

rmw_zenoh

An alternative open-source middleware for ROS 2

Phani Gangula

Senior Solutions Architect

phani@zettascale.tech

Background

Zenoh: The RMW Alternate

Zenoh has been selected as the first non-DDS protocol to be natively supported in ROS2

Intrinsic is already implementing Zenoh RMW as the major contribution for next ROS2 release

The full report is available [here](#)

2023-09 ROS 2 RMW alternate

Abstract.....	1
User Challenges with DDS.....	2
DDS has a fully-connected graph of participants.....	2
DDS uses UDP multicast for discovery.....	3
DDS can have difficulty transferring large data.....	3
DDS can struggle on some WiFi networks.....	4
DDS tends to have complex tuning parameters.....	4
Vendor specific non-standard DDS extensions.....	4
Next Steps.....	5
Requirements gathering.....	5
User Survey.....	5
Demographics.....	6
Technical Data.....	6
Alternative middleware options.....	8
Requirements.....	9
Comparative analysis of currently available middlewares.....	11
Complete list of investigated middlewares.....	11
Performance.....	13
Middlewares X Requirements.....	13
Conclusion.....	14
Appendix A.....	15

Conclusion

Requirements from ROS 2 users were gathered, and middleware options that are available today were investigated. The research has concluded that Zenoh best meets the requirements, and will be chosen as an alternative middleware. Zenoh was also the most-recommended alternative by users. It can be viewed as a modern version of the TCPROS implementation, and meets most of the ROS 2 requirements. There are still a number of design decisions to be made regarding this implementation; those details will be discussed on <https://discourse.ros.org> as development begins.



Zenoh stands-out as the protocol that satisfies essentially all requirements identified by the survey

Middleware options

<https://docs.google.com/spreadsheets/d/18SgD-aFJAIDus5cYMN-va9greQgXq-ErVEsWuv40Kjw/edit#gid=1189402529>

Zenoh





Zenoh

Pub/Sub/Query protocol that **Unifies data** in **motion**, data at **rest** and **computations** from embedded microcontrollers up the data centre

Provides **location-transparent** abstractions for **high performance pub/sub** and **distributed queries** across heterogeneous systems

Built-in support for zero-copy and shared memory

Runs Everywhere


Written in Rust for security, safety and performance

Native libraries and API bindings for many programming languages, e.g., Rust, C/C++, Python, JS, REST, C# and Kotlin

Built-in support Shared Memory and Zero Copy

Supports network technologies from transport layer down-to the data link. Currently runs on, TCP/IP, UDP/IP, QUIC, Serial, Bluetooth, OpenThreadX, Unix Sockets.

Available on embedded and extremely constrained devices and networks — 5 bytes minimal overhead




Any Topology


Peer-to-peer

Clique and mesh topologies

Clique



Mesh




Any Topology


Peer-to-peer

Clique and mesh topologies

Clique




Mesh




Brokered

Clients communicate through
a router or a peer



Brokered



Routed

Any Topology

Peer-to-peer

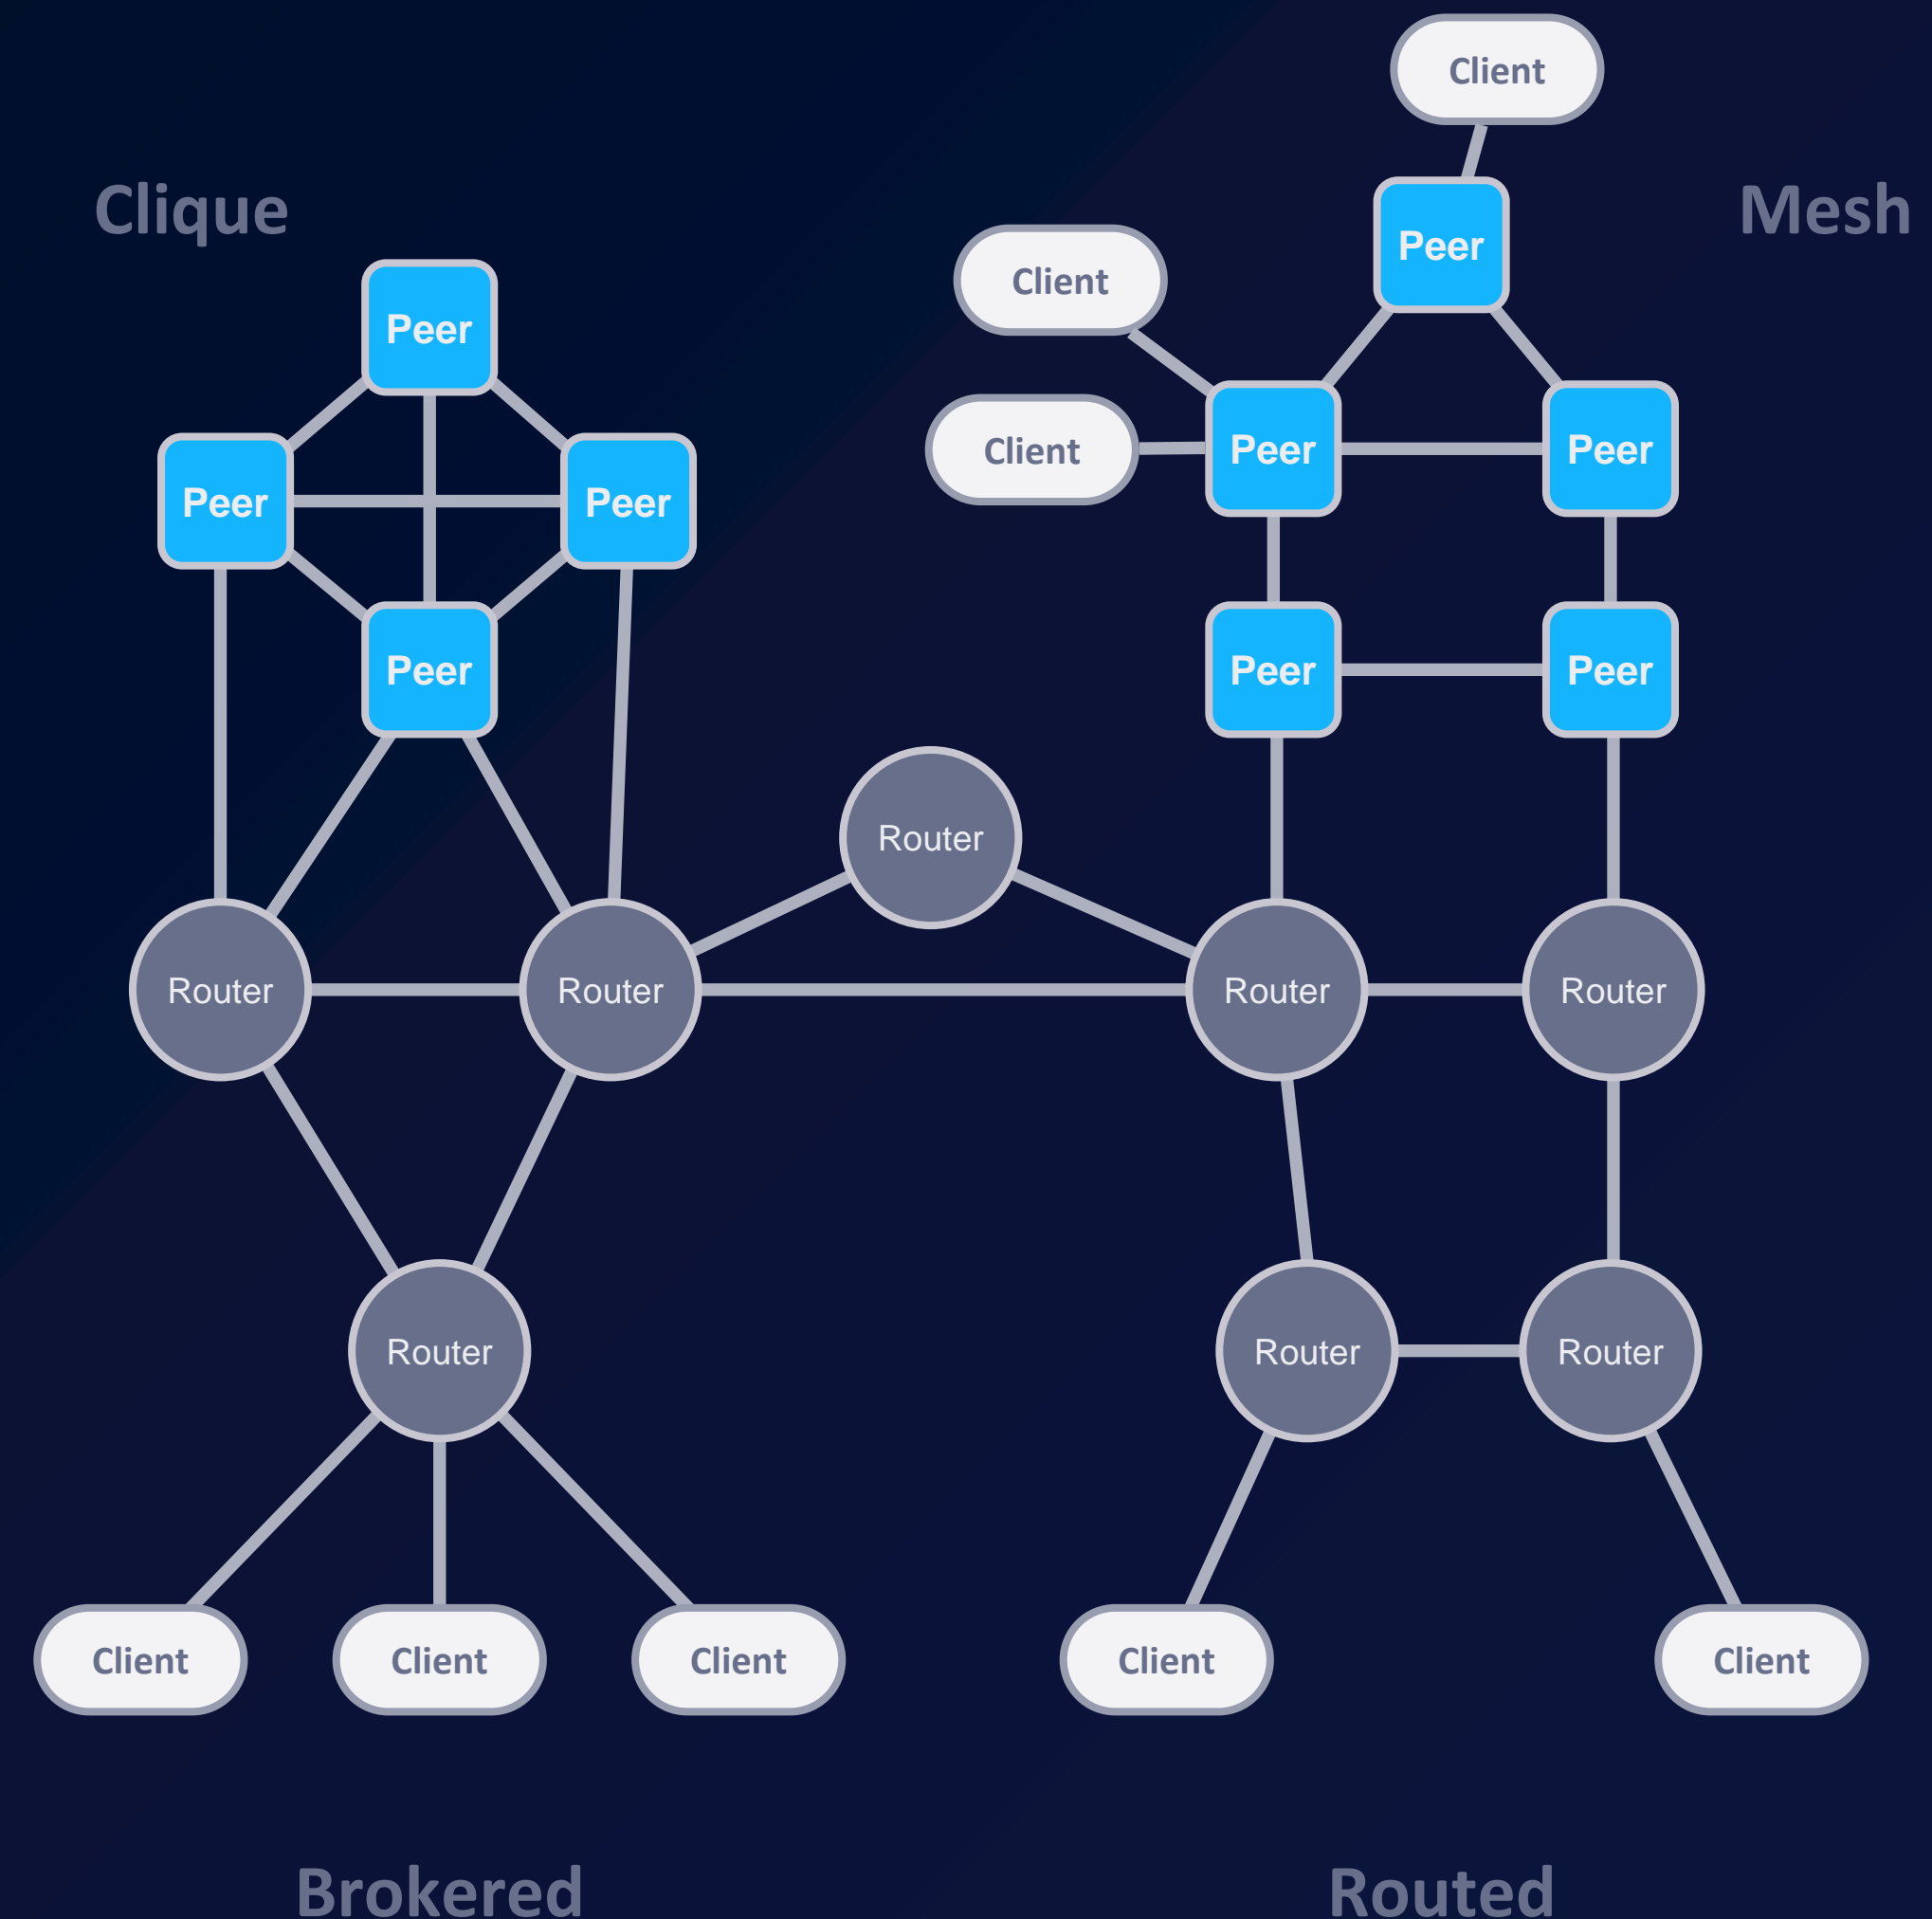
Clique and mesh topologies

Brokered



Clients communicate through a router or a peer

Routed

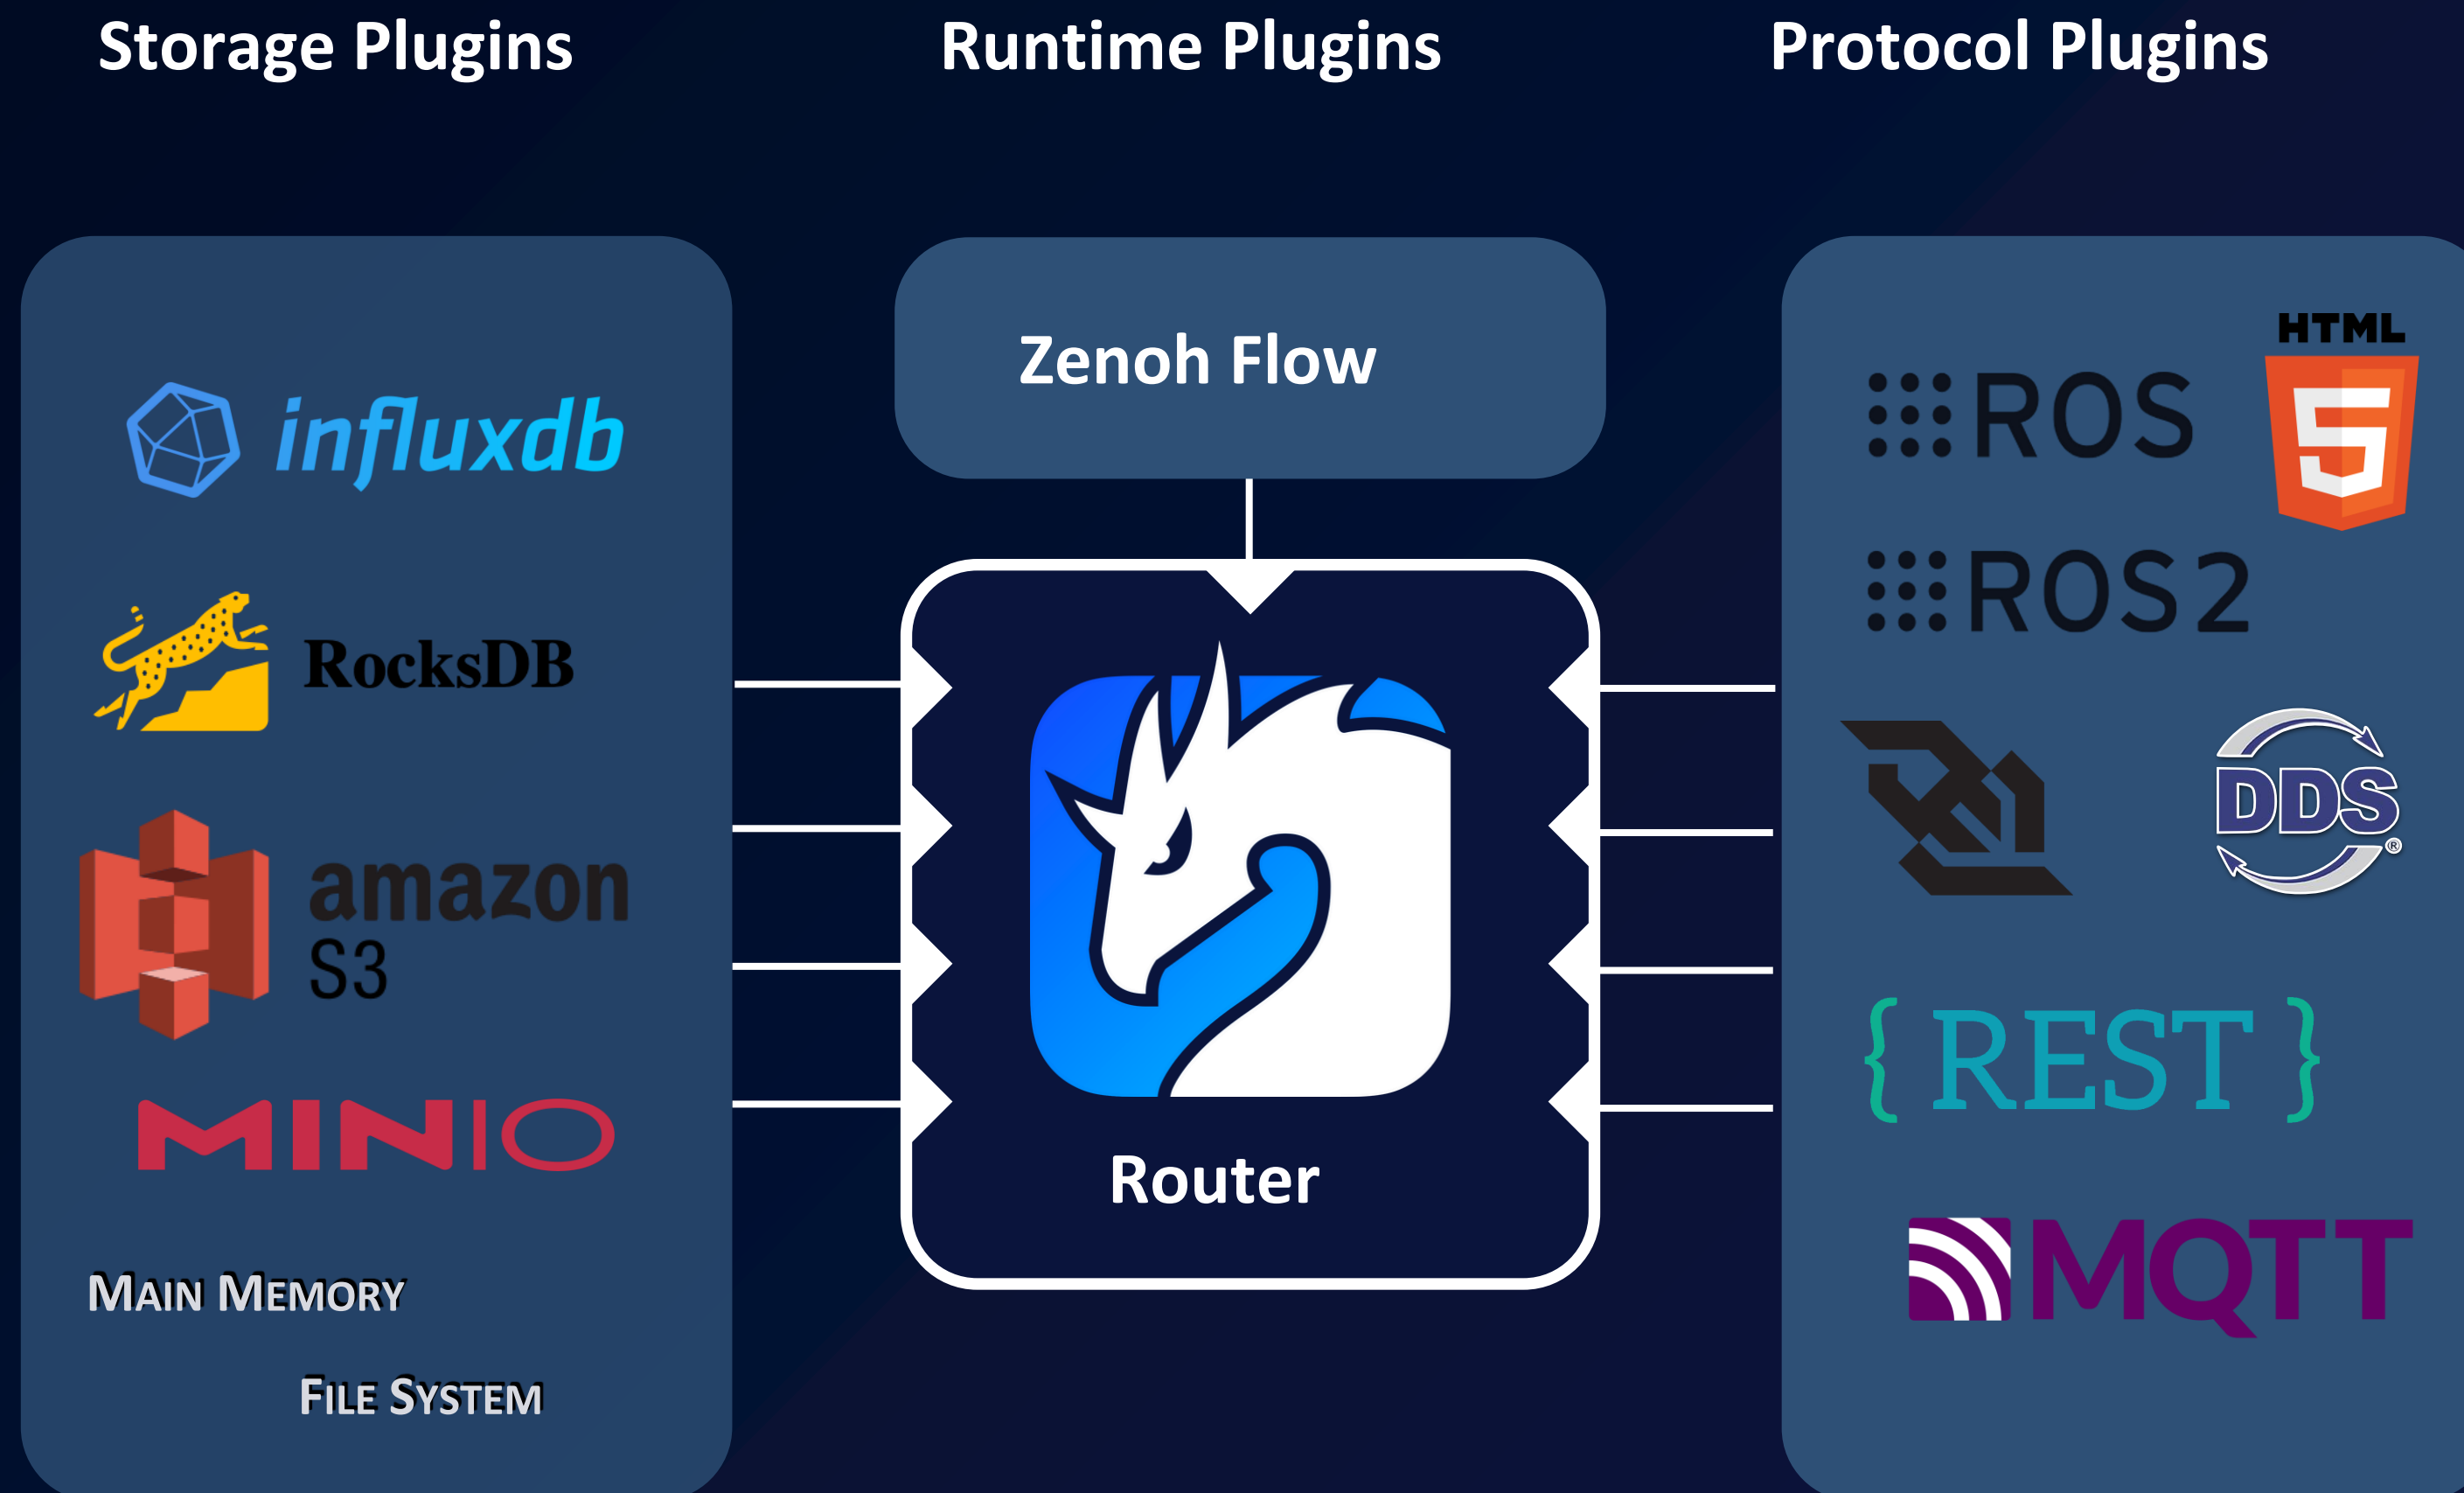
Routers forward data to and from peers and clients



Topology in Perspective



Plug-Ins



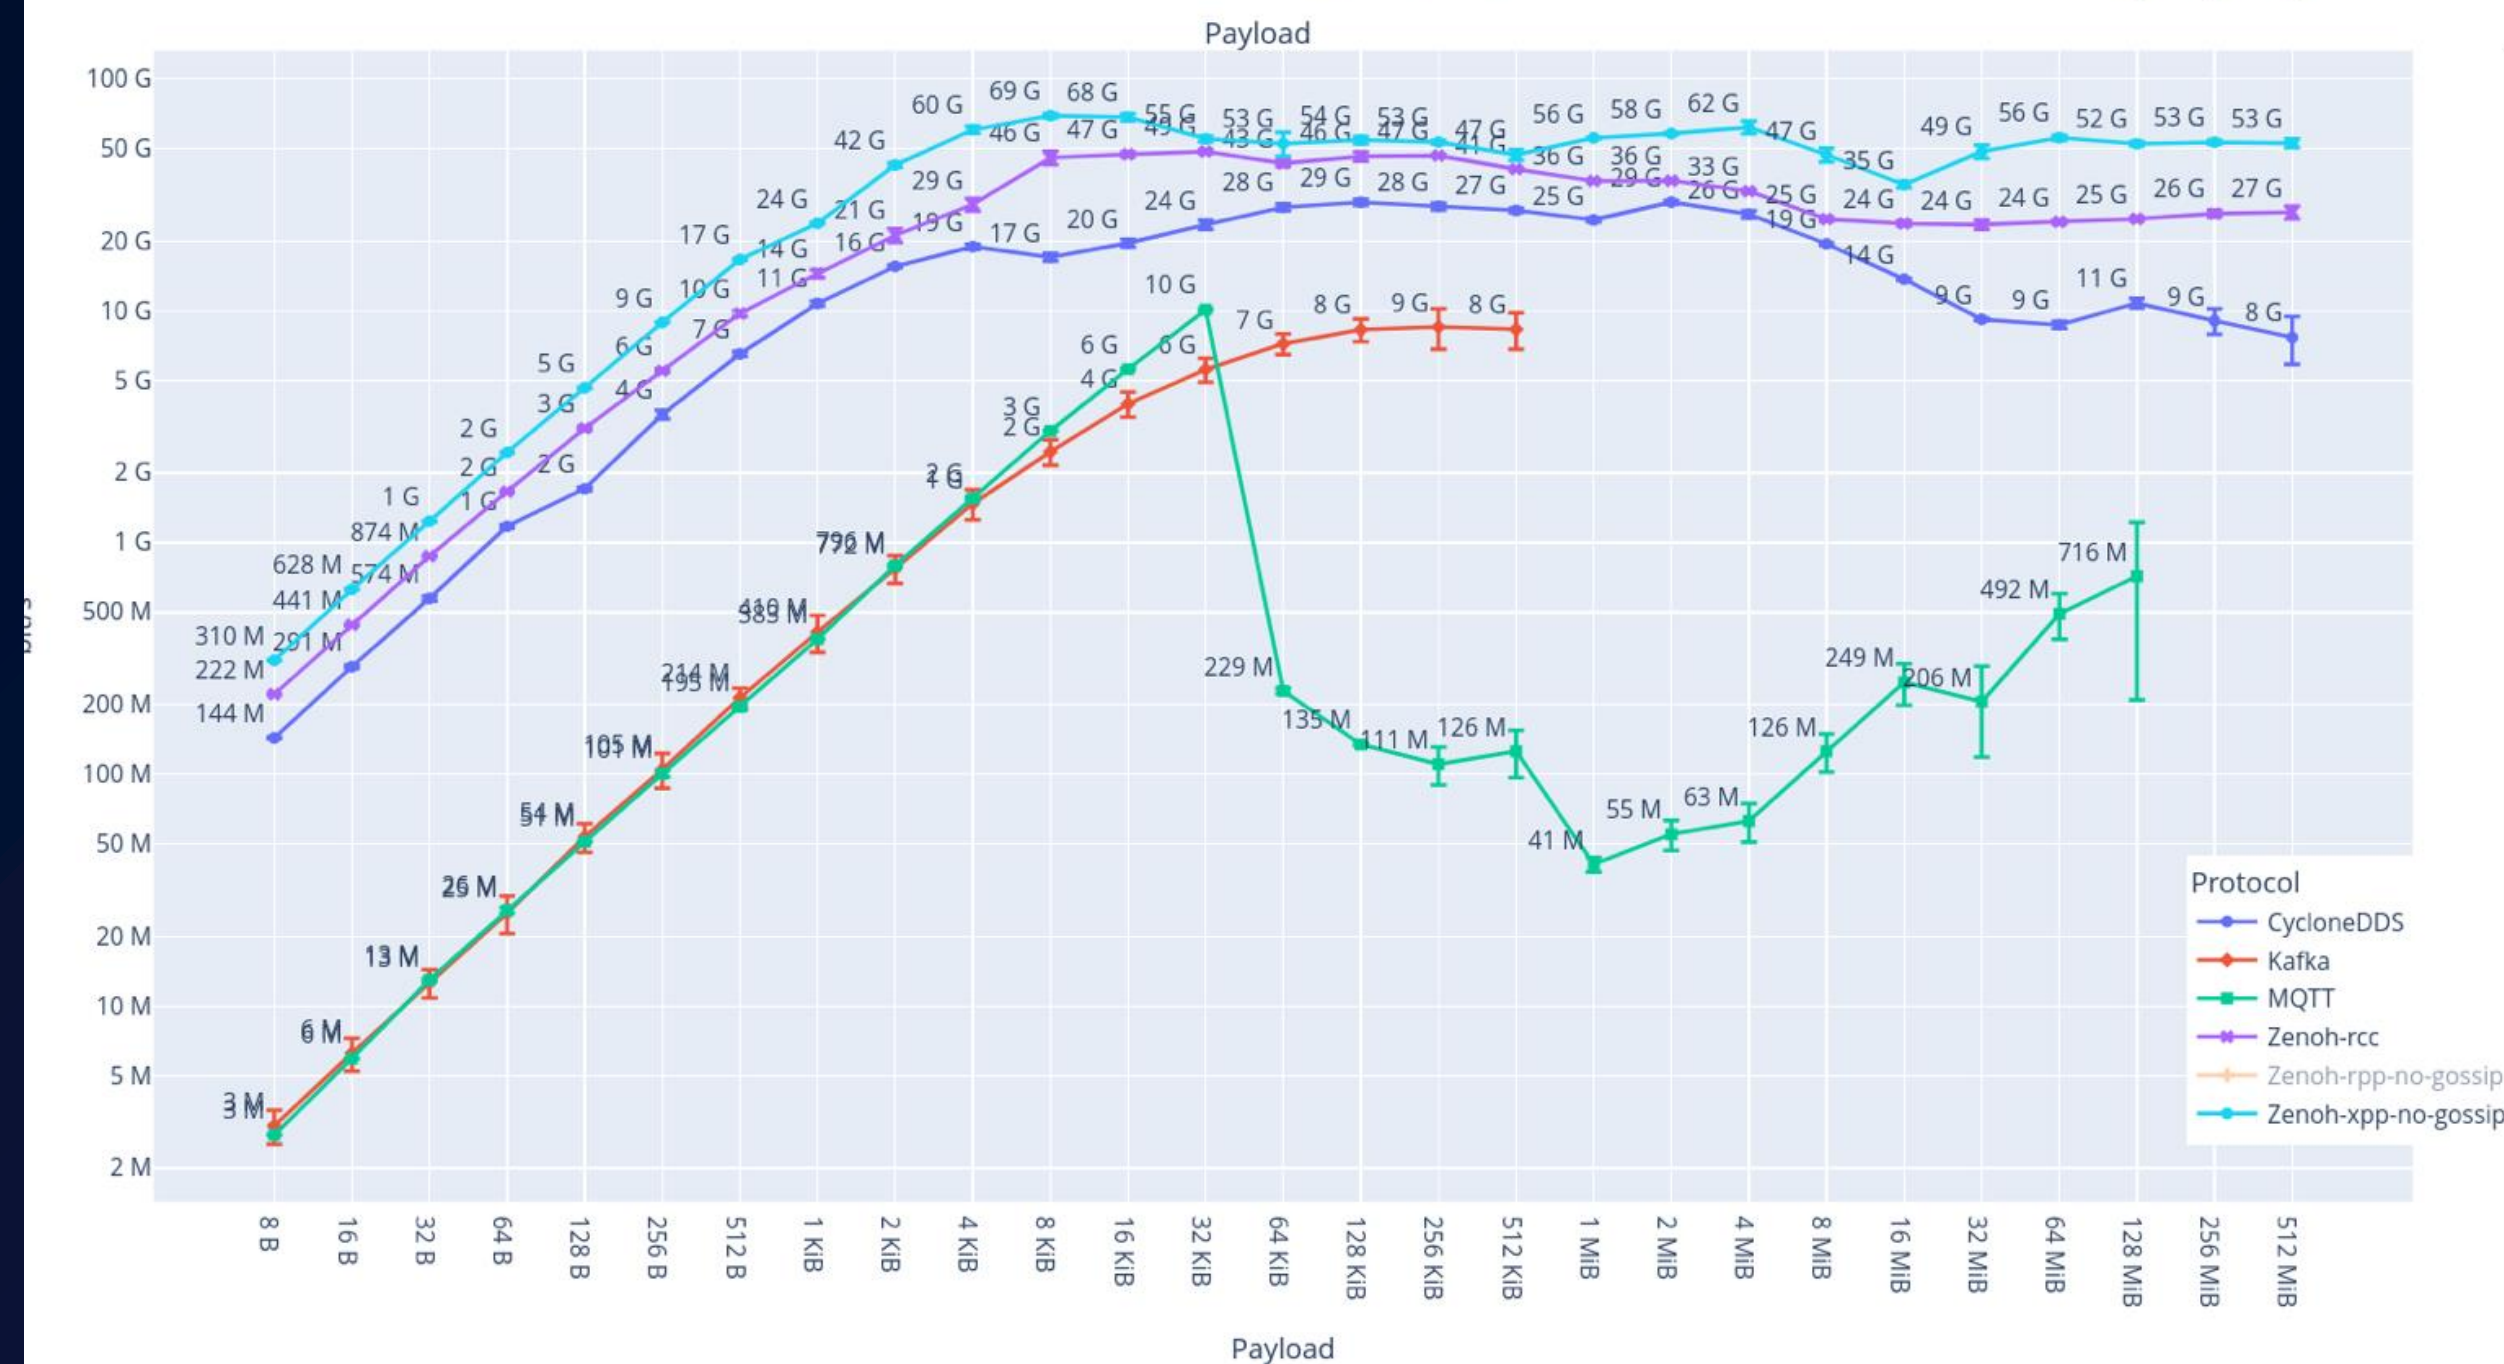

Zenoh vs DDS, MQTT & Kafka

Zenoh can deliver at peak performance of ~70Gbps at 8Kb payload:

- 3.3x higher than DDS
- 23x higher than Kafka
- 35x than MQTT (higher for larger payload)

Zenoh's latency 10us, 6us with ultra-low latency support

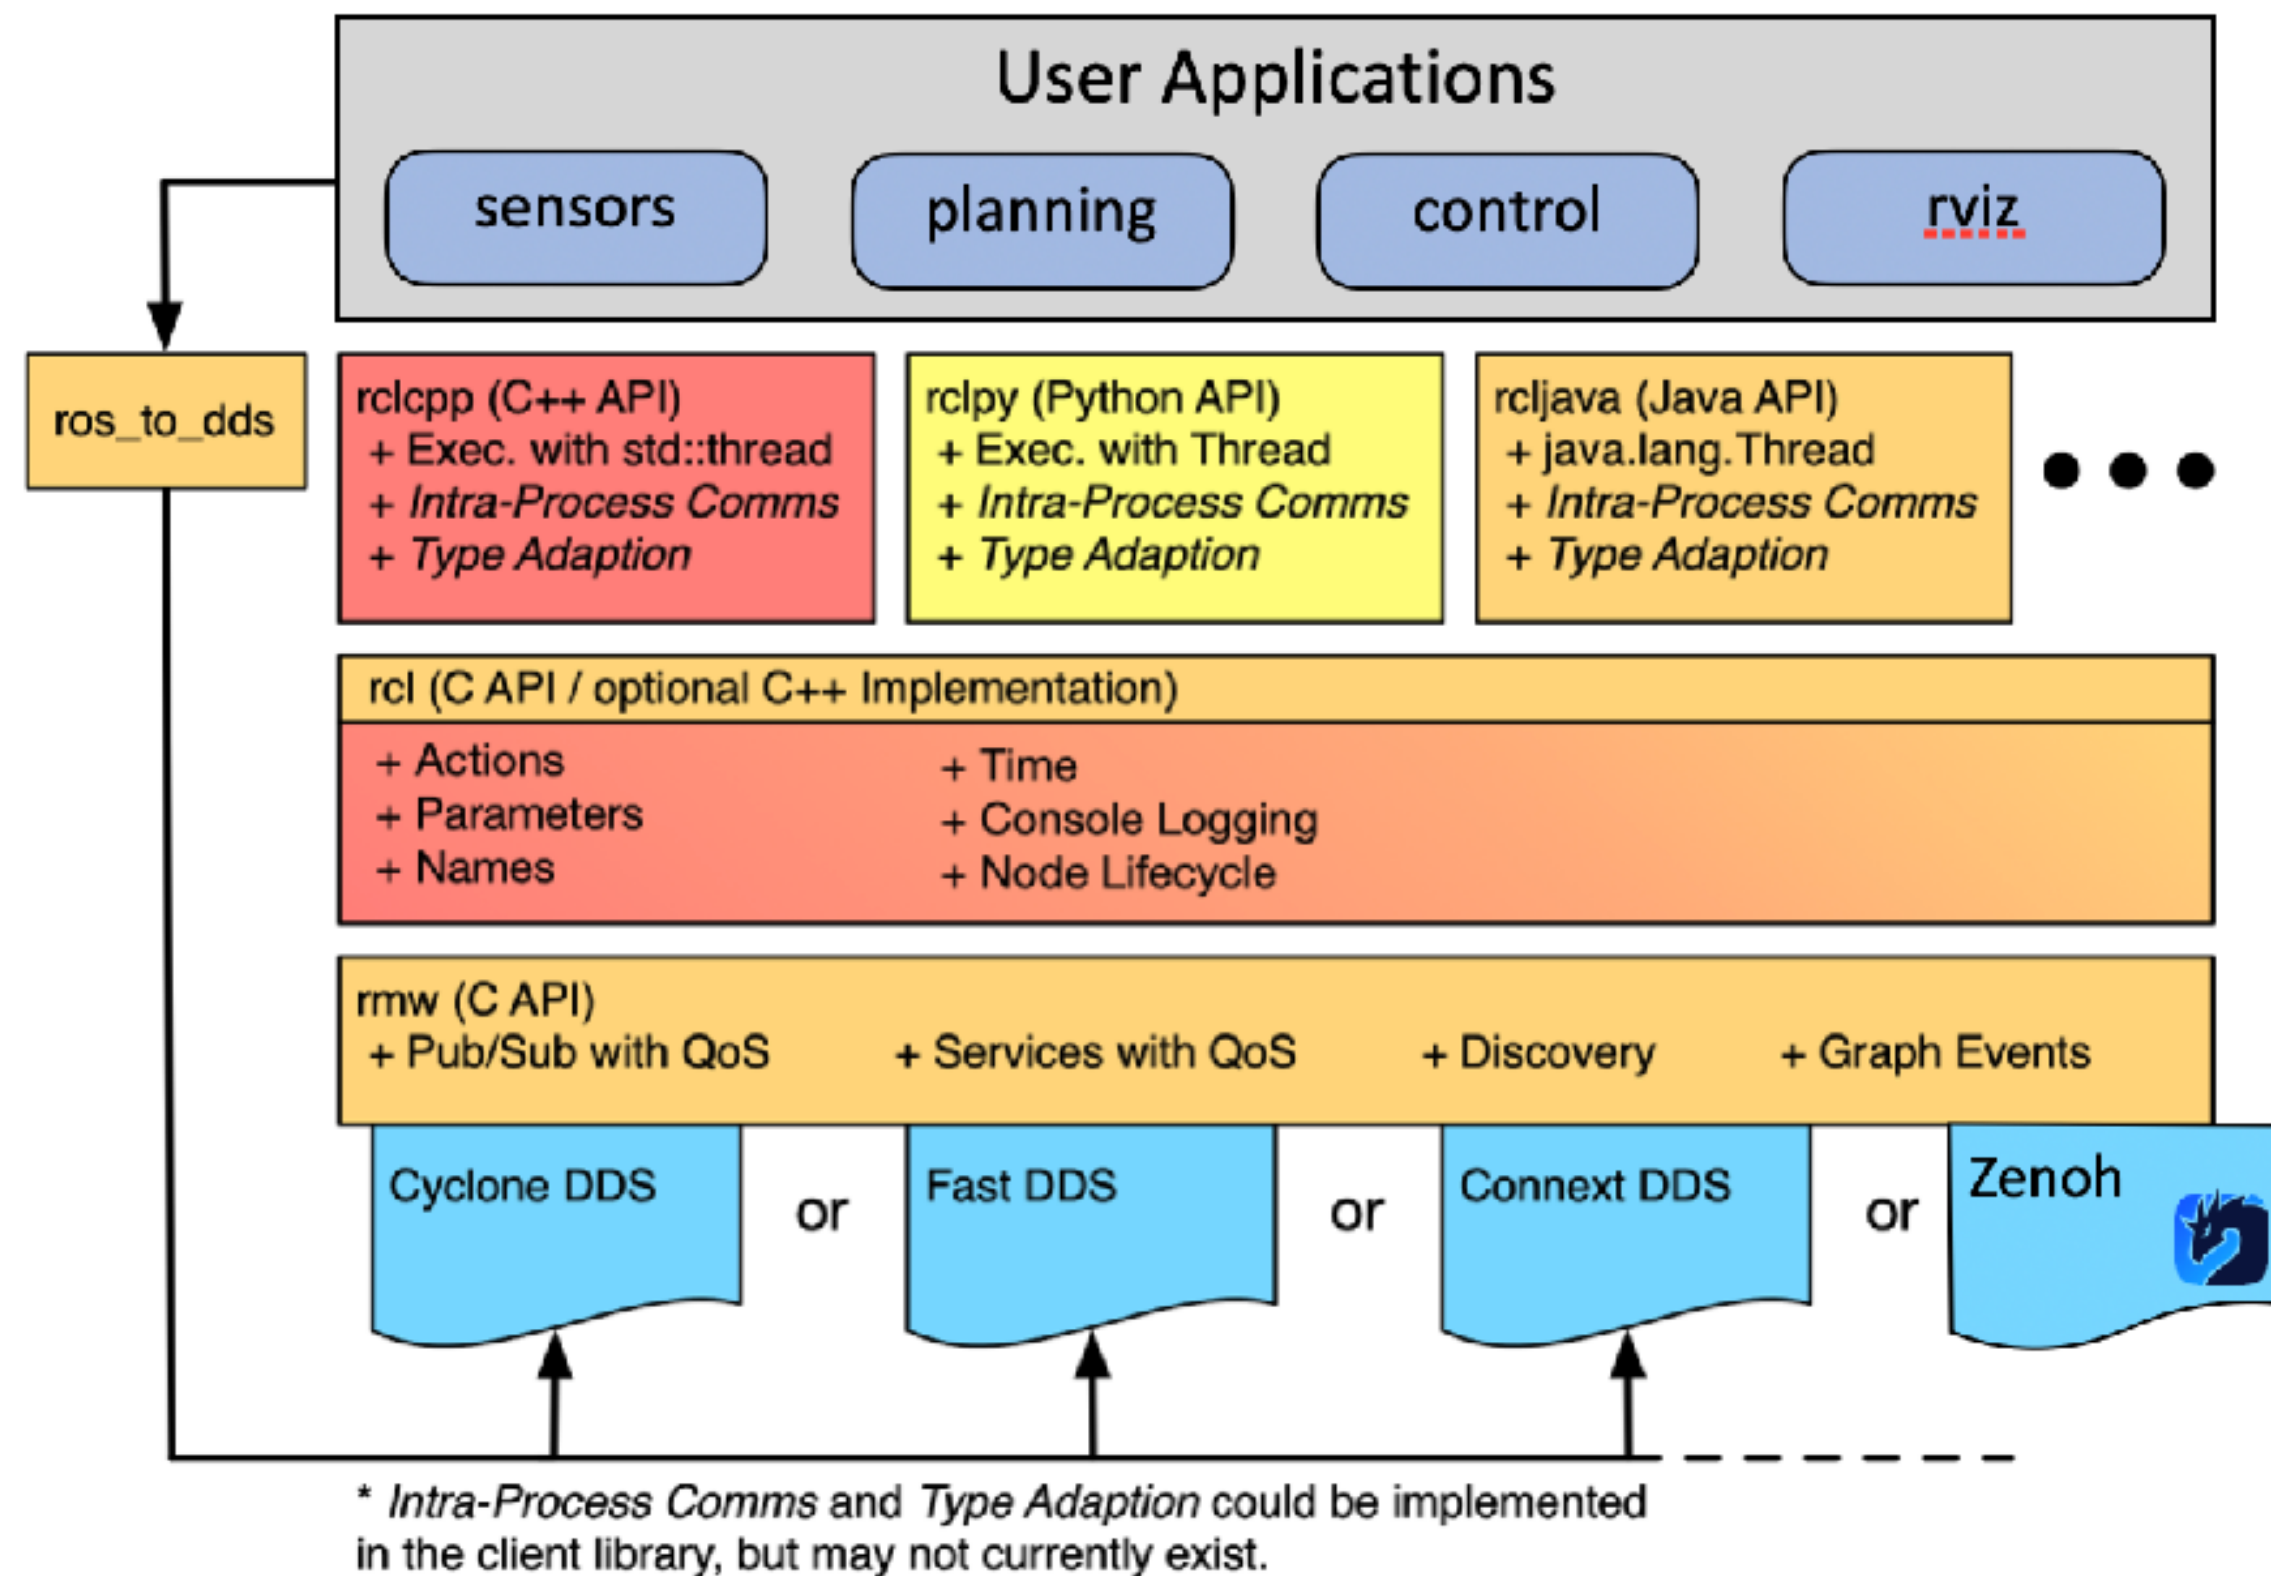
- 25us for MQTT
- 75us for Kafka
- 8us DDS




rmw_zenoh

An RMW implementation based on Zenoh

ROS 2 —Modular Architecture




https://github.com/ros2/rmw_zenoh



Zenoh Router

```
> ros2 run rmw_zenoh_cpp zenohd
```




- ↔ Endpoints Gossip discovery via the loopback
- Peer-to-peer communication via the loopback

Local Communication

- The router listen on tcp/0.0.0.0:7447
- Each Node connect to the router on tcp/127.0.0.1:7447
- The router acts as a broker for endpoints discovery
- Nodes establish peer-to-peer connections via 127.0.0.1
- Select the transport for local communication independently than that for R2X
- Leverage SHM Optimisation (coming up)

Zenoh Router

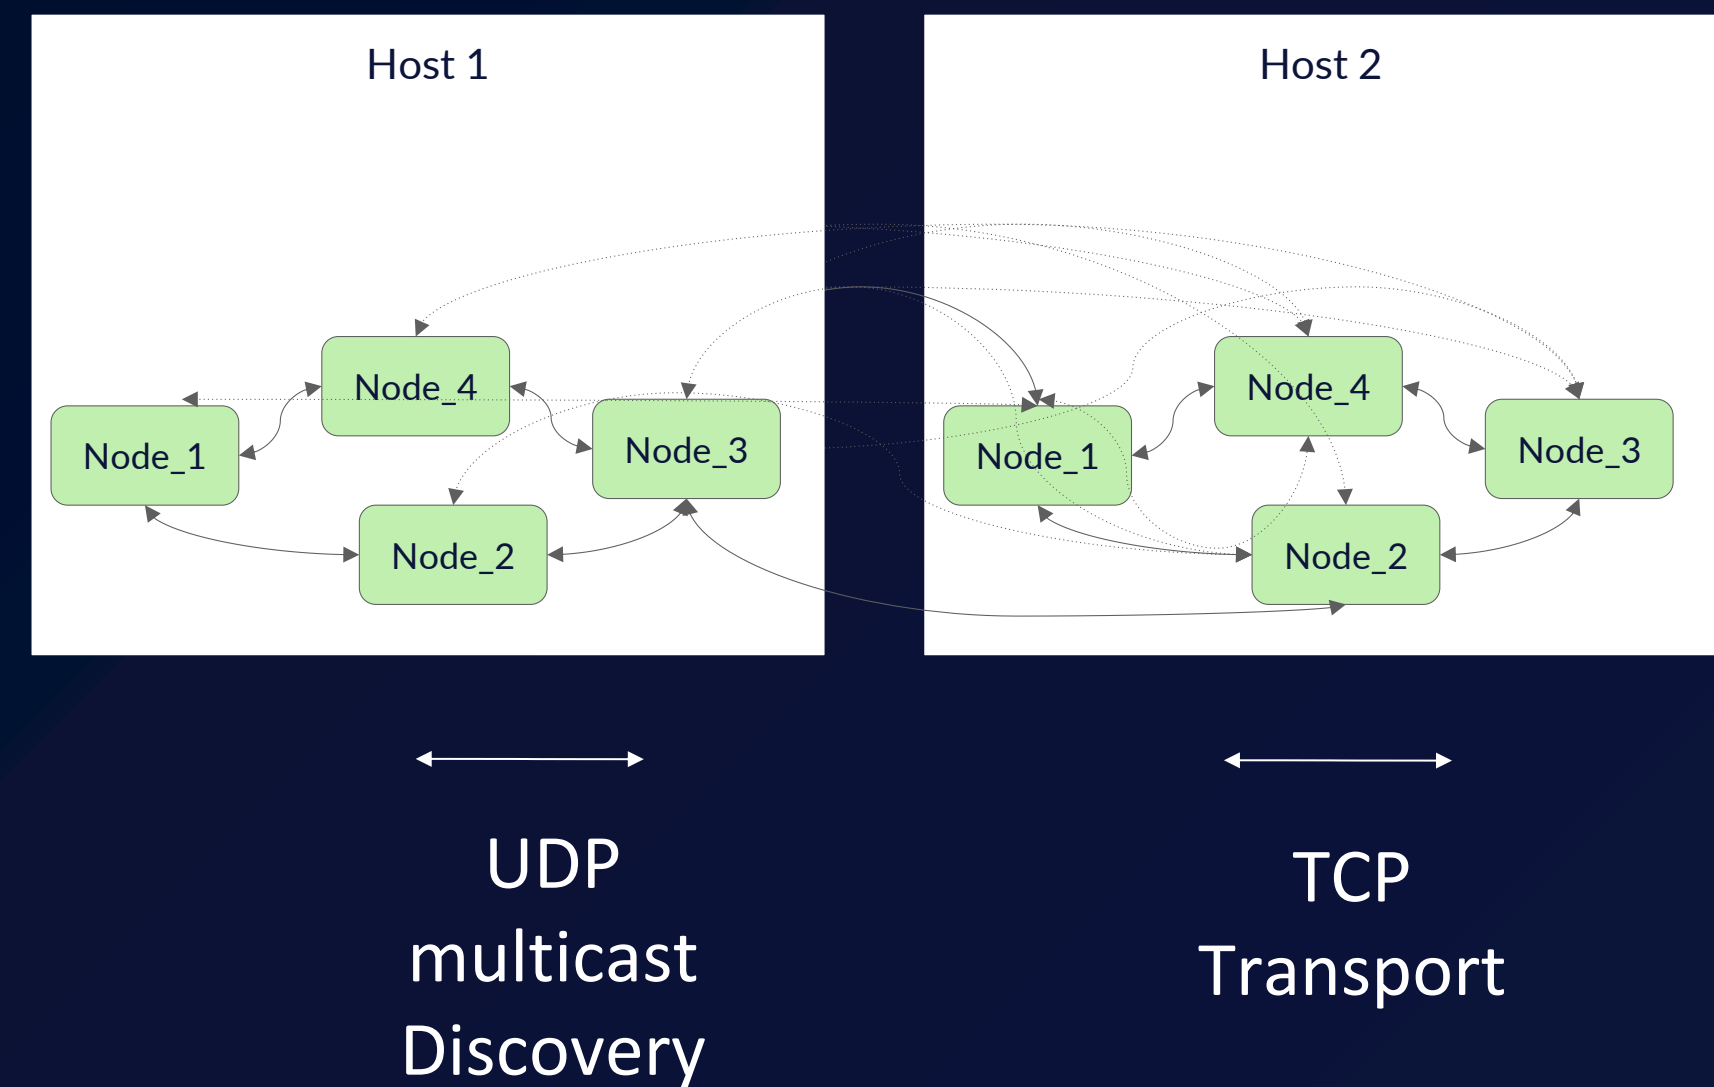



External Communications (R2X):

- May use different transport for R2X, i.e. TCP, TLS, QUIC, etc.
- Transparent Batching and compression
- Access Control and Downsampling
- Smaller surface of attack

Other configurable topologies

https://github.com/ros2/rmw_zenoh



Router for discovery,
but peer-to-peer communications

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
```


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


Transport reliably over many network hops

ROUTER_CONFIG.jsn5 - NeobotixTest - Visual Studio Code

Tiix: guest@jazzy: ~/ws_rmw_zench

2/2

+

+

+

+

+

+

File

Panels

Help

Move Camera

Select

Focus Camera

Measure

2D Point Estimate

Publish Point

Nav2 Goal

Displays

Global Options

Fixed Frame: map

Background Color: 40; 40; 40

Frame Rate: 30

Global Status: Ok

Grid: ☒

RobotModel: ☐

TF: ☒

LaserScan: ☒

Bumper Hit: ☒

Map: ☒

Amcl Particle Swarm: ☒

Add

Duplicate

Remove

Reset

Navigation 2

Navigation: active

Localization: active

Feedback: reached

ETA: 0 s

Distance remaining: 0.06 m

Time taken: 3 s

Recoveries: 0

Pause

Reset

Waypoint / Nav Through Poses Mode

Tools for WP-Following

Send WP

Load WP

Reset WP

Neobotix

video.m...

LIVE

PLAYING

192.168.50.202/rtsp/vid001

Julien's laptop

ROX

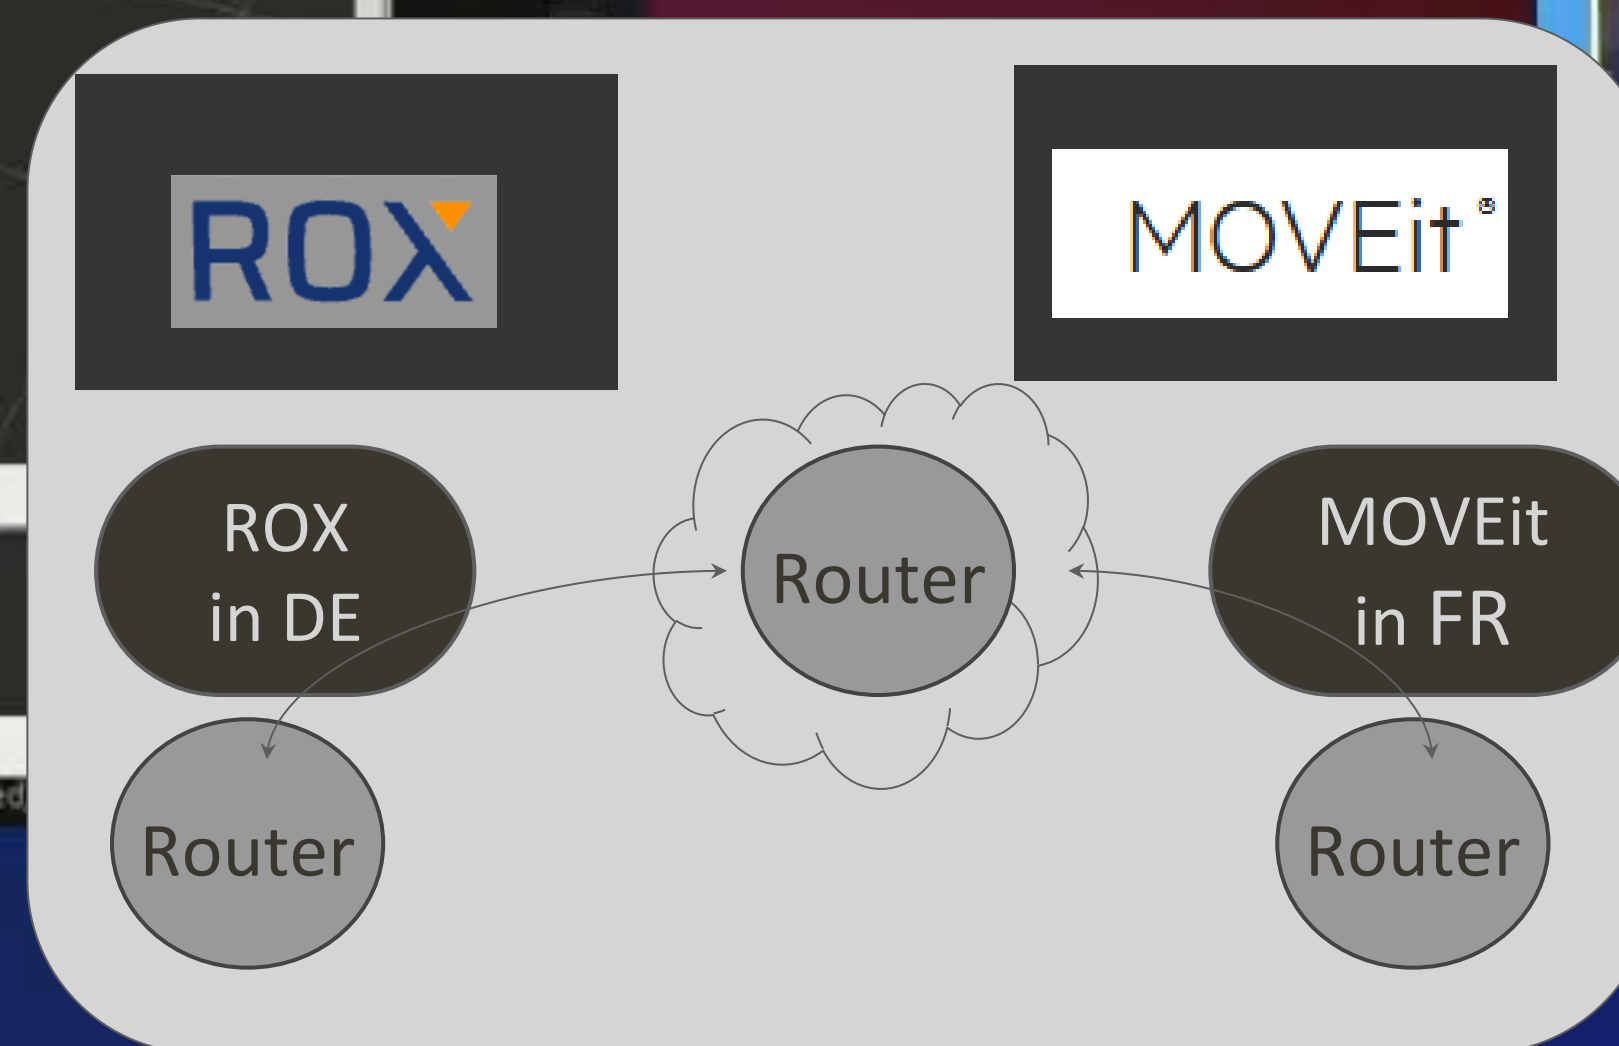
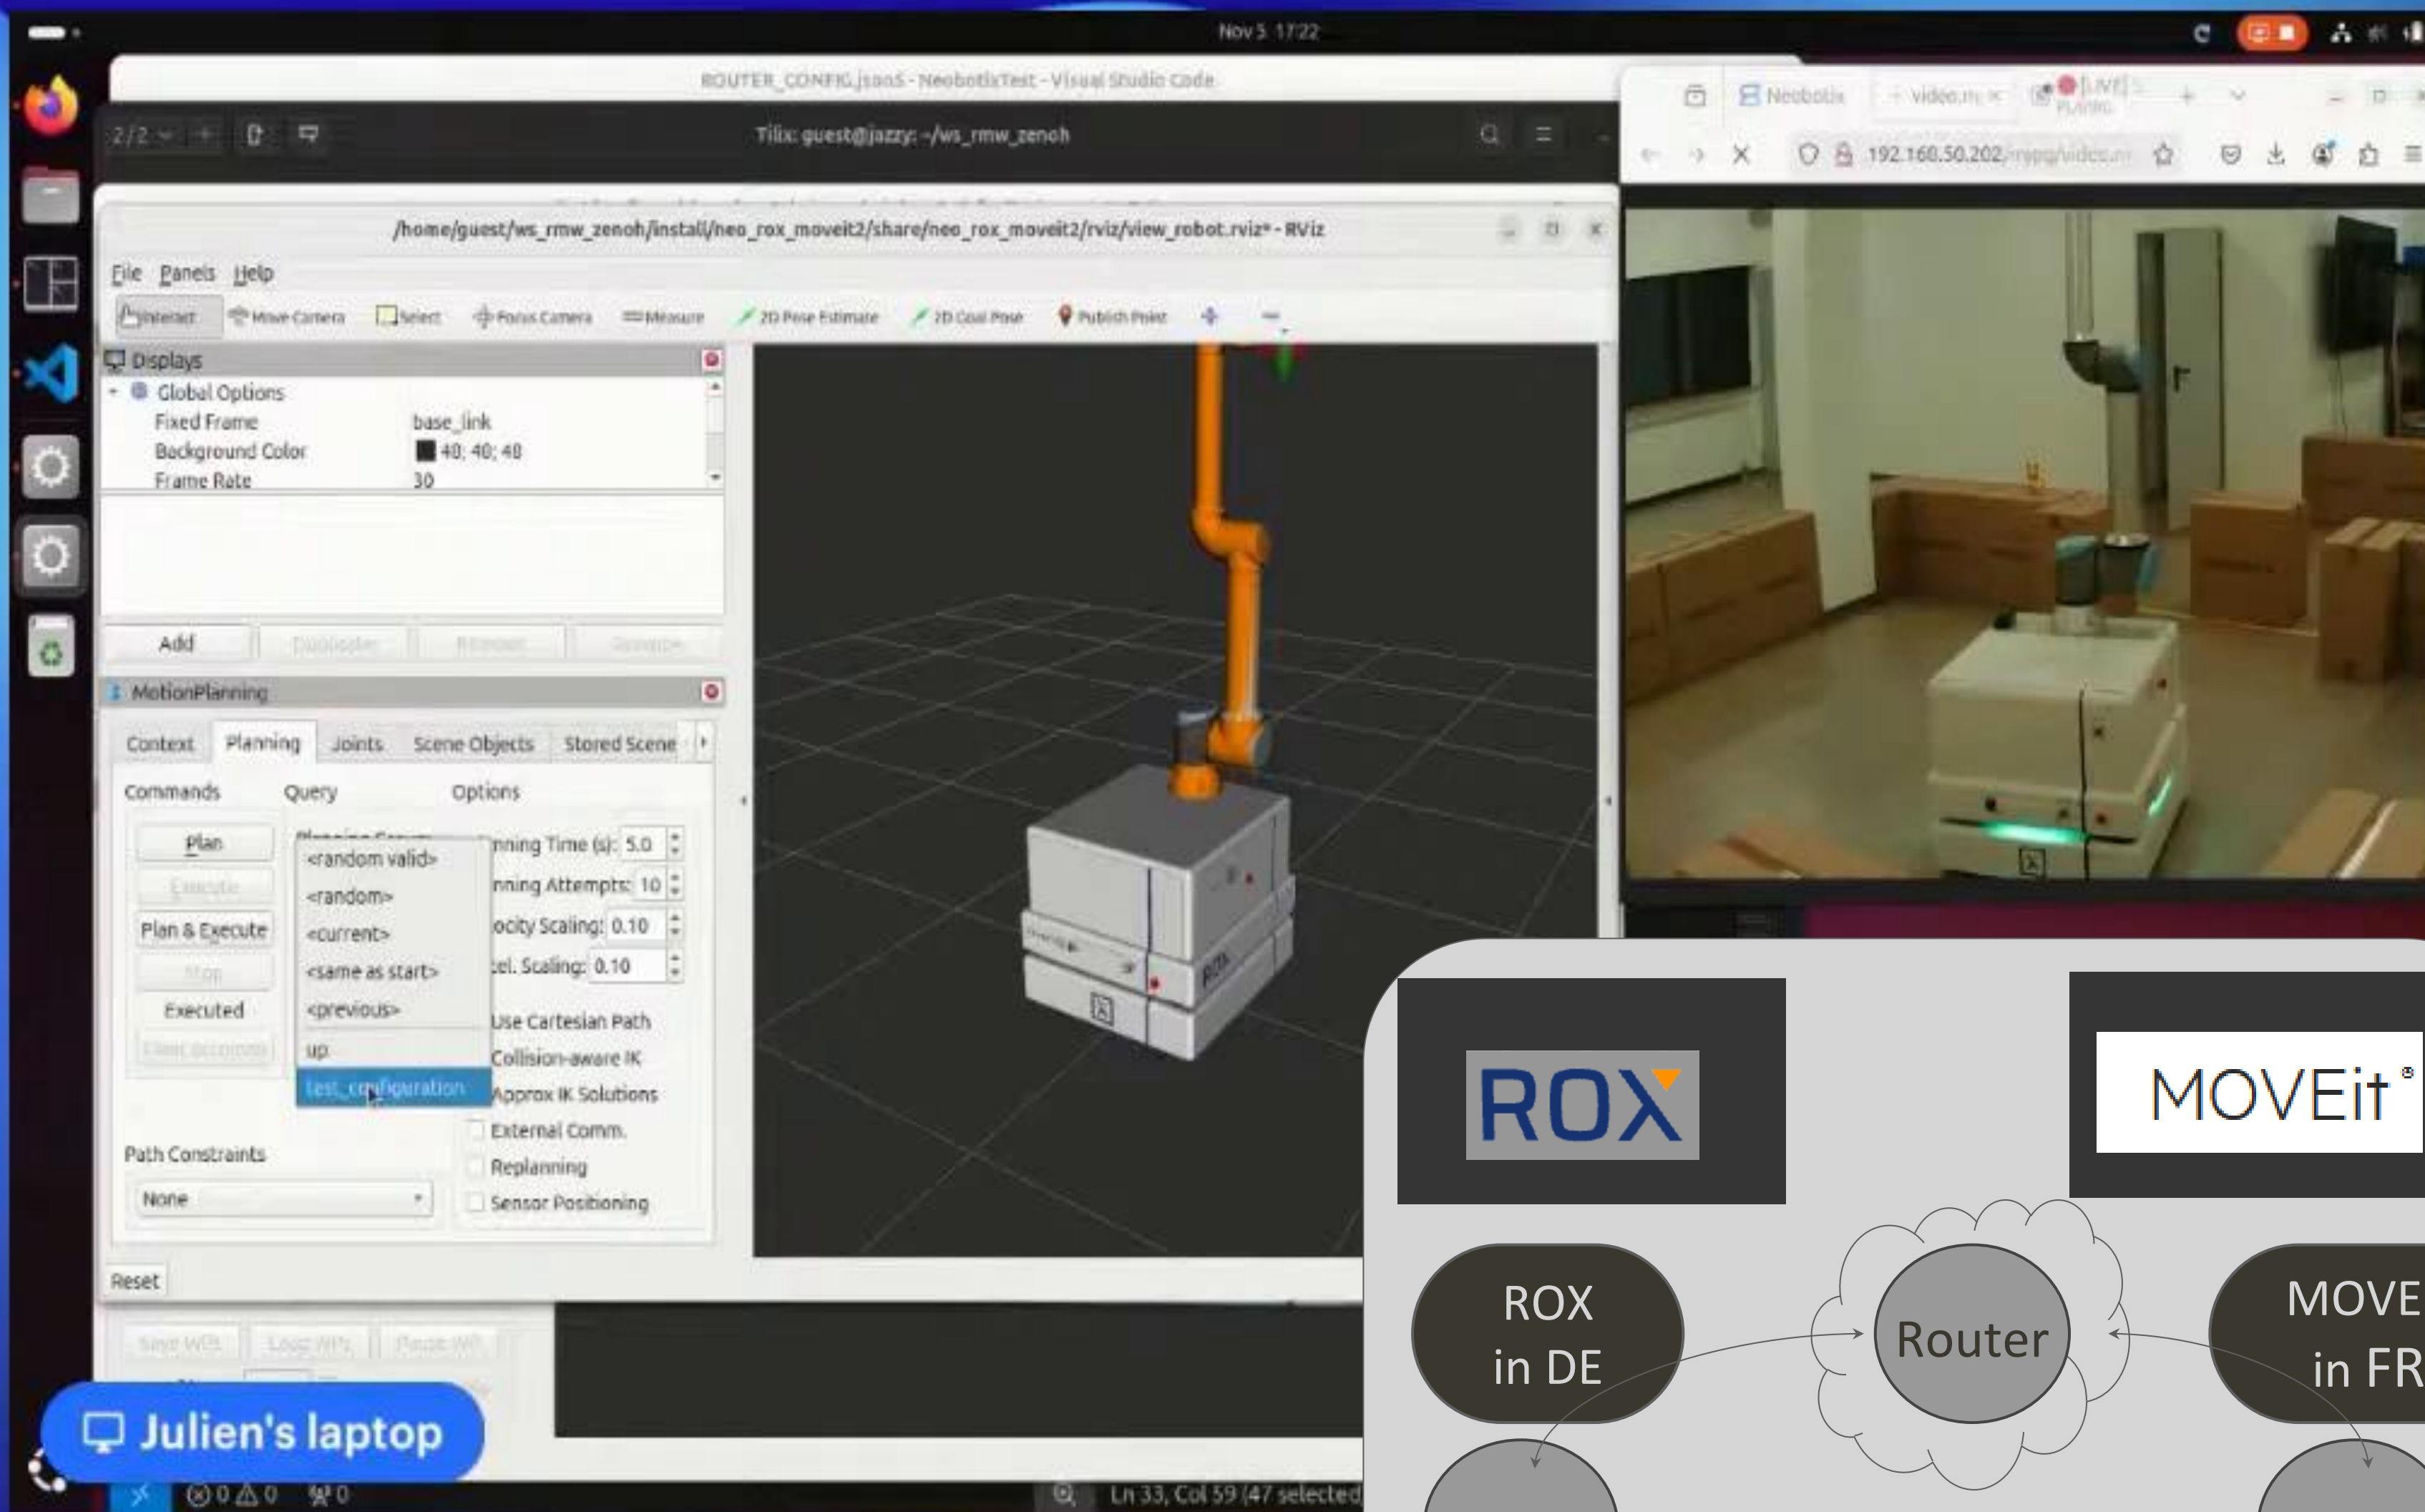
RViz

Router

Router

Router





Downsample when needed



```
2024-10-16T18:56:51.257913Z DEBUG acc-0 ThreadId(84) zenoh_link_
pted TCP connection on [::ffff:127.0.0.1]:7447: [::ffff:127.0.0.1]
2024-10-16T18:56:52.770492 DEBUG acc-0 ThreadId(84) zenoh_link_
pted TCP connection on [::ffff:192.200.40.12]:7447: [::ffff:192.200.40.12]
2024-10-16T18:56:53.262829Z DEBUG acc-0 ThreadId(84) zenoh_link_
pted TCP connection on [::ffff:127.0.0.1]:7447: [::ffff:127.0.0.1]
2024-10-16T18:56:59.285744Z DEBUG acc-0 ThreadId(84) zenoh_link_
pted TCP connection on [::ffff:127.0.0.1]:7447: [::ffff:127.0.0.1]
2024-10-16T18:57:03.267460Z DEBUG acc-0 ThreadId(84) zenoh_link_
pted TCP connection on [::ffff:127.0.0.1]:7447: [::ffff:127.0.0.1]
2024-10-16T18:57:07.269792Z DEBUG acc-0 ThreadId(84) zenoh_link_
pted TCP connection on [::ffff:127.0.0.1]:7447: [::ffff:127.0.0.1]
2024-10-16T18:57:11.275589Z DEBUG acc-0 ThreadId(84) zenoh_link_
pted TCP connection on [::ffff:127.0.0.1]:7447: [::ffff:127.0.0.1]
2024-10-16T18:57:15.279290Z DEBUG acc-0 ThreadId(84) zenoh_link_
pted TCP connection on [::ffff:127.0.0.1]:7447: [::ffff:127.0.0.1]
2024-10-16T18:57:19.282012Z DEBUG acc-0 ThreadId(84) zenoh_link_
pted TCP connection on [::ffff:127.0.0.1]:7447: [::ffff:127.0.0.1]
```

```
min: 0.032s max: 0.036s std dev: 0.00151s window: 20
average rate: 30.136
min: 0.027s max: 0.035s std dev: 0.00250s window: 20
average rate: 30.020
min: 0.028s max: 0.044s std dev: 0.00445s window: 20
average rate: 30.140
min: 0.029s max: 0.038s std dev: 0.00227s window: 20
average rate: 29.885
min: 0.017s max: 0.059s std dev: 0.00924s window: 20
average rate: 29.945
min: 0.030s max: 0.037s std dev: 0.00100s window: 20
```


Known limitations



- `rc1cpp::shutdown()` must explicitly be called before program termination.
- Router must be manually started (for now).
- Liveliness and deadline QoS events not supported.

Thank You

Patience, persistence and perspiration
make an unbeatable combination for
success.

