

KIARA Demonstration

Context

FI-PPP

European programme for
Internet-enabled innovation



FI-WARE

Technology Foundation of
FI-PPP



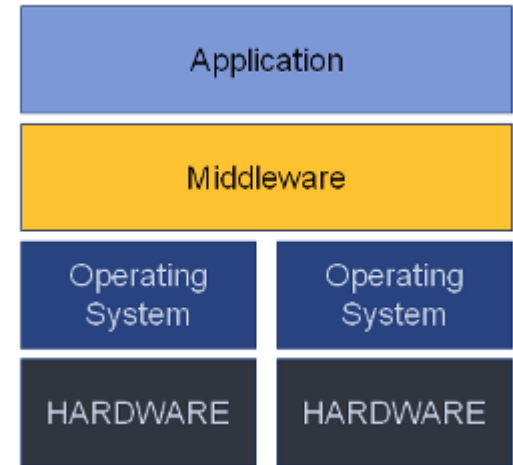
KIARA

Advanced Middleware



KIARA - An advanced Middleware

- Message oriented communication middleware
- Simplifies communication between distributed heterogeneous systems
 - Systems with different operating systems
 - Applications using different programming languages and paradigms
- Abstracts network-layer with common API



KIARA - Advanced Features

- Data-structures used in an application can be dynamically mapped to IDL definitions at run-time
- No extra code generation (skeleton/stub) required to use middleware
- Embedded compiler (LLVM) generates highly optimized code
- Negotiation of optimal communication mechanisms, protocols, and data representations to be used between two peers

KIARA - Advanced Features

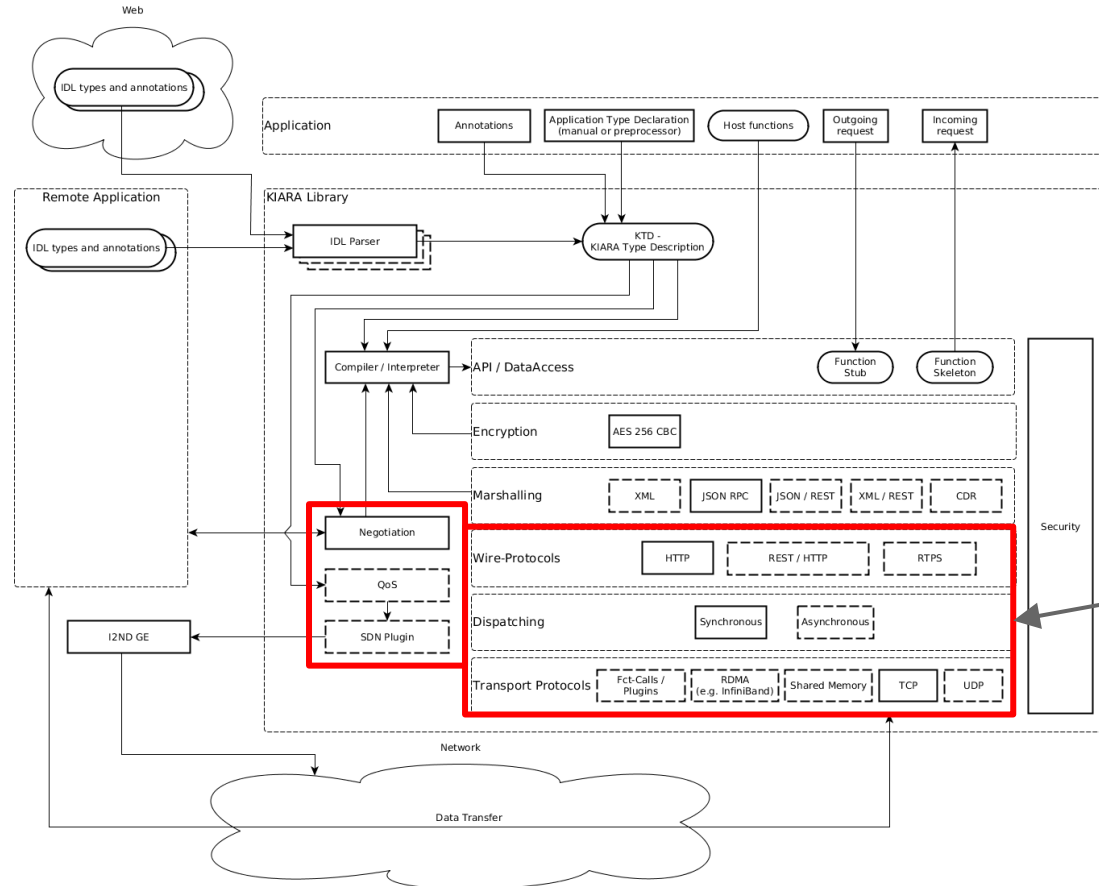
- Offers multiple communication paradigms like **Request/Reply** or **Publish/Subscribe**.
- *Secure by Design* approach Applications can declare their security needs in the form of security policies (security rules) and apply them to data structures and service at development time or even later during deployment definitions

KIARA - First Benchmarks

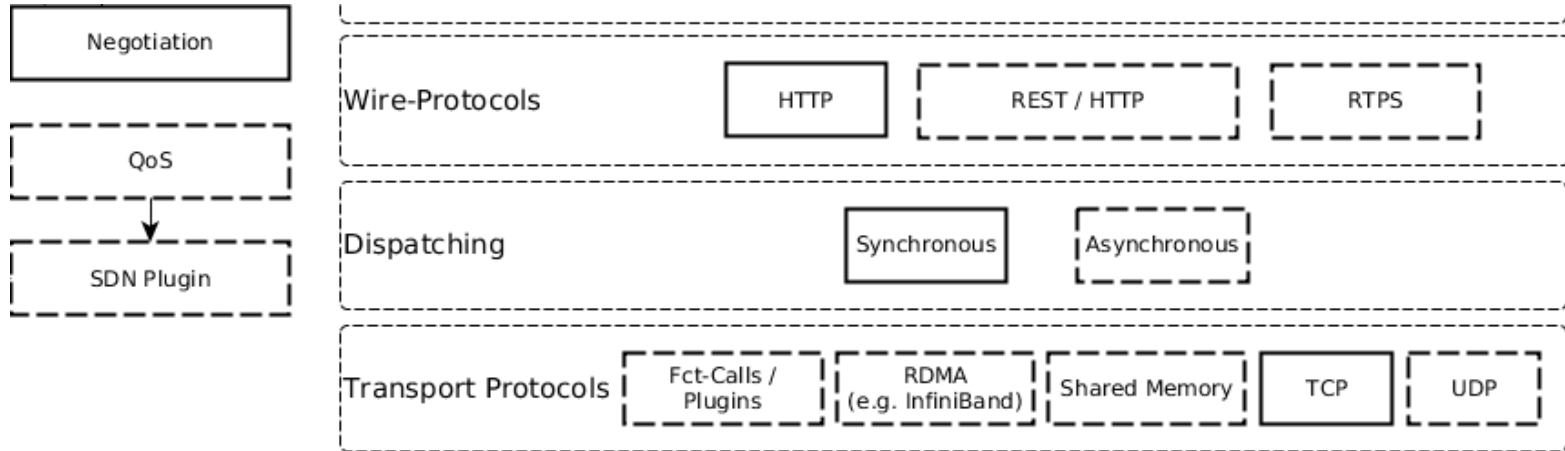
Localhost	Network transfer and serialization							
	Kiara ortecd		BoostTyped		Thrift 0.9.1		Ice 3.5.1	
	AvgLat	SSD	AvgLat	SSD	AvgLat	SSD	AvgLat	SSD
hurricane/L	10.74	0.08	21.40	1.31	14.84	0.47	16.77	1.59
hurricane/W	11.55	0.69	86.35	3.41	16.50	0.51	18.10	1.12
uragan/W	10.10	0.45	77.75	8.24	14.30	1.08	18.55	3.61
fpga/L	18.20	0.19	37.11	0.38	30.36	0.46	27.75	1.96
ghost/L	15.46	1.77	32.96	0.15	25.71	1.29	27.14	0.52
currywurst/L	21.43	0.18	39.40	0.21	32.81	0.37	33.74	1.22
bratwurst/L	13.29	0.16	24.11	0.21	19.60	0.31	30.43	0.59
Client → server								
ghost/L → fpga/L	178.23	6.73	190.73	5.10	195.67	10.50	182.67	8.79
hurricane/W → uragan/W	139.55	6.53	266.80	9.85	116.90	1.65	173.95	9.08
uragan/W → hurricane/W	130.40	9.05	254.25	12.36	127.50	6.38	169.15	8.81
currywurst/L → bratwurst/L	120.86	1.33	139.78	2.12	107.17	1.73	153.88	0.55
bratwurst/L → currywurst/L	120.35	0.45	142.49	0.99	108.16	2.09	152.12	1.23

Message-Size: 470 Byte
 Messages per Test: 10'000
 Measured Tests: 20
 Unit of Time: micro-seconds
 AvgLat: average
 latency
 SSD: sample standard deviation

KIARA - Architecture



KIARA - Architecture Detailed View



KIARA - Responsibilities of ICCLab

- Offer transport capabilities to upper layers
- Offer functionality to negotiate transport, QoS and security parameters

KIARA - Transport Stack

Offers SCALN - a **S**ystem **C**all **A**bstraction **L**ayer for **N**etworking

- `bind / unbind`
- `connect / disconnect`
- `send / receive`
- `register_callback / callback`
- `get_context / set_context`
- `get_session / set_session`
- `get_configuration / set_configuration`

Offers C-Bindings of SCALN

Support of additional languages planned

Transport Stack - Underlying Techn.

ZeroMQ: High-performance asynchronous messaging library aimed at use in scalable distributed or concurrent applications.

Why:

- very lightweight library
- implements several communication patterns that can be leveraged
- has very well defined clean documentation and API (reusable messaging stack)

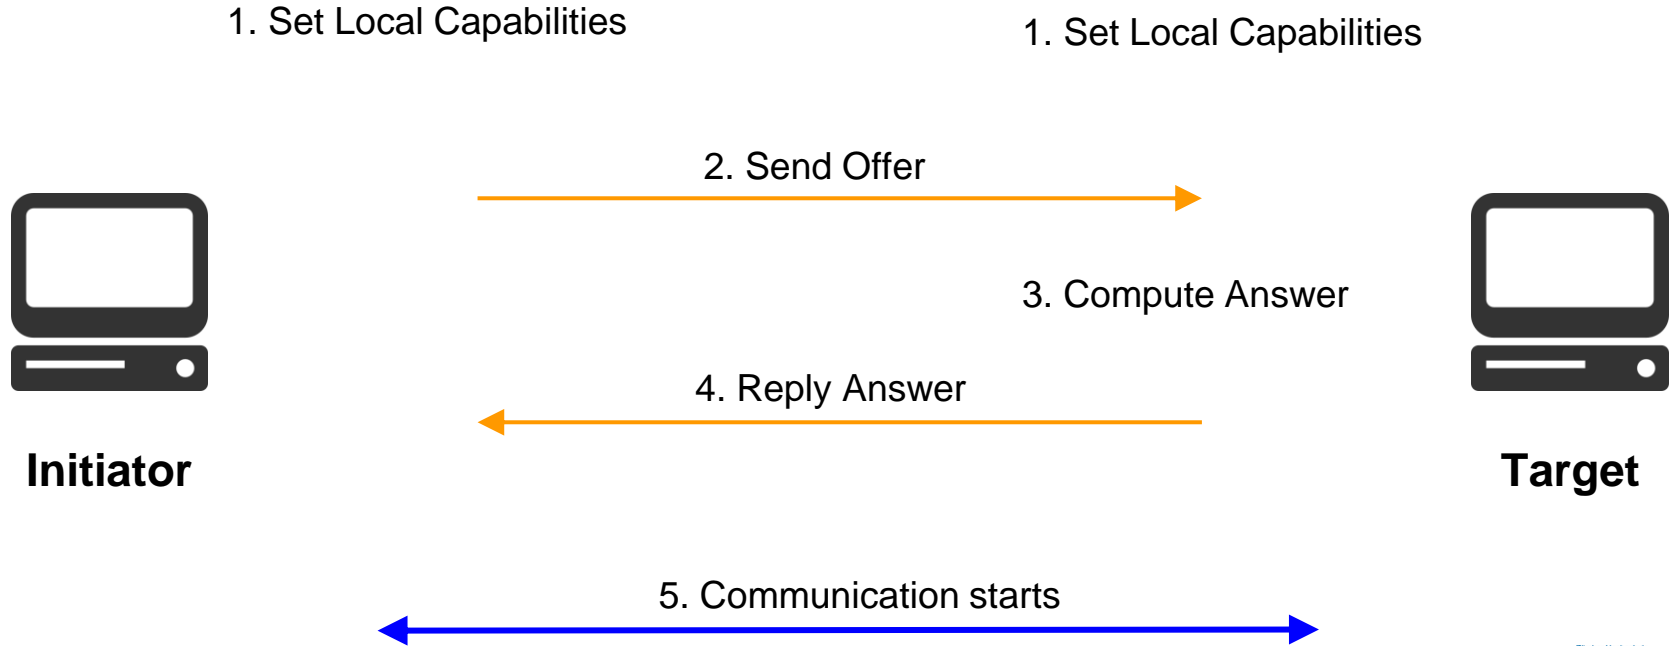
Transport Stack - Underlying Techn.

InfiniBand: Switched fabric computer network communications technology used in high-performance computing and enterprise data centers.

Why: - For high throughput / low latency use cases
- Offer RDMA (Remote Direct Memory

Access)

KIARA - Negotiation



KIARA - QoS

- If the network has a SDN controller with the KIARA SDN application, we can assure/offer:
 - Bandwidth - certain amount of bandwidth with rate-limiters
 - Path - policy based privileged network paths
 - RTT - using monitoring to ensure max Round Trip Time
 - TOD - configure the communication according to the type of the devices along the path
- Transport with timestamp information (deadline) by using the Real Time Publish Subscribe (RTPS) protocol

KIARA Demo - Code Examples

Use different application types

```
KT_Configuration conf;  
conf.set_application_type( KT_STREAM );  
  
KT_WEBSERVER);  
  
KT_PUBLISHSUBSCRIBE);  
KT_REQUESTREPLY);
```

Use different communication technologies

```
KT_Connection* conn = new KT_Zeromq();  
  
KT_InfiniBand();  
  
KT_Boost();
```

KIARA Transport- Client Code Example

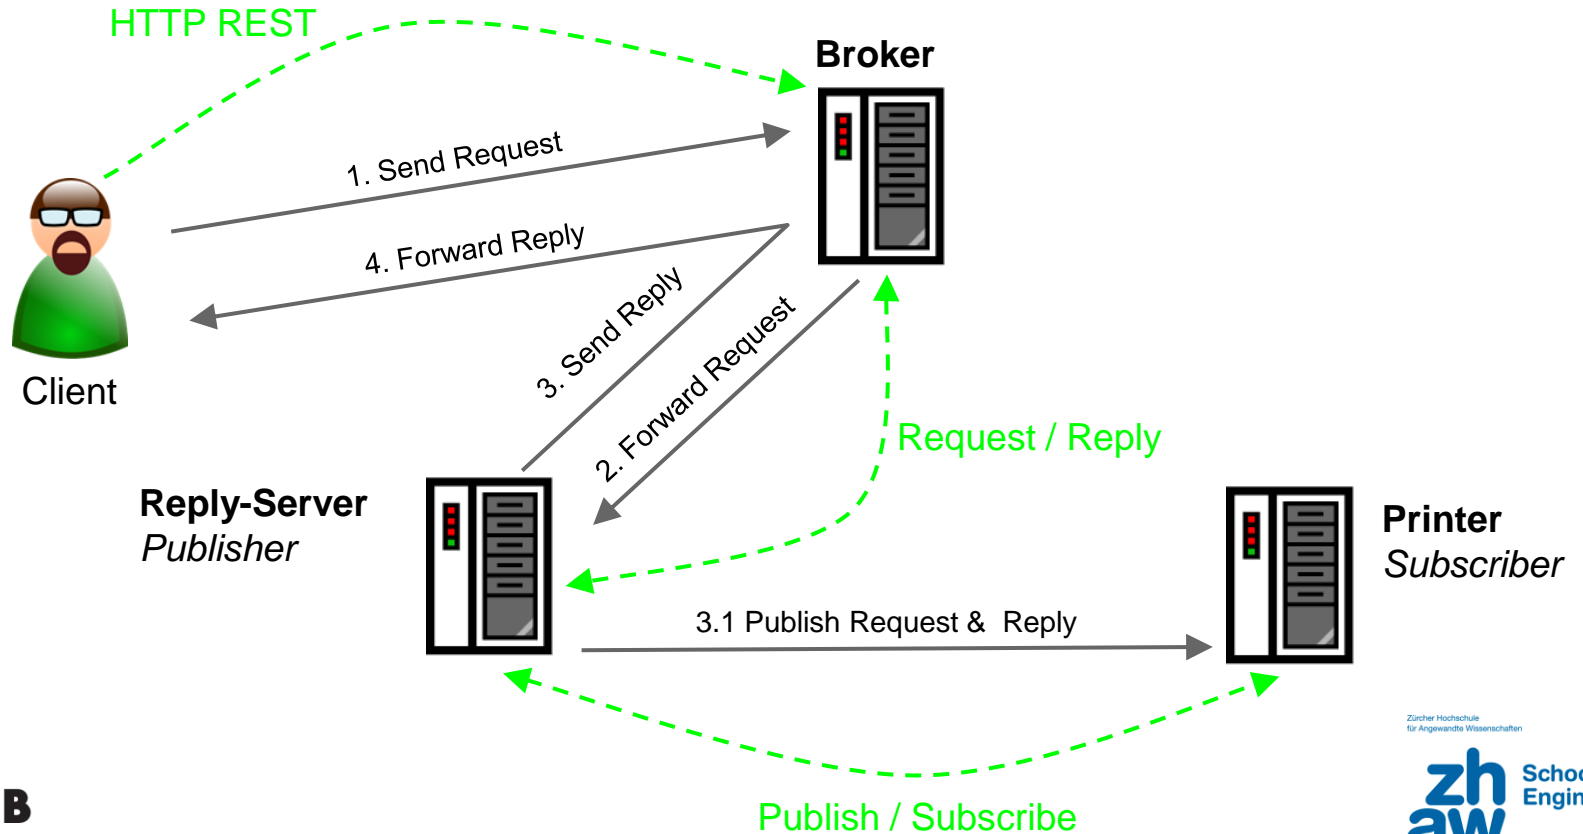
1. Create and configure connection

```
KT_Configuration config;  
config.set_application_type ( KT_REQUESTREPLY / KT_PUBLISHSUBSCRIBE );  
  
config.set_host( KT_TCP, "192.168.100.1", 5555 );  
  
KT_Connection *connection = new KT_Zeromq() / new KT_InfiniBand() ;  
connection->set_configuration(config);
```

2. Use connection to send and receive messages

```
connection->connect(...);  
connection->send(...);  
connection->recv(...);  
connection->disconnect(...);
```


KIARA - Demo Setup



DEMO

Send PUT Request to <http://160.85.4.249:8080>

```
$ curl http://160.85.4.249:8080 -X PUT -d "some text"
```

We're Done.



Questions?

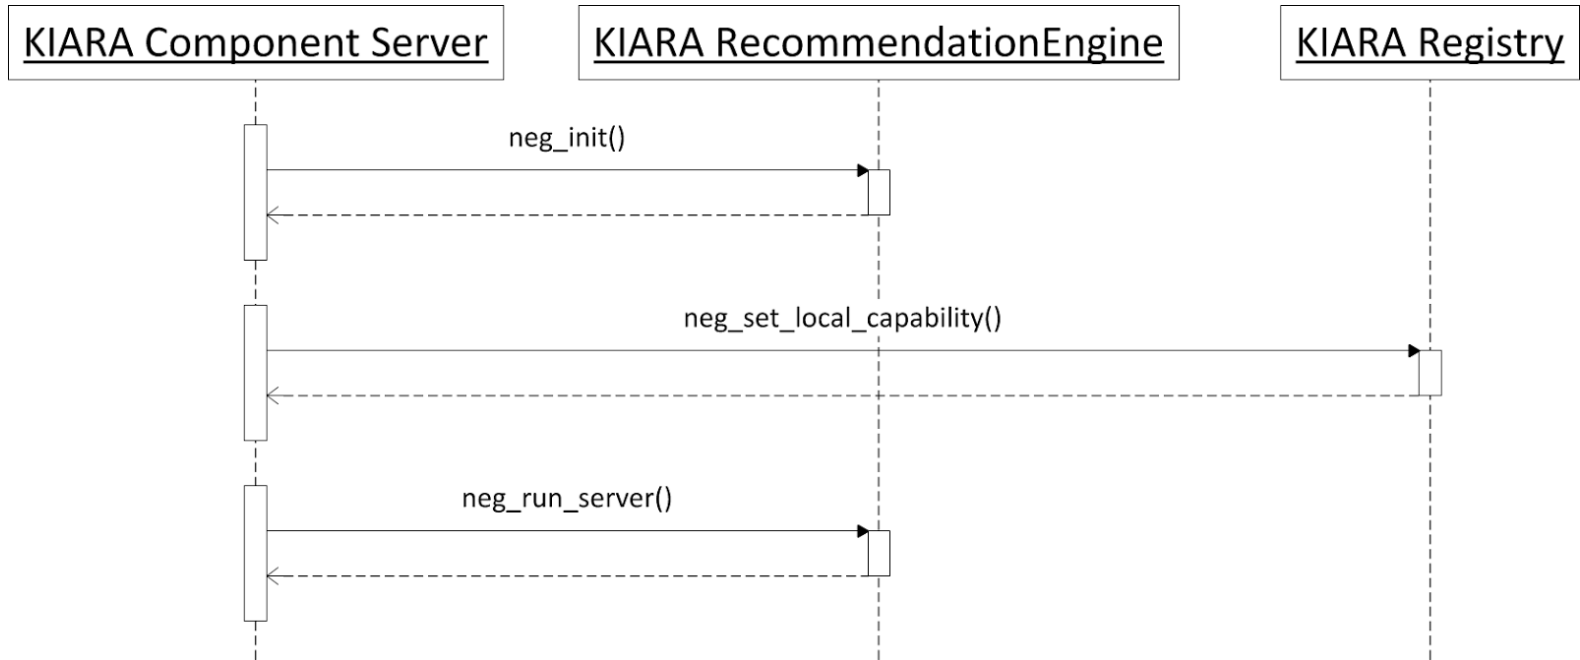
quikkmeme.com



Backup Slides

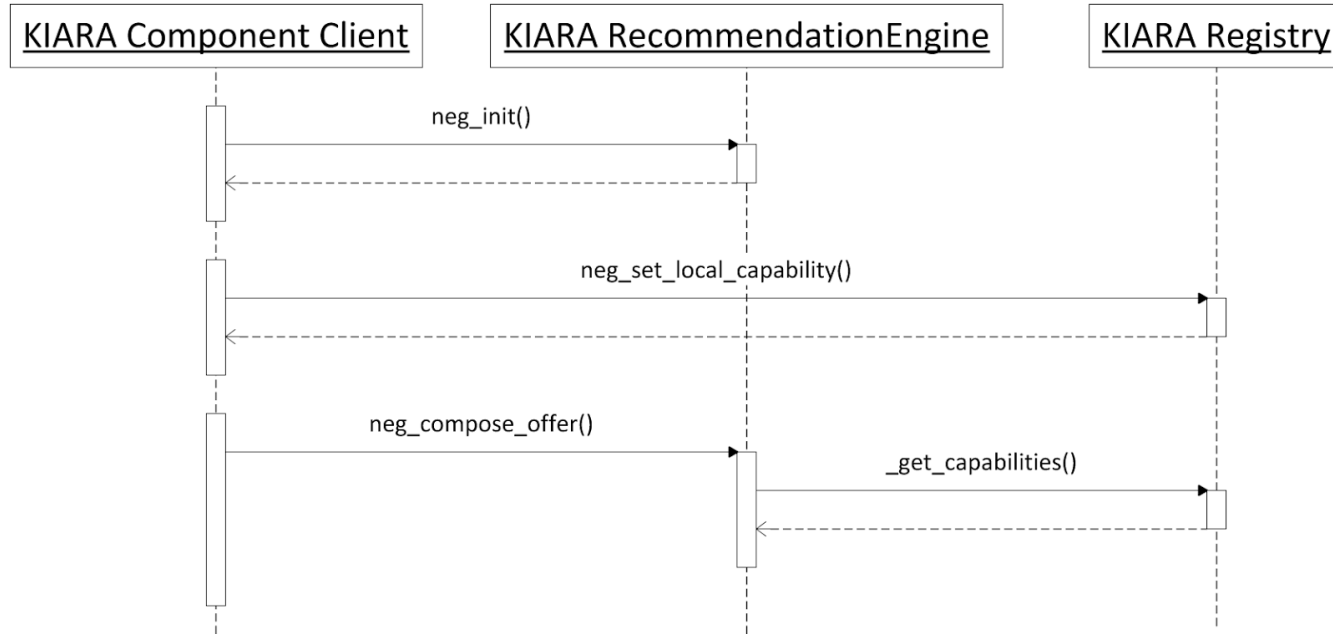
KIARA Negotiation Sequence

1: - Server sets local capabilities



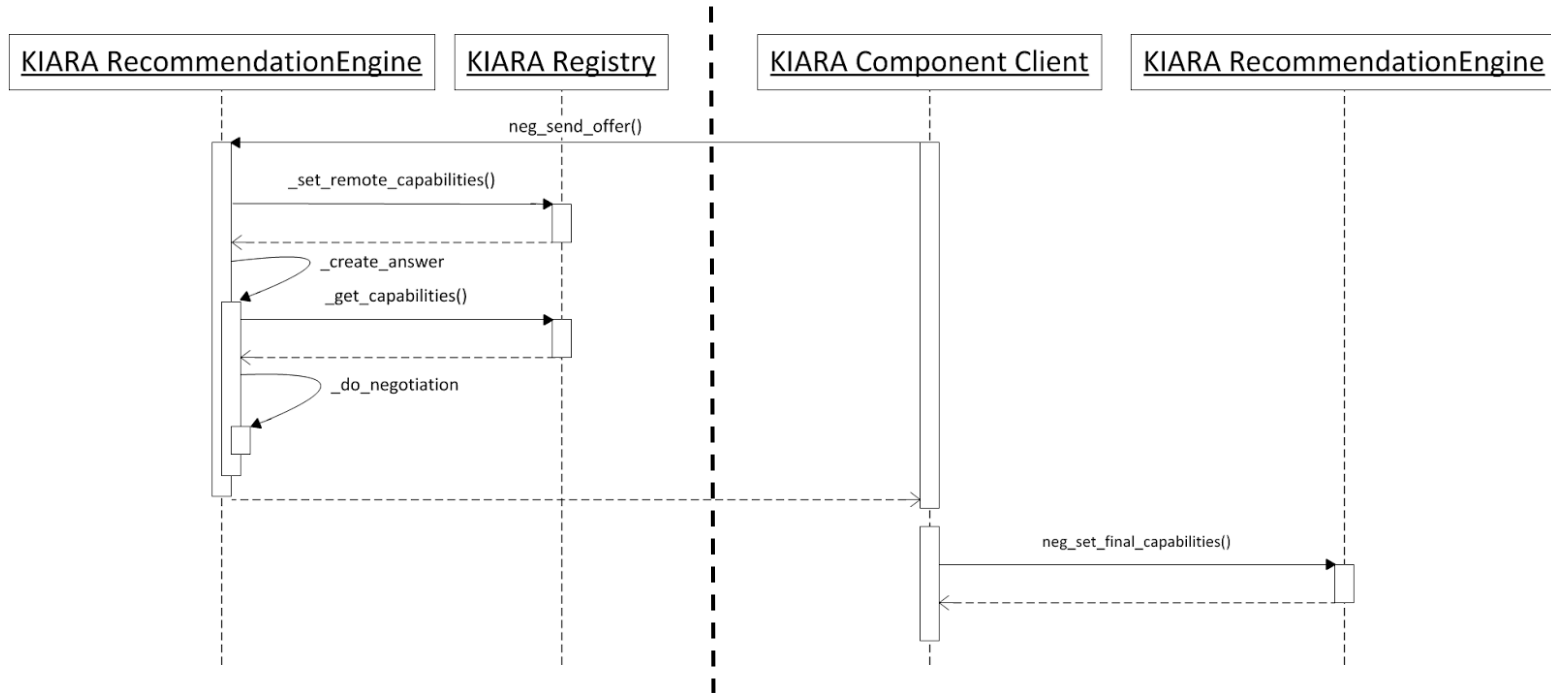
KIARA Negotiation Sequence

- 2:
- Client sets local capabilities
 - Client composes offer

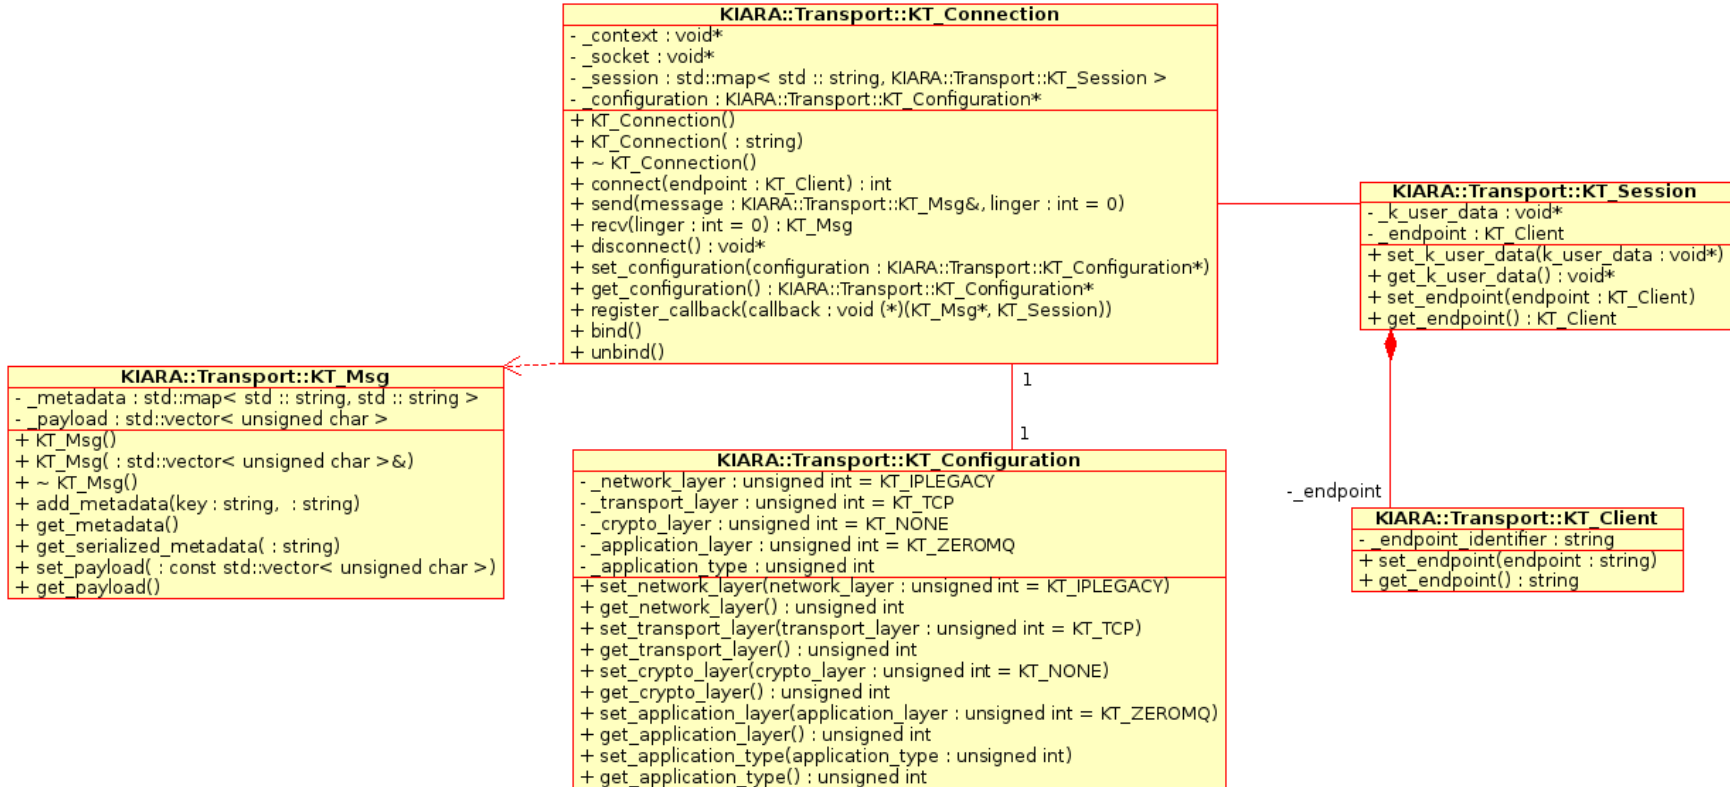


KIARA Negotiation Sequence

- 3: - Client and Server negotiate transport and QoS settings



Internal Transport architecture



KIARA Transport - Code Beispiele

Because nobody wants to write all this for just a webserver

```
kt_configuration_t* config = kt_configuration_create();
kt_configuration_set_network_layer( config, KT_IPLEGACY );
kt_configuration_set_application_layer( config, KT_ZEROMQ );
kt_configuration_set_transport_layer( config, KT_TCP );
kt_configuration_set_application_type( config, KT_REQUESTREPLY );
kt_configuration_set_hostname( config, "localhost" );
kt_configuration_set_port( config, 5555 );
```

KIARA - Negotiation

Goal: Negotiate application and communication settings

1. Both, Initiator (client) and target (server) set their local capabilities
2. Initiator composes an offer out of his local capabilities and sends an offer to the target
3. The target compares the remote capabilities with the local ones and sends back an answer
4. The initiator sets the negotiated answer as final and starts the communication

Optional step 2: The initiator fetches the capabilities of the target