

# rmw\_zenoh

An alternative open-source middleware for ROS 2

Yadunund Vijay

Julien Enoch

ROSCon 2024, Odense





Yadunund  
Vijay

**Senior Software Engineer @  
Intrinsic**



@yadunund



@yadunundvijay



Julien  
Enoch

**Senior Solution Architect @  
ZettaScale**



@JEnoch

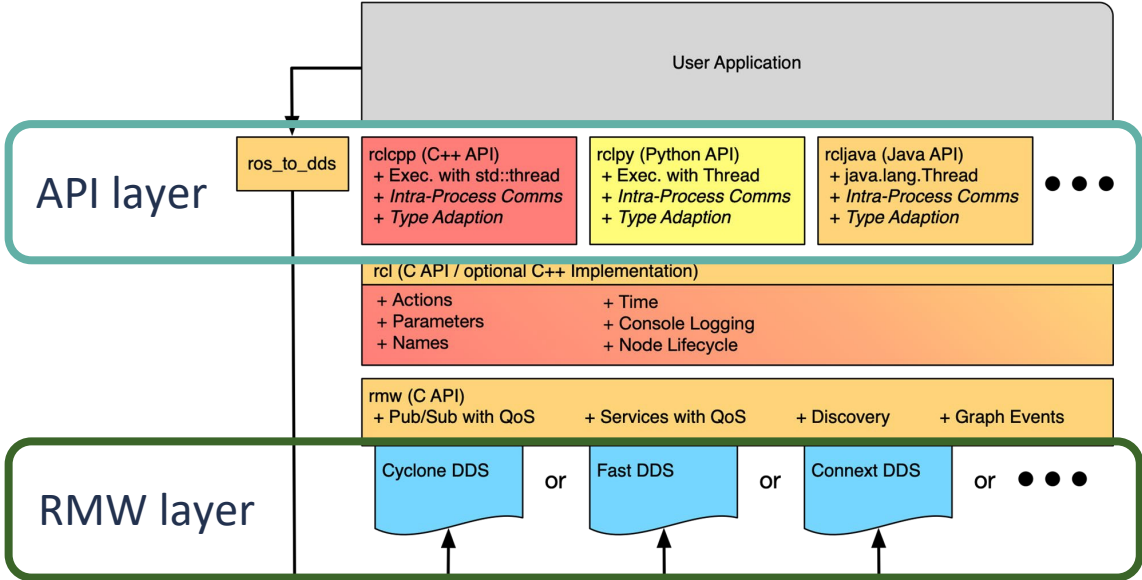


@julienenoach



# ROS 2 has a modular architecture

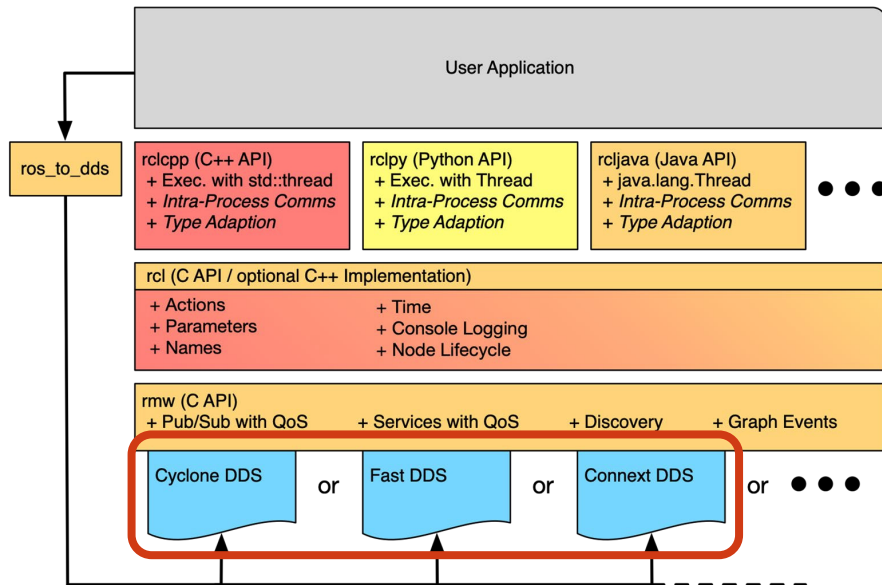
with a runtime-swappable middleware layer!



\* Intra-Process Comms and Type Adaption could be implemented in the client library, but may not currently exist.



But all tier-1 middlewares are DDS-based



\* *Intra-Process Comms* and *Type Adaption* could be implemented in the client library, but may not currently exist.



credible

flexible

reliable

# For good reasons...

[https://design.ros2.org/articles/ros\\_on\\_dds.html](https://design.ros2.org/articles/ros_on_dds.html)

<https://design.ros2.org>

## Technical Credibility

DDS has an extensive list of varied installations which are typically mission critical. DDS has been used in:

- battleships
- large utility installations like dams
- financial systems
- space systems
- flight systems
- train switchboard systems

and many other equally important and varied scenarios. These successful use cases lend credibility to DDS's design being both reliable and flexible.

Not only has DDS met the needs of these use cases, but after talking with users of DDS (in this case government and NASA employees who are also users of ROS), they have all praised its reliability and flexibility. Those same users will note that the flexibility of DDS comes at the cost of complexity. The complexity of the API and configuration of DDS is something that ROS would need to address.

The DDS wire specification ([DDSI-RTSPS](#)) is extremely flexible, allowing it to be used for reliable, high level systems integration as well as real-time applications on embedded devices. Several of the DDS vendors have special implementations of DDS for embedded systems which boast specs related to library size and memory footprint on the scale of tens or hundreds of kilobytes. Since DDS is implemented, by default, on UDP, it does not depend on a reliable transport or hardware for communication. This means that DDS has to reinvent the reliability wheel (basically TCP plus or minus some features), but in exchange DDS gains portability and control over the behavior. Control over several parameters of reliability, what DDS calls Quality of Service (QoS), gives maximum flexibility in controlling the behavior of communication. For example, if you are concerned about latency, like for soft real-time, you can basically tune DDS to be just a UDP blaster. In another scenario you might need something that behaves like TCP, but needs to be more tolerant to long dropouts, and with DDS all of these things can be controlled by changing the QoS parameters.

\* *Intra-Process Comms* and *Type Adaption* could be implemented in the client library, but may not currently exist.

ROS™



## But we learnt of some challenges

that lead to poor out-of-the box experience with ROS 2.

discovery issues

large payloads

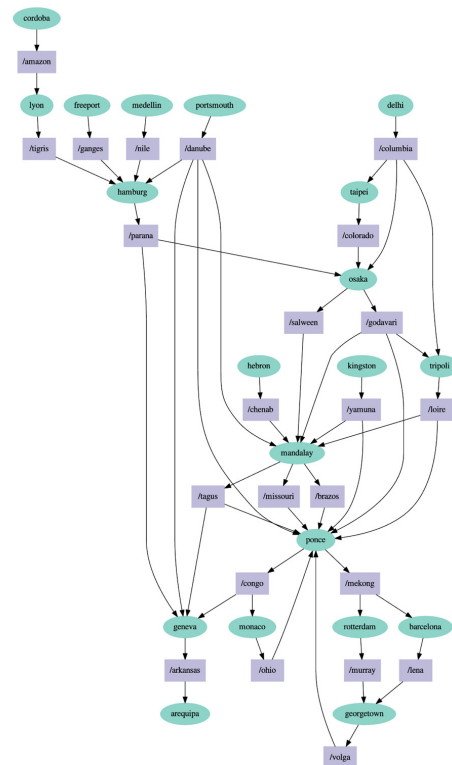
parameter tuning

containers, firewalls and VPNs



# Challenges are inherent to DDS

- **Peer-to-peer only**  
All participants discovers details from all other participants.
- **Heavy discovery protocol**  
 $O(N^2)$  discovery messages.  
N is number of Participants, Topics, Readers, Writers.
- **Over UDP (multicast and unicast)**  
Fragile to message losses, especially with large payloads.  
Retransmissions  $\Rightarrow$  more traffic  $\Rightarrow$  more losses  
No NAT/Firewall traversal  
No optimization nor hardware acceleration such as for TCP.



# Some must-haves for the new middleware



## **Pub/Sub**

For data-centric distributed systems

## **Transport**

Ability to send multi-megabyte payloads reliably.

## **Low latency**

Especially for sending small messages at a very high frequency.

## **Resilient**

Robust against network disconnections.

## **Discovery**

Built-in discovery and the ability to restart discovery without restarting all the nodes.

## **Security**

Access control, authentication and encryption.

<https://discourse.ros.org/t/ros-2-alternative-middleware-report/33771>





# Alternative middlewares considered

 OpenCyphal

*TCPROS*

 eCAL





*LCM*



 kafka

*NNg*

 Zenoh

 ROS™

Requirements		Middleware Options										
Item	Level	Cyphal	DDS	eCal	Kafka	LCM	MQTT	MNG	OPC-UA	TCPROS	Zenoh	ZeroMQ
Configure which interface to use	Must	Yes (IP address)	Yes	Yes (via standard Linux configuration)	Yes	Yes	Yes, implementation dependant. See mosquitto			Indirectly through the	Yes, configuration variable	Yes (both tcp and udp during zmq_bind())
Ability to send multi-megabyte messages	Must	Y	Borderline; in UDP mode									Should be since TCP, see <a href="https://github.com/jeffbass/imagezmq#why-use-imagezmq">https://github.com/jeffbass/imagezmq#why-use-imagezmq</a>
Ability to send fast small messages	Must	Y									Yes, relying on this analysis	Yes
Restart discovery without restarting all nodes	Must	Yes since the channels are statically configured.	Yes	Yes (UDP multicast discovery)	Yes	Yes	Yes	N/A (discovery is not builtin)	N/A	No	Yes: Discovery happens when a zenoh session is created in each application. It looks for other peers/routers.	N/A

**Zenoh**

<https://discourse.ros.org/t/ros-2-alternative-middleware-report/33771>

Matches survey suggestion

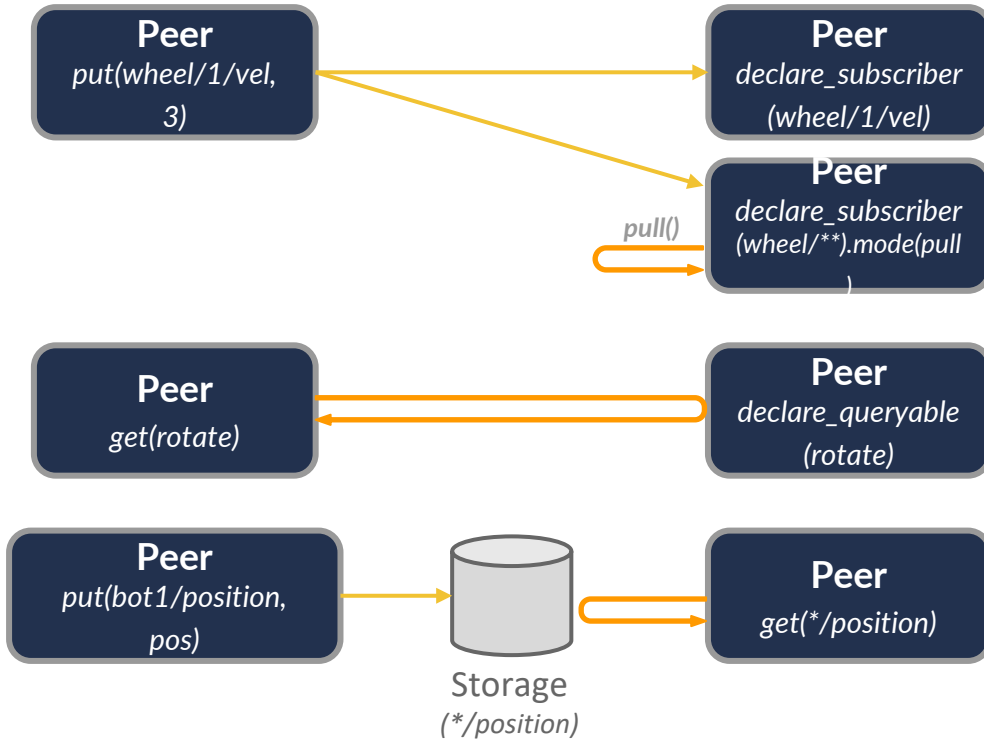


# Zenoh

Zero overhead network protocol



# Pub/Sub/Query protocol



Unifies data in motion, data at rest and computations from embedded microcontrollers up to the data center.


# Zenoh Extensions



**Publication Cache**

**Querying Subscriber**


**Liveliness Tokens:  
for discovery and  
supervision of anything**



# Runs everywhere



**Native libraries** and **API bindings** for many programming languages.



Over various **network technologies**: from **transport layer** to **data link**.

With support of TCP, UDP, TLS, QUIC, serial...

On **embedded** and **constrained devices**

# Any topology

## Peer-to-peer

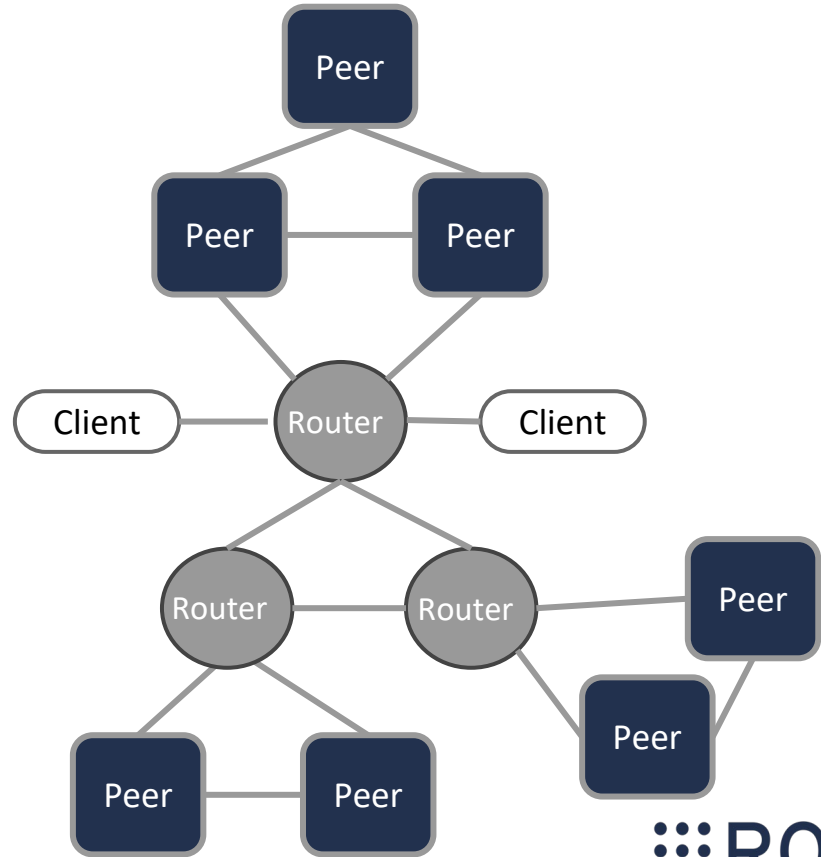
Clique and mesh topologies

## Brokered

Clients communicate through a router or a peer

## Routed

Routers forward data and requests between peers and clients



# rmw\_zenoh

An RMW implementation based on Zenoh

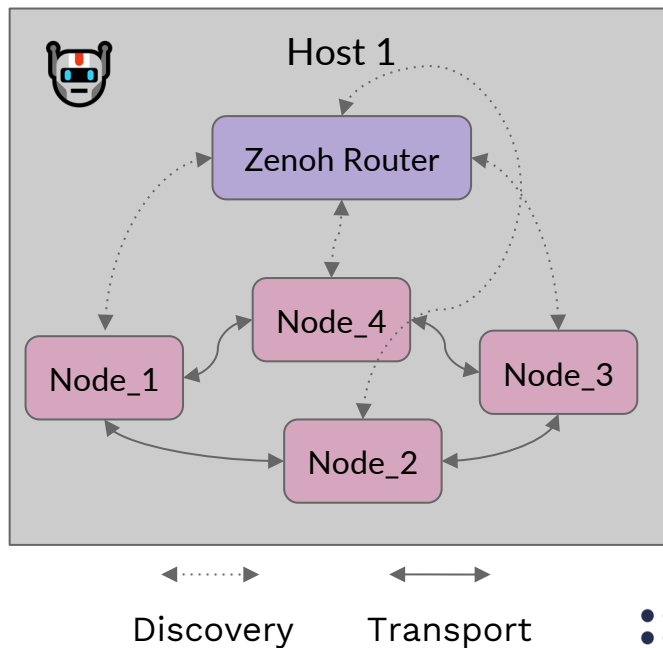


[https://github.com/ros2/rmw\\_zenoh](https://github.com/ros2/rmw_zenoh)

# rmw\_zenoh

- RMW implementation written with zenoh-c(pp) binding.
- TCP for discovery and transport.
- No QoS mismatches.
- Default topology
  - ❑ Discovery is brokered by the Zenoh router
  - ❑ Data transmission is P2P.
  - ❑ Discovery range is localhost only

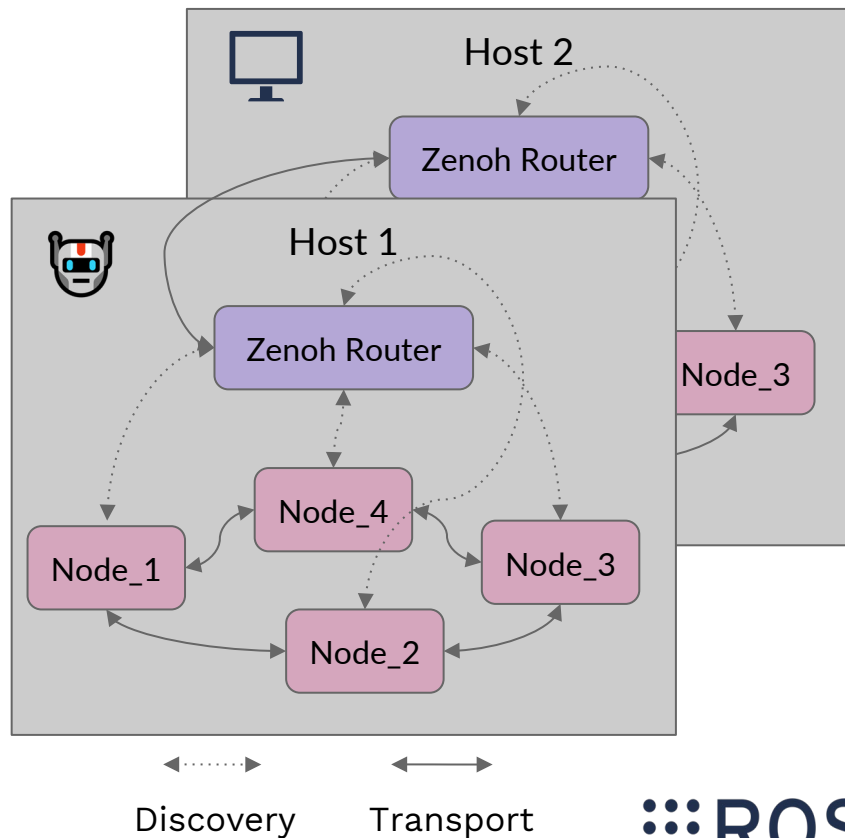
```
export RMW_IMPLEMENTATION=rmw_zenoh_cpp
```



[https://github.com/ros2/rmw\\_zenoh](https://github.com/ros2/rmw_zenoh)


# rmw\_zenoh

- Default multi-host topology
  - Different hosts connect through routers
    - Brokered data transfer
    - Opportunity to downsample




[https://github.com/ros2/rmw\\_zenoh](https://github.com/ros2/rmw_zenoh)

# Other configurable topologies



←·····→  
TCP  
Discovery

↔  
TCP  
Transport



←·····→  
UDP  
multicast  
Discovery

↔  
TCP  
Transport

[https://github.com/ros2/rmw\\_zenoh](https://github.com/ros2/rmw_zenoh)

# Let's talk about the router

- **Is it similar ROS 1's roscore?**

Yes, but it does a lot more.

- **What if the router crashes?**

No impact on running Nodes.

ROS Daemon still present for graph cache.

Just restart the router! **No need to re-launch your Nodes.**

- **Is the router mandatory?**

No. You can configure Zenoh for UDP multicast discovery.

Router for discovery,  
but peer-to-peer  
communications

1.

# If router crashes, peer-to-peer communications remain

```
guest@jazzy: ~/ws_rmw_zenoh
guest@jazzy:~/ws_rmw_zenoh$ ros2 run rmw_zenoh_cpp rmw_zenohd
```

```
guest@jazzy: ~/ws_rmw_zenoh
guest@jazzy:~/ws_rmw_zenoh$ ros2 run demo_nodes_cpp talker
```

```
guest@jazzy: ~/ws_rmw_zenoh
guest@jazzy:~/ws_rmw_zenoh$ ros2 run demo_nodes_cpp listener
```

# 2.

Discovery is robust

# 10 Nodes, 1000 Topics discovered in less than 1s over WiFi

```
guest@jazzy: ~/ws_rmw_zenoh
guest@jazzy:~/ws_rmw_zenoh$ ros2 run rmw_zenoh_cpp rmw_zenohd
```

```
root@colima:/ros_ws# grep -A3 "connect:" zenoh_confs/ROUTER_CONFIG.json5
connect: {
  endpoints: [
    "tcp/192.200.40.12:7447"
  ],
root@colima:/ros_ws#
```



```
guest@jazzy: ~/ws_rmw_zenoh
guest@jazzy:~/ws_rmw_zenoh$ for i in {1..10}; do ros2 run cpp_pubsub talker -r __node:=N_$i & sleep 0.1 ; done
```

```
root@colima:/ros_ws# ros2 topic list --no-daemon
```

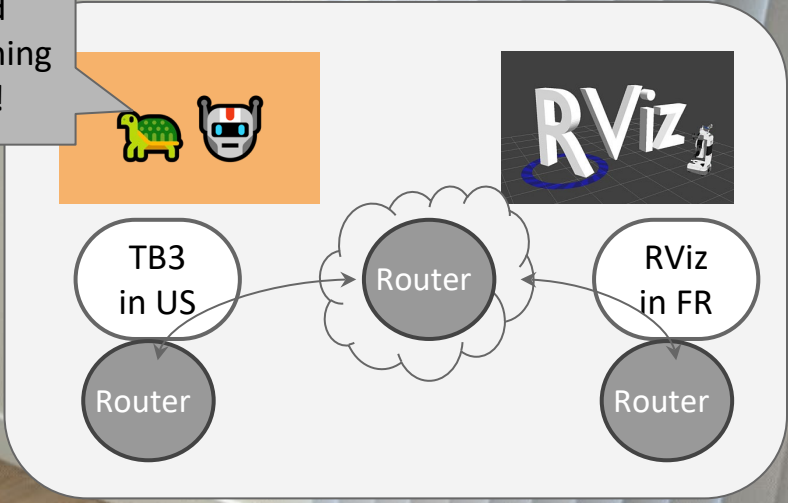


# 3.

Transport reliably over  
many network hops



nav2 and  
drivers running  
on RPi 4!




# 4.

Downsample when  
needed




```
2024-10-16T18:56:51.257911Z [INFO] acc=0 ThreadId(84) zench_link_
pted TCP connection on [::ffff:127.0.0.1]:7447: [::ffff:127.0.0.
2024-10-16T18:56:52.776449Z [INFO] acc=0 ThreadId(84) zench_link_
pted TCP connection on [::ffff:192.200.40.12]:7447: [::ffff:192.
2024-10-16T18:56:53.267829Z [INFO] acc=0 ThreadId(84) zench_link_
pted TCP connection on [::ffff:127.0.0.1]:7447: [::ffff:127.0.0.
2024-10-16T18:56:59.283744Z [INFO] acc=0 ThreadId(84) zench_link_
pted TCP connection on [::ffff:127.0.0.1]:7447: [::ffff:127.0.0.
2024-10-16T18:57:03.267460Z [INFO] acc=0 ThreadId(84) zench_link_
pted TCP connection on [::ffff:127.0.0.1]:7447: [::ffff:127.0.0.
2024-10-16T18:57:07.269792Z [INFO] acc=0 ThreadId(84) zench_link_
pted TCP connection on [::ffff:127.0.0.1]:7447: [::ffff:127.0.0.
2024-10-16T18:57:11.275569Z [INFO] acc=0 ThreadId(84) zench_link_
pted TCP connection on [::ffff:127.0.0.1]:7447: [::ffff:127.0.0.
2024-10-16T18:57:15.279290Z [INFO] acc=0 ThreadId(84) zench_link_
pted TCP connection on [::ffff:127.0.0.1]:7447: [::ffff:127.0.0.
2024-10-16T18:57:19.282812Z [INFO] acc=0 ThreadId(84) zench_link_
pted TCP connection on [::ffff:127.0.0.1]:7447: [::ffff:127.0.0.
```

```
min: 0.832s max: 0.836s std dev: 0.00151s window: 20
average rate: 30.136
min: 0.827s max: 0.838s std dev: 0.00250s window: 20
average rate: 30.826
min: 0.828s max: 0.844s std dev: 0.00445s window: 20
average rate: 30.146
min: 0.829s max: 0.838s std dev: 0.00227s window: 20
average rate: 29.885
min: 0.817s max: 0.859s std dev: 0.00324s window: 20
average rate: 29.945
min: 0.838s max: 0.837s std dev: 0.00100s window: 20
```



# Known limitations

- `rclcpp::shutdown()` must explicitly be called before program termination.
- Router must be manually started (for now).
- Liveliness and deadline QoS events not supported.




# Road to Tier-1 status

[https://github.com/ros2/rmw\\_zenoh/issues/265](https://github.com/ros2/rmw_zenoh/issues/265)

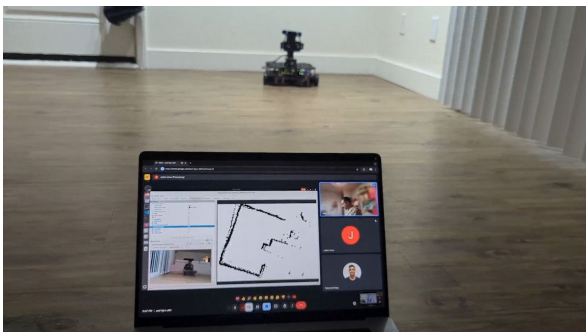
- Targeted for Kilted Kaiju release.
- Make all system tests pass.
- Windows support.
- SROS2 integration.

We're testing  
extensively




nav2, camera driver running on RPi 4  
RViz opened on laptop over wifi

Open-RMF with  
34 nodes 82  
topics and 251  
services with an  
additional RViz  
window open on a  
laptop over wifi.



TB3 in US, RViz in Paris, via a router in  
Cloud



Moveit2




[https://github.com/ros2/rmw\\_zenoh](https://github.com/ros2/rmw_zenoh)

# Try rmw\_zenoh

**Run your existing applications  
with it!**

**Benchmark it!**

**And tell us...**



# Questions?




special thanks to




**Chris Lalancette**  
clalancette




**Morgan Quigley**  
codebot



**Franco Cipollone**  
francocipollone




**Geoffrey Biggs**  
gbiggs



**methylDragon**  
methylDragon




**Esteve Fernandez**  
estev



**Andreas Bihlmaier**  
andreasBihlmaier



**Alejandro Hernández Cordero**



**Steven Palma**  
imstevenprwork



**Yuyuan Yuan**  
YuanYuYuan



**ChenYing Kuo (CY)**  
evshary



**Luca Cominardi**  
Mallets

And all contributors  
to eclipse-zenoh :

