

Remote Procedure Call as a Managed System Service

Jingrong Chen*, Yongji Wu*, Shihan Lin, Yechen Xu, Xinhao Kong,
Thomas Anderson, Matthew Lentz, Xiaowei Yang, Danyang Zhuo


*Equal contributions



Remote Procedure Calls Widely Used



Distributed Data Store



Network Filesystem



Data Analytics Framework



Cluster Orchestrator



Consensus Protocol



TensorFlow



Deep Learning System

Remote Procedure Calls Widely Used

In Google's datacenter, RPCs

- generate **>95%** of application traffic[1]
- spend **~10%** of its CPU cycles[2]

- [1] Aequitas: Admission Control for Performance-Critical RPCs in Datacenters, SIGCOMM '22
- [2] Profiling a Warehouse-Scale Computer, ISCA '15

Remote Procedure Calls Widely Used

In Google's datacenter, RPCs

- generate **>95%** of application traffic[1]
- spend **~10%** of its CPU cycles[2]


Performance is always a key design goal of RPC

- [1] Aequitas: Admission Control for Performance-Critical RPCs in Datacenters, SIGCOMM '22
- [2] Profiling a Warehouse-Scale Computer, ISCA '15

Remote Procedure Calls: Implementation

Remote Procedure Calls: Implementation


- 1 Write protocol specification



```
Service KVStore
  Message GetReq
    bytes key
  Message Entry
    bytes? value
  Func Get(GetReq) -> Entry
```

Remote Procedure Calls: Implementation

- 1 Write protocol specification




```
Service KVStore
  Message GetReq
    bytes key
  Message Entry
    bytes? value
  Func Get(GetReq) -> Entry
```



Remote Procedure Calls: Implementation


1 Write protocol specification

 `Service KVStore`
`Message GetReq`
`bytes key`
`Message Entry`
`bytes? value`
`Func Get(GetReq) -> Entry`

2 Protocol compiler generates stub code




3 App compiles with the stub and RPC library



Remote Procedure Calls: Implementation

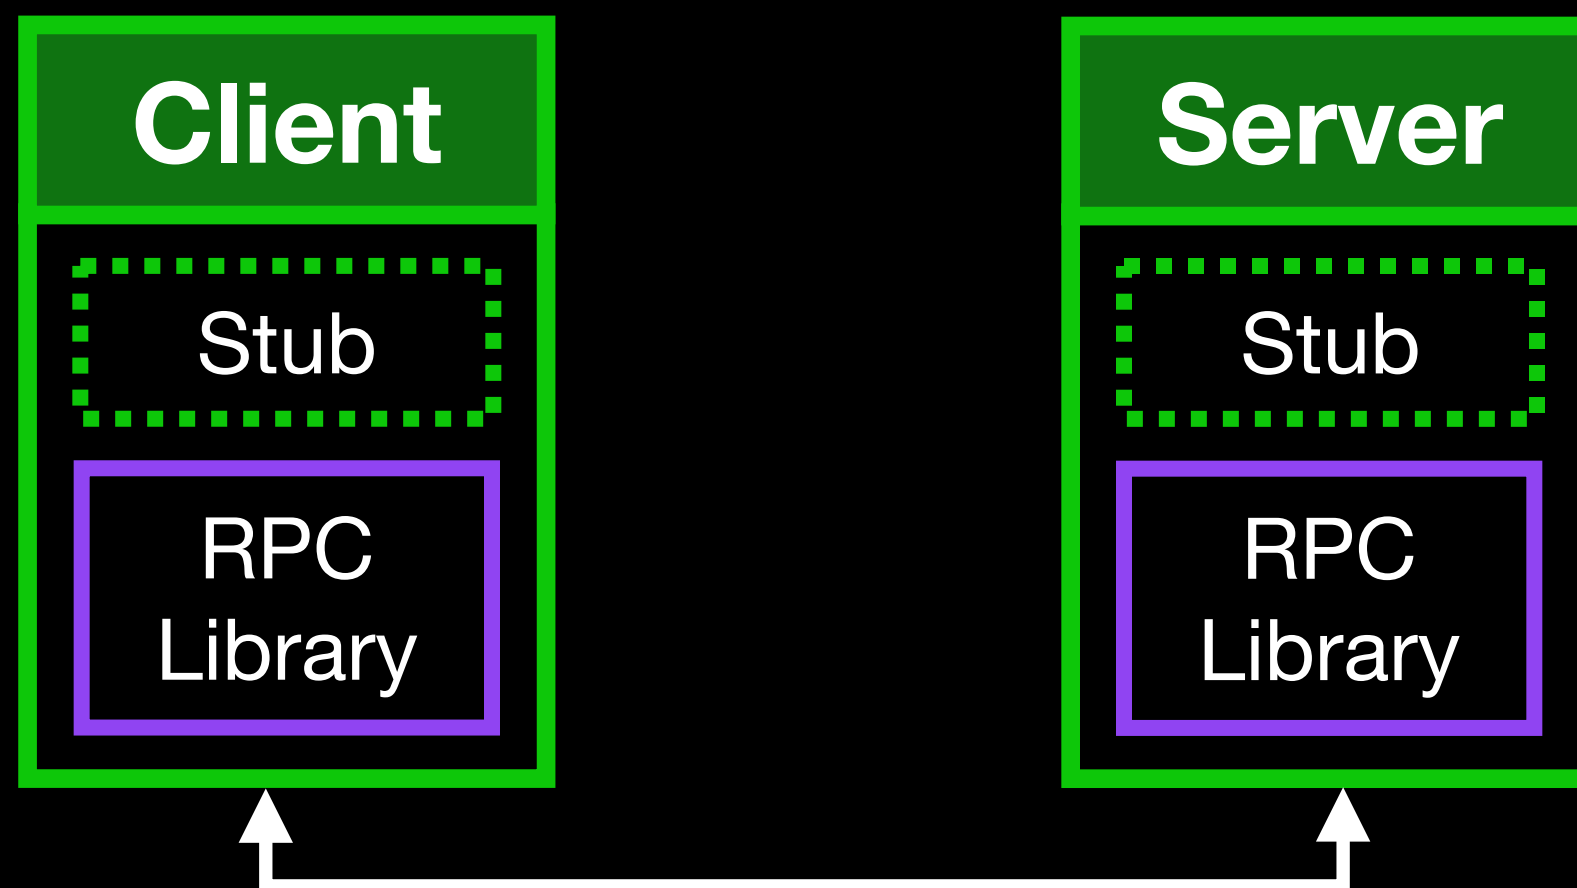
1 Write protocol specification

 `Service KVStore`
`Message GetReq`
`bytes key`
`Message Entry`
`bytes? value`
`Func Get(GetReq) -> Entry`

2 Protocol compiler generates stub code



3 App compiles with the stub and RPC library



Andrew D. Birrel and Bruce Jay Nelson,
Implementing Remote Procedure Calls,
ACM Transactions on Computer Systems
2(1):39-59, February **1984**

Remote Procedure Calls: Implementation


1 Write protocol specification

```
Service KVStore
Message GetReq
  bytes key
Message Entry
  bytes? value
Func Get(GetReq) -> Entry
```

2 Protocol compiler generates stub code



3 App compiles with the stub and RPC library




Andrew D. Birrel and Bruce Jay Nelson,
Implementing Remote Procedure Calls,
ACM Transactions on Computer Systems
2(1):39-59, February 1984

RPC-as-a-library
gRPC, Thrift, eRPC
Cap'n Proto, rpclib, XML-RPC
brpc, tarpc, tonic...

The Need for Manageability


The Need for Manageability



Observability

e.g., How many RPCs? RPC Latency?

The Need for Manageability



Observability


e.g., How many RPCs? RPC Latency?



Policy Enforcement

e.g., Prioritize certain RPCs?

The Need for Manageability



Observability

e.g., How many RPCs? RPC Latency?



Policy Enforcement


e.g., Prioritize certain RPCs?



Upgradability

e.g., Fix vulnerabilities while app running?

The Need for Manageability




Observability

e.g., How many RPCs? RPC Latency?




Policy Enforcement

e.g., Prioritize certain RPCs?




Upgradability


e.g., Fix vulnerabilities while app running?



envoy




LINKERD



Istio

Manageability in a feature-rich RPC library?



Observability

e.g., How many RPCs? RPC Latency?



Policy Enforcement


e.g., Prioritize certain RPCs?



Upgradability

e.g., Fix vulnerabilities while app running?

Manageability in a feature-rich RPC library?




Observability

e.g., How many RPCs? RPC Latency? -----> *YES*



Policy Enforcement


e.g., Prioritize certain RPCs?



Upgradability

e.g., Fix vulnerabilities while app running?

Manageability in a feature-rich RPC library?




Observability

e.g., How many RPCs? RPC Latency? -----> *YES*



Policy Enforcement


e.g., Prioritize certain RPCs? -----> *NO* Mandatory policies?



Upgradability

e.g., Fix vulnerabilities while app running?

Manageability in a feature-rich RPC library?




Observability

e.g., How many RPCs? RPC Latency? -----> *YES*



Policy Enforcement


e.g., Prioritize certain RPCs? -----> *NO* Mandatory policies?



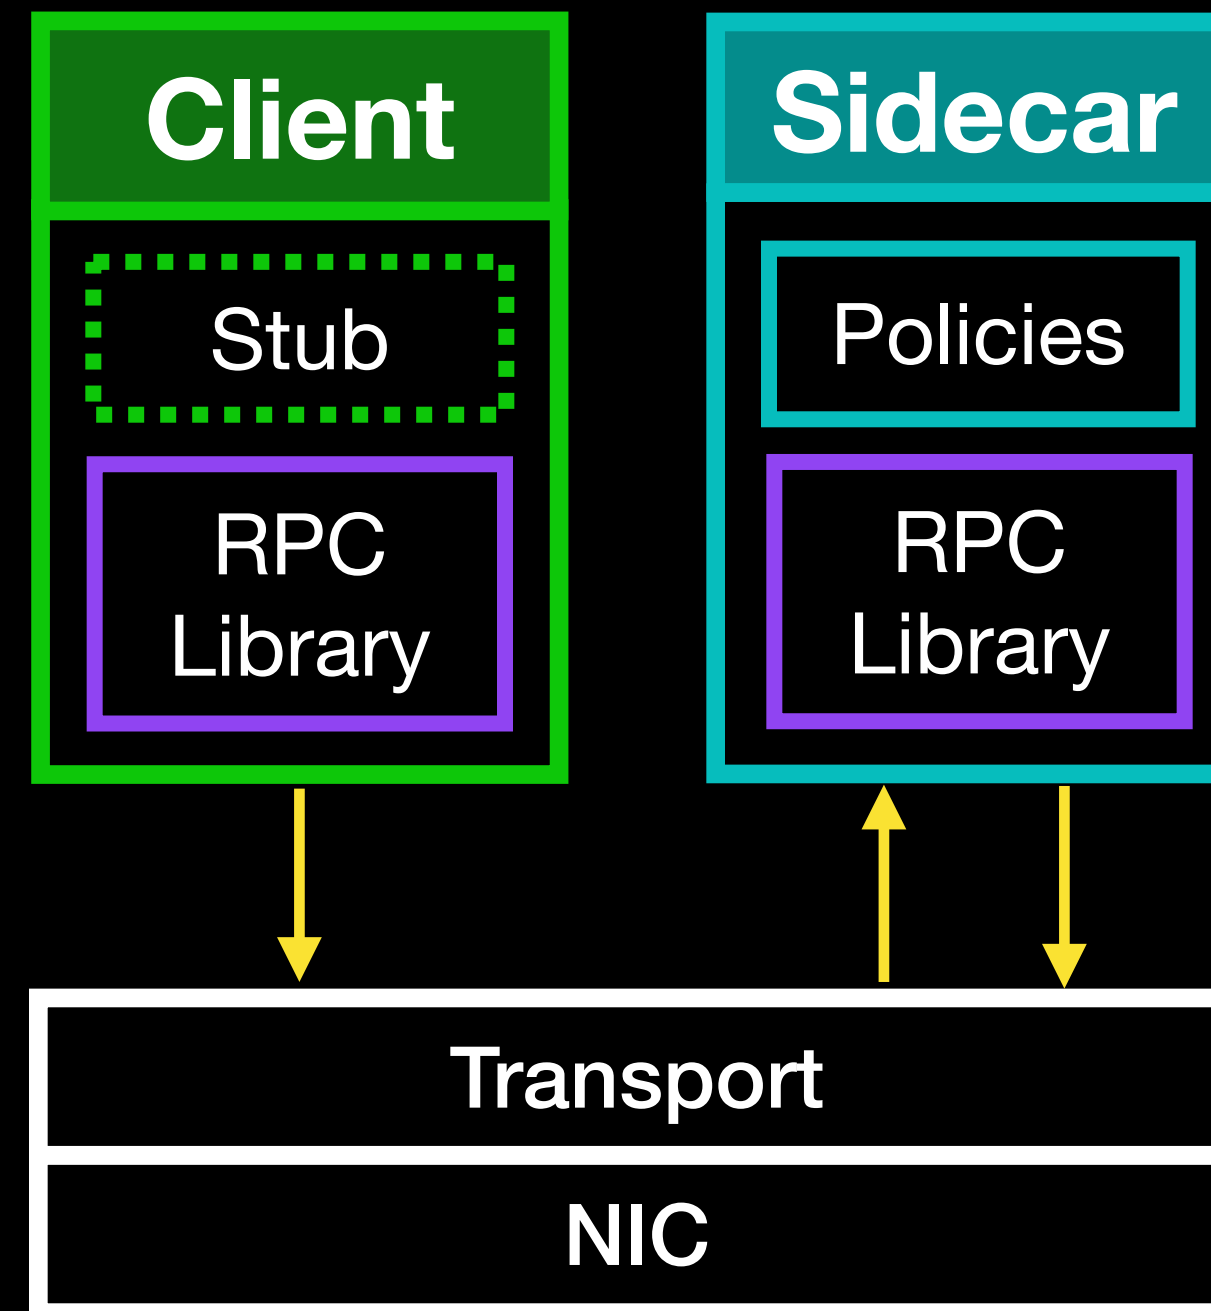
Upgradability

e.g., Fix vulnerabilities while app running? -----> *Currently NO* Upgradability?

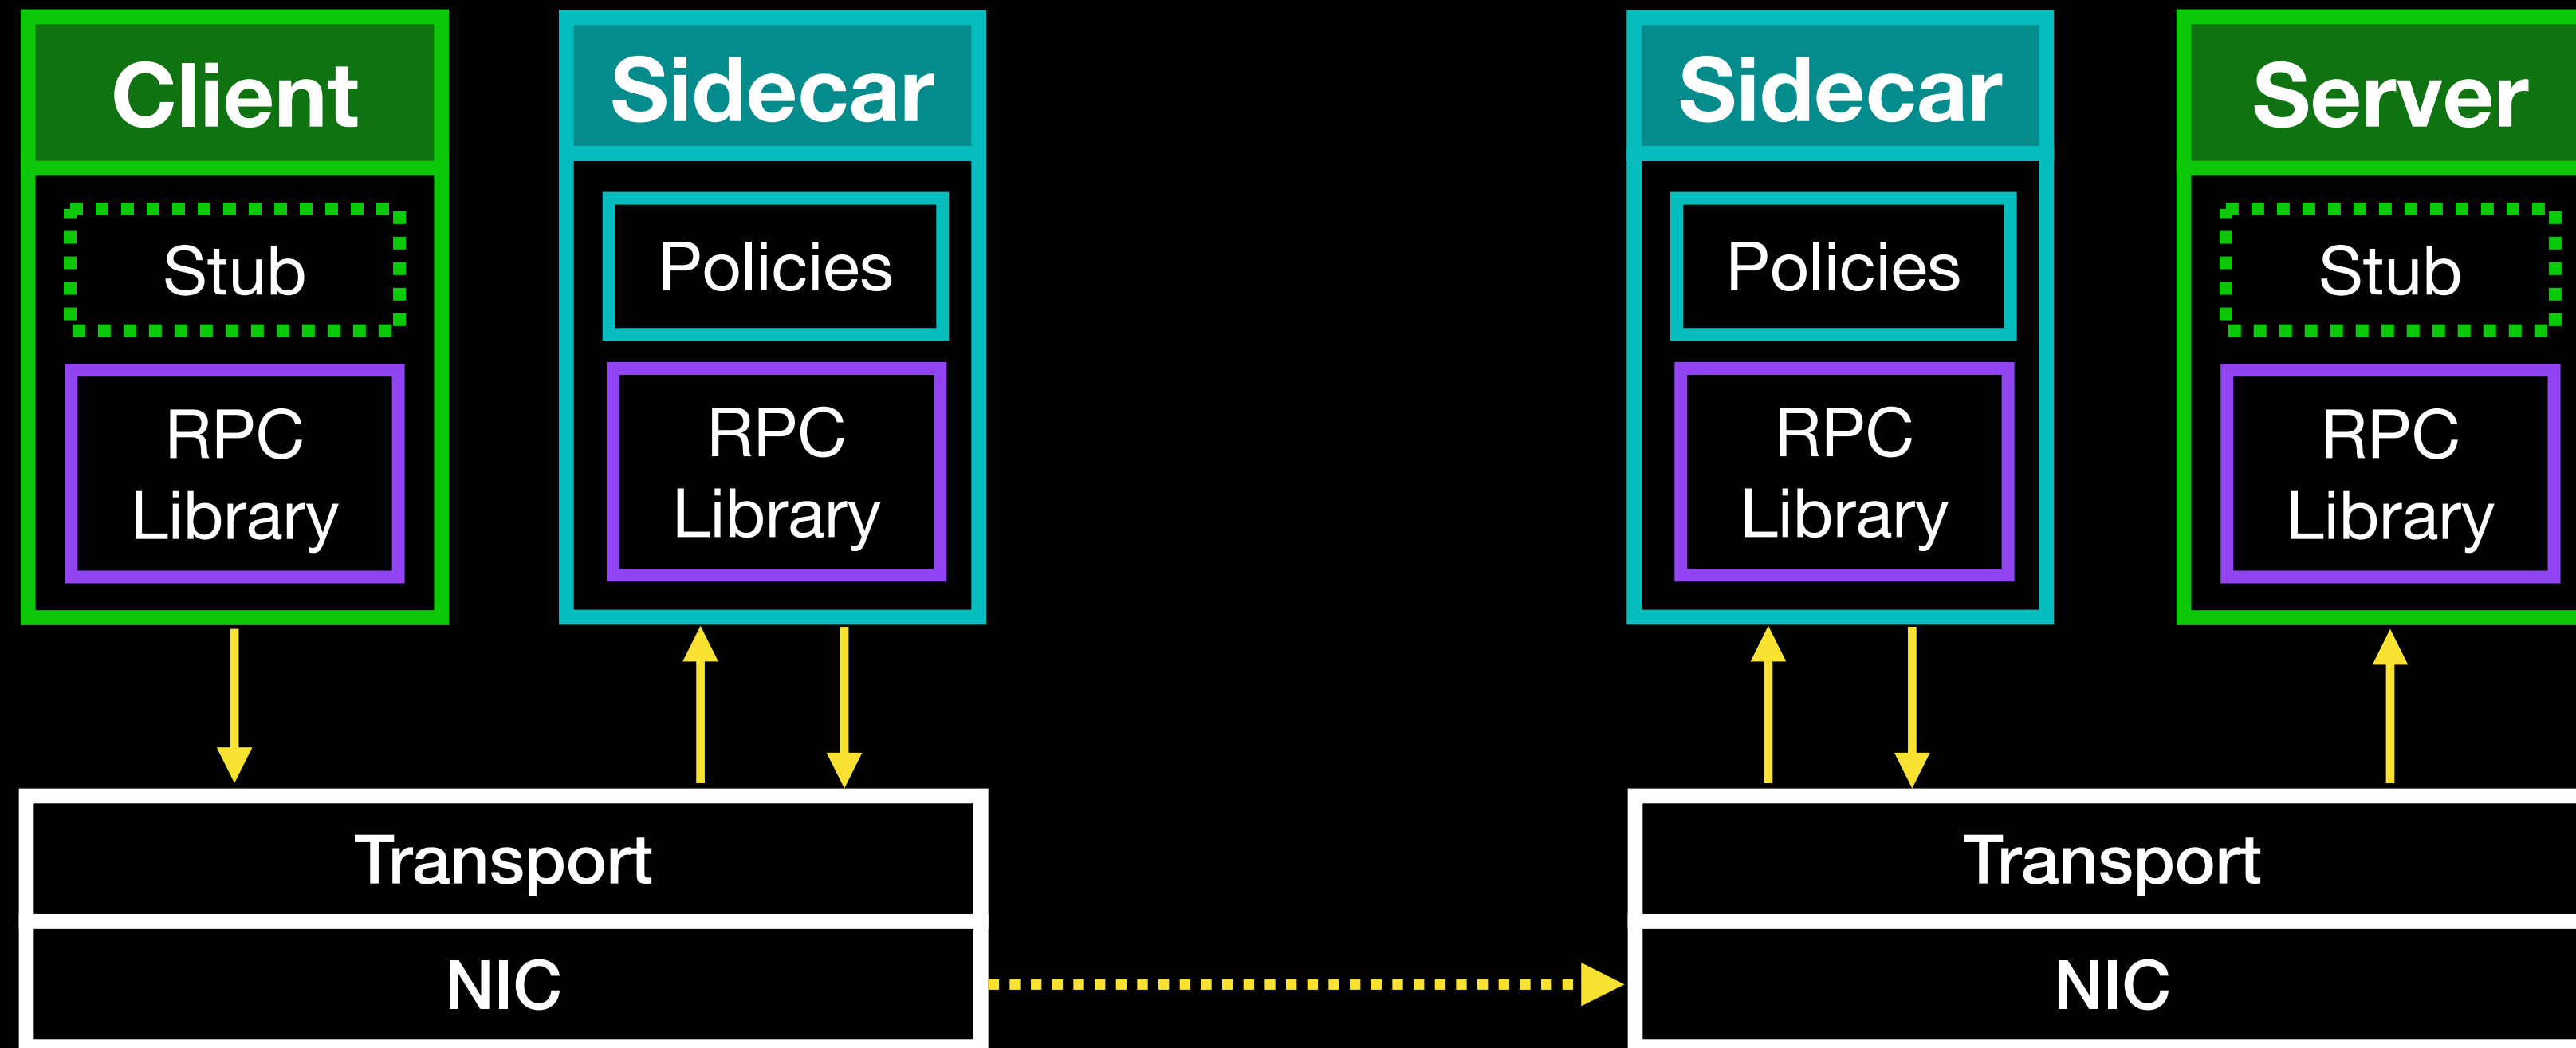
Current Solution to Policy Enforcement



Current Solution to Policy Enforcement



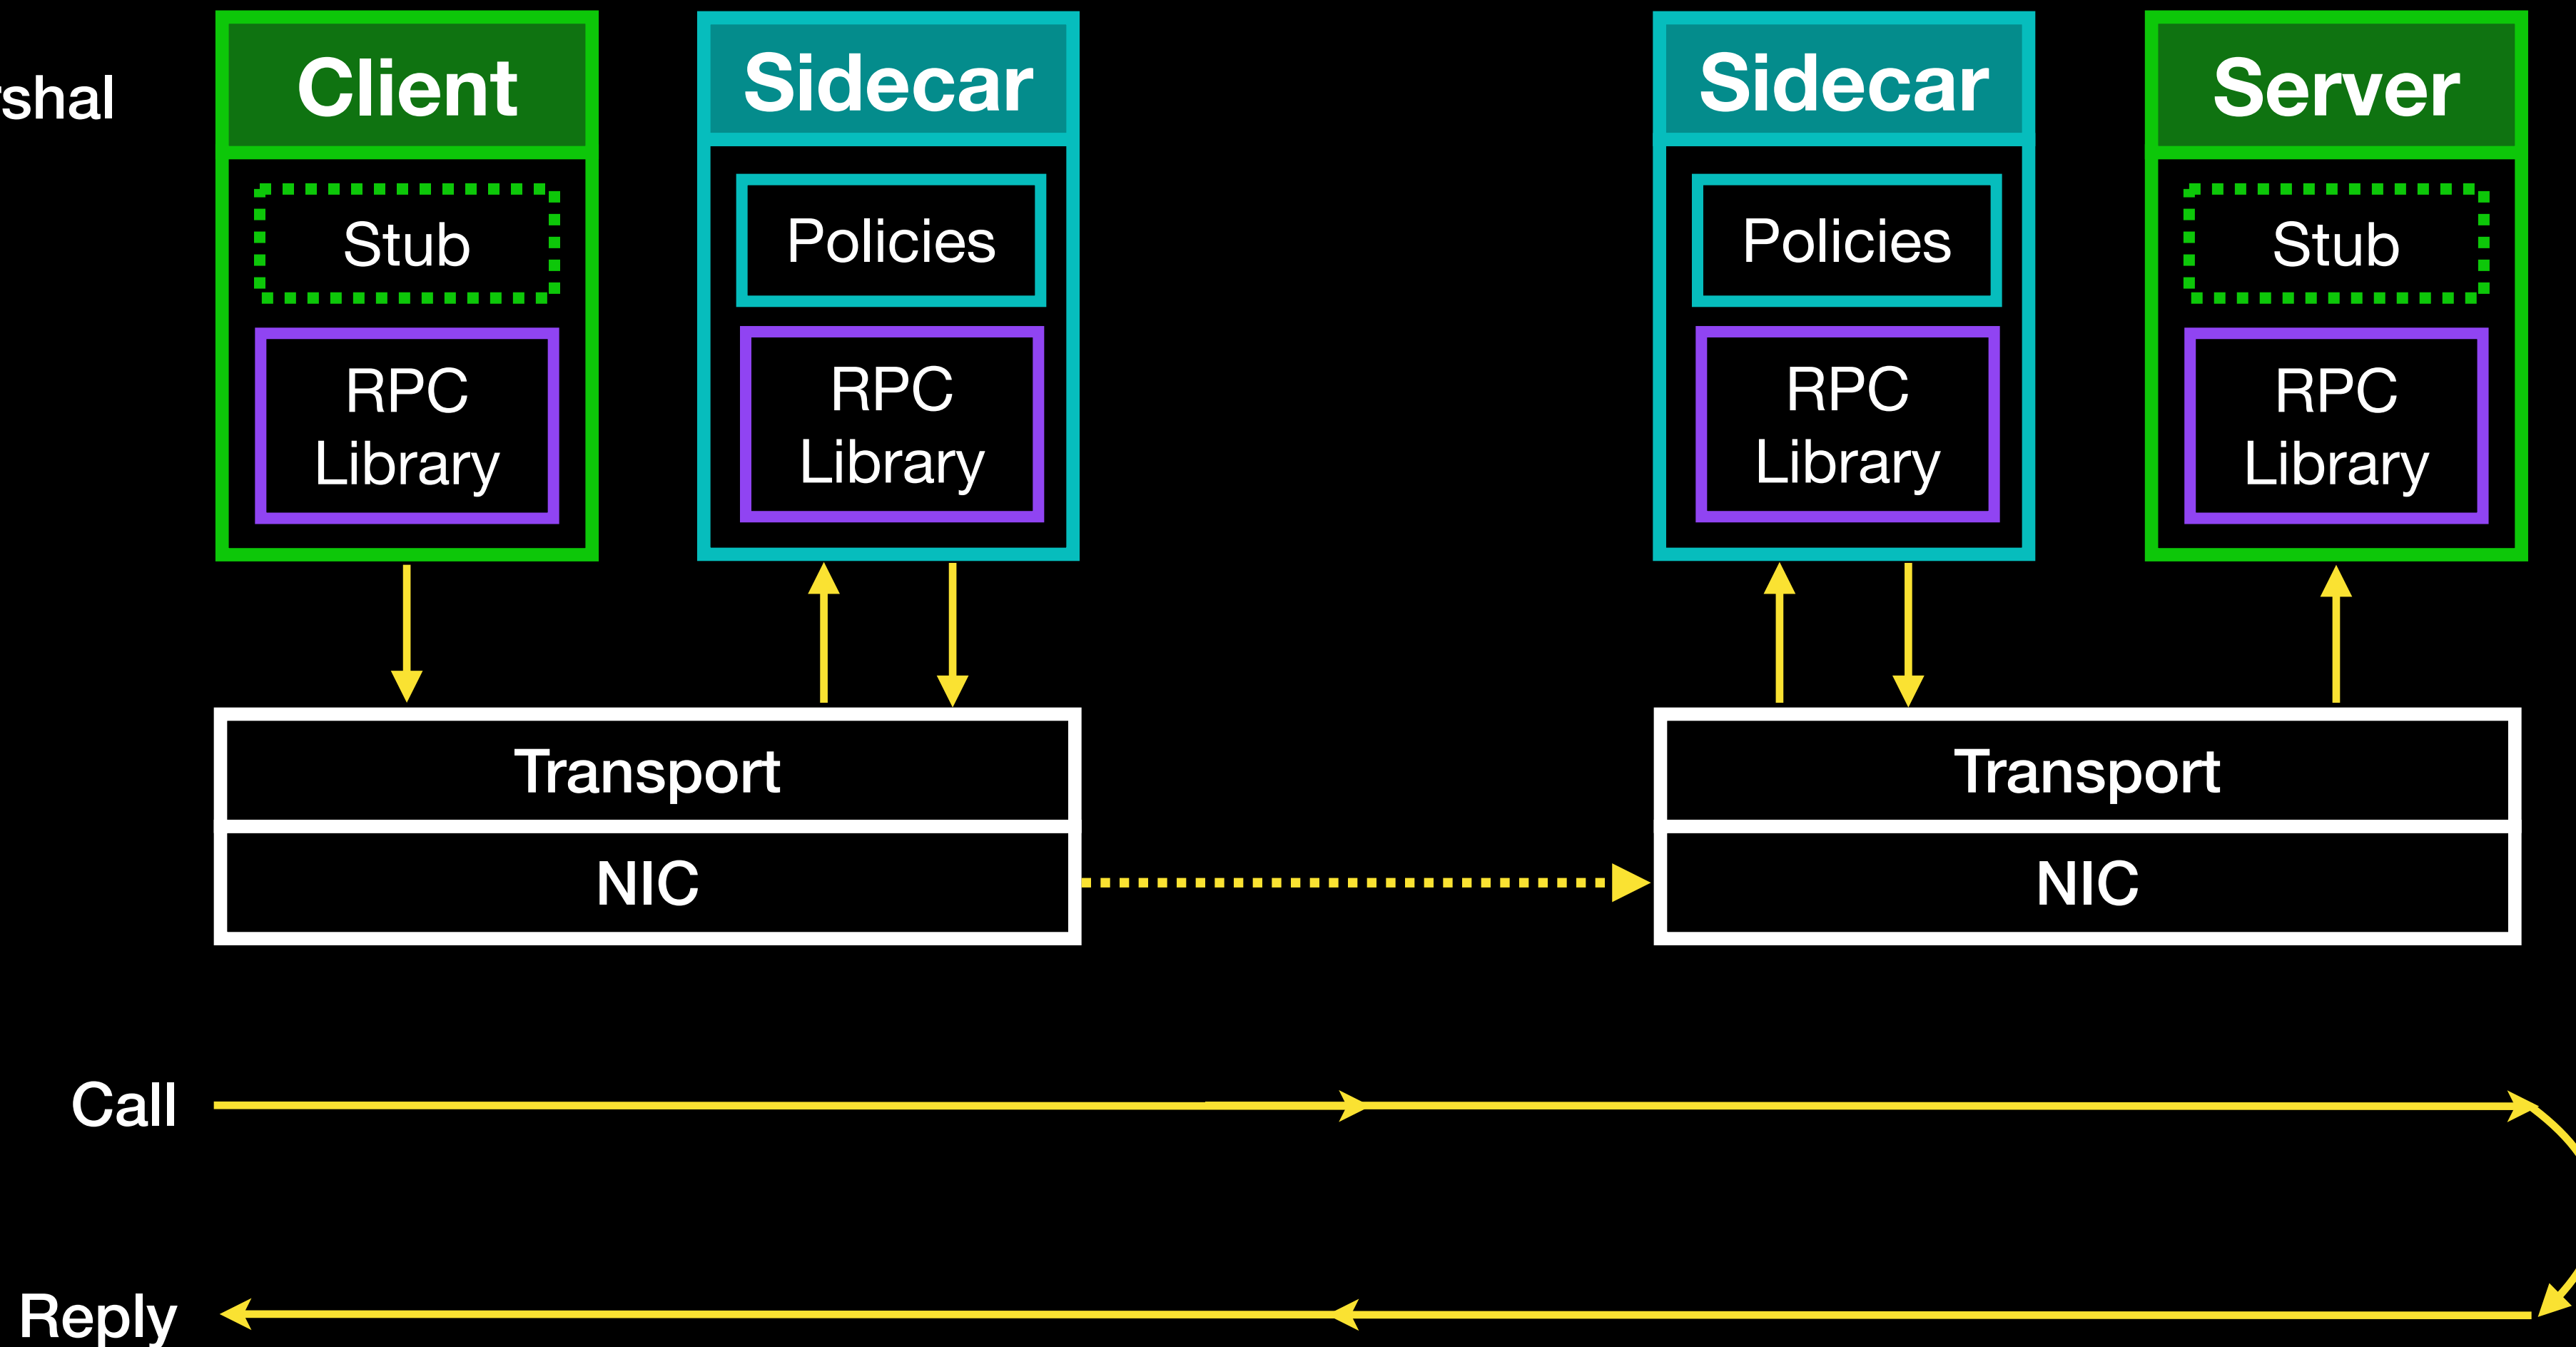
Current Solution to Policy Enforcement



Current Solution to Policy Enforcement

M Marshal

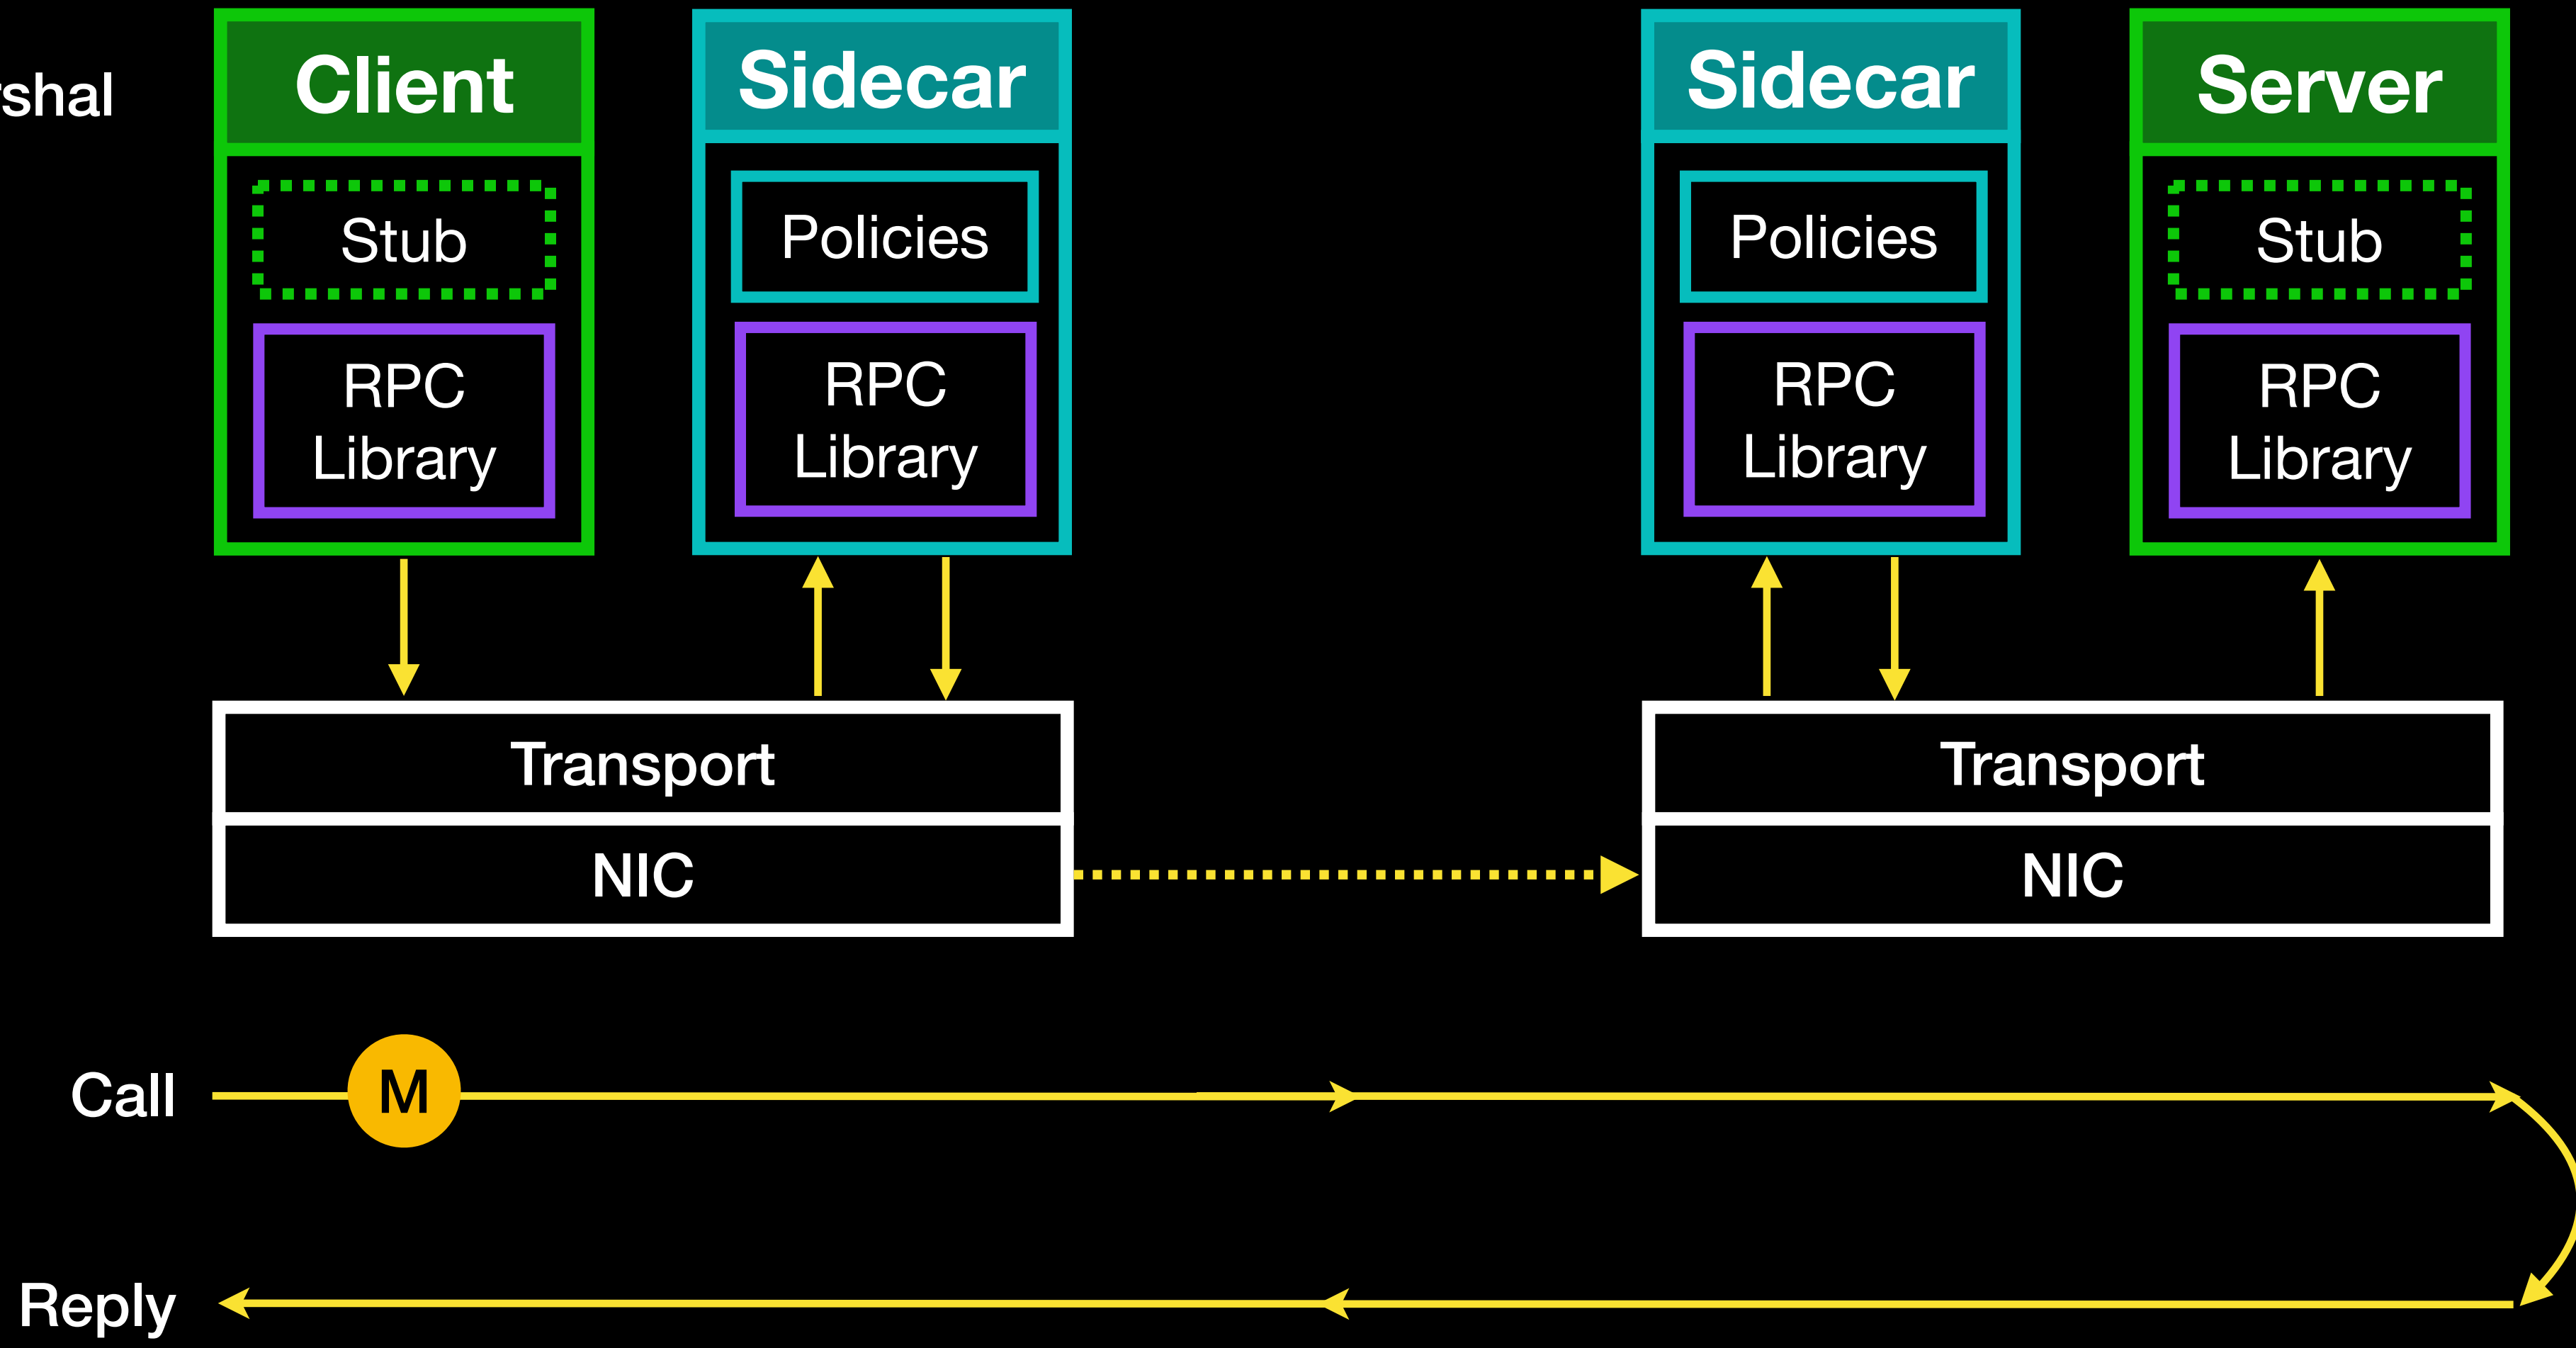
U Unmarshal



Current Solution to Policy Enforcement

M Marshal

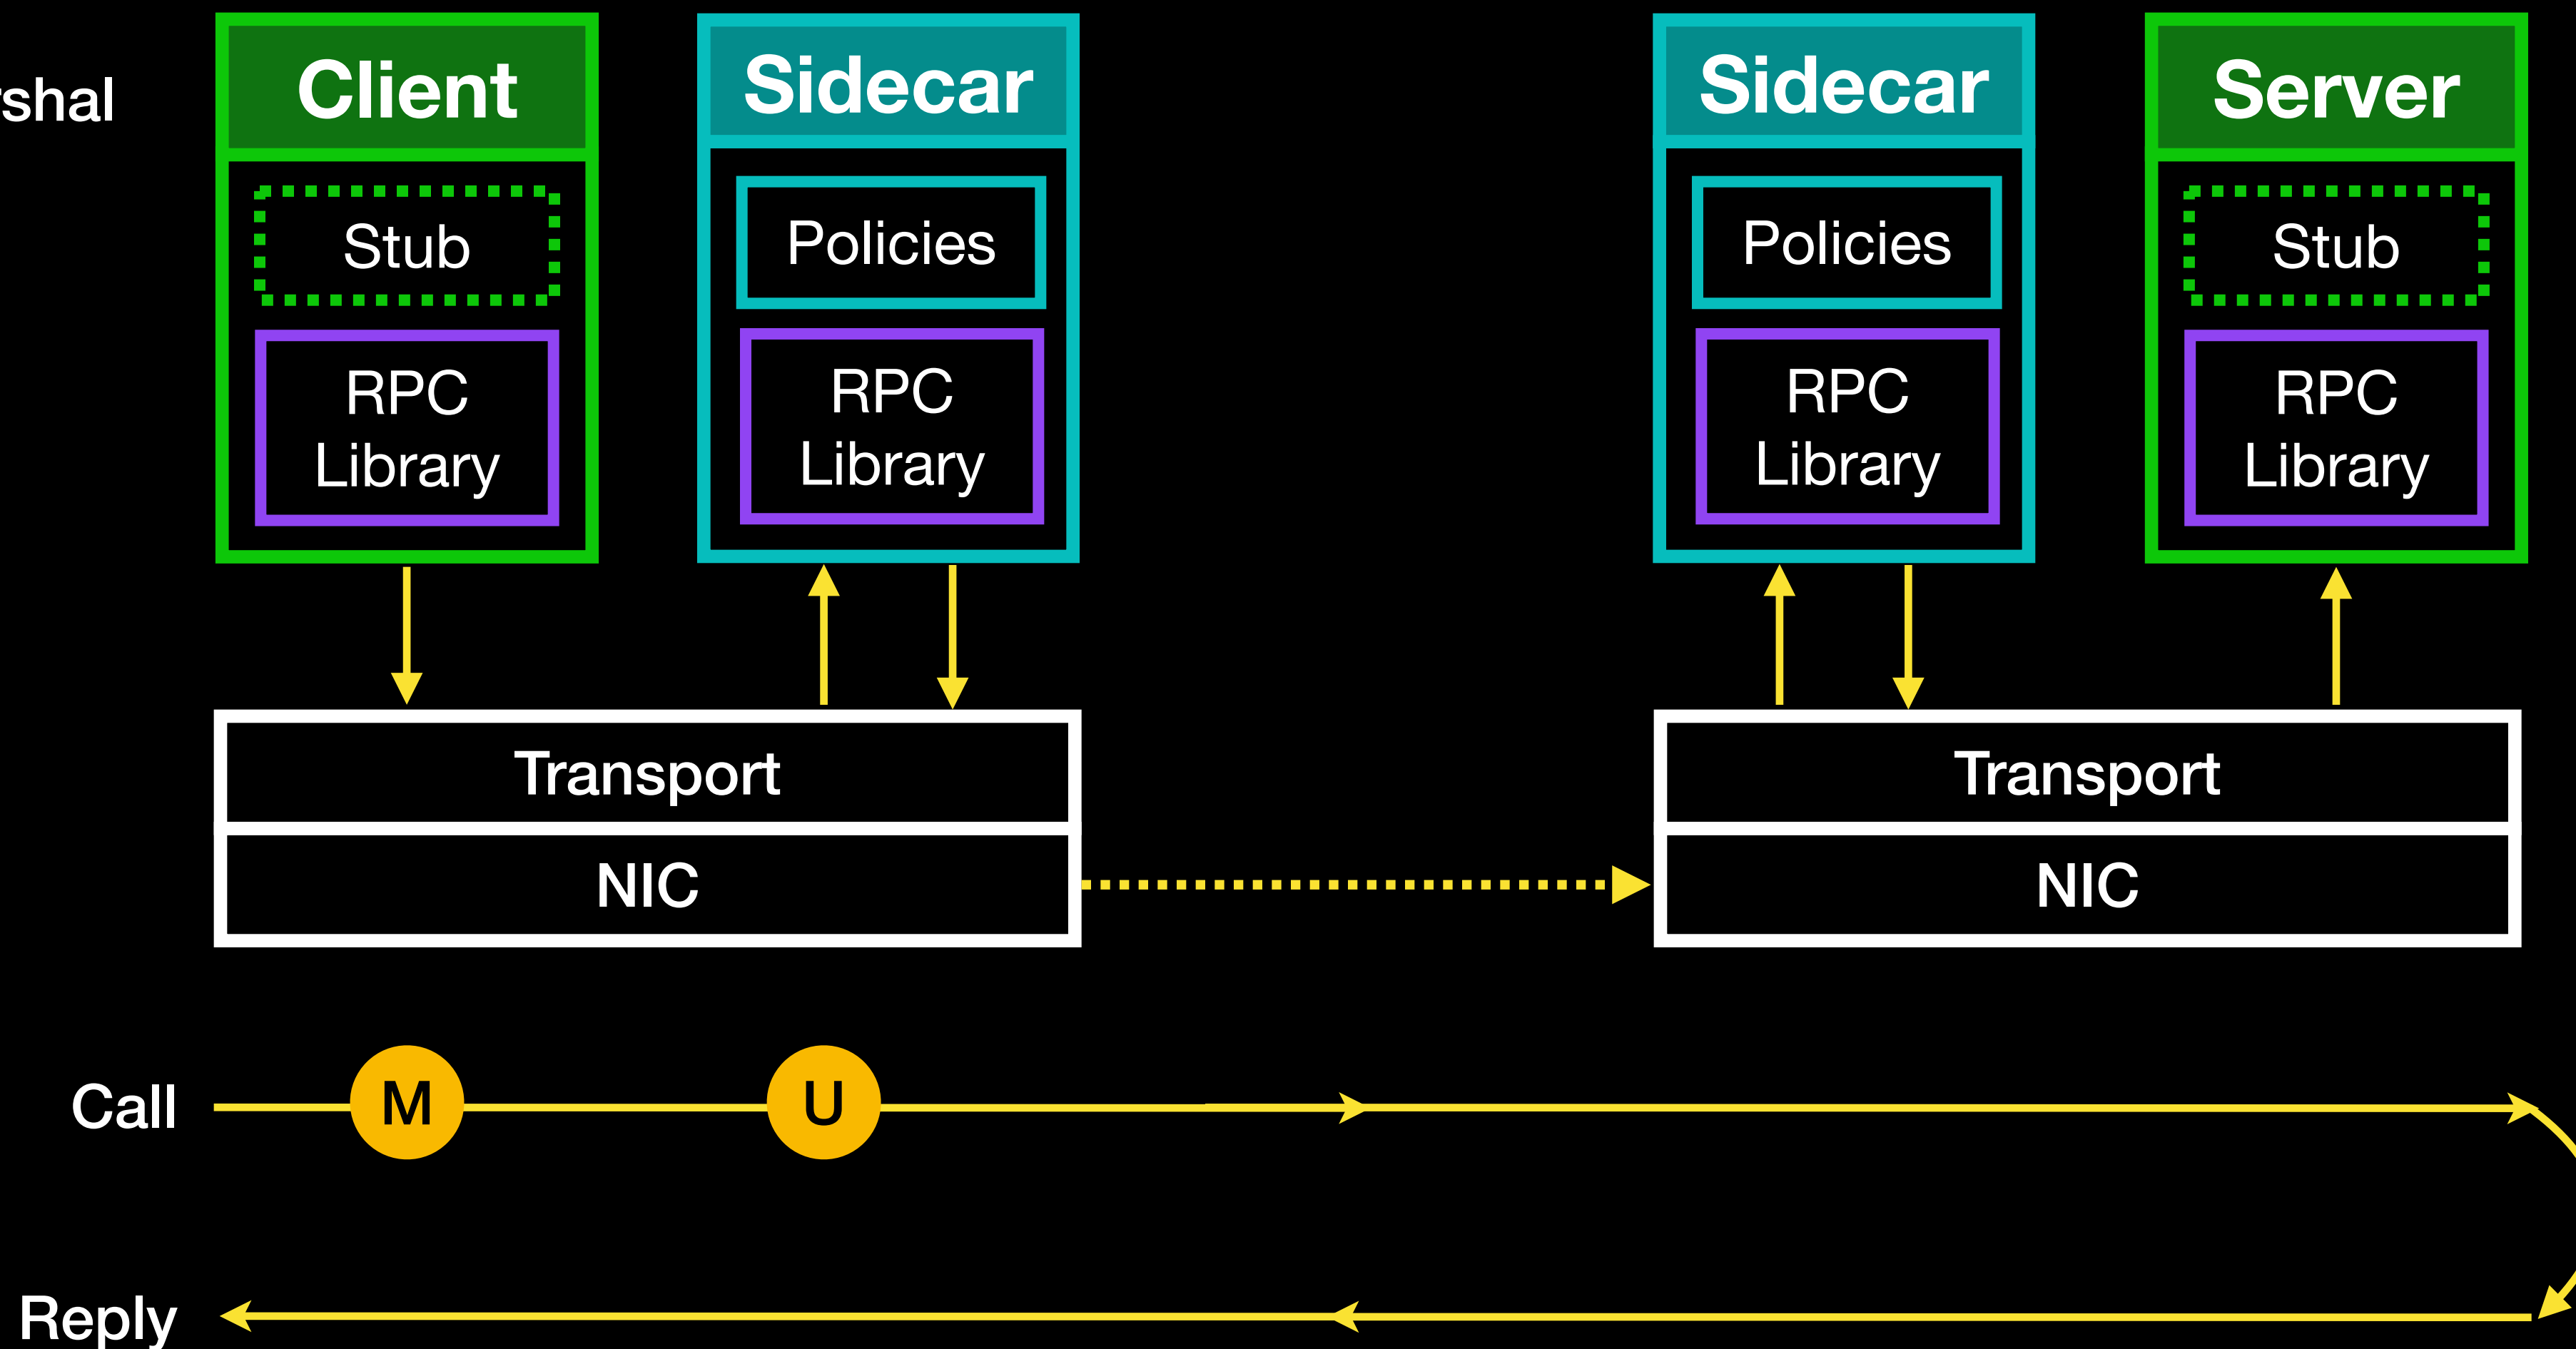
U Unmarshal



Current Solution to Policy Enforcement

M Marshal

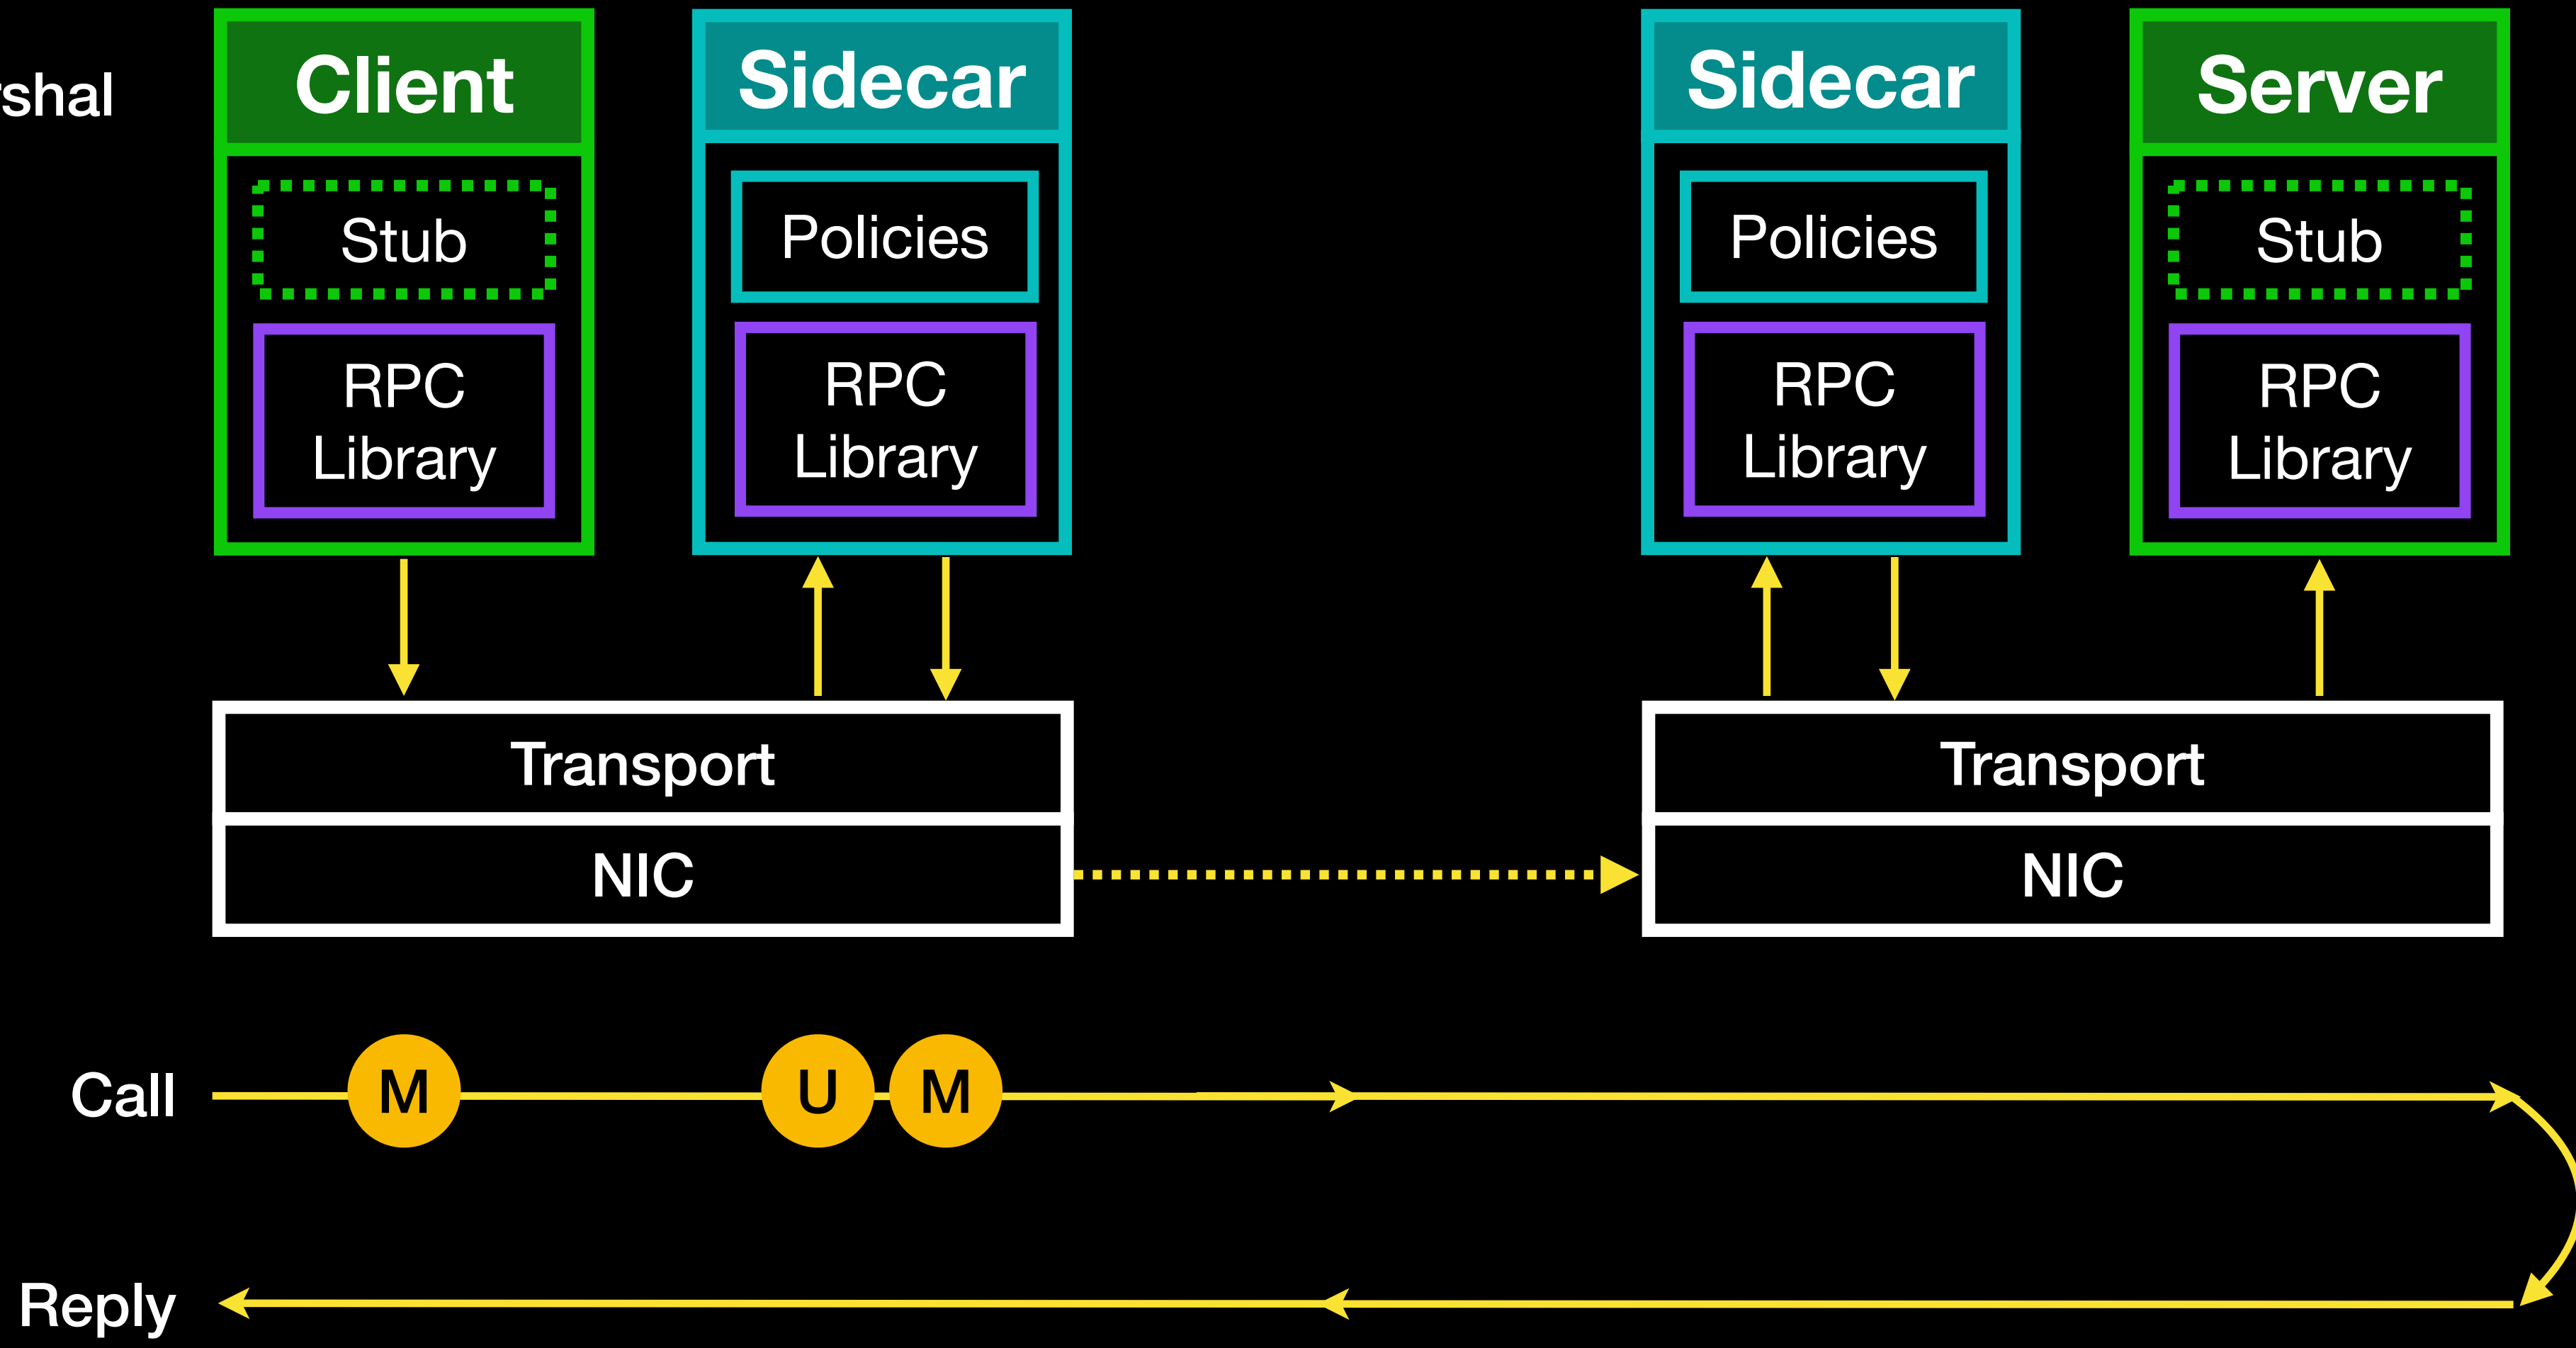
U Unmarshal



Current Solution to Policy Enforcement

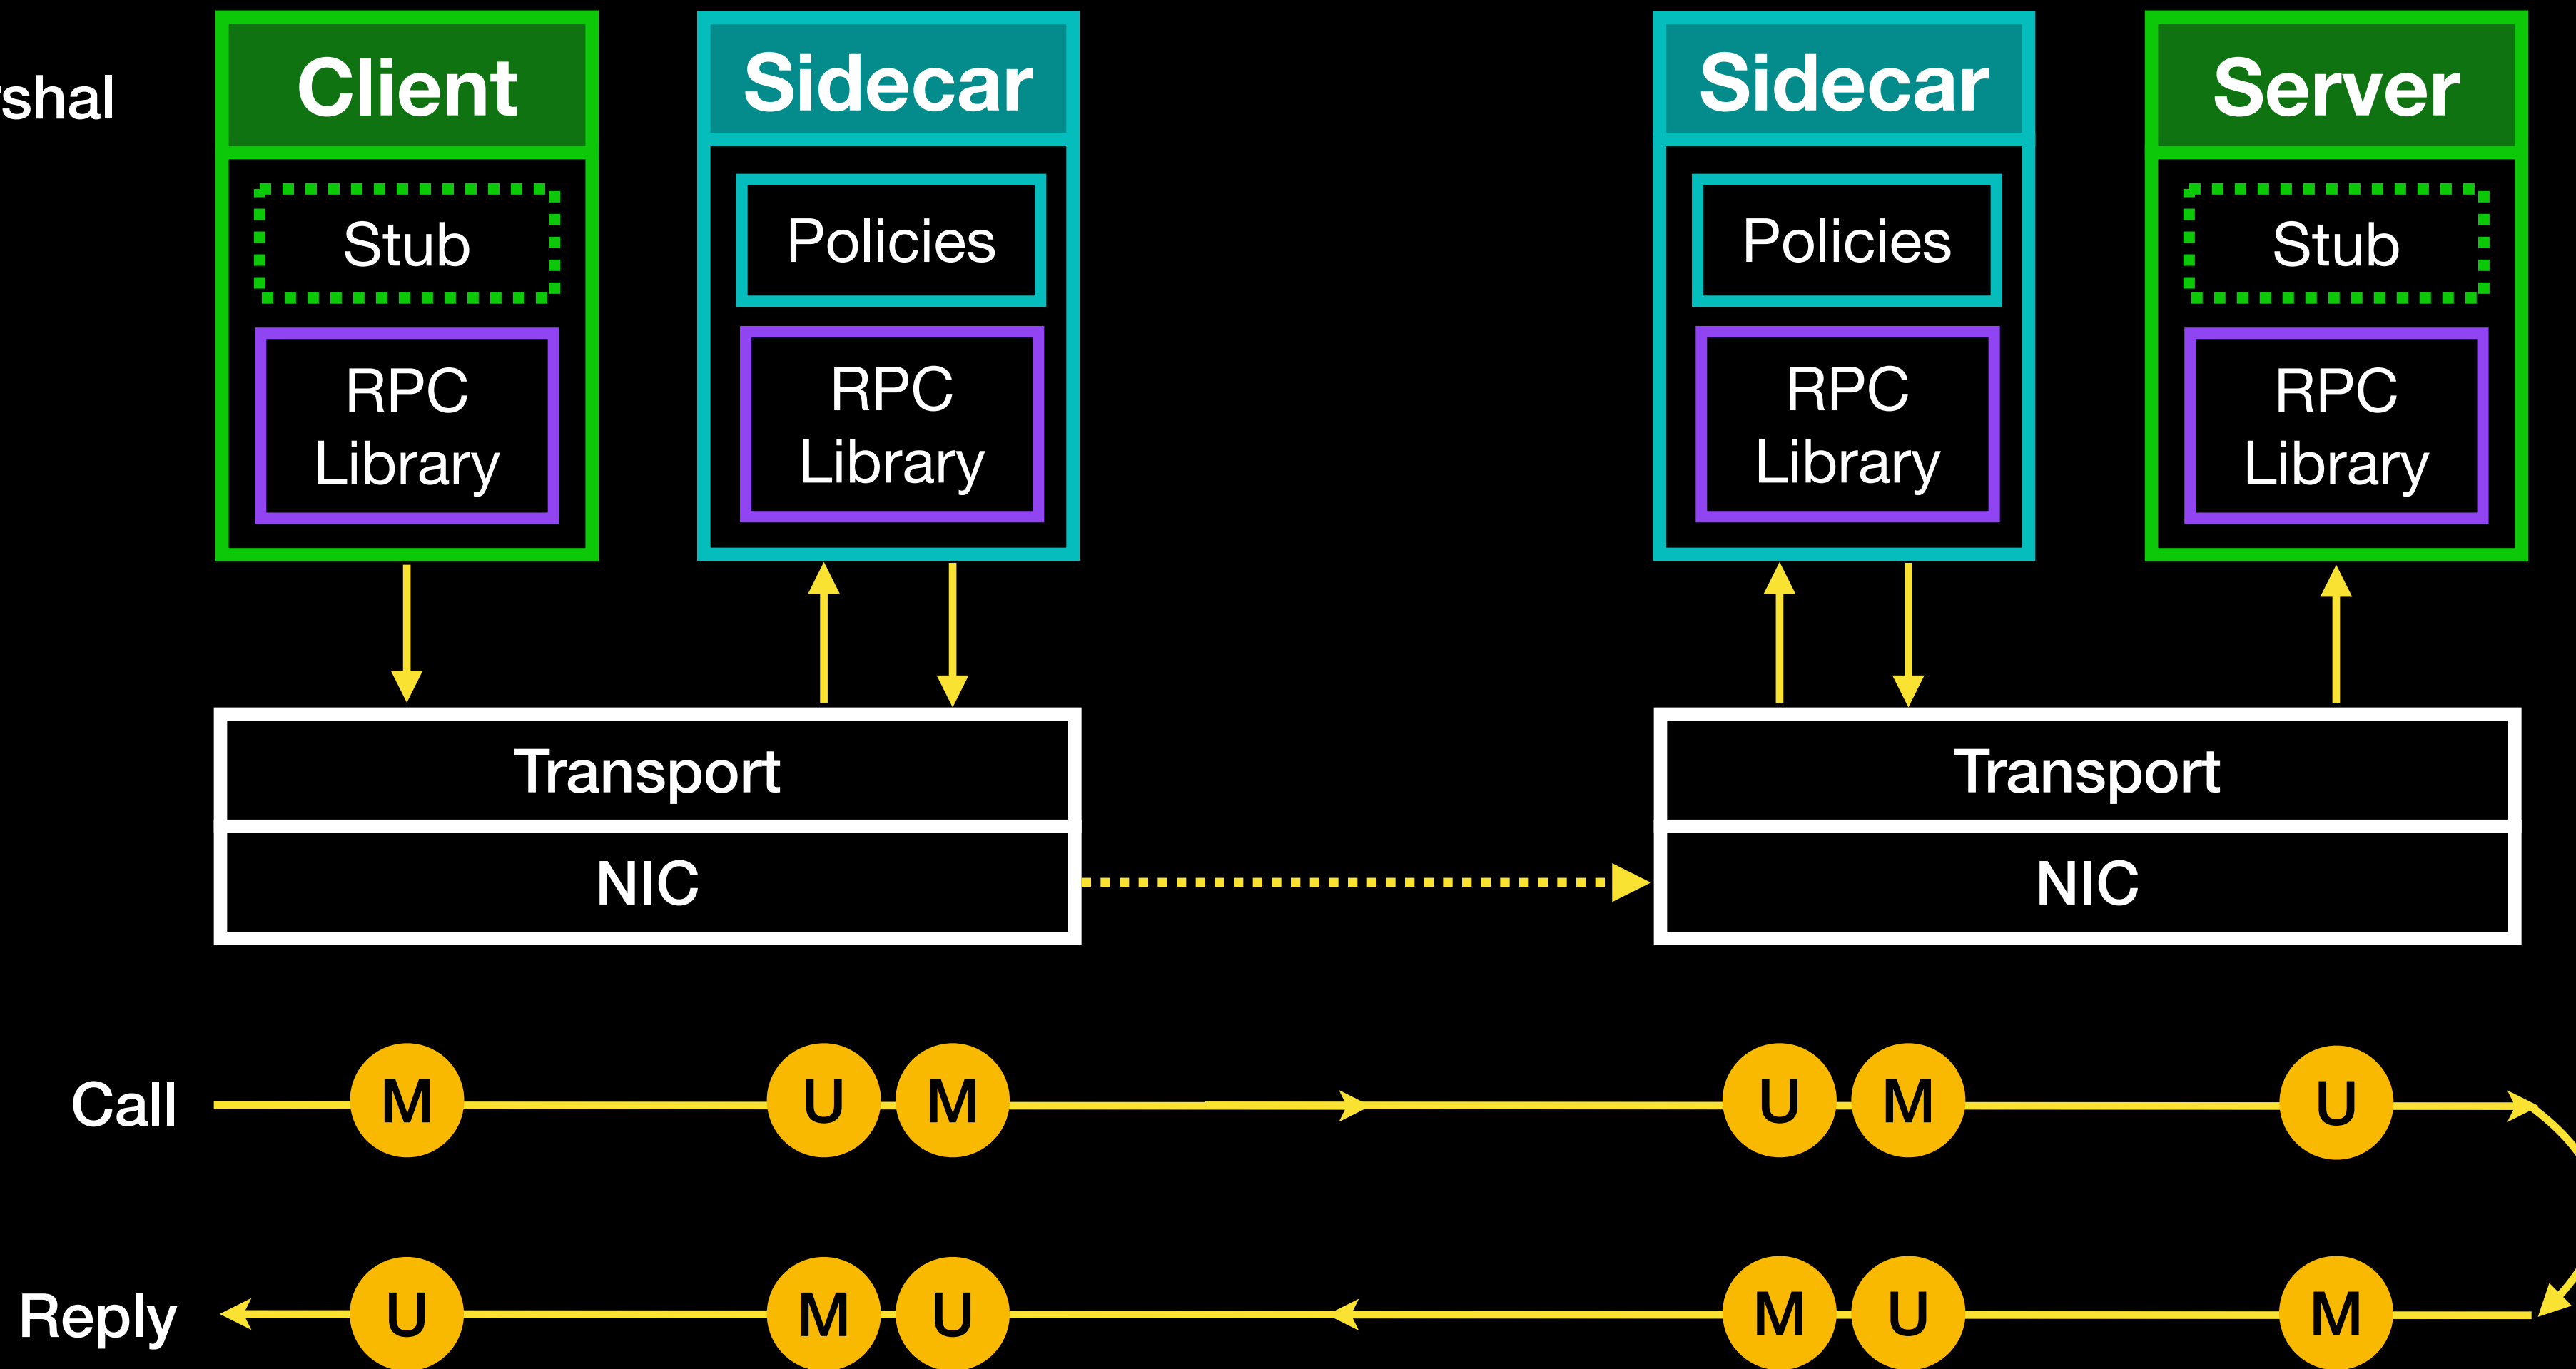
M Marshal

U Unmarshal



Current Solution to Policy Enforcement


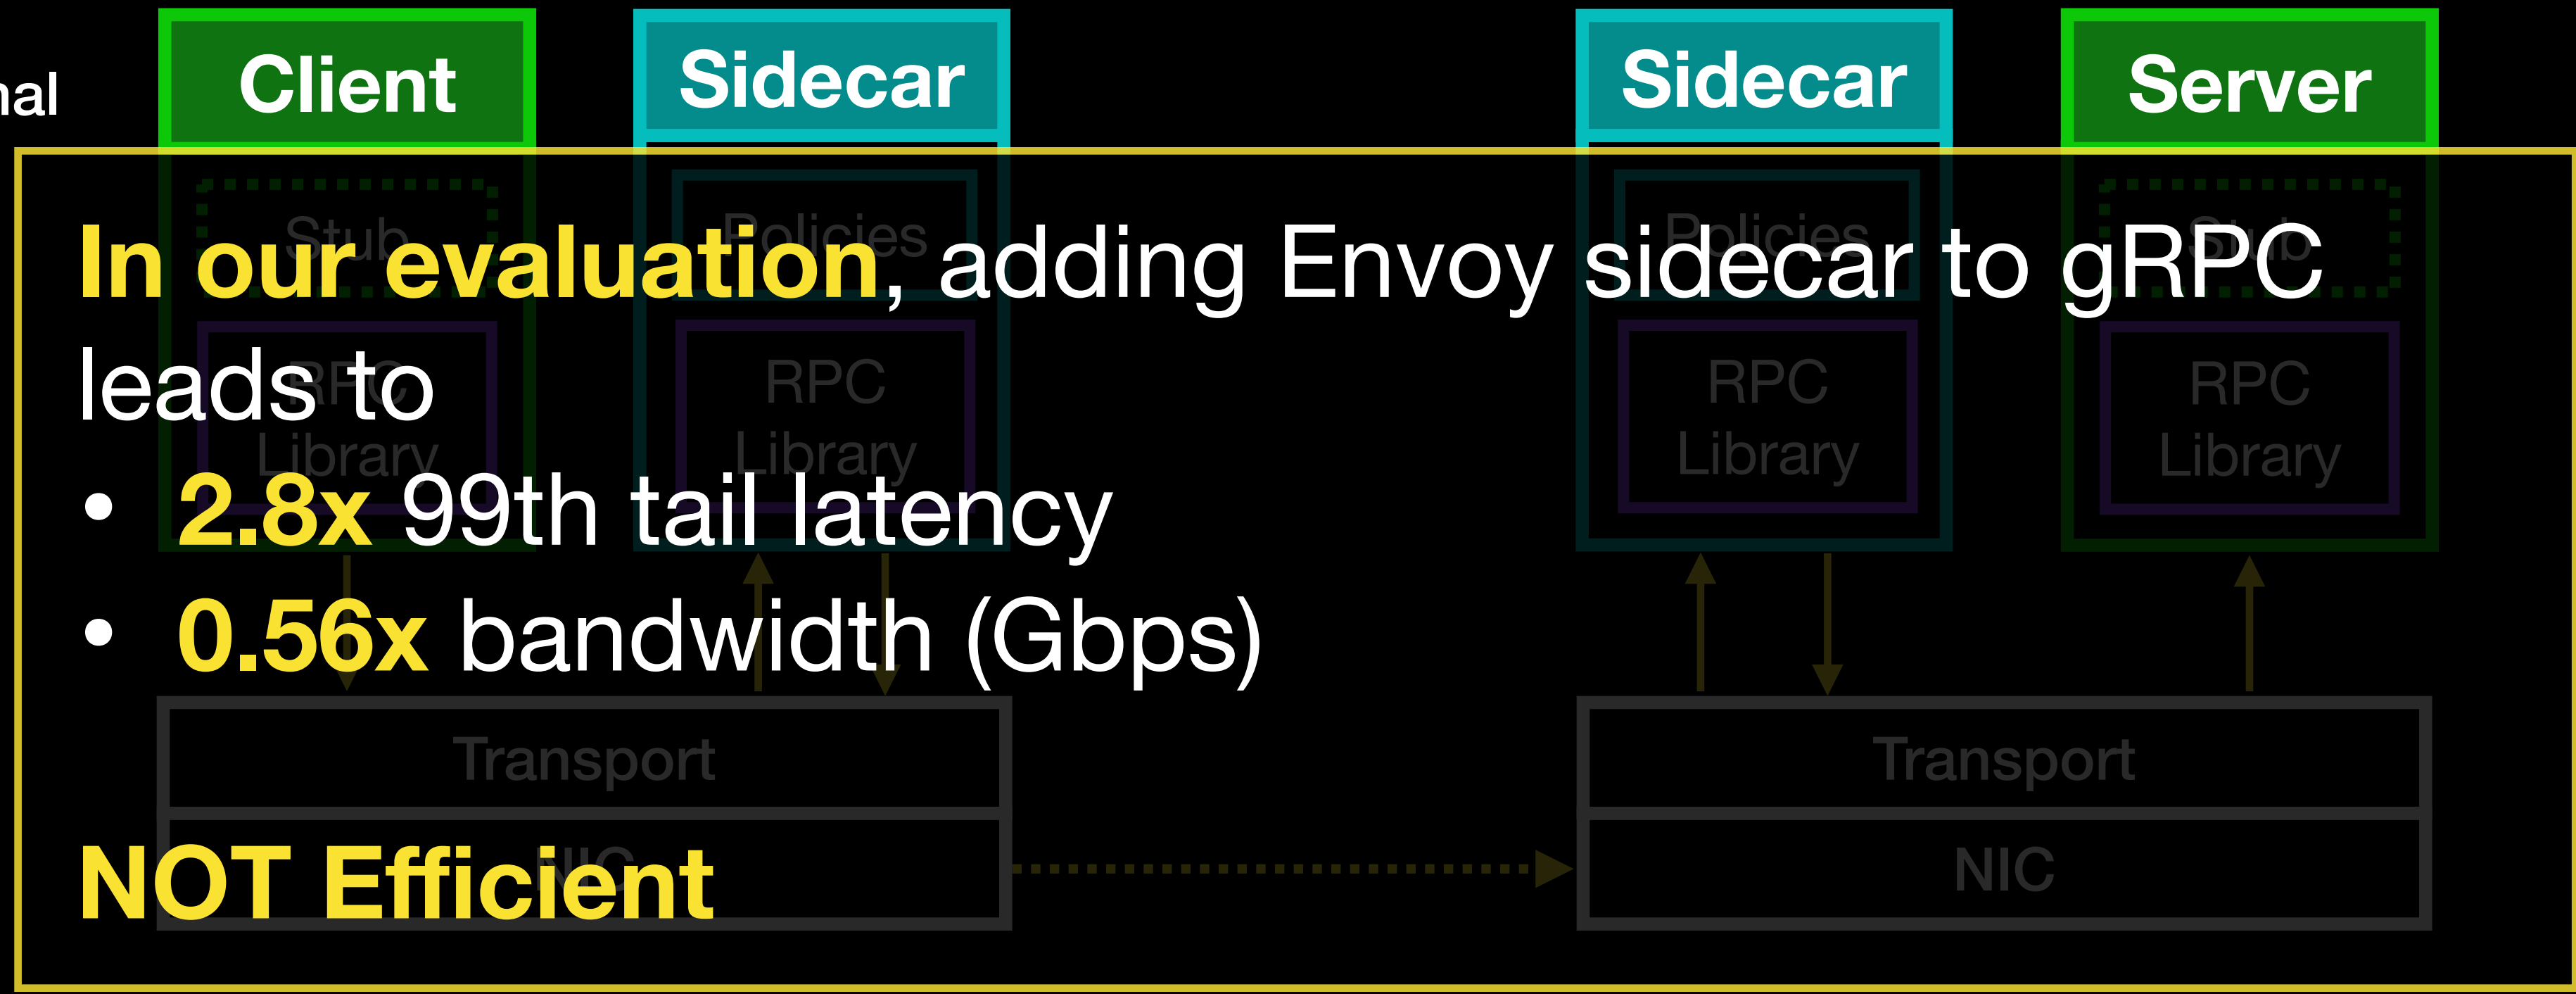
- M** Marshal
- U** Unmarshal



Current Solution to Policy Enforcement

M Marshal

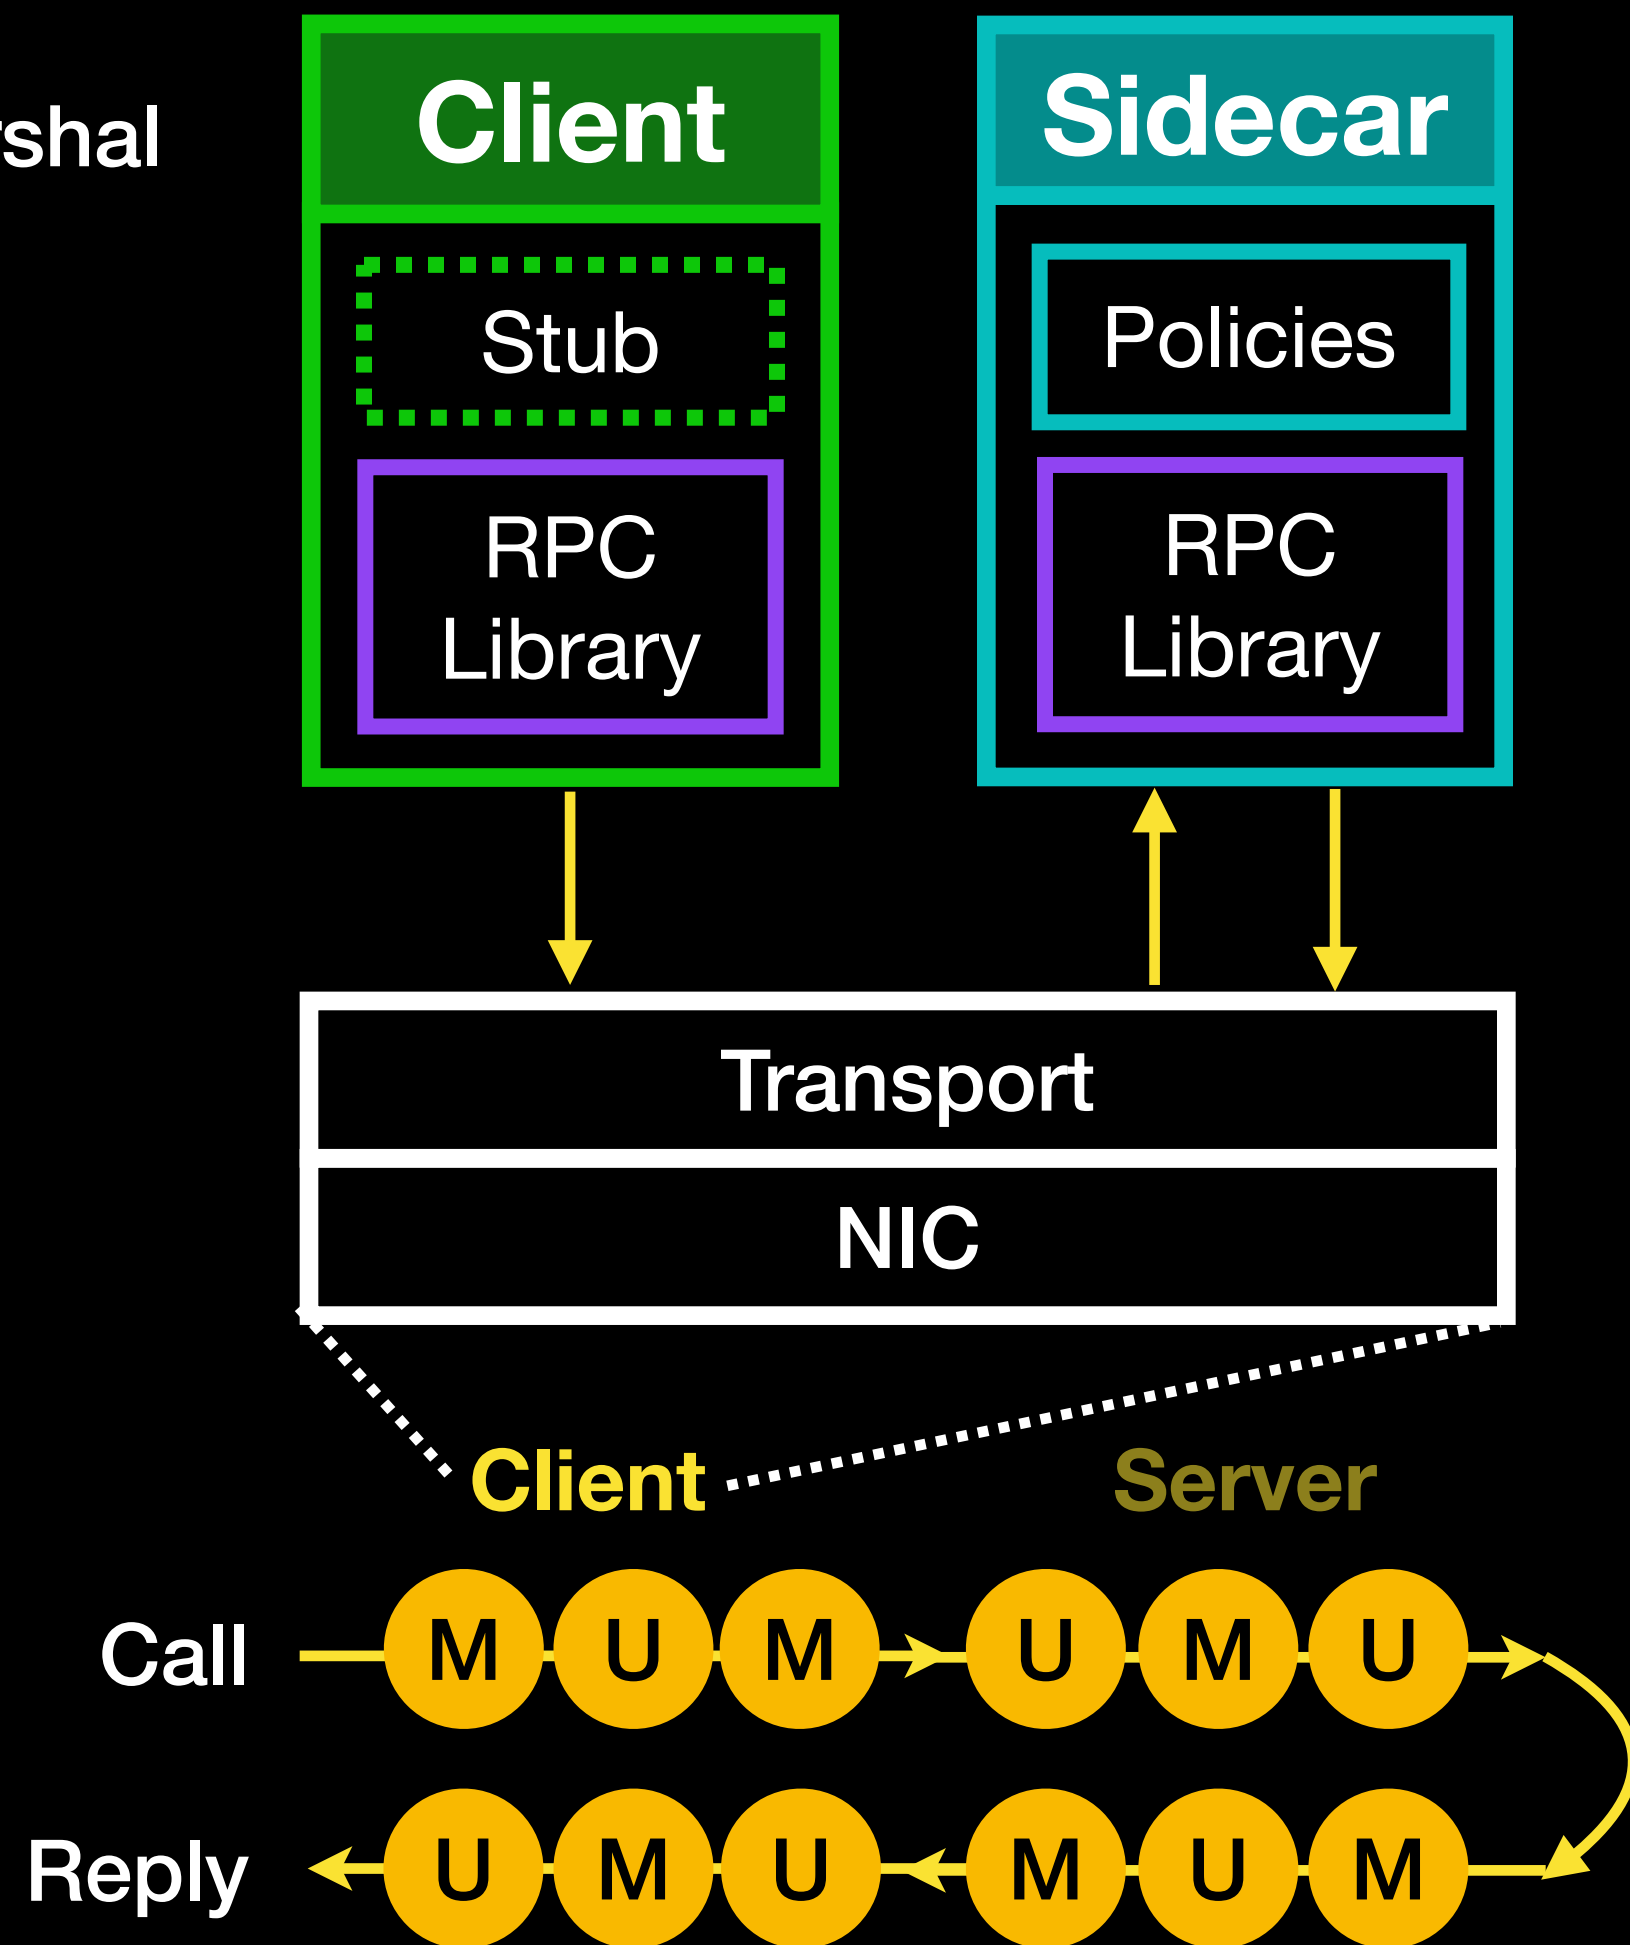
U Unmarshal



RPC-as-a-Library Limitation

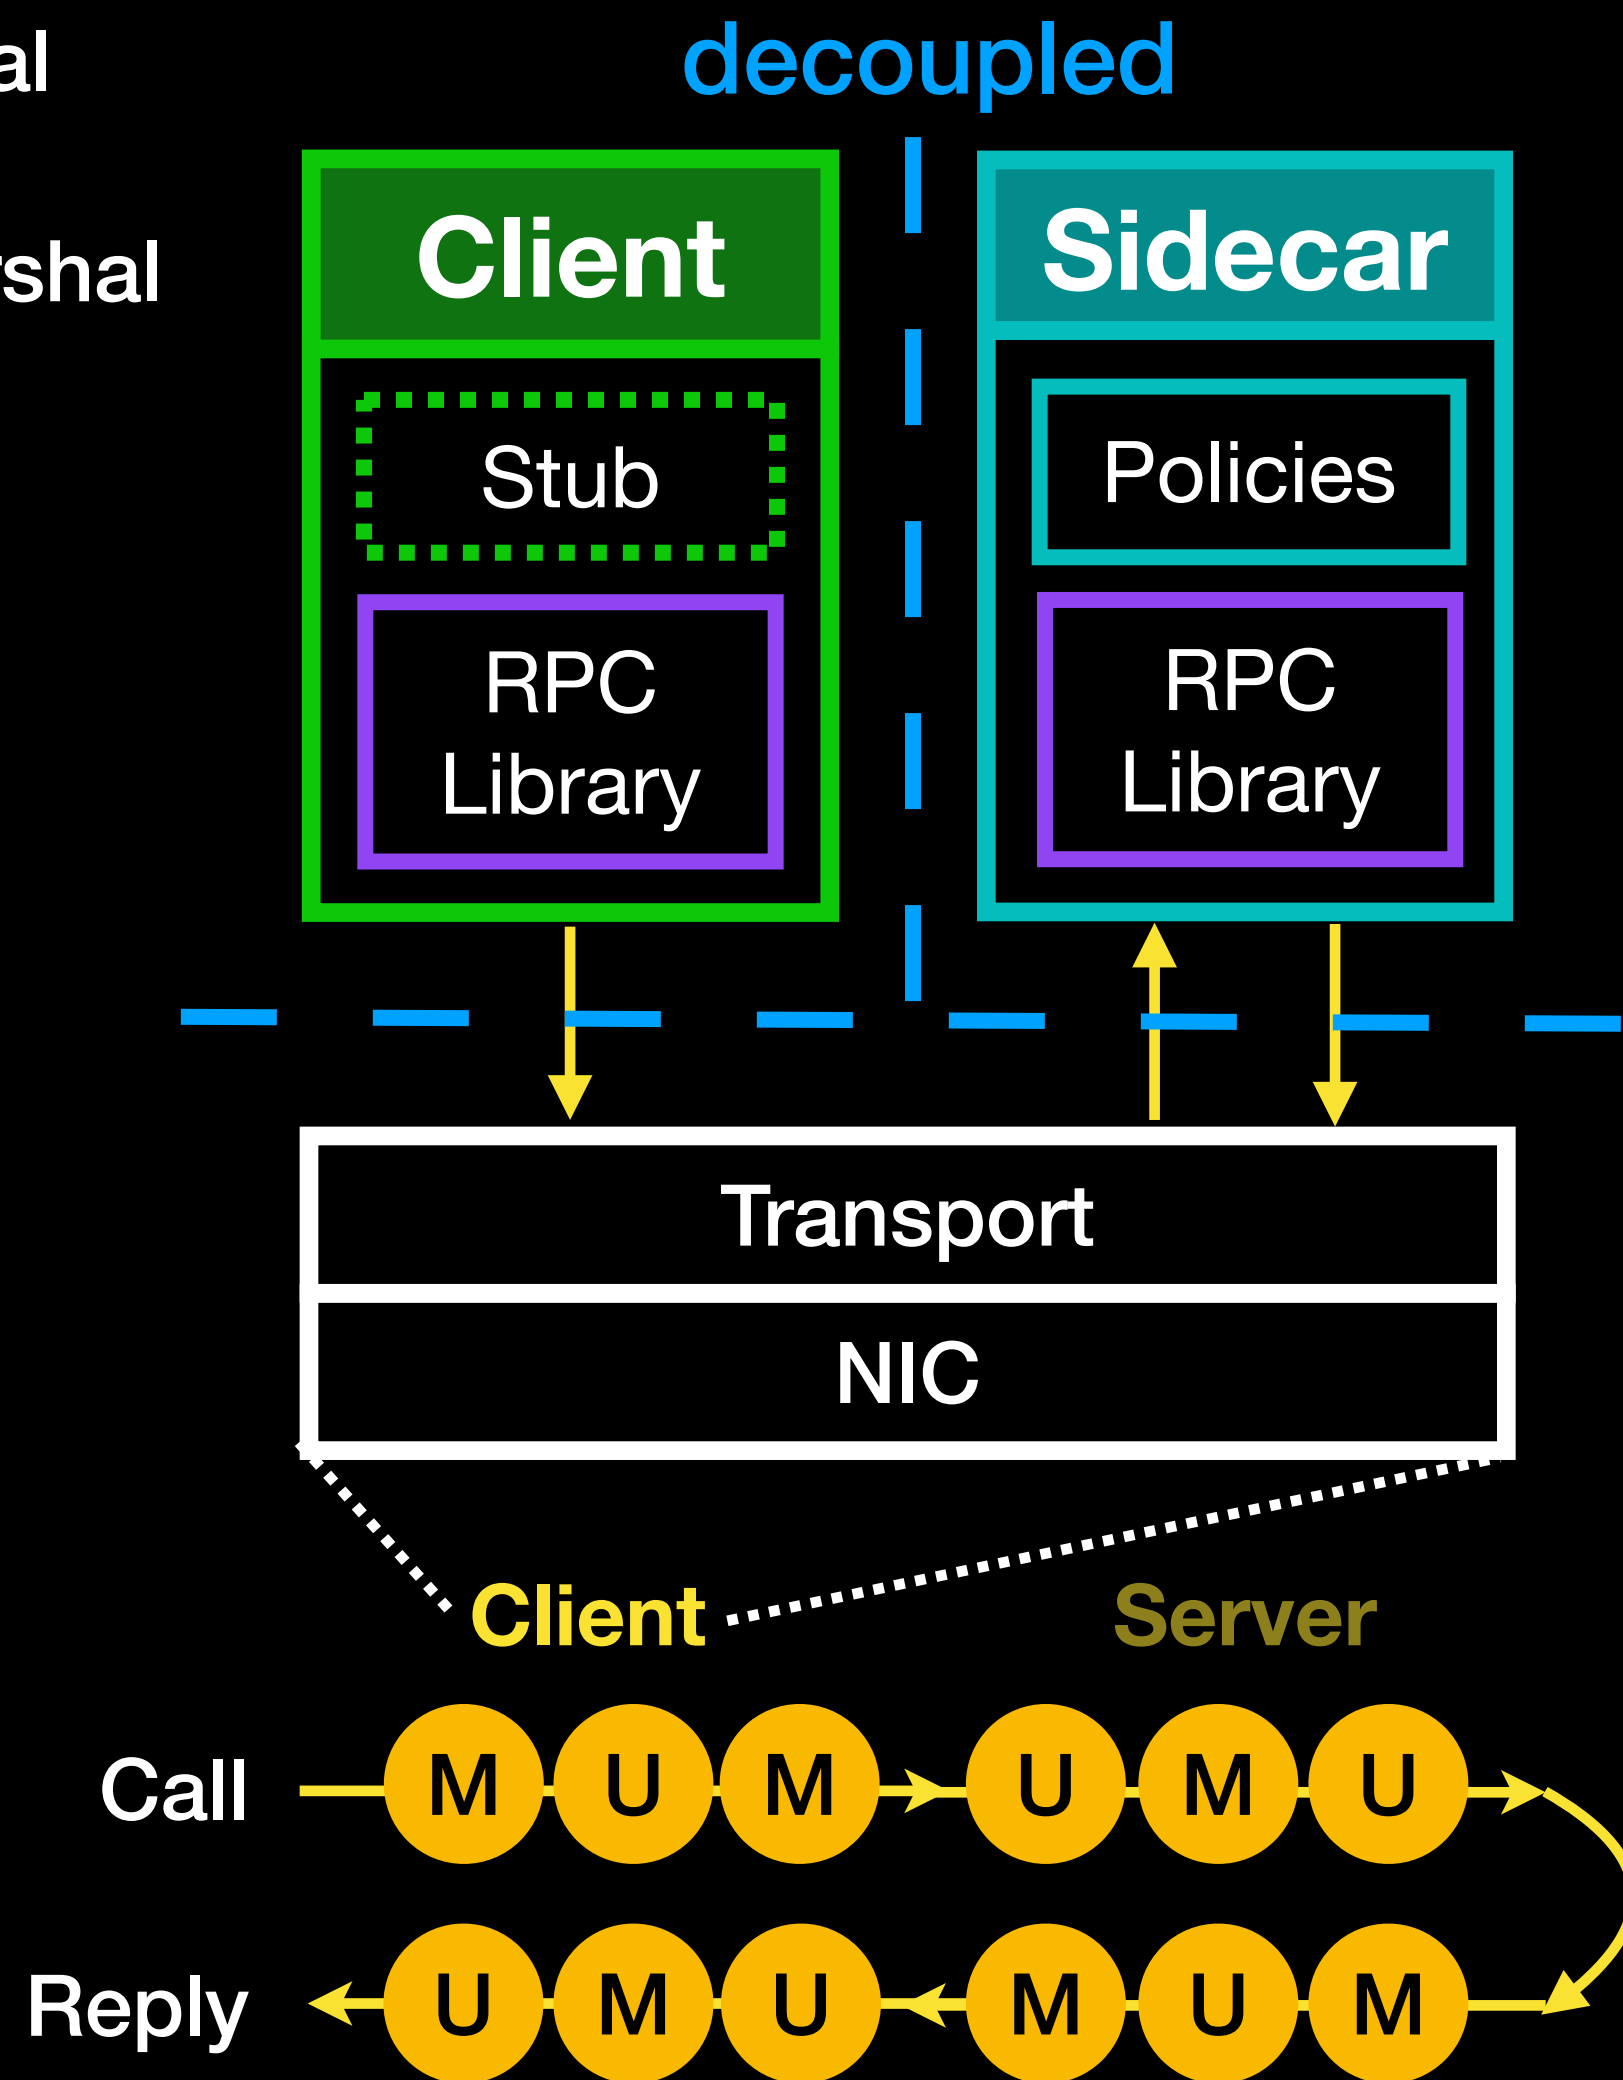
M Marshal

U Unmarshal



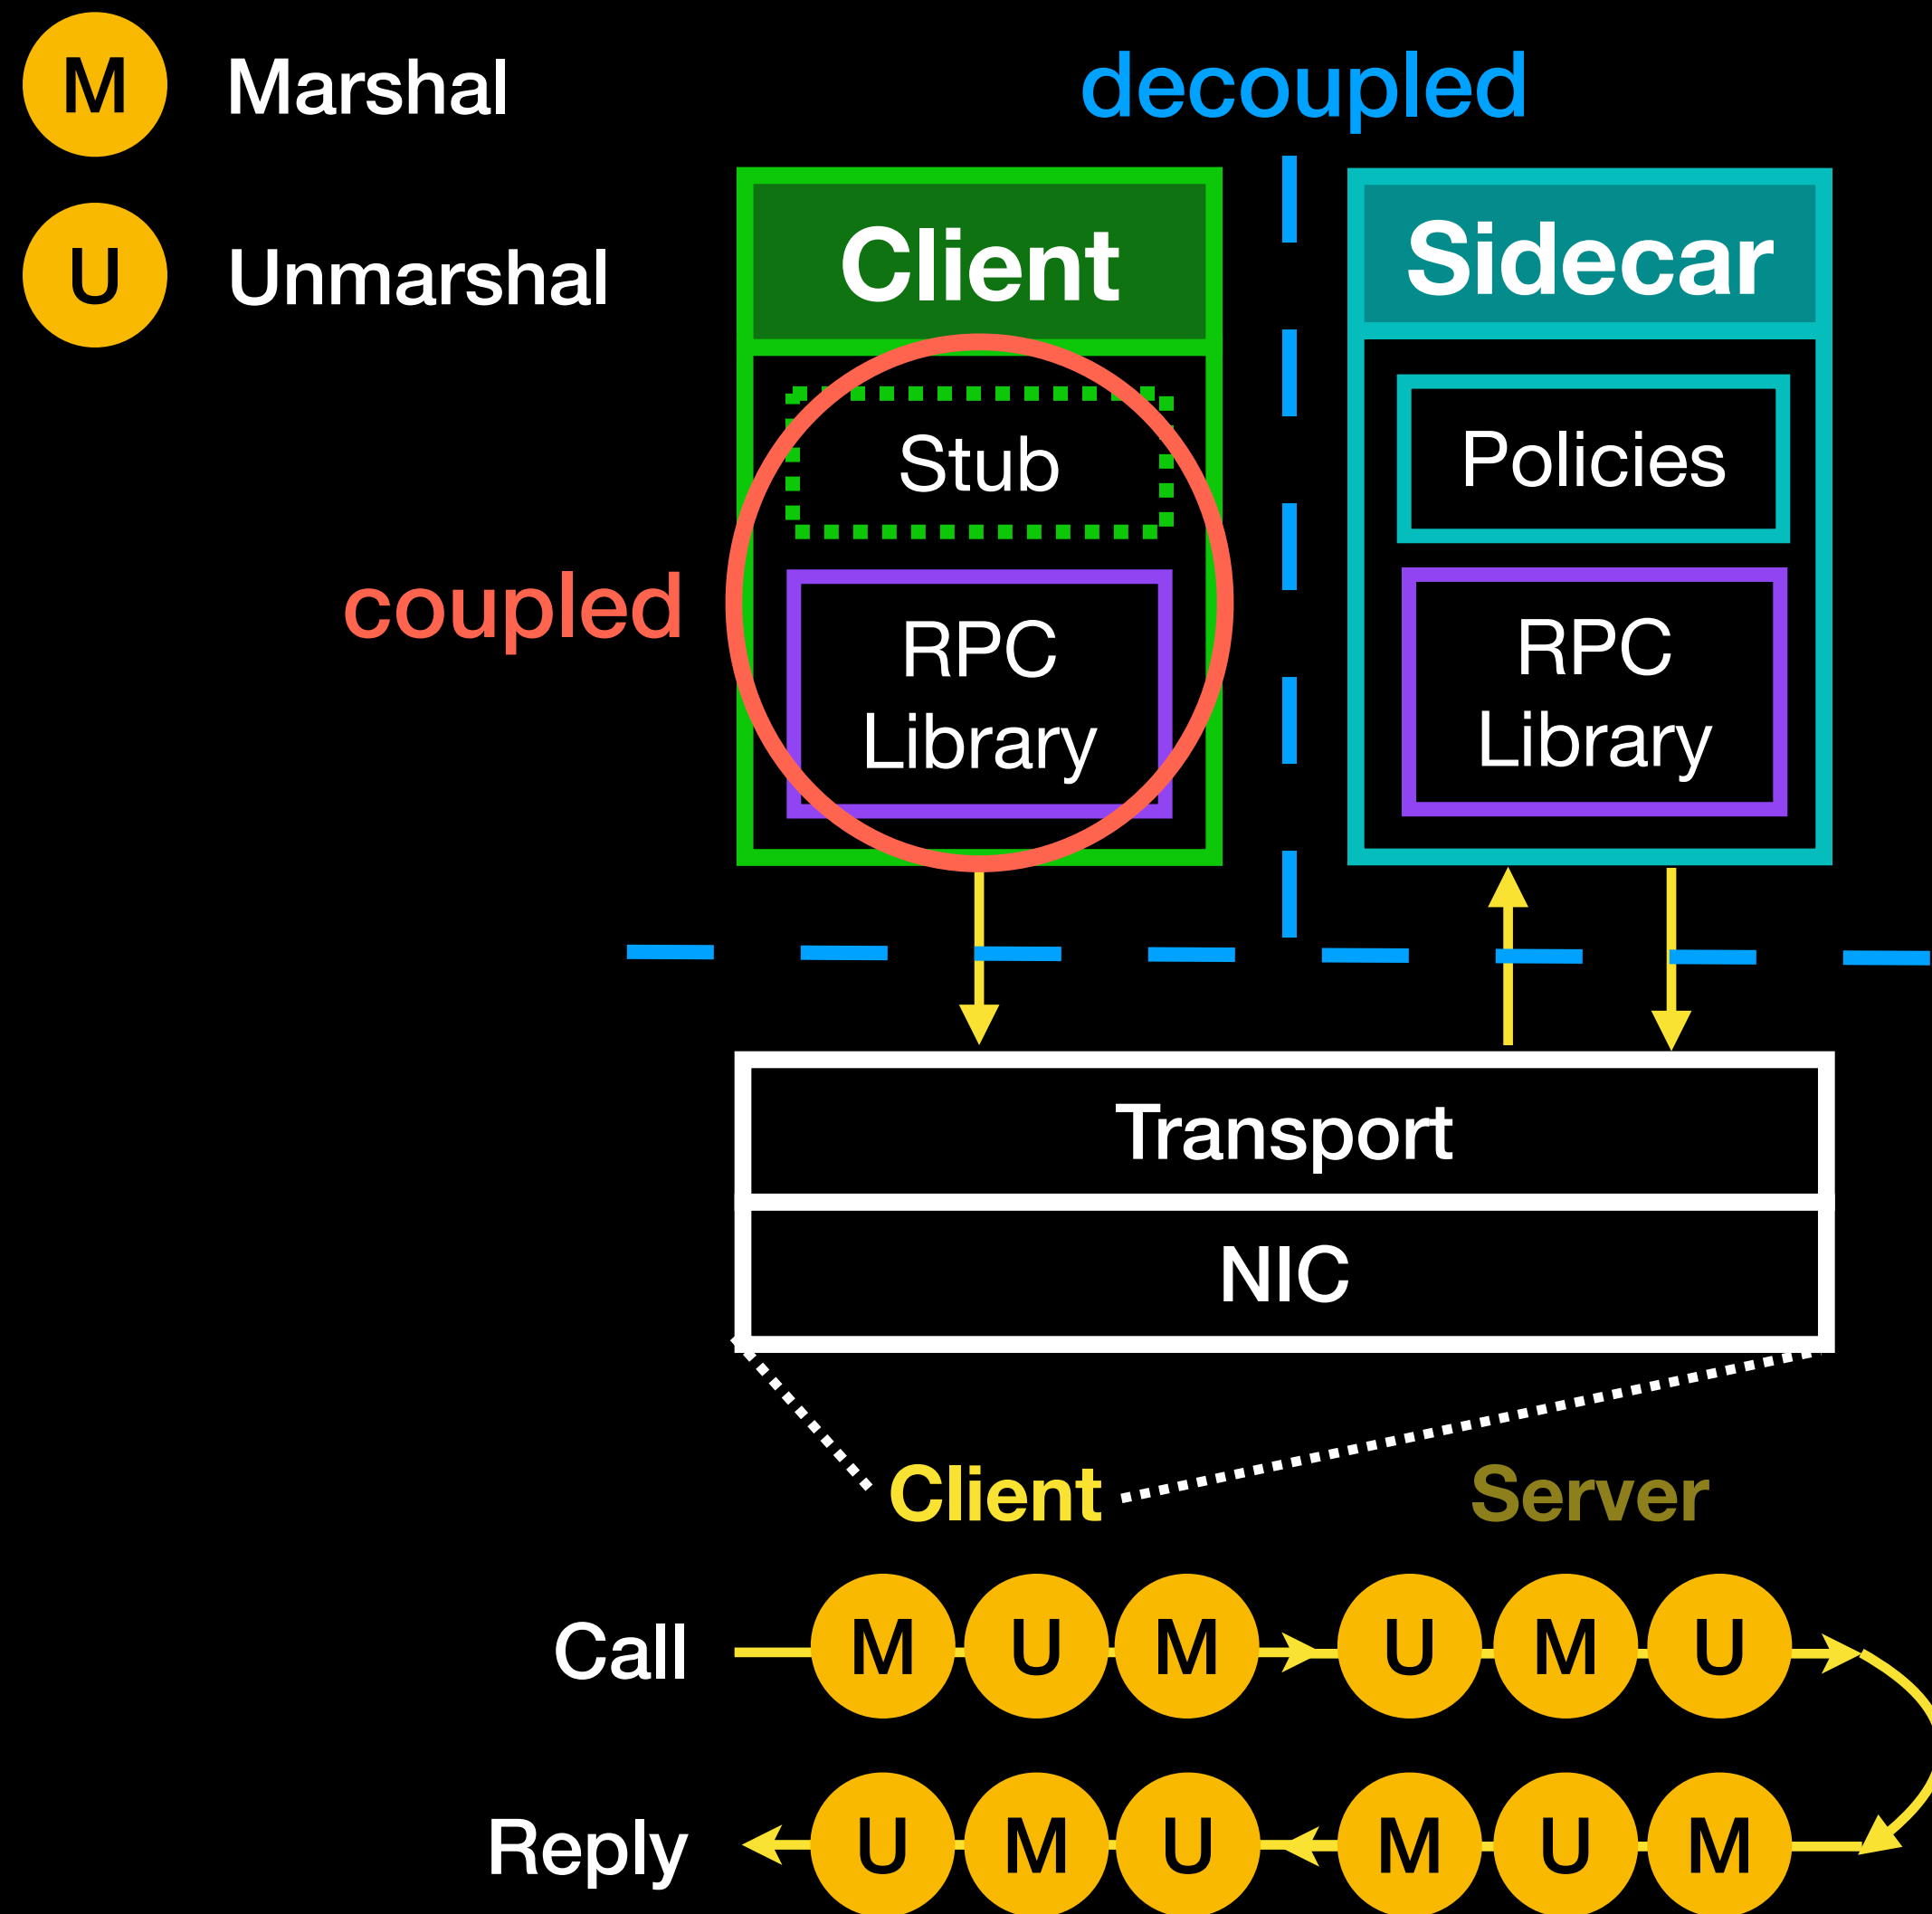
RPC-as-a-Library Limitation

- M** Marshal
- U** Unmarshal



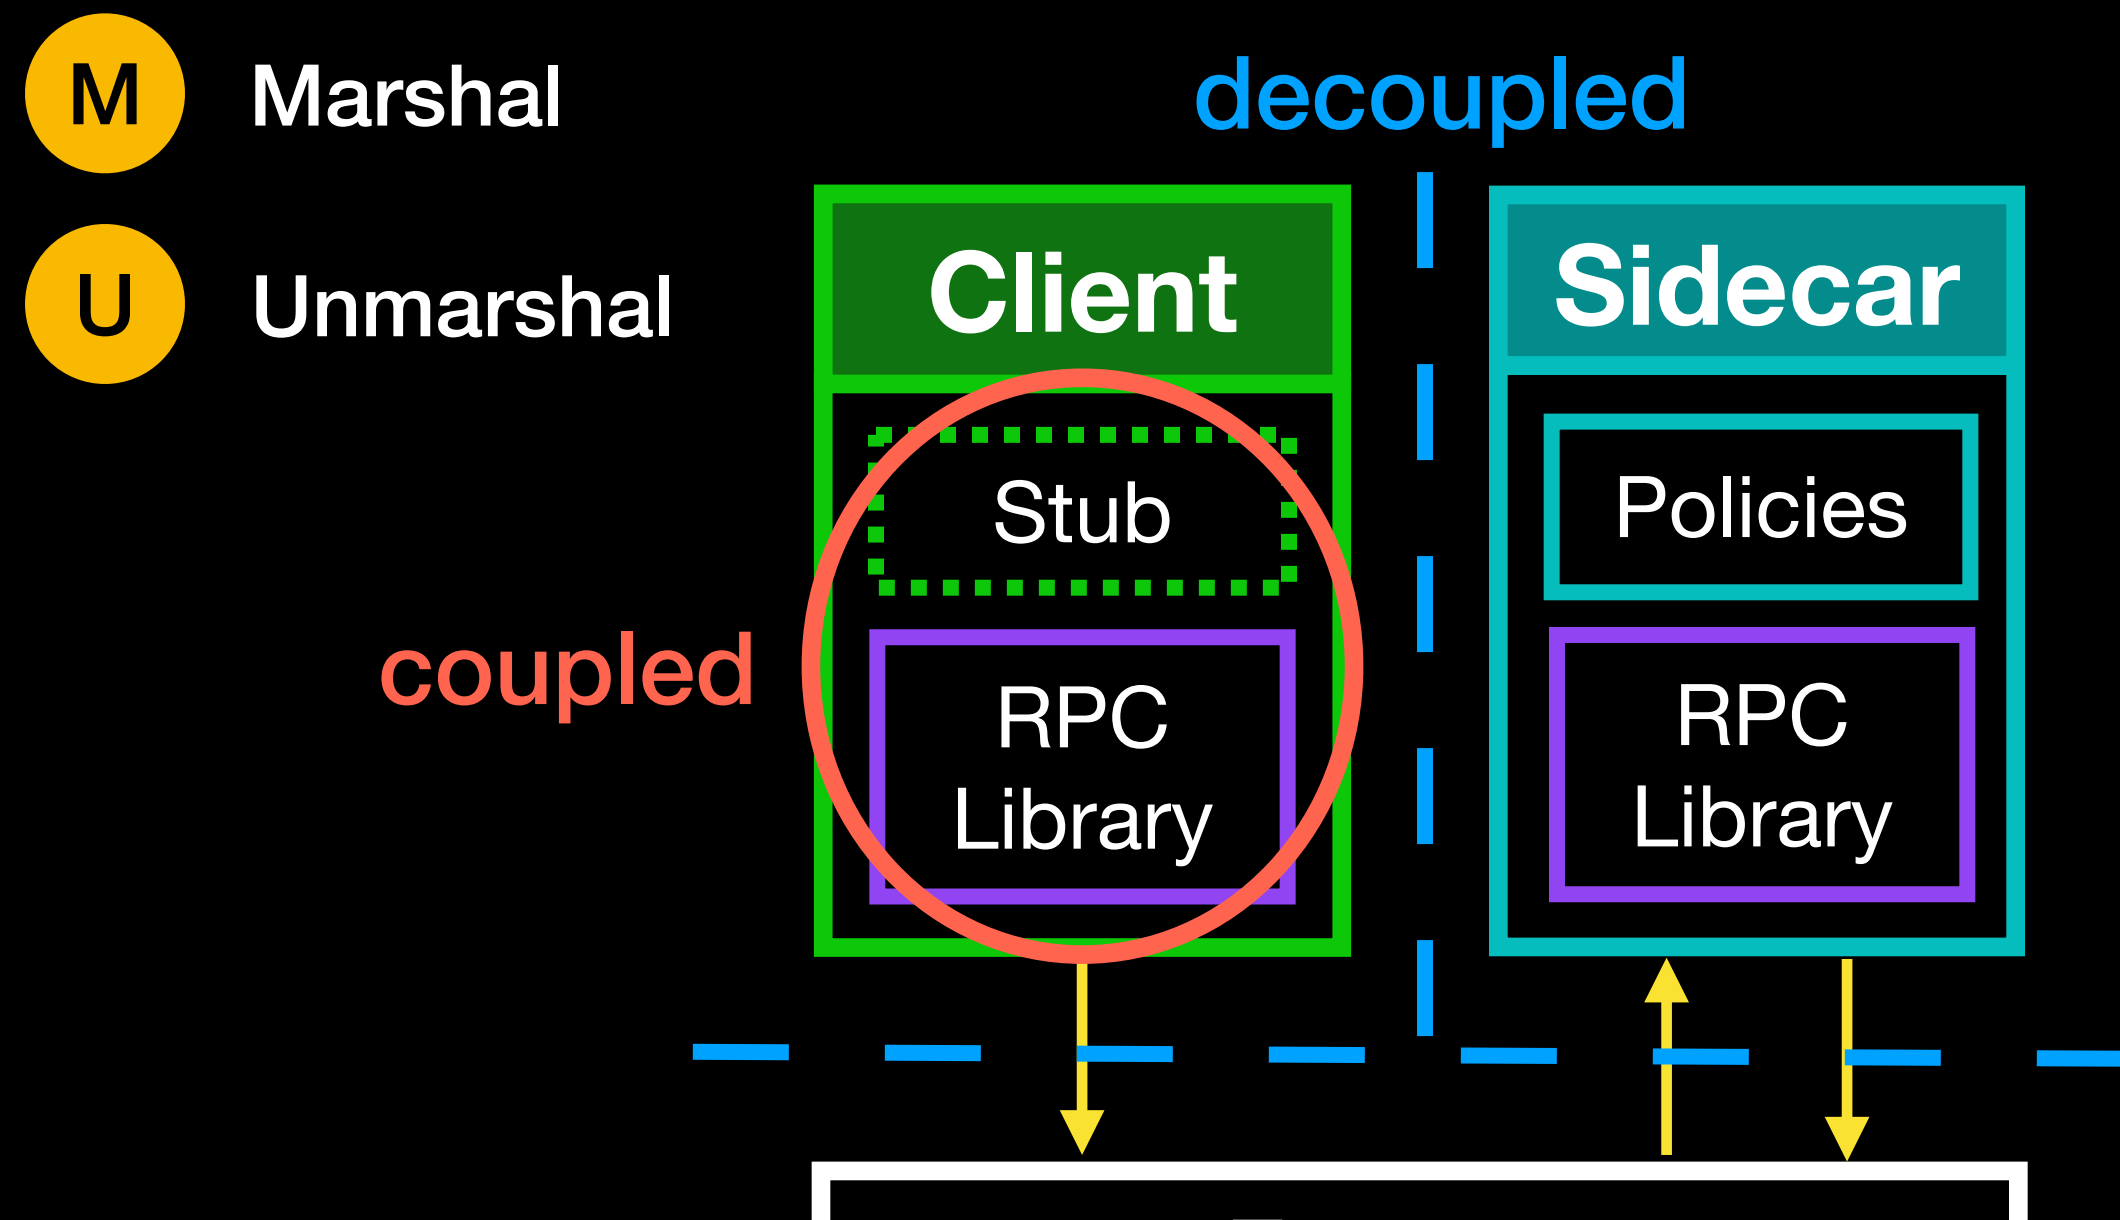
- RPC library and sidecar are **weakly coupled**
 - prevent from cross-layer optimization
 - operate/coupled at L4

RPC-as-a-Library Limitation



- RPC library and sidecar are **weakly coupled**
 - prevent from cross-layer optimization
 - operate/coupled at L4
- RPC Library and app are **strongly coupled**
 - Difficult to upgrade RPC library

RPC-as-a-Library Limitation



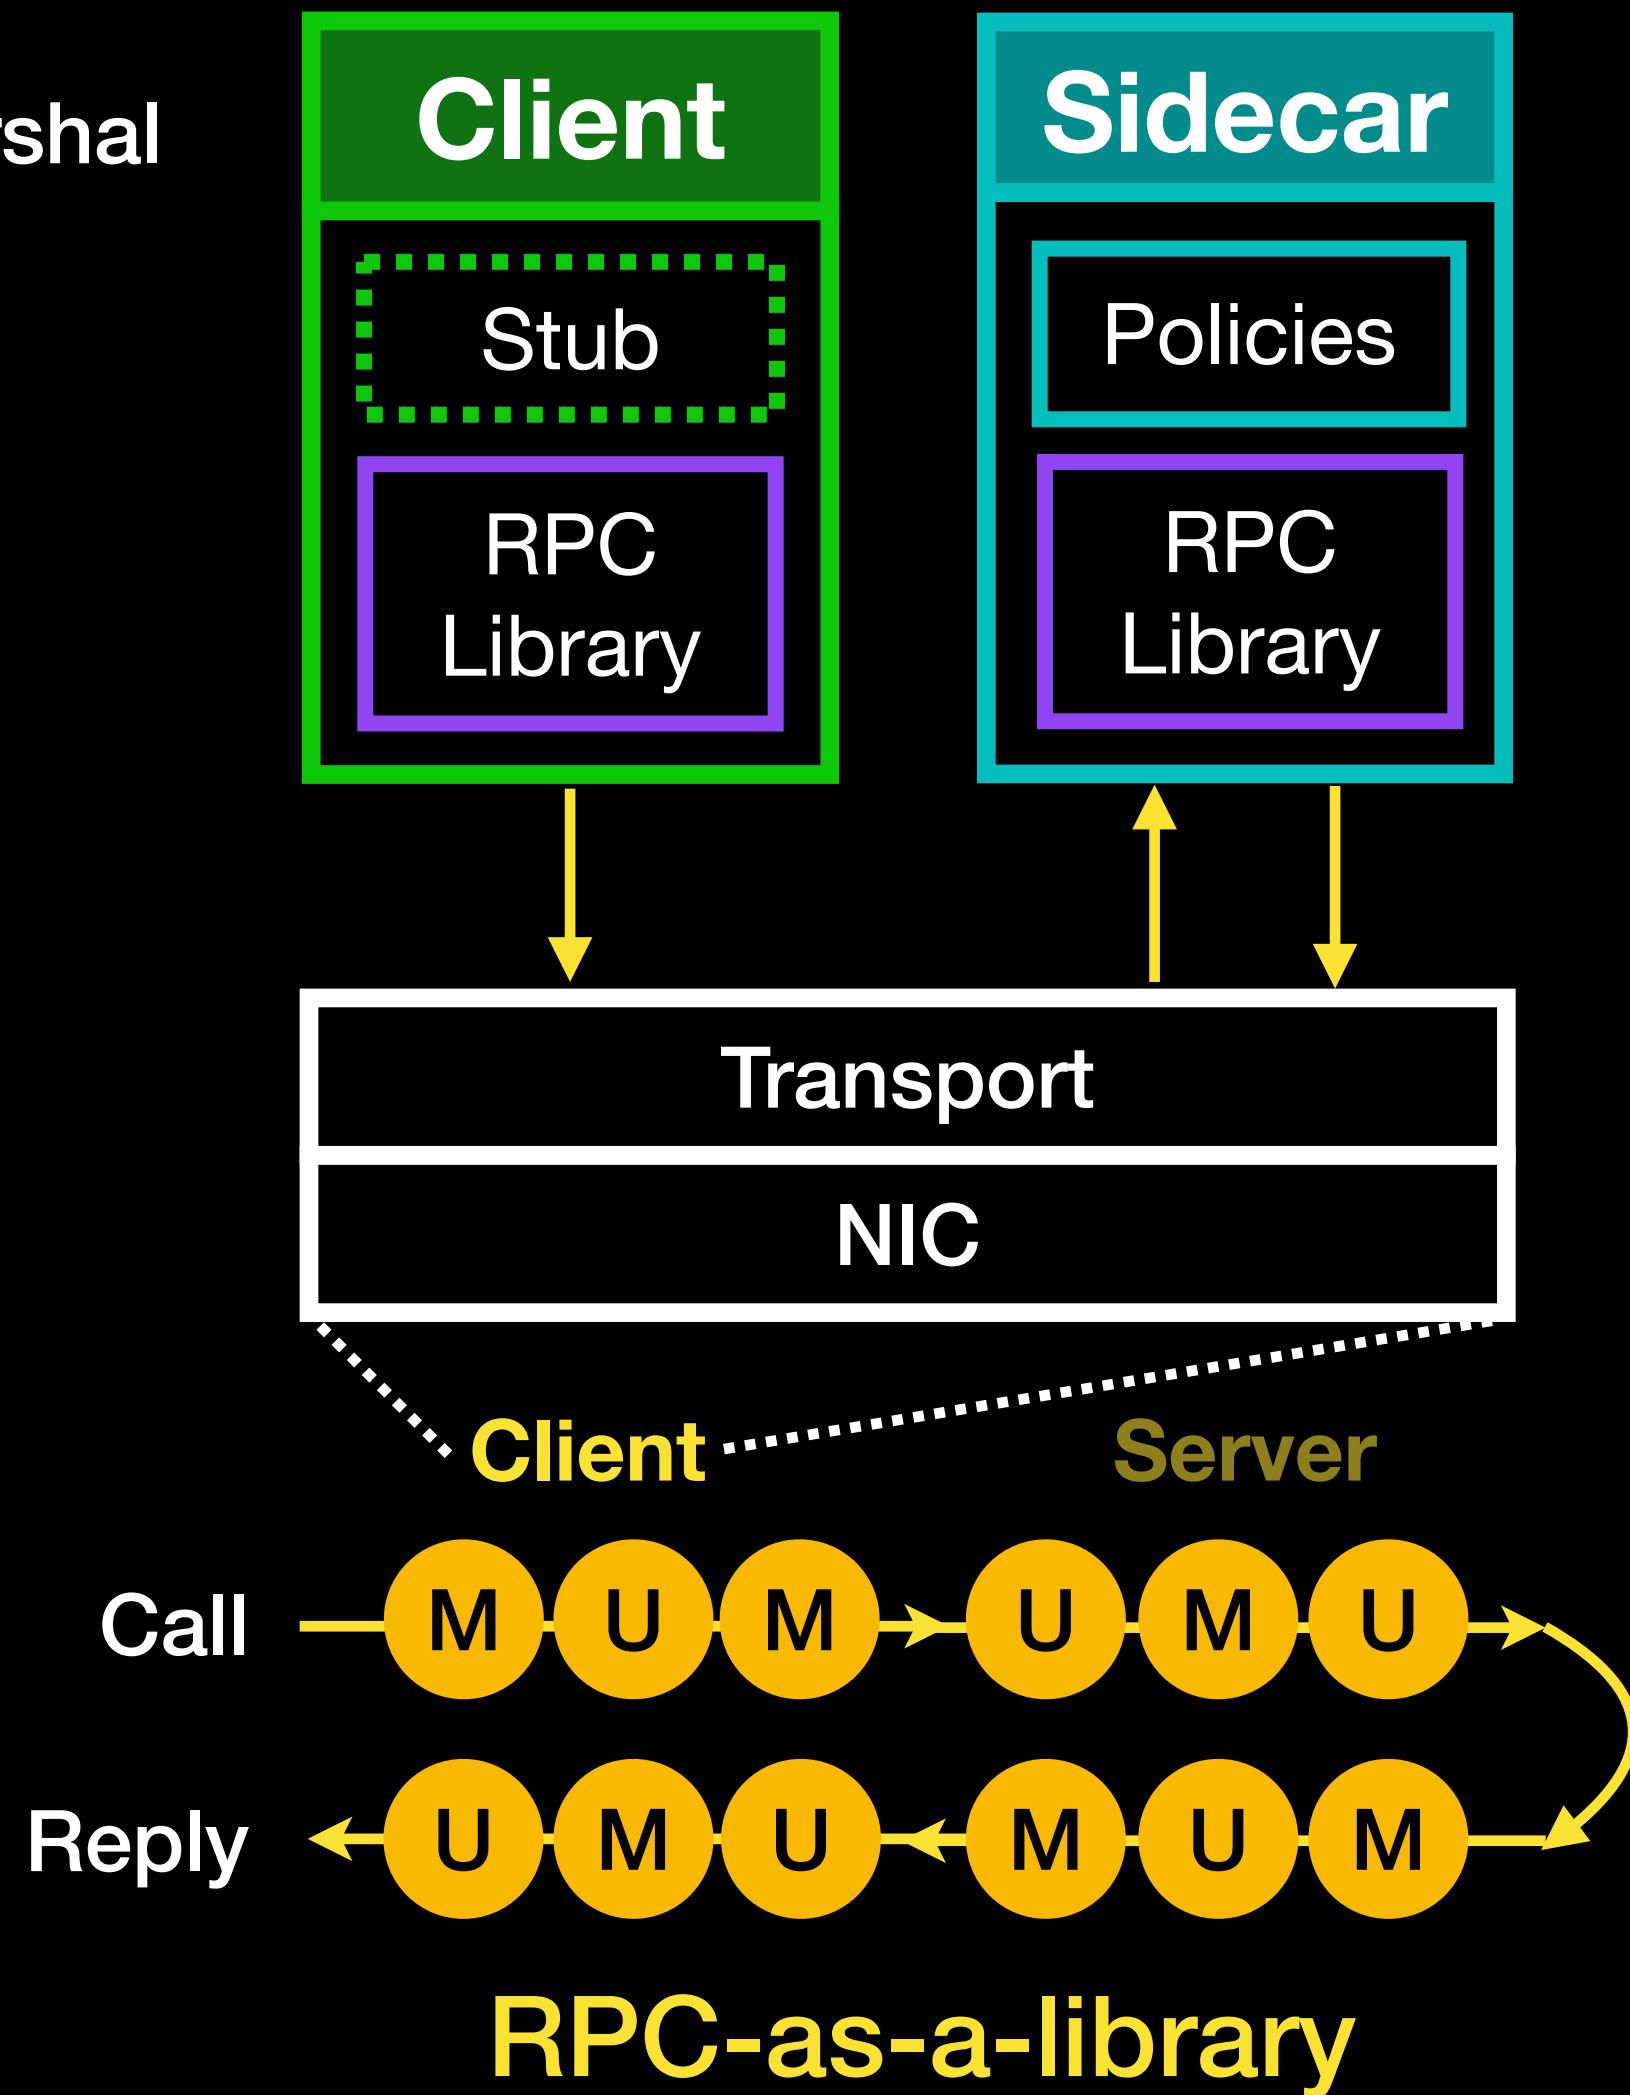
- RPC library and sidecar are **weakly coupled**
 - prevent from cross-layer optimization
 - operate/coupled at L4
- RPC Library and app are **strongly coupled**
 - Difficult to upgrade RPC library

We want

- strong coupling: operate at L7
- weak coupling: most of the functionalities extracted into a **separate service**

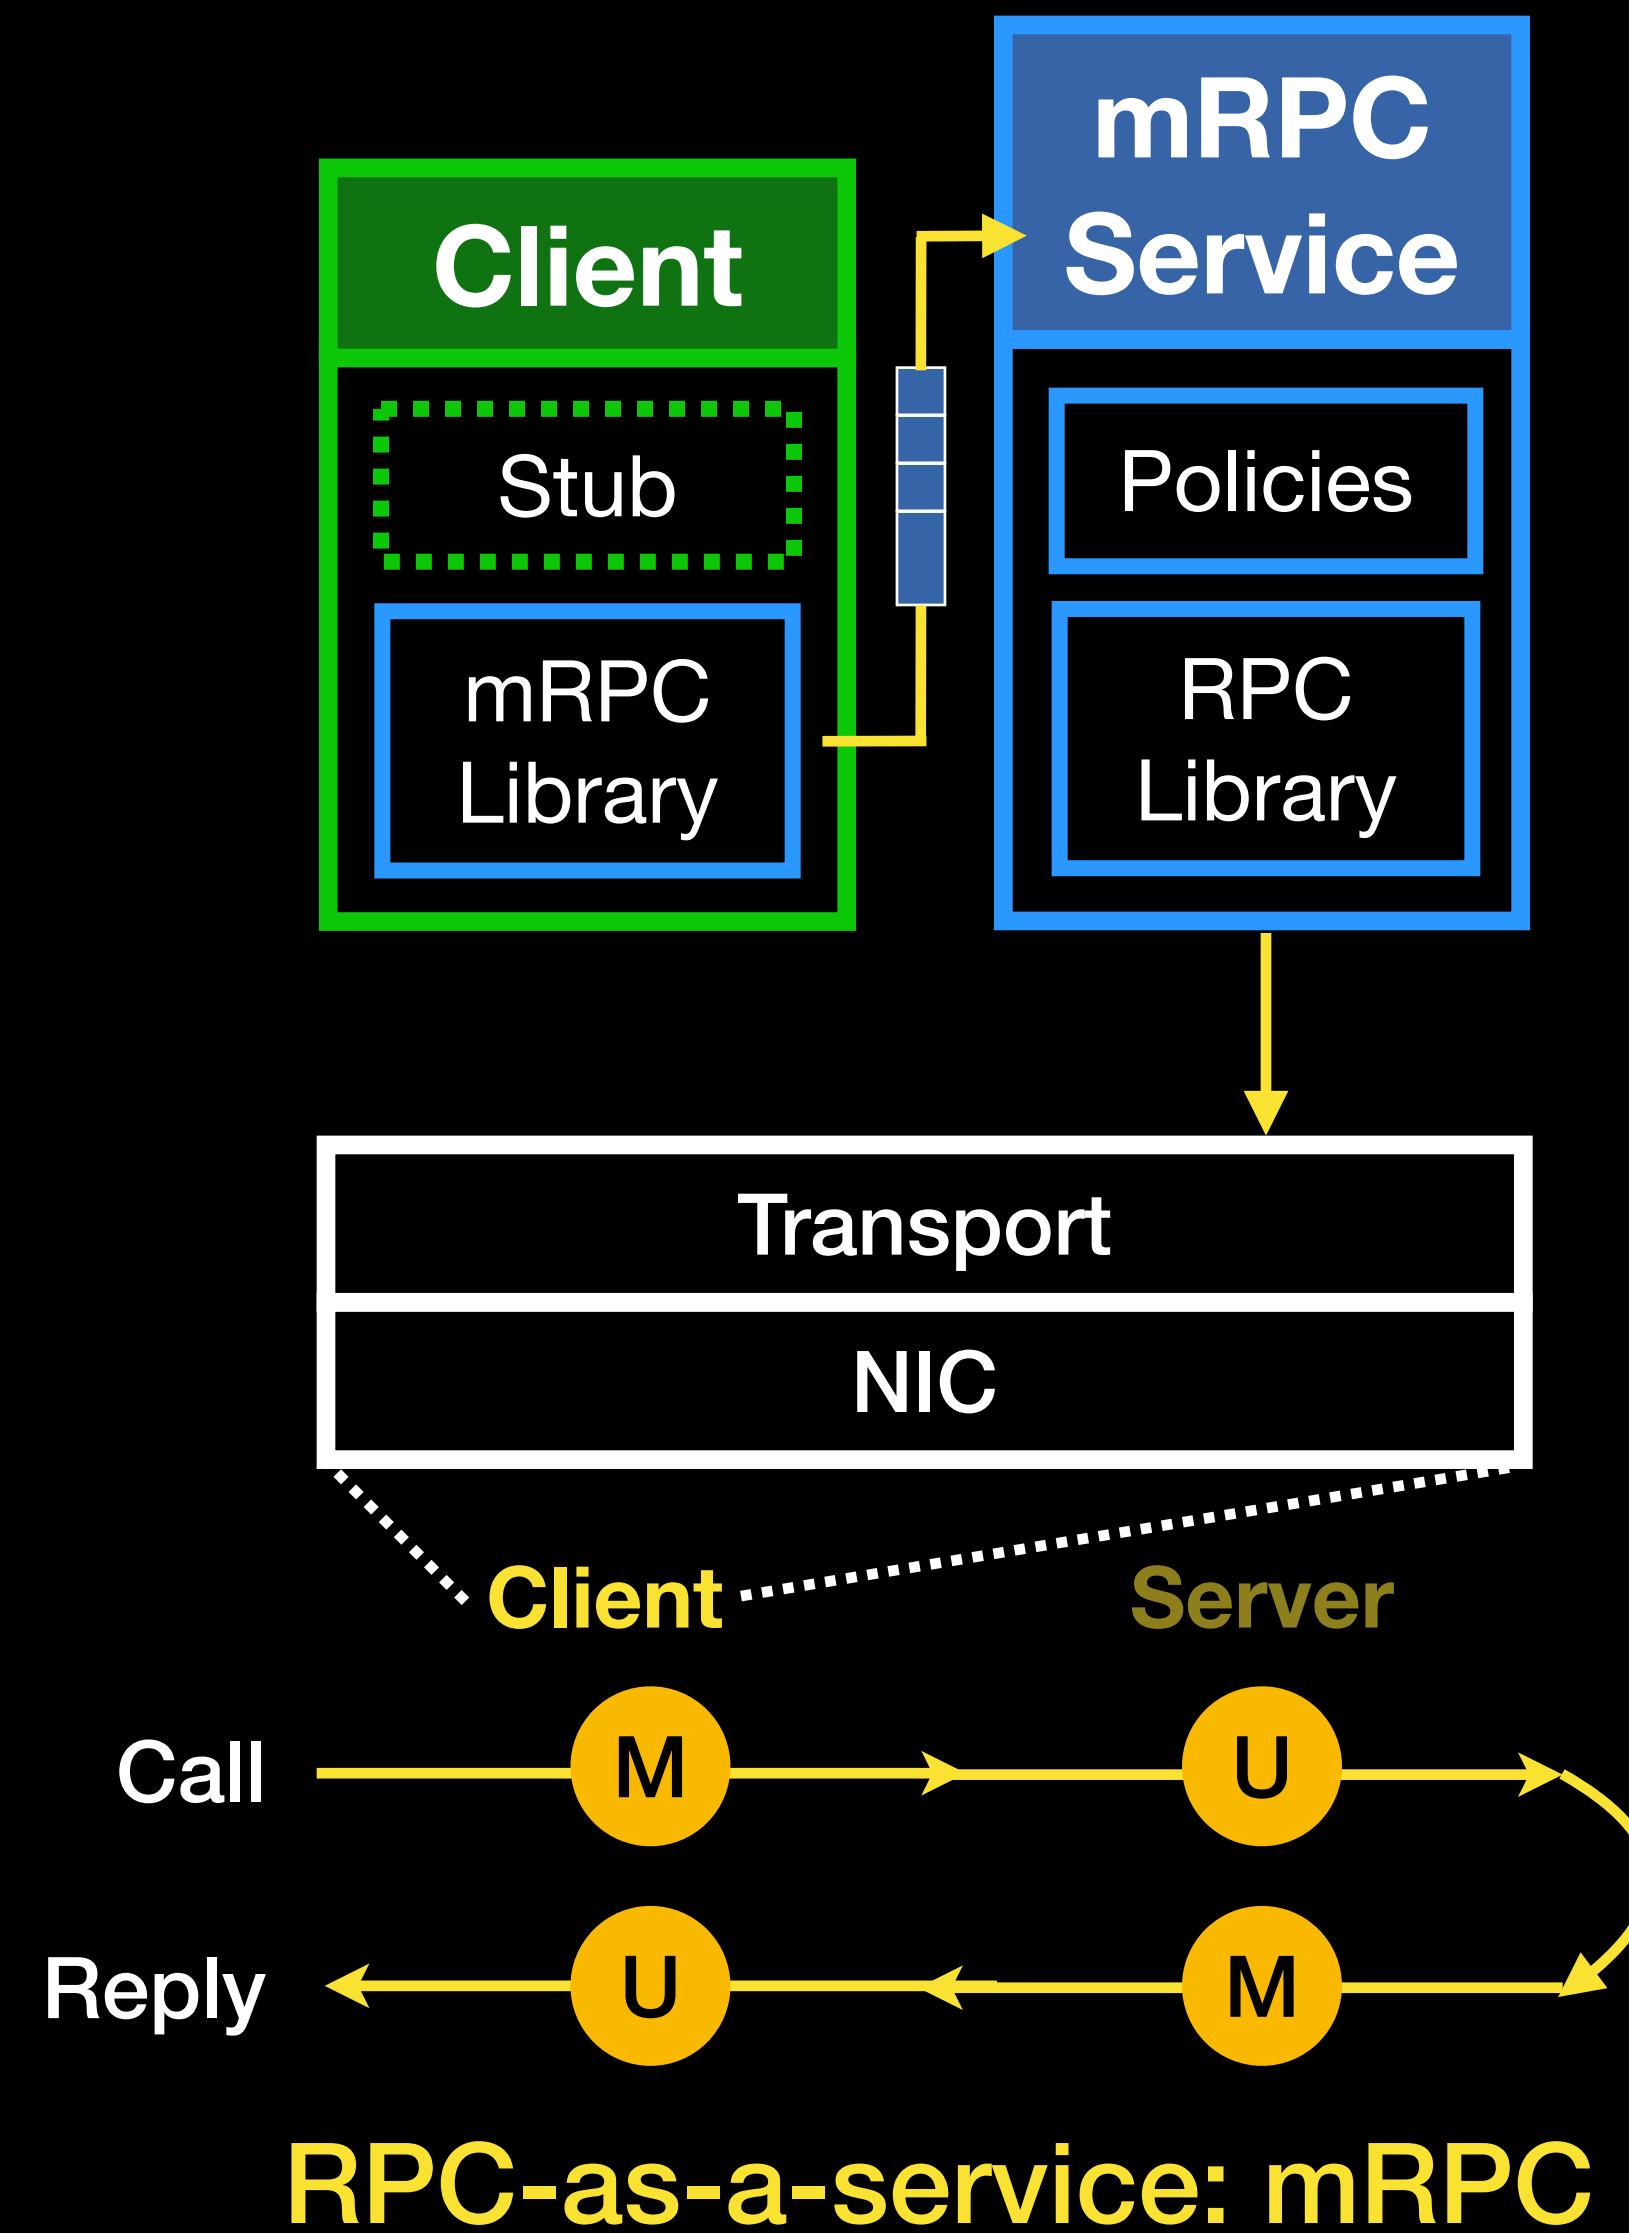

mRPC Overview

- M** Marshal
- U** Unmarshal



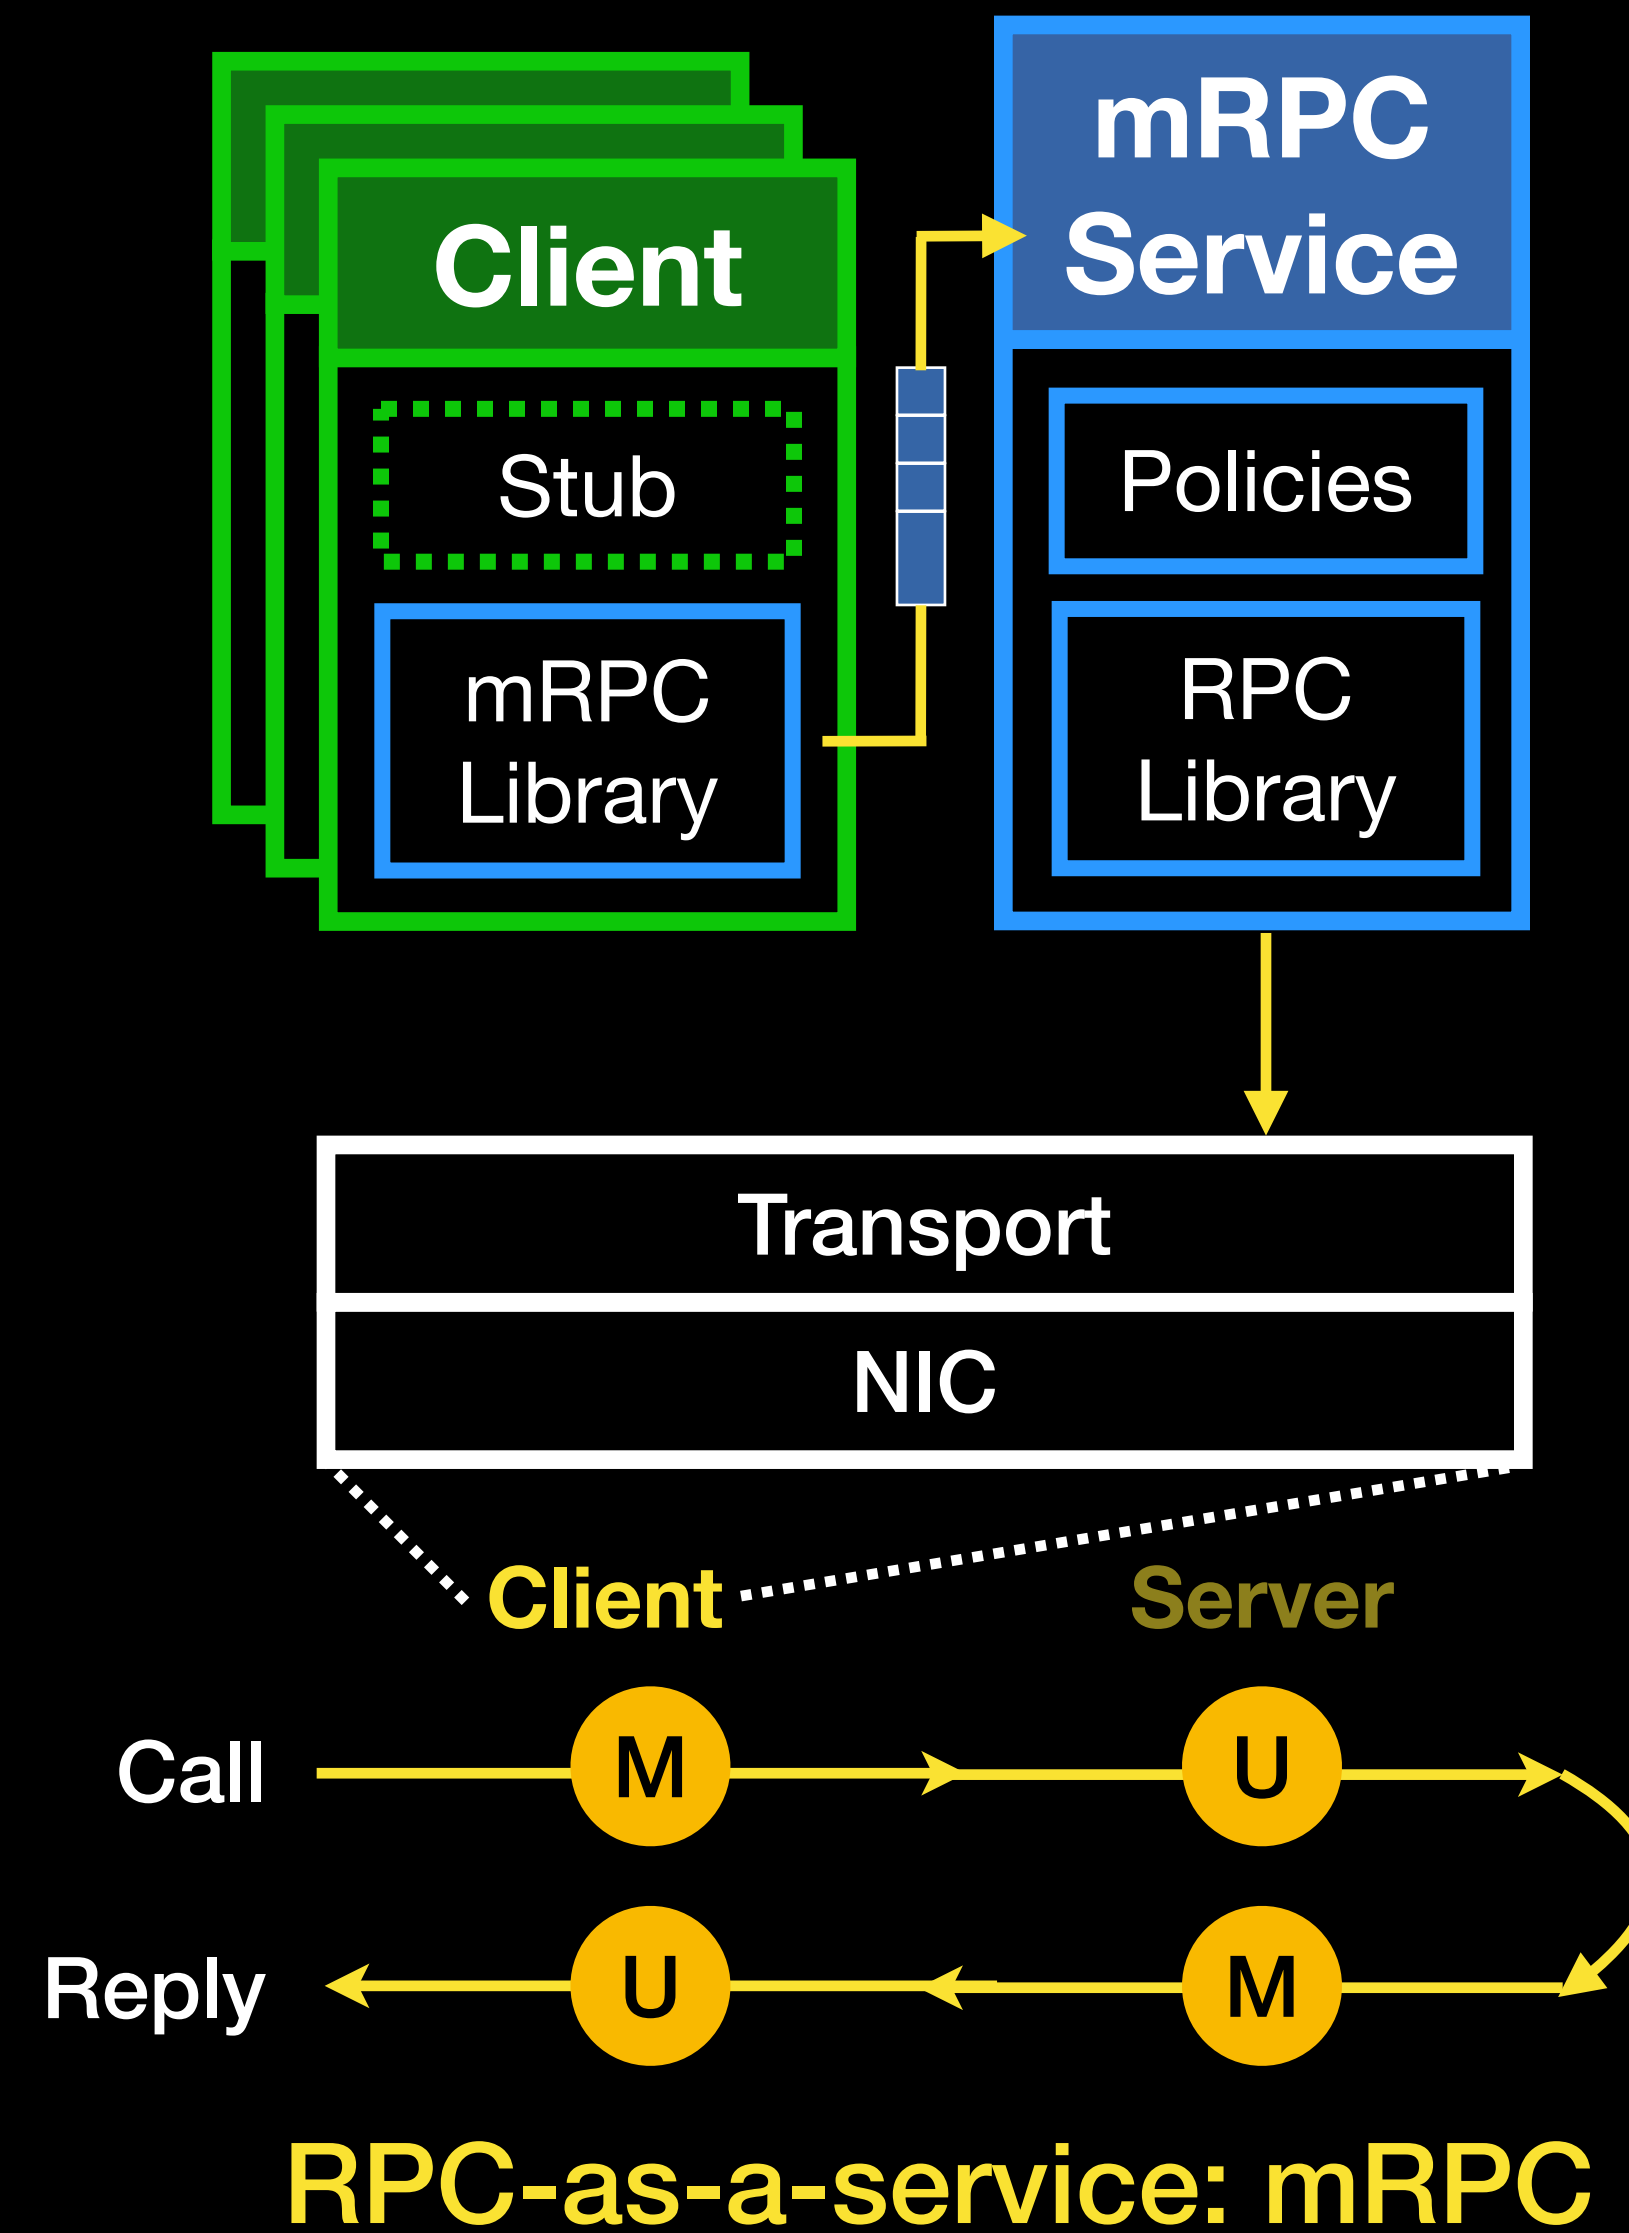
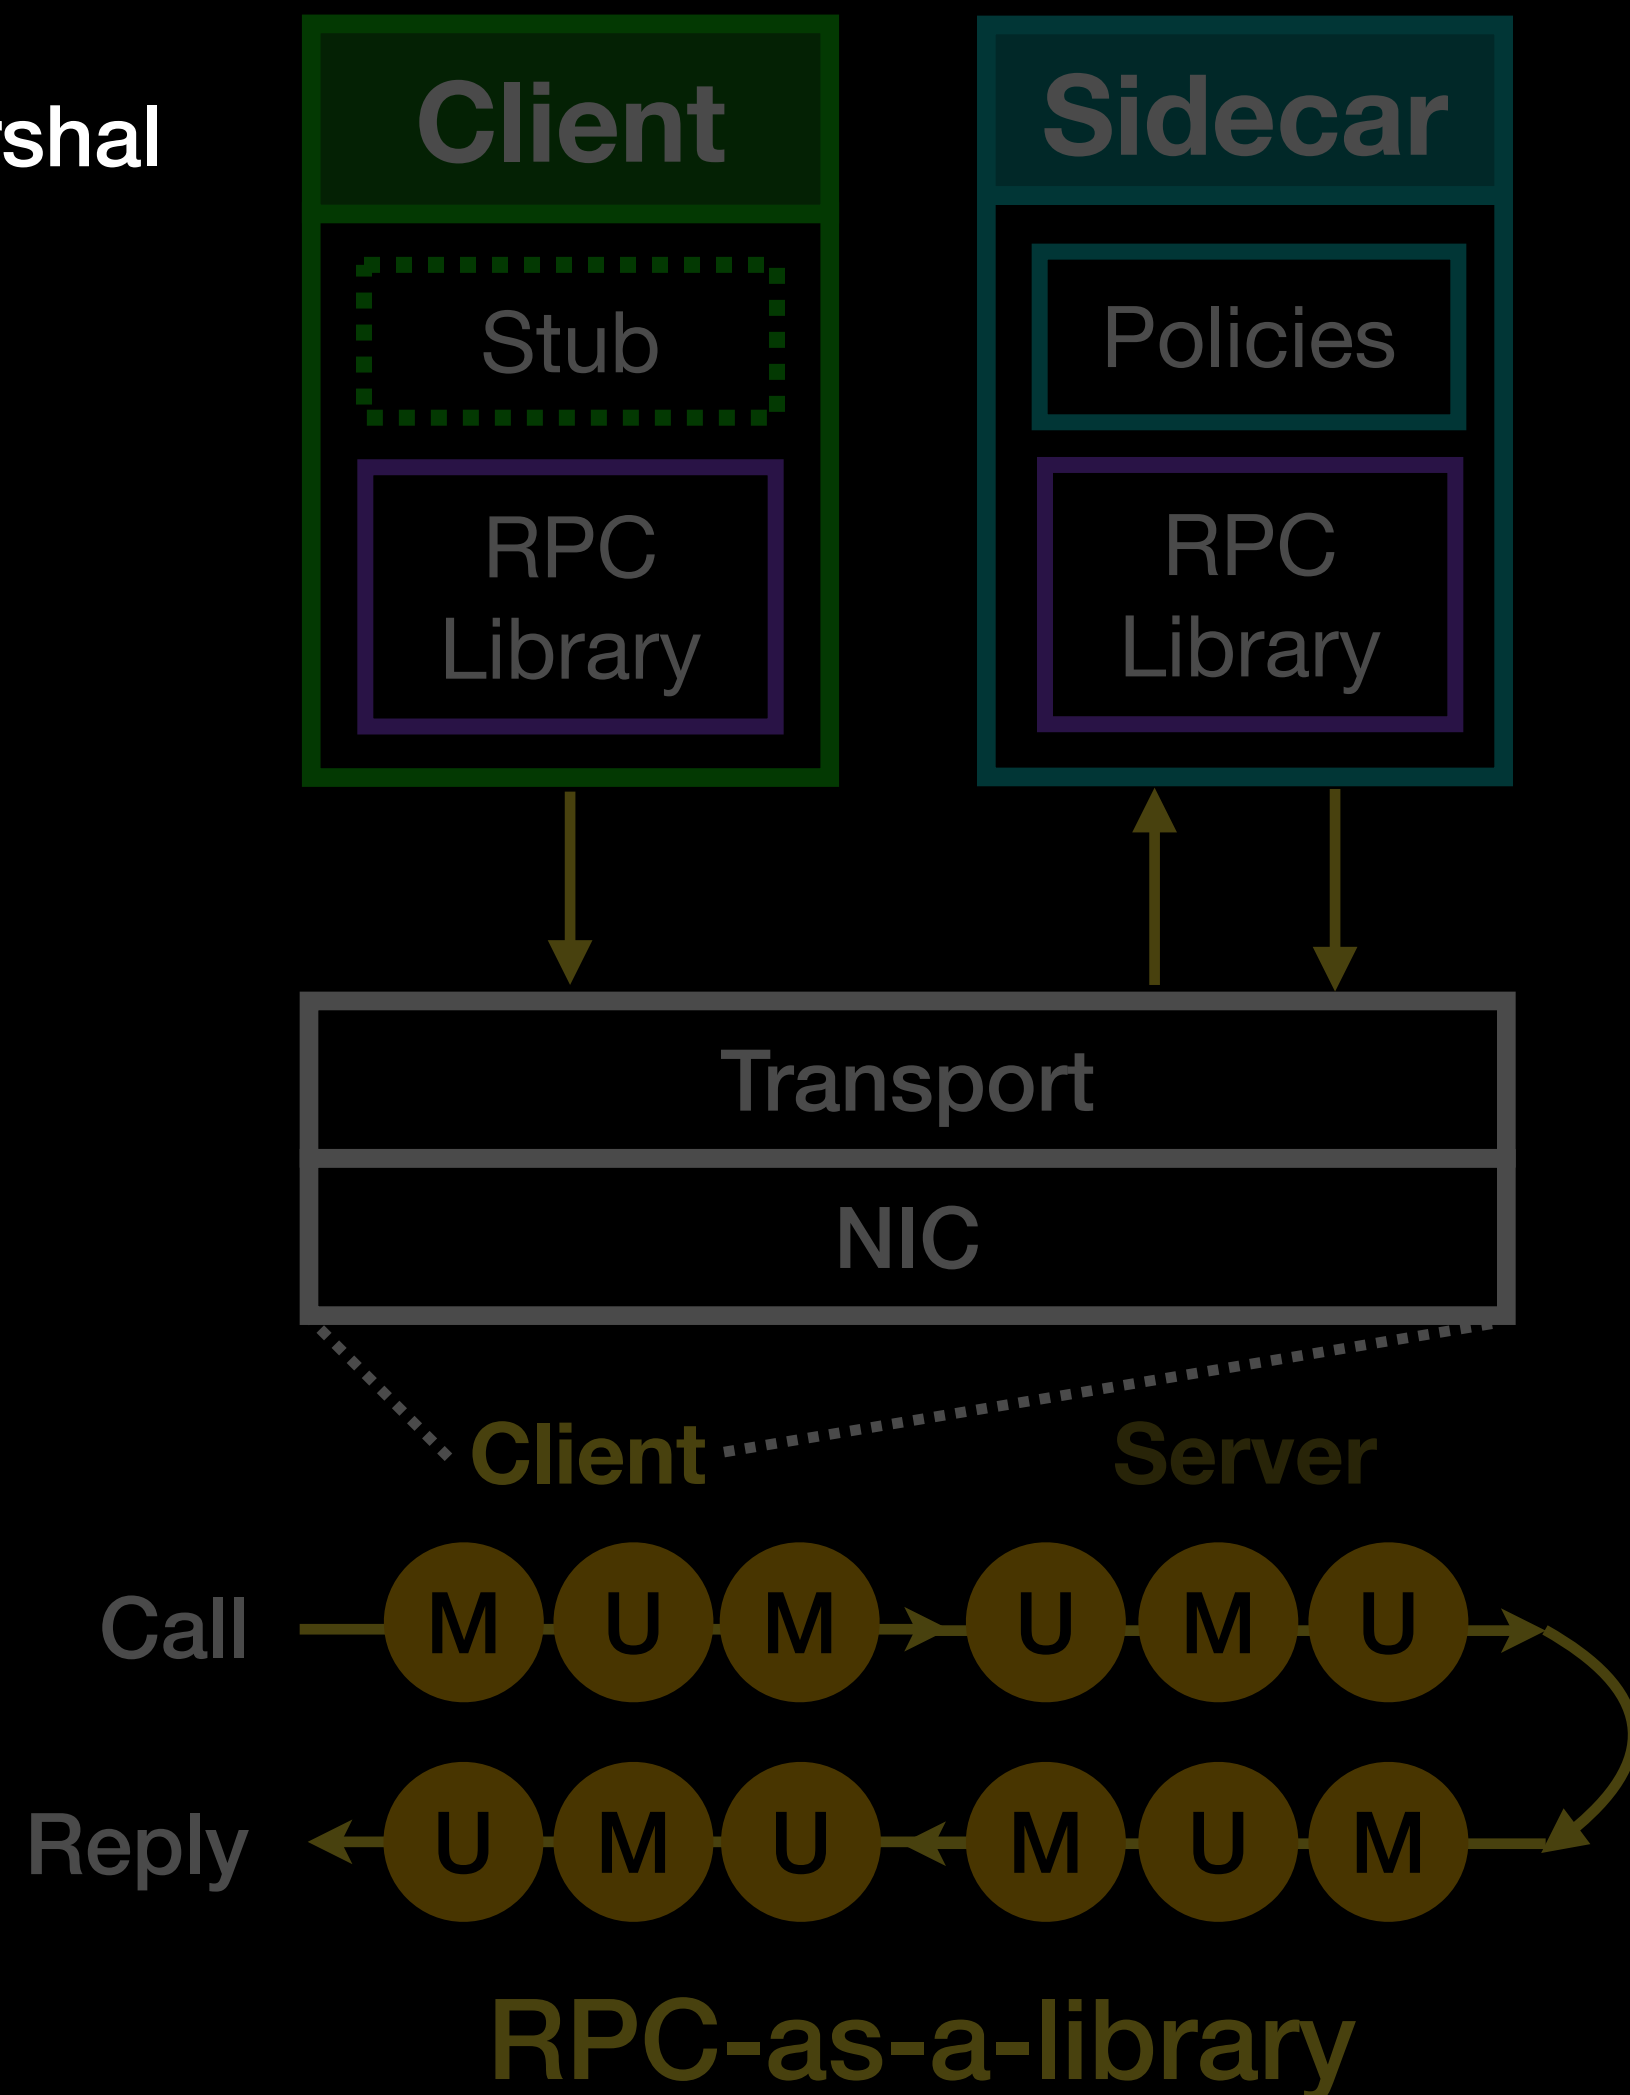
mRPC Overview

- M** Marshal
- U** Unmarshal



mRPC Overview

- M** Marshal
- U** Unmarshal



Challenges

How to support new applications with new RPC specifications at runtime?

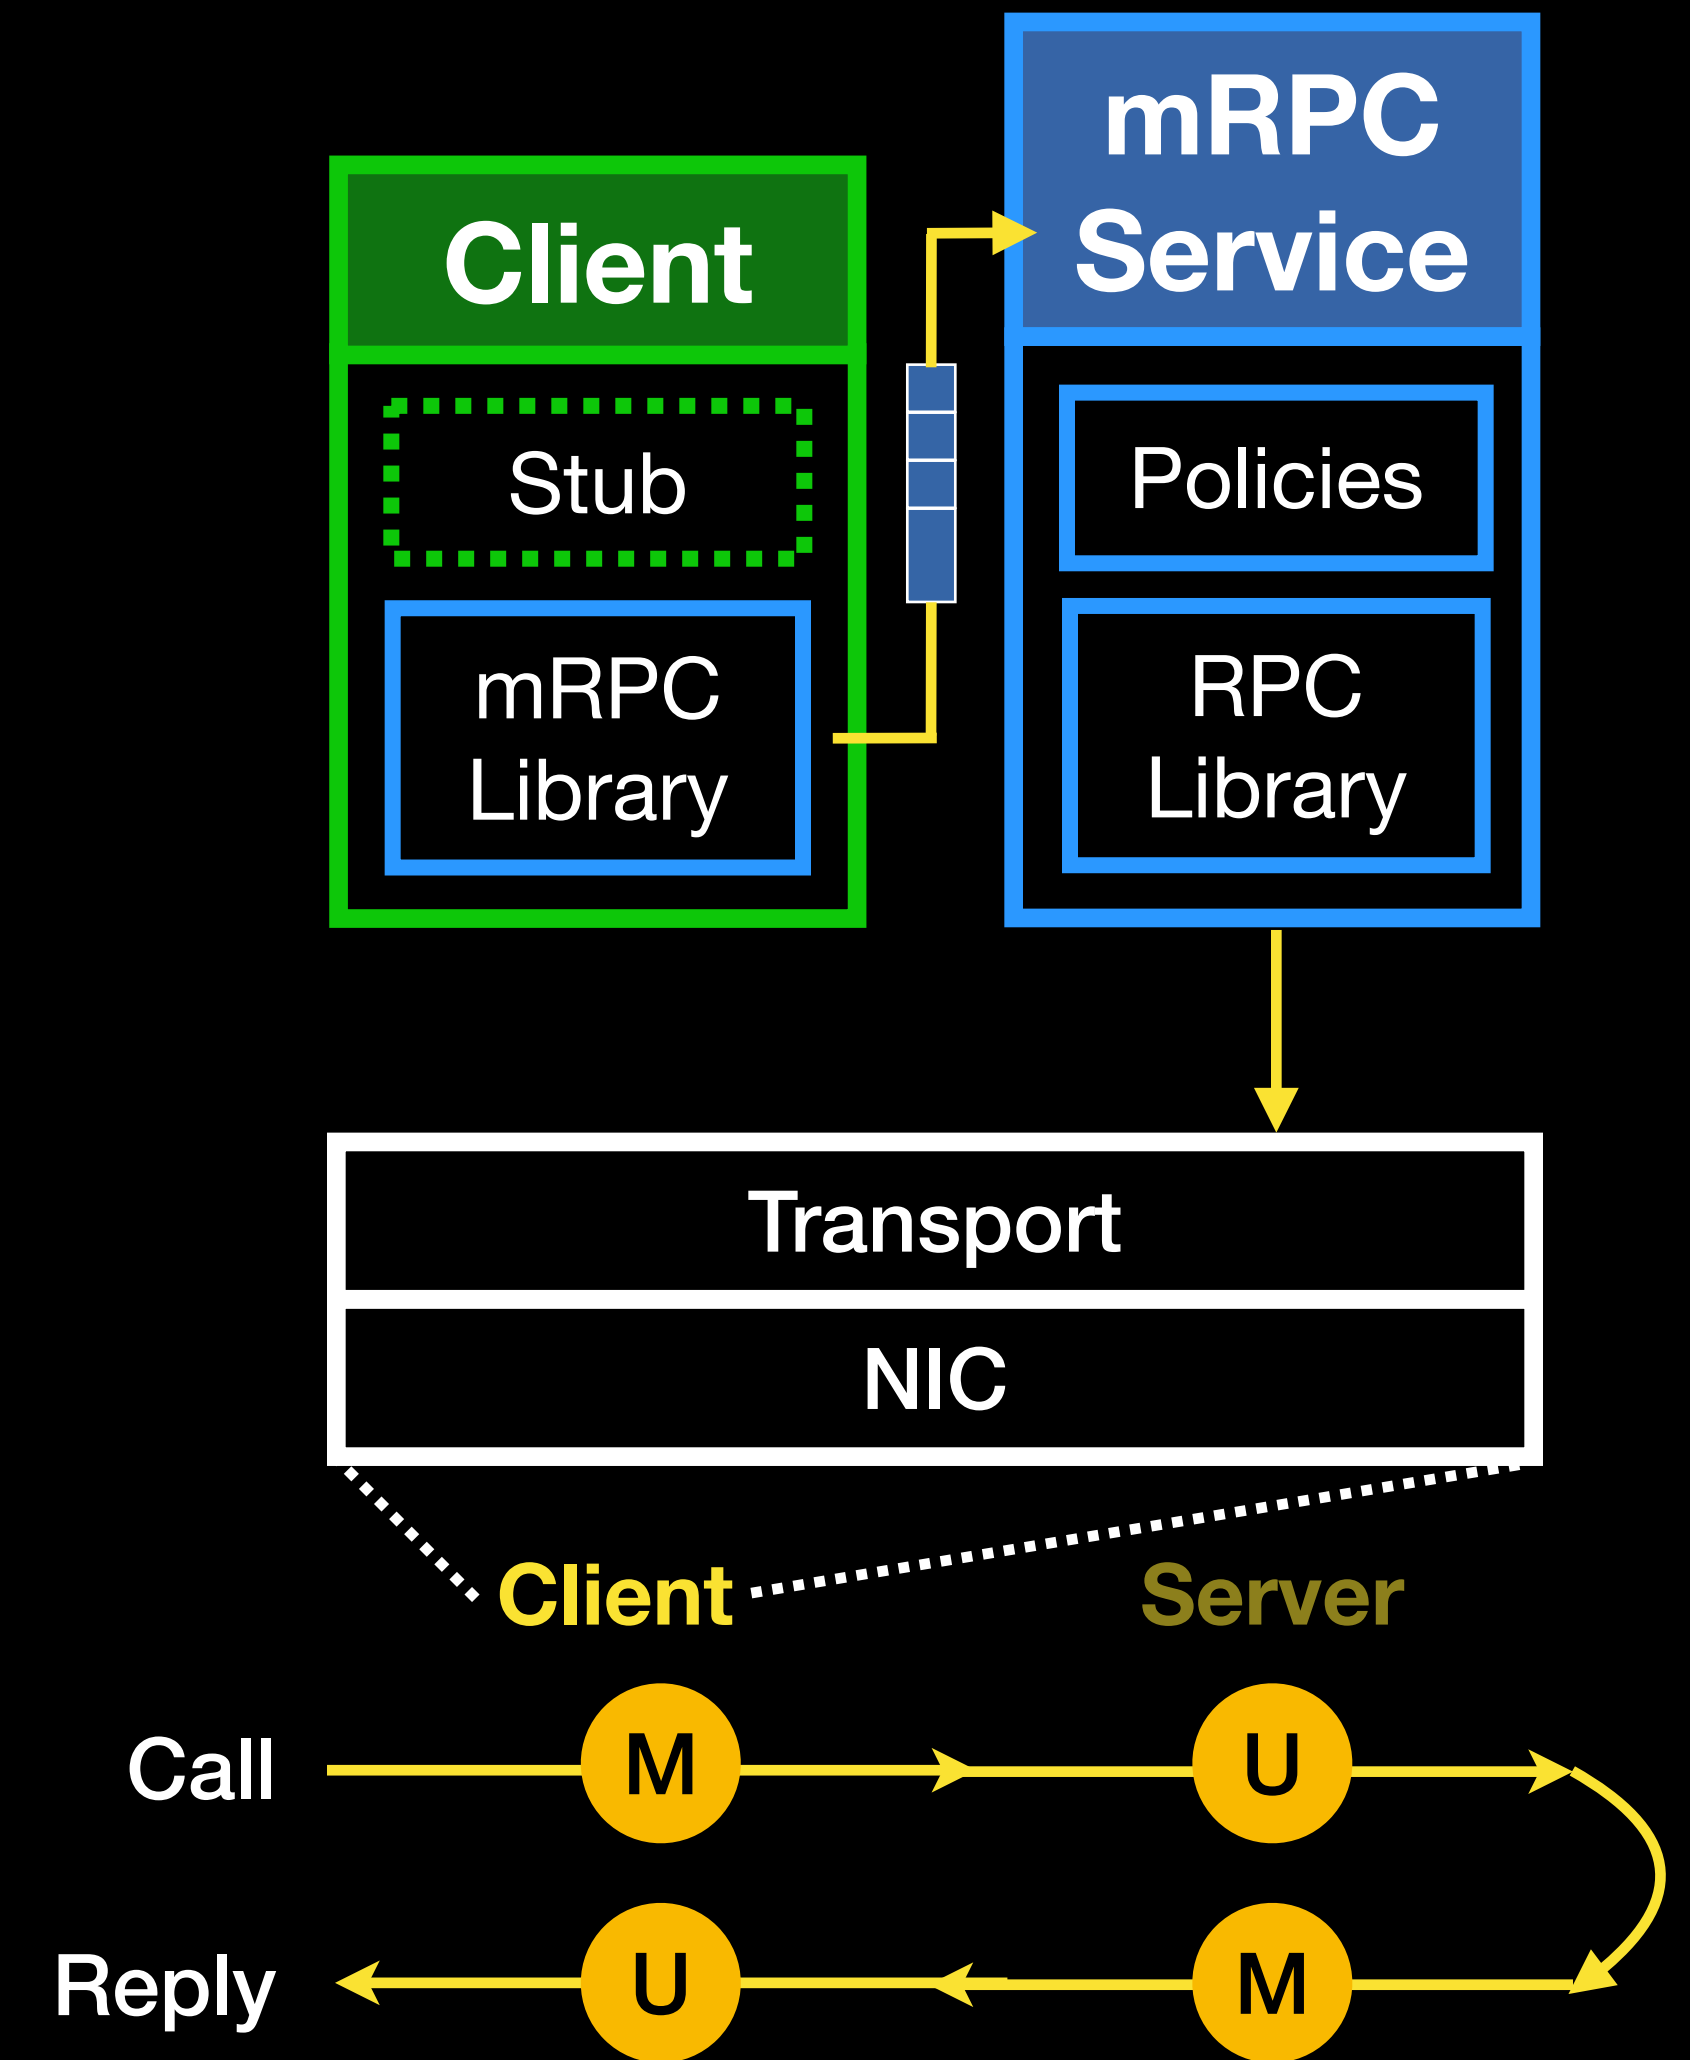
How to enforce policies with efficiency and security?

How to live upgrade RPC implementations without disrupting other applications?

Dynamic Binding

Memory Management

Live Upgrade



RPC-as-a-service: mRPC

Challenges

How to support new applications with new RPC specifications at runtime?

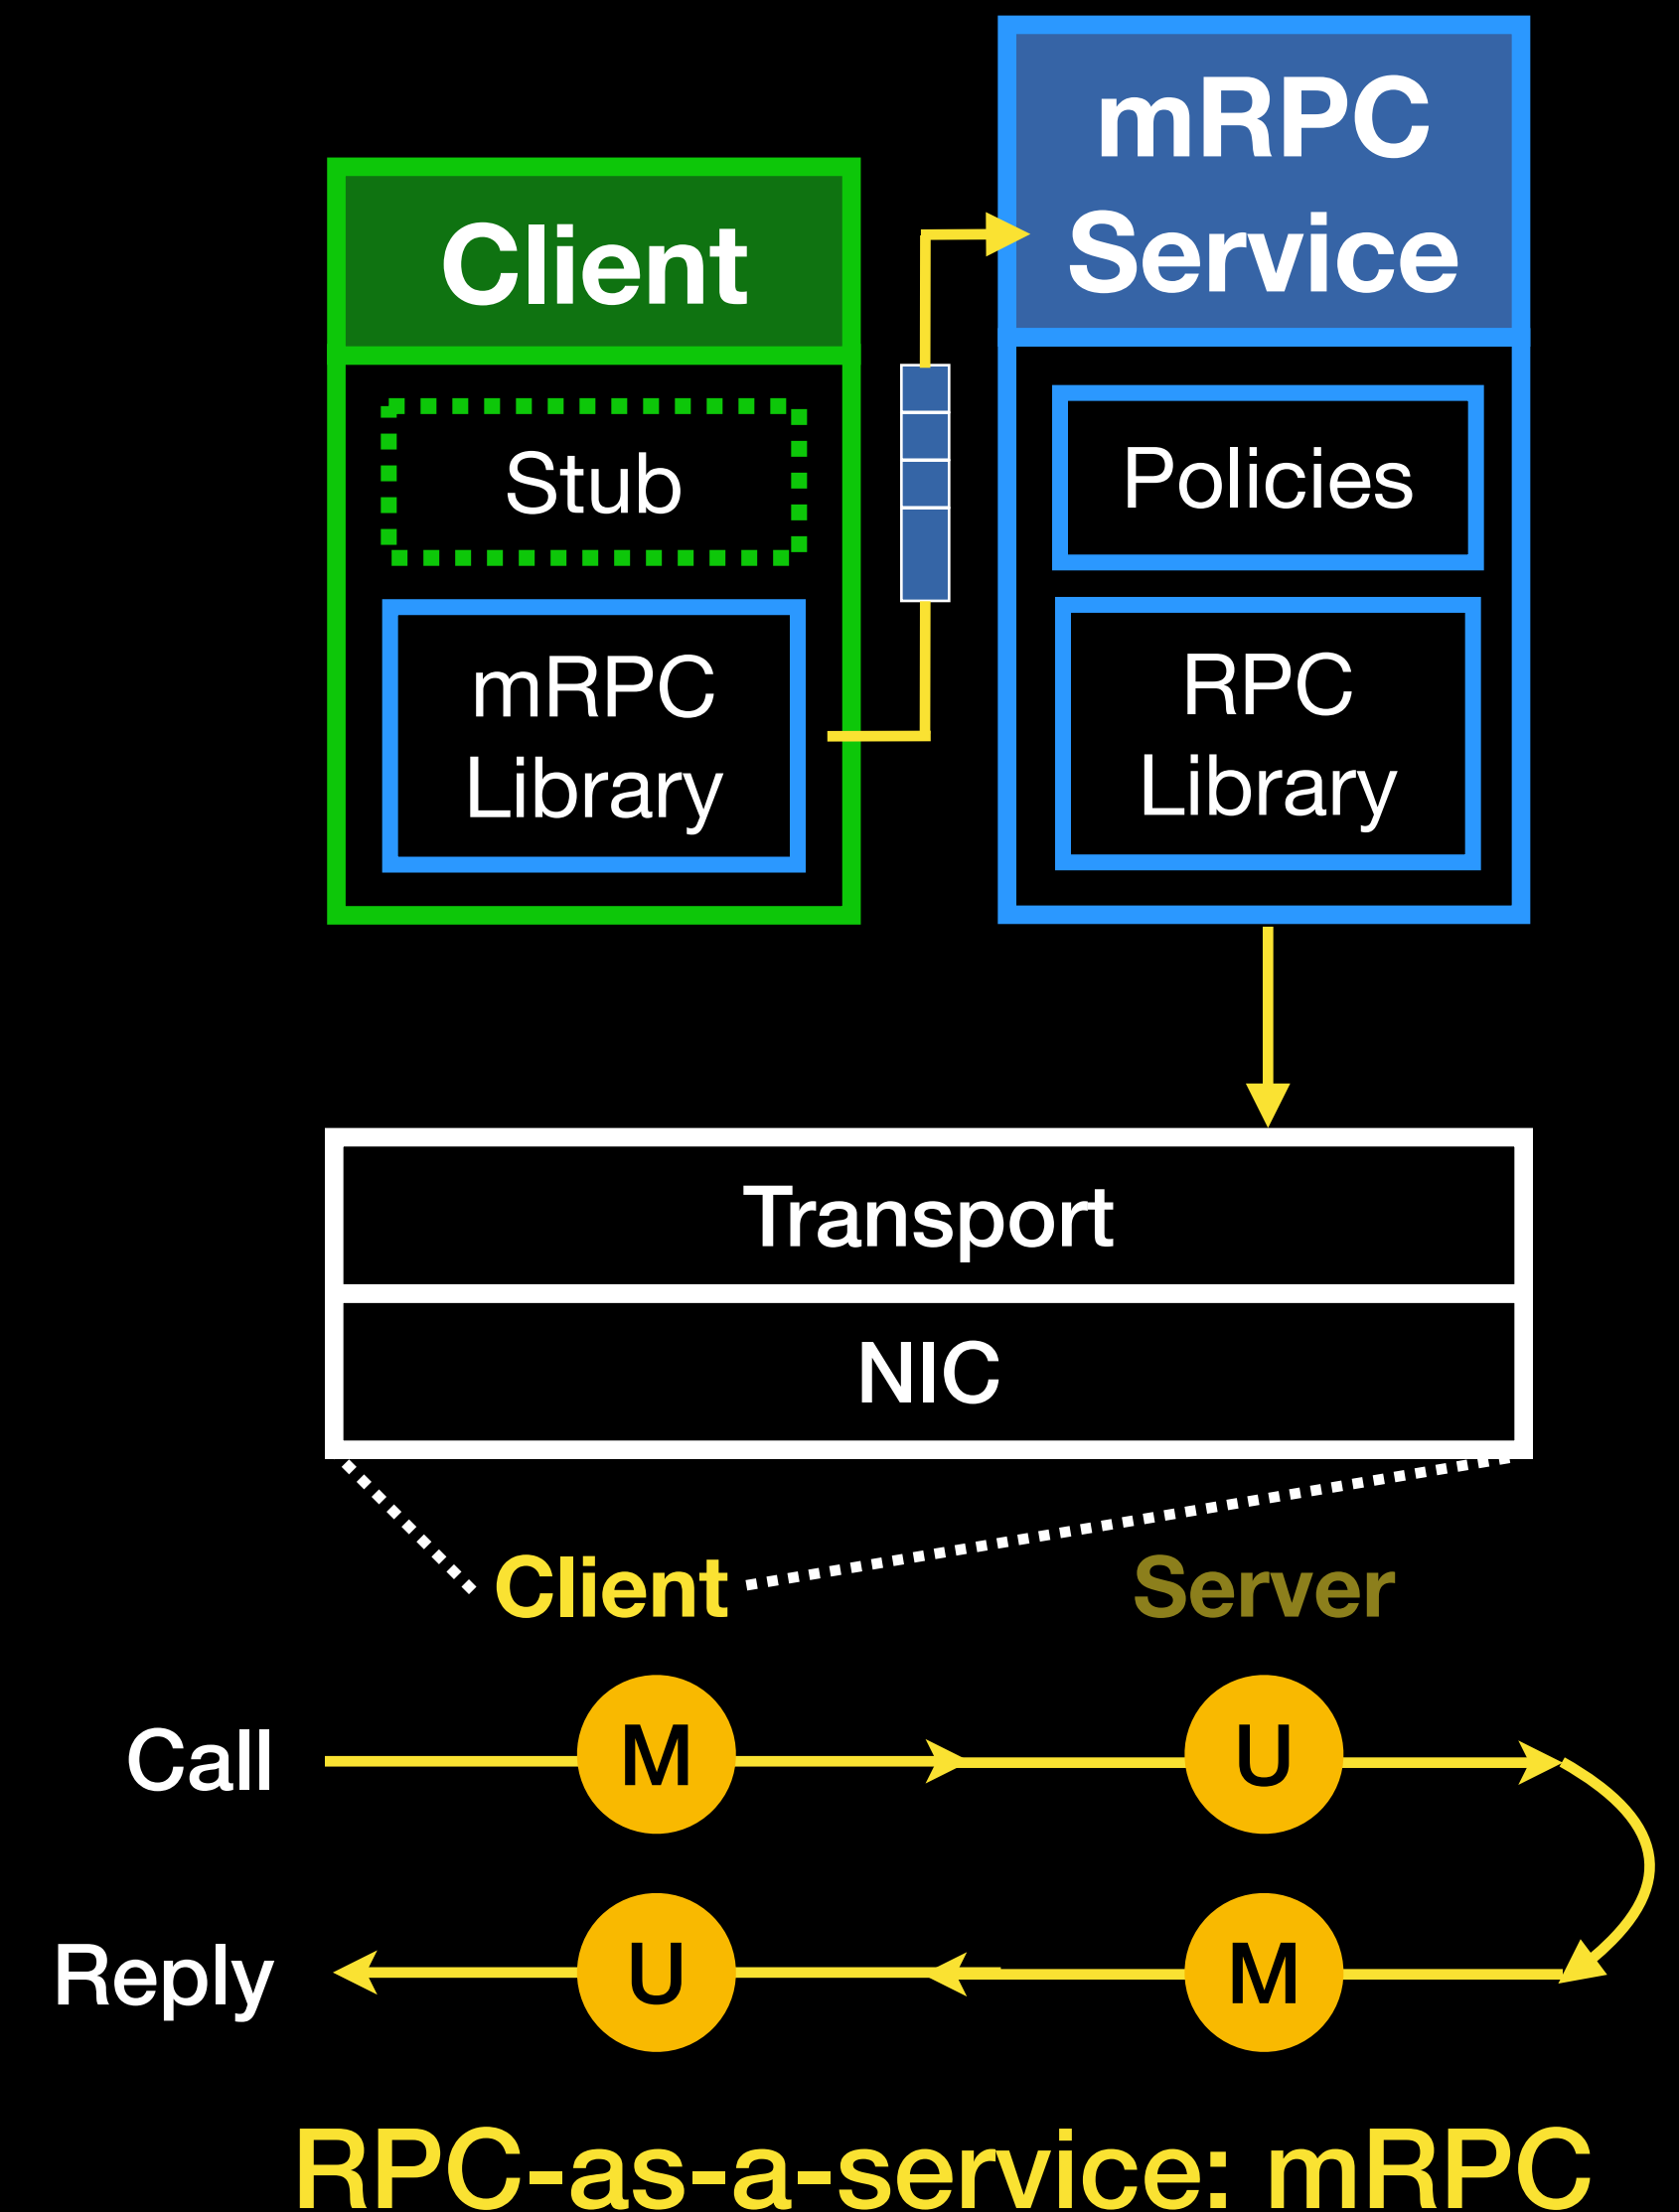
How to enforce policies with efficiency and security?

How to live upgrade RPC implementations without disrupting other applications?

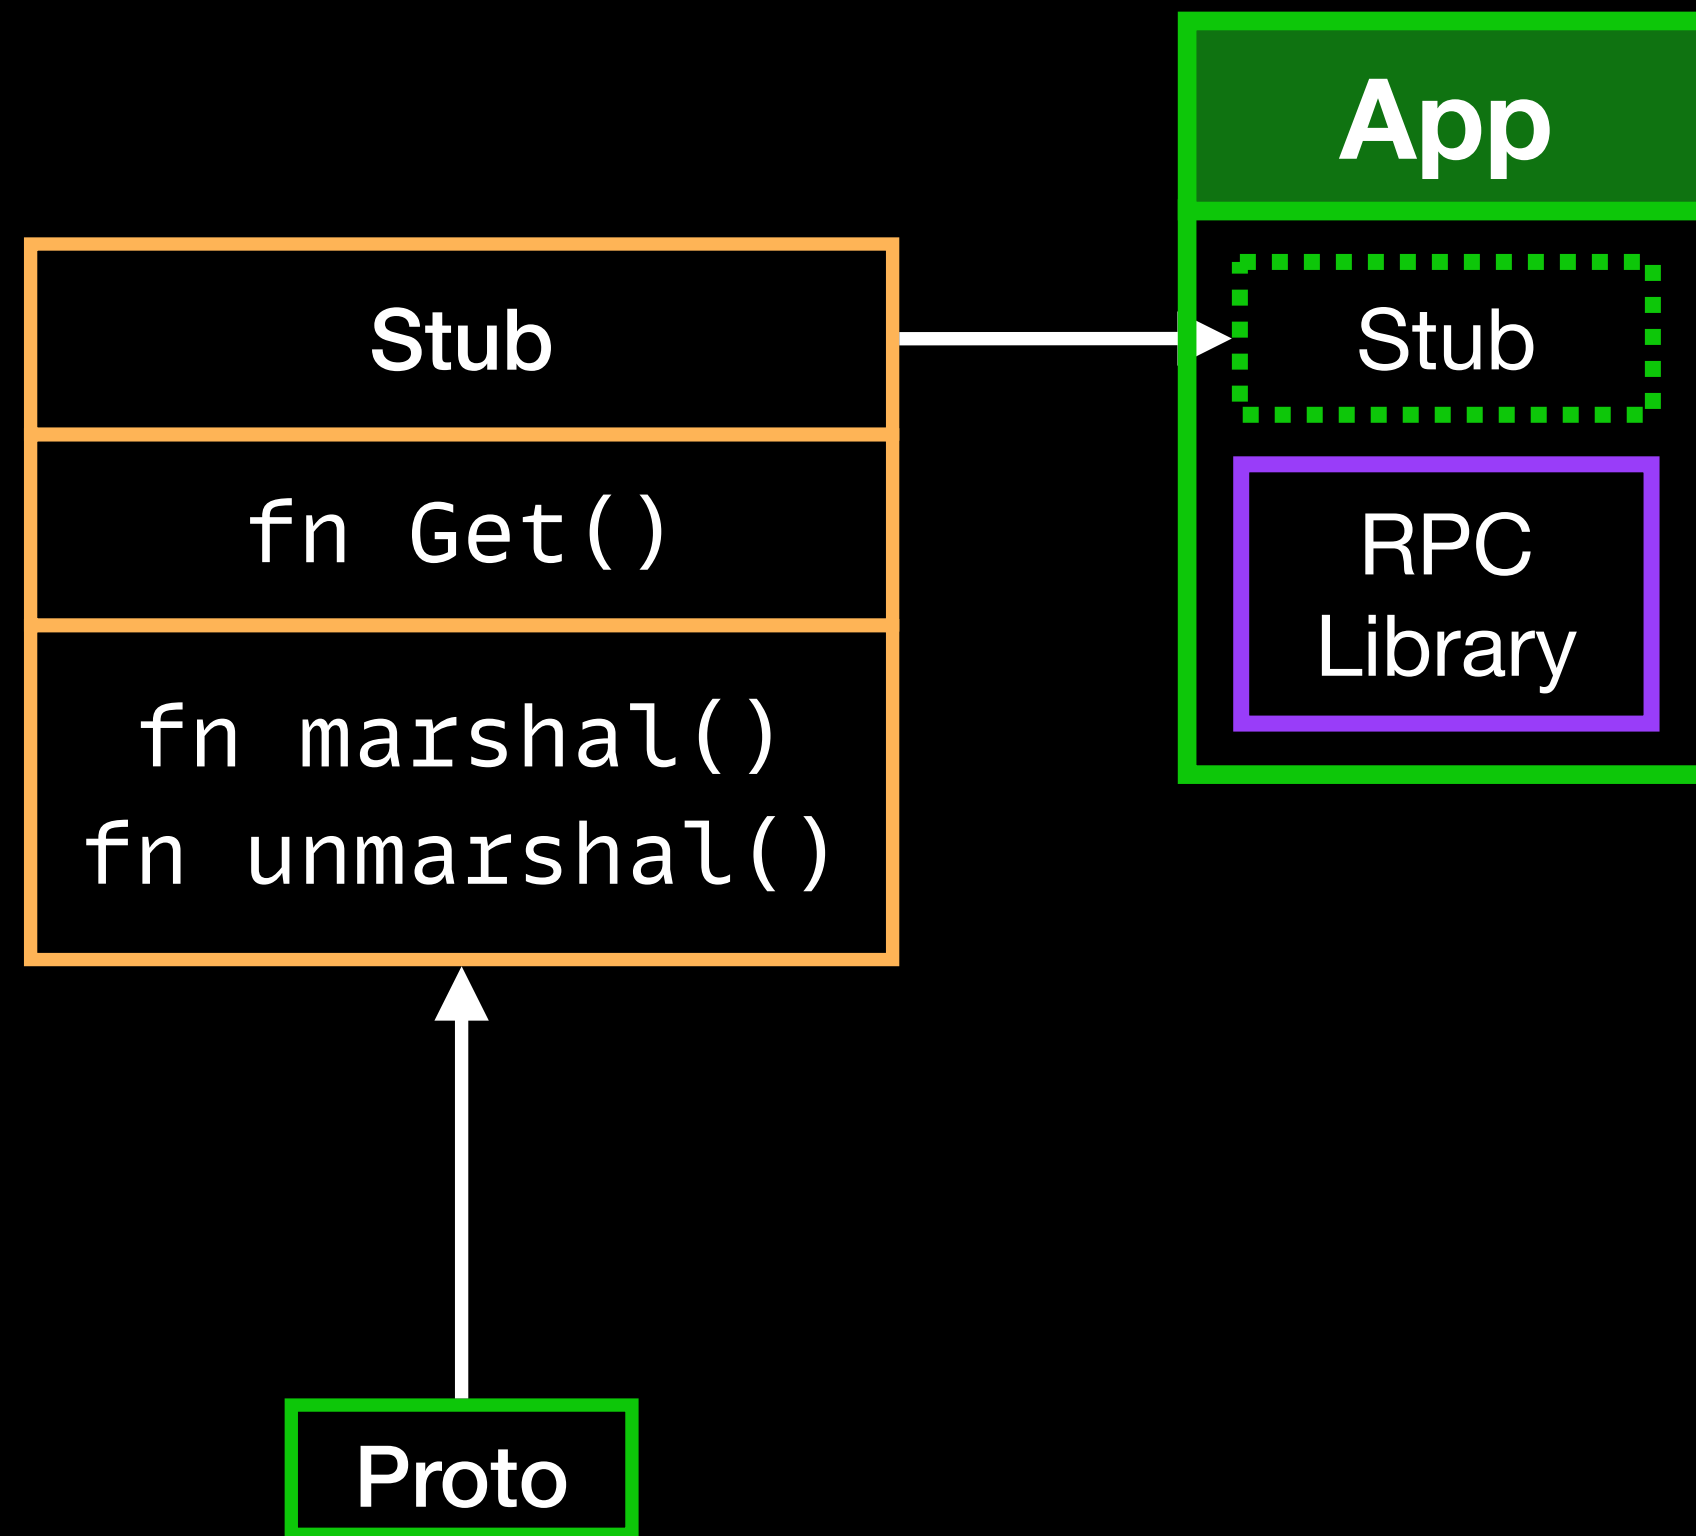
Dynamic Binding

Memory Management

Live Upgrade

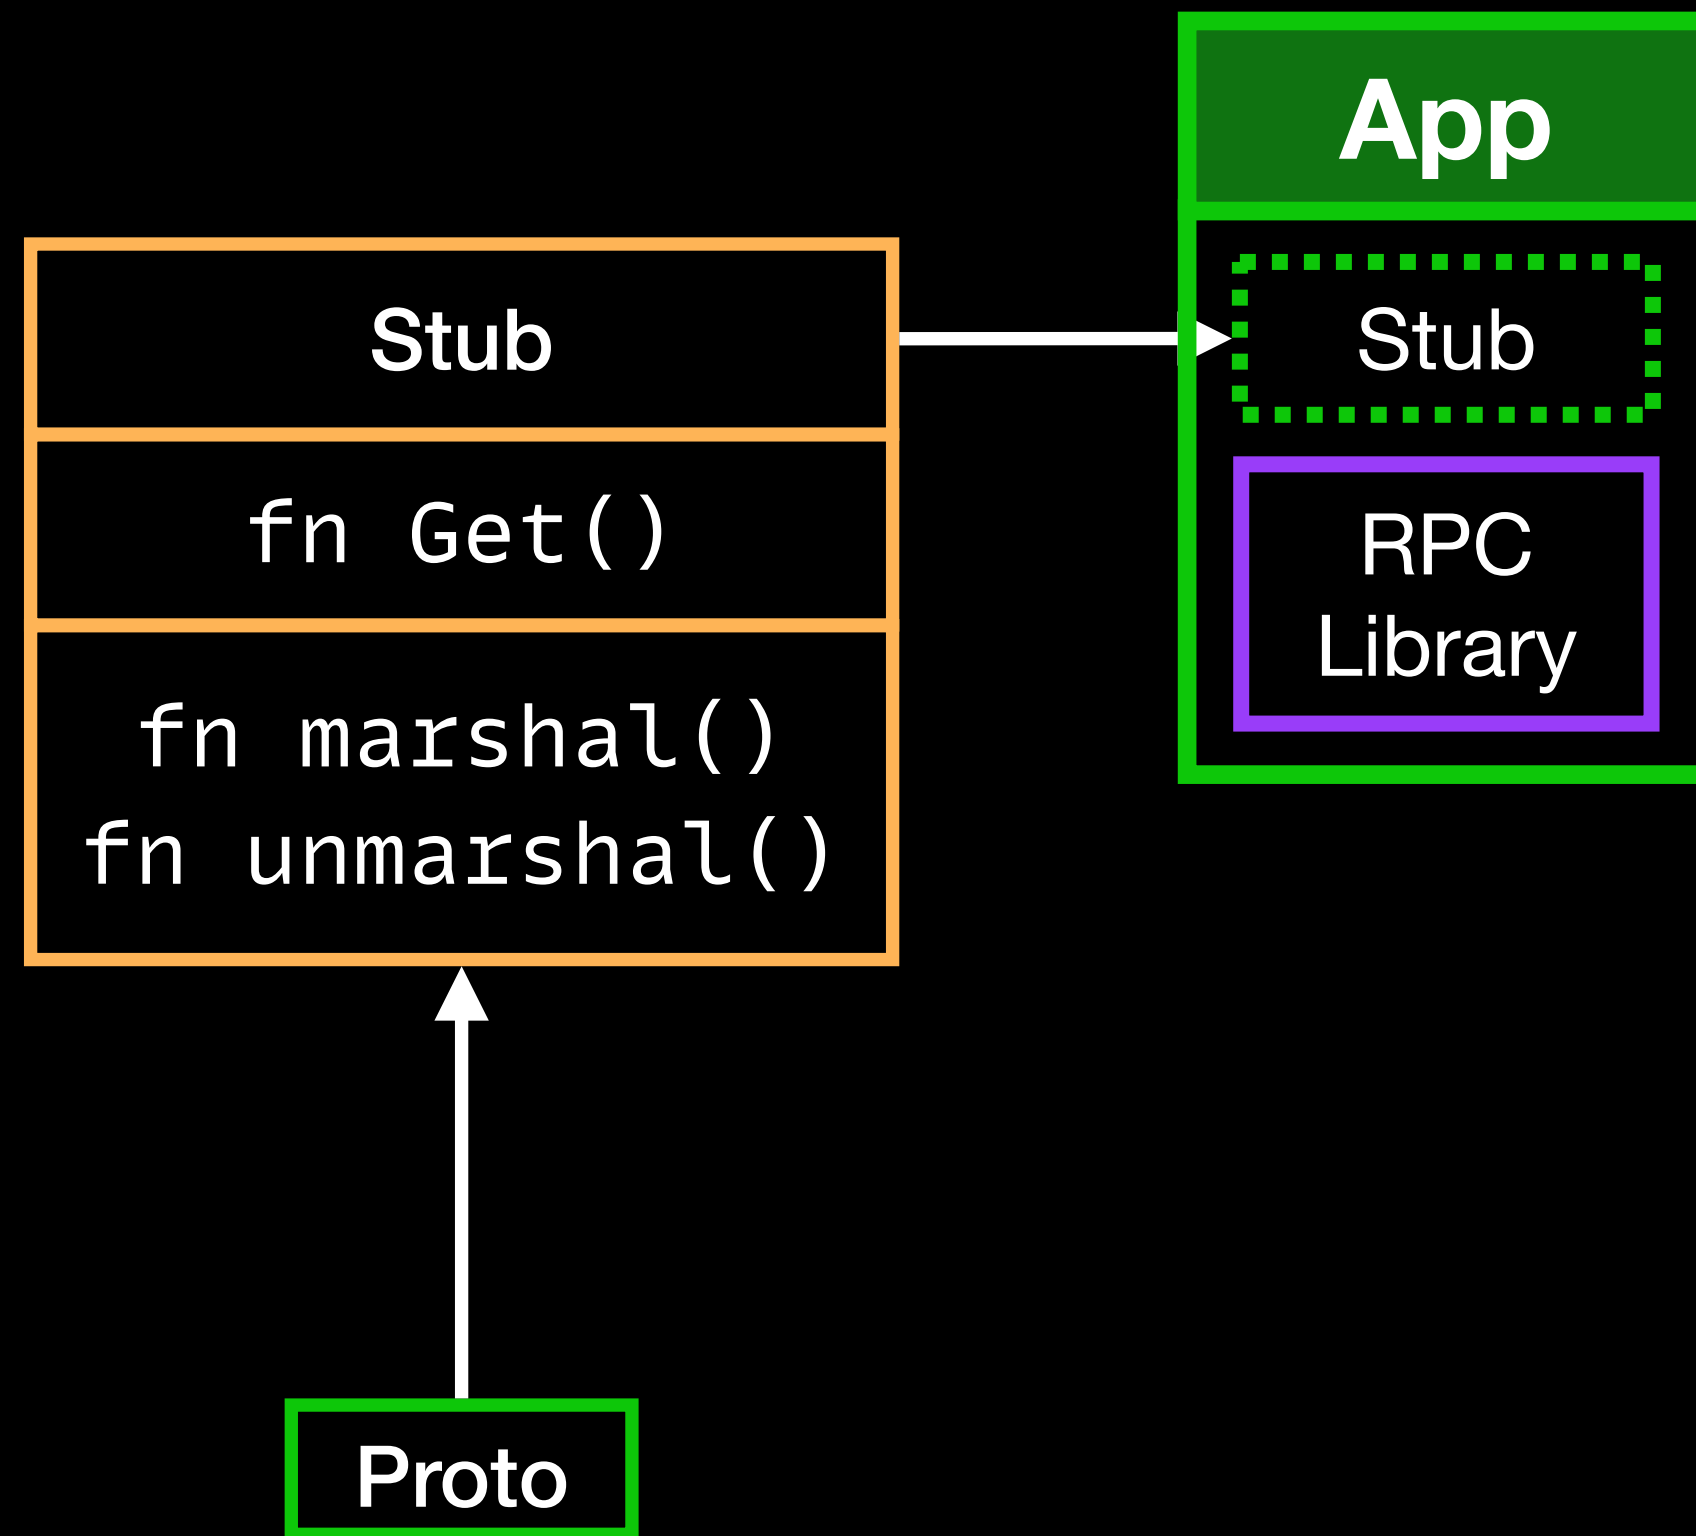


Traditional RPC Libraries



```
Service KVStore
Func Get(GetReq) -> Entry
```

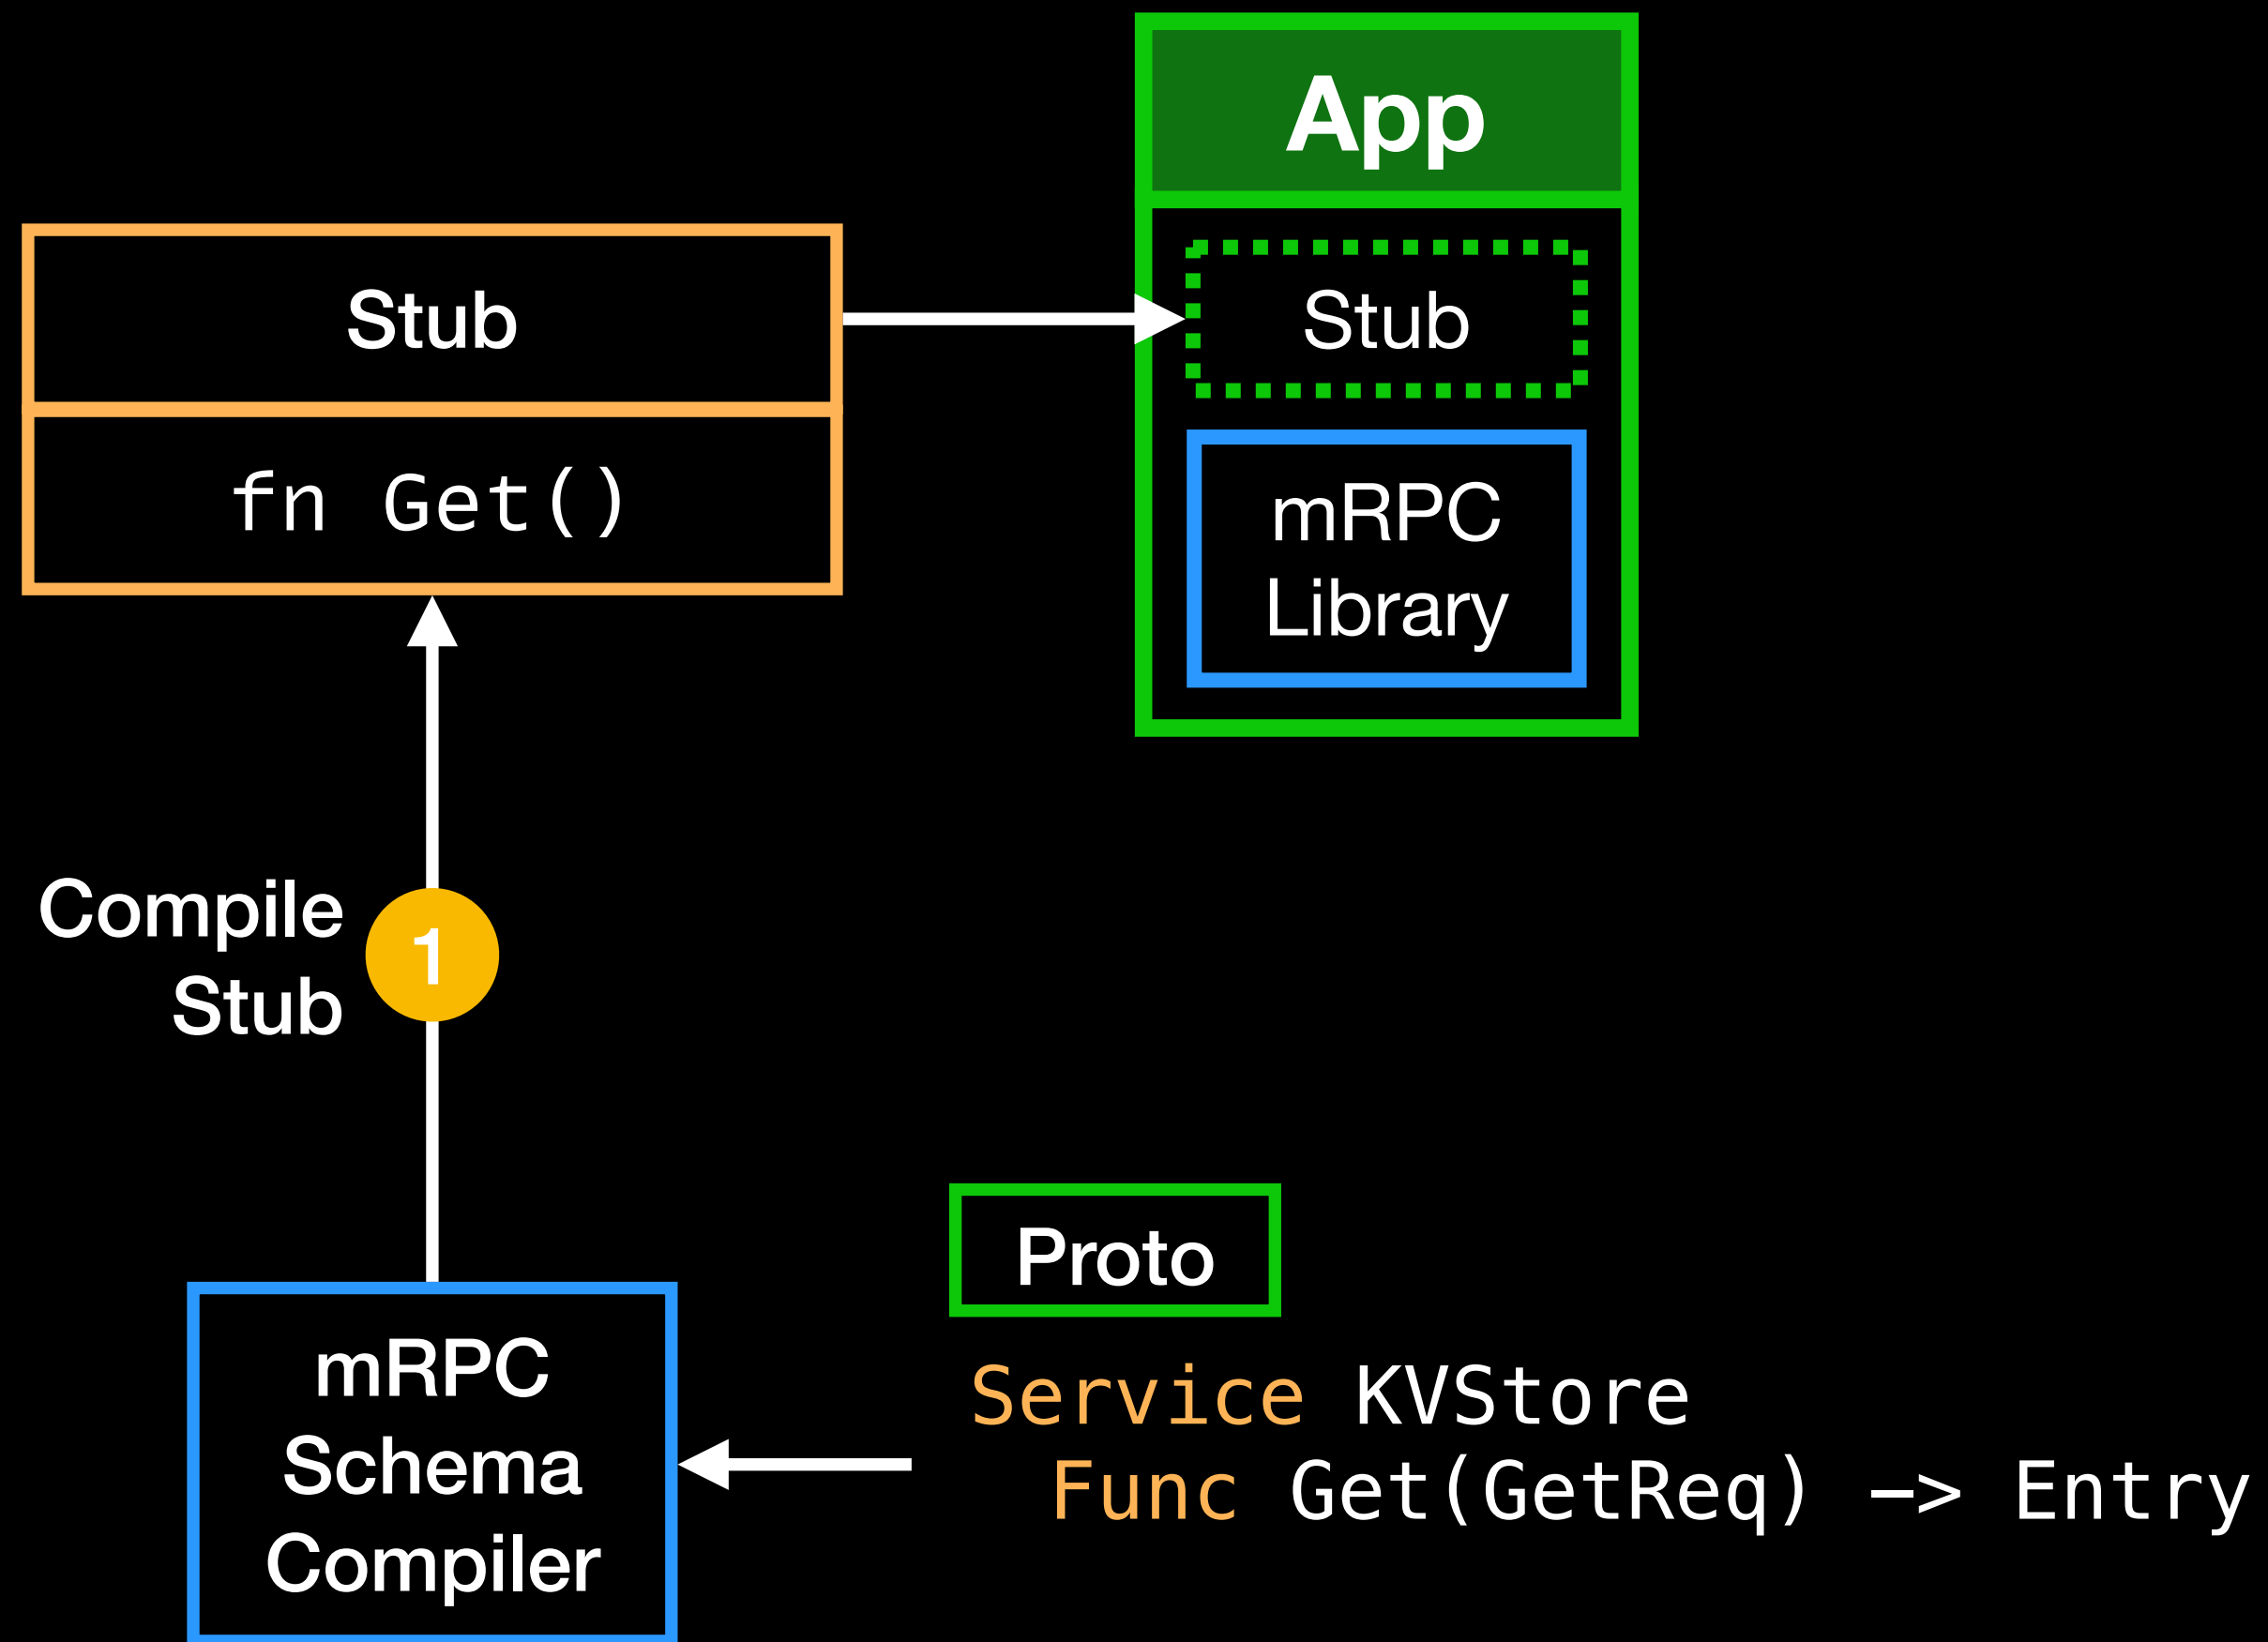
Traditional RPC Libraries



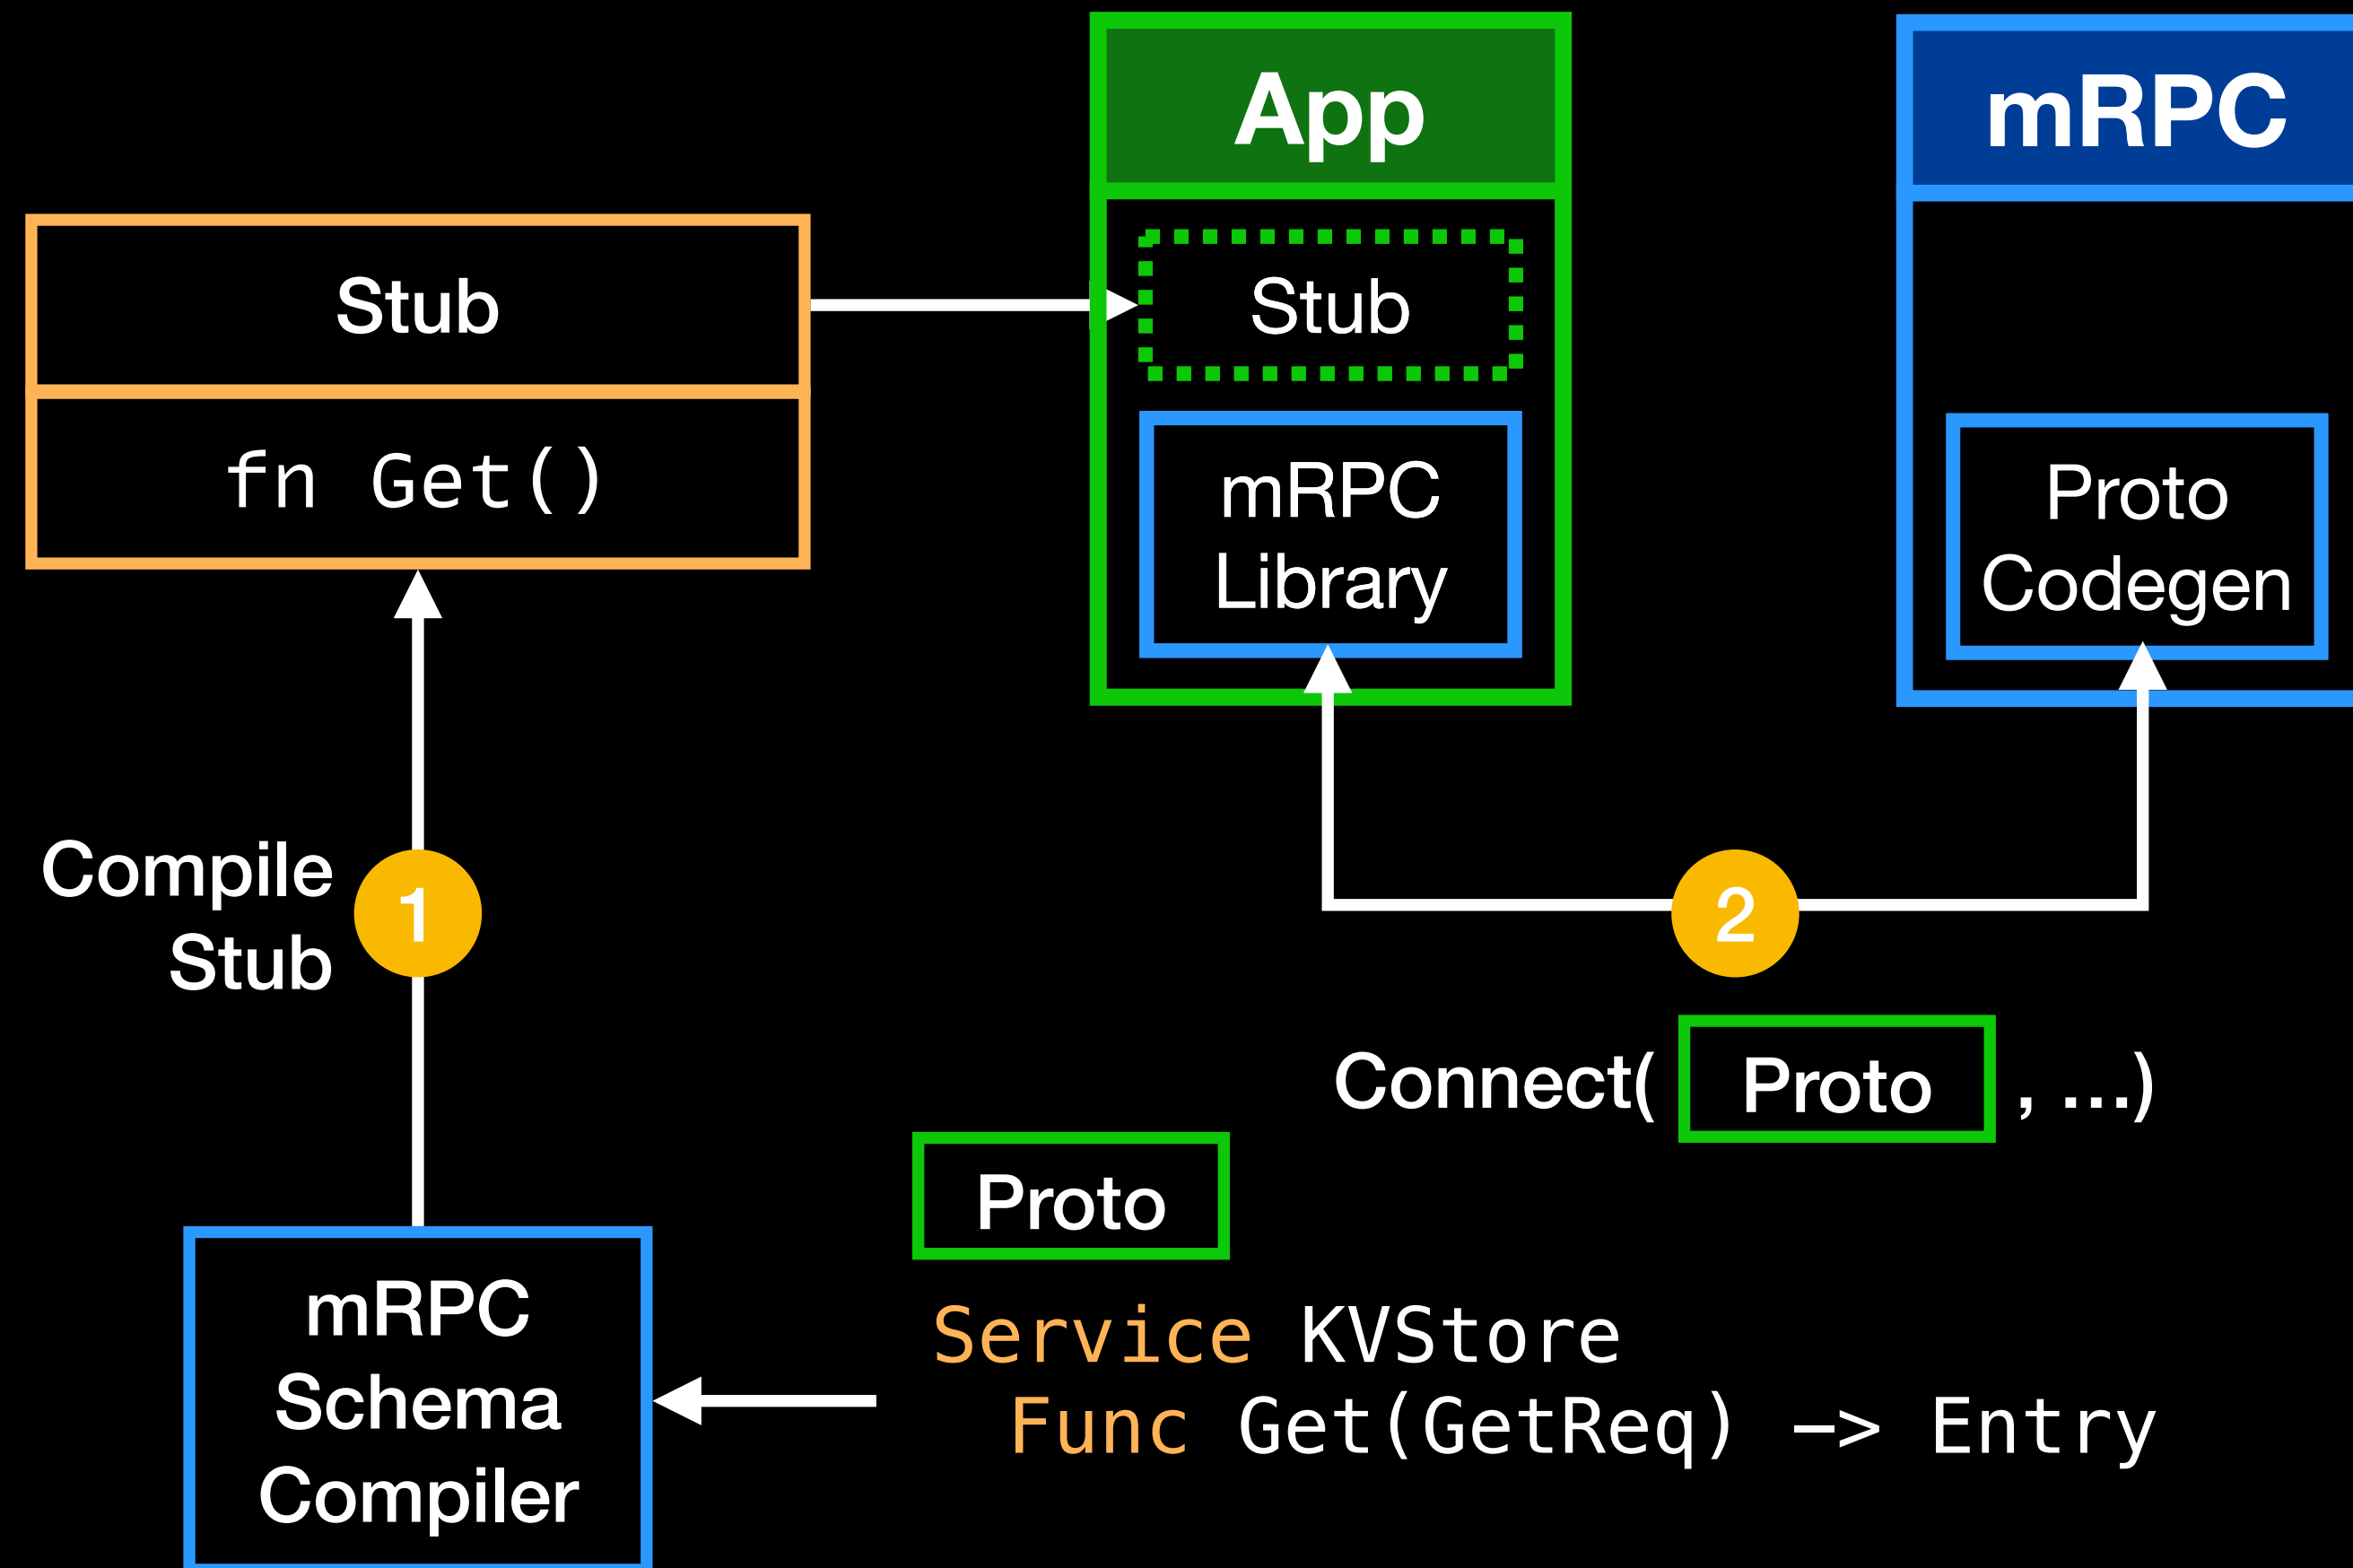
```
Service KVStore  
Func Get(GetReq) -> Entry
```

In traditional RPC libraries, marshal/unmarshal and service methods code will be generated as a stub and loaded into user applications as a library

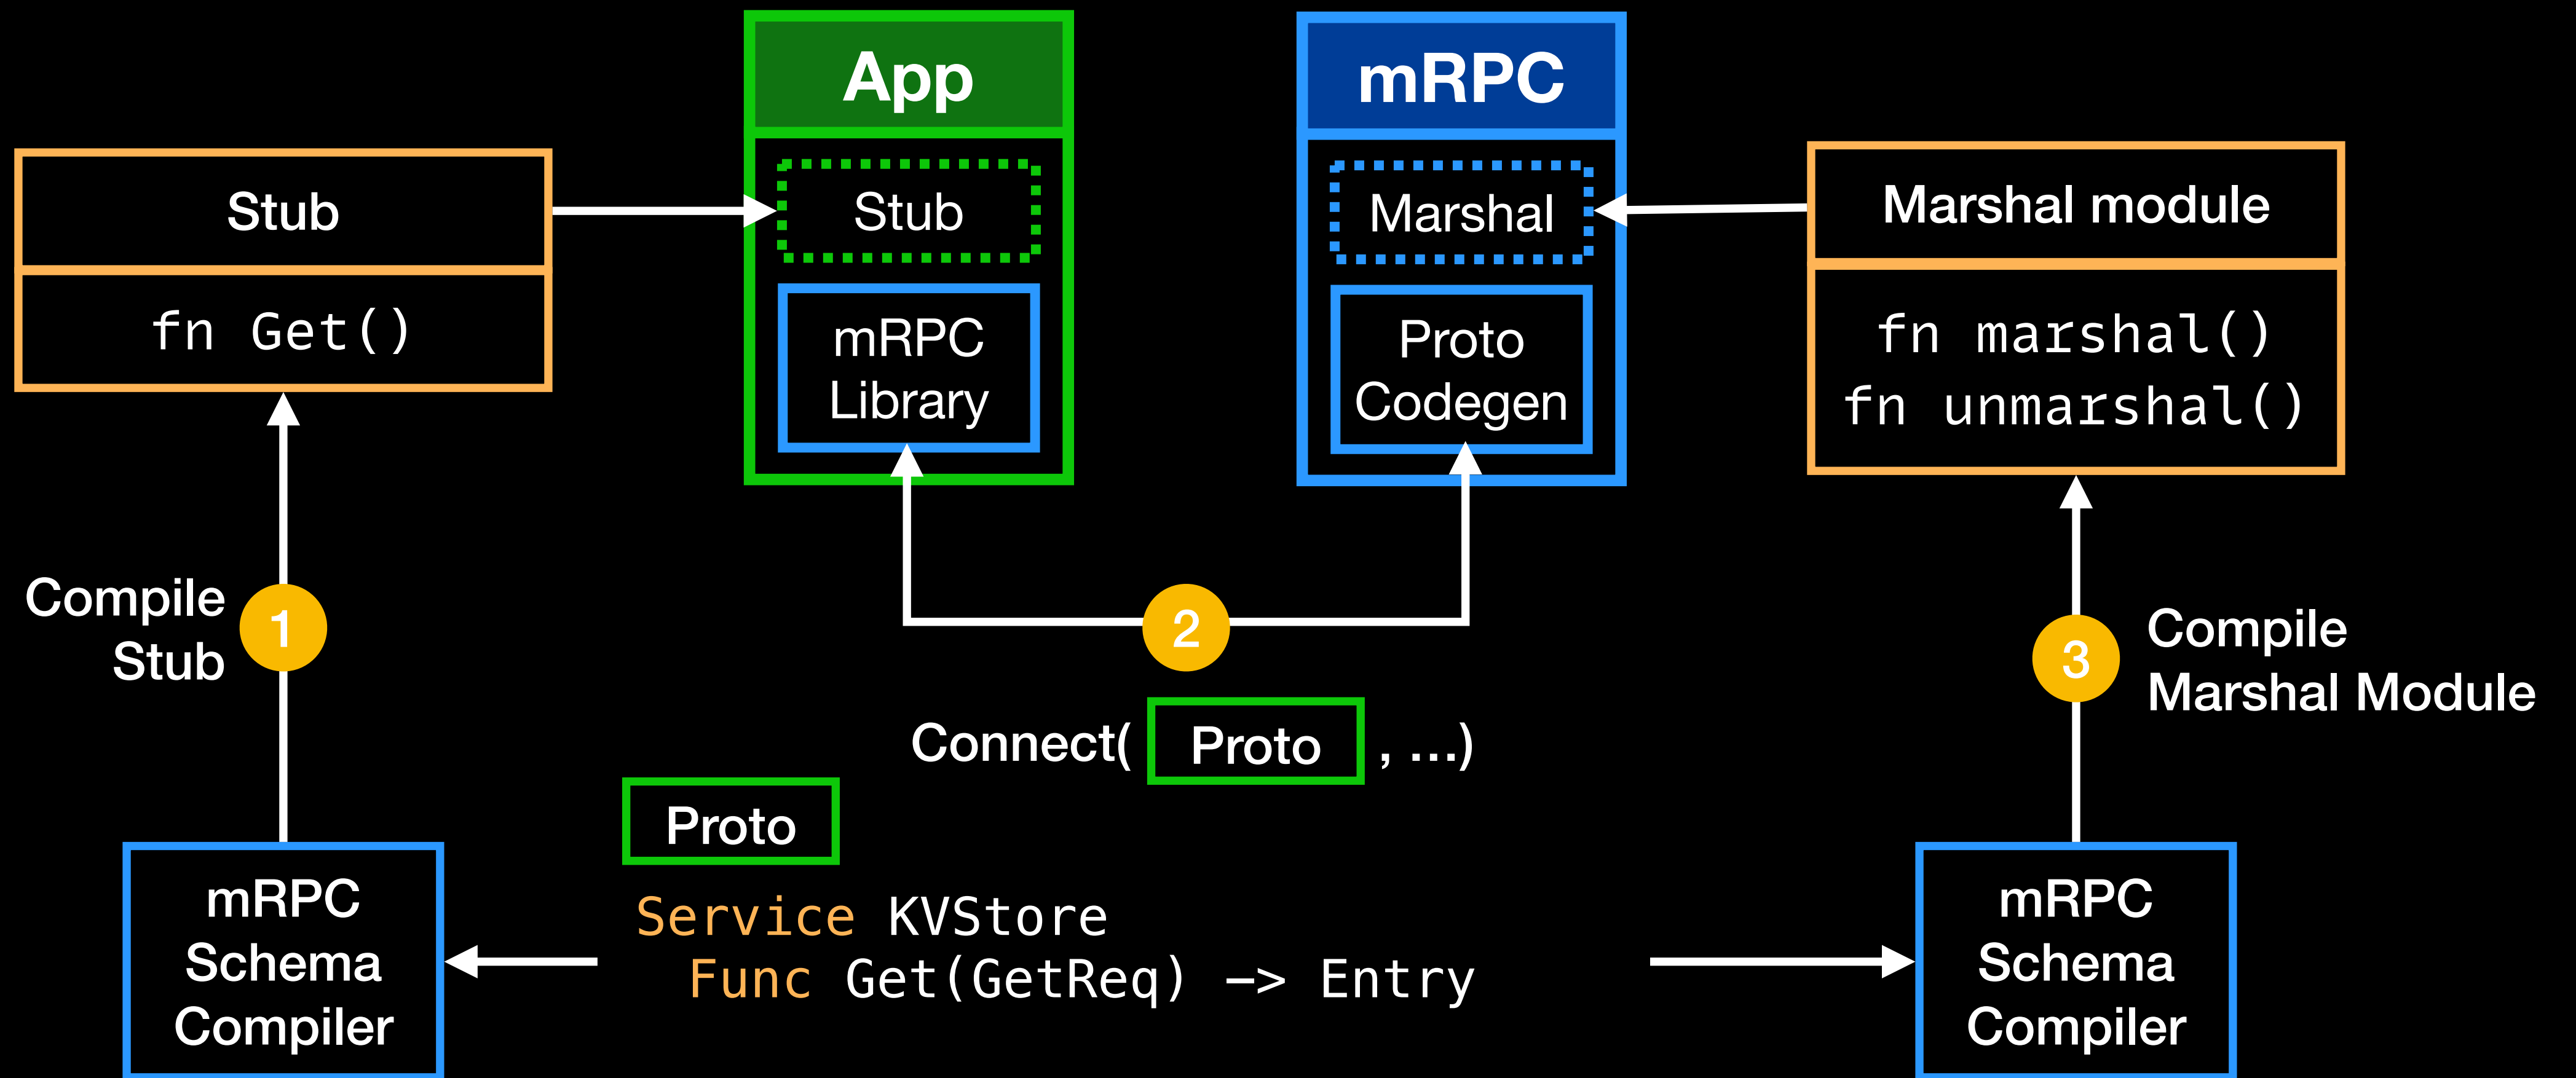
mRPC's Solution



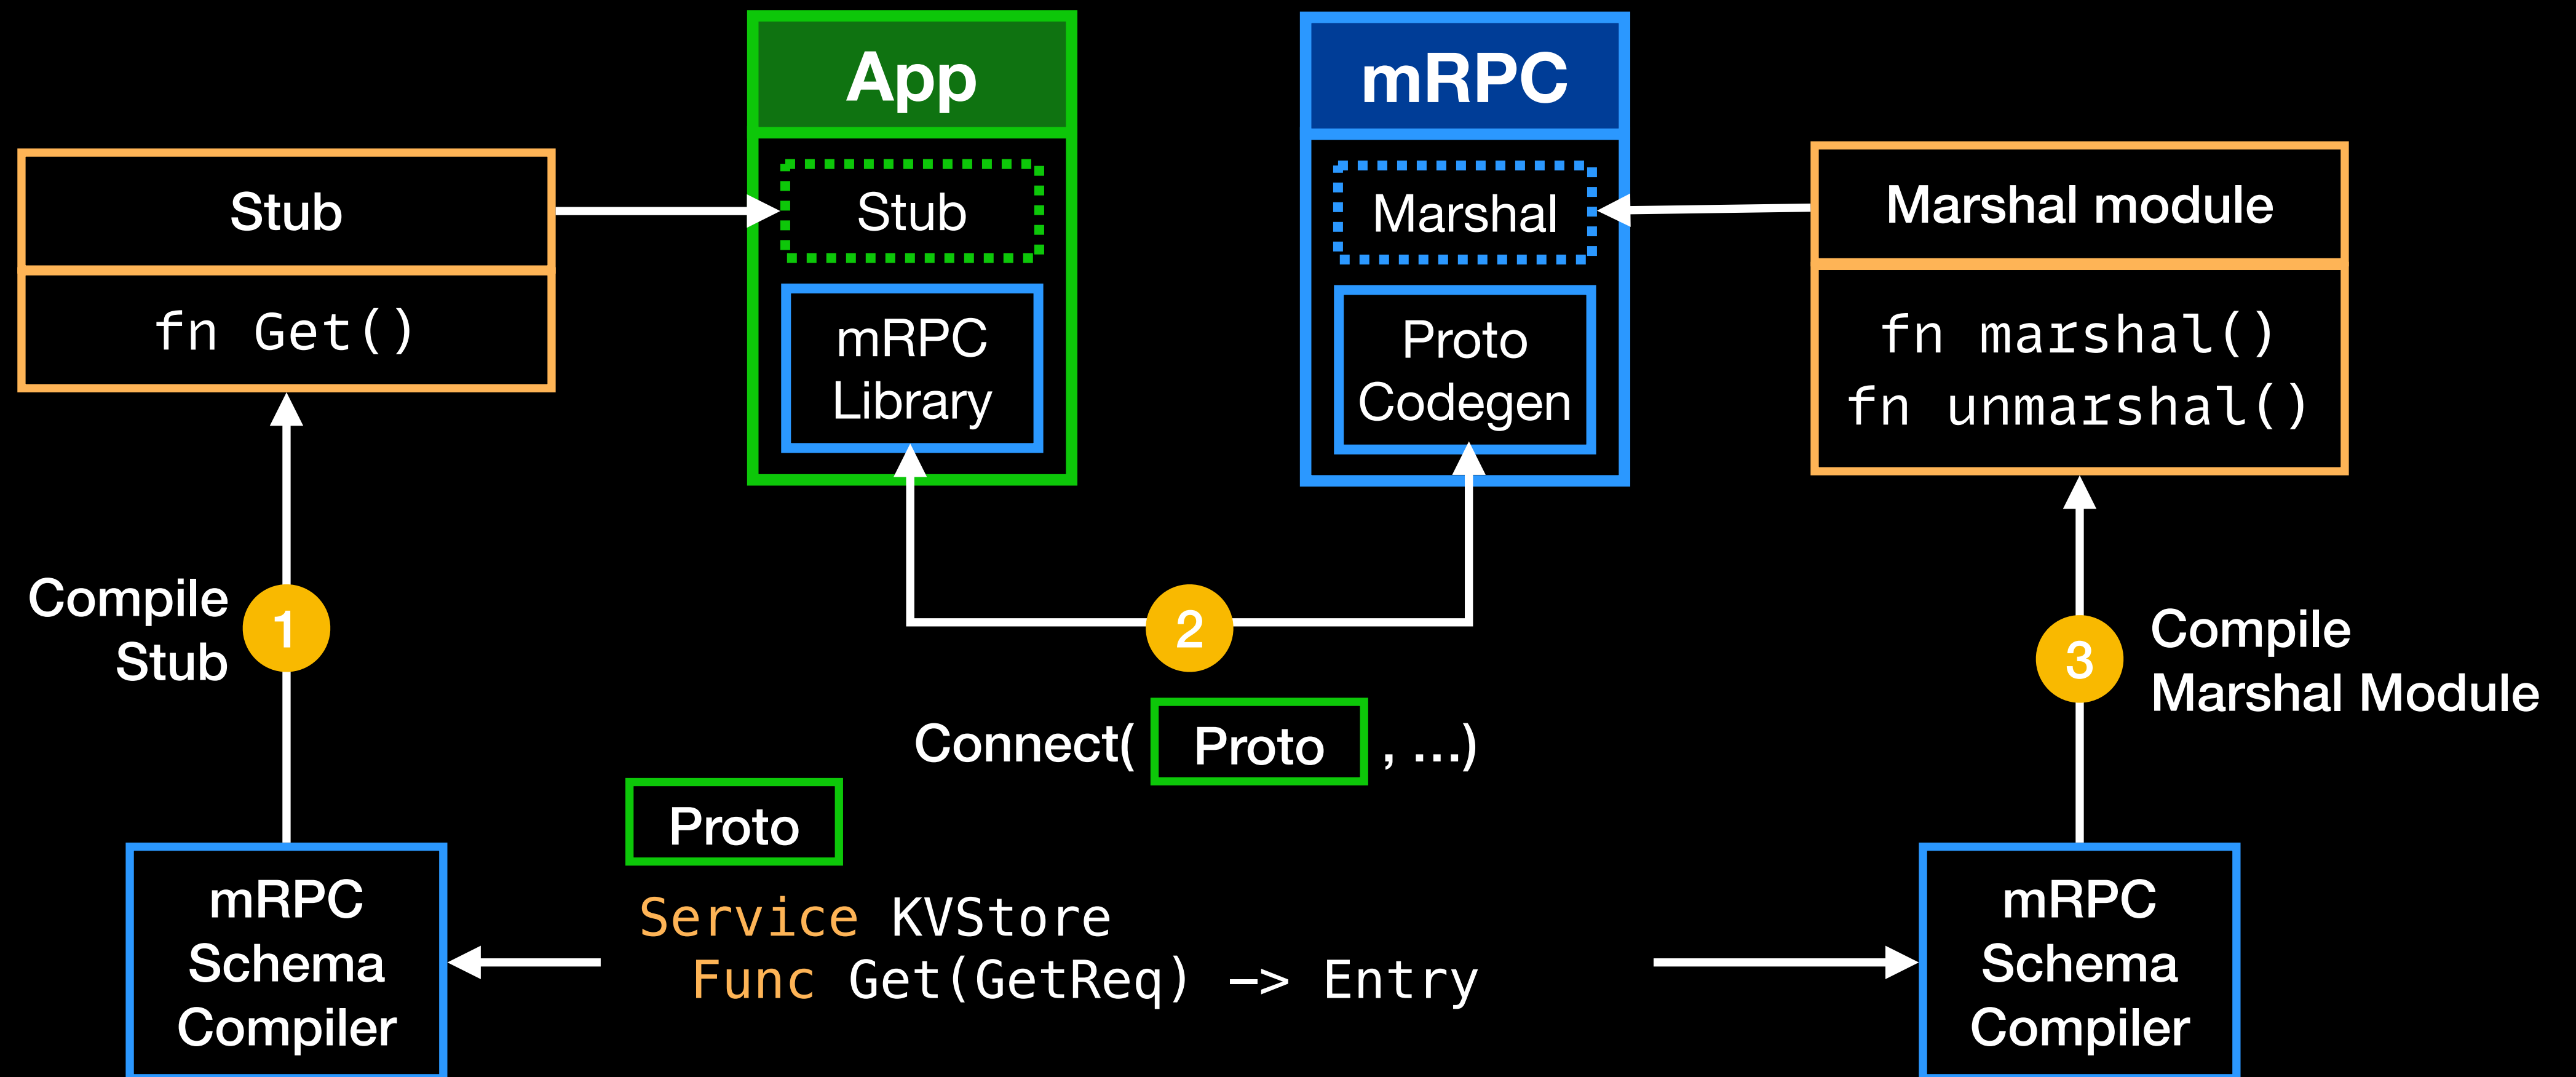
mRPC's Solution



mRPC's Solution

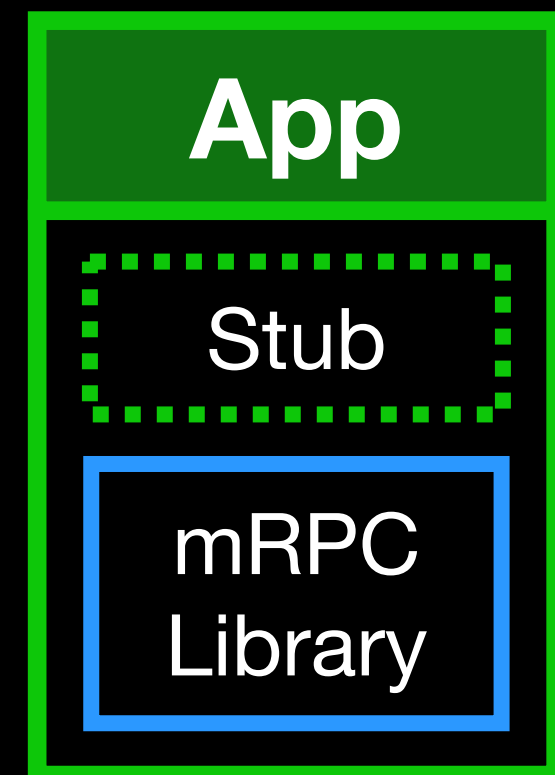


mRPC's Solution



In mRPC, marshal/unmarshal code are decoupled from user stub, and generated/loaded by mRPC service instead



Challenge #2: Memory Management



```
Service KVStore  
Func Get(GetReq) -> Entry
```



Challenge #2: Memory Management



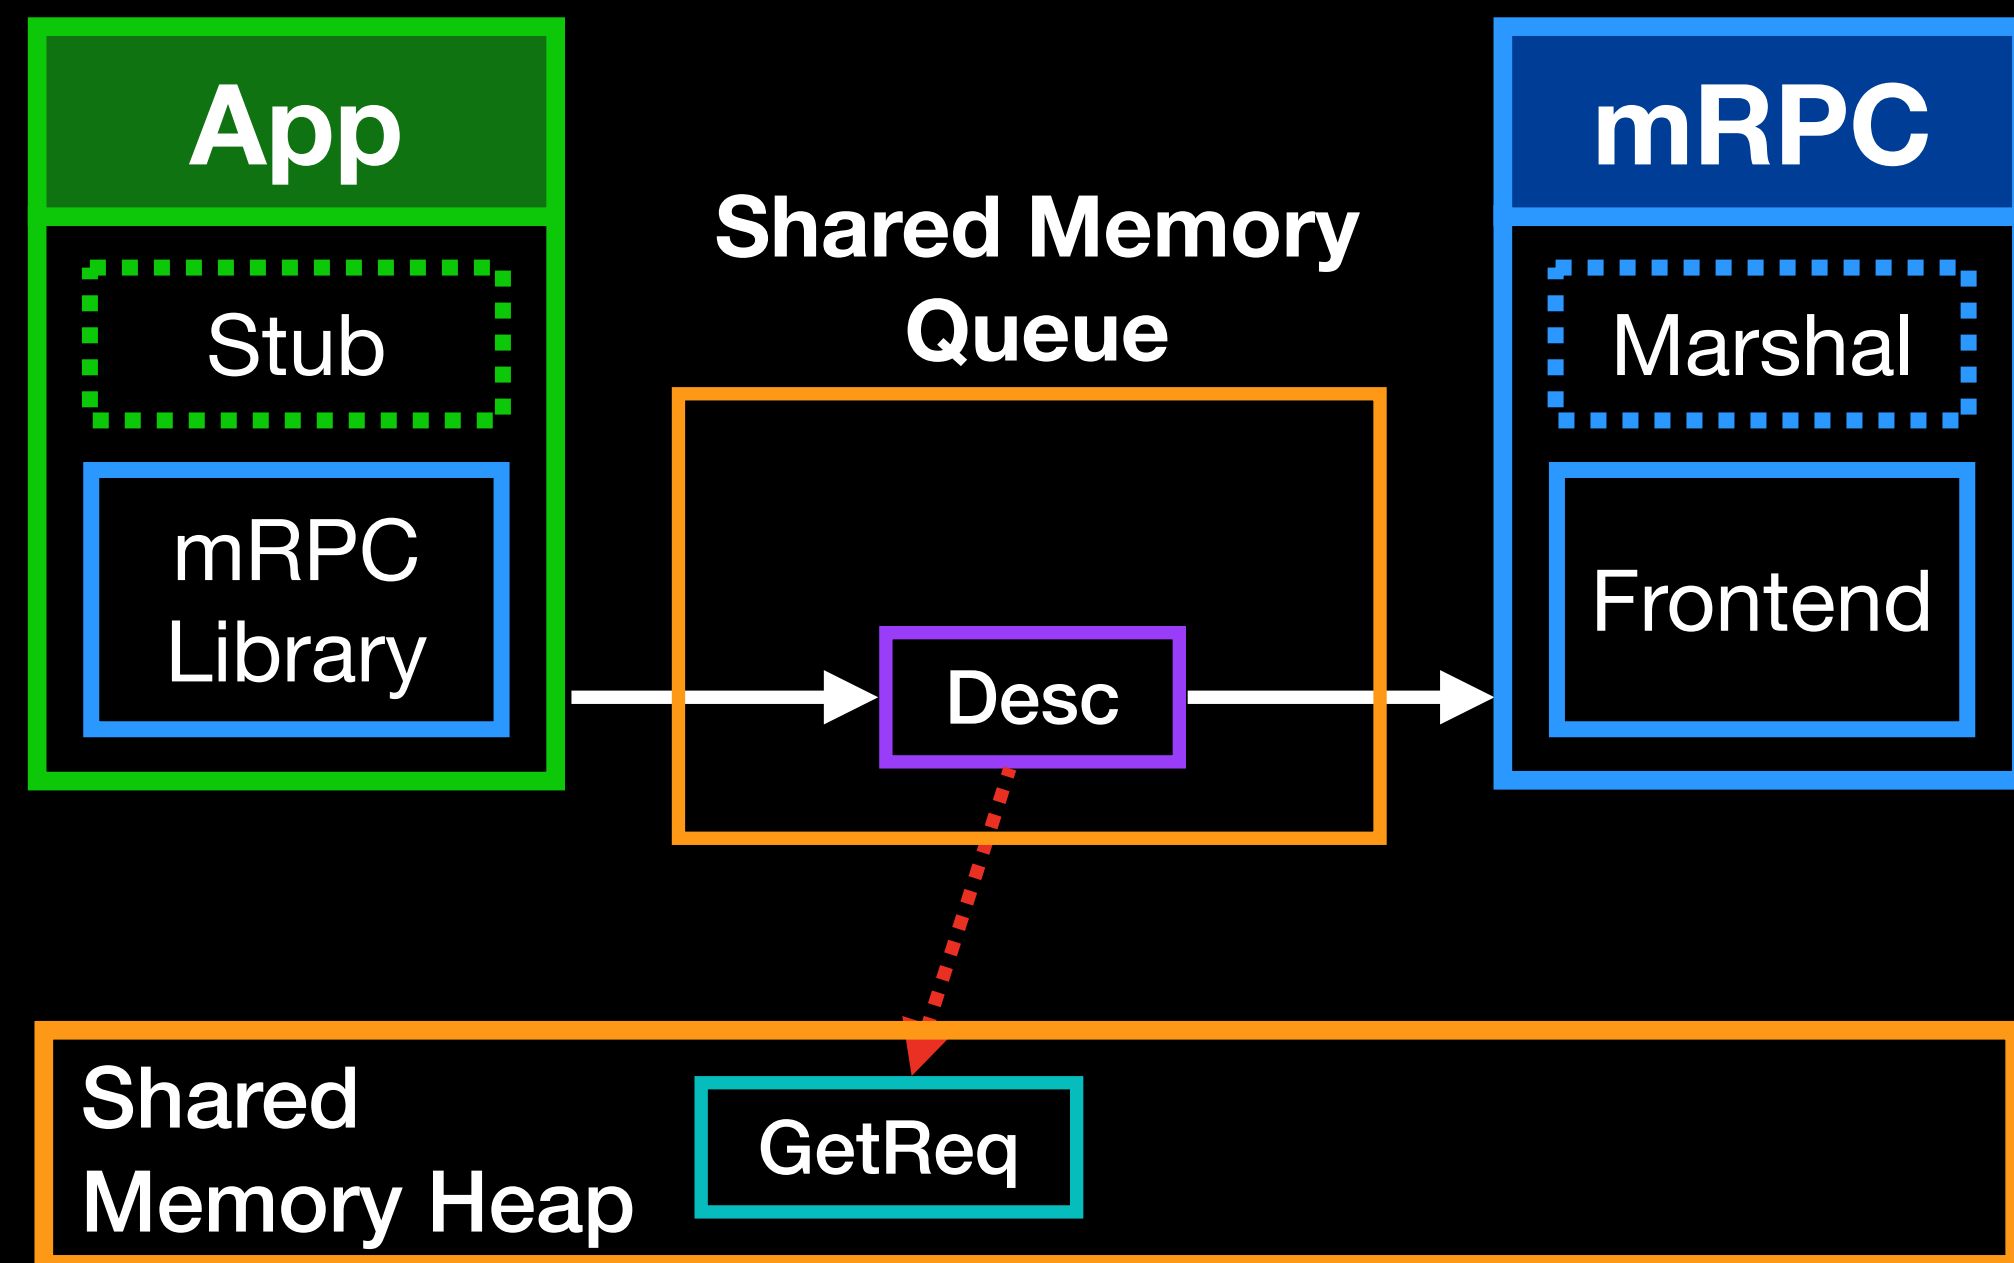
```
Service KVStore  
Func Get(GetReq) -> Entry
```




RPC messages are allocated on shared memory heap.

Accessed by both the application and the mRPC service.

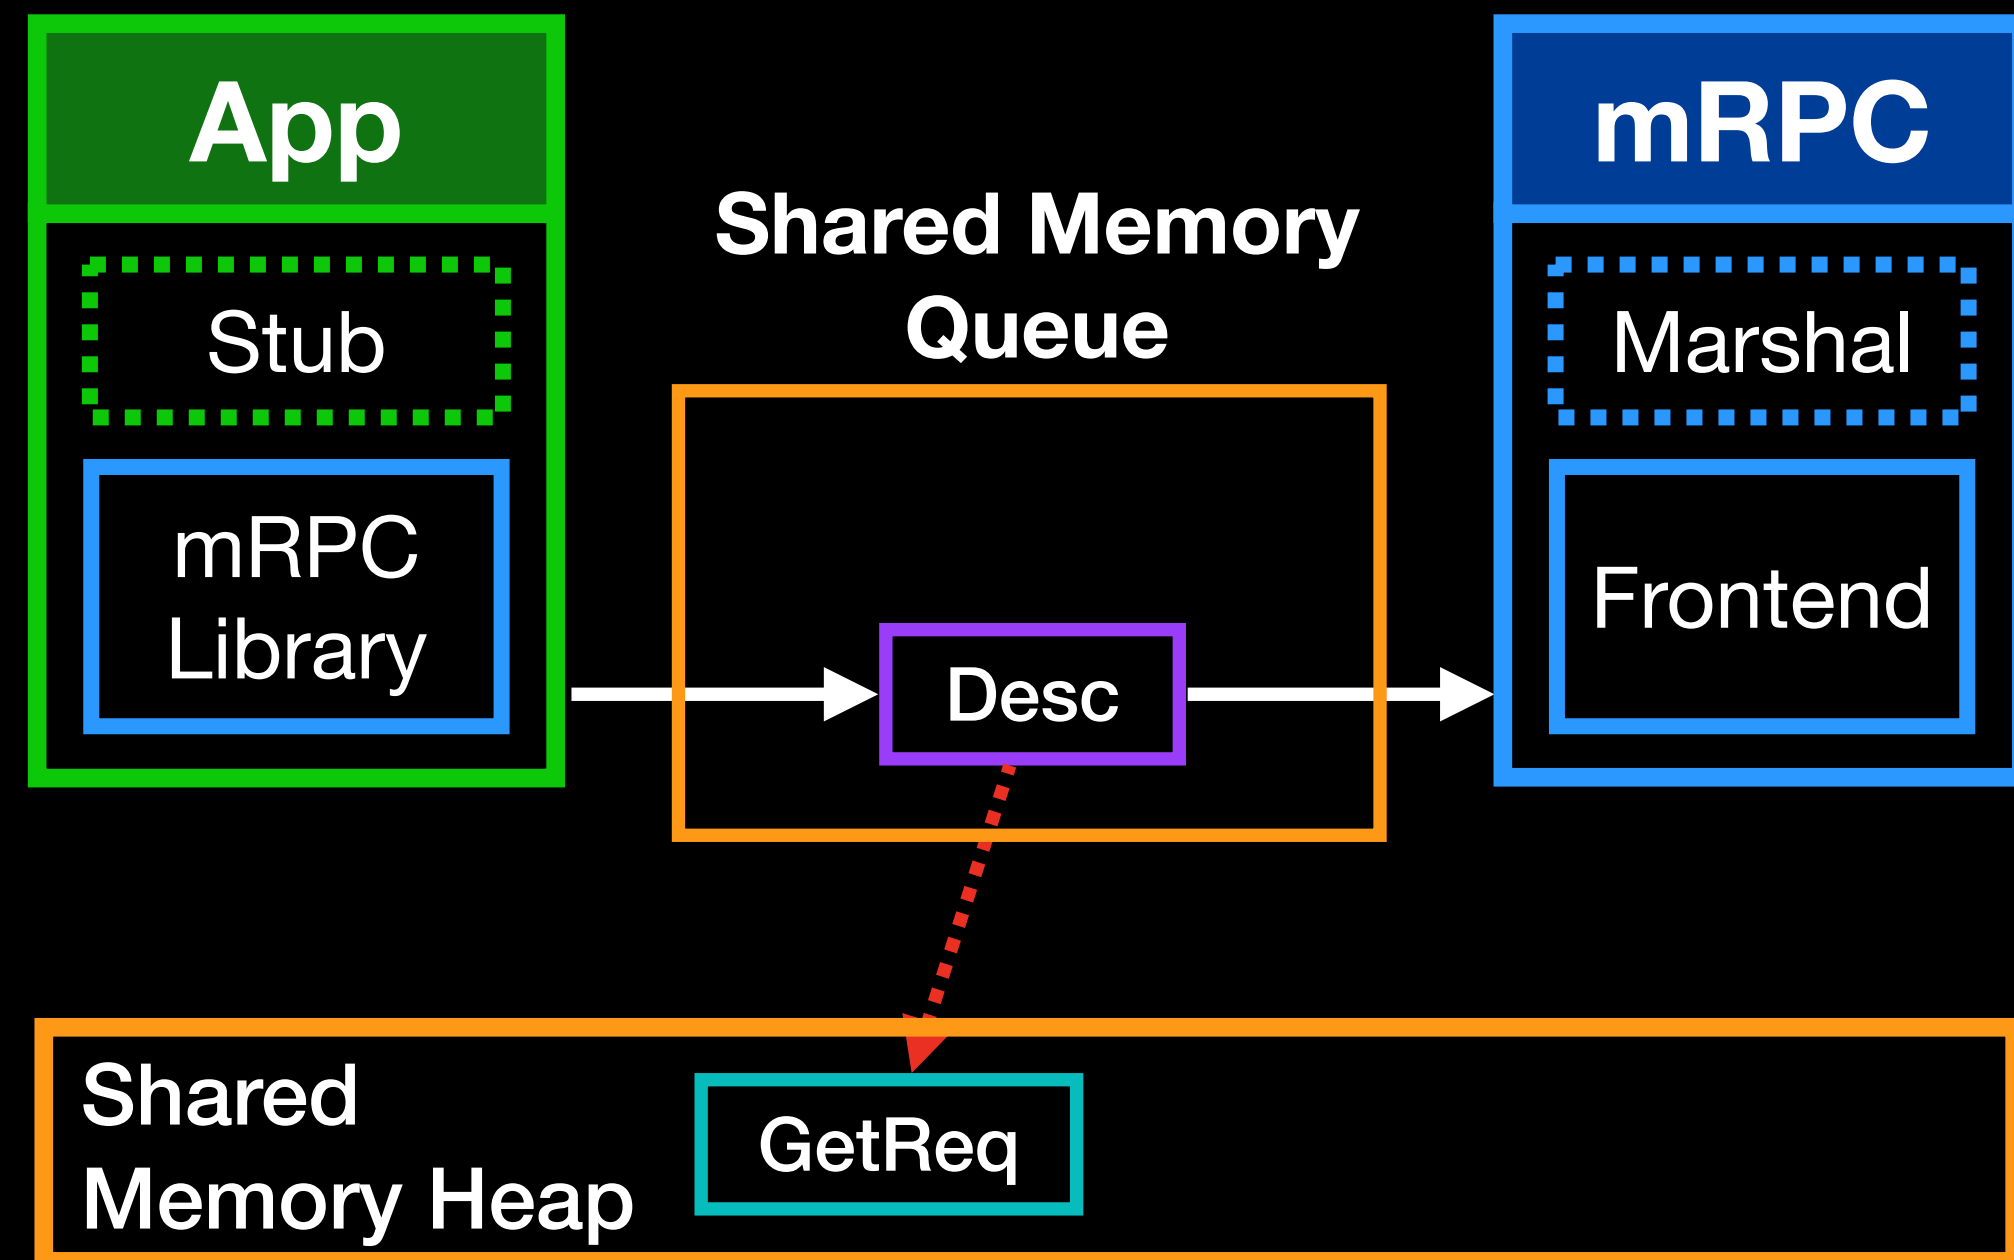
Memory Management




Service KVStore
Func Get(GetReq) -> Entry



Memory Management

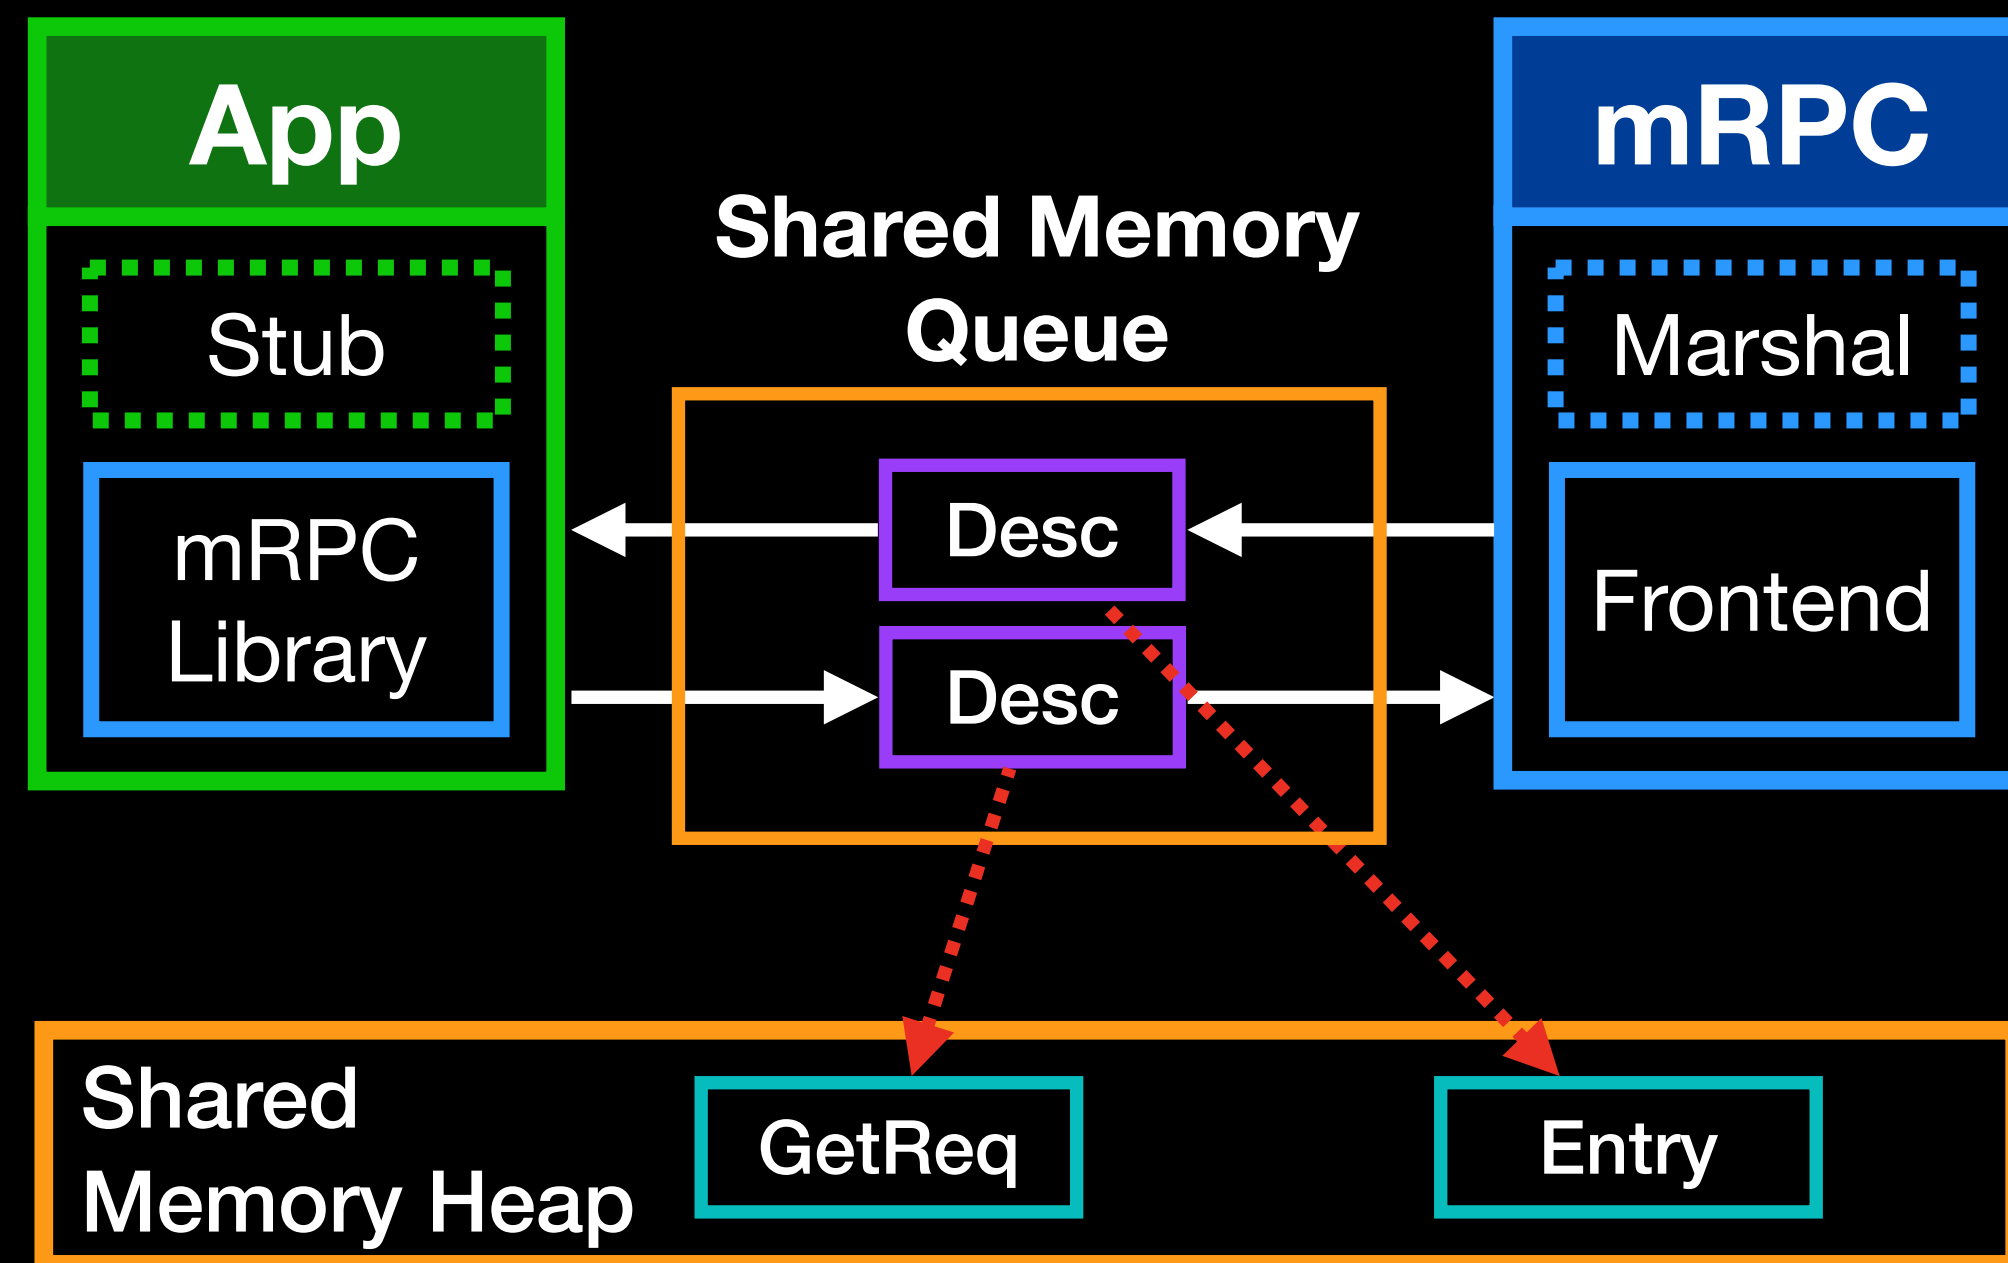


```
Service KVStore  
Func Get(GetReq) -> Entry
```




A shared memory queue is used to pass RPC descriptors

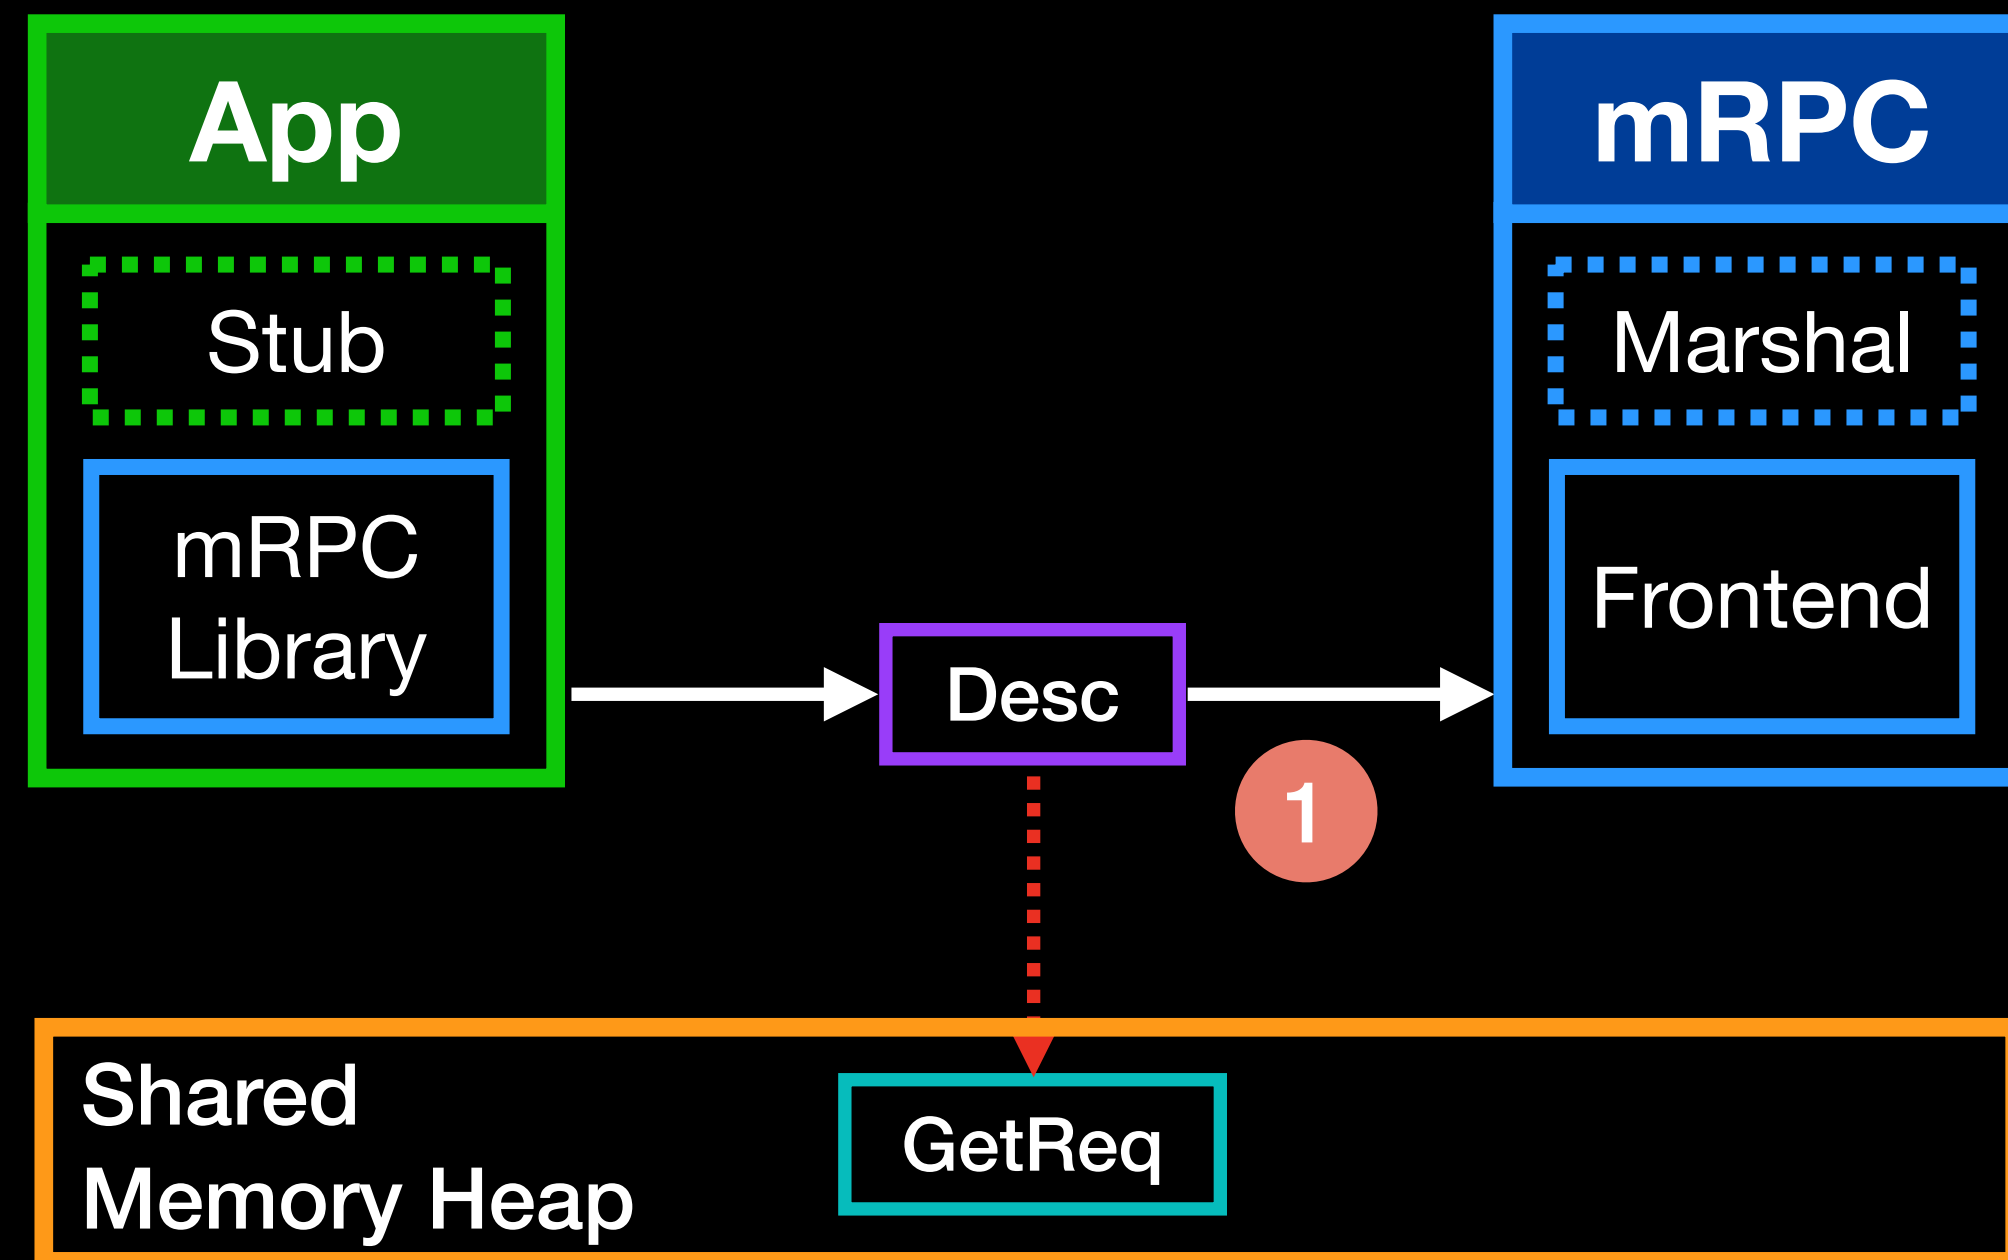
Memory Management



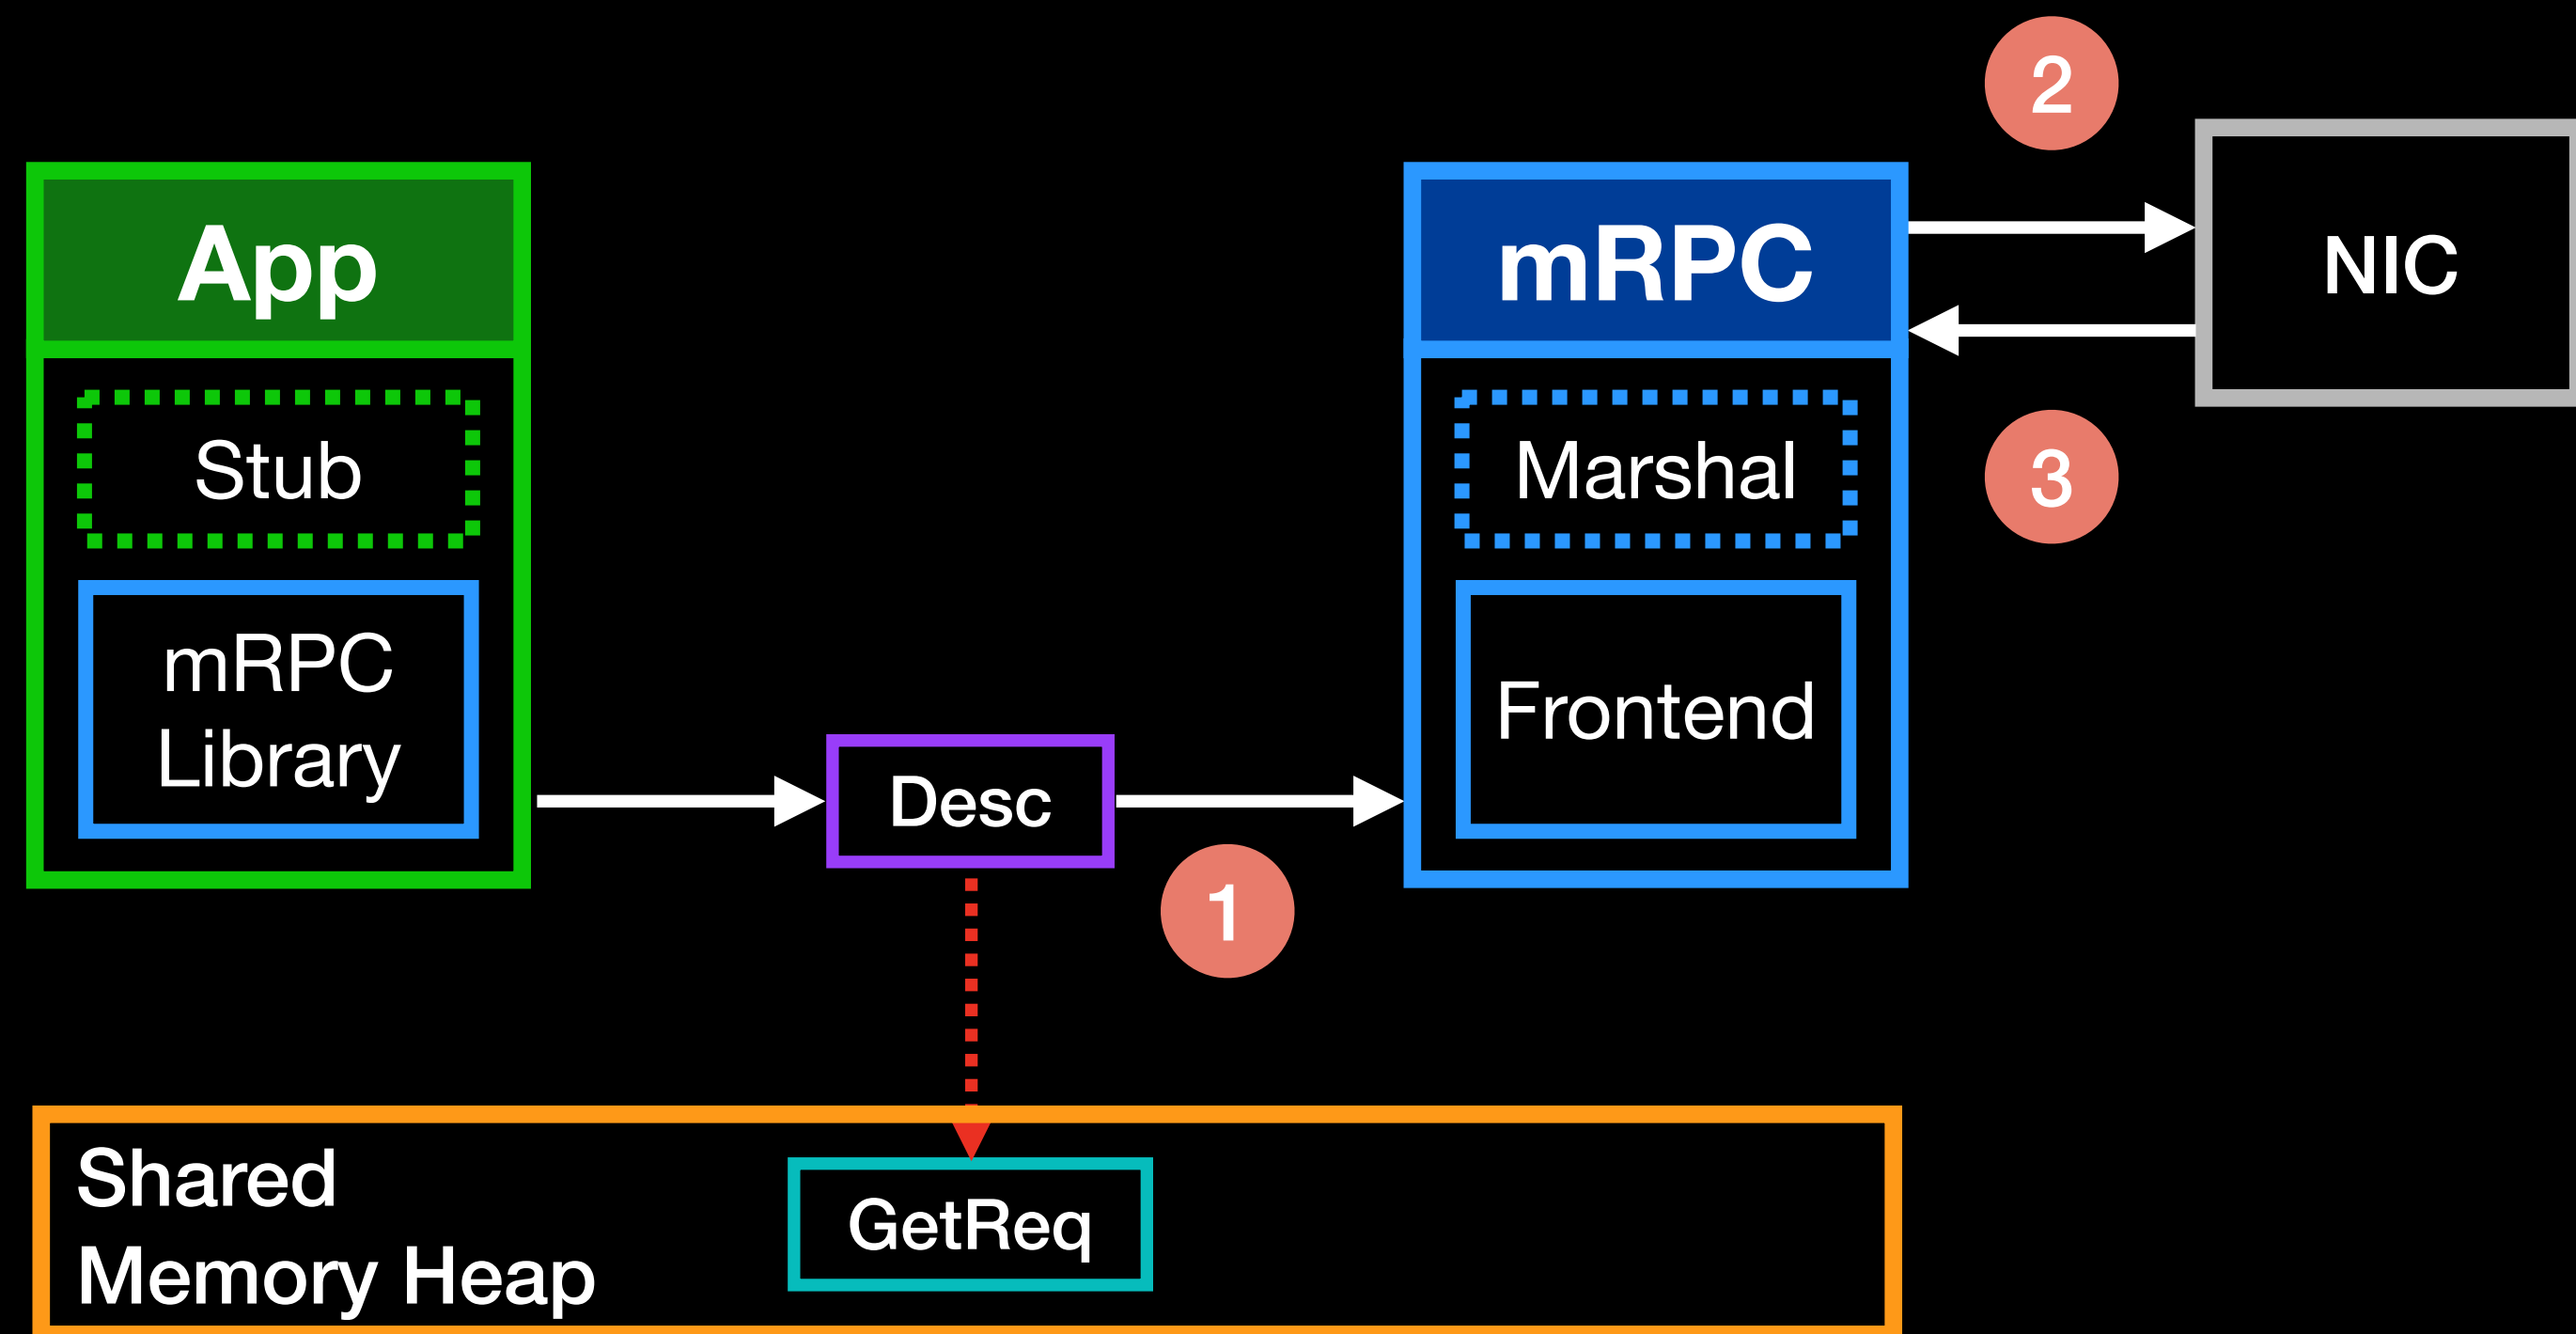
```
Service KVStore  
Func Get(GetReq) -> Entry
```



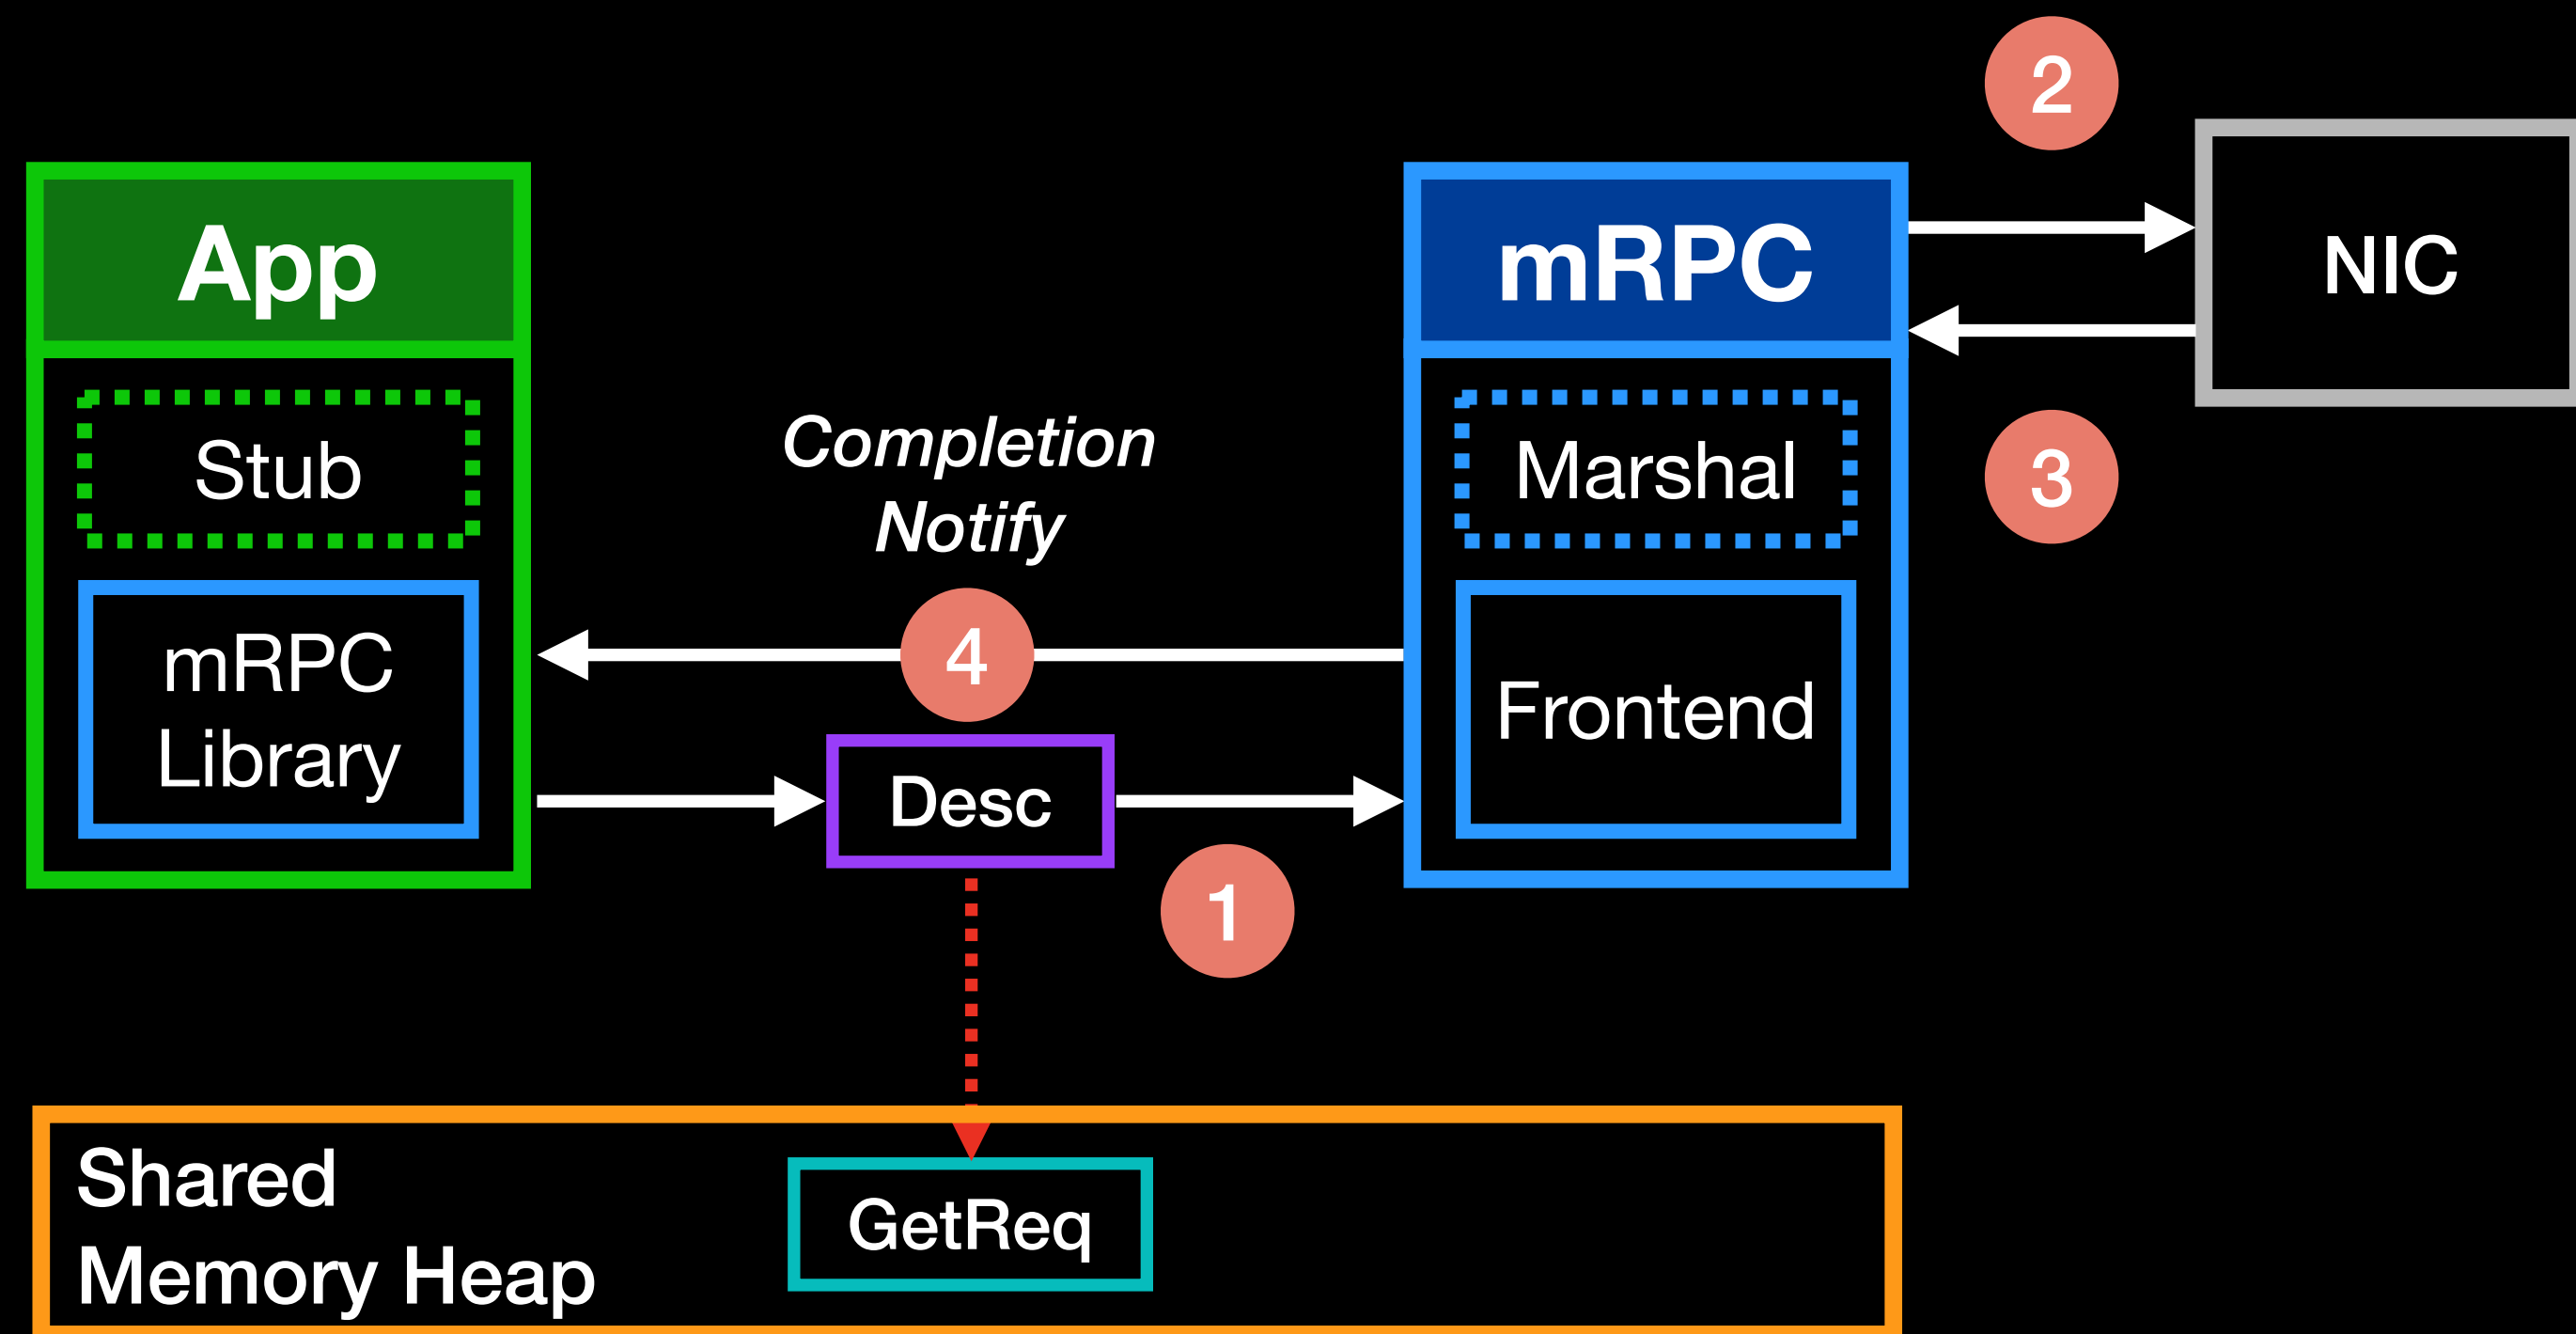
Memory Management (Outgoing Message)



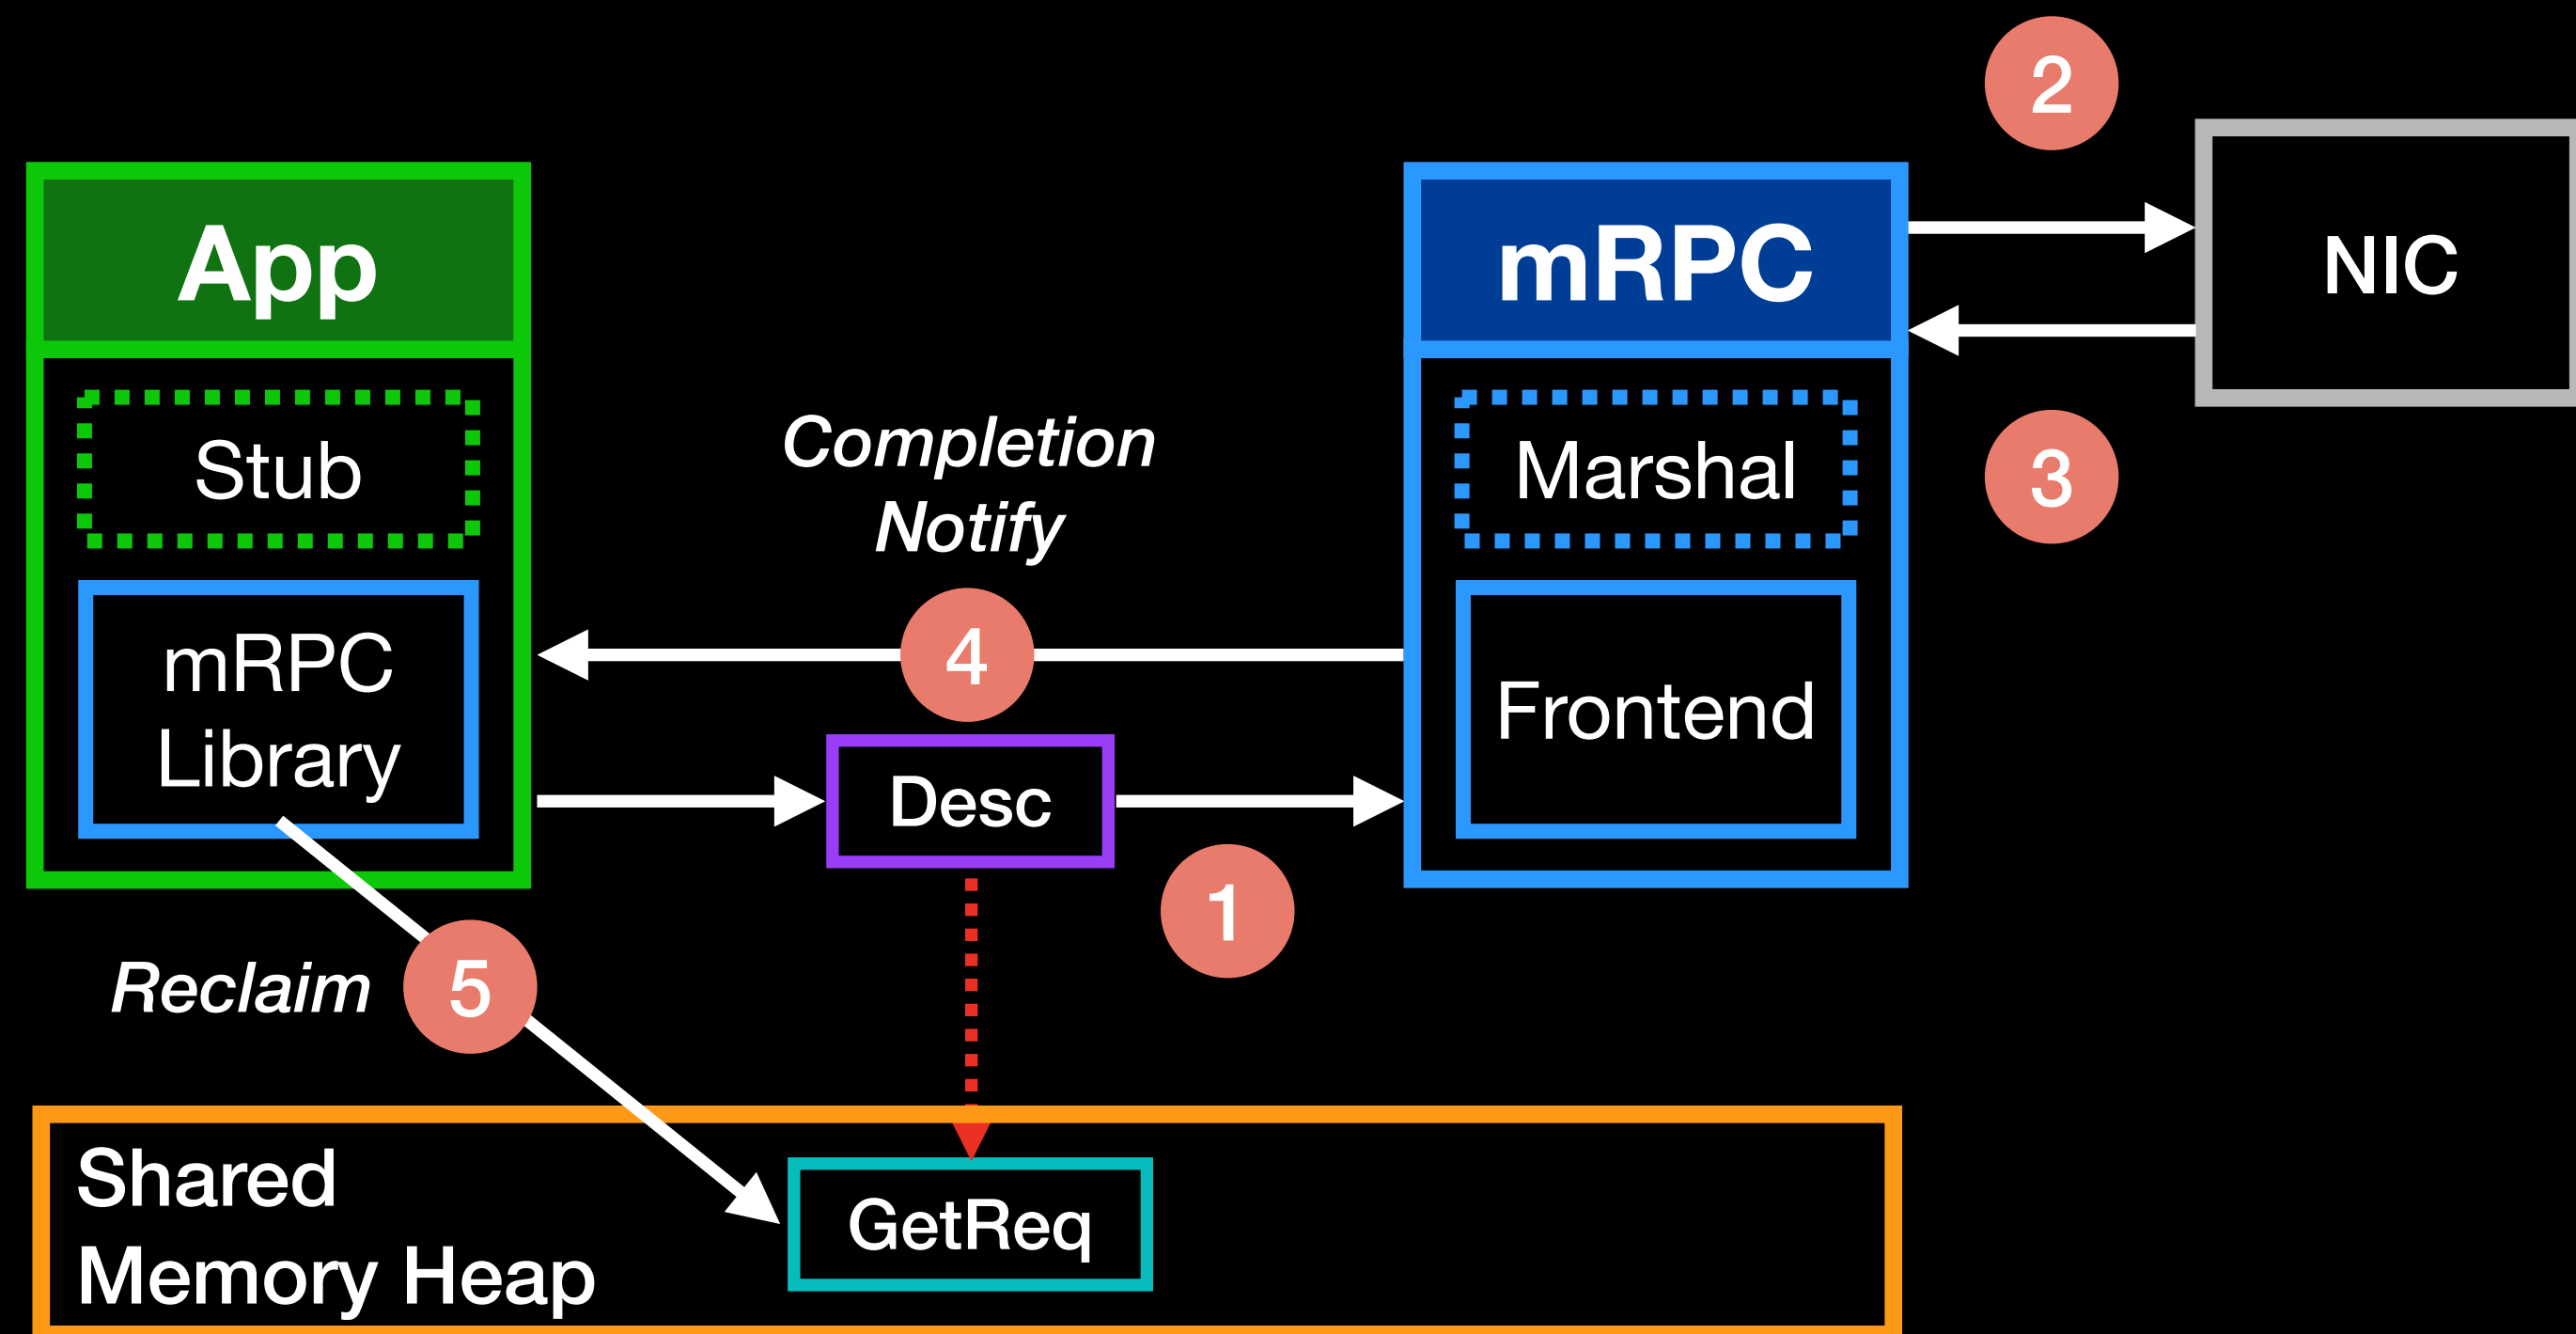
Memory Management (Outgoing Message)



Memory Management (Outgoing Message)





Memory Management (Outgoing Message)



Policy Enforcement



Processing Flow / Time



Shared
Heap

Policy Enforcement

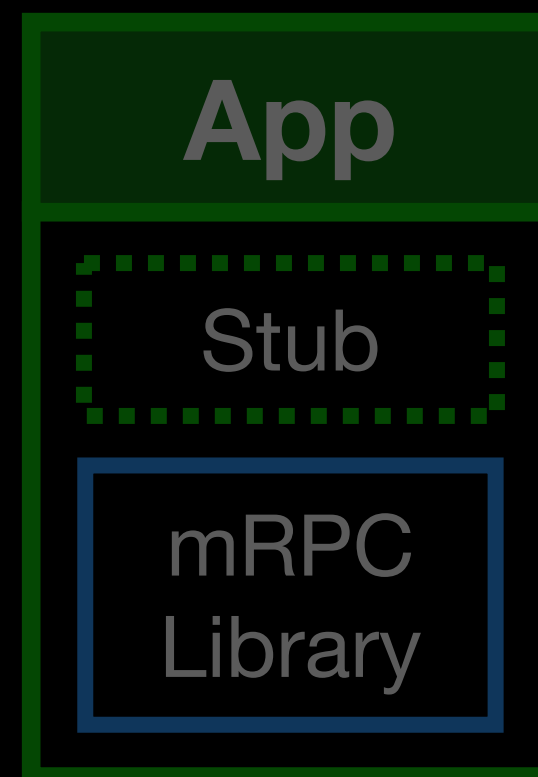
Processing Flow / Time




Shared
Heap

Policy Enforcement

Processing Flow / Time

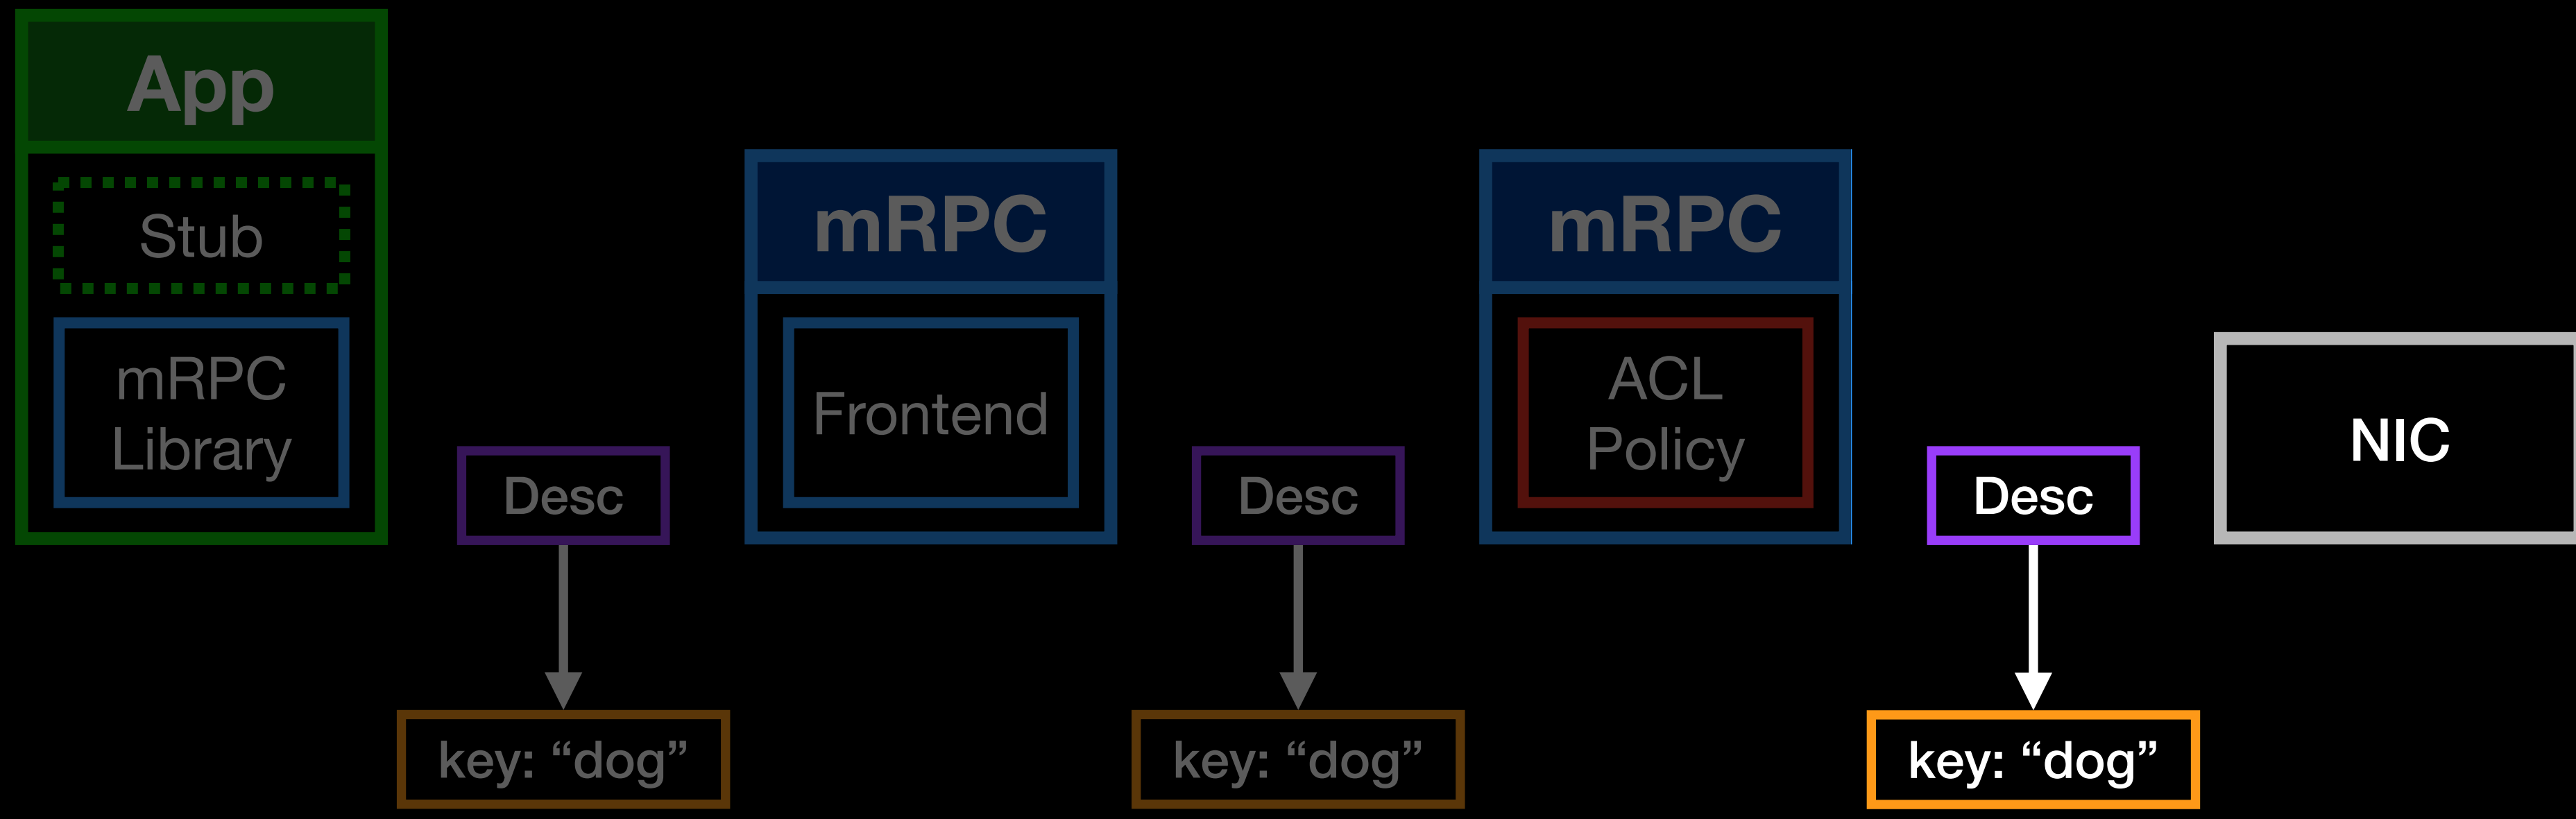


Shared
Heap



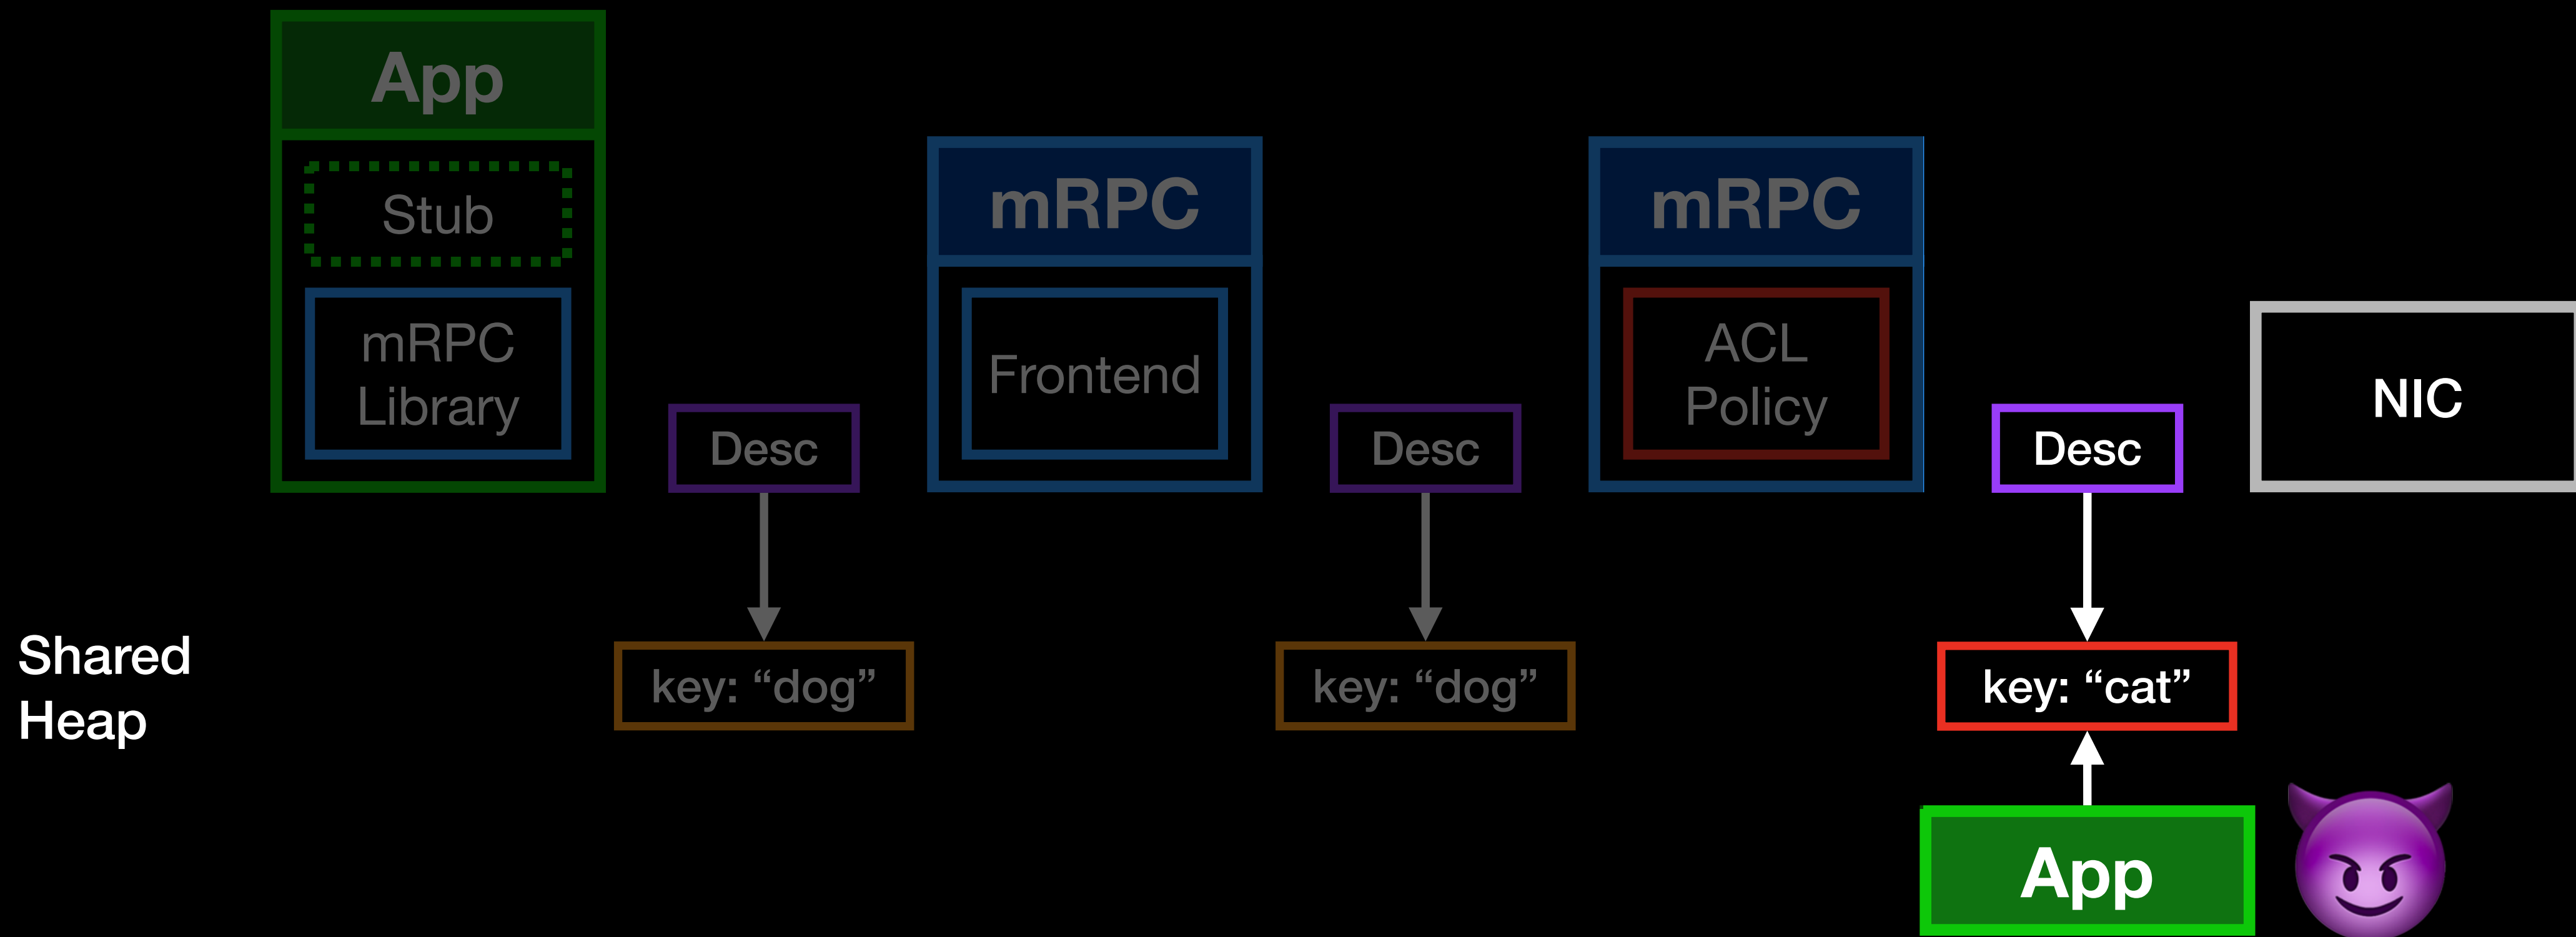
Policy Enforcement

Processing Flow / Time



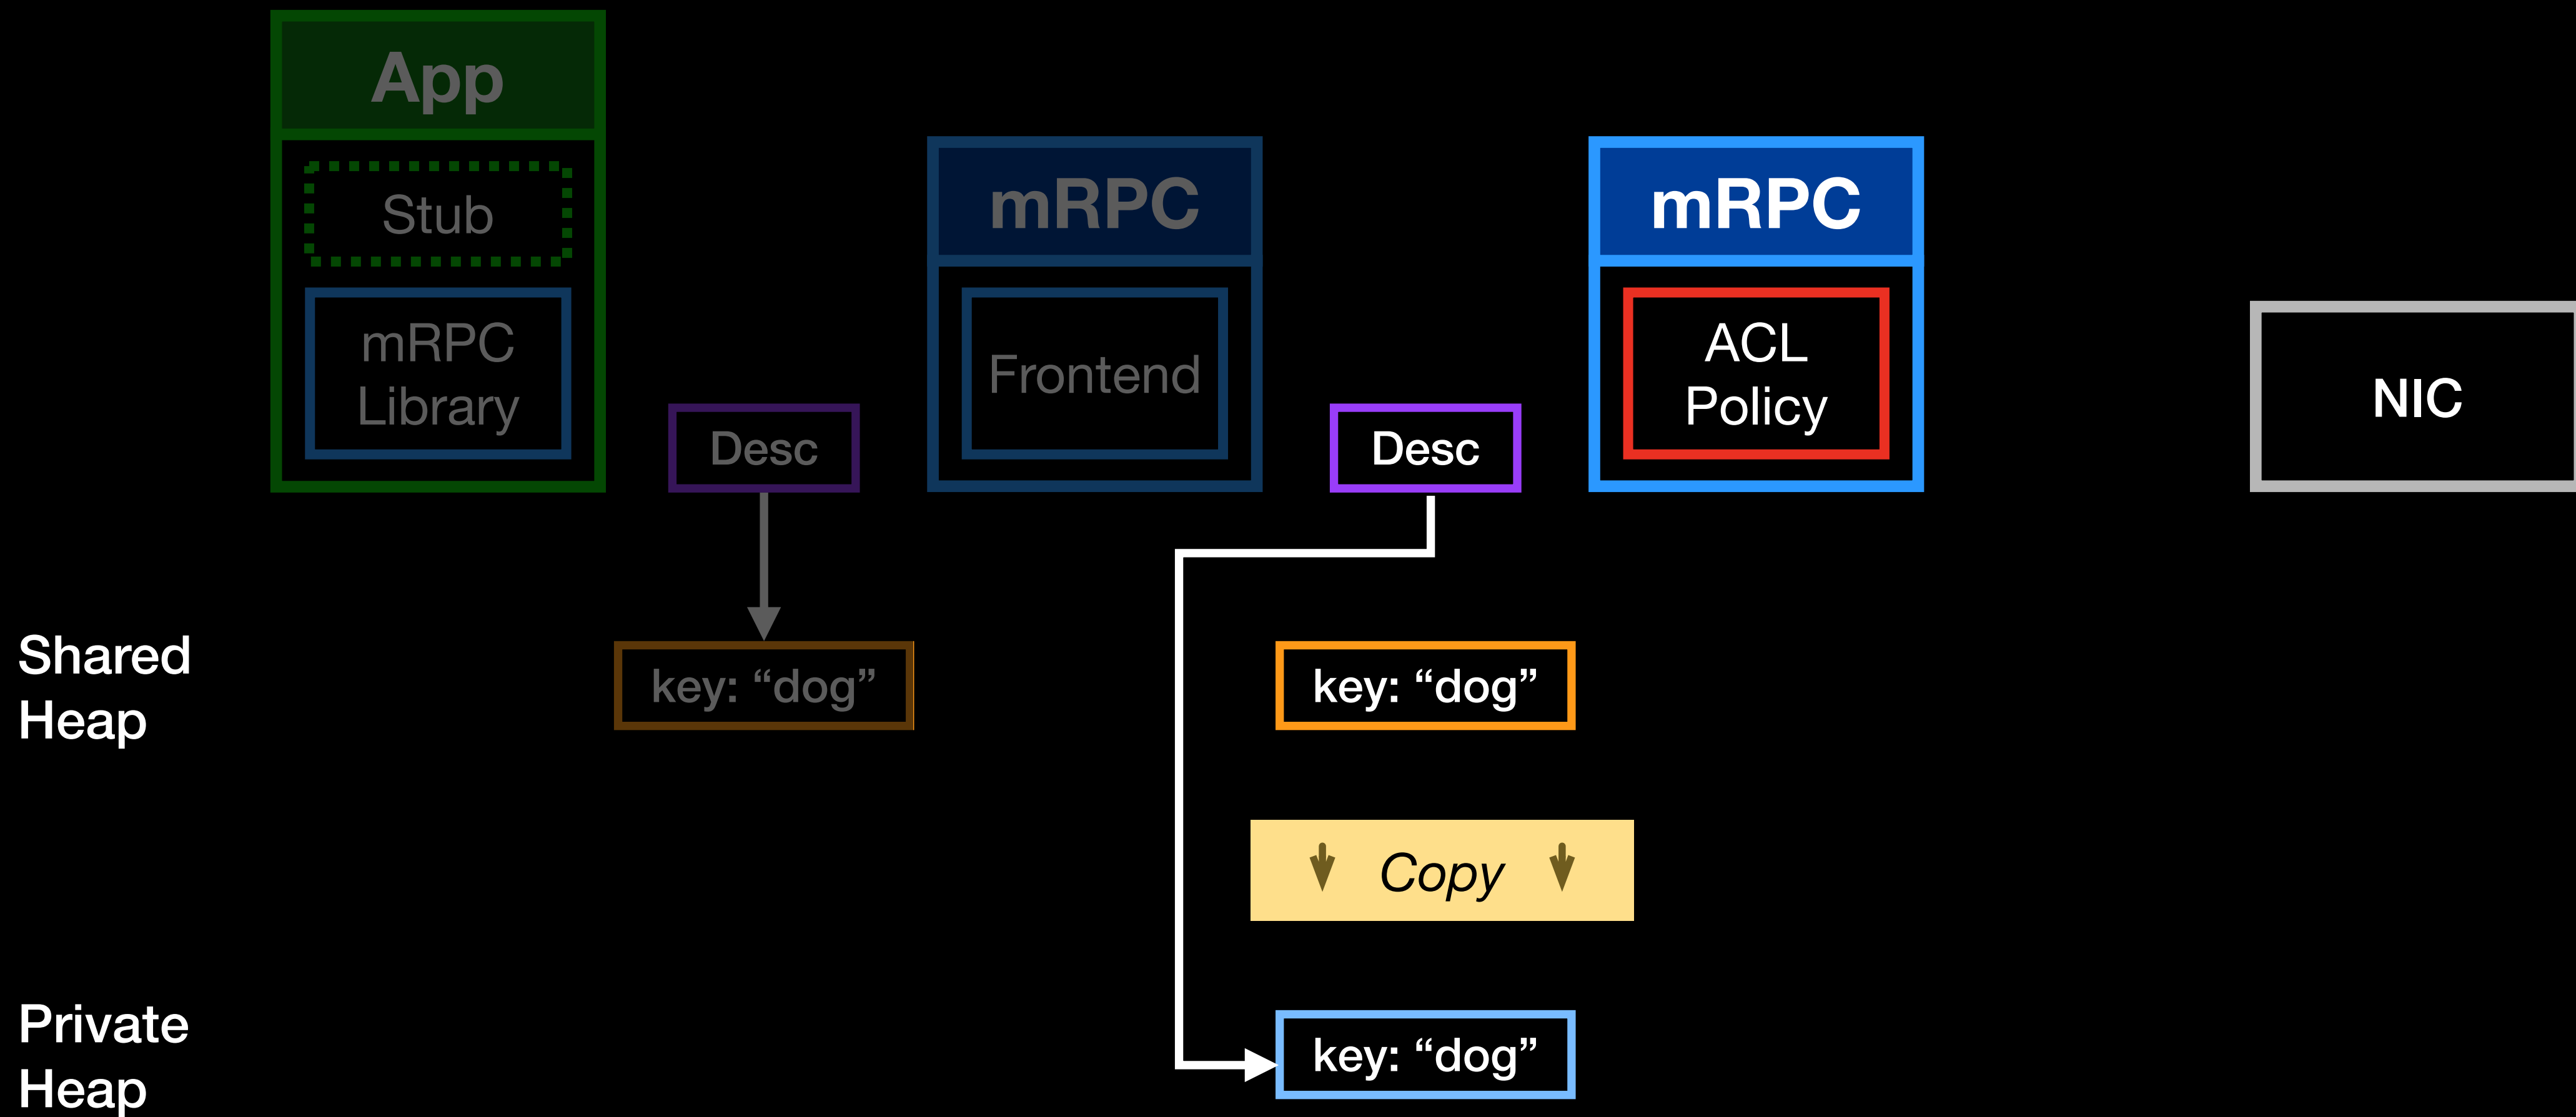
Policy Enforcement

Processing Flow / Time



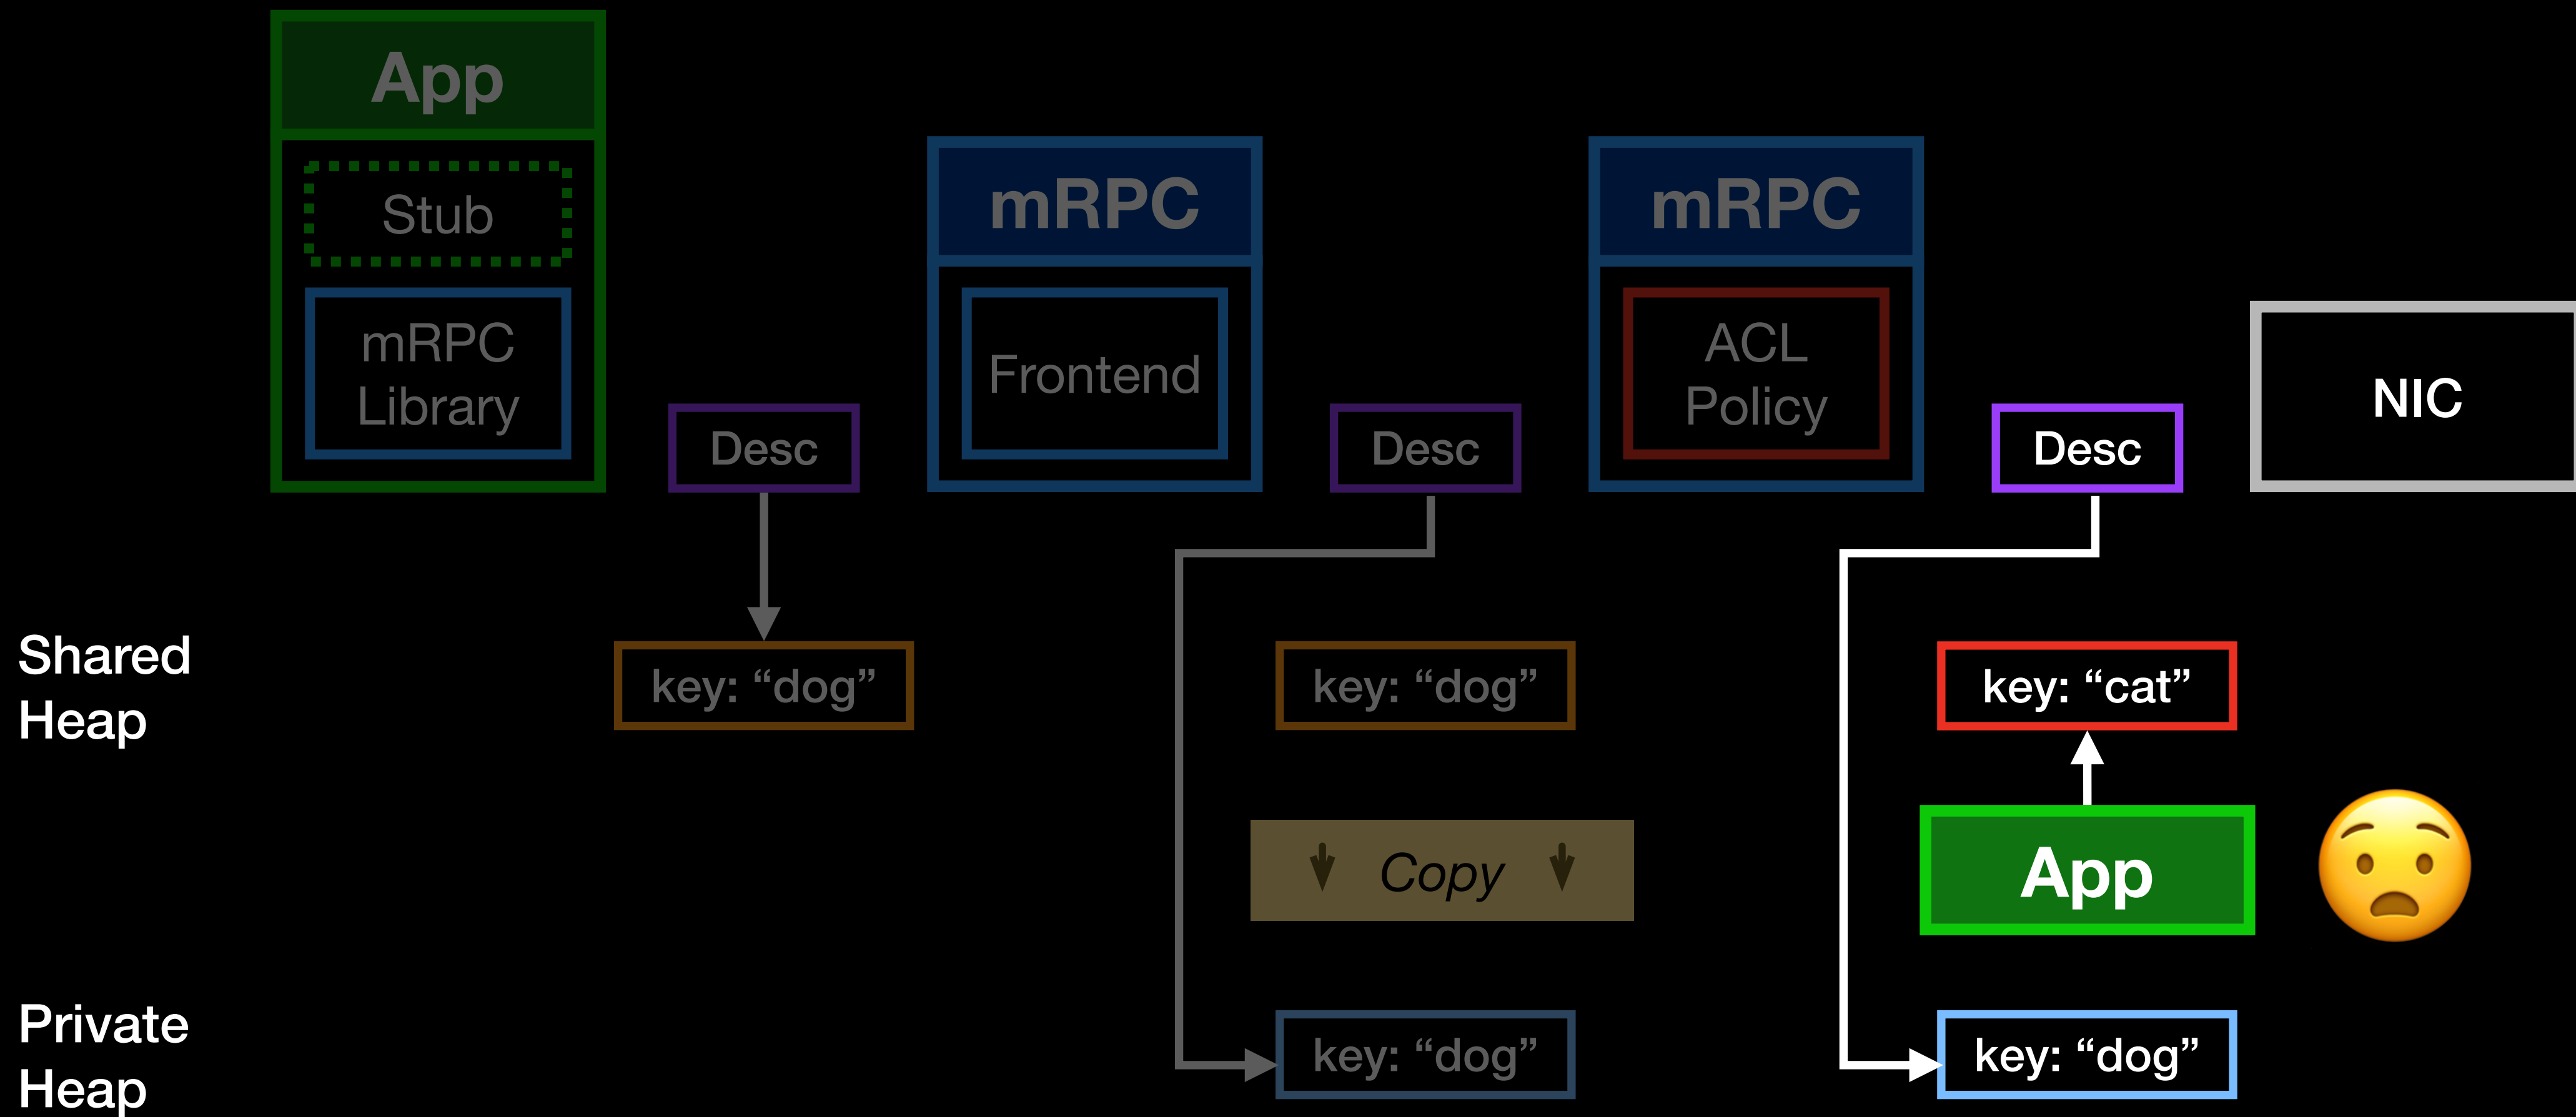
Policy Enforcement

Processing Flow / Time



Policy Enforcement

Processing Flow / Time



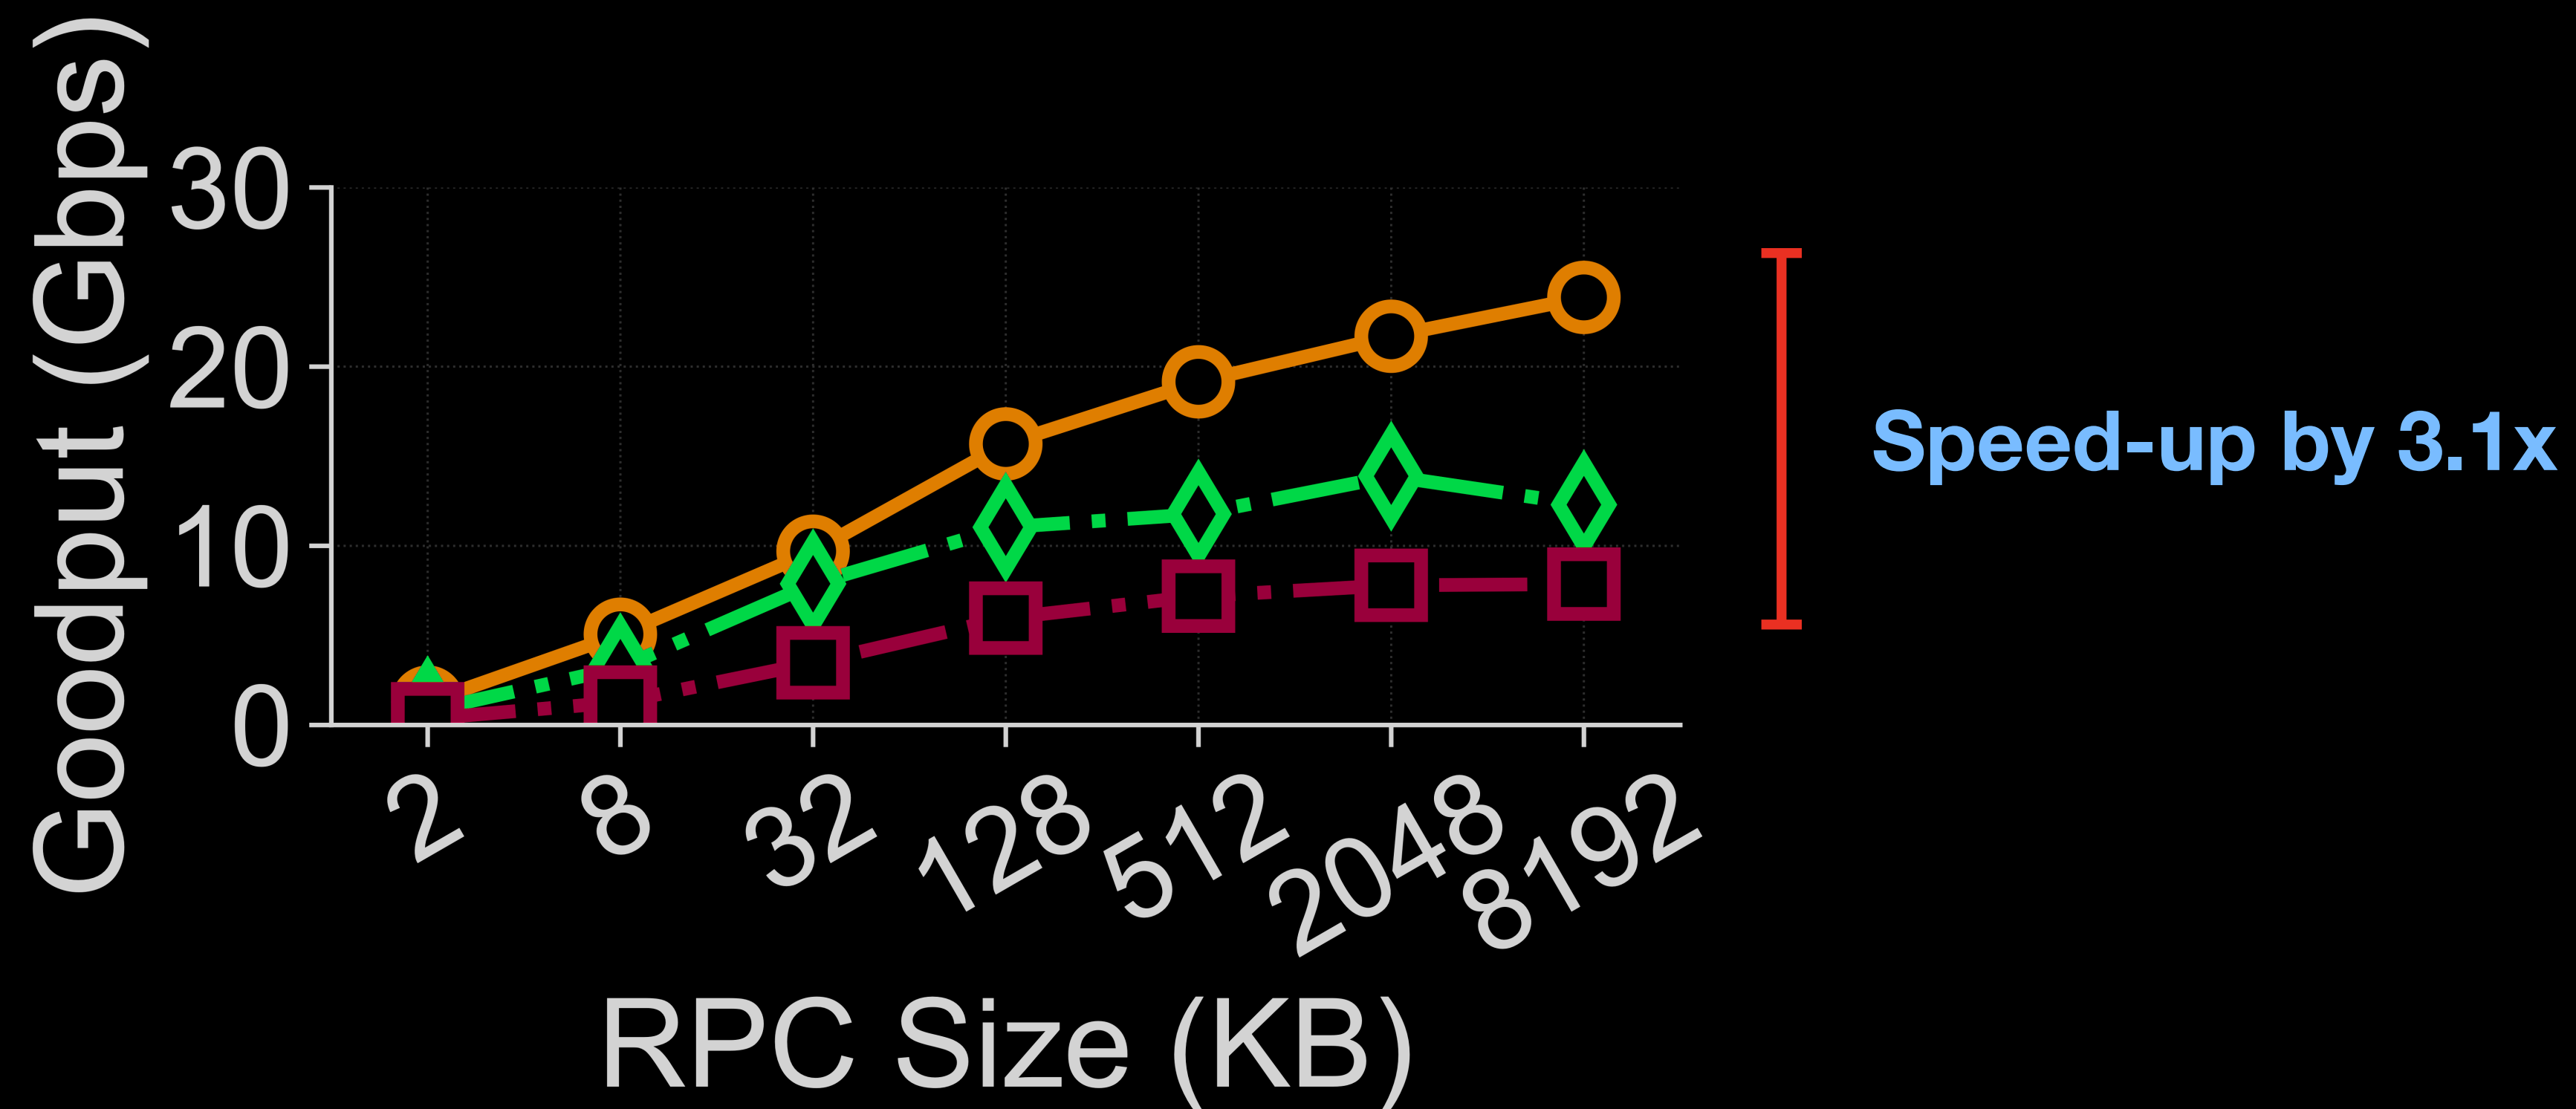
Evaluation

Does mRPC deliver smaller latency and higher goodput compared to existing solutions?


Does mRPC enforce policy efficiently?

Can mRPC improve real-world application's performance?

Evaluation: Large RPC Goodput



- TCP transport
- Keep 128 concurrent RPCs to hide latency




Evaluation: Small RPC Latency

	Median Latency (μ s)	P99 Latency (μ s)
eRPC	3.6	4.1
mRPC	7.6	8.7
eRPC + Proxy	11.3	15.6
mRPC + NullPolicy	7.9	9.1

I Speed-up by 1.7x

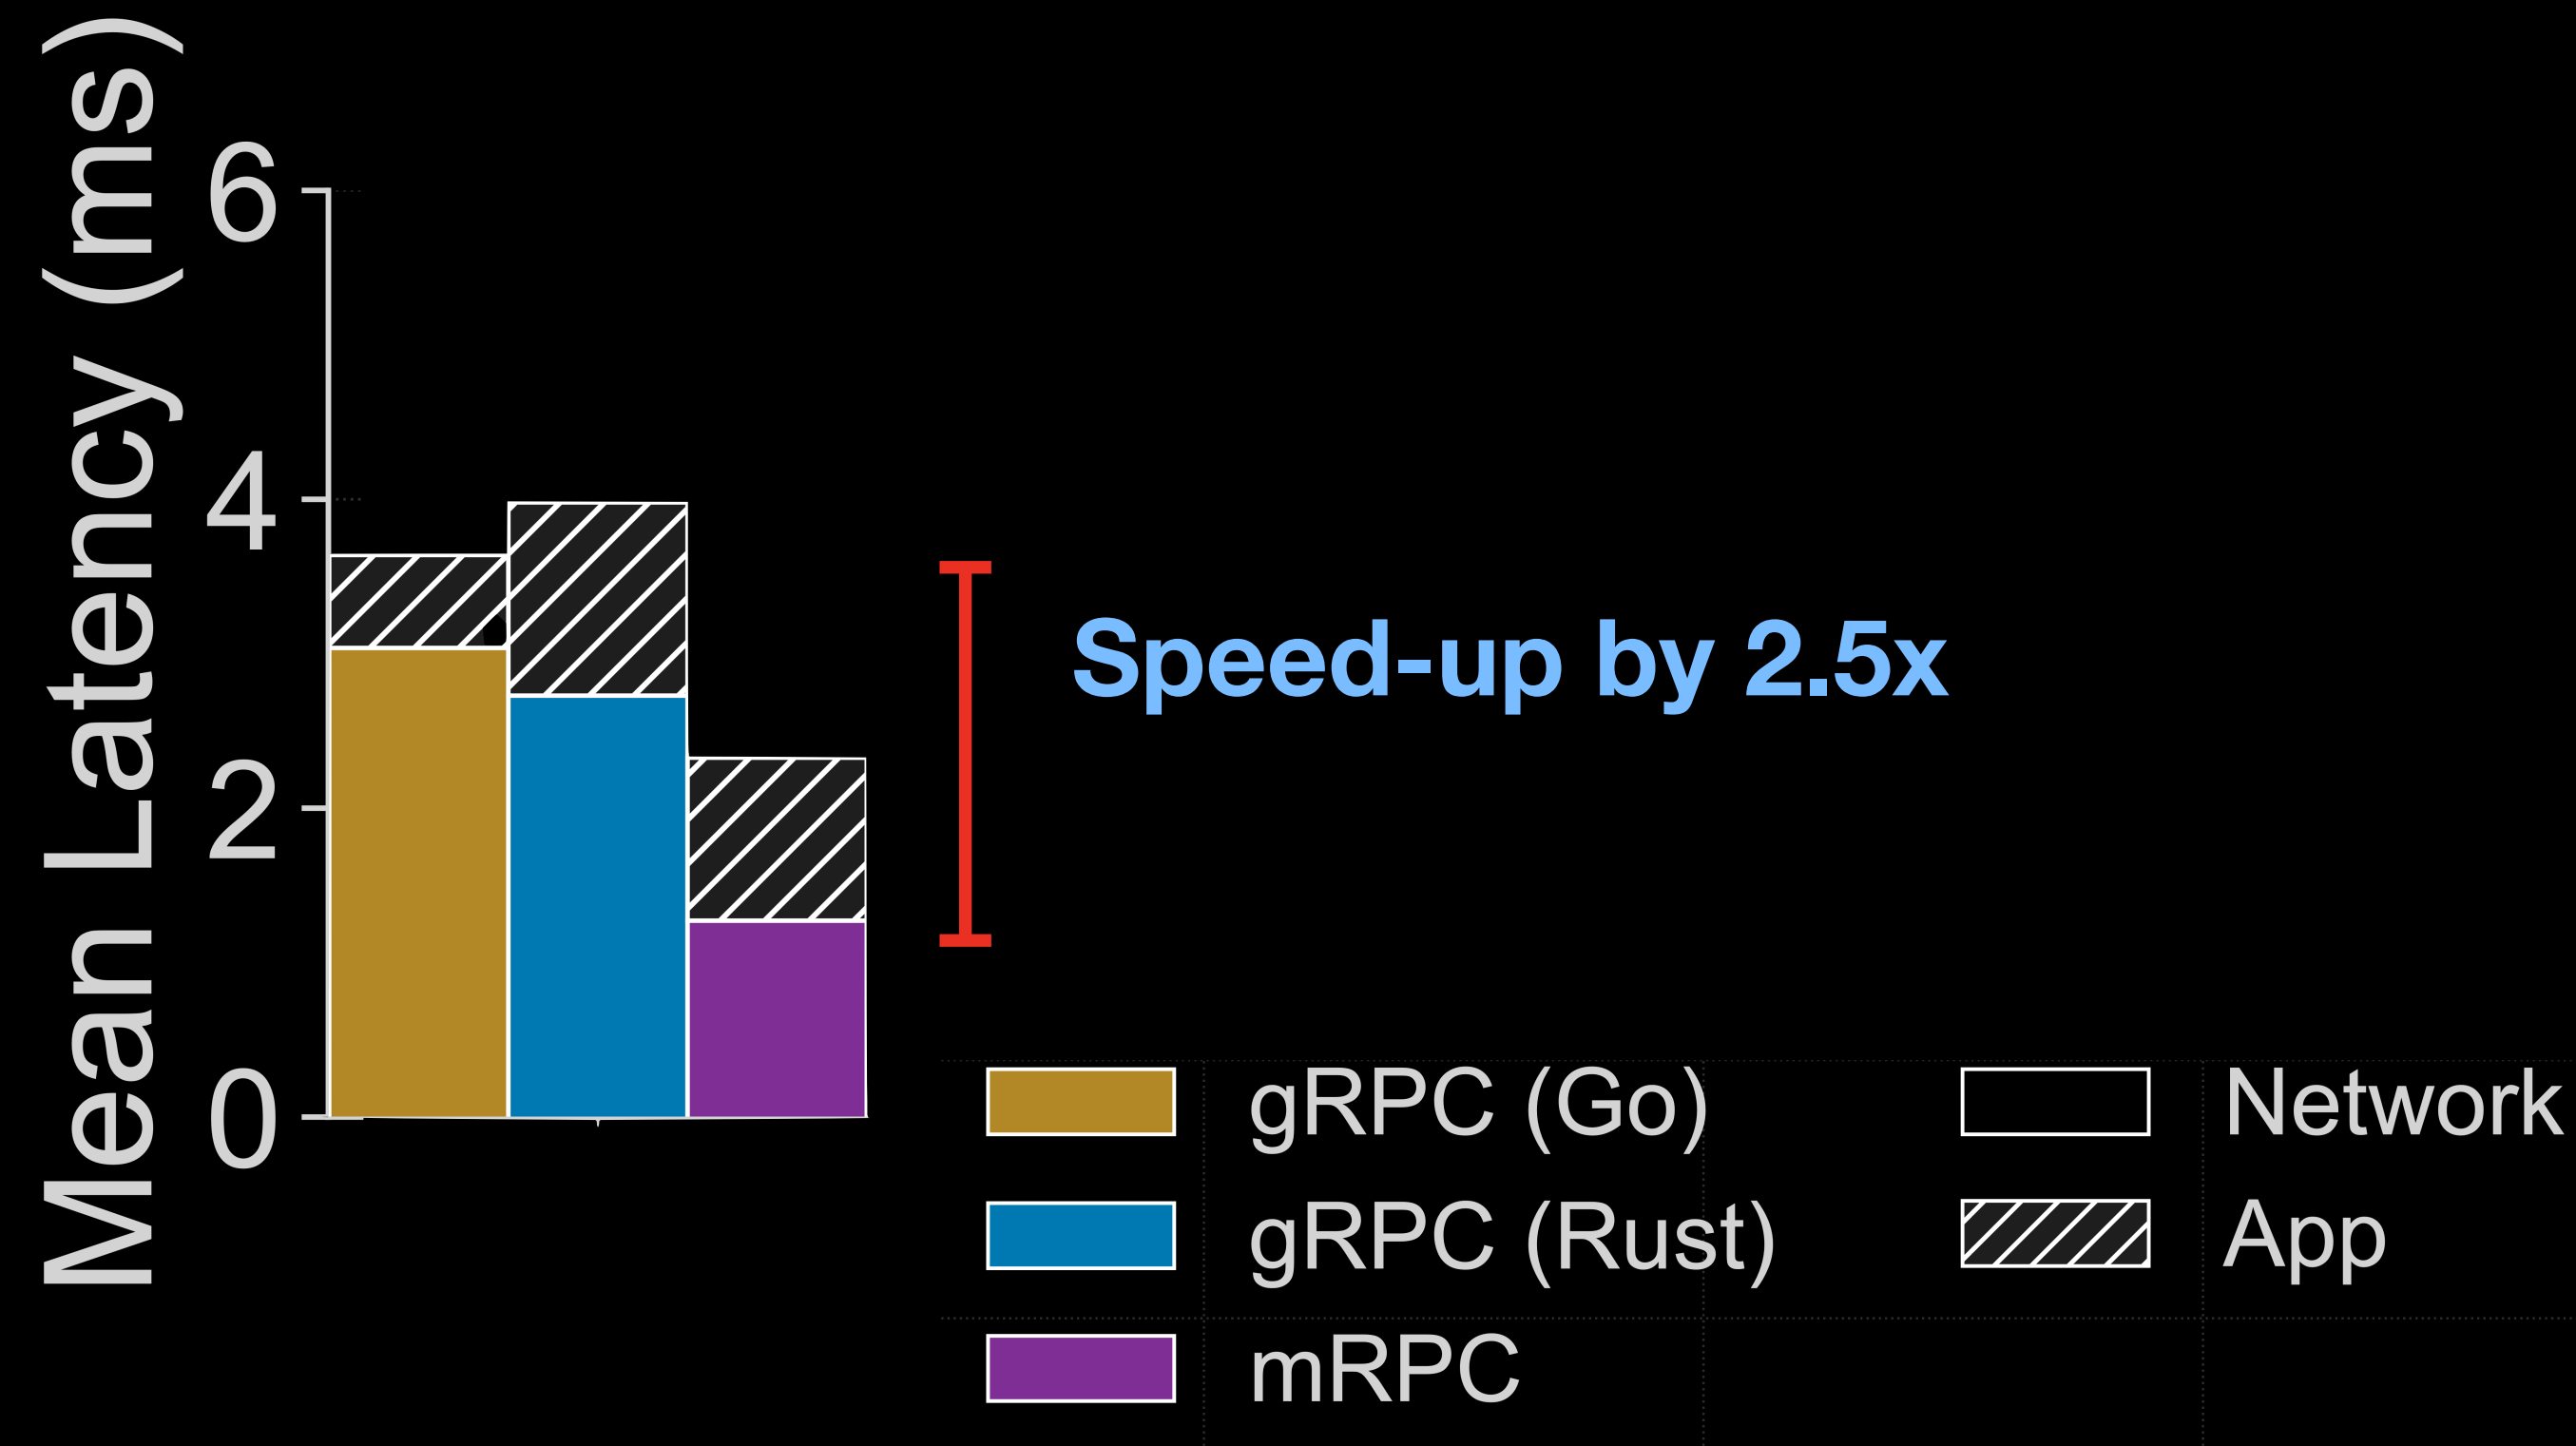
- RDMA transport
- 64-byte RPC requests, 8-byte replies

Evaluation: Policy Enforcement



- Filter RPCs based on string matching on one field
- 1% requests will not pass

Evaluation: DeathStarBench



- TCP transport
- Measured over 250 secs @ 20 reqs/sec

Summary

RPC-as-a-library cannot meet both **manageability** and **efficiency**

mRPC: Reimagined RPC as a managed system service

Efficient policy enforcement

Upgrade of RPC implementation *without* shutting down user applications



<https://github.com/phoenix-dataplane/phoenix>