

# Reactor Netty Reference Guide

Stephane Maldini, Violeta Georgieva

Version 1.0.0-M2

# Table of Contents

1. About the Documentation .....	1
1.1. Latest Version and Copyright Notice .....	1
1.2. Contributing to the Documentation .....	1
1.3. Getting Help .....	1
2. Getting Started .....	2
2.1. Introducing Reactor Netty .....	2
2.2. Prerequisites .....	2
2.3. Understanding the BOM and versioning scheme .....	2
2.4. Getting Reactor Netty .....	3
3. TCP Server .....	7
3.1. Starting and Stopping .....	7
3.2. Writing Data .....	8
3.3. Consuming Data .....	9
3.4. Lifecycle Callbacks .....	10
3.5. TCP-level Configurations .....	11
3.6. SSL/TLS .....	15
3.7. Metrics .....	16
4. TCP Client .....	19
4.1. Connect and Disconnect .....	19
4.2. Writing Data .....	20
4.3. Consuming Data .....	21
4.4. Lifecycle Callbacks .....	22
4.5. TCP-level Configurations .....	23
4.6. Connection Pool .....	27
4.7. SSL and TLS .....	30
4.8. Proxy Support .....	31
4.9. Metrics .....	32
5. HTTP Server .....	36
5.1. Starting and Stopping .....	36
5.2. Routing HTTP .....	37
5.3. Writing Data .....	40
5.4. Consuming Data .....	43
5.5. TCP-level Configuration .....	47
5.6. SSL and TLS .....	49
5.7. HTTP Access Log .....	50
5.8. HTTP/2 .....	51
5.9. Metrics .....	53
6. HTTP Client .....	57

6.1. Connect . . . . .	57
6.2. Writing Data . . . . .	59
6.3. Consuming Data . . . . .	63
6.4. TCP-level Configuration . . . . .	66
6.5. SSL andTLS . . . . .	68
6.6. Retry Strategies . . . . .	69
6.7. Metrics . . . . .	69
7. UDP Server . . . . .	74
7.1. Starting and Stopping . . . . .	74
7.2. Writing Data . . . . .	75
7.3. Consuming Data . . . . .	76
7.4. Lifecycle Callbacks . . . . .	77
7.5. Connection Configuration . . . . .	78
7.6. Metrics . . . . .	82
8. UDP Client . . . . .	85
8.1. Connecting and Disconnecting . . . . .	85
8.2. Writing Data . . . . .	86
8.3. Consuming Data . . . . .	87
8.4. Lifecycle Callbacks . . . . .	88
8.5. Connection Configuration . . . . .	89
8.6. Metrics . . . . .	93

# Chapter 1. About the Documentation

This section provides a brief overview of [Reactor Netty](#) reference documentation. You do not need to read this guide in a linear fashion. Each piece stands on its own, though they often refer to other pieces.

## 1.1. Latest Version and Copyright Notice

The [Reactor Netty](#) reference guide is available as [HTML](#) documents. The latest copy is available at <https://projectreactor.io/docs/netty/release/reference/index.html>

Copies of this document may be made for your own use and for distribution to others, provided that you do not charge any fee for such copies and further provided that each copy contains this [Copyright Notice](#), whether distributed in print or electronically.

## 1.2. Contributing to the Documentation

The reference guide is written in [Asciidoc](#), and you can find its sources at <https://github.com/reactor/reactor-netty/tree/master/docs/asciidoc>.

If you have an improvement, we will be happy to get a pull request from you!

We recommend that you check out a local copy of the repository so that you can generate the documentation by using the [asciidoctor](#) Gradle task and checking the rendering. Some of the sections rely on included files, so [GitHub](#) rendering is not always complete.

## 1.3. Getting Help

There are several ways to reach out for help with [Reactor Netty](#). You can:

- Get in touch with the community on [Gitter](#).
- Ask a question on stackoverflow.com at [reactor-netty](#).
- Report bugs in [Github](#) issues. The repository is the following: [reactor-netty](https://github.com/reactor/reactor-netty).



All of [Reactor Netty](#) is open source, [including this documentation](#).

# Chapter 2. Getting Started

This section contains information that should help you get going with [Reactor Netty](#). It includes the following information:

- [Introducing Reactor Netty](#)
- [Prerequisites](#)
- [Understanding the BOM and versioning scheme](#)
- [Getting Reactor Netty](#)

## 2.1. Introducing Reactor Netty

Suited for Microservices Architecture, [Reactor Netty](#) offers backpressure-ready network engines for [HTTP](#) (including Websockets), [TCP](#), and [UDP](#).

## 2.2. Prerequisites

[Reactor Netty](#) runs on [Java 8](#) and above.

It has transitive dependencies on:

- Reactive Streams v1.0.3
- Reactor Core v3.x
- Netty v4.1.x

## 2.3. Understanding the BOM and versioning scheme

[Reactor Netty](#) is part of the [Project Reactor BOM](#) (since the [Aluminium](#) release train). This curated list groups artifacts that are meant to work well together, providing the relevant versions despite potentially divergent versioning schemes in these artifacts.



The versioning scheme has changed between 0.9.x and 1.0.x (Dysprosium and Europium).

Artifacts follow a versioning scheme of [MAJOR.MINOR.PATCH-QUALIFIER](#) while the BOM is versioned using a CalVer inspired scheme of [YYYY.MINOR.PATCH-QUALIFIER](#), where:

- **MAJOR** is the current generation of Reactor, where each new generation can bring fundamental changes to the structure of the project (which might imply a more significant migration effort)
- **YYYY** is the year of the first GA release in a given release cycle (like 1.0.0 for 1.0.x)
- **.MINOR** is a 0-based number incrementing with each new release cycle
  - in the case of projects, it generally reflects wider changes and can indicate a moderate migration effort
  - in the case of the BOM it allows discerning between release cycles in case two get first

released the same year

- **.PATCH** is a 0-based number incrementing with each service release
- **-QUALIFIER** is a textual qualifier, which is omitted in the case of GA releases (see below)

The first release cycle to follow that convention is thus **2020.0.x**, codename **Europium**. The scheme uses the following qualifiers (note the use of dash separator), in order:

- **-M1..-M9**: milestones (we don't expect more than 9 per service release)
- **-RC1..-RC9**: release candidates (we don't expect more than 9 per service release)
- **-SNAPSHOT**: snapshots
- *no qualifier* for GA releases



Snapshots appear higher in the order above because, conceptually, they're always "the freshest pre-release" of any given PATCH. Even though the first deployed artifact of a PATCH cycle will always be a **-SNAPSHOT**, a similarly named but more up-to-date snapshot would also get released after eg. a milestone or between release candidates.

Each release cycle is also given a codename, in continuity with the previous codename-based scheme, which can be used to reference it more informally (like in discussions, blog posts, etc...). The codenames represent what would traditionally be the MAJOR.MINOR number. They (mostly) come from the [Periodic Table of Elements](#), in increasing alphabetical order.



Up until Dysprosium, the BOM was versioned using a release train scheme with a codename followed by a qualifier, and the qualifiers were slightly different. For example: Aluminium-RELEASE (first GA release, would now be something like YYYY.0.0), Bismuth-M1, Californium-SR1 (service release would now be something like YYYY.0.1), Dysprosium-RC1, Dysprosium-BUILD-SNAPSHOT (after each patch, we'd go back to the same snapshot version. would now be something like YYYY.0.X-SNAPSHOT so we get 1 snapshot per PATCH)

## 2.4. Getting Reactor Netty

As [mentioned earlier](#), the easiest way to use **Reactor Netty** in your core is to use the **BOM** and add the relevant dependencies to your project. Note that, when adding such a dependency, you must omit the version so that the version gets picked up from the **BOM**.

However, if you want to force the use of a specific artifact's version, you can specify it when adding your dependency as you usually would. You can also forego the **BOM** entirely and specify dependencies by their artifact versions.

### 2.4.1. Maven Installation

The **BOM** concept is natively supported by **Maven**. First, you need to import the **BOM** by adding the following snippet to your **pom.xml**. If the top section (**dependencyManagement**) already exists in your pom, add only the contents.

```
<dependencyManagement> ①
  <dependencies>
    <dependency>
      <groupId>io.projectreactor</groupId>
      <artifactId>reactor-bom</artifactId>
      <version>Dysprosium-SR10</version> ②
      <type>pom</type>
      <scope>import</scope>
    </dependency>
  </dependencies>
</dependencyManagement>
```

① Notice the `dependencyManagement` tag. This is in addition to the regular `dependencies` section.

② As of this writing, `Dysprosium-SR10` is the latest version of the `BOM`. Check for updates at <https://github.com/reactor/reactor/releases>.

Next, add your dependencies to the relevant reactor projects, as usual (except without a `<version>`). The following listing shows how to do so:

```
<dependencies>
  <dependency>
    <groupId>io.projectreactor.netty</groupId>
    <artifactId>reactor-netty</artifactId> ①
    ②
  </dependency>
</dependencies>
```

① Dependency on `Reactor Netty`

② No version tag here

## 2.4.2. Gradle Installation

The `BOM` concept is supported in Gradle since version 5. The following listing shows how to import the `BOM` and add a dependency to `Reactor Netty`:

```

dependencies {
    // import a BOM
    implementation platform('io.projectreactor:reactor-bom:Dysprosium-SR10') ①

    // define dependencies without versions
    implementation 'io.projectreactor.netty:reactor-netty' ②
}

```

① As of this writing, **Dysprosium-SR10** is the latest version of the **BOM**. Check for updates at <https://github.com/reactor/reactor/releases>.

② There is no third `:` separated section for the version. It is taken from the **BOM**.

### 2.4.3. Milestones and Snapshots

Milestones and developer previews are distributed through the **Spring Milestones** repository rather than **Maven Central**. To add it to your build configuration file, use the following snippet:

#### *Milestones in Maven*

```

<repositories>
    <repository>
        <id>spring-milestones</id>
        <name>Spring Milestones Repository</name>
        <url>https://repo.spring.io/milestone</url>
    </repository>
</repositories>

```

For Gradle, use the following snippet:

#### *Milestones in Gradle*

```

repositories {
    maven { url 'https://repo.spring.io/milestone' }
    mavenCentral()
}

```

Similarly, snapshots are also available in a separate dedicated repository (for both Maven and Gradle):

*-SNAPSHOTs in Maven*

```
<repositories>
  <repository>
    <id>spring-snapshots</id>
    <name>Spring Snapshot Repository</name>
    <url>https://repo.spring.io/snapshot</url>
  </repository>
</repositories>
```

*-SNAPSHOTs in Gradle*

```
repositories {
  maven { url 'https://repo.spring.io/snapshot' }
  mavenCentral()
}
```

# Chapter 3. TCP Server

Reactor Netty provides an easy to use and configure `TcpServer`. It hides most of the Netty functionality that is needed to create a TCP server and adds Reactive Streams backpressure.

## 3.1. Starting and Stopping

To start a TCP server, you must create and configure a `TcpServer` instance. By default, the `host` is configured for any local address, and the system picks up an ephemeral port when the `bind` operation is invoked. The following example shows how to create and configure a `TcpServer` instance:

```
././reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/tcp/server/create/Application.java
```

```
import reactor.netty.DisposableServer;
import reactor.netty.tcp.TcpServer;

public class Application {

    public static void main(String[] args) {
        DisposableServer server =
            TcpServer.create() ①
                .bindNow(); ②

        server.onDispose()
            .block();
    }
}
```

① Creates a `TcpServer` instance that is ready for configuring.

② Starts the server in a blocking fashion and waits for it to finish initializing.

The returned `DisposableServer` offers a simple server API, including `disposeNow()`, which shuts the server down in a blocking fashion.

### 3.1.1. Host and Port

To serve on a specific `host` and `port`, you can apply the following configuration to the TCP server:

`././reactor-netty-examples/src/main/java/reactor/netty/examples/documentation/tcp/server/address/Application.java`

```
import reactor.netty.DisposableServer;
import reactor.netty.tcp.TcpServer;

public class Application {

    public static void main(String[] args) {
        DisposableServer server =
            TcpServer.create()
                .host("localhost") ①
                .port(8080)         ②
                .bindNow();

        server.onDispose()
            .block();
    }
}
```

① Configures the TCP server host

② Configures the TCP server port

## 3.2. Writing Data

In order to send data to a connected client, you must attach an I/O handler. The I/O handler has access to `NettyOutbound` to be able to write data. The following example shows how to attach an I/O handler:

`././reactor-netty-examples/src/main/java/reactor/netty/examples/documentation/tcp/server/send/Application.java`

```
import reactor.core.publisher.Mono;
import reactor.netty.DisposableServer;
import reactor.netty.tcp.TcpServer;

public class Application {

    public static void main(String[] args) {
        DisposableServer server =
            TcpServer.create()
                .handle((inbound, outbound) ->
outbound.sendString(Mono.just("hello"))) ①
                .bindNow();

        server.onDispose()
            .block();
    }
}
```

① Sends `hello` string to the connected clients

### 3.3. Consuming Data

In order to receive data from a connected client, you must attach an I/O handler. The I/O handler has access to `NettyInbound` to be able to read data. The following example shows how to use it:

`././reactor-netty-examples/src/main/java/reactor/netty/examples/documentation/tcp/server/read/Application.java`

```
import reactor.netty.DisposableServer;
import reactor.netty.tcp.TcpServer;

public class Application {

    public static void main(String[] args) {
        DisposableServer server =
            TcpServer.create()
                .handle((inbound, outbound) -> inbound.receive().then())
①
                .bindNow();

        server.onDispose()
            .block();
    }
}
```

① Receives data from the connected clients

## 3.4. Lifecycle Callbacks

The following lifecycle callbacks are provided to let you extend the [TCP](#) server:

- [doOnBind](#): Invoked when the server channel is about to bind.
- [doOnBound](#): Invoked when the server channel is bound.
- [doOnConnection](#): Invoked when a remote client is connected
- [doOnUnbound](#): Invoked when the server channel is unbound.

The following example uses the [doOnConnection](#) callback:

././reactor-netty-examples/src/main/java/reactor/netty/examples/documentation/tcp/server/lifecycle/Application.java

```
import io.netty.handler.timeout.ReadTimeoutHandler;
import reactor.netty.DisposableServer;
import reactor.netty.tcp.TcpServer;
import java.util.concurrent.TimeUnit;

public class Application {

    public static void main(String[] args) {
        DisposableServer server =
            TcpServer.create()
                .doOnConnection(conn ->
                    conn.addHandler(new ReadTimeoutHandler(10,
TimeUnit.SECONDS))) ①
                .bindNow();

        server.onDispose()
            .block();
    }
}
```

① Netty pipeline is extended with `ReadTimeoutHandler` when a remote client is connected.

## 3.5. TCP-level Configurations

This section describes three kinds of configuration that you can use at the TCP level:

- [Setting Channel Options](#)
- [Using a Wire Logger](#)
- [Using an Event Loop Group](#)

### 3.5.1. Setting Channel Options

By default, the `TCP` server is configured with the following options:

```
../../../../reactor-netty-core/src/main/java/reactor/netty/tcp/TcpServerBind.java
```

```
TcpServerBind() {  
    Map<ChannelOption<?>, Boolean> childOptions = new HashMap<>(2);  
    childOptions.put(ChannelOption.AUTO_READ, false);  
    childOptions.put(ChannelOption.TCP_NODELAY, true);  
    this.config = new TcpServerConfig(  
        Collections.singletonMap(ChannelOption.SO_REUSEADDR, true),  
        childOptions,  
        () -> new InetSocketAddress(DEFAULT_PORT));  
}
```

If additional options are necessary or changes to the current options are needed, you can apply the following configuration:

```
../../../../reactor-netty-examples/src/main/java/reactor/netty/examples/documentation/tcp/server/channeloptions/Application.java
```

```
import io.netty.channel.ChannelOption;  
import reactor.netty.DisposableServer;  
import reactor.netty.tcp.TcpServer;  
  
public class Application {  
  
    public static void main(String[] args) {  
        DisposableServer server =  
            TcpServer.create()  
                .option(ChannelOption.CONNECT_TIMEOUT_MILLIS, 10000)  
                .bindNow();  
  
        server.onDispose()  
            .block();  
    }  
}
```

You can find more about [Netty](#) channel options at the following links:

- [ChannelOption](#)
- [Socket Options](#)

### 3.5.2. Using a Wire Logger

Reactor Netty provides wire logging for when the traffic between the peers has to be inspected. By default, wire logging is disabled. To enable it, you must set the logger `reactor.netty.tcp.TcpServer` level to `DEBUG` and apply the following configuration:

././reactor-netty-examples/src/main/java/reactor/netty/examples/documentation/tcp/server/wiretap/Application.java

```
import reactor.netty.DisposableServer;
import reactor.netty.tcp.TcpServer;

public class Application {

    public static void main(String[] args) {
        DisposableServer server =
            TcpServer.create()
                .wiretap(true) ①
                .bindNow();

        server.onDispose()
            .block();
    }
}
```

① Enables the wire logging

### 3.5.3. Using an Event Loop Group

By default, the [TCP](#) server uses an “Event Loop Group,” where the number of the worker threads equals the number of processors available to the runtime on initialization (but with a minimum value of 4). When you need a different configuration, you can use one of the [LoopResource#create](#) methods.

The default configuration for the [Event Loop Group](#) is the following:

`../../../../reactor-netty-core/src/main/java/reactor/netty/ReactorNetty.java`

```
* Default worker thread count, fallback to available processor
* (but with a minimum value of 4)
*/
public static final String IO_WORKER_COUNT = "reactor.netty.ioWorkerCount";
/**
 * Default selector thread count, fallback to -1 (no selector thread)
 */
public static final String IO_SELECT_COUNT = "reactor.netty.ioSelectCount";
/**
 * Default worker thread count for UDP, fallback to available processor
 * (but with a minimum value of 4)
*/
public static final String UDP_IO_THREAD_COUNT =
"reactor.netty.udp.ioThreadCount";
/**
 * Default quiet period that guarantees that the disposal of the underlying
LoopResources
 * will not happen, fallback to 2 seconds.
*/
public static final String SHUTDOWN QUIET_PERIOD =
"reactor.netty.ioShutdownQuietPeriod";
/**
 * Default maximum amount of time to wait until the disposal of the underlying
LoopResources
 * regardless if a task was submitted during the quiet period, fallback to 15
seconds.
*/
public static final String SHUTDOWN_TIMEOUT = "reactor.netty.ioShutdownTimeout";

/**
 * Default value whether the native transport (epoll, kqueue) will be preferred,
 * fallback it will be preferred when available
*/
public static final String NATIVE = "reactor.netty.native";
```

If changes to the these settings are needed, you can apply the following configuration:

```
../../../../reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/tcp/server/eventloop/Application.java
```

```
import reactor.netty.DisposableServer;
import reactor.netty.resources.LoopResources;
import reactor.netty.tcp.TcpServer;

public class Application {

    public static void main(String[] args) {
        LoopResources loop = LoopResources.create("event-loop", 1, 4, true);

        DisposableServer server =
            TcpServer.create()
                .runOn(loop)
                .bindNow();

        server.onDispose()
            .block();
    }
}
```

## 3.6. SSL/TLS

When you need SSL or TLS, you can apply the configuration shown in the next listing. By default, if `OpenSSL` is available, `SslProvider.OPENSSL` provider is used as a provider. Otherwise `SslProvider.JDK` is used. Switching the provider can be done through `SslContextBuilder` or by setting `-Dio.netty.handler.ssl.noOpenSsl=true`.

The following example uses `SslContextBuilder`:

```
././reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/tcp/server/security/Application.java
```

```
import io.netty.handler.ssl.SslContextBuilder;
import reactor.netty.DisposableServer;
import reactor.netty.tcp.TcpServer;
import java.io.File;

public class Application {

    public static void main(String[] args) {
        File cert = new File("certificate.crt");
        File key = new File("private.key");

        SslContextBuilder sslContextBuilder = SslContextBuilder.forServer(cert,
key);

        DisposableServer server =
            TcpServer.create()
                .secure(spec -> spec.sslContext(sslContextBuilder))
                .bindNow();

        server.onDispose()
            .block();
    }
}
```

## 3.7. Metrics

The TCP server supports built-in integration with [Micrometer](#). It exposes all metrics with a prefix of `reactor.netty.tcp.server`.

The following table provides information for the TCP server metrics:

metric name	type	description
reactor.netty.tcp.server.data.received	DistributionSummary	Amount of the data received, in bytes
reactor.netty.tcp.server.data.sent	DistributionSummary	Amount of the data sent, in bytes
reactor.netty.tcp.server.errors	Counter	Number of errors that occurred
reactor.netty.tcp.server.tls.handshake.time	Timer	Time spent for TLS handshake

These additional metrics are also available:

[ByteBufAllocator](#) metrics

metric name	type	description
reactor.netty.bytebufallocator.used.heap.memory	Gauge	The number of the bytes of the heap memory
reactor.netty.bytebufallocator.used.direct.memory	Gauge	The number of the bytes of the direct memory
reactor.netty.bytebufallocator.used.heap.arenas	Gauge	The number of heap arenas (when <code>PooledByteBufAllocator</code> )
reactor.netty.bytebufallocator.used.direct.arenas	Gauge	The number of direct arenas (when <code>PooledByteBufAllocator</code> )
reactor.netty.bytebufallocator.used.threadlocal.caches	Gauge	The number of thread local caches (when <code>PooledByteBufAllocator</code> )
reactor.netty.bytebufallocator.used.tiny.cache.size	Gauge	The size of the tiny cache (when <code>PooledByteBufAllocator</code> )
reactor.netty.bytebufallocator.used.small.cache.size	Gauge	The size of the small cache (when <code>PooledByteBufAllocator</code> )
reactor.netty.bytebufallocator.used.normal.cache.size	Gauge	The size of the normal cache (when <code>PooledByteBufAllocator</code> )
reactor.netty.bytebufallocator.used.chunk.size	Gauge	The chunk size for an arena (when <code>PooledByteBufAllocator</code> )

The following example enables that integration:

```
./../reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/tcp/server/metrics/Application.java
```

```
import reactor.netty.DisposableServer;
import reactor.netty.tcp.TcpServer;

public class Application {

    public static void main(String[] args) {
        DisposableServer server =
            TcpServer.create()
                .metrics(true) ①
                .bindNow();

        server.onDispose()
            .block();
    }
}
```

① Enables the built-in integration with Micrometer

When TCP server metrics are needed for an integration with a system other than `Micrometer` or you want to provide your own integration with `Micrometer`, you can provide your own metrics recorder,

as follows:

```
././reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/tcp/server/metrics/custom/Application.java
```

```
import reactor.netty.DisposableServer;
import reactor.netty.channel.ChannelMetricsRecorder;
import reactor.netty.tcp.TcpServer;

import java.net.SocketAddress;
import java.time.Duration;

public class Application {

    public static void main(String[] args) {
        DisposableServer server =
            TcpServer.create()
                .metrics(true, CustomChannelMetricsRecorder::new) ①
                .bindNow();

        server.onDispose()
            .block();
    }
}
```

① Enables TCP server metrics and provides `ChannelMetricsRecorder` implementation.

# Chapter 4. TCP Client

Reactor Netty provides the easy-to-use and easy-to-configure `TcpClient`. It hides most of the Netty functionality that is needed in order to create a `TCP` client and adds Reactive Streams backpressure.

## 4.1. Connect and Disconnect

To connect the `TCP` client to a given endpoint, you must create and configure a `TcpClient` instance. By default, the `host` is `localhost` and the `port` is `12012`. The following example shows how to create a `TcpClient`:

```
./../reactor-netty-  
examples/src/main/java/reactor/netty/examples/documentation/tcp/client/create/Application.java
```

```
import reactor.netty.Connection;  
import reactor.netty.TcpClient;  
  
public class Application {  
  
    public static void main(String[] args) {  
        Connection connection =  
            TcpClient.create()          ①  
                .connectNow(); ②  
  
        connection.onDispose()  
            .block();  
    }  
}
```

① Creates a `TcpClient` instance that is ready for configuring.

② Connects the client in a blocking fashion and waits for it to finish initializing.

The returned `Connection` offers a simple connection API, including `disposeNow()`, which shuts the client down in a blocking fashion.

### 4.1.1. Host and Port

To connect to a specific `host` and `port`, you can apply the following configuration to the `TCP` client. The following example shows how to do so:

`././reactor-netty-examples/src/main/java/reactor/netty/examples/documentation/tcp/client/address/Application.java`

```
import reactor.netty.Connection;
import reactor.netty.tcp.TcpClient;

public class Application {

    public static void main(String[] args) {
        Connection connection =
            TcpClient.create()
                .host("example.com") ①
                .port(80)             ②
                .connectNow();

        connection.onDispose()
            .block();
    }
}
```

① Configures the **TCP** host

② Configures the **TCP** port

## 4.2. Writing Data

To send data to a given endpoint, you must attach an I/O handler. The I/O handler has access to **NettyOutbound** to be able to write data.

`././reactor-netty-examples/src/main/java/reactor/netty/examples/documentation/tcp/client/send/Application.java`

```
import reactor.core.publisher.Mono;
import reactor.netty.Connection;
import reactor.netty.tcp.TcpClient;

public class Application {

    public static void main(String[] args) {
        Connection connection =
            TcpClient.create()
                .host("example.com")
                .port(80)
                .handle((inbound, outbound) ->
outbound.sendString(Mono.just("hello")) ①
                .connectNow();

        connection.onDispose()
            .block();
    }
}
```

① Sends `hello` string to the endpoint.

## 4.3. Consuming Data

To receive data from a given endpoint, you must attach an I/O handler. The I/O handler has access to `NettyInbound` to be able to read data. The following example shows how to do so:

`././reactor-netty-examples/src/main/java/reactor/netty/examples/documentation/tcp/client/read/Application.java`

```
import reactor.netty.Connection;
import reactor.netty.tcp.TcpClient;

public class Application {

    public static void main(String[] args) {
        Connection connection =
            TcpClient.create()
                .host("example.com")
                .port(80)
                .handle((inbound, outbound) -> inbound.receive().then())
①
                .connectNow();

        connection.onDispose()
            .block();
    }
}
```

① Receives data from a given endpoint

## 4.4. Lifecycle Callbacks

The following lifecycle callbacks are provided to let you extend the `TCP` client.

- `doOnConnect`: Invoked when the channel is about to connect.
- `doOnConnected`: Invoked after the channel has been connected.
- `doOnDisconnected`: Invoked after the channel has been disconnected.

The following example uses the `doOnConnected` callback:

```
././reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/tcp/client/lifecycle/Application.java
```

```
import io.netty.handler.timeout.ReadTimeoutHandler;
import reactor.netty.Connection;
import reactor.netty.tcp.TcpClient;
import java.util.concurrent.TimeUnit;

public class Application {

    public static void main(String[] args) {
        Connection connection =
            TcpClient.create()
                .host("example.com")
                .port(80)
                .doOnConnected(conn ->
                    conn.addHandler(new ReadTimeoutHandler(10,
TimeUnit.SECONDS))) ①
                .connectNow();

        connection.onDispose()
            .block();
    }
}
```

① Netty pipeline is extended with `ReadTimeoutHandler` when the channel has been connected.

## 4.5. TCP-level Configurations

This section describes three kinds of configuration that you can use at the TCP level:

- [Channel Options](#)
- [Wire Logger](#)
- [Event Loop Group](#)

### 4.5.1. Channel Options

By default, the `TCP` client is configured with the following options:

```
../../../../reactor-netty-core/src/main/java/reactor/netty/tcp/TcpClientConnect.java
```

```
TcpClientConnect(ConnectionProvider provider) {
    this.config = new TcpClientConfig(
        provider,
        Collections.singletonMap(ChannelOption.AUTO_READ, false),
        () ->
    AddressUtils.createUnresolved(NetUtil.LOCALHOST.getHostAddress(), DEFAULT_PORT));
}
```

If additional options are necessary or changes to the current options are needed, you can apply the following configuration:

```
../../../../reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/tcp/client/channeloptions/Application.java
```

```
import io.netty.channel.ChannelOption;
import reactor.netty.Connection;
import reactor.netty.tcp.TcpClient;

public class Application {

    public static void main(String[] args) {
        Connection connection =
            TcpClient.create()
                .host("example.com")
                .port(80)
                .option(ChannelOption.CONNECT_TIMEOUT_MILLIS, 10000)
                .connectNow();

        connection.onDispose()
            .block();
    }
}
```

You can find more about [Netty](#) channel options at the following links:

- [ChannelOption](#)
- [Socket Options](#)

#### 4.5.2. Wire Logger

Reactor Netty provides wire logging for when the traffic between the peers has to be inspected. By default, wire logging is disabled. To enable it, you must set the logger `reactor.netty.tcp.TcpClient` level to `DEBUG` and apply the following configuration:

`././reactor-netty-examples/src/main/java/reactor/netty/examples/documentation/tcp/client/wiretap/Application.java`

```
import reactor.netty.Connection;
import reactor.netty.tcp.TcpClient;

public class Application {

    public static void main(String[] args) {
        Connection connection =
            TcpClient.create()
                .wiretap(true) ①
                .host("example.com")
                .port(80)
                .connectNow();

        connection.onDispose()
            .block();
    }
}
```

① Enables the wire logging

#### 4.5.3. Event Loop Group

By default the [TCP](#) client uses an “Event Loop Group”, where the number of the worker threads equals the number of processors available to the runtime on initialization (but with a minimum value of 4). When you need a different configuration, you can use one of the [LoopResource#create](#) methods.

The following listing shows the default configuration for the Event Loop Group:

`../../../../reactor-netty-core/src/main/java/reactor/netty/ReactorNetty.java`

```
* Default worker thread count, fallback to available processor
* (but with a minimum value of 4)
*/
public static final String IO_WORKER_COUNT = "reactor.netty.ioWorkerCount";
/**
 * Default selector thread count, fallback to -1 (no selector thread)
 */
public static final String IO_SELECT_COUNT = "reactor.netty.ioSelectCount";
/**
 * Default worker thread count for UDP, fallback to available processor
 * (but with a minimum value of 4)
*/
public static final String UDP_IO_THREAD_COUNT =
"reactor.netty.udp.ioThreadCount";
/**
 * Default quiet period that guarantees that the disposal of the underlying
LoopResources
 * will not happen, fallback to 2 seconds.
*/
public static final String SHUTDOWN QUIET PERIOD =
"reactor.netty.ioShutdownQuietPeriod";
/**
 * Default maximum amount of time to wait until the disposal of the underlying
LoopResources
 * regardless if a task was submitted during the quiet period, fallback to 15
seconds.
*/
public static final String SHUTDOWN_TIMEOUT = "reactor.netty.ioShutdownTimeout";

/**
 * Default value whether the native transport (epoll, kqueue) will be preferred,
 * fallback it will be preferred when available
*/
public static final String NATIVE = "reactor.netty.native";
```

If you need changes to the these settings, you can apply the following configuration:

```
../../../../reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/tcp/client/eventloop/Application.java
```

```
import reactor.netty.Connection;
import reactor.netty.resources.LoopResources;
import reactor.netty.tcp.TcpClient;

public class Application {

    public static void main(String[] args) {
        LoopResources loop = LoopResources.create("event-loop", 1, 4, true);

        Connection connection =
            TcpClient.create()
                .host("example.com")
                .port(80)
                .runOn(loop)
                .connectNow();

        connection.onDispose()
            .block();
    }
}
```

## 4.6. Connection Pool

By default, the [TCP](#) client uses a “fixed” connection pool with [500](#) as the maximum number of channels, [45s](#) as the pending acquire timeout and [1000](#) as the maximum number of the registered requests for acquire to keep in the pending queue. This means that the implementation creates a new channel if someone tries to acquire a channel but none is in the pool. When the maximum number of the channels in the pool is reached, new tries to acquire a channel are delayed until a channel is returned to the pool again. The implementation uses [FIFO](#) order for channels in the pool. By default, there is no idle time specified for the channels in the pool.

If you need to disable the connection pool, you can apply the following configuration:

*././reactor-netty-examples/src/main/java/reactor/netty/examples/documentation/tcp/client/pool/Application.java*

```
import reactor.netty.Connection;
import reactor.netty.tcp.TcpClient;

public class Application {

    public static void main(String[] args) {
        Connection connection =
            TcpClient.newConnection()
                .host("example.com")
                .port(80)
                .connectNow();

        connection.onDispose()
            .block();
    }
}
```

If you need to specify an idle time for the channels in the connection pool, you can apply the following configuration:

```
././reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/tcp/client/pool/config/Application.java
```

```
import reactor.netty.Connection;
import reactor.netty.resources.ConnectionProvider;
import reactor.netty.tcp.TcpClient;
import java.time.Duration;

public class Application {

    public static void main(String[] args) {
        ConnectionProvider provider =
            ConnectionProvider.builder("fixed")
                .maxConnections(50)
                .pendingAcquireTimeout(Duration.ofMillis(30000))
                .maxIdleTime(Duration.ofMillis(60))
                .build();

        Connection connection =
            TcpClient.create(provider)
                .host("example.com")
                .port(80)
                .connectNow();

        connection.onDispose()
            .block();
    }
}
```

 When you expect a high load, be cautious with a connection pool with a very high value for maximum connections. You might experience `reactor.netty.http.client.PrematureCloseException` exception with a root cause "Connect Timeout" due to too many concurrent connections opened/acquired.

#### 4.6.1. Metrics

The pooled `ConnectionProvider` supports built-in integration with `Micrometer`. It exposes all metrics with a prefix of `reactor.netty.connection.provider`.

Pooled `ConnectionProvider` metrics

metric name	type	description
<code>reactor.netty.connection.provider.total.connections</code>	Gauge	The number of all connections, active or idle
<code>reactor.netty.connection.provider.active.connections</code>	Gauge	The number of the connections that have been successfully acquired and are in active use

metric name	type	description
reactor.netty.connection.provider.idle.connections	Gauge	The number of the idle connections
reactor.netty.connection.provider.pending.connections	Gauge	The number of requests that are waiting for a connection

The following example enables that integration:

```
./../reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/tcp/client/pool/metrics/Application.java

import reactor.netty.Connection;
import reactor.netty.resources.ConnectionProvider;
import reactor.netty.tcp.TcpClient;

public class Application {

    public static void main(String[] args) {
        ConnectionProvider provider =
            ConnectionProvider.builder("fixed")
                .maxConnections(50)
                .metrics(true) ①
                .build();

        Connection connection =
            TcpClient.create(provider)
                .host("example.com")
                .port(80)
                .connectNow();

        connection.onDispose()
            .block();
    }
}
```

① Enables the built-in integration with Micrometer

## 4.7. SSL and TLS

When you need SSL or TLS, you can apply the following configuration. By default, if `OpenSSL` is available, the `SslProvider.OPENSSL` provider is used as a provider. Otherwise, the provider is `SslProvider.JDK`. You can switch the provider by using `SslContextBuilder` or by setting `-Dio.netty.handler.ssl.noOpenSsl=true`.

The following example uses `SslContextBuilder`:

././reactor-netty-  
examples/src/main/java/reactor/netty/examples/documentation/tcp/client/security/Application.java

```
import io.netty.handler.ssl.SslContextBuilder;
import reactor.netty.Connection;
import reactor.netty.tcp.TcpClient;

public class Application {

    public static void main(String[] args) {
        SslContextBuilder sslContextBuilder = SslContextBuilder.forClient();

        Connection connection =
            TcpClient.create()
                .host("example.com")
                .port(443)
                .secure(spec -> spec.sslContext(sslContextBuilder))
                .connectNow();

        connection.onDispose()
            .block();
    }
}
```

## 4.8. Proxy Support

The TCP client supports the proxy functionality provided by Netty and provides a way to specify “non proxy hosts” through the [ProxyProvider](#) builder. The following example uses [ProxyProvider](#):

```
././reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/tcp/client/proxy/Application.java
```

```
import reactor.netty.Connection;
import reactor.netty.transport.ProxyProvider;
import reactor.netty.tcp.TcpClient;

public class Application {

    public static void main(String[] args) {
        Connection connection =
            TcpClient.create()
                .host("example.com")
                .port(80)
                .proxy(spec -> spec.type(ProxyProvider.Proxy.SOCKS4)
                    .host("proxy")
                    .port(8080)
                    .nonProxyHosts("localhost"))
                .connectNow();

        connection.onDispose()
            .block();
    }
}
```

## 4.9. Metrics

The TCP client supports built-in integration with [Micrometer](#). It exposes all metrics with a prefix of `reactor.netty.tcp.client`.

The following table provides information for the TCP client metrics:

metric name	type	description
reactor.netty.tcp.client.data.received	DistributionSummary	Amount of the data received, in bytes
reactor.netty.tcp.client.data.sent	DistributionSummary	Amount of the data sent, in bytes
reactor.netty.tcp.client.errors	Counter	Number of errors that occurred
reactor.netty.tcp.client.tls.handshake.time	Timer	Time spent for TLS handshake
reactor.netty.tcp.client.connect.time	Timer	Time spent for connecting to the remote address
reactor.netty.tcp.client.address.resolver	Timer	Time spent for resolving the address

These additional metrics are also available:

## Pooled ConnectionProvider metrics

metric name	type	description
reactor.netty.connection.provider.total.connections	Gauge	The number of all connections, active or idle
reactor.netty.connection.provider.active.connections	Gauge	The number of the connections that have been successfully acquired and are in active use
reactor.netty.connection.provider.idle.connections	Gauge	The number of the idle connections
reactor.netty.connection.provider.pending.connections	Gauge	The number of requests that are waiting for a connection

## ByteBufAllocator metrics

metric name	type	description
reactor.netty.bytebufallocator.used.heap.memory	Gauge	The number of the bytes of the heap memory
reactor.netty.bytebufallocator.used.direct.memory	Gauge	The number of the bytes of the direct memory
reactor.netty.bytebufallocator.used.heap.arenas	Gauge	The number of heap arenas (when PooledByteBufAllocator)
reactor.netty.bytebufallocator.used.direct.arenas	Gauge	The number of direct arenas (when PooledByteBufAllocator)
reactor.netty.bytebufallocator.used.threadlocal.caches	Gauge	The number of thread local caches (when PooledByteBufAllocator)
reactor.netty.bytebufallocator.used.tiny.cache.size	Gauge	The size of the tiny cache (when PooledByteBufAllocator)
reactor.netty.bytebufallocator.used.small.cache.size	Gauge	The size of the small cache (when PooledByteBufAllocator)
reactor.netty.bytebufallocator.used.normal.cache.size	Gauge	The size of the normal cache (when PooledByteBufAllocator)
reactor.netty.bytebufallocator.used.chunk.size	Gauge	The chunk size for an arena (when PooledByteBufAllocator)

The following example enables that integration:

`././reactor-netty-examples/src/main/java/reactor/netty/examples/documentation/tcp/client/metrics/Application.java`

```
import reactor.netty.Connection;
import reactor.netty.tcp.TcpClient;

public class Application {

    public static void main(String[] args) {
        Connection connection =
            TcpClient.create()
                .host("example.com")
                .port(80)
                .metrics(true) ①
                .connectNow();

        connection.onDispose()
            .block();
    }
}
```

① Enables the built-in integration with Micrometer

When TCP client metrics are needed for an integration with a system other than [Micrometer](#) or you want to provide your own integration with [Micrometer](#), you can provide your own metrics recorder, as follows:

`././reactor-netty-examples/src/main/java/reactor/netty/examples/documentation/tcp/client/metrics/custom/Application.java`

```
import reactor.netty.Connection;
import reactor.netty.channel.ChannelMetricsRecorder;
import reactor.netty.tcp.TcpClient;

import java.net.SocketAddress;
import java.time.Duration;

public class Application {

    public static void main(String[] args) {
        Connection connection =
            TcpClient.create()
                .host("example.com")
                .port(80)
                .metrics(true, CustomChannelMetricsRecorder::new) ①
                .connectNow();

        connection.onDispose()
            .block();
    }
}
```

① Enables TCP client metrics and provides `ChannelMetricsRecorder` implementation.

# Chapter 5. HTTP Server

Reactor Netty provides the easy-to-use and easy-to-configure `HttpServer` class. It hides most of the Netty functionality that is needed in order to create a HTTP server and adds Reactive Streams backpressure.

## 5.1. Starting and Stopping

To start an HTTP server, you must create and configure a `HttpServer` instance. By default, the `host` is configured for any local address, and the system picks up an ephemeral port when the `bind` operation is invoked. The following example shows how to create an `HttpServer` instance:

```
././reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/http/server/create/Application.java
```

```
import reactor.netty.DisposableServer;
import reactor.netty.http.server.HttpServer;

public class Application {

    public static void main(String[] args) {
        DisposableServer server =
            HttpServer.create()    ①
                .bindNow(); ②

        server.onDispose()
            .block();
    }
}
```

① Creates an `HttpServer` instance ready for configuring.

② Starts the server in a blocking fashion and waits for it to finish initializing.

The returned `DisposableServer` offers a simple server API, including `disposeNow()`, which shuts the server down in a blocking fashion.

### 5.1.1. Host and Port

To serve on a specific `host` and `port`, you can apply the following configuration to the `HTTP` server:

```
././reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/http/server/address/Application.java
```

```
import reactor.netty.DisposableServer;
import reactor.netty.http.server.HttpServer;

public class Application {

    public static void main(String[] args) {
        DisposableServer server =
            HttpServer.create()
                .host("localhost") ①
                .port(8080)         ②
                .bindNow();

        server.onDispose()
            .block();
    }
}
```

① Configures the `HTTP` server host

② Configures the `HTTP` server port

## 5.2. Routing HTTP

Defining routes for the `HTTP` server requires configuring the provided `HttpServerRoutes` builder. The following example shows how to do so:

```
././reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/http/server/routing/Application.java
```

```
import reactor.core.publisher.Mono;
import reactor.netty.DisposableServer;
import reactor.netty.http.server.HttpServer;

public class Application {

    public static void main(String[] args) {
        DisposableServer server =
            HttpServer.create()
                .route(routes ->
                    routes.get("/hello",           ①
                               (request, response) ->
                    response.sendString(Mono.just("Hello World!")))
                    .post("/echo",               ②
                           (request, response) ->
                    response.send(request.receive().retain()))
                    .get("/path/{param}",        ③
                          (request, response) ->
                    response.sendString(Mono.just(request.param("param"))))
                    .ws("/ws",                  ④
                         (wsInbound, wsOutbound) ->
                    wsOutbound.send(wsInbound.receive().retain())))
                .bindNow();
    }

    server.onDispose()
        .block();
}
}
```

① Serves a `GET` request to `/hello` and returns `Hello World!`

② Serves a `POST` request to `/echo` and returns the received request body as a response.

③ Serves a `GET` request to `/path/{param}` and returns the value of the path parameter.

④ Serves websocket to `/ws` and returns the received incoming data as outgoing data.



The server routes are unique and only the first matching in order of declaration is invoked.

### 5.2.1. SSE

The following code shows how you can configure the `HTTP` server to serve `Server-Sent Events`:

```
././reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/http/server/sse/Application.java
```

```

import com.fasterxml.jackson.databind.ObjectMapper;
import io.netty.buffer.ByteBuf;
import io.netty.buffer.ByteBufAllocator;
import org.reactivestreams.Publisher;
import reactor.core.publisher.Flux;
import reactor.netty.DisposableServer;
import reactor.netty.http.server.HttpServer;
import reactor.netty.http.server.HttpServerRequest;
import reactor.netty.http.server.HttpServerResponse;

import java.io.ByteArrayOutputStream;
import java.nio.charset.Charset;
import java.time.Duration;
import java.util.function.BiFunction;

public class Application {

    public static void main(String[] args) {
        DisposableServer server =
            HttpServer.create()
                .route(routes -> routes.get("/sse", serveSse()))
                .bindNow();

        server.onDispose()
            .block();
    }

    /**
     * Prepares SSE response
     * The "Content-Type" is "text/event-stream"
     * The flushing strategy is "flush after every element" emitted by the
     provided Publisher
     */
    private static BiFunction<HttpServerRequest, HttpServerResponse,
Publisher<Void>> serveSse() {
        Flux<Long> flux = Flux.interval(Duration.ofSeconds(10));
        return (request, response) ->
            response.sse()
                .send(flux.map(Application::toByteBuf), b -> true);
    }

    /**
     * Transforms the Object to ByteBuf following the expected SSE format.
     */
    private static ByteBuf toByteBuf(Object any) {
        ByteArrayOutputStream out = new ByteArrayOutputStream();
        try {
            out.write("data: ".getBytes(Charset.defaultCharset()));
            MAPPER.writeValue(out, any);
            out.write("\n\n".getBytes(Charset.defaultCharset()));
        }

```

```

        catch (Exception e) {
            throw new RuntimeException(e);
        }
        return ByteBufAllocator.DEFAULT
            .buffer()
            .writeBytes(out.toByteArray());
    }

    private static final ObjectMapper MAPPER = new ObjectMapper();
}

```

## 5.2.2. Static Resources

The following code shows how you can configure the [HTTP](#) server to serve static resources:

```

././reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/http/server/staticresources/Application.java

import reactor.netty.DisposableServer;
import reactor.netty.http.server.HttpServer;

import java.net.URISyntaxException;
import java.nio.file.Path;
import java.nio.file.Paths;

public class Application {

    public static void main(String[] args) throws URISyntaxException {
        Path file =
Paths.get(Application.class.getResource("/logback.xml").toURI());
        DisposableServer server =
HttpServer.create()
    .route(routes -> routes.file("/index.html", file))
    .bindNow();

        server.onDispose()
    .block();
    }
}

```

## 5.3. Writing Data

To send data to a connected client, you must attach an I/O handler by using either [handle\(...\)](#) or [route\(...\)](#). The I/O handler has access to [HttpServerResponse](#), to be able to write data. The following example uses the [handle\(...\)](#) method:

```
../../../../reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/http/server/send/Application.java
```

```
import reactor.core.publisher.Mono;
import reactor.netty.DisposableServer;
import reactor.netty.http.server.HttpServer;

public class Application {

    public static void main(String[] args) {
        DisposableServer server =
            HttpServer.create()
                .handle((request, response) ->
        response.sendString(Mono.just("hello")) ①
                .bindNow();

        server.onDispose()
            .block();
    }
}
```

① Sends `hello` string to the connected clients

### 5.3.1. Adding Headers and Other Metadata

When you send data to the connected clients, you may need to send additional headers, cookies, status code, and other metadata. You can provide this additional metadata by using `HttpServerResponse`. The following example shows how to do so:

```
././reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/http/server/send/headers/Application.j
ava
```

```
import io.netty.handler.codec.http.HttpHeaderNames;
import io.netty.handler.codec.http.HttpResponseStatus;
import reactor.core.publisher.Mono;
import reactor.netty.DisposableServer;
import reactor.netty.http.server.HttpServer;

public class Application {

    public static void main(String[] args) {
        DisposableServer server =
            HttpServer.create()
                .route(routes ->
                    routes.get("/hello",
                        (request, response) ->
                            response.status(HttpResponseStatus.OK)

                    .header(HttpHeaderNames.CONTENT_LENGTH, "12")
                        .sendString(Mono.just("Hello
World!")))
                .bindNow();

        server.onDispose()
            .block();
    }
}
```

### 5.3.2. Compression

You can configure the **HTTP** server to send a compressed response, depending on the request header **Accept-Encoding**.

**Reactor Netty** provides three different strategies for compressing the outgoing data:

- **compress(boolean)**: Depending on the boolean that is provided, the compression is enabled (**true**) or disabled (**false**).
- **compress(int)**: The compression is performed once the response size exceeds the given value (in bytes).
- **compress(BiPredicate<HttpServerRequest, HttpServerResponse>)**: The compression is performed if the predicate returns **true**.

The following example uses the **compress** method (set to **true**) to enable compression:

```
../../../../reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/http/server/compression/Application.j
ava
```

```
import reactor.netty.DisposableServer;
import reactor.netty.http.server.HttpServer;

import java.net.URISyntaxException;
import java.nio.file.Path;
import java.nio.file.Paths;

public class Application {

    public static void main(String[] args) throws URISyntaxException {
        Path file =
Paths.get(Application.class.getResource("/logback.xml").toURI());
        DisposableServer server =
            HttpServer.create()
                .compress(true)
                .route(routes -> routes.file("/index.html", file))
                .bindNow();

        server.onDispose()
            .block();
    }
}
```

## 5.4. Consuming Data

To receive data from a connected client, you must attach an I/O handler by using either `handle(...)` or `route(...)`. The I/O handler has access to `HttpServerRequest`, to be able to read data.

The following example uses the `handle(...)` method:

`././reactor-netty-examples/src/main/java/reactor/netty/examples/documentation/http/server/read/Application.java`

```
import reactor.netty.DisposableServer;
import reactor.netty.http.server.HttpServer;

public class Application {

    public static void main(String[] args) {
        DisposableServer server =
            HttpServer.create()
                .handle((request, response) -> request.receive().then())
①
                .bindNow();

        server.onDispose()
            .block();
    }
}
```

① Receives data from the connected clients

#### 5.4.1. Reading Headers, URI Params, and other Metadata

When you receive data from the connected clients, you might need to check request headers, parameters, and other metadata. You can obtain this additional metadata by using `HttpServerRequest`. The following example shows how to do so:

```
../../../../reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/http/server/read/headers/Application.j
ava
```

```
import reactor.core.publisher.Mono;
import reactor.netty.DisposableServer;
import reactor.netty.http.server.HttpServer;

public class Application {

    public static void main(String[] args) {
        DisposableServer server =
            HttpServer.create()
                .route(routes ->
                    routes.get("/{param}",
                        (request, response) -> {
                            if (request.requestHeaders().contains("Some-
Header")) {
                                return
                            response.sendString(Mono.just(request.param("param")));
                        }
                        return response.sendNotFound();
                    }))
                .bindNow();

        server.onDispose()
            .block();
    }
}
```

## Obtaining the Remote (Client) Address

In addition to the metadata that you can obtain from the request, you can also receive the **host (server)** address, the **remote (client)** address and the **scheme**. Depending on the chosen factory method, you can retrieve the information directly from the channel or by using the **Forwarded** or **X-Forwarded-\*** HTTP request headers. The following example shows how to do so:

```
../../../../reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/http/server/clientaddress/Application.j
ava
```

```
import reactor.core.publisher.Mono;
import reactor.netty.DisposableServer;
import reactor.netty.http.server.HttpServer;

public class Application {

    public static void main(String[] args) {
        DisposableServer server =
            HttpServer.create()
                .forwarded(true) ①
                .route(routes ->
                    routes.get("/clientip",
                        (request, response) ->
                            response.sendString(Mono.just(request.remoteAddress()) ②
                                .getHostString())))
                .bindNow();

        server.onDispose()
            .block();
    }
}
```

① Specifies that the information about the connection is to be obtained from the [Forwarded](#) and [X-Forwarded-\\*](#) HTTP request headers, if possible.

② Returns the address of the remote (client) peer.

## 5.4.2. HTTP Request Decoder

By default, [Netty](#) configures some restrictions for the incoming requests, such as:

- The maximum length of the initial line.
- The maximum length of all headers.
- The maximum length of the content or each chunk.

For more information, see [HttpRequestDecoder](#) and [HttpServerUpgradeHandler](#)

By default, the [HTTP](#) server is configured with the following settings:

```
../../../../reactor-netty-http/src/main/java/reactor/netty/http/HttpDecoderSpec.java
```

```
public static final int DEFAULT_MAX_INITIAL_LINE_LENGTH = 4096;  
public static final int DEFAULT_MAX_HEADER_SIZE = 8192;  
public static final int DEFAULT_MAX_CHUNK_SIZE = 8192;  
public static final boolean DEFAULT_VALIDATE_HEADERS = true;  
public static final int DEFAULT_INITIAL_BUFFER_SIZE = 128;
```

```
../../../../reactor-netty-http/src/main/java/reactor/netty/http/server/HttpRequestDecoderSpec.java
```

```
public static final int DEFAULT_H2C_MAX_CONTENT_LENGTH = 0;
```

When you need to change these default settings, you can configure the **HTTP** server as follows:

```
../../../../reactor-netty-
```

```
examples/src/main/java/reactor/netty/examples/documentation/http/server/requestdecoder/Application.java
```

```
import reactor.core.publisher.Mono;  
import reactor.netty.DisposableServer;  
import reactor.netty.http.server.HttpServer;  
  
public class Application {  
  
    public static void main(String[] args) {  
        DisposableServer server =  
            HttpServer.create()  
                .httpRequestDecoder(spec -> spec.maxHeaderSize(16384))  
        ①  
                .handle((request, response) ->  
        response.sendString(Mono.just("hello"))  
                .bindNow();  
  
        server.onDispose()  
            .block();  
    }  
}
```

① The maximum length of all headers will be **16384**. When this value is exceeded, a **TooLongFrameException** is raised.

## 5.5. TCP-level Configuration

When you need to change configuration on the TCP level, you can use the following snippet to

extend the default [TCP](#) server configuration:

```
././reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/http/server/channeloptions/Application.java
```

```
import io.netty.channel.ChannelOption;
import reactor.netty.DisposableServer;
import reactor.netty.http.server.HttpServer;

public class Application {

    public static void main(String[] args) {
        DisposableServer server =
            HttpServer.create()
                .option(ChannelOption.CONNECT_TIMEOUT_MILLIS, 10000)
                .bindNow();

        server.onDispose()
            .block();
    }
}
```

See [TCP Server](#) for more detail about TCP-level configuration.

### 5.5.1. Wire Logger

[Reactor Netty](#) provides wire logging for when you need to inspect the traffic between the peers. By default, wire logging is disabled. To enable it, you must set the logger `reactor.netty.http.server.HttpServer` level to `DEBUG` and apply the following configuration:

`././reactor-netty-examples/src/main/java/reactor/netty/examples/documentation/http/server/wiretap/Application.java`

```
import reactor.netty.DisposableServer;
import reactor.netty.http.server.HttpServer;

public class Application {

    public static void main(String[] args) {
        DisposableServer server =
            HttpServer.create()
                .wiretap(true) ①
                .bindNow();

        server.onDispose()
            .block();
    }
}
```

① Enables the wire logging

## 5.6. SSL and TLS

When you need SSL or TLS, you can apply the configuration shown in the next example. By default, if `OpenSSL` is available, `SslProvider.OPENSSL` provider is used as a provider. Otherwise `SslProvider.JDK` is used. You can switch the provider by using `SslContextBuilder` or by setting `-Dio.netty.handler.ssl.noOpenSsl=true`.

The following example uses `SslContextBuilder`:

```
././reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/http/server/security/Application.java
```

```
import io.netty.handler.ssl.SslContextBuilder;
import reactor.netty.DisposableServer;
import reactor.netty.http.server.HttpServer;
import java.io.File;

public class Application {

    public static void main(String[] args) {
        File cert = new File("certificate.crt");
        File key = new File("private.key");

        SslContextBuilder sslContextBuilder = SslContextBuilder.forServer(cert,
key);

        DisposableServer server =
            HttpServer.create()
                .secure(spec -> spec.sslContext(sslContextBuilder))
                .bindNow();

        server.onDispose()
            .block();
    }
}
```

## 5.7. HTTP Access Log

The current logging support provides only the [Common Log Format](#).

You can use `-Dreactor.netty.http.server.accessLogEnabled=true` to enable the [HTTP](#) access log. By default, it is disabled.

You can use the following configuration (for Logback or similar logging frameworks) to have a separate [HTTP](#) access log file:

```
<appender name="accessLog" class="ch.qos.logback.core.FileAppender">
    <file>access_log.log</file>
    <encoder>
        <pattern>%msg%n</pattern>
    </encoder>
</appender>
<appender name="async" class="ch.qos.logback.classic.AsyncAppender">
    <appender-ref ref="accessLog" />
</appender>

<logger name="reactor.netty.http.server.AccessLog" level="INFO"
additivity="false">
    <appender-ref ref="async"/>
</logger>
```

## 5.8. HTTP/2

By default, the `HTTP` server supports `HTTP/1.1`. If you need `HTTP/2`, you can get it through configuration. In addition to the protocol configuration, if you need `H2` but not `H2C (cleartext)`, you must also configure SSL.



As that protocol is not supported “out-of-the-box” by JDK8, you need an additional dependency to a native library that supports it—for example, `netty-tcnative-boringssl-static`.

The following listing presents a simple `H2` example:

```
./../reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/http/server/http2/H2Application.java
```

```
import io.netty.handler.ssl.SslContextBuilder;
import reactor.core.publisher.Mono;
import reactor.netty.DisposableServer;
import reactor.netty.http.HttpProtocol;
import reactor.netty.http.server.HttpServer;
import java.io.File;

public class H2Application {

    public static void main(String[] args) {
        File cert = new File("certificate.crt");
        File key = new File("private.key");

        SslContextBuilder sslContextBuilder = SslContextBuilder.forServer(cert,
key);

        DisposableServer server =
            HttpServer.create()
                .port(8080)
                .protocol(HttpProtocol.H2) ①
                .secure(spec -> spec.sslContext(sslContextBuilder)) ②
                .handle((request, response) ->
response.sendString(Mono.just("hello")))
                .bindNow();

        server.onDispose()
            .block();
    }
}
```

① Configures the server to support only **HTTP/2**

② Configures **SSL**

The application should now behave as follows:

```
$ curl --http2 https://localhost:8080 -i
HTTP/2 200

hello
```

The following listing presents a simple **H2C** example:

```
./../reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/http/server/http2/H2CApplication.java
```

```
import reactor.core.publisher.Mono;
import reactor.netty.DisposableServer;
import reactor.netty.http.HttpProtocol;
import reactor.netty.http.server.HttpServer;

public class H2CApplication {

    public static void main(String[] args) {
        DisposableServer server =
            HttpServer.create()
                .port(8080)
                .protocol(HttpProtocol.H2C)
                .handle((request, response) ->
        response.sendString(Mono.just("hello")))
                .bindNow();

        server.onDispose()
            .block();
    }
}
```

The application should now behave as follows:

```
$ curl --http2-prior-knowledge http://localhost:8080 -i
HTTP/2 200

hello
```

## 5.9. Metrics

The HTTP server supports built-in integration with [Micrometer](#). It exposes all metrics with a prefix of `reactor.netty.http.server`.

The following table provides information for the HTTP server metrics:

metric name	type	description
reactor.netty.http.server.data.received	DistributionSummary	Amount of the data received, in bytes
reactor.netty.http.server.data.sent	DistributionSummary	Amount of the data sent, in bytes
reactor.netty.http.server.errors	Counter	Number of errors that occurred

<b>metric name</b>	<b>type</b>	<b>description</b>
reactor.netty.http.server.data.received.time	Timer	Time spent in consuming incoming data
reactor.netty.http.server.data.send.time	Timer	Time spent in sending outgoing data
reactor.netty.http.server.response.time	Timer	Total time for the request/response

These additional metrics are also available:

#### [ByteBufAllocator](#) metrics

<b>metric name</b>	<b>type</b>	<b>description</b>
reactor.netty.bytebuf_allocator.used.heap.memory	Gauge	The number of the bytes of the heap memory
reactor.netty.bytebuf_allocator.used.direct.memory	Gauge	The number of the bytes of the direct memory
reactor.netty.bytebuf_allocator.used.heap.arenas	Gauge	The number of heap arenas (when <a href="#">PooledByteBufAllocator</a> )
reactor.netty.bytebuf_allocator.used.direct.arenas	Gauge	The number of direct arenas (when <a href="#">PooledByteBufAllocator</a> )
reactor.netty.bytebuf_allocator.used.threadlocal.caches	Gauge	The number of thread local caches (when <a href="#">PooledByteBufAllocator</a> )
reactor.netty.bytebuf_allocator.used.tiny.cache.size	Gauge	The size of the tiny cache (when <a href="#">PooledByteBufAllocator</a> )
reactor.netty.bytebuf_allocator.used.small.cache.size	Gauge	The size of the small cache (when <a href="#">PooledByteBufAllocator</a> )
reactor.netty.bytebuf_allocator.used.normal.cache.size	Gauge	The size of the normal cache (when <a href="#">PooledByteBufAllocator</a> )
reactor.netty.bytebuf_allocator.used.chunk.size	Gauge	The chunk size for an arena (when <a href="#">PooledByteBufAllocator</a> )

The following example enables that integration:

././reactor-netty-examples/src/main/java/reactor/netty/examples/documentation/http/server/metrics/Application.java

```
import io.micrometer.core.instrument.Metrics;
import io.micrometer.core.instrument.config.MeterFilter;
import reactor.core.publisher.Mono;
import reactor.netty.DisposableServer;
import reactor.netty.http.server.HttpServer;

public class Application {

    public static void main(String[] args) {
        Metrics.globalRegistry ①
            .config()

        .meterFilter(MeterFilter.maximumAllowableTags("reactor.netty.http.server", "URI",
100, MeterFilter.deny()));

        DisposableServer server =
            HttpServer.create()
                .metrics(true, s -> {
                    if (s.startsWith("/stream/")) { ②
                        return "/stream/{n}";
                    }
                    else if (s.startsWith("/bytes/")) {
                        return "/bytes/{n}";
                    }
                    return s;
                }) ③
                .route(r ->
                    r.get("/stream/{n}",
                        (req, res) ->
                            res.sendString(Mono.just(req.param("n"))))
                    .get("/bytes/{n}",
                        (req, res) ->
                            res.sendString(Mono.just(req.param("n")))))
                .bindNow();

        server.onDispose()
            .block();
    }
}
```

① Applies upper limit for the meters with **URI** tag

② Templatized URIs will be used as an URI tag value when possible

③ Enables the built-in integration with Micrometer

In order to avoid a memory and CPU overhead of the enabled metrics, it is important to convert the real URIs to templated URIs when possible. Without a conversion to a template-like form, each distinct URI leads to the creation of a distinct tag, which takes a lot of memory for the metrics.

Always apply an upper limit for the meters with URI tags. Configuring an upper limit on the number of meters can help in cases when the real URIs cannot be templated. You can find more information at [maximumAllowableTags](#).

When HTTP server metrics are needed for an integration with a system other than [Micrometer](#) or you want to provide your own integration with [Micrometer](#), you can provide your own metrics recorder, as follows:

```
././reactor-netty-  
examples/src/main/java/reactor/netty/examples/documentation/http/server/metrics/custom/Application.java
```

```
import reactor.core.publisher.Mono;  
import reactor.netty.DisposableServer;  
import reactor.netty.channel.ChannelMetricsRecorder;  
import reactor.netty.http.server.HttpServer;  
  
import java.net.SocketAddress;  
import java.time.Duration;  
  
public class Application {  
  
    public static void main(String[] args) {  
        DisposableServer server =  
            HttpServer.create()  
                .metrics(true, CustomHttpServerMetricsRecorder::new) ①  
                .route(r ->  
                    r.get("/stream/{n}",  
                          (req, res) ->  
                            res.sendString(Mono.just(req.param("n"))))  
                    .get("/bytes/{n}",  
                          (req, res) ->  
                            res.sendString(Mono.just(req.param("n")))))  
                .bindNow();  
  
        server.onDispose()  
            .block();  
    }  
}
```

① Enables HTTP server metrics and provides [HttpServerMetricsRecorder](#) implementation.

# Chapter 6. HTTP Client

Reactor Netty provides the easy-to-use and easy-to-configure [HttpClient](#). It hides most of the Netty functionality that is required to create a [HTTP](#) client and adds Reactive Streams backpressure.

## 6.1. Connect

To connect the [HTTP](#) client to a given [HTTP](#) endpoint, you must create and configure a [HttpClient](#) instance. The following example shows how to do so:

```
./../reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/http/client/connect/Application.java
```

```
import reactor.netty.http.client.HttpClient;
import reactor.netty.client.HttpClientResponse;

public class Application {

    public static void main(String[] args) {
        HttpClient client = HttpClient.create(); ①

        client.get()                      ②
            .uri("http://example.com/") ③
            .response()                 ④
            .block();
    }
}
```

① Creates a [HttpClient](#) instance ready for configuring.

② Specifies that [GET](#) method will be used.

③ Specifies the path.

④ Obtains the response [HttpClientResponse](#)

The following example uses [WebSocket](#):

././reactor-netty-examples/src/main/java/reactor/netty/examples/documentation/http/client/websocket/Application.java

```
import io.netty.buffer.Unpooled;
import io.netty.util.CharsetUtil;
import reactor.core.publisher.Flux;
import reactor.netty.http.client.HttpClient;

public class Application {

    public static void main(String[] args) {
        HttpClient client = HttpClient.create();

        client.websocket()
            .uri("wss://echo.websocket.org")
            .handle((inbound, outbound) -> {
                inbound.receive()
                    .asString()
                    .take(1)
                    .subscribe(System.out::println);

                final byte[] msgBytes =
"hello".getBytes(CharsetUtil.ISO_8859_1);
                return outbound.send(Flux.just(Unpooled.wrappedBuffer(msgBytes),
Unpooled.wrappedBuffer(msgBytes)))
                    .neverComplete();
            })
            .blockLast();
    }
}
```

### 6.1.1. Host and Port

In order to connect to a specific host and port, you can apply the following configuration to the [HTTP client](#):

```
../../../../reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/http/client/address/Application.java
```

```
import reactor.netty.http.client.HttpClient;
import reactor.netty.http.client.HttpClientResponse;

public class Application {

    public static void main(String[] args) {
        HttpClient client =
            HttpClient.create()
                .host("example.com") ①
                .port(80);             ②

        client.get()
            .uri("/")
            .response()
            .block();
    }
}
```

① Configures the **HTTP** host

② Configures the **HTTP** port

## 6.2. Writing Data

To send data to a given **HTTP** endpoint, you can provide a **Publisher** by using the `send(Publisher)` method. By default, `Transfer-Encoding: chunked` is applied for those **HTTP** methods for which a request body is expected. `Content-Length` provided through request headers disables `Transfer-Encoding: chunked`, if necessary. The following example sends `hello`:

`././reactor-netty-examples/src/main/java/reactor/netty/examples/documentation/http/client/send/Application.java`

```
import reactor.core.publisher.Mono;
import reactor.netty.ByteBufFlux;
import reactor.netty.http.client.HttpClient;
import reactor.netty.http.client.HttpClientResponse;

public class Application {

    public static void main(String[] args) {
        HttpClient client = HttpClient.create();

        client.post()
            .uri("http://example.com/")
            .send(ByteBufFlux.fromString(Mono.just("hello"))) ①
            .response()
            .block();
    }
}
```

① Sends a `hello` string to the given `HTTP` endpoint

### 6.2.1. Adding Headers and Other Metadata

When sending data to a given `HTTP` endpoint, you may need to send additional headers, cookies and other metadata. You can use the following configuration to do so:

```
../../../../reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/http/client/send/headers/Application.java
```

```
import io.netty.handler.codec.http.HttpHeaderNames;
import reactor.core.publisher.Mono;
import reactor.netty.ByteBufFlux;
import reactor.netty.http.client.HttpClient;
import reactor.netty.http.client.HttpClientResponse;

public class Application {

    public static void main(String[] args) {
        HttpClient client =
            HttpClient.create()
                .headers(h -> h.set(HttpHeaderNames.CONTENT_LENGTH, 5));
①

        client.post()
            .uri("http://example.com/")
            .send(ByteBufFlux.fromString(Mono.just("hello")))
            .response()
            .block();
    }
}
```

① Disables `Transfer-Encoding: chunked` and provides `Content-Length` header.

## Compression

You can enable compression on the `HTTP` client, which means the request header `Accept-Encoding` is added to the request headers. The following example shows how to do so:

```
../../../../reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/http/client/compression/Application.java
```

```
import reactor.netty.http.client.HttpClient;
import reactor.netty.http.client.HttpClientResponse;

public class Application {

    public static void main(String[] args) {
        HttpClient client =
            HttpClient.create()
                .compress(true);

        client.get()
            .uri("http://example.com/")
            .response()
            .block();
    }
}
```

## Auto-Redirect Support

You can configure the [HTTP](#) client to enable auto-redirect support.

Reactor Netty provides two different strategies for auto-redirect support:

- [followRedirect\(boolean\)](#): Specifies whether HTTP auto-redirect support is enabled for statuses `301|302|307|308`.
- [followRedirect\(BiPredicate<HttpClientRequest, HttpClientResponse>\)](#): Enables auto-redirect support if the supplied predicate matches.

The following example uses [followRedirect\(true\)](#):

```
../../../../reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/http/client/redirect/Application.java
```

```
import reactor.netty.http.client.HttpClient;
import reactor.netty.http.client.HttpClientResponse;

public class Application {

    public static void main(String[] args) {
        HttpClient client =
            HttpClient.create()
                .followRedirect(true);

        client.get()
            .uri("http://example.com/")
            .response()
            .block();
    }
}
```

## 6.3. Consuming Data

To receive data from a given `HTTP` endpoint, you can use one of the methods from `HttpClient.ResponseReceiver`. The following example uses the `responseContent` method:

`././reactor-netty-examples/src/main/java/reactor/netty/examples/documentation/http/client/read/Application.java`

```
import reactor.netty.http.client.HttpClient;

public class Application {

    public static void main(String[] args) {
        HttpClient client = HttpClient.create();

        client.get()
            .uri("http://example.com/")
            .responseContent() ①
            .aggregate()        ②
            .asString()         ③
            .block();
    }
}
```

① Receives data from a given [HTTP](#) endpoint

② Aggregates the data

③ Transforms the data as string

### 6.3.1. Reading Headers and Other Metadata

When receiving data from a given [HTTP](#) endpoint, you can check response headers, status code, and other metadata. You can obtain this additional metadata by using [HttpClientResponse](#). The following example shows how to do so.

```
../../../../reactor-netty-
```

```
examples/src/main/java/reactor/netty/examples/documentation/http/client/read/status/Application.jav  
a
```

```
import reactor.netty.http.client.HttpClient;  
  
public class Application {  
  
    public static void main(String[] args) {  
        HttpClient client = HttpClient.create();  
  
        client.get()  
            .uri("http://example.com/")  
            .responseSingle((resp, bytes) -> {  
                System.out.println(resp.status()); ①  
                return bytes.asString();  
            })  
            .block();  
    }  
}
```

① Obtains the status code.

### 6.3.2. HTTP Response Decoder

By default, **Netty** configures some restrictions for the incoming responses, such as:

- The maximum length of the initial line.
- The maximum length of all headers.
- The maximum length of the content or each chunk.

For more information, see [HttpResponseDecoder](#)

By default, the **HTTP** client is configured with the following settings:

```
../../../../reactor-netty-http/src/main/java/reactor/netty/http/HttpDecoderSpec.java
```

```
public static final int DEFAULT_MAX_INITIAL_LINE_LENGTH = 4096;  
public static final int DEFAULT_MAX_HEADER_SIZE          = 8192;  
public static final int DEFAULT_MAX_CHUNK_SIZE          = 8192;  
public static final boolean DEFAULT_VALIDATE_HEADERS     = true;  
public static final int DEFAULT_INITIAL_BUFFER_SIZE     = 128;
```

```
../../../../reactor-netty-http/src/main/java/reactor/netty/http/client/HttpResponseDecoderSpec.java
```

```
public static final boolean DEFAULT_FAIL_ON_MISSING_RESPONSE = false;
public static final boolean DEFAULT_PARSE_HTTP_AFTER_CONNECT_REQUEST = false;

/***
 * The maximum length of the content of the HTTP/2.0 clear-text upgrade request.
 * By default the client will allow an upgrade request with up to 65536 as
 * the maximum length of the aggregated content.
 */
public static final int DEFAULT_H2C_MAX_CONTENT_LENGTH = 65536;
```

When you need to change these default settings, you can configure the [HTTP](#) client as follows:

```
../../../../reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/http/client/respondedecoder/Application.java
```

```
import reactor.netty.http.client.HttpClient;

public class Application {

    public static void main(String[] args) {
        HttpClient client =
            HttpClient.create()
                .httpResponseDecoder(spec -> spec.maxHeaderSize(16384));
①

        client.get()
            .uri("http://example.com/")
            .responseContent()
            .aggregate()
            .asString()
            .block();
    }
}
```

① The maximum length of all headers will be [16384](#). When this value is exceeded, a [TooLongFrameException](#) is raised.

## 6.4. TCP-level Configuration

When you need configurations on a TCP level, you can use the following snippet to extend the default [TCP](#) client configuration (add an option, bind address etc.):

```
../../../../reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/http/client/channeloptions/Application
.java
```

```
import io.netty.channel.ChannelOption;
import reactor.netty.http.client.HttpClient;
import java.net.InetSocketAddress;

public class Application {

    public static void main(String[] args) {
        HttpClient client =
            HttpClient.create()
                .bindAddress(() -> new InetSocketAddress("host", 1234))
                .option(ChannelOption.CONNECT_TIMEOUT_MILLIS, 10000);

        String response =
            client.get()
                .uri("http://example.com/")
                .responseContent()
                .aggregate()
                .asString()
                .block();

        System.out.println("Response " + response);
    }
}
```

See [TCP Client](#) for more about **TCP** level configurations.

#### 6.4.1. Wire Logger

Reactor Netty provides wire logging for when the traffic between the peers needs to be inspected. By default, wire logging is disabled. To enable it, you must set the logger `reactor.netty.http.client.HttpClient` level to `DEBUG` and apply the following configuration:

```
././reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/http/client/wiretap/Application.java
```

```
import reactor.netty.http.client.HttpClient;
import reactor.netty.http.client.HttpClientResponse;

public class Application {

    public static void main(String[] args) {
        HttpClient client =
            HttpClient.create()
                .wiretap(true); ①

        client.get()
            .uri("http://example.com/")
            .response()
            .block();
    }
}
```

① Enables the wire logging

## 6.5. SSL and TLS

When you need SSL or TLS, you can apply the configuration shown in the next example. By default, if `OpenSSL` is available, a `SslProvider.OPENSSL` provider is used as a provider. Otherwise a `SslProvider.JDK` provider is used. You can switch the provider by using `SslContextBuilder` or by setting `-Dio.netty.handler.ssl.noOpenSsl=true`. The following example uses `SslContextBuilder`:

```
././reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/http/client/security/Application.java
```

```
import io.netty.handler.ssl.SslContextBuilder;
import reactor.netty.http.client.HttpClient;
import reactor.netty.http.client.HttpClientResponse;

public class Application {

    public static void main(String[] args) {
        SslContextBuilder sslContextBuilder = SslContextBuilder.forClient();

        HttpClient client =
            HttpClient.create()
                .secure(spec -> spec.sslContext(sslContextBuilder));

        client.get()
            .uri("https://example.com/")
            .response()
            .block();
    }
}
```

## 6.6. Retry Strategies

By default, the [HTTP](#) client retries the request once if it was aborted on the [TCP](#) level.

## 6.7. Metrics

The HTTP client supports built-in integration with [Micrometer](#). It exposes all metrics with a prefix of [reactor.netty.http.client](#).

The following table provides information for the HTTP client metrics:

metric name	type	description
reactor.netty.http.client.data.received	DistributionSummary	Amount of the data received, in bytes
reactor.netty.http.client.data.sent	DistributionSummary	Amount of the data sent, in bytes
reactor.netty.http.client.errors	Counter	Number of errors that occurred
reactor.netty.http.client.tls.handshake.time	Timer	Time spent for TLS handshake
reactor.netty.http.client.connect.time	Timer	Time spent for connecting to the remote address

<b>metric name</b>	<b>type</b>	<b>description</b>
reactor.netty.http.client.address.resolver	Timer	Time spent for resolving the address
reactor.netty.http.client.data.received.time	Timer	Time spent in consuming incoming data
reactor.netty.http.client.data.sent.time	Timer	Time spent in sending outgoing data
reactor.netty.http.client.response.time	Timer	Total time for the request/response

These additional metrics are also available:

#### Pooled [ConnectionProvider](#) metrics

<b>metric name</b>	<b>type</b>	<b>description</b>
reactor.netty.connection.provider.total.connections	Gauge	The number of all connections, active or idle
reactor.netty.connection.provider.active.connections	Gauge	The number of the connections that have been successfully acquired and are in active use
reactor.netty.connection.provider.idle.connections	Gauge	The number of the idle connections
reactor.netty.connection.provider.pending.connections	Gauge	The number of requests that are waiting for a connection

#### ByteBufAllocator metrics

<b>metric name</b>	<b>type</b>	<b>description</b>
reactor.netty.bytebuf_allocator.used.heap.memory	Gauge	The number of the bytes of the heap memory
reactor.netty.bytebuf_allocator.used.direct.memory	Gauge	The number of the bytes of the direct memory
reactor.netty.bytebuf_allocator.used.heap.arenas	Gauge	The number of heap arenas (when <a href="#">PooledByteBufAllocator</a> )
reactor.netty.bytebuf_allocator.used.direct.arenas	Gauge	The number of direct arenas (when <a href="#">PooledByteBufAllocator</a> )
reactor.netty.bytebuf_allocator.used.threadlocal.caches	Gauge	The number of thread local caches (when <a href="#">PooledByteBufAllocator</a> )
reactor.netty.bytebuf_allocator.used.tiny.cache.size	Gauge	The size of the tiny cache (when <a href="#">PooledByteBufAllocator</a> )
reactor.netty.bytebuf_allocator.used.small.cache.size	Gauge	The size of the small cache (when <a href="#">PooledByteBufAllocator</a> )
reactor.netty.bytebuf_allocator.used.normal.cache.size	Gauge	The size of the normal cache (when <a href="#">PooledByteBufAllocator</a> )

metric name	type	description
reactor.netty.bytebufallocator.used.chunk.size	Gauge	The chunk size for an arena (when <code>PooledByteBufAllocator</code> )

The following example enables that integration:

././reactor-netty-examples/src/main/java/reactor/netty/examples/documentation/http/client/metrics/Application.java

```
import io.micrometer.core.instrument.Metrics;
import io.micrometer.core.instrument.config.MeterFilter;
import reactor.netty.http.client.HttpClient;

public class Application {

    public static void main(String[] args) {
        Metrics.globalRegistry ①
            .config()

        .meterFilter(MeterFilter.maximumAllowableTags("reactor.netty.http.client", "URI",
100, MeterFilter.deny()));

        HttpClient client =
            HttpClient.create()
                .metrics(true, s -> {
                    if (s.startsWith("/stream/")) { ②
                        return "/stream/{n}";
                    }
                    else if (s.startsWith("/bytes/")) {
                        return "/bytes/{n}";
                    }
                    return s;
                }); ③

        client.get()
            .uri("http://httpbin.org/stream/2")
            .responseContent()
            .blockLast();

        client.get()
            .uri("http://httpbin.org/bytes/1024")
            .responseContent()
            .blockLast();
    }
}
```

① Applies upper limit for the meters with **URI** tag

② Templatized URIs will be used as an URI tag value when possible

③ Enables the built-in integration with Micrometer

**i** In order to avoid a memory and CPU overhead of the enabled metrics, it is important to convert the real URIs to templated URIs when possible. Without a conversion to a template-like form, each distinct URI leads to the creation of a distinct tag, which takes a lot of memory for the metrics.

**i** Always apply an upper limit for the meters with URI tags. Configuring an upper limit on the number of meters can help in cases when the real URIs cannot be templated. You can find more information at [maximumAllowableTags](#).

When HTTP client metrics are needed for an integration with a system other than [Micrometer](#) or you want to provide your own integration with [Micrometer](#), you can provide your own metrics recorder, as follows:

```
././reactor-netty-examples/src/main/java/reactor/netty/examples/documentation/http/client/metrics/custom/Application.java
```

```
import reactor.netty.channel.ChannelMetricsRecorder;
import reactor.netty.http.client.HttpClient;
import reactor.netty.http.client.HttpClientResponse;

import java.net.SocketAddress;
import java.time.Duration;

public class Application {

    public static void main(String[] args) {
        HttpClient client =
            HttpClient.create()
                .metrics(true, CustomHttpClientMetricsRecorder::new); ①

        client.get()
            .uri("https://httpbin.org/stream/2")
            .response()
            .block();
    }
}
```

① Enables HTTP client metrics and provides [HttpClientMetricsRecorder](#) implementation.

# Chapter 7. UDP Server

Reactor Netty provides the easy-to-use and easy-to-configure [UdpServer](#). It hides most of the Netty functionality that is required to create a [UDP](#) server and adds [Reactive Streams](#) backpressure.

## 7.1. Starting and Stopping

To start a UDP server, a [UdpServer](#) instance has to be created and configured. By default, the host is configured to be [localhost](#) and the port is [12012](#). The following example shows how to create and start a UDP server:

```
./../reactor-netty-  
examples/src/main/java/reactor/netty/examples/documentation/udp/server/create/Application.java
```

```
import reactor.netty.Connection;  
import reactor.netty.UdpServer;  
import java.time.Duration;  
  
public class Application {  
  
    public static void main(String[] args) {  
        Connection server =  
            UdpServer.create()  
                .bindNow(Duration.ofSeconds(30)); ①  
                .bindNow(Duration.ofSeconds(30)); ②  
  
        server.onDispose()  
            .block();  
    }  
}
```

① Creates a [UdpServer](#) instance that is ready for configuring.

② Starts the server in a blocking fashion and waits for it to finish initializing.

The returned [Connection](#) offers a simple server API, including [disposeNow\(\)](#), which shuts the server down in a blocking fashion.

### 7.1.1. Host and Port

In order to serve on a specific host and port, you can apply the following configuration to the [UDP](#) server:

`././reactor-netty-examples/src/main/java/reactor/netty/examples/documentation/udp/server/address/Application.java`

```
import reactor.netty.Connection;
import reactor.netty.udp.UdpServer;
import java.time.Duration;

public class Application {

    public static void main(String[] args) {
        Connection server =
            UdpServer.create()
                .host("localhost") ①
                .port(8080)         ②
                .bindNow(Duration.ofSeconds(30));

        server.onDispose()
            .block();
    }
}
```

① Configures the **UDP** server host

② Configures the **UDP** server port

## 7.2. Writing Data

To send data to the remote peer, you must attach an I/O handler. The I/O handler has access to **UdpOutbound**, to be able to write data. The following example shows how to send **hello**:

```
././reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/udp/server/send/Application.java
```

```
import io.netty.buffer.ByteBuf;
import io.netty.buffer.Unpooled;
import io.netty.channel.socket.DatagramPacket;
import io.netty.util.CharsetUtil;
import reactor.core.publisher.Mono;
import reactor.netty.Connection;
import reactor.netty.udp.UdpServer;

import java.time.Duration;

public class Application {

    public static void main(String[] args) {
        Connection server =
            UdpServer.create()
                .handle((in, out) ->
                    out.sendObject(
                        in.receiveObject()
                            .map(o -> {
                                if (o instanceof DatagramPacket) {
                                    DatagramPacket p = (DatagramPacket) o;
                                    ByteBuf buf =
                                        Unpooled.copiedBuffer("hello", CharsetUtil.UTF_8);
                                    return new DatagramPacket(buf,
                                        p.sender()); ①
                                }
                            })
                            .else {
                                return Mono.error(new
                                    Exception("Unexpected type of the message: " + o));
                            }
                        })))
                .bindNow(Duration.ofSeconds(30));

        server.onDispose()
            .block();
    }
}
```

① Sends a `hello` string to the remote peer

## 7.3. Consuming Data

To receive data from a remote peer, you must attach an I/O handler. The I/O handler has access to `UdpInbound`, to be able to read data. The following example shows how to consume data:

```
./../reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/udp/server/read/Application.java
```

```
import io.netty.channel.socket.DatagramPacket;
import reactor.core.publisher.Mono;
import reactor.netty.Connection;
import reactor.netty.udp.UdpServer;

import java.time.Duration;

public class Application {

    public static void main(String[] args) {
        Connection server =
            UdpServer.create()
                .handle((in, out) ->
                    out.sendObject(
                        in.receiveObject()
                            .map(o -> {
                                if (o instanceof DatagramPacket) {
                                    DatagramPacket p = (DatagramPacket) o;
                                    return new
                                        DatagramPacket(p.content().retain(), p.sender()); ①
                                }
                                else {
                                    return Mono.error(new
                                        Exception("Unexpected type of the message: " + o));
                                }
                            })))
                .bindNow(Duration.ofSeconds(30));

        server.onDispose()
            .block();
    }
}
```

① Receives data from the remote peer

## 7.4. Lifecycle Callbacks

The following lifecycle callbacks are provided to let you extend the **UDP** server:

- **doOnBind**: Invoked when the server channel is about to bind.
- **doOnBound**: Invoked when the server channel is bound.
- **doOnUnbound**: Invoked when the server channel is unbound.

The following example uses the **doOnBound** method:

```
././reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/udp/server/lifecycle/Application.java
```

```
import io.netty.handler.codec.LineBasedFrameDecoder;
import reactor.netty.Connection;
import reactor.netty.udp.UdpServer;
import java.time.Duration;

public class Application {

    public static void main(String[] args) {
        Connection server =
            UdpServer.create()
                .doOnBound(conn -> conn.addHandler(new
LineBasedFrameDecoder(8192))) ①
                .bindNow(Duration.ofSeconds(30));

        server.onDispose()
            .block();
    }
}
```

① Netty pipeline is extended with `LineBasedFrameDecoder` when the server channel is bound.

## 7.5. Connection Configuration

This section describes three kinds of configuration that you can use at the UDP level:

- [Channel Options](#)
- [Wire Logger](#)
- [Event Loop Group](#)

### 7.5.1. Channel Options

By default, the `UDP` server is configured with the following options:

```
././reactor-netty-core/src/main/java/reactor/netty/udp/UdpServerBind.java
```

```
UdpServerBind() {
    this.config = new UdpServerConfig(
        Collections.singletonMap(ChannelOption.AUTO_READ, false),
        () -> new InetSocketAddress(NetUtil.LOCALHOST, DEFAULT_PORT));
}
```

If you need additional options or need to change the current options, you can apply the following

configuration:

```
././reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/udp/server/channeloptions/Application.java
```

```
import io.netty.channel.ChannelOption;
import reactor.netty.Connection;
import reactor.netty.udp.UdpServer;
import java.time.Duration;

public class Application {

    public static void main(String[] args) {
        Connection server =
            UdpServer.create()
                .option(ChannelOption.CONNECT_TIMEOUT_MILLIS, 10000)
                .bindNow(Duration.ofSeconds(30));

        server.onDispose()
            .block();
    }
}
```

For more information about Netty channel options, see the following links:

- [ChannelOption](#)
- [Socket Options](#)

### 7.5.2. Wire Logger

Reactor Netty provides wire logging for when the traffic between the peers has to be inspected. By default, wire logging is disabled. To enable it, you must set the logger `reactor.netty.udp.UdpServer` level to `DEBUG` and apply the following configuration:

`././reactor-netty-examples/src/main/java/reactor/netty/examples/documentation/udp/server/wiretap/Application.java`

```
import reactor.netty.Connection;
import reactor.netty.udp.UdpServer;
import java.time.Duration;

public class Application {

    public static void main(String[] args) {
        Connection server =
            UdpServer.create()
                .wiretap(true) ①
                .bindNow(Duration.ofSeconds(30));

        server.onDispose()
            .block();
    }
}
```

① Enables the wire logging

### 7.5.3. Event Loop Group

By default, the UDP server uses “Event Loop Group,” where the number of the worker threads equals the number of processors available to the runtime on initialization (but with a minimum value of 4). When you need a different configuration, you can use one of the [LoopResource#create](#) methods.

The default configuration for the “Event Loop Group” is the following:

`../../../../reactor-netty-core/src/main/java/reactor/netty/ReactorNetty.java`

```
* Default worker thread count, fallback to available processor
* (but with a minimum value of 4)
*/
public static final String IO_WORKER_COUNT = "reactor.netty.ioWorkerCount";
/**
 * Default selector thread count, fallback to -1 (no selector thread)
 */
public static final String IO_SELECT_COUNT = "reactor.netty.ioSelectCount";
/**
 * Default worker thread count for UDP, fallback to available processor
 * (but with a minimum value of 4)
*/
public static final String UDP_IO_THREAD_COUNT =
"reactor.netty.udp.ioThreadCount";
/**
 * Default quiet period that guarantees that the disposal of the underlying
LoopResources
 * will not happen, fallback to 2 seconds.
*/
public static final String SHUTDOWN QUIET_PERIOD =
"reactor.netty.ioShutdownQuietPeriod";
/**
 * Default maximum amount of time to wait until the disposal of the underlying
LoopResources
 * regardless if a task was submitted during the quiet period, fallback to 15
seconds.
*/
public static final String SHUTDOWN_TIMEOUT = "reactor.netty.ioShutdownTimeout";

/**
 * Default value whether the native transport (epoll, kqueue) will be preferred,
 * fallback it will be preferred when available
*/
public static final String NATIVE = "reactor.netty.native";
```

If you need changes to these settings, you can apply the following configuration:

```
././reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/udp/server/eventloop/Application.java
```

```
import reactor.netty.Connection;
import reactor.netty.resources.LoopResources;
import reactor.netty.udp.UdpServer;
import java.time.Duration;

public class Application {

    public static void main(String[] args) {
        LoopResources loop = LoopResources.create("event-loop", 1, 4, true);

        Connection server =
            UdpServer.create()
                .runOn(loop)
                .bindNow(Duration.ofSeconds(30));

        server.onDispose()
            .block();
    }
}
```

## 7.6. Metrics

The UDP server supports built-in integration with [Micrometer](#). It exposes all metrics with a prefix of `reactor.netty.udp.server`.

The following table provides information for the UDP server metrics:

metric name	type	description
reactor.netty.udp.server.data.received	DistributionSummary	Amount of the data received, in bytes
reactor.netty.udp.server.data.sent	DistributionSummary	Amount of the data sent, in bytes
reactor.netty.udp.server.errors	Counter	Number of errors that occurred

These additional metrics are also available:

[ByteBufAllocator](#) metrics

metric name	type	description
reactor.netty.bytebufallocator.used.heap.memory	Gauge	The number of the bytes of the heap memory
reactor.netty.bytebufallocator.used.direct.memory	Gauge	The number of the bytes of the direct memory

metric name	type	description
reactor.netty.bytebufallocator.used.heap.arenas	Gauge	The number of heap arenas (when <a href="#">PooledByteBufAllocator</a> )
reactor.netty.bytebufallocator.used.direct.arenas	Gauge	The number of direct arenas (when <a href="#">PooledByteBufAllocator</a> )
reactor.netty.bytebufallocator.used.threadlocal.caches	Gauge	The number of thread local caches (when <a href="#">PooledByteBufAllocator</a> )
reactor.netty.bytebufallocator.used.tiny.cache.size	Gauge	The size of the tiny cache (when <a href="#">PooledByteBufAllocator</a> )
reactor.netty.bytebufallocator.used.small.cache.size	Gauge	The size of the small cache (when <a href="#">PooledByteBufAllocator</a> )
reactor.netty.bytebufallocator.used.normal.cache.size	Gauge	The size of the normal cache (when <a href="#">PooledByteBufAllocator</a> )
reactor.netty.bytebufallocator.used.chunk.size	Gauge	The chunk size for an arena (when <a href="#">PooledByteBufAllocator</a> )

The following example enables that integration:

```
./././reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/udp/server/metrics/Application.java
```

```
import reactor.netty.Connection;
import reactor.netty.udp.UdpServer;

import java.time.Duration;

public class Application {

    public static void main(String[] args) {
        Connection server =
            UdpServer.create()
                .metrics(true) ①
                .bindNow(Duration.ofSeconds(30));

        server.onDispose()
            .block();
    }
}
```

① Enables the built-in integration with Micrometer

When UDP server metrics are needed for an integration with a system other than [Micrometer](#) or you want to provide your own integration with [Micrometer](#), you can provide your own metrics recorder, as follows:

`././reactor-netty-examples/src/main/java/reactor/netty/examples/documentation/udp/server/metrics/custom/Application.java`

```
import reactor.netty.Connection;
import reactor.netty.channel.ChannelMetricsRecorder;
import reactor.netty.udp.UdpServer;

import java.net.SocketAddress;
import java.time.Duration;

public class Application {

    public static void main(String[] args) {
        Connection server =
            UdpServer.create()
                .metrics(true, CustomChannelMetricsRecorder::new) ①
                .bindNow(Duration.ofSeconds(30));

        server.onDispose()
            .block();
    }
}
```

① Enables UDP server metrics and provides `ChannelMetricsRecorder` implementation.

# Chapter 8. UDP Client

Reactor Netty provides the easy-to-use and easy-to-configure [UdpClient](#). It hides most of the Netty functionality that is required to create a [UDP](#) client and adds Reactive Streams backpressure.

## 8.1. Connecting and Disconnecting

To connect the UDP client to a given endpoint, you must create and configure a [UdpClient](#) instance. By default, the host is configured for [localhost](#) and the port is [12012](#). The following example shows how to create and connect a UDP client:

```
./../reactor-netty-  
examples/src/main/java/reactor/netty/examples/documentation/udp/client/create/Application.java
```

```
import reactor.netty.Connection;  
import reactor.netty.UdpClient;  
import java.time.Duration;  
  
public class Application {  
  
    public static void main(String[] args) {  
        Connection connection =  
            UdpClient.create()  
                .connectNow(Duration.ofSeconds(30)); ①  
                .connectNow(Duration.ofSeconds(30)); ②  
  
        connection.onDispose()  
            .block();  
    }  
}
```

① Creates a [UdpClient](#) instance that is ready for configuring.

② Connects the client in a blocking fashion and waits for it to finish initializing.

The returned [Connection](#) offers a simple connection API, including [disposeNow\(\)](#), which shuts the client down in a blocking fashion.

### 8.1.1. Host and Port

To connect to a specific [host](#) and [port](#), you can apply the following configuration to the [UDP](#) client:

`././reactor-netty-examples/src/main/java/reactor/netty/examples/documentation/udp/client/address/Application.java`

```
import reactor.netty.Connection;
import reactor.netty.udp.UdpClient;
import java.time.Duration;

public class Application {

    public static void main(String[] args) {
        Connection connection =
            UdpClient.create()
                .host("example.com") ①
                .port(80)             ②
                .connectNow(Duration.ofSeconds(30));

        connection.onDispose()
            .block();
    }
}
```

① Configures the `host` to which this client should connect

② Configures the `port` to which this client should connect

## 8.2. Writing Data

To send data to a given peer, you must attach an I/O handler. The I/O handler has access to `UdpOutbound`, to be able to write data.

The following example shows how to send `hello`:

`././reactor-netty-examples/src/main/java/reactor/netty/examples/documentation/udp/client/send/Application.java`

```
import reactor.core.publisher.Mono;
import reactor.netty.Connection;
import reactor.netty.udp.UdpClient;

import java.time.Duration;

public class Application {

    public static void main(String[] args) {
        Connection connection =
            UdpClient.create()
                .host("example.com")
                .port(80)
                .handle((udpInbound, udpOutbound) ->
        udpOutbound.sendString(Mono.just("hello"))) ①
                .connectNow(Duration.ofSeconds(30));

        connection.onDispose()
            .block();
    }
}
```

① Sends `hello` string to the remote peer.

## 8.3. Consuming Data

To receive data from a given peer, you must attach an I/O handler. The I/O handler has access to `UdpInbound`, to be able to read data. The following example shows how to consume data:

```
././reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/udp/client/read/Application.java
```

```
import reactor.netty.Connection;
import reactor.netty.udp.UdpClient;
import java.time.Duration;

public class Application {

    public static void main(String[] args) {
        Connection connection =
            UdpClient.create()
                .host("example.com")
                .port(80)
                .handle((udpInbound, udpOutbound) ->
        udpInbound.receive().then() ①
            .connectNow(Duration.ofSeconds(30));

        connection.onDispose()
            .block();
    }
}
```

① Receives data from a given peer

## 8.4. Lifecycle Callbacks

The following lifecycle callbacks are provided to let you extend the **UDP** client:

- **doOnConnect**: Invoked when the channel is about to connect.
- **doOnConnected**: Invoked after the channel has been connected.
- **doOnDisconnected**: Invoked after the channel has been disconnected.

The following example uses the **doOnConnected** method:

././reactor-netty-examples/src/main/java/reactor/netty/examples/documentation/udp/client/lifecycle/Application.java

```
import io.netty.handler.codec.LineBasedFrameDecoder;
import reactor.netty.Connection;
import reactor.netty.UdpClient;
import java.time.Duration;

public class Application {

    public static void main(String[] args) {
        Connection connection =
            UdpClient.create()
                .host("example.com")
                .port(80)
                .doOnConnected(conn -> conn.addHandler(new
LineBasedFrameDecoder(8192))) ①
                .connectNow(Duration.ofSeconds(30));

        connection.onDispose()
            .block();
    }
}
```

① The Netty pipeline is extended with `LineBasedFrameDecoder` when the channel has been connected.

## 8.5. Connection Configuration

This section describes three kinds of configuration that you can use at the UDP level:

- [Channel Options](#)
- [Wire Logger](#)
- [Event Loop Group](#)

### 8.5.1. Channel Options

By default, the `UDP` client is configured with the following options:

```
../../../../reactor-netty-core/src/main/java/reactor/netty/udp/UdpClientConnect.java
```

```
UdpClientConnect() {  
    this.config = new UdpClientConfig(  
        ConnectionProvider.newConnection(),  
        Collections.singletonMap(ChannelOption.AUTO_READ, false),  
        () -> new InetSocketAddress(NetUtil.LOCALHOST, DEFAULT_PORT));  
}
```

If you need additional options or need to change the current options, you can apply the following configuration:

```
../../../../reactor-netty-  
examples/src/main/java/reactor/netty/examples/documentation/udp/client/channeloptions/Application.  
java
```

```
import io.netty.channel.ChannelOption;  
import reactor.netty.Connection;  
import reactor.netty.udp.UdpClient;  
import java.time.Duration;  
  
public class Application {  
  
    public static void main(String[] args) {  
        Connection connection =  
            UdpClient.create()  
                .host("example.com")  
                .port(80)  
                .option(ChannelOption.CONNECT_TIMEOUT_MILLIS, 10000)  
                .connectNow(Duration.ofSeconds(30));  
  
        connection.onDispose()  
            .block();  
    }  
}
```

You can find more about Netty channel options at the following links:

- [ChannelOption](#)
- [Socket Options](#)

### 8.5.2. Wire Logger

Reactor Netty provides wire logging for when the traffic between the peers has to be inspected. By default, wire logging is disabled. To enable it, you must set the logger `reactor.netty.udp.UdpClient` level to `DEBUG` and apply the following configuration:

././reactor-netty-examples/src/main/java/reactor/netty/examples/documentation/udp/client/wiretap/Application.java

```
import reactor.netty.Connection;
import reactor.netty.udp.UdpClient;
import java.time.Duration;

public class Application {

    public static void main(String[] args) {
        Connection connection =
            UdpClient.create()
                .host("example.com")
                .port(80)
                .wiretap(true) ①
                .connectNow(Duration.ofSeconds(30));

        connection.onDispose()
            .block();
    }
}
```

① Enables the wire logging

### 8.5.3. Event Loop Group

By default, the UDP client uses “Event Loop Group,” where the number of the worker threads equals the number of processors available to the runtime on initialization (but with a minimum value of 4). When you need a different configuration, you can use one of the [LoopResources#create](#) methods.

The following listing shows the default configuration for the “Event Loop Group”:

`../../../../reactor-netty-core/src/main/java/reactor/netty/ReactorNetty.java`

```
* Default worker thread count, fallback to available processor
* (but with a minimum value of 4)
*/
public static final String IO_WORKER_COUNT = "reactor.netty.ioWorkerCount";
/**
 * Default selector thread count, fallback to -1 (no selector thread)
 */
public static final String IO_SELECT_COUNT = "reactor.netty.ioSelectCount";
/**
 * Default worker thread count for UDP, fallback to available processor
 * (but with a minimum value of 4)
*/
public static final String UDP_IO_THREAD_COUNT =
"reactor.netty.udp.ioThreadCount";
/**
 * Default quiet period that guarantees that the disposal of the underlying
LoopResources
 * will not happen, fallback to 2 seconds.
*/
public static final String SHUTDOWN QUIET PERIOD =
"reactor.netty.ioShutdownQuietPeriod";
/**
 * Default maximum amount of time to wait until the disposal of the underlying
LoopResources
 * regardless if a task was submitted during the quiet period, fallback to 15
seconds.
*/
public static final String SHUTDOWN_TIMEOUT = "reactor.netty.ioShutdownTimeout";

/**
 * Default value whether the native transport (epoll, kqueue) will be preferred,
 * fallback it will be preferred when available
*/
public static final String NATIVE = "reactor.netty.native";
```

If you need changes to the these settings, you can apply the following configuration:

```
././reactor-netty-
examples/src/main/java/reactor/netty/examples/documentation/udp/client/eventloop/Application.java
```

```
import reactor.netty.Connection;
import reactor.netty.resources.LoopResources;
import reactor.netty.udp.UdpClient;
import java.time.Duration;

public class Application {

    public static void main(String[] args) {
        LoopResources loop = LoopResources.create("event-loop", 1, 4, true);

        Connection connection =
            UdpClient.create()
                .host("example.com")
                .port(80)
                .runOn(loop)
                .connectNow(Duration.ofSeconds(30));

        connection.onDispose()
            .block();
    }
}
```

## 8.6. Metrics

The UDP client supports built-in integration with [Micrometer](#). It exposes all metrics with a prefix of `reactor.netty.udp.client`.

The following table provides information for the UDP client metrics:

metric name	type	description
reactor.netty.udp.client.data.received	DistributionSummary	Amount of the data received, in bytes
reactor.netty.udp.client.data.sent	DistributionSummary	Amount of the data sent, in bytes
reactor.netty.udp.client.errors	Counter	Number of errors that occurred
reactor.netty.udp.client.connect.time	Timer	Time spent for connecting to the remote address
reactor.netty.udp.client.address.resolver	Timer	Time spent for resolving the address

These additional metrics are also available:

[ByteBufAllocator](#) metrics

<b>metric name</b>	<b>type</b>	<b>description</b>
reactor.netty.bytebufallocator.used.heap.memory	Gauge	The number of the bytes of the heap memory
reactor.netty.bytebufallocator.used.direct.memory	Gauge	The number of the bytes of the direct memory
reactor.netty.bytebufallocator.used.heap.arenas	Gauge	The number of heap arenas (when <a href="#">PooledByteBufAllocator</a> )
reactor.netty.bytebufallocator.used.direct.arenas	Gauge	The number of direct arenas (when <a href="#">PooledByteBufAllocator</a> )
reactor.netty.bytebufallocator.used.threadlocal.caches	Gauge	The number of thread local caches (when <a href="#">PooledByteBufAllocator</a> )
reactor.netty.bytebufallocator.used.tiny.cache.size	Gauge	The size of the tiny cache (when <a href="#">PooledByteBufAllocator</a> )
reactor.netty.bytebufallocator.used.small.cache.size	Gauge	The size of the small cache (when <a href="#">PooledByteBufAllocator</a> )
reactor.netty.bytebufallocator.used.normal.cache.size	Gauge	The size of the normal cache (when <a href="#">PooledByteBufAllocator</a> )
reactor.netty.bytebufallocator.used.chunk.size	Gauge	The chunk size for an arena (when <a href="#">PooledByteBufAllocator</a> )

The following example enables that integration:

././reactor-netty-examples/src/main/java/reactor/netty/examples/documentation/udp/client/metrics/Application.java

```
import reactor.netty.Connection;
import reactor.netty.udp.UdpClient;

import java.time.Duration;

public class Application {

    public static void main(String[] args) {
        Connection connection =
            UdpClient.create()
                .host("example.com")
                .port(80)
                .metrics(true) ①
                .connectNow(Duration.ofSeconds(30));

        connection.onDispose()
            .block();
    }
}
```

① Enables the built-in integration with Micrometer

When UDP client metrics are needed for an integration with a system other than **Micrometer** or you want to provide your own integration with **Micrometer**, you can provide your own metrics recorder, as follows:

`././reactor-netty-examples/src/main/java/reactor/netty/examples/documentation/udp/client/metrics/custom/Application.java`

```
import reactor.netty.Connection;
import reactor.netty.channel.ChannelMetricsRecorder;
import reactor.netty.udp.UdpClient;

import java.net.SocketAddress;
import java.time.Duration;

public class Application {

    public static void main(String[] args) {
        Connection connection =
            UdpClient.create()
                .host("example.com")
                .port(80)
                .metrics(true, CustomChannelMetricsRecorder::new) ①
                .connectNow(Duration.ofSeconds(30));

        connection.onDispose()
            .block();
    }
}
```

① Enables UDP client metrics and provides `ChannelMetricsRecorder` implementation.