

ROS2: Intestines & Advanced Features

Course: Robotic Programming Environments

Michał Drwięga

michal.drwiega@pwr.edu.pl

www.mdrwiega.com/edu/rpe

Department of Cybernetics and Robotics
Wrocław University of Science and Technology



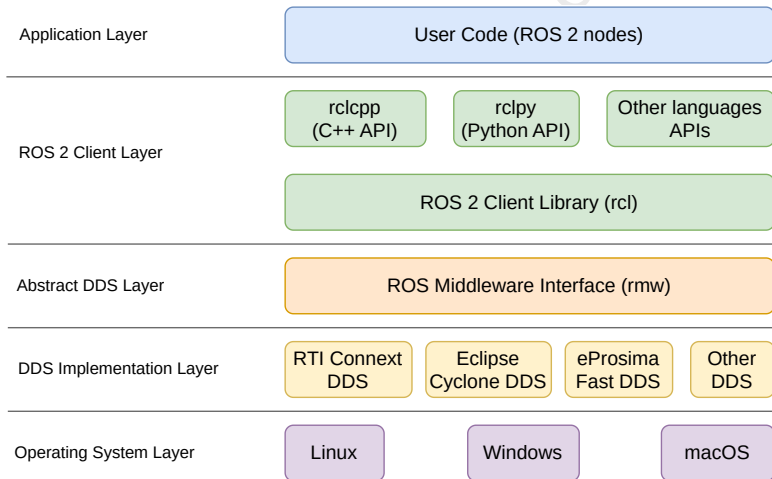
Wrocław University
of Science and Technology

Content

- 1 ROS 2 Architecture
- 2 DDS
- 3 Zenoh
- 4 ROS 2: Environmental Variables
- 5 Nodes execution
- 6 IPC & Composable Nodes

ROS 2 Architecture

ROS 2: Layered Architecture



ROS Client Library (RCL)

The ROS Client Library (RCL) is a **language-agnostic API** in ROS 2 that provides the core functionality for creating nodes, publishers, subscribers, services, and timers.

Key Features:

- Node creation and management.
- Publishers and subscribers for message passing.
- Services and clients for request/response communication.
- Timers and callbacks for periodic tasks.
- Abstracts underlying middleware (RMW) from language-specific implementations.

Usage:

- Available in multiple languages through `rclpy` (Python), `rclcpp` (C++).

ROS Middleware (RMW)

The ROS Middleware (RMW) is an **abstraction layer** in ROS 2 that connects the ROS Client Library (RCL) with the underlying communication middleware (DDS, ROS 1 bridge, etc.).

Key Features:

- Abstracts underlying communication protocol from RCL.
- Handles message transport between nodes.
- Supports multiple middleware implementations (e.g., Fast DDS, Cyclone DDS, RTI Connext).
- Enables cross-language communication between nodes.

Usage:

- RCL calls are translated into RMW calls, which interface with DDS or other middlewares.
- Developers can switch middleware implementations without changing application code.


DDS

For educational use only
©2026 Michał Drwiega

Data Distribution Service (DDS)

• Middleware for real-time systems

- The DDS Middleware is a software layer that abstracts the Application from the details of the operating system, network transport, and low-level data formats.
- High-performance, scalable communication layer
- Used as the core middleware in ROS 2



1

¹<https://www.dds-foundation.org/>


Data Distribution Service (DDS)

- **Data definition**

- Utilizes **Interface Description Language (IDL)** for defining message structures.
- Enables strongly-typed message interfaces

- **Security and control**

- Access control and data flow policies
- Built-in encryption for secure communication
- Supports flexible **Quality of Service (QoS)** for reliability, liveliness, and security.



Data Distribution Service (DDS)


- Implements a publish-subscribe communication pattern.
- DDS is an industry standard, implemented by vendors like RTI, eProsima, and Eclipse.³
- Provides dynamic discovery

³<https://docs.ros.org/>

What is Discovery?


Process where ROS 2 nodes **automatically find each other** and establish communication without manual configuration. ⁴

Simple Discovery (Default)



Peer-to-peer

Discovery Server (Alternative)




Client-server

⁴<https://docs.ros.org/en/foxy/Tutorials/Advanced/Discovery-Server/Discovery-Server.html>

ROS 2 - Standard DDS Discovery

The **Simple Discovery Protocol** (SDP) is the default mechanism defined in the DDS standard.



How it works:


Participants send periodic multicast announcements
All nodes share discovery info with ALL other nodes

Limitations


- **Scaling issue:** Packet count increases *exponentially* with nodes
- **Multicast required:** May not work on WiFi, corporate/school networks

ROS 2 - Centralized Discovery Server

A **Discovery Server** (Fast DDS) acts as a rendezvous point where nodes register and discover each other.



- **Clients** connect to server(s)
- Share only **necessary** discovery info




5

- No multicast required
- Works over WiFi, cloud, WAN


⁵<https://docs.ros.org/en/foxy/Tutorials/Advanced/Discovery-Server/Discovery-Server.html>

Redundancy




Connect to multiple servers for fault tolerance

Backup Server



Persists discovery database to file (`--backup`)

Partitioning



Virtual partitions via server selection

Discovery Server v2 Filtering

Uses **topic information** to decide if nodes need to discover each other, dramatically reducing network traffic [citation:4].

ROS 2 - Using Discovery Server

1. Start Discovery Server

```
1 # Terminal 1
2 fastdds discovery --server-id 0 \
3   --ip-address 0.0.0.0 --port 11811
```

Default port: 11811

2. Configure Nodes

```
1 # Terminal 2 (listener)
2 export ROS_DISCOVERY_SERVER="127.0.0.1:11811"
3   "
4 ros2 run demo_nodes_cpp listener
5 # Terminal 3 (talker)
6 export ROS_DISCOVERY_SERVER="127.0.0.1:11811"
7   "
8 ros2 run demo_nodes_cpp talker
```

Multiple Servers (Redundancy)

```
1 # Start two servers
2 fastdds discovery --server-id 0 -p 11811
3 fastdds discovery --server-id 1 -p 11888
4
5 # Connect to both
6 export ROS_DISCOVERY_SERVER=\
7   "127.0.0.1:11811;127.0.0.1:11888"
```

If one server fails, discovery still works! [citation:4]

ROS 2 Daemon

For CLI tools (ros2 topic list) with Discovery Server:

```
1 export ROS_SUPER_CLIENT=True
```

Super client receives all discovery info

Built-in Security

ROS 2 leverages DDS Security plugins to provide **authentication**, **encryption**, and **access control** without additional software.

Authentication

Verify participant identity
PKI-based (x.509 certs)

Access Control

Define permissions
Topic-level restrictions

Encryption

AES-GCM authenticated
Protect data in transit

ROS 2 - DDS Vendors Comparison

Feature	Fast DDS	Cyclone DDS	RTI Connex
Vendor	eProsima	Eclipse Foundation	RTI
License	Apache 2	Eclipse Public License v2.0	Commercial / Research
ROS 2 Status	Default RMW	Full support	Full support
RMW package	rmw_fastrtps_cpp	rmw_cyclonedds_cpp	rmw_connextdds

Performance characteristics vary by use case

Metric	Fast DDS	Cyclone DDS	RTI Connex
Latency (small messages)	Medium	Lowest	Low
Throughput (large data)	High	Medium	High
CPU utilization	Medium	Lowest	Higher
Memory footprint	Medium	Smallest	Largest
Discovery time (many nodes)	2s	1s	1.5s
Worst-case latency bound	Higher	Medium	More predictable

Fast DDS

Best for: Complex systems, feature-rich apps

Cyclone DDS

Best for: Real-time control, resource-constrained

RTI Connex

Best for: Mission-critical, enterprise

Limitations of DDS in ROS 2

• Network Constraints

- Designed primarily for LAN environments
- Issues with Wi-Fi, NAT traversal, and WAN/cloud communication

• Discovery Overhead

- Heavy reliance on multicast-based discovery
- Increased traffic and slower startup in large systems

• Resource Usage

- Higher CPU and memory consumption
- Less suitable for embedded or resource-constrained devices

• Configuration Complexity

- Requires careful QoS tuning
- Vendor-specific behavior (e.g., Fast DDS vs Cyclone DDS)

• Debugging Difficulty

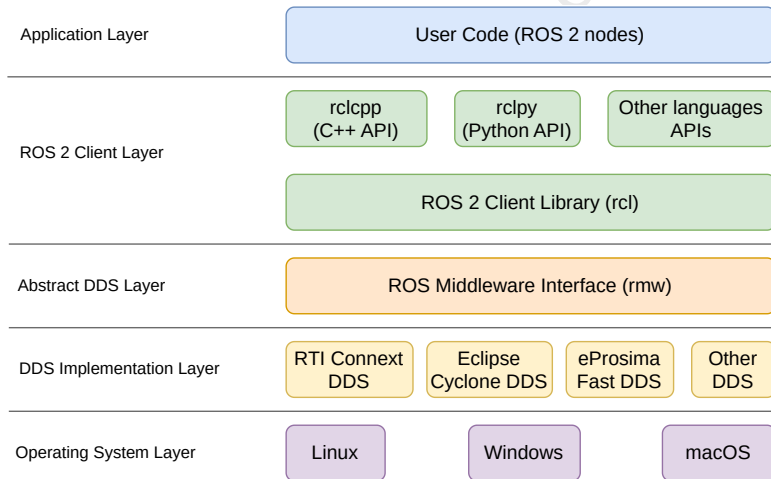
- Non-deterministic issues due to QoS and discovery
- Hard to diagnose communication failures

Motivation: These challenges drive interest in alternative middleware.

Zenoh

For educational use only
©2026 Michał Drwiega

ROS 2: Layered Architecture



What is Zenoh? Zenoh is a data-centric communication middleware designed for distributed systems spanning embedded devices, robots, edge, and cloud.

Core Concepts:

- Unified **publish/subscribe**, **query**, and **storage**
- Data addressed via **key expressions** (not fixed topics)
- Transport-agnostic: UDP, TCP, QUIC, shared memory

Design Goals:

- Low latency and low overhead
- Efficient operation on constrained devices
- Seamless LAN, WAN, and cloud communication

Zenoh vs DDS in Robotics

Aspect	DDS	Zenoh
Communication model	Data-centric pub/sub	Data-centric pub/sub + query + storage
Scope	Primarily LAN / local networks	LAN, WAN, cloud, embedded
Transport	UDP multicast / TCP unicast	UDP, TCP, QUIC, shared memory
Discovery	Automatic (multicast-based)	Flexible (supports WAN / NAT)
Latency	Low, deterministic in LAN	Low, but optimized for WAN and intermittent links
Resource usage	Moderate (requires DDS middleware)	Lightweight, low-overhead for embedded devices
Fault tolerance	Local network only	WAN-aware, handles temporary disconnections

What is PicoZenoh?

- Minimal client implementation of Zenoh protocol
- Designed for **resource-constrained devices** (MCUs, embedded systems)
- Part of the **Zenoh ecosystem**

Key Features

- Very small footprint (tens of kB)
- Written in C for portability
- Supports publish/subscribe and query mechanisms
- Works over multiple transports (UDP, serial, etc.)

Example Use Cases

- Sensor nodes publishing data
- Microcontrollers in robotic systems

ROS 2: Environmental Variables

What is Domain ID?

A **logical partition** that isolates ROS 2 networks. Nodes with different Domain IDs cannot communicate .

Why Use Domain ID?

- Multiple robots in same network
- Separate development/testing
- Avoid interference between teams

Usage

```
1 # Set before running nodes
2 export ROS_DOMAIN_ID=42
3 ros2 run demo_nodes_cpp talker
4
5 # Range: 0-232 (inclusive)
6 # 0 = default
7 # Avoid common IDs (e.g., 1, 2)
```

Important

Domain ID determines UDP port numbers: $7400 + \text{DOMAIN_ID}$ (e.g., ID=42 → port 7442)

ROS 2 - Essential Environment Variables

Core Communication

ROS_DOMAIN_ID

0-232 (default: 0)
Isolates ROS 2 networks

ROS_DISCOVERY_SERVER

IP:port of discovery server
Default: (none)

ROS_SUPER_CLIENT

True/False (CLI tools)
Default: False

DDS/RMW Configuration

RMW_IMPLEMENTATION

rmw_fastrtps_cpp
rmw_cyclonedds_cpp
rmw_connextdds

FASTRTPS_DEFAULT_PROFILES_FILE

XML config file path

CYCLONEDDS_URI

Cyclone DDS config file

Security

ROS_SECURITY_ENABLE

true/false

ROS_SECURITY_STRATEGY

Enforce or Permissive
Default: Permissive

ROS_SECURITY_KEYSTORE

Path to keystore directory

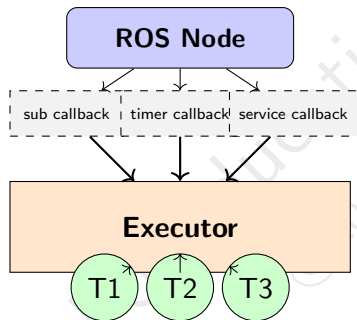
Examples

```
1 # Isolate robot fleet
2 export ROS_DOMAIN_ID=42
3
4 # Switch DDS vendor
5 export RMW_IMPLEMENTATION=
6     rmw_cyclonedds_cpp
7
8 # Enable security
9 export ROS_SECURITY_ENABLE=true
```

Nodes execution

ROS 2 - What is an Executor?

Execution manager that uses one or more threads to invoke callbacks when messages or events arrive





What it does

- Waits for incoming messages/events
- Routes them to appropriate callbacks
- Manages thread(s) for execution

Callback Types

- Subscription callbacks
- Timer callbacks
- Service/action callbacks


Executors



When the `spin()` is called, the current thread starts querying the rcl and middleware layers for incoming messages and other events and calls the corresponding callback functions until the node shuts down. An incoming message is not stored in a queue on the Client Library layer but kept in the middleware until it is taken for processing by a callback function. A wait set is used to inform the Executor about available messages on the middleware layer, with one binary flag per queue. ⁶

⁶<https://docs.ros.org/en/iron/Concepts/Intermediate/About-Executors.html>

Executor types



The **Multi-Threaded Executor** creates a configurable number of threads to allow for processing multiple messages or events in parallel.

The **Static Single-Threaded Executor** optimizes the runtime costs for scanning the structure of a node in terms of subscriptions, timers, service servers, action servers, etc. It performs this scan only once when the node is added, while the other two executors regularly scan for such changes.

7

⁷<https://docs.ros.org/en/iron/Concepts/Intermediate/About-Executors.html>

ROS 2 - Executor Variants

Feature	Static Executor	Dynamic Executor
Callback collection	At creation	Continuously
Performance	Higher	Medium
Flexibility	Lower	Higher
Use case	Fixed node graph	Dynamic changes

Static Executor

- Collects callbacks **once** at start
- No runtime overhead for collection
- Best for fixed-pipeline robots

Dynamic Executor

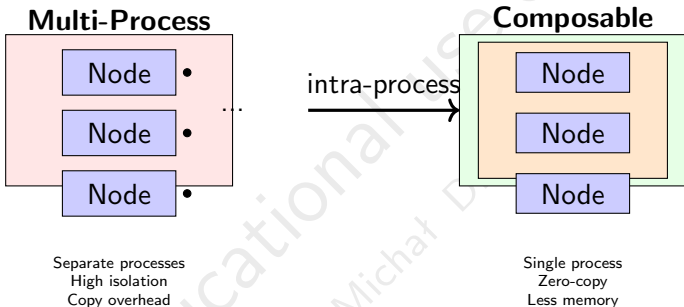
- Continuously checks for new callbacks
- Can add/remove nodes at runtime
- Better for reconfigurable systems

Executors - code example

```
1 rclpy.init()
2 try:
3     node = Node()
4     executor = SingleThreadedExecutor()
5     executor.add_node(node)
6     try:
7         executor.spin()
8     finally:
9         executor.shutdown()
10        node.destroy_node()
11 finally:
12    rclpy.shutdown()
```

IPC & Composable Nodes

ROS 2 - Composable Nodes




Key Takeaway

Same node code - only deployment changes!

Nodes are compiled as shared libraries, loaded into container

Camera Processing Pipeline - All in One Process

Process: camera_pipeline



Intra-process Communication

- Messages passed via **shared memory**
- **Zero-copy** for large data
- No serialization overhead

Traditional Multi-process

- Serialize → copy → deserialize
- Higher latency
- Higher memory usage

ROS 2 - Component Container

Component Container - special node that **loads multiple components** (nodes) at runtime

Command Line

```
1 # Start container
2 ros2 run rclcpp_components
   component_container
3
4 # Load components into container
5 ros2 component load /container \
6   composition composition::Talker
7
8 ros2 component load /container \
9   composition composition::Listener
```

C++ Component

```
1 #include "rclcpp/rclcpp.hpp"
2 #include "rclcpp_components/
   register_node_macro.hpp"
3
4 class MyComponent : public rclcpp::Node {
5 public:
6   MyComponent(const rclcpp::NodeOptions&
7     options)
8     : Node("my_component", options) {
9     // node initialization
10  }
11 };
12 RCLCPP_COMPONENTS_REGISTER_NODE(
   MyComponent)
```

Load/unload nodes dynamically without restarting the process!

ROS 2 - Multi-Process vs Composable

Aspect	Multi-Process	Composable
Isolation	High (crash isolated)	Low (crash kills all)
Latency	Higher (serialization)	Lower (zero-copy)
Memory	Higher (per-process)	Lower (shared)
CPU overhead	Higher	Lower
Data passing	Copy	Shared memory
Debugging	Easier	Harder
Startup time	Slower	Faster

When to use composable nodes

- **Large data** - Images, point clouds
- **Low latency** - Real-time control
- **Limited resources** - Embedded systems
- **Tightly coupled** - Pipeline stages

Summary

- 1 ROS 2 Architecture
- 2 DDS
- 3 Zenoh
- 4 ROS 2: Environmental Variables
- 5 Nodes execution
- 6 IPC & Composable Nodes

Thank you!

For educational use only
©2026 Michał Drwiega