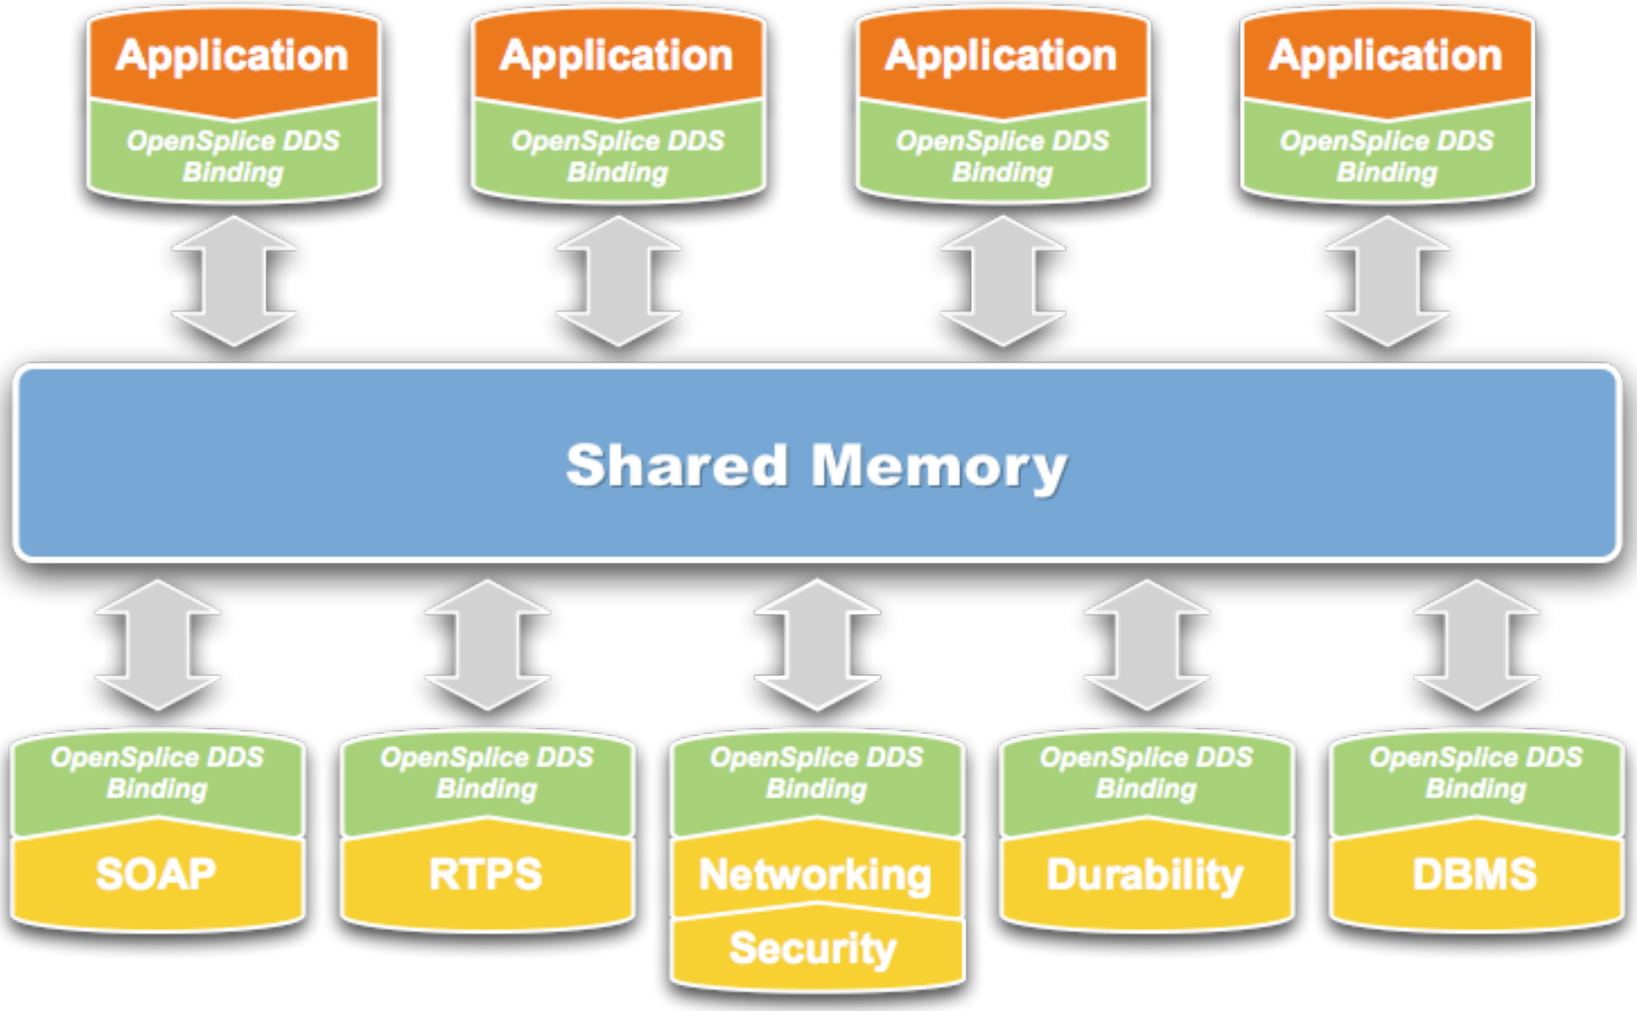

OpenSplice DDS

PRISMTECH

Real-Time Networking Technology

Architecture

- **Network-channels**
  - Priority bands
- **Network-partitions**
  - Multicast Groups
- **Traffic-shaping**
  - Burst/Throughput

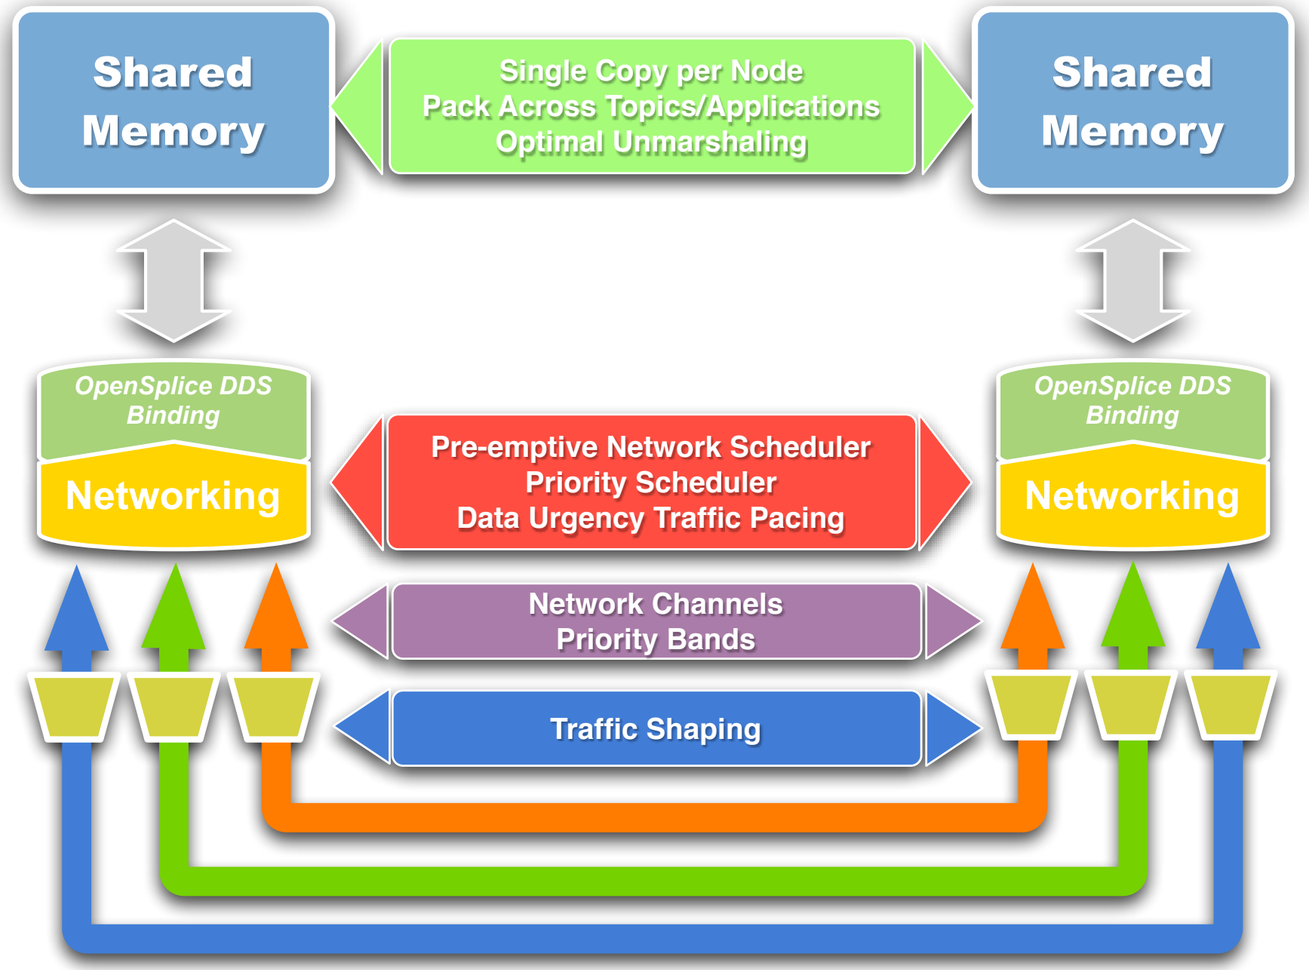

Scalability and Efficiency

- Single shared library for applications & services
- Ring-fenced shared memory segment
- Data urgency driven network-packing

Fault-Tolerance

- Active Channels
- Fall back on next highest priority active channel

Determinism & Safety

- Preemptive network-scheduler
- Data importance based network-channel selection
- Partition based multicast-group selection
- Managed critical network-resource

OpenSplice DDS

PRISMTECH

Durable Data Technology

Architecture

- **Fault-Tolerant Data Availability**
  - Transient – on memory
  - Persistent – on disk
- **Partitioning**
  - DDS Partitions
- **Alignment**
  - Dedicated Channels

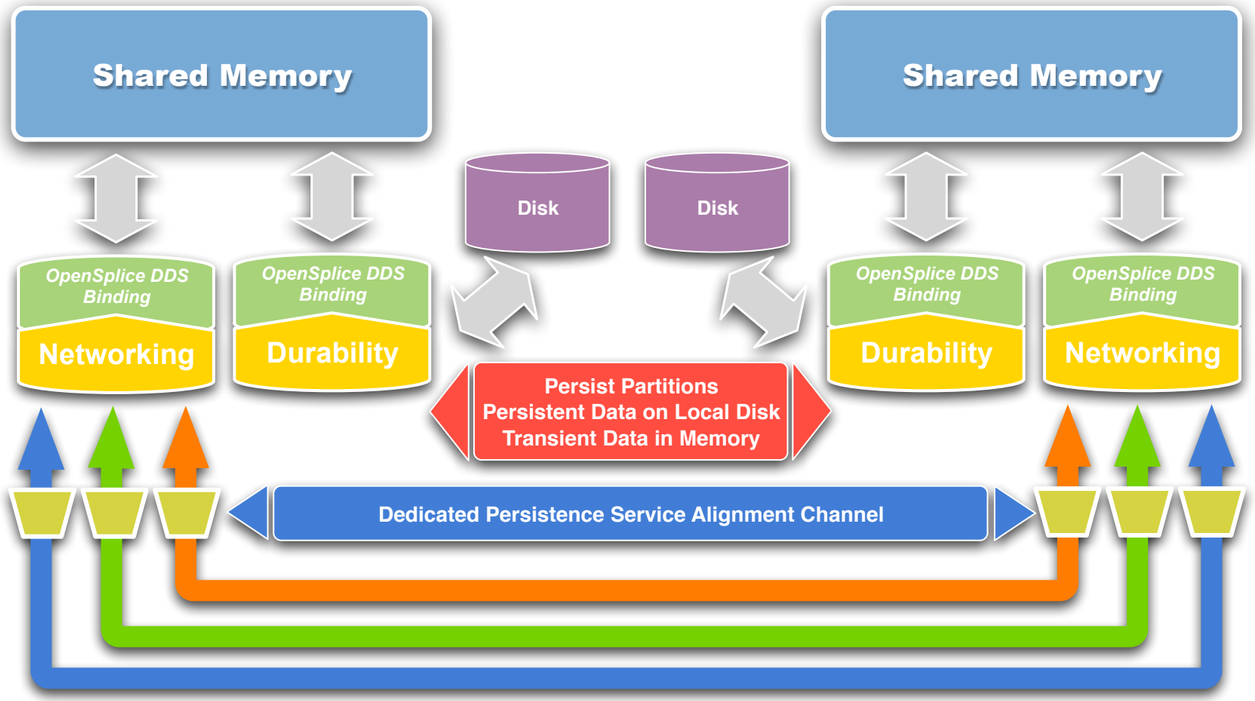

Goal

- **Transient QoS**. Keep **state-data** outside the scope/lifecycle of its publishers
- **Persistence QoS.** Keep **persistent settings** to outlive the system downtime

Features

- **Fault-tolerant availability** of non-volatile data
- Efficient delivery of initial data to **late-joining applications**
- **Pluggable** Durability Service
- **Automatic alignment** of replicated durability-services

OpenSplice DDS

PRISMTECH

OpenSplice | DDS

Delivering Performance, Openness, and Freedom

Playing with OpenSplice DDS

Online Resources

OpenSplice DDS

Delivering Performance, Openness, and Freedom

* http://www.opensplice.com/
* emailto:opensplicedds@prismtech.com

slideshare

Present Yourself

* http://www.slideshare.net/angelo.corsaro

webex™

* http://bit.ly/1Sreg

twitter

* http://twitter.com/acorsaro/

YouTube

* http://www.youtube.com/OpenSpliceTube

Blogger

* http://opensplice.blogspot.com