MIDDLEWARE TRENDS AND MARKET LEADERS 2011

A. Dworak

P. Charrue, F. Ehm, W. Sliwinski, M. Sobczak

Controls Group, Beams Department

CERN

ICALEPCS 2011

14TH OCTOBER 2011, GRENOBLE

Middleware definition and types

Middleware – software that allows communication between other software components running on one or more machines.

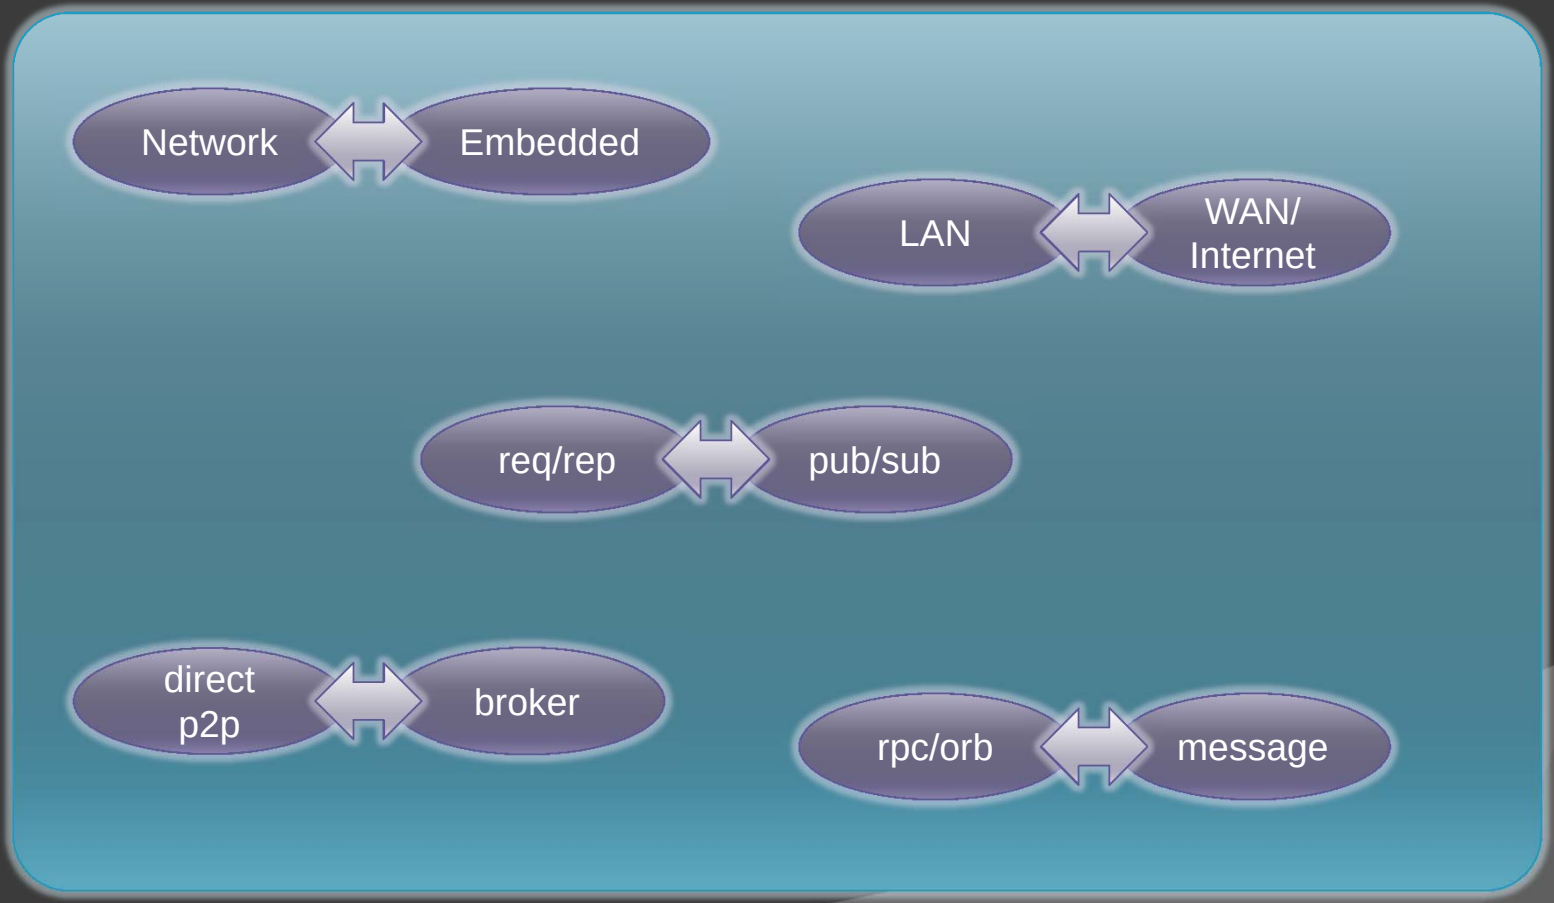

CERN Controls System

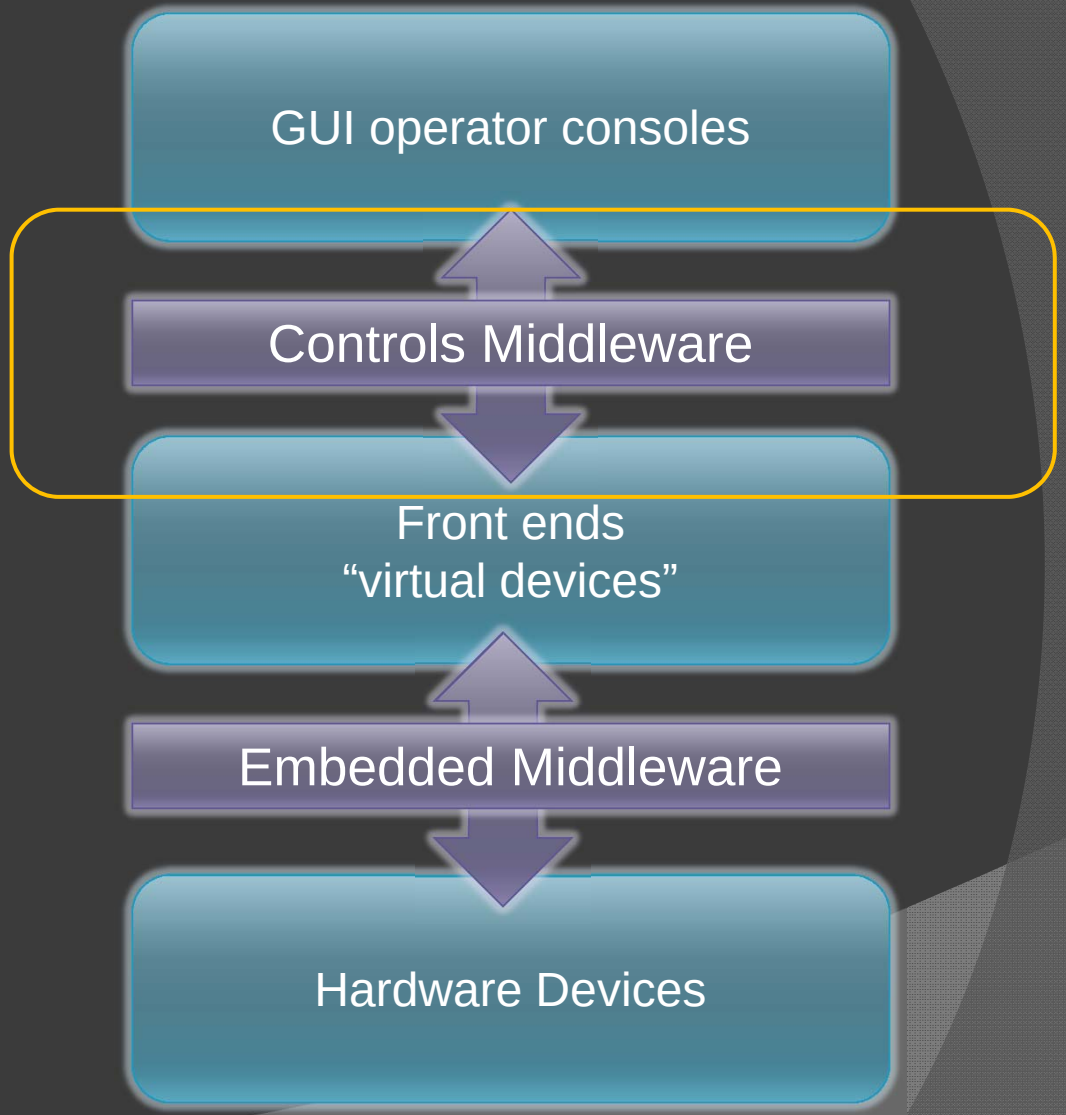

CERN Controls Middleware

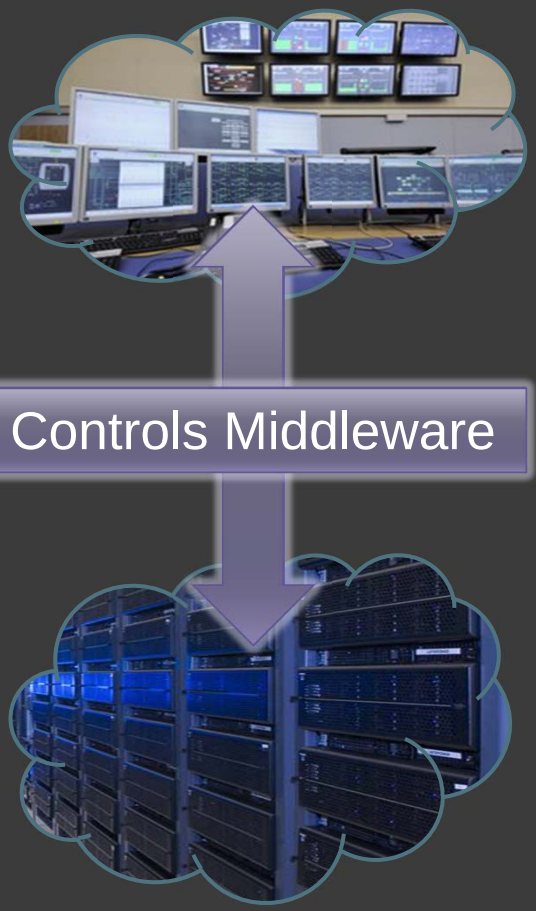

Linux/Windows

C++/Java process

Controls Middleware API

CORBA

TCP/IP

CORBA

Controls Middleware API

C++ process

RT Linux/LynxOS

CERN Controls Middleware

Problems:

- C++ and Java implementations differ ☹
- Heavy in memory usage ☹
- Complex error prone API ☹
- No direct support for pub/sub ☹
- Blocking issues (JacORB) ☹
- Shrinking community ☹
- Lack of new releases and bug fixes ☹

Long LHC shutdown at the end of 2012
A unique possibility!
Instead of patching let's provide a new one!

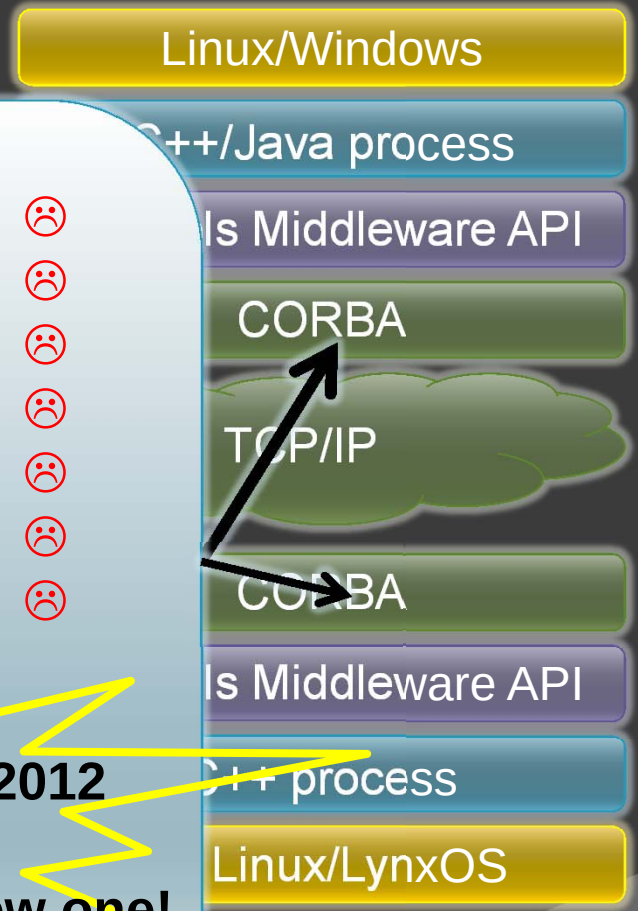

CERN Middleware Requirements

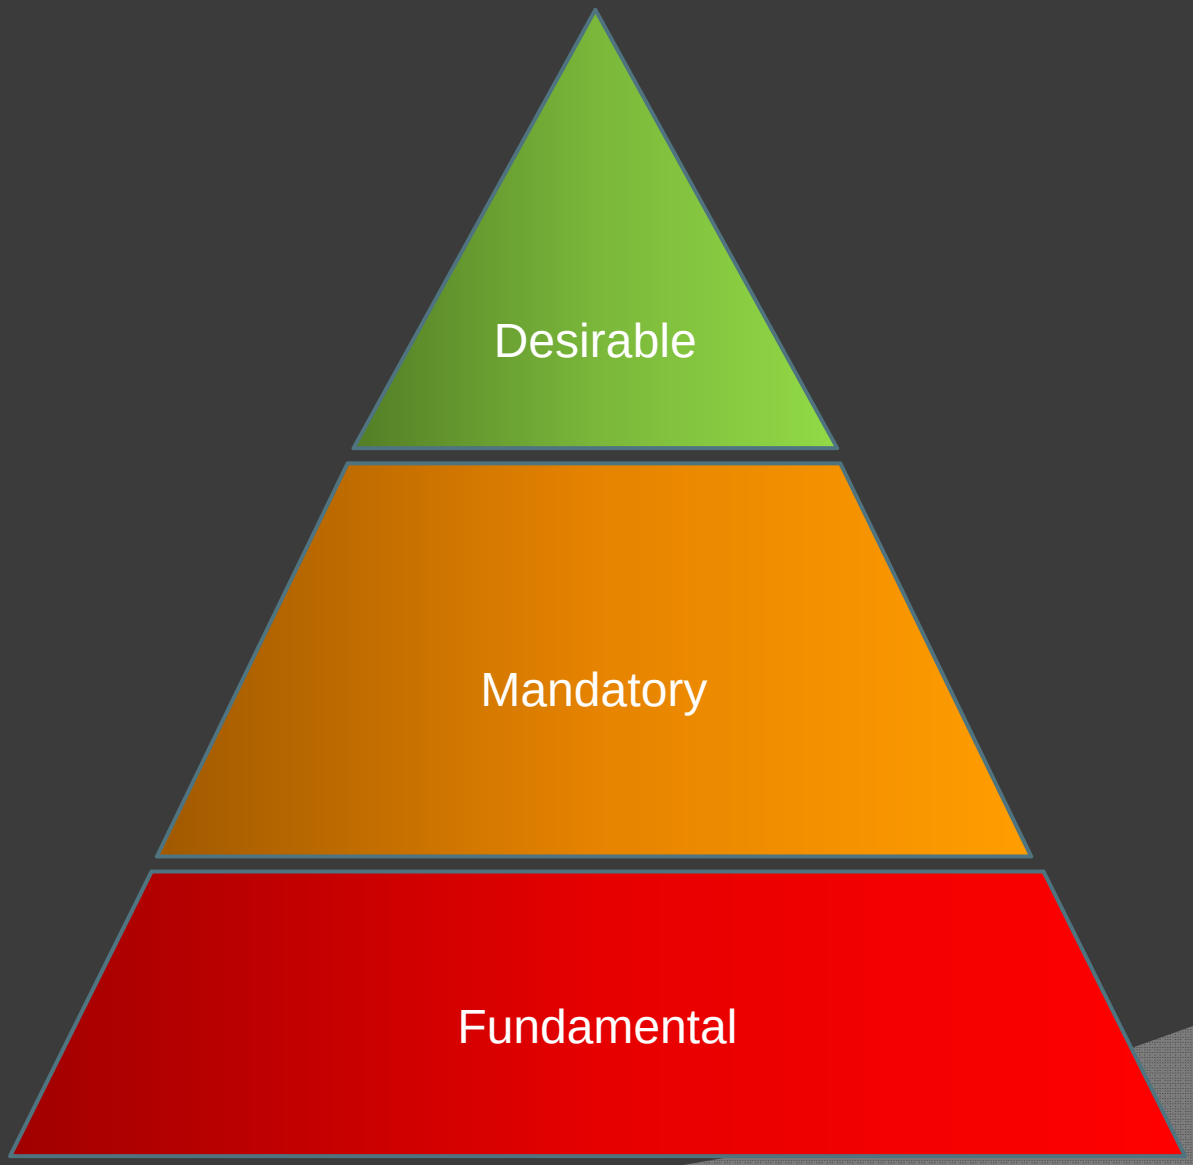

CERN Middleware Requirements

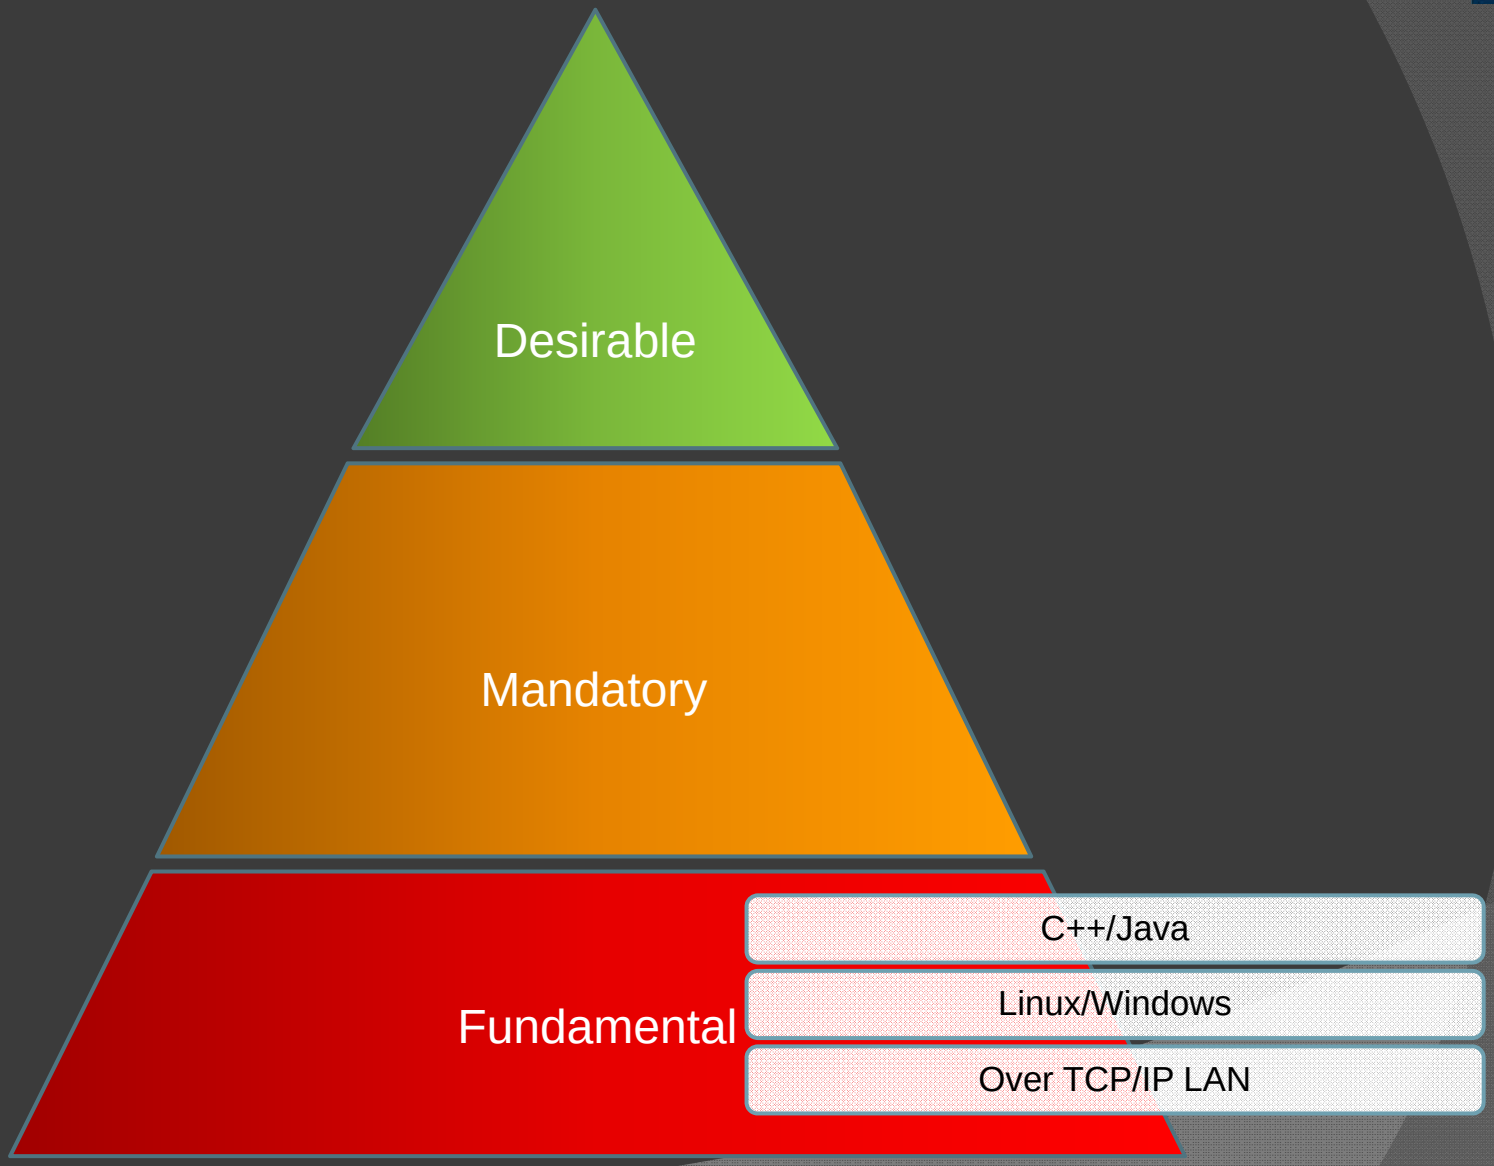

CERN Middleware Requirements

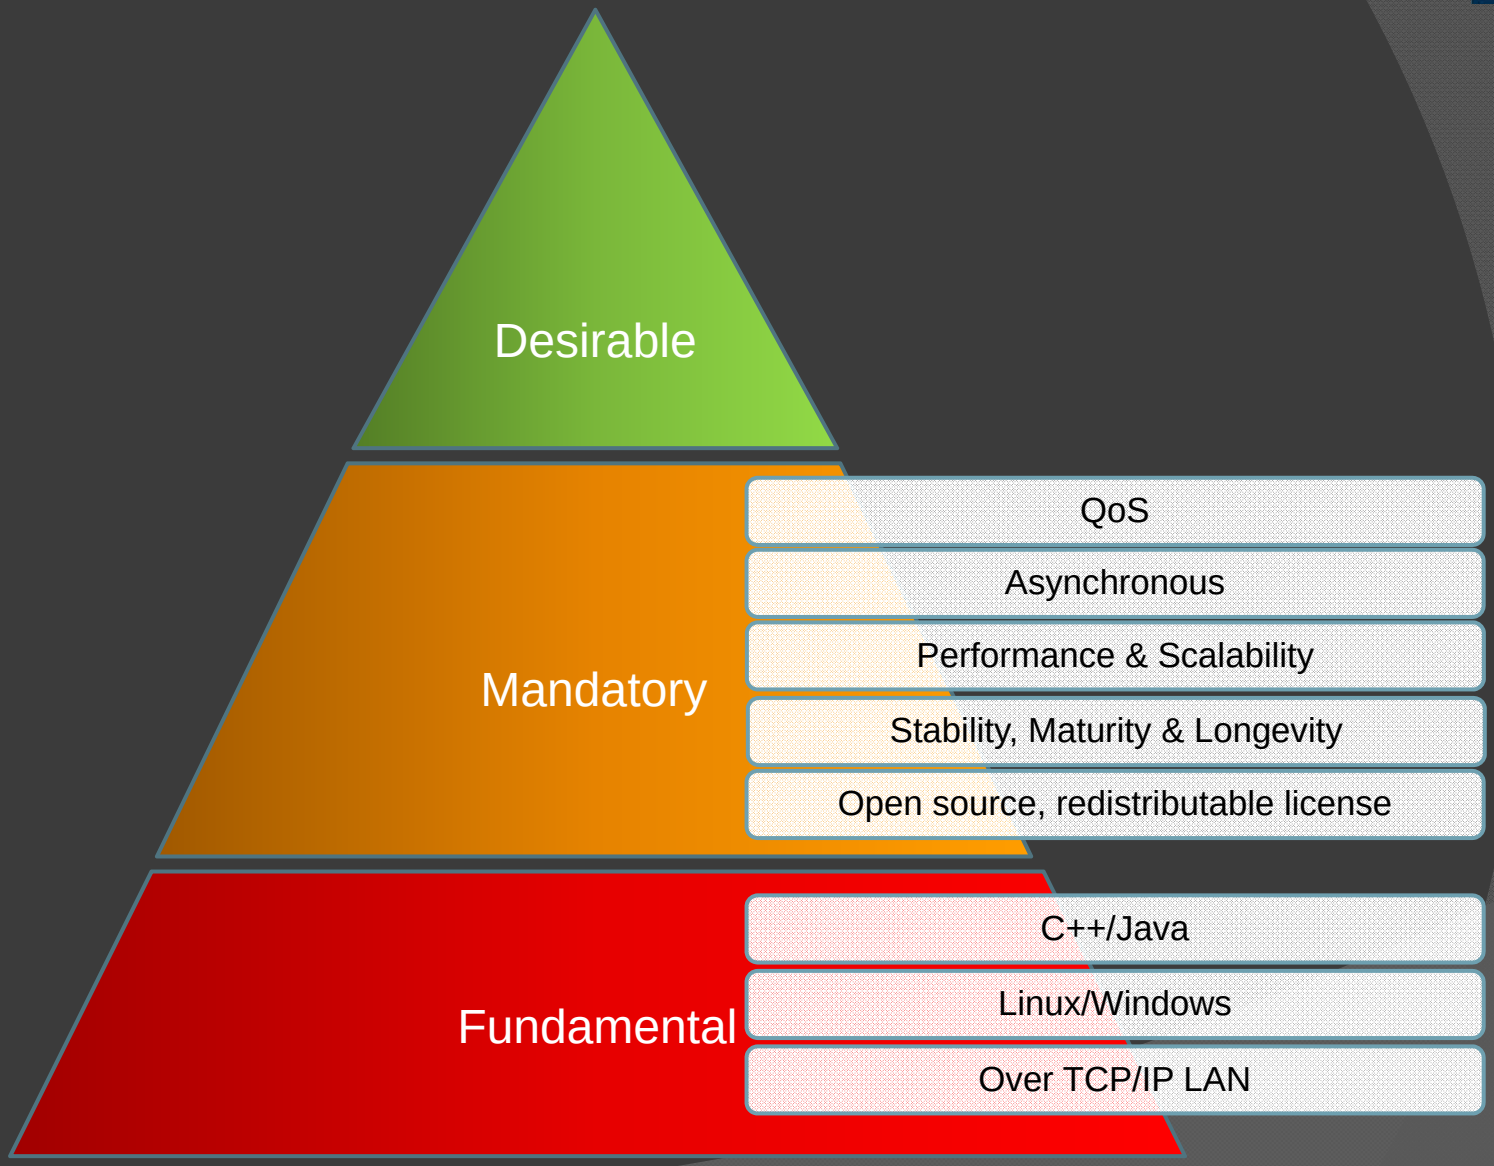

CERN Middleware Requirements

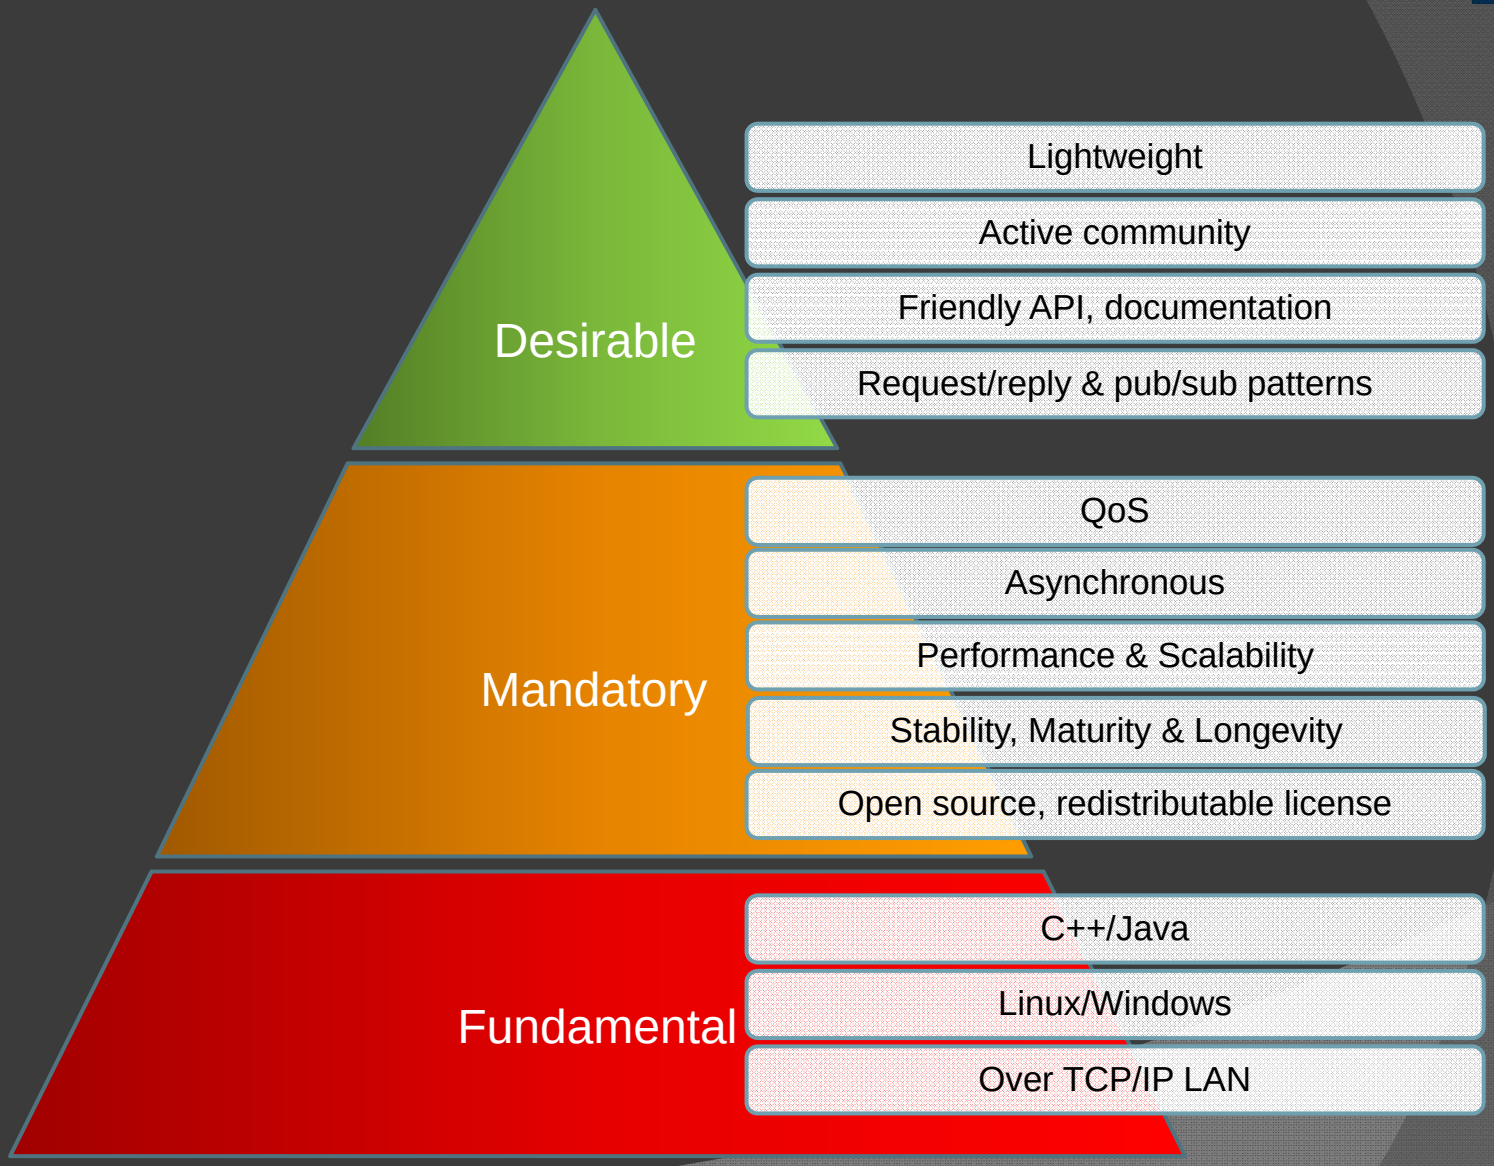

How did we evaluate –> our criteria

Appearance

- **Creators**
  - specification
  - documentation
- **Users**
  - forums
  - bug reports
- **Internet**

Simple usage

- **Download**
  - licensing
- **Compile**
  - LynxOS & gcc 2.95
- **Run examples**

Testing

- Communication patterns
- Performance
- QoS
- Exceptional situations

CRITERIA

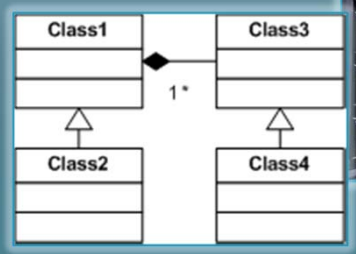

API, look & feel, documentation

resources, binary size, memory

Community, maturity

Communications patterns

QoS

performance

Surprising how many great new products!

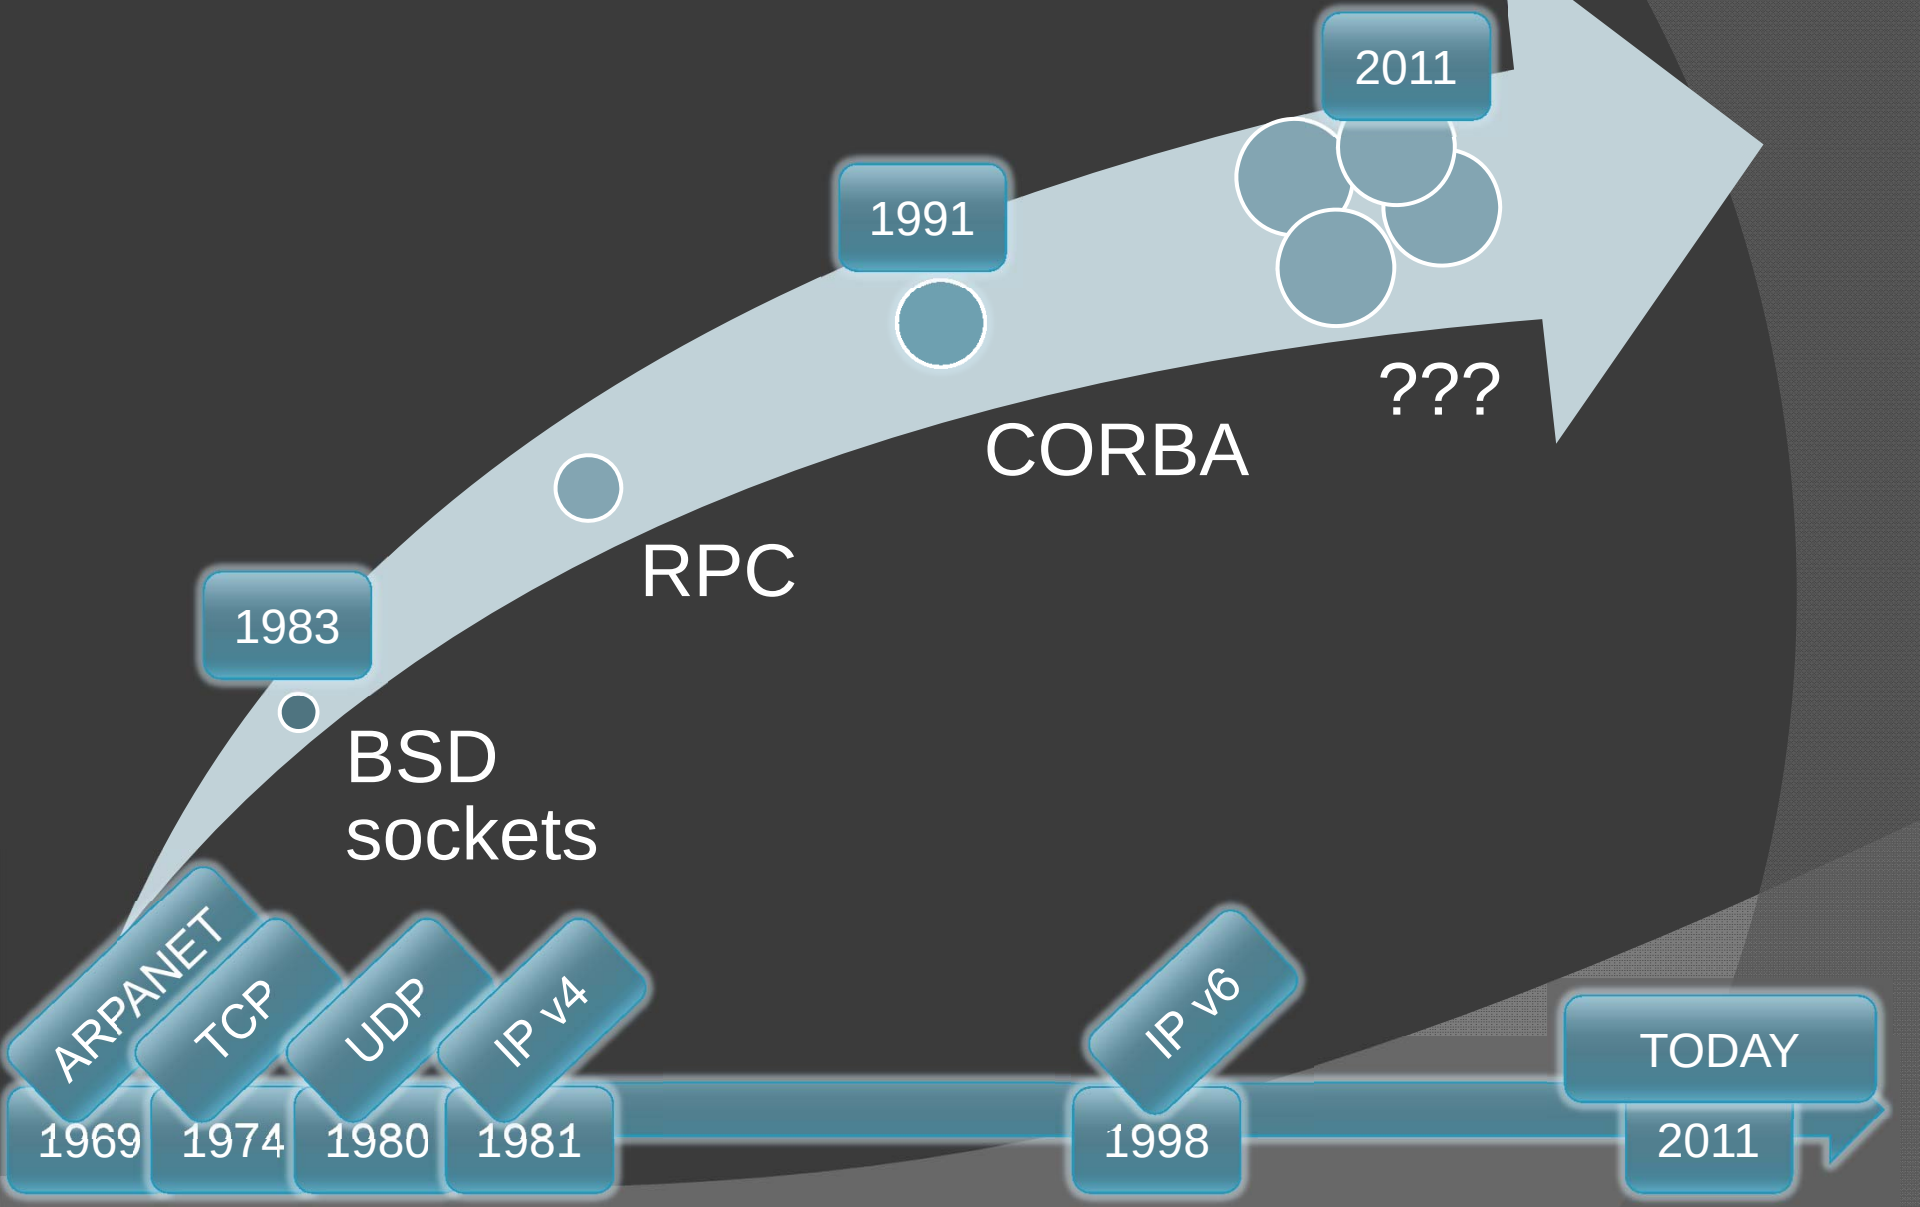

Evaluated middleware products

All opinions are based only on our knowledge and evaluation. Each of the products, depending on the requirements, may constitute a good solution.

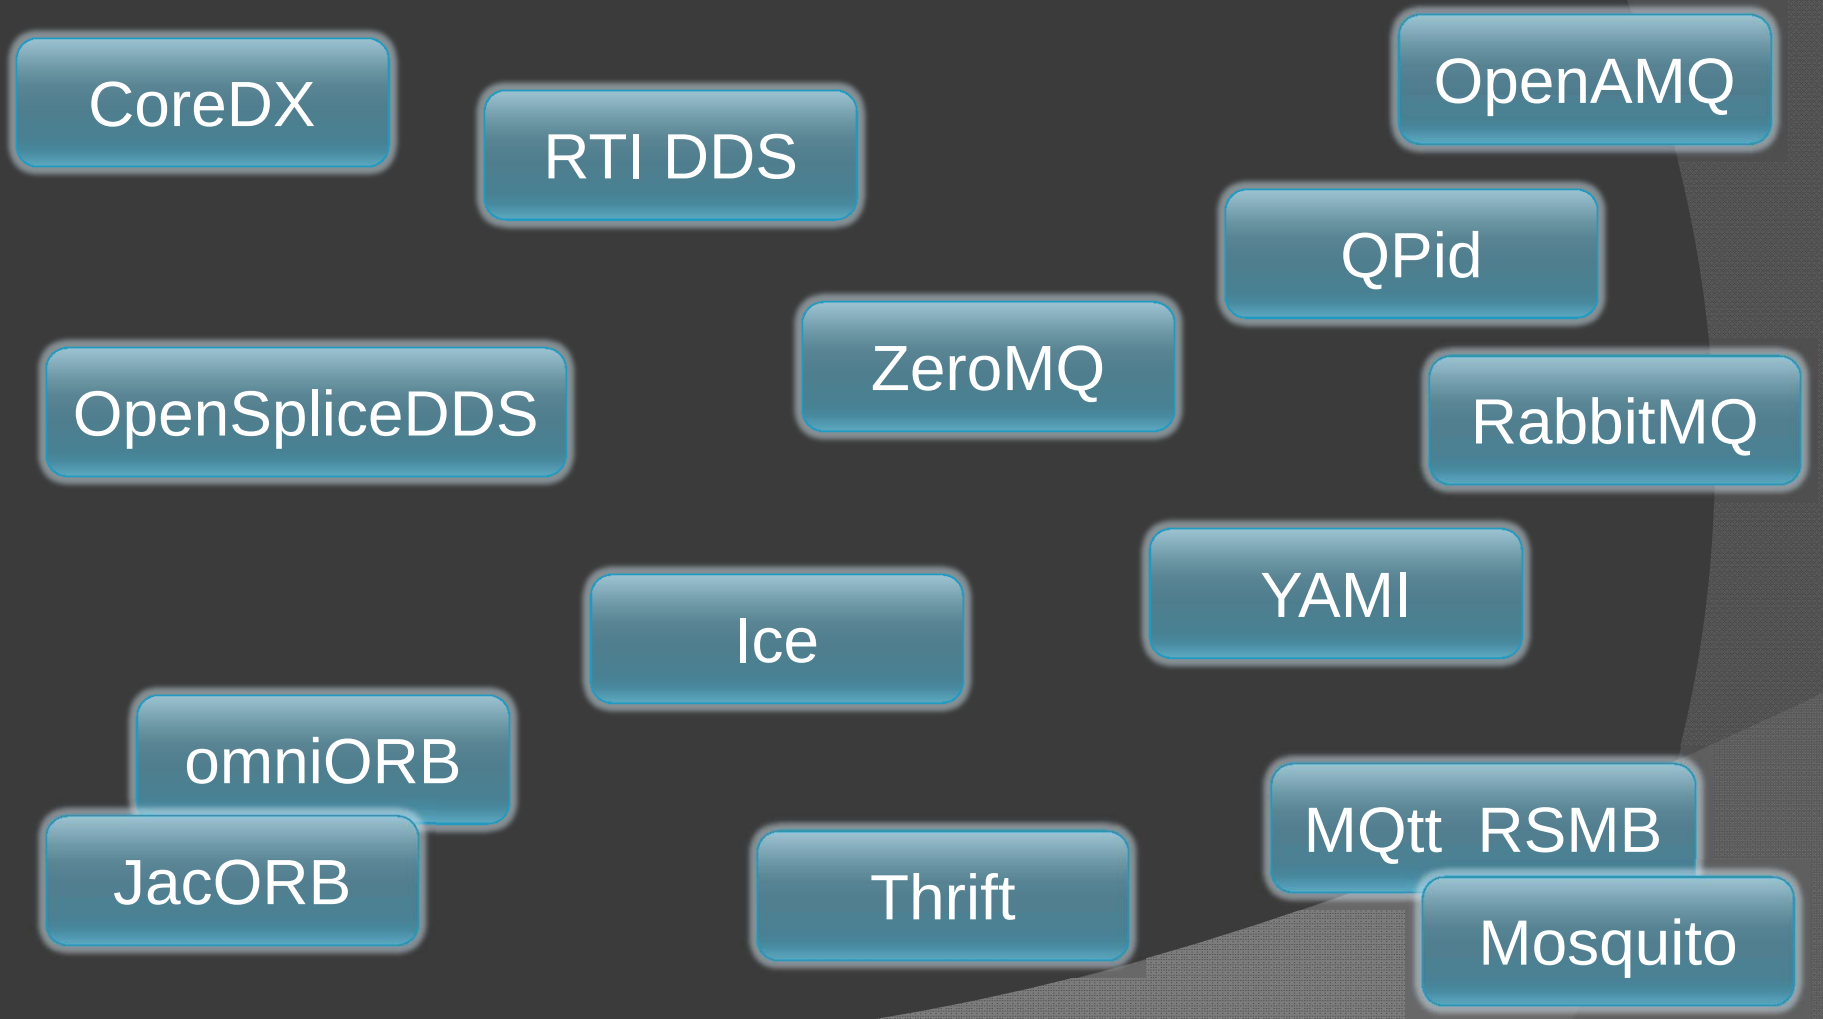

CORBA (omniORB, JacORB)

Object-oriented communication platform standardized by OMG. It is over 20 years old, and the standard is well established but...

- Complex, hard to learn and use ☹
- Differences between implementations ☹
- Big memory footprint ☹
- Shrinking community ☹
- Seldom updates and bug fixes ☹

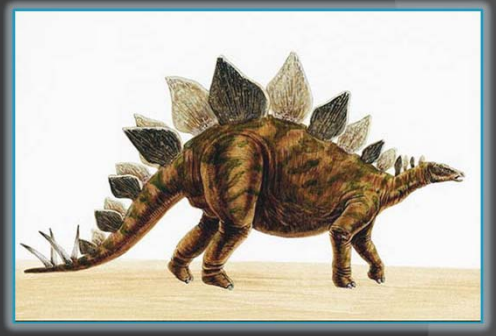

Dying out

Complex

Heavy

Shrinking community

Ice

Object-oriented middleware by ZeroC, people who left CORBA to fix its problems. Conceptually similar to CORBA but...

- Better design and implementation ☺
- Modern easier to use API and IDL mapping ☺
- Concise C++ and Java implementation ☺
- Support for versioning ☺
- Active community ☺

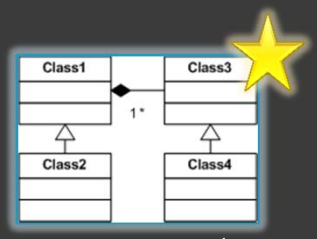

Well designed

Active community

Thrift

Middleware by Facebook (moved to Apache)

- Lightweight, free of unnecessary dependencies ☺
- Modern, clean API ☺
- Active community ☺
- Ongoing development, still incomplete ☹
- No documentation, empty tutorial ☹

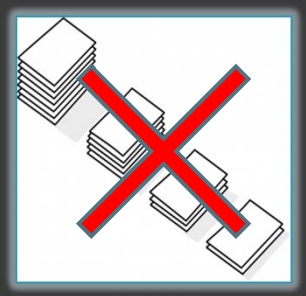

Poor documentation

Work still in progress

YAMI4

Message-oriented middleware developed by one of our former colleagues. Used by some CERN Controls services.

- Lightweight, free of unnecessary dependencies ☺
- Well designed, modern API ☺
- Asynchronous req/rep and pub/sub patterns ☺
- Dynamic type specification (paid in performance) ☺☹
- Small community ☹

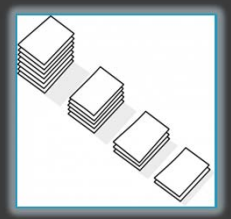

Lightweight

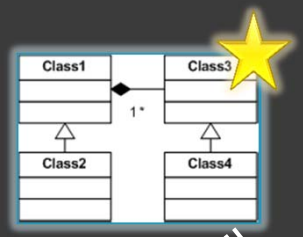

Well designed

Small community

AMQP – OpenAMQ, Qpid, RabbitMQ

Middleware where messages are distributed by a broker. A few independent implementations of the broker and clients.

- Big community ☺
- Designed to support pub/sub ☺
- Protocol standard issues ☹
- Broker – single point of failure ☹
- Broker – additional hop, slower communication ☹
- Broker – additional, non compliant monitoring tools ☹

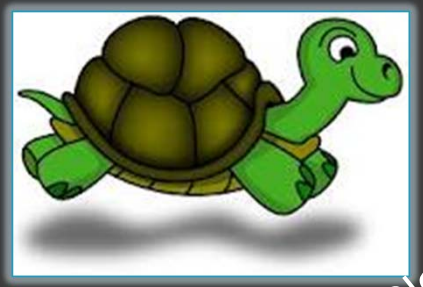

Slower for req/rep

Broker issues

Active community

DDS – RTI, CoreDX, OpenSplice

Data-oriented middleware with p2p communication. Standard defined by OMG, with a few compatible implementations.

- Big industrial and military community ☺
- Possible compatibility with CORBA ☺
- Designed to support pub/sub ☺
- Req/rep possible but needs 2 channels ☹
- Steep learning curve ☹
- Complex API ☹

ZeroMQ

Message-oriented middleware by iMatix (they dropped OpenAMQ.)
API resembles BSD sockets but it is so much more.

- Simple yet powerful API ☺
- Support for in-proc, inter-proc, TCP, PGM ☺
- Lightweight, free of unnecessary dependencies ☺
- Many communication patterns ☺
- May be used as a concurrency framework ☺
- Active community ☺

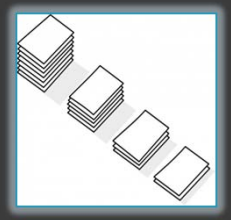

Lightweight

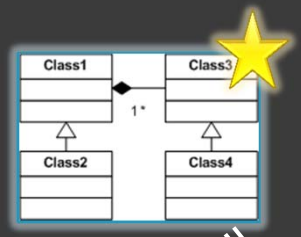

Well designed

Active community

Performance requirements reviewed

Instead of estimating

- We asked our users what they will need
- Gathered statistics from the current system

- Most demanding users identified
- Definition of a few performance tests

Performance tests, reliable req/rep

4000 msg/sec
Payload = 4B

5 msg/sec
Payload = 10MB

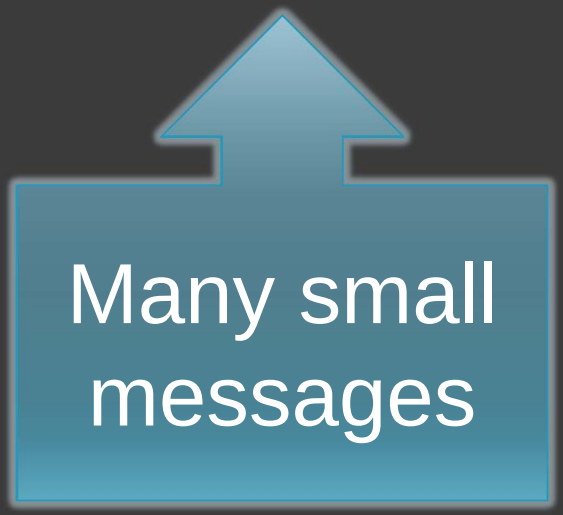

Performance tests, reliable req/rep

4000 msg/sec
Payload = 4B

5 msg/sec
Payload = 10MB

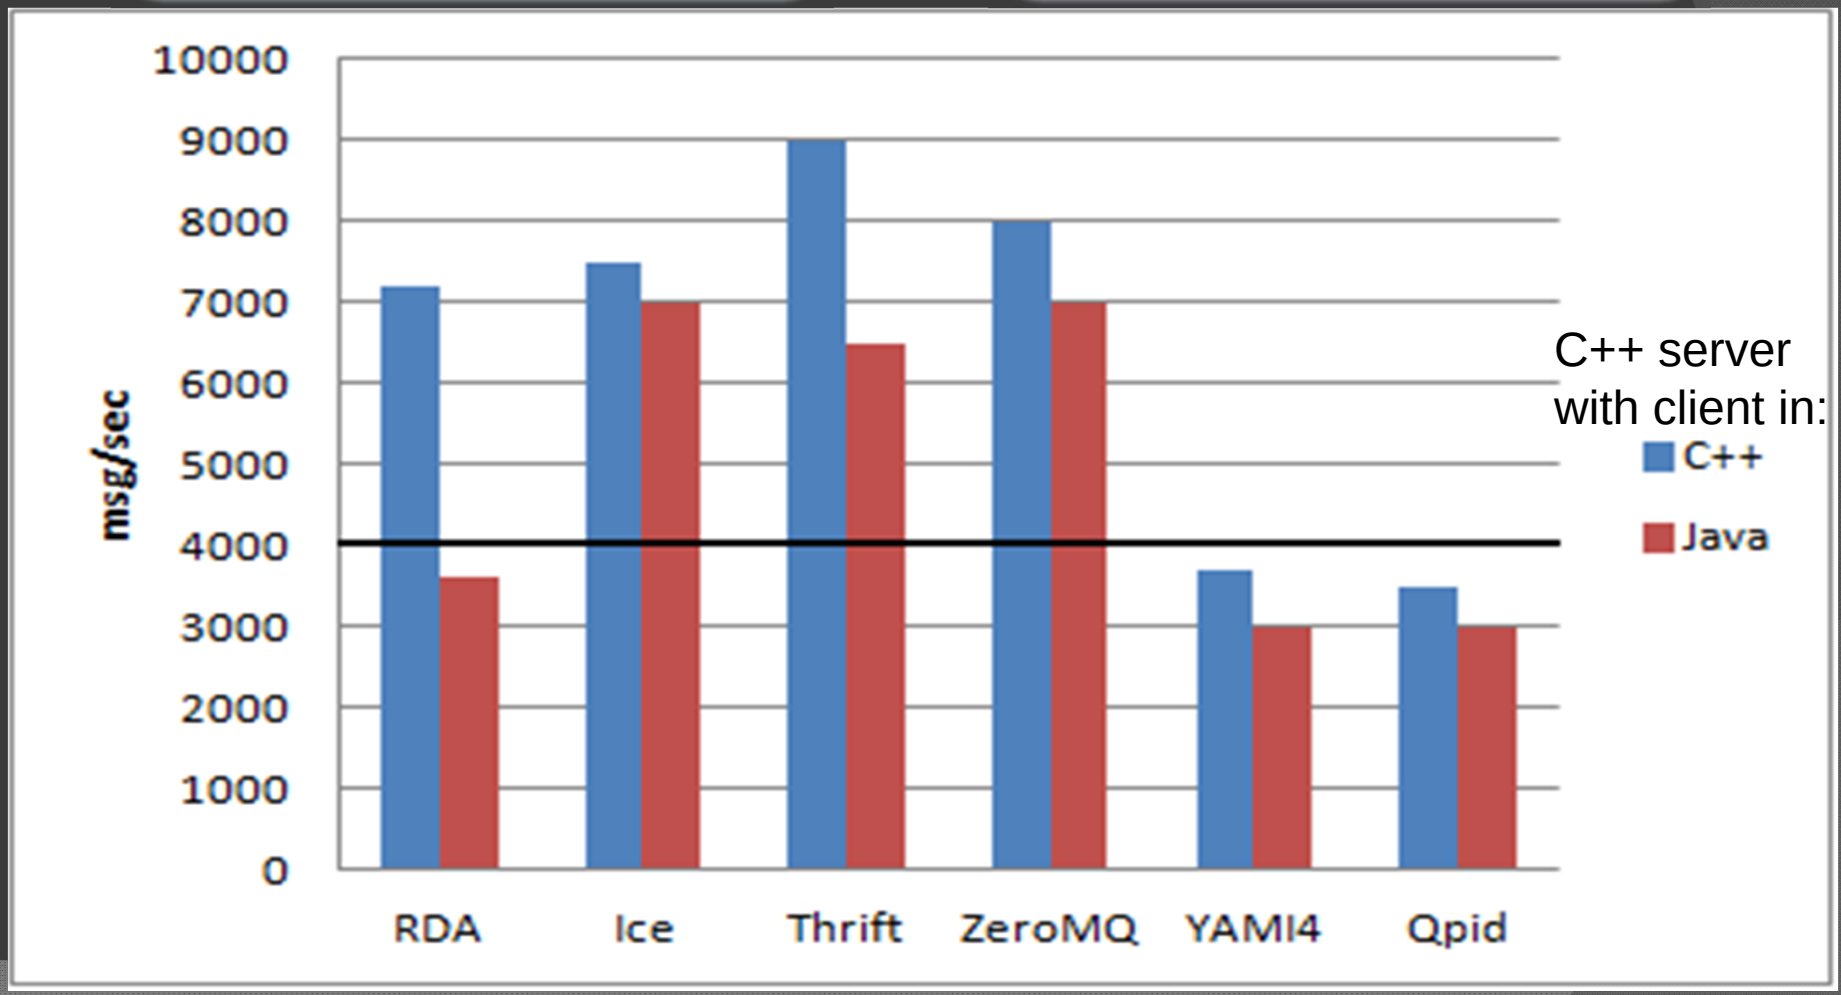

Performance tests, reliable req/rep

4000 msg/sec
Payload = 4B

5 msg/sec
Payload = 10MB

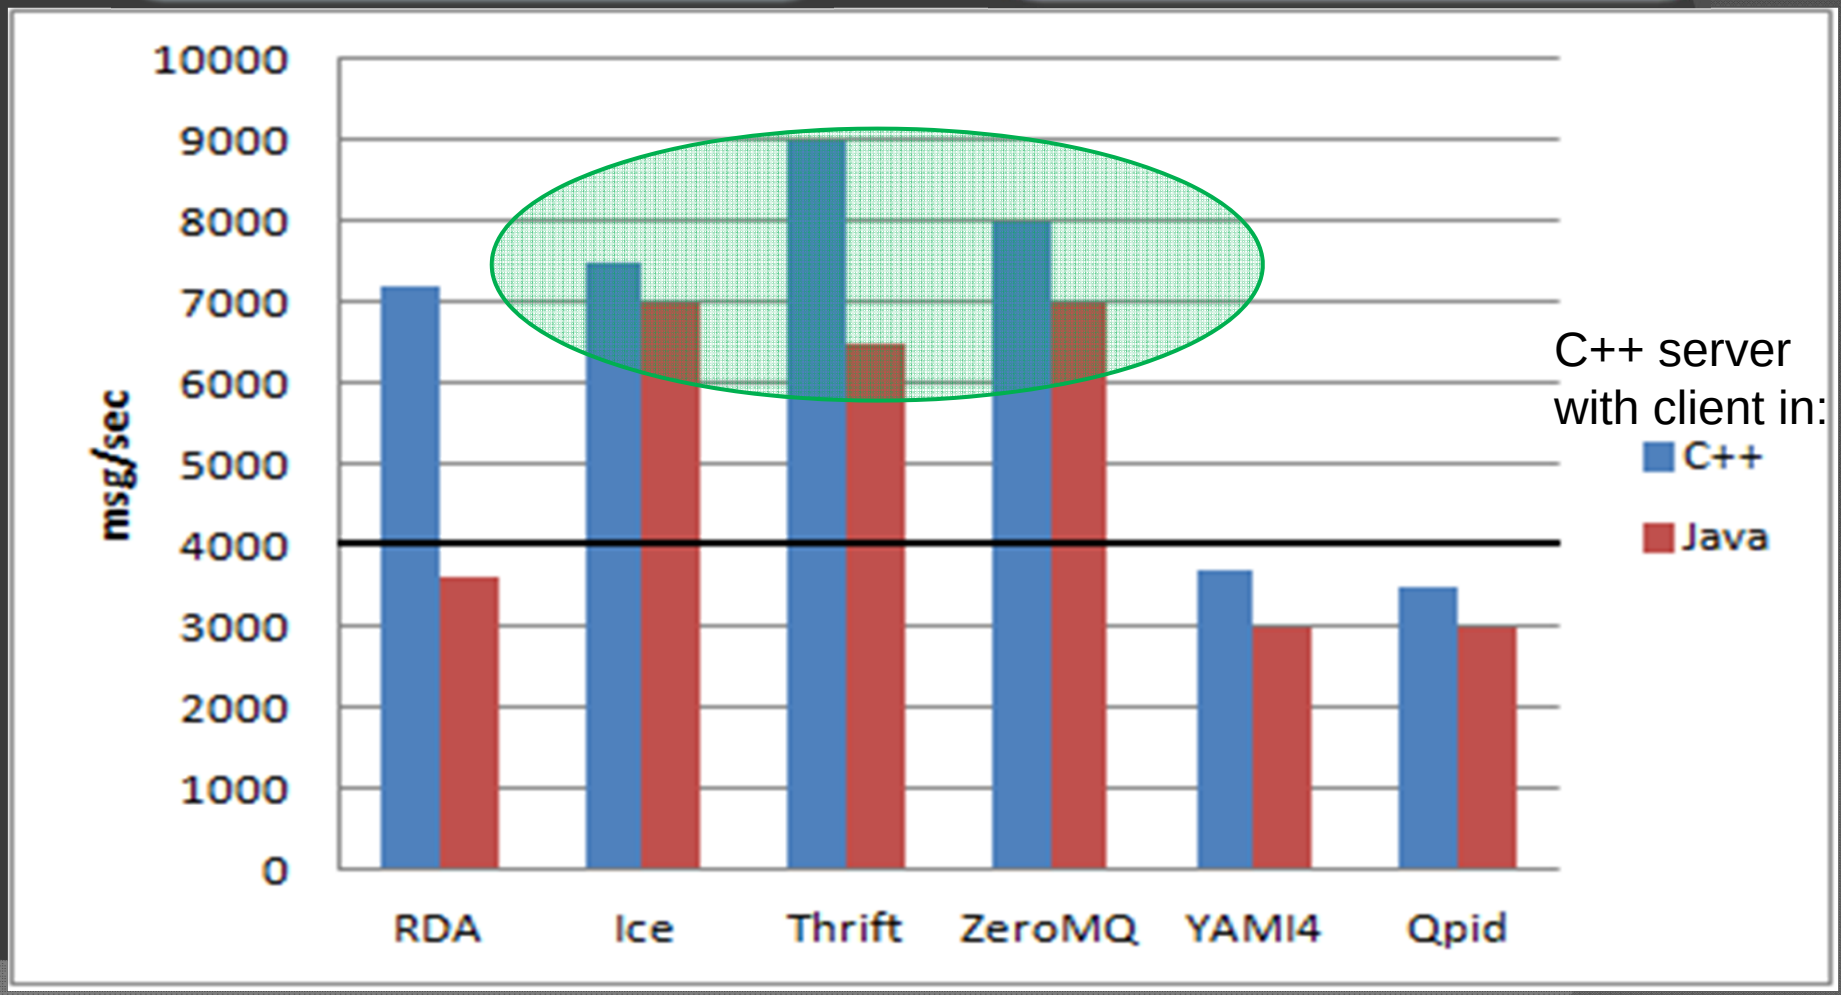

Performance tests, reliable req/rep

4000 msg/sec
Payload = 4B

5 msg/sec
Payload = 10MB

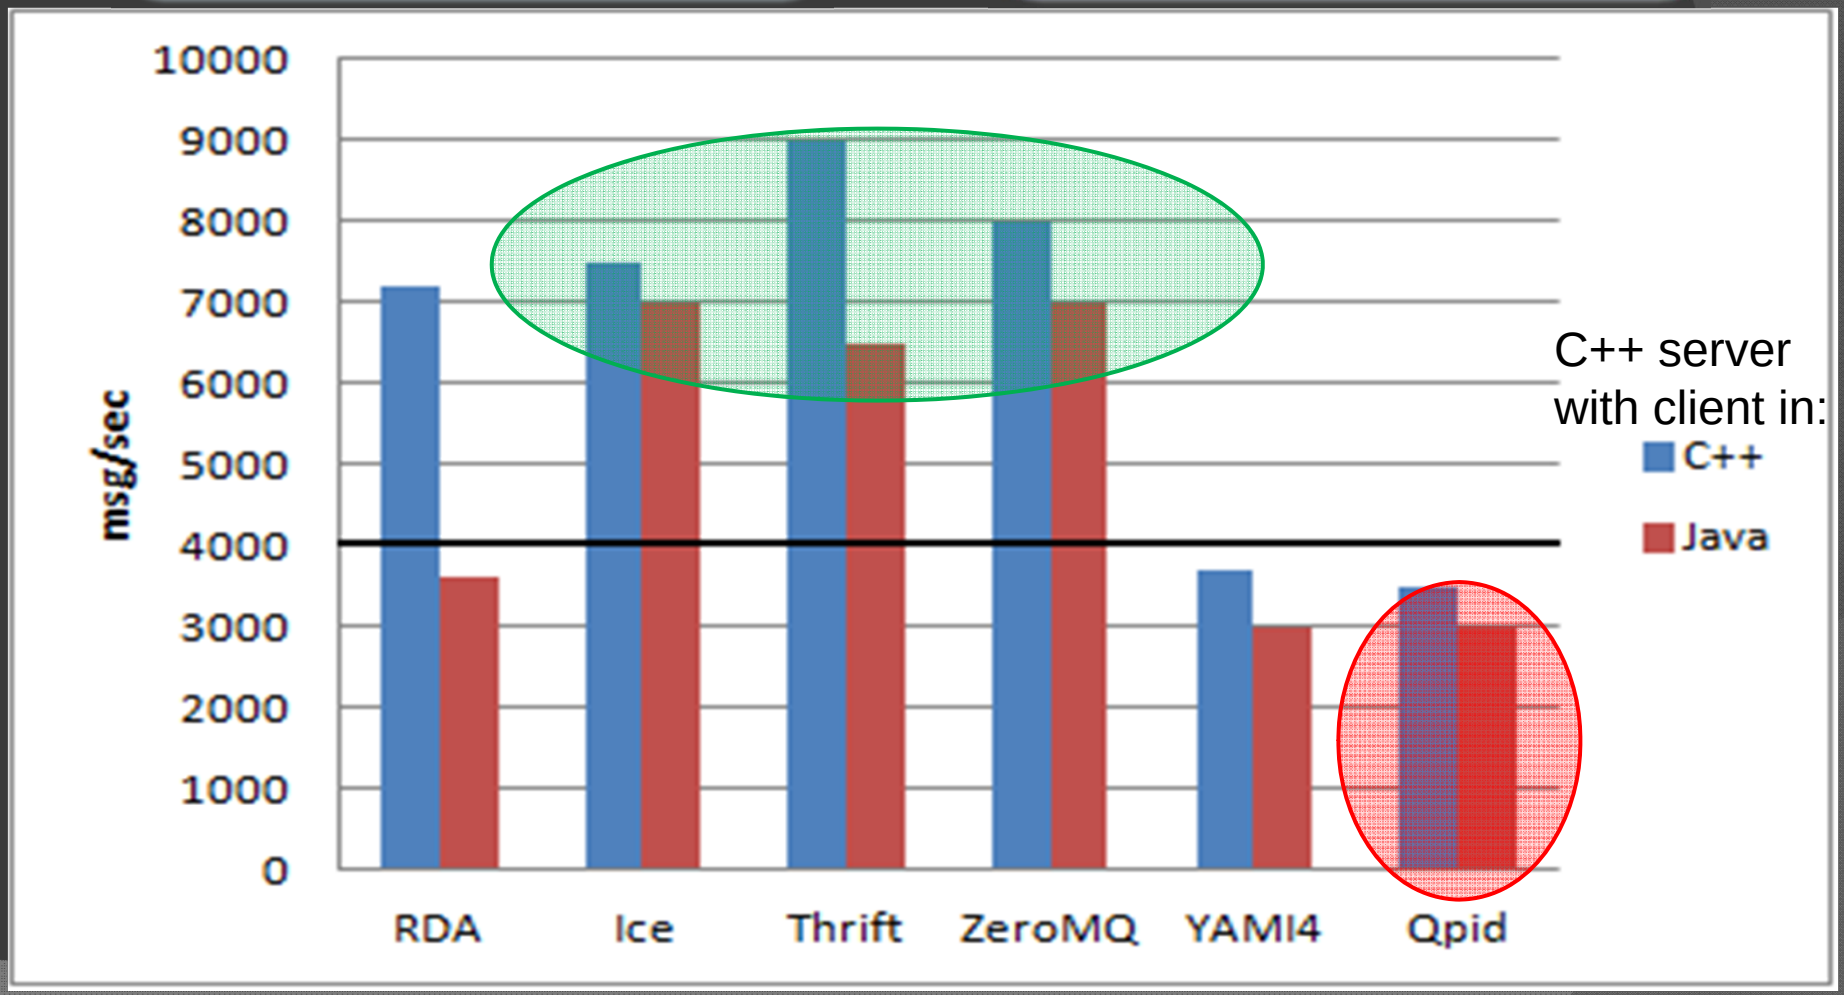

Performance tests, reliable req/rep

4000 msg/sec
Payload = 4B

5 msg/sec
Payload = 10MB

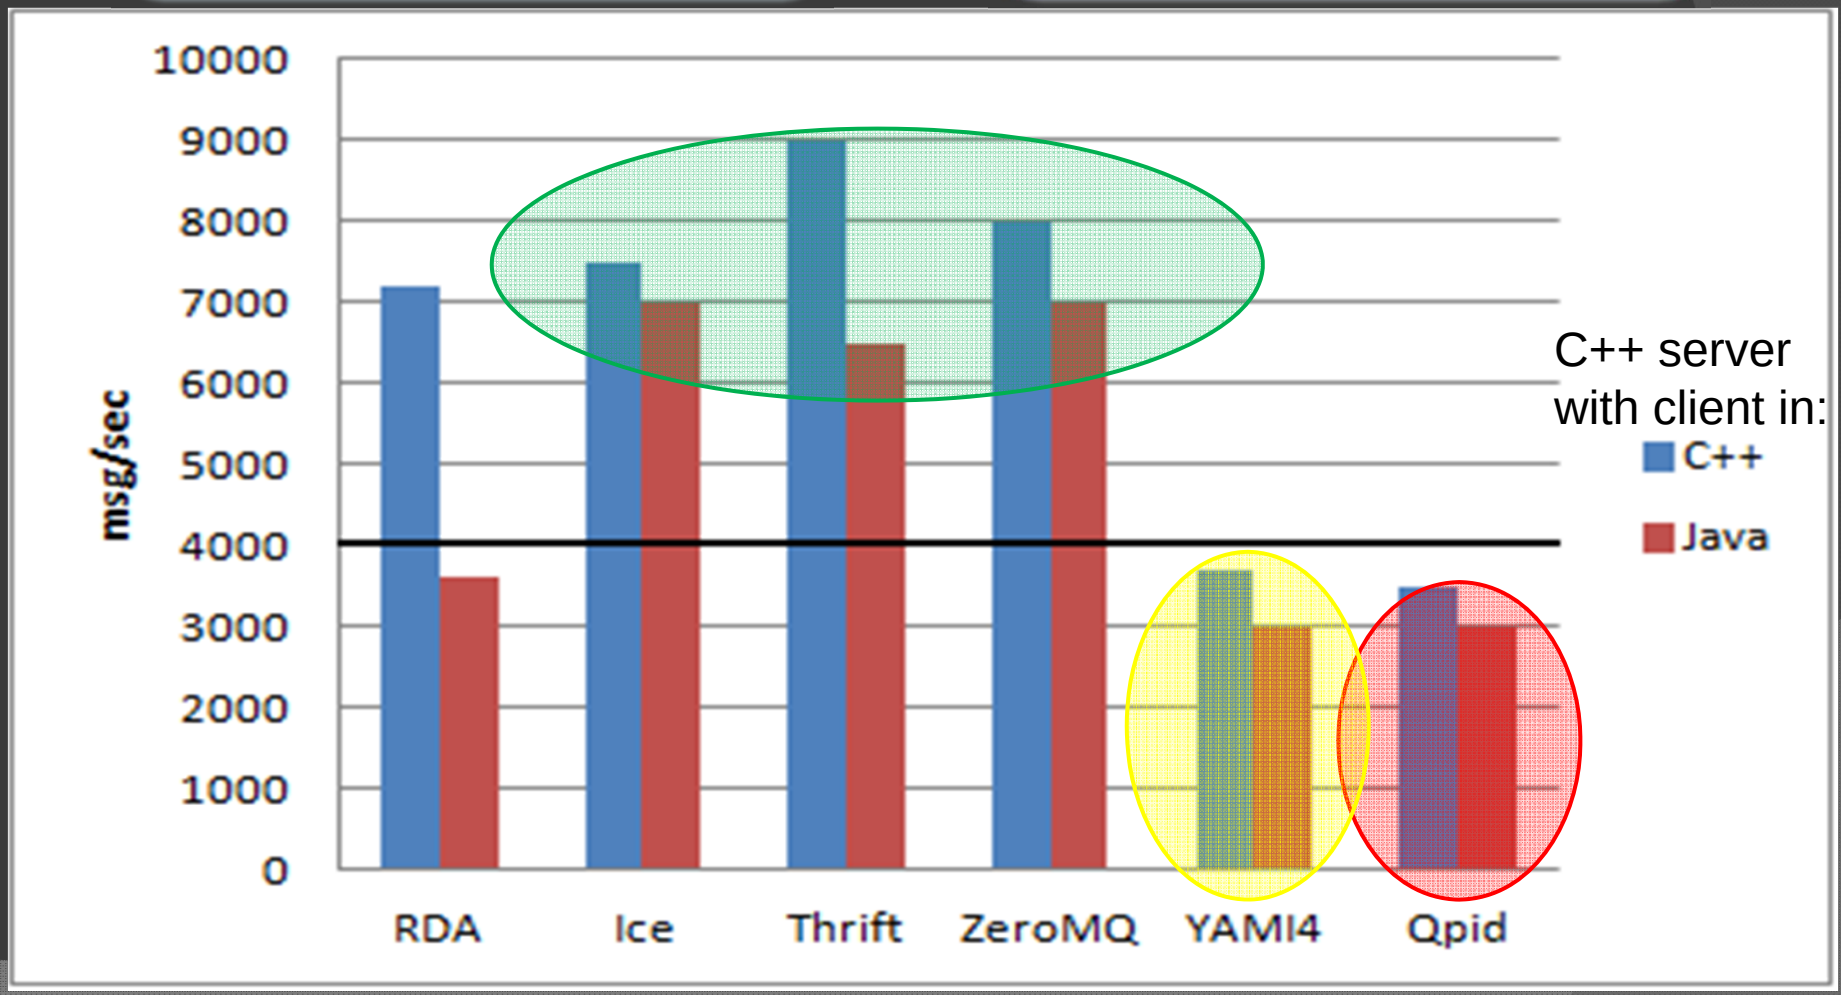

Performance tests, reliable req/rep

4000 msg/sec
Payload = 4B

5 msg/sec
Payload = 10MB

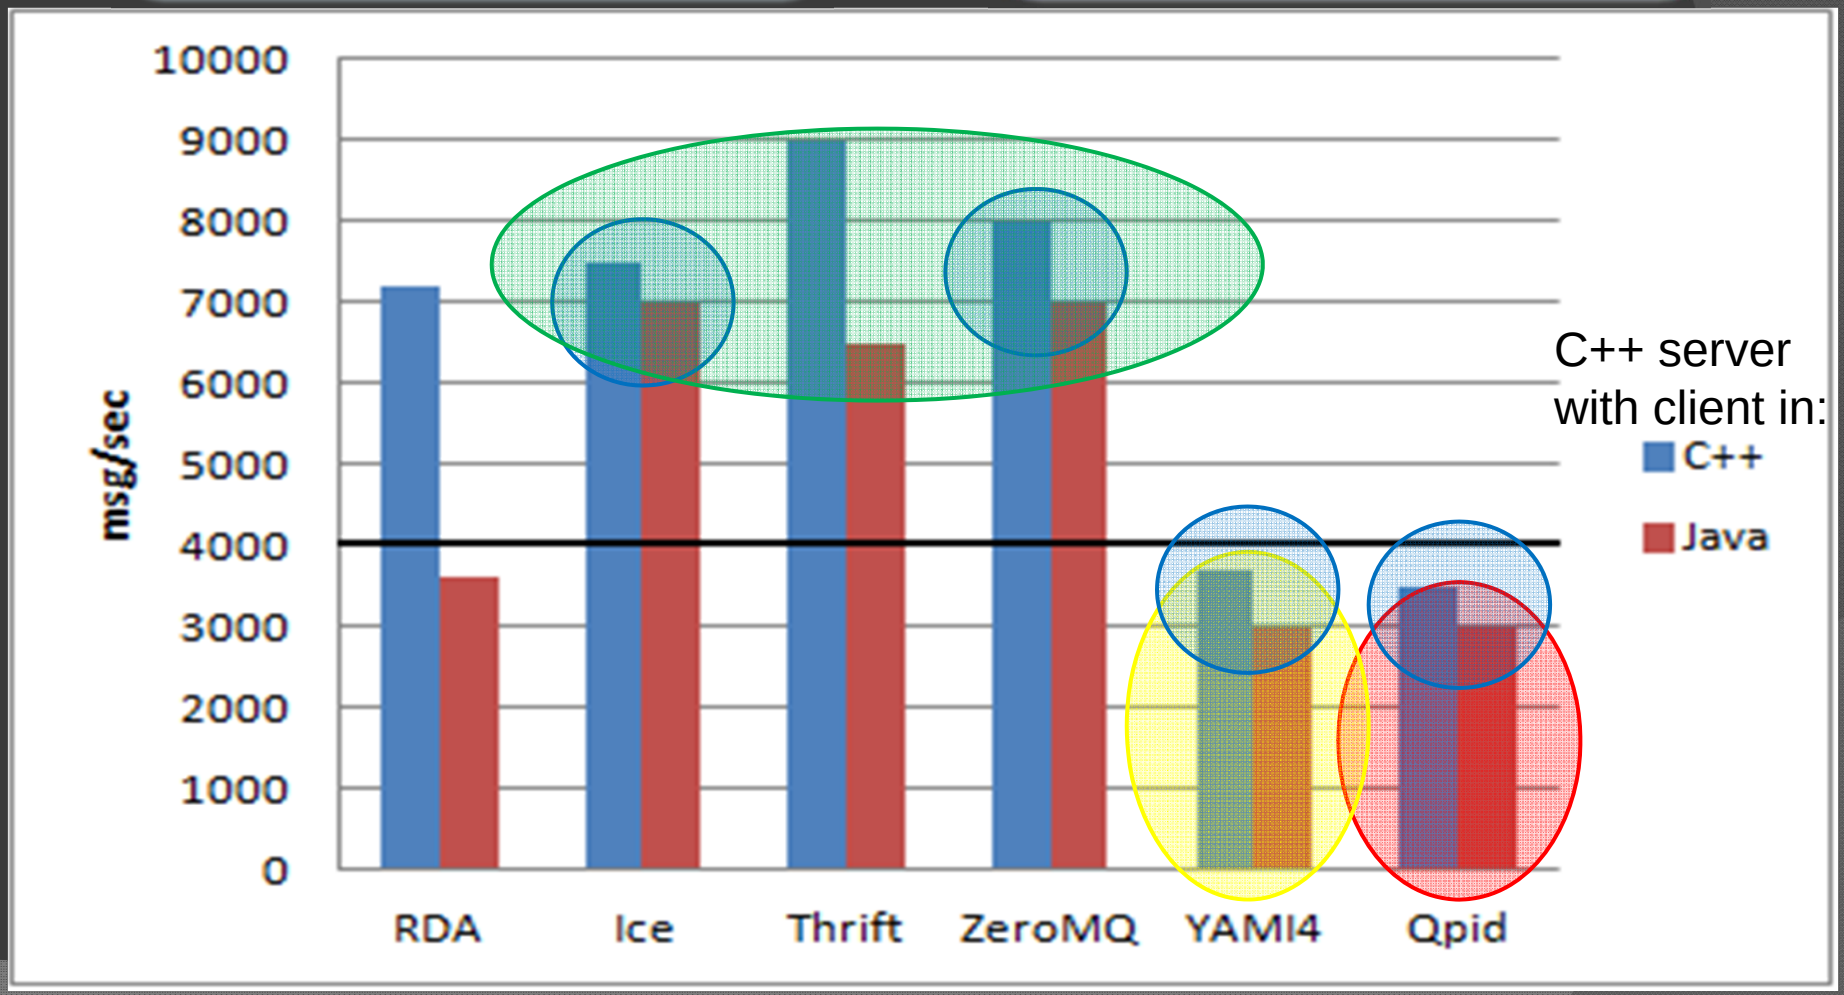

Performance/Scalability tests, reliable pub/sub

400 msg x 8 B
10 clients, <50ms

30 msg x 8 B
10 clients, <20ms

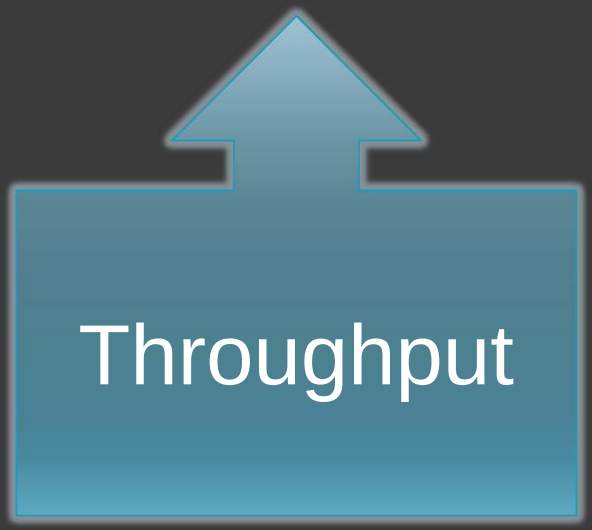

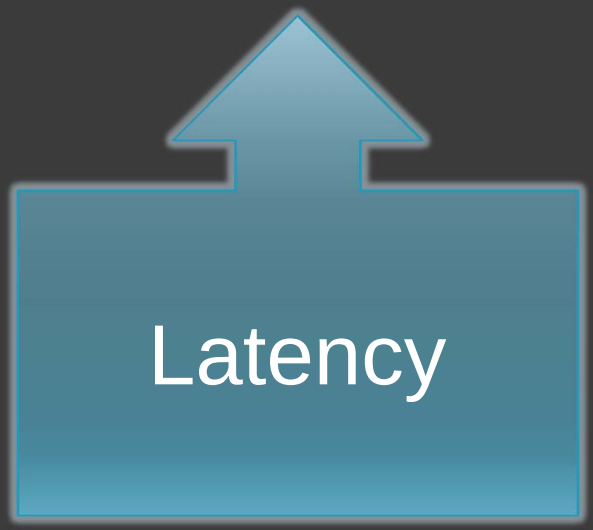

Performance/Scalability tests, reliable pub/sub

400 msg x 8 B
10 clients, <50ms

30 msg x 8 B
10 clients, <20ms

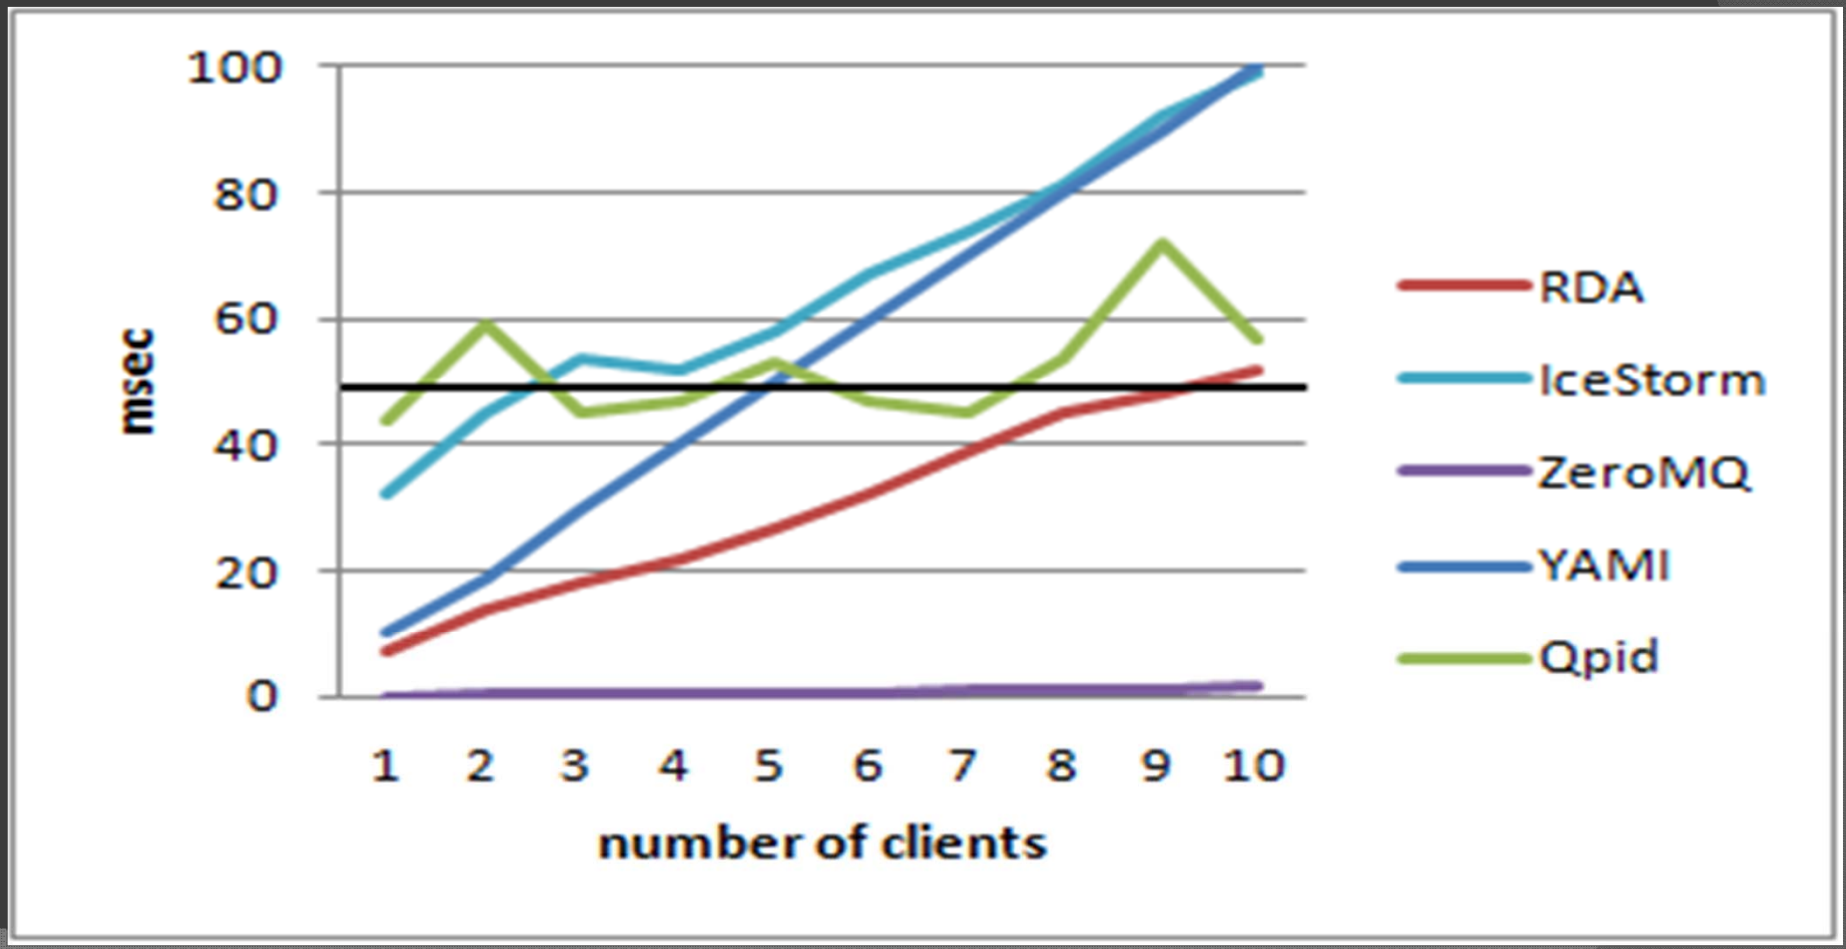

Performance/Scalability tests, reliable pub/sub

400 msg x 8 B
10 clients, <50ms

30 msg x 8 B
10 clients, <20ms

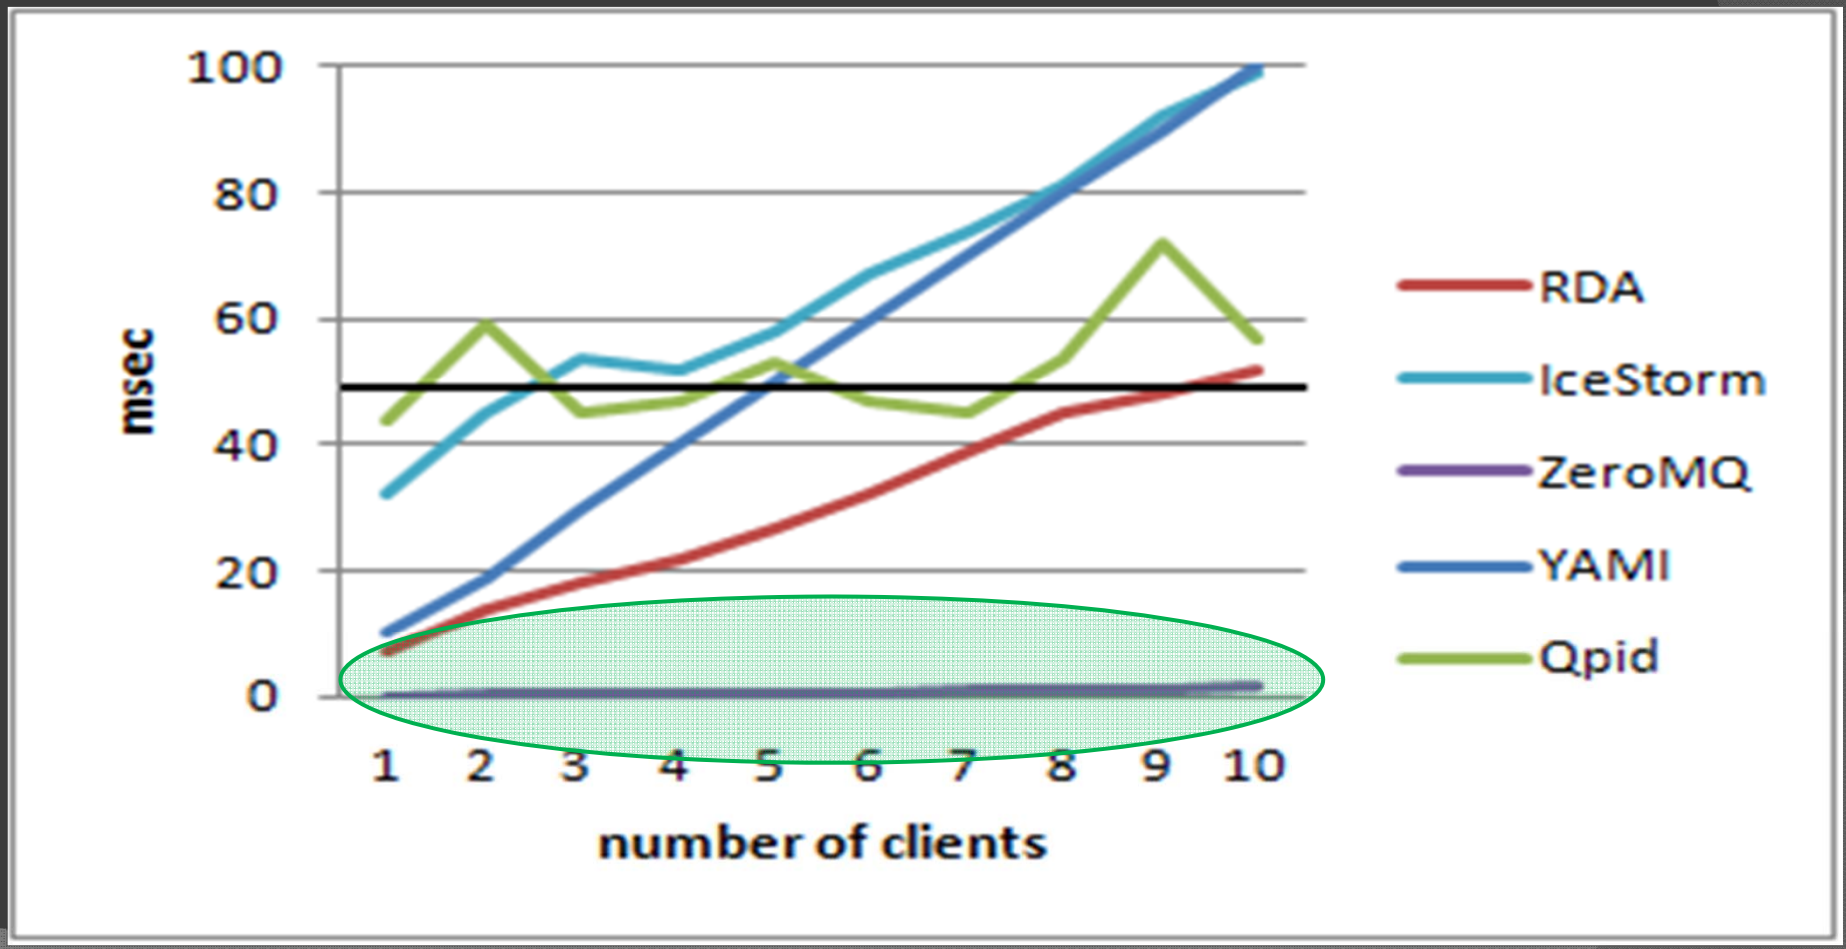

Performance/Scalability tests, reliable pub/sub

400 msg x 8 B
10 clients, <50ms

30 msg x 8 B
10 clients, <20ms

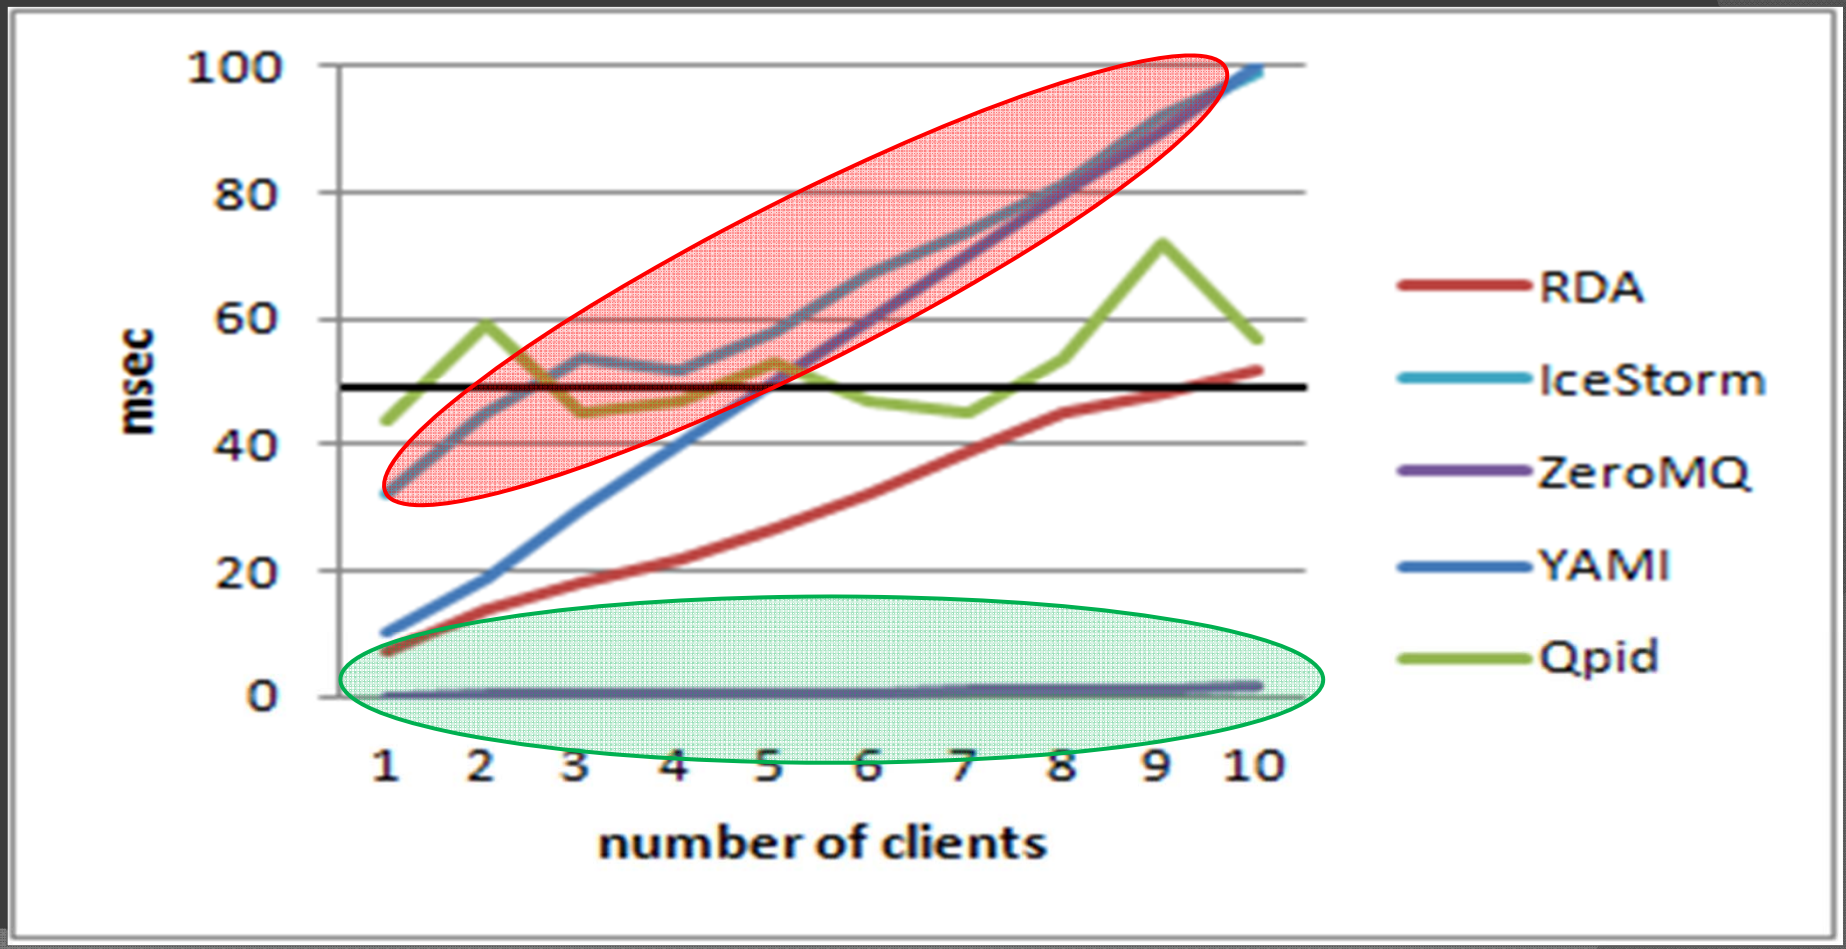

Products comparison (according to the criteria)

	Sync, async & msg patterns	QoS	Dependencies & memory f-p	Performance	Look & feel, API, docs	Community & maturity	Score
ZeroMQ	✓	✓	✓	✓	✓	✓	6
Ice	✓	✓	✗	✓	✓	✓	5
YAMI4	✓	✓	✓	✗	✓	✗	4
RTI	✗	✓	✗	✓	✗	✓	3
Qpid	✗	✓	✗	✗	✓	✓	3
CORBA	✗	✓	✗	✓	✗	✗	2
Thrift	✗	✗	✓	✓	✗	✗	2

Products comparison (according to the criteria)

	Sync, async & msg patterns	QoS	Dependencies & memory f-p	Performance	Look & feel, API, docs	Community & maturity	Score
ZeroMQ	✓	✓	✓	✓	✓	✓	6
Ice	✓	✓	✗	✓	✓	✓	5
YAMI4	✓	✓	✓	✗	✓	✗	4
RTI	✗	✓	✗	✓	✗	✓	3
Qpid	✗	✓	✗	✗	✓	✓	3
CORBA	✗	✓	✗	✓	✗	✗	2
Thrift	✗	✗	✓	✓	✗	✗	2

Products comparison (according to the criteria)

	Sync, async & msg patterns	QoS	Dependencies & memory f-p	Performance	Look & feel, API, docs	Community & maturity	Score
ZeroMQ	✓	✓	✓	✓	✓	✓	6
Ice	✓	✓	✗	✓	✓	✓	5
YAMI4	✓	✓	✓	✗	✓	✗	4
RTI	✗	✓	✗	✓	✗	✓	3
Qpid	✗	✓	✗	✗	✓	✓	3
CORBA	✗	✓	✗	✓	✗	✗	2
Thrift	✗	✗	✓	✓	✗	✗	2

Conclusions

- Several good middleware solutions available.
- The choice is dictated by the most critical requirements for any given application. Not easy → QoS policies and performance matter, but also ease of use, community, …
- Concerning CERN Controls Middleware…
  - Prototype with the most promising candidates:

  - Deploy the new middleware before the long accelerator shutdown at the end of 2012