

OpenspliceDDS Architecture white paper

Version 1.0 - August 31th, 2026



ZETTASCALE TECHNOLOGIES Les Algorithmes 91190, Saint Aubin France

• **Web:** zettascale.tech • **Email:** contact@zettascale.tech • **LinkedIn:** [@zettascaletech](https://www.linkedin.com/company/zettascaletech)

CONFIDENTIAL - This document may contain legally privileged or confidential information. If you are not the recipient or believe you have received this message by mistake, we would be grateful if you would notify the issuer and destroy the content without communicating it to others or reproducing it.



*Innovation shapes the vision.
Hard work brings it to life.
Determination turns it into success."*
— **ZettaScale Technology**

About ZettaScale

ZettaScale's mission is to bring to every connected human and machine the unconstrained freedom to communicate, compute and store — anywhere, at any scale, efficiently and securely. We strive for excellence as engineers and humans -- we firmly believe that we can't separate "the engineer" from "the human". As such we embrace and promote the Seven Knightly Virtues of Courage, Justice, Mercy, Generosity, Faith, Nobility and Hope.



Contents

1. Context	Error! Bookmark not defined.
Table of Contents.....	Error! Bookmark not defined.
Overview	1
The High-level Architecture	2
Federated Architecture with Shared Memory and Network Scheduling.....	2
Simple Standalone/Single Process Architecture.....	5
Configuration Deployments.....	5
DDS Persistent Profile and Durability Service.....	7
Support/Maintenance	7
References and For More Information	7



OpenSpliceDDS is part of the ZettaScale DDS offering - it is a fully featured DDS implementation targeted for enterprise platforms and devices. It is fully interoperable with Cyclone DDS, Vortex Link and Vortex Insight products. The OpenSpliceDDS product offering is made of the Core offering and a set of commercially available add-ons and extra features. The OpenSpliceDDS Core is a full implementation of the Object Management Group (OMG) Data Distribution Service for Real-time Systems (henceforth DDS) rev1.4 standard (DCPS profiles) and the OMG-DDSI-RTPS v2.3 interoperable wire-protocol.

Overview

In wide distributed systems the Network and CPUs (Central Processing Units) are considered the most critical resources. As such, they are deemed the main bottleneck of the overall system performance.

In Real-time systems, CPUs orchestration and scheduling are controlled by the Operation System scheduler who has very limited capabilities to control the Network, neither the Computer's Network Interface Cards nor the Network Routers in use. This lack of network scheduling is problematic when you need to prioritize your data flows.

To take advantage of multicore computing architectures and to address the lack of network scheduling that can manage the data at the DDS level based on its Urgency and Importance, it is crucial to come up with a DDS architecture that can Federate all the applications running on the same computing unit and orchestrate their data dissemination based on their needs and the Quality of Services (QoSs).

Associating QoS to data enables the DDS infrastructure to pre-empt the low-priority traffic in favor of the one with the highest-priority and the most urgent. On the other hand, when there is little processing to rationalize, factorize and federate, the classical way to implement DDS as a set of libraries you link your application code within a Single Process can be good enough. The Single process architecture can be typically recommended when only a singleton application is running on a given computer.



The High-level Architecture

The OpenSpliceDDS architecture can operate and get deployed in two, fully interoperable, modes:

In a Federated mode, for complex architectures where each computer can potentially host several DDS aware applications.

In a library based Standalone mode, when only one application is using DDS and there is nothing, or little to federate or orchestrate.

DEPLOYMENT MODES




Figure 1: OpenSpliceDDS Deployment Modes

Federated Architecture with Shared Memory and Network Scheduling

To rationalize the memory resources when several DDS aware applications are running on the same node, OpenSpliceDDS architecture can support shared memory to minimize the need to keep the data in the address space of each application. The shared memory segment is common to all the local DDS applications; it can be seen as an in-memory real-time database that can be queried using SQL (Structured Query Language). The Federated architecture with shared memory option has also the advantage to have an ultra-low latency inter-core communication.

When configuring OpenSpliceDDS for the Federated deployment architecture, data is physically stored only once on a machine's federation. A smart administration still provides each subscriber within the federation with his own private 'view' on the data space. This allows a reader's data cache to be perceived as an individual 'database' that can be content-



filtered, queried, etc. (using the content-subscription profile as supported by the OpenSpliceDDS Core)[1].

OpenSpliceDDS also has a unique architecture in the market that offers a **DDS Network Scheduler** in order to:

- Organize and classify the data based on their Importance and Urgency, and creates network channels for each urgent and important data stream class. These channels are called priority lanes.
- Pre-empt the less urgent data streams and allocate the network bandwidth to the highest priority, most up-to-date, and most urgent data stream, in compliance with the QoSs attached to each data stream.
- Perform traffic shaping to adapt the amount of data to be sent on the network with the network bandwidth available.
- Confine the data to the subnetworks where it is actually needed when the system is logically and physically partitioned.

QUALITY-OF-SERVICE DRIVEN NETWORKING




Figure 2: OpenSpliceDDS Network Scheduler and its Network priority Lanes

The DDS Network Scheduler allows you to have maximal nodal scalability. It can be configured to appear as the unique proxy of all the local applications on the network. The system scalability is, therefore, depending on the number of machines rather than being correlated to the number of DDS participant applications, as it is the case with the other DDS implementations in the market or with the standalone deployment mode.

The Network Scheduler creates as many network channels as needed and helps you to enforce message priority even on non-priority-preserving transports, such as TCP/IP or UDP/IP.



LOGICAL-TO-PHYSICAL TRAFFIC CONFINEMENT

The DDS Global data space as well as the applications that produce them can be organized in logical groups known as DDS partitions. The Non-OpenSpliceDDS implementations use the same multicast address to spread the data everywhere and rely on the DDSI protocol stack to discard the unwanted data on every node.




Figure 3: DDS Logical to Physical Partitioning Mapping

For instance, if you are building an air surveillance-like system, your system can be organized in four geographical partitions, North, South, East, and West. Each region needs to have a sub-system that detects intrusions and need to react locally and immediately.

When a Radar of given geographical region detects and publishes the position of a given track in the air, you don't want that information to get spread over the entire network and waste your critical network bandwidth and resources and consume unnecessarily CPUs time when the subscribing application that is handling the threat is in the same region as the publisher.

OpenSpliceDDS Network Scheduler allows you to associate different multicast and unicast addresses to the same DDS logical partition to confine, physically, the traffic to the subnetwork where there is real interest in the produced data.

The federated architecture with shared memory results in an extremely low footprint, excellent scalability, and optimal performance when compared to other DDS implementations where each reader and writer are seen as 'communication-endpoints', each one with their own storage.



Simple Standalone/Single Process Architecture

This deployment allows the DDS applications and OpenSpliceDDS middleware libraries to be simply linked together into a single process. One of the advantages of this deployment option is that there is no need to pre-configure a shared memory segment as dynamic heap memory is utilized by the standalone process. Along with the 'zero-configuration', there is no need to start up the federation's services as all functionality is 'bundled' within the application process. The middleware and the data lifecycle are tied to that of the application. Standalone/Single Process deployment is the default architecture provided within OpenSpliceDDS and allows for easy deployment with minimal configuration required for a running DDS system. This kind of simple deployment can suffice when you have one DDS aware application per machine and when there is nothing to federate. It can be particularly useful for embedded platforms.

Recognizing the difficulty in some scenarios of choosing the right architecture upfront in the system design and development cycle, OpenSpliceDDS allows you to run exactly the same application, meaning the same executable application, in different architectural styles by simply changing a runtime deployment descriptor.

Adaptable Deployments

OpenSpliceDDS is highly configurable, allowing system architects to select the Federated Deployment or Standalone deployment modes, even at Runtime.

Both deployment modes support a configurable and extensible set of services, providing functionality such as:

- **Network Scheduling** - providing QoS-driven real-time networking based on multiple reliable priority lanes, as described earlier.
- **Durability Service** - providing fault-tolerant storage for both real-time state data as well as persistent settings.
- **Record and Replay Service**, to record any application topics or DDS built-in topics and replay them later on with similar or different QoSs and properties.
- **Remote Control and Monitoring SOAP Service** - providing remote web-based access using the SOAP protocol from various OpenSpliceDDS tools.
- **Connector service** to connect the DDS Global Data Space to potentially any ODBC Relational Data Base; Web Browser and javascript application through Node.js; and Matlab, Simulink and Labview tools.

When deployed with Federated architecture all the applications collocated on the same machine will share the same memory segment including the DDS services such as Durability and Record and Replay services.




Figure 4: OpenSpliceDDS Federated architecture

When deployed in the standalone architecture, all the DDS infrastructure including the services are linked with the application business logic code as a single process.




Figure 5: OpenSpliceDDS Standalone architecture

OpenSpliceDDS is the only DDS implementation that gives you the choice of selecting, as well as easily adapting, the deployment architecture that best suits your needs.



DDS Persistent Profile and Durability Service

The OMG-DDS Persistence profile is an essential part of the OMG Data Distribution Service OMG Standard. The Persistence profile is implemented by the so-called Durability Service that needs to be fault tolerance and high availability. Unlike most competing DDS implementations, our OpenSpliceDDS Durability Service is included as a standard feature in OpenSpliceDDS.

OpenSpliceDDS Durability Service provides information management capabilities by assuring high availability of non-volatile information via fault-tolerant persistence medias, such as disks.

OpenSpliceDDS Durability Service implementation have unique features that provide the following benefits:

- **Allows late-joining of applications to retrieve the non-volatile data** and the latest system state without reliance on publisher-presence.
- **Application transparent**, as no application code is required to manage it.
- **Location Transparent**, fault-tolerant and replication of non-volatile data. You can deploy as many Durability Service instances in your DDS domain as needed, from any place in the network.
- **Offers the Split-brain syndrome management**, through the supports of powerful merge policies of the non-volatile data sets to recovery from situation where the system's components get disconnected from each other during a long period of time which leads to two non-aligned states of the same data.

OpenSpliceDDS has two durability service implementations. A default implementation that is part of the OpenSpliceDDS Core, and a new commercial add-on called Lite Durability Service implementation. Lite Durability Service is designed for ultra-large-scale systems where applications are frequently disconnected and reconnected or have intermittent connection.

Support/Maintenance

ZettaScale provides world class support, delivering a timely, reliable service to ensure every customer's business success. We offer a wide range of Support and Maintenance Programs for OpenSpliceDDS which can be tailored to customer's best-fit requirements. OpenSpliceDDS Shared Memory/Federated deployment and Standalone/Single Process deployment are supported by annual support programs.

References and For More Information

For further information regarding OpenSpliceDDS, please e-mail: contact@zettascale.tech or visit: www.zettascale.tech



OpenSpliceDDS

DATA-CENTRIC MIDDLEWARE FOR MISSION-CRITICAL SYSTEMS



contact@zettascale.tech | www.zettascale.tech

Copyright © 2026 ZettaScale Technology. All Rights Reserved.