
tapir documentation

Release 0.x

Adam Warski

Mar 08, 2021

1	Code teaser	3
2	Other sttp projects	5
3	Sponsors	7
4	Table of contents	9
4.1	Quickstart	9
4.2	Examples	9
4.3	Goals of the project	11
4.4	Basics	11
4.5	Inputs/outputs	12
4.6	Status codes	16
4.7	Codecs	19
4.8	Custom types	21
4.9	Validation	25
4.10	Working with JSON	27
4.11	Forms	31
4.12	Authentication	33
4.13	Streaming support	33
4.14	Web sockets	34
4.15	Datatypes integrations	36
4.16	ZIO integration	38
4.17	Running as an akka-http server	39
4.18	Running as an http4s server	42
4.19	Running as a Finatra server	44
4.20	Running as a Play server	46
4.21	Running as a Vert.X server	48
4.22	Server options	52
4.23	Server logic	52
4.24	Error handling	56
4.25	Debugging	58
4.26	Using as an sttp client	58
4.27	Using as a Play client	59
4.28	Generating OpenAPI documentation	61
4.29	Generating AsyncAPI documentation	63
4.30	Testing	64

4.31	Generate endpoint definitions from an OpenAPI YAML	65
4.32	Other interpreters	66
4.33	Creating your own Tapir	67
4.34	Troubleshooting	67
4.35	Contributing	67

With tapir, you can describe HTTP API endpoints as immutable Scala values. Each endpoint can contain a number of input parameters, error-output parameters, and normal-output parameters. An endpoint specification can be interpreted as:

- a server, given the “business logic”: a function, which computes output parameters based on input parameters. Currently supported:
 - *Akka HTTP* Routes/Directives.
 - *Http4s* `HttpRoutes[F]`
 - *Finatra* `http.Controller`
 - *Play* `Route`
- a client, which is a function from input parameters to output parameters. Currently supported:
 - *sttp*.
 - *Play*.
- documentation. Currently supported:
 - *OpenAPI*.
 - *AsyncAPI*.

Tapir is licensed under Apache2, the source code is [available on GitHub](#).

Depending on how you prefer to explore the library, take a look at one of the [examples](#) or read on for a more detailed description of how tapir works!

Tapir is available for Scala 2.12 and 2.13 on the JVM. The client interpreter is also available for Scala.JS.

CHAPTER 1

Code teaser

```
import sttp.tapir._
import sttp.tapir.generic.auto._
import sttp.tapir.json.circe._
import io.circe.generic.auto._

type Limit = Int
type AuthToken = String
case class BooksFromYear(genre: String, year: Int)
case class Book(title: String)

// Define an endpoint

val booksListing: Endpoint[(BooksFromYear, Limit, AuthToken), String, List[Book], _
  ↳ Any] =
  endpoint
    .get
    .in(("books" / path[String]("genre") / path[Int]("year")).mapTo(BooksFromYear))
    .in(query[Limit]("limit").description("Maximum number of books to retrieve"))
    .in(header[AuthToken]("X-Auth-Token"))
    .errorOut(stringBody)
    .out(jsonBody[List[Book]])

// Generate OpenAPI documentation

import sttp.tapir.docs.openapi.OpenAPIDocsInterpreter
import sttp.tapir.openapi.circe.yaml._

val docs = OpenAPIDocsInterpreter.toOpenAPI(booksListing, "My Bookshop", "1.0")
println(docs.toYaml)

// Convert to akka-http Route
```

(continues on next page)

(continued from previous page)

```
import sttp.tapir.server.akkahttp.AkkaHttpServerInterpreter
import akka.http.scaladsl.server.Route
import scala.concurrent.Future

def bookListingLogic(bfy: BooksFromYear,
                    limit: Limit,
                    at: AuthToken): Future[Either[String, List[Book]]] =
  Future.successful(Right(List(Book("The Sorrows of Young Werther"))))
val booksListingRoute: Route = AkkaHttpServerInterpreter
  .toRoute(booksListing)((bookListingLogic _).tupled)

// Convert to sttp Request

import sttp.tapir.client.sttp.SttpClientInterpreter
import sttp.client3._

val booksListingRequest: Request[DecodeResult[Either[String, List[Book]]], Any] =
  ↪SttpClientInterpreter
    .toRequest(booksListing, Some(uri"http://localhost:8080"))
    .apply((BooksFromYear("SF", 2016), 20, "xyz-abc-123"))
```

CHAPTER 2

Other sttp projects

sttp is a family of Scala HTTP-related projects, and currently includes:

- [sttp client](#): the Scala HTTP client you always wanted!
- sttp tapir: this project
- [sttp model](#): simple HTTP model classes (used by client & tapir)

CHAPTER 3

Sponsors

Development and maintenance of sttp tapir is sponsored by [SoftwareMill](#), a software development and consulting company. We help clients scale their business through software. Our areas of expertise include backends, distributed systems, blockchain, machine learning and data analytics.



4.1 Quickstart

To use tapir, add the following dependency to your project:

```
"com.softwaremill.sttp.tapir" %% "tapir-core" % "0.17.15"
```

This will import only the core classes needed to create endpoint descriptions. To generate a server or a client, you will need to add further dependencies.

Many of tapir functionalities come as builder methods in the main package, hence it's easiest to work with tapir if you import the main package entirely, i.e.:

```
import sttp.tapir._
```

If you don't have it already, you'll also need partial unification enabled in the compiler (alternatively, you'll need to manually provide type arguments in some cases). In sbt, this is:

```
scalacOptions += "-Ypartial-unification"
```

Finally, type:

```
endpoint.
```

and see where auto-complete gets you!

4.2 Examples

The `examples` sub-project contains a number of runnable tapir usage examples:

- Hello world server, using akka-http
- Hello world server, using http4s

- Separate error & success outputs, using akka-http
- Multiple endpoints, exposing OpenAPI/Swagger documentation, using akka-http
- Multiple endpoints, exposing OpenAPI/Swagger documentation, using http4s
- Multiple endpoints, with the description coupled with server logic, using akka-http
- Reporting errors in a custom format when a query/path/.. parameter cannot be decoded
- Using custom types in endpoint descriptions
- Multipart form upload, using akka-http
- Basic Authentication, using akka-http
- Partial server logic (authentication), using akka-http
- Authorization using Github (OAuth2) with http4s
- Books example
- ZIO example, using http4s
- ZIO example with environment, using http4s
- ZIO partial server logic example, using http4s
- Streaming body, using akka-http
- Streaming body, using http4s + fs2
- Web sockets server, using akka-http
- Web sockets client, using akka-http
- Web sockets server, using http4s

4.2.1 Other examples

To see an example project using Tapir, check out this [Todo-Backend](#) using tapir and http4s.

For a quick start using http4s and tapir you can use a [gitter8 template](#). A new project can be created using: `sbt new https://codeberg.org/wegtam/http4s-tapir.g8.git`

4.2.2 Blogs, articles

- Designing tapir's WebSockets support
- Three easy endpoints
- tAPIr's Endpoint meets ZIO's IO
- Describe, then interpret: HTTP endpoints using tapir
- Functional pancakes

4.2.3 Videos

- ScalaWorld 2019: Designing Programmer-Friendly APIs

4.3 Goals of the project

- programmer-friendly, human-comprehensible types, that you are not afraid to write down
- (also inferencable by IntelliJ)
- discoverable API through standard auto-complete
- separate “business logic” from endpoint definition & documentation
- as simple as possible to generate a server, client & docs
- based purely on case class-based, immutable and reusable data structures
- first-class OpenAPI support. Provide as much or as little detail as needed.
- reasonably type safe: only, and as much types to safely generate the server/client/docs

4.3.1 Similar projects

There’s a number of similar projects from which tapir draws inspiration:

- [endpoints](#)
- [typedapi](#)
- [rho](#)
- [typed-schema](#)
- [guardrail](#)

4.4 Basics

An endpoint is represented as a value of type `Endpoint[I, E, O, S]`, where:

- `I` is the type of the input parameters
- `E` is the type of the error-output parameters
- `O` is the type of the output parameters
- `S` is the type of streams that are used by the endpoint’s inputs/outputs

Input/output parameters (`I`, `E` and `O`) can be:

- of type `Unit`, when there’s no input/output of the given type
- a single type
- a tuple of types

Hence, an empty, initial endpoint (`tapir.endpoint`), with no inputs and no outputs, from which all other endpoints are derived has the type:

```
import sttp.tapir._
val endpoint: Endpoint[Unit, Unit, Unit, Any] = ???
```

An endpoint which accepts two parameters of types `UUID` and `Int`, upon error returns a `String`, and on normal completion returns a `User`, would have the type:

```
import sttp.tapir._

val userEndpoint: Endpoint[(UUID, Int), String, User, Any] = ???
```

You can think of an endpoint as a function, which takes input parameters of type `I` and returns a result of type `Either[E, O]`, where inputs or outputs can contain streaming bodies of type `S`.

4.4.1 Infallible endpoints

Note that the empty `endpoint` description maps no values to either error and success outputs, however errors are still represented and allowed to occur. If you would prefer to use an endpoint description, where errors can not happen, use `infallibleEndpoint: Endpoint[Unit, Nothing, Unit, Nothing]`. This might be useful when interpreting endpoints *as a client*.

4.4.2 Defining an endpoint

The description of an endpoint is an immutable case class, which includes a number of methods:

- the `name`, `description`, etc. methods allow modifying the endpoint information, which will then be included in the endpoint documentation
- the `get`, `post` etc. methods specify the HTTP method which the endpoint should support
- the `in`, `errorOut` and `out` methods allow adding a new input/output parameter
- `mapIn`, `mapInTo`, ... methods allow mapping the current input/output parameters to another value or to a case class

An important note on mapping: in tapir, all mappings are bi-directional. That's because each mapping can be used to generate a server or a client, as well as in many cases can be used both for input and for output.

4.4.3 Next

Read on about describing [endpoint inputs/outputs](#).

4.5 Inputs/outputs

An input is described by an instance of the `EndpointInput` trait, and an output by an instance of the `EndpointOutput` trait. Some inputs can be used both as inputs and outputs; then, they additionally implement the `EndpointIO` trait.

Each input or output can yield/accept a value (but doesn't have to).

For example, `query[Int]("age") : EndpointInput[Int]` describes an input, which is the `age` parameter from the URI's query, and which should be coded (using the string-to-integer [codec](#)) as an `Int`.

The `tapir` package contains a number of convenience methods to define an input or an output for an endpoint. For inputs, these are:

- `path[T]`, which captures a path segment as an input parameter of type `T`
- any string, which will be implicitly converted to a fixed path segment. Path segments can be combined with the `/` method, and don't map to any values (have type `EndpointInput[Unit]`)
- `paths`, which maps to the whole remaining path as a `List[String]`

- `query[T] (name)` captures a query parameter with the given name
- `queryParams` captures all query parameters, represented as `QueryParams`
- `cookie[T] (name)` captures a cookie from the `Cookie` header with the given name
- `extractFromRequest` extracts a value from the request. This input is only used by server interpreters, ignored by documentation interpreters. Client interpreters ignore the provided value. It can also be used to access the original request through the `underlying: Any` field.

For both inputs/outputs:

- `header[T] (name)` captures a header with the given name
- `headers` captures all headers, represented as `List[Header]`
- `cookies` captures cookies from the `Cookie` header and represents them as `List[Cookie]`
- `setCookie (name)` captures the value & metadata of the a `Set-Cookie` header with a matching name
- `setCookies` captures cookies from the `Set-Cookie` header and represents them as `List[SetCookie]`
- `stringBody`, `plainBody[T]`, `jsonBody[T]`, `rawBinaryBody[R]`, `binaryBody[R, T]`, `formBody[T]`, `multipartBody[T]` captures the body
- `streamBody[S]` captures the body as a stream: only a client/server interpreter supporting streams of type `S` can be used with such an endpoint

For outputs:

- `statusCode` maps to the status code of the response
- `statusCode (code)` maps to a fixed status code of the response

4.5.1 Combining inputs and outputs

Endpoint inputs/outputs can be combined in two ways. However they are combined, the values they represent always accumulate into tuples of values.

First, inputs/outputs can be combined using the `.and` method. Such a combination results in an input/output, which maps to a tuple of the given types. This combination can be assigned to a value and re-used in multiple endpoints. As all other values in tapir, endpoint input/output descriptions are immutable. For example, an input specifying two query parameters, `start` (mandatory) and `limit` (optional) can be written down as:

```
import sttp.tapir._
import sttp.tapir.generic.auto._
import sttp.tapir.json.circe._
import io.circe.generic.auto._
import java.util.UUID

case class User(name: String)

val paging: EndpointInput[(UUID, Option[Int])] =
  query[UUID]("start").and(query[Option[Int]]("limit"))

// we can now use the value in multiple endpoints, e.g.:
val listUsersEndpoint: Endpoint[(UUID, Option[Int]), Unit, List[User], Any] =
  endpoint.in("user" / "list").in(paging).out(jsonBody[List[User]])
```

Second, inputs can be combined by calling the `in`, `out` and `errorOut` methods on `Endpoint` multiple times. Each time such a method is invoked, it extends the list of inputs/outputs. This can be useful to separate different groups of

parameters, but also to define template-endpoints, which can then be further specialized. For example, we can define a base endpoint for our API, where all paths always start with `/api/v1.0`, and errors are always returned as a json:

```
import sttp.tapir._
import sttp.tapir.generic.auto._
import sttp.tapir.json.circe._
import io.circe.generic.auto._

case class ErrorInfo(message: String)

val baseEndpoint: Endpoint[Unit, ErrorInfo, Unit, Any] =
  endpoint.in("api" / "v1.0").errorOut(jsonBody[ErrorInfo])
```

Thanks to the fact that inputs/outputs accumulate, we can use the base endpoint to define more inputs, for example:

```
case class Status(uptime: Long)

val statusEndpoint: Endpoint[Unit, ErrorInfo, Status, Any] =
  baseEndpoint.in("status").out(jsonBody[Status])
```

The above endpoint will correspond to the `api/v1.0/status` path.

4.5.2 Mapping over input/output values

Inputs/outputs can also be mapped over. As noted before, all mappings are bi-directional, so that they can be used both when interpreting an endpoint as a server, and as a client, as well as both in input and output contexts.

There's a couple of ways to map over an input/output. First, there's the `map[II] (f: I => II) (g: II => I)` method, which accepts functions which provide the mapping in both directions. For example:

```
import sttp.tapir._
import java.util.UUID

case class Paging(from: UUID, limit: Option[Int])

val paging: EndpointInput[Paging] =
  query[UUID]("start").and(query[Option[Int]]("limit"))
  .map(input => Paging(input._1, input._2))(paging => (paging.from, paging.limit))
```

Next, you can use `mapDecode[II] (f: I => DecodeResult[II]) (g: II => I)`, to handle cases where decoding (mapping a low-level value to a higher-value one) can fail. There's a couple of failure reasons, captured by the alternatives of the `DecodeResult` trait.

Mappings can also be done given an `Mapping[I, II]` instance. More on that in the section on [codecs](#).

Creating a mapping between a tuple and a case class is a common operation, hence there's also a `mapTo(CaseClassCompanion)` method, which automatically provides the functions to construct/deconstruct the case class:

```
val paging: EndpointInput[Paging] =
  query[UUID]("start").and(query[Option[Int]]("limit"))
  .mapTo(Paging)
```

Mapping methods can also be called on an endpoint (which is useful if inputs/outputs are accumulated, for example). The `Endpoint.mapIn`, `Endpoint.mapInTo` etc. have the same signatures as the ones above.

4.5.3 Describing input/output values using annotations

Inputs and outputs can also be built for case classes using annotations. For example, for the case class `User`

```
import sttp.tapir.annotations._

case class User(
  @query
  name: String,
  @cookie
  sessionId: Long
)
```

endpoint input can be generated using macro `sttp.tapir.annotations.deriveEndpointInput[User]` which is equivalent to

```
import sttp.tapir._

val userInput: EndpointInput[User] =
  query[String]("user").and(cookie[Long]("sessionId")).mapTo(User)
```

Following annotations are available in package `sttp.tapir.annotations` for describing both input and output values:

- `@header` captures a header with the same name as name of annotated field in a case class. This annotation can also be used with optional parameter `@header("headerName")` in order to capture a header with name "headerName" if a name of header is different from name of annotated field in a case class
- `@headers` captures all headers. Can only be applied to fields represented as `List[Header]`
- `@cookies` captures all cookies. Can only be applied to fields represented as `List[Cookie]`
- `@jsonbody` captures JSON body of request or response. Can only be applied to field if there is implicit JSON Codec instance from `String` to target type
- `@xmlbody` captures XML body of request or response. Also requires implicit XML Codec instance from `String` to target type

Following annotations are only available for describing input values:

- `@query` captures a query parameter with the same name as name of annotated field in a case class. The same as annotation `@header` it has optional parameter to specify alternative name for query parameter
- `@params` captures all query parameters. Can only be applied to fields represented as `QueryParams`
- `@cookie` captures a cookie with the same name as name of annotated field in a case class. The same as annotation `@header` it has optional parameter to specify alternative name for cookie
- `@path` captures a path segment. Can only be applied to field of a case class if this case class is annotated by annotation `@endpointInput`. For example,

```
import sttp.tapir.annotations._

@endpointInput("books/{year}/{genre}")
case class Book(
  @path
  genre: String,
  @path
  year: Int,
  @query
```

(continues on next page)

(continued from previous page)

```
name: String
)
```

Annotation `@endpointInput` specifies endpoint path. In order to capture one segment of this path it must be surrounded in curly braces

- `@apikey` wraps any other input and designates it as an API key. Can only be used with another annotations
- `@basic` extracts data from the Authorization header. Can only be applied for field represented as `UsernamePassword`
- `@bearer` extracts data from the Authorization header removing the Bearer prefix.

Following annotations are only available for describing output values:

- `@setCookie` sends value in header `Set-Cookie`. The same as annotation `@header` it has optional parameter to specify alternative name for cookie. Can only be applied for field represented as `CookieValueWithMeta`
- `@setCookies` sends several `Set-Cookie` headers. Can only be applied for field represented as `List[Cookie]`
- `@statusCode` sets status code for response. Can only be applied for field represented as `StatusCode`

4.5.4 Path matching

By default (as with all other types of inputs), if no path input/path segments are defined, any path will match.

If any path input/path segment is defined, the path must match *exactly* - any remaining path segments will cause the endpoint not to match the request. For example, `endpoint.in("api")` will match `/api`, `/api/`, but won't match `/`, `/api/users`.

To match only the root path, use an empty string: `endpoint.in("")` will match `http://server.com/` and `http://server.com`.

To match a path prefix, first define inputs which match the path prefix, and then capture any remaining part using paths, e.g.: `endpoint.in("api" / "download").in(paths)`.

4.5.5 Next

Read on about [status codes](#).

4.6 Status codes

To provide a (varying) status code of a server response, use the `statusCode` output, which maps to a value of type `sttp.model.StatusCode`. The companion object contains known status codes as constants. This type of output is used only when interpreting the endpoint as a server. If your endpoint returns varying status codes which you would like to have listed in documentation use `statusCode.description(code1, "code1 description").description(code2, "code2 description")` output.

Alternatively, a fixed status code can be specified using the `statusCode(code)` output.

4.6.1 Dynamic status codes

It is also possible to specify how status codes map to different outputs. All mappings should have a common supertype, which is also the type of the output. These mappings are used to determine the status code when interpreting an endpoint as a server, as well as when generating documentation and to deserialise client responses to the appropriate type, basing on the status code.

For example, below is a specification for an endpoint where the error output is a sealed trait `ErrorInfo`; such a specification can then be refined and reused for other endpoints:

```
import sttp.tapir._
import sttp.tapir.json.circe._
import sttp.model.StatusCode
import io.circe.generic.auto._

sealed trait ErrorInfo
case class NotFound(what: String) extends ErrorInfo
case class Unauthorized(realm: String) extends ErrorInfo
case class Unknown(code: Int, msg: String) extends ErrorInfo
case object NoContent extends ErrorInfo

// here we are defining an error output, but the same can be done for regular outputs
val baseEndpoint = endpoint.errorOut(
  oneOf[ErrorInfo](
    statusMapping(StatusCode.NotFound, jsonBody[NotFound].description("not found")),
    statusMapping(StatusCode.Unauthorized, jsonBody[Unauthorized].description(
      ↪ "unauthorized")),
    statusMapping(StatusCode.NoContent, emptyOutput.map(_ => NoContent) (_ => ())),
    statusDefaultMapping(jsonBody[Unknown].description("unknown"))
  )
)
```

Each mapping, defined using the `statusMapping` method is a case class, containing the output description as well as the status code. Moreover, default mappings can be defined using `statusDefaultMapping`:

- for servers, the default status code for error outputs is 400, and for normal outputs 200 (unless a `statusCode` is used in the nested output)
- for clients, a default mapping is a catch-all.

Both `statusMapping` and `statusDefaultMapping` return a value of type `StatusMapping`. A list of these values can be dynamically assembled (e.g. using a default set of cases, plus endpoint-specific mappings), and provided to `oneOf`.

4.6.2 Status mapping and type erasure

Type erasure may prevent a status mapping from working properly. The following example will fail at compile time because `Right[NotFound]` and `Right[BadRequest]` will become `Right[Any]`:

```
import sttp.tapir._
import sttp.tapir.json.circe._
import sttp.tapir.generic.auto._
import sttp.model.StatusCode
import io.circe.generic.auto._

case class ServerError(what: String)
```

(continues on next page)

(continued from previous page)

```
sealed trait UserError
case class BadRequest(what: String) extends UserError
case class NotFound(what: String) extends UserError

val baseEndpoint = endpoint.errorOut(
  oneOf[Either[ServerError, UserError]](
    statusMapping(StatusCode.NotFound, jsonBody[Right[ServerError, NotFound]]),
    ↳description("not found")),
    statusMapping(StatusCode.BadRequest, jsonBody[Right[ServerError, BadRequest]]),
    ↳description("unauthorized")),
    statusMapping(StatusCode.InternalServerError, jsonBody[Left[ServerError,
    ↳UserError]]).description("unauthorized")),
  )
)
// error: Constructing statusMapping of type scala.util.Right[repl.MdocSession.App.
↳ServerError,repl.MdocSession.App.NotFound] is not allowed because of type erasure.
↳Using a runtime-class-based check it isn't possible to verify that the input
↳matches the desired class. Please use statusMappingClassMatcher,
↳statusMappingValueMatcher or statusMappingFromMatchType instead
//      statusMappingValueMatcher(StatusCode.NotFound, jsonBody[Right[ServerError,
↳NotFound]].description("not found")) {
//      ^
// error: Constructing statusMapping of type scala.util.Right[repl.MdocSession.App.
↳ServerError,repl.MdocSession.App.BadRequest] is not allowed because of type erasure.
↳ Using a runtime-class-based check it isn't possible to verify that the input
↳matches the desired class. Please use statusMappingClassMatcher,
↳statusMappingValueMatcher or statusMappingFromMatchType instead
//      statusMappingValueMatcher(StatusCode.BadRequest, jsonBody[Right[ServerError,
↳BadRequest]].description("unauthorized")) {
//      ^
// error: Constructing statusMapping of type scala.util.Left[repl.MdocSession.App.
↳ServerError,repl.MdocSession.App.UserError] is not allowed because of type erasure.
↳Using a runtime-class-based check it isn't possible to verify that the input
↳matches the desired class. Please use statusMappingClassMatcher,
↳statusMappingValueMatcher or statusMappingFromMatchType instead
//      statusMappingValueMatcher(StatusCode.InternalServerError,
↳jsonBody[Left[ServerError, UserError]].description("unauthorized")) {
//      ^
//
```

The solution is therefore to handwrite a function checking that a `val` (of type `Any`) is of the correct type:

```
val baseEndpoint = endpoint.errorOut(
  oneOf[Either[ServerError, UserError]](
    statusMappingValueMatcher(StatusCode.NotFound, jsonBody[Right[ServerError,
    ↳NotFound]].description("not found")) {
      case Right(NotFound(_)) => true
    },
    statusMappingValueMatcher(StatusCode.BadRequest, jsonBody[Right[ServerError,
    ↳BadRequest]].description("unauthorized")) {
      case Right(BadRequest(_)) => true
    },
    statusMappingValueMatcher(StatusCode.InternalServerError,
    ↳jsonBody[Left[ServerError, UserError]].description("unauthorized")) {
      case Left(ServerError(_)) => true
    }
  )
)
```

Of course you could use `statusMappingValueMatcher` to do runtime filtering for other purpose than solving type erasure.

In the case of solving type erasure, writing by hand partial function to match value against composition of case class and sealed trait can be repetitive. To make that more easy, we provide an **experimental** typeclass - `MatchType` - so you can automatically derive that partial function:

```
import sttp.tapir.typelevel.MatchType

val baseEndpoint = endpoint.errorOut(
  oneOf[Either[ServerError, UserError]](
    statusMappingFromMatchType(StatusCode.NotFound, jsonBody[Right[ServerError, _]
    ↳ NotFound].description("not found")),
    statusMappingFromMatchType(StatusCode.BadRequest, jsonBody[Right[ServerError, _]
    ↳ BadRequest].description("unauthorized")),
    statusMappingFromMatchType(StatusCode.InternalServerError, _
    ↳ jsonBody[Left[ServerError, UserError]].description("unauthorized"))
  )
)
```

4.6.3 Server interpreters

Unless specified otherwise, successful responses are returned with the 200 OK status code, and errors with 400 Bad Request. For exception and decode failure handling, see *error handling*.

4.6.4 Next

Read on about [codecs](#).

4.7 Codecs

4.7.1 Mappings

The `Mapping[L, H]` trait defines a bi-directional mapping between values of type `L` and values of type `H`.

Low-level values of type `L` can be **decoded** to a higher-level value of type `H`. The decoding can fail; this is represented by a result of type `DecodeResult.Failure`. Failures might occur due to format errors, wrong arity, exceptions, or validation errors. Validators can be added through the `validate` method.

High-level values of type `H` can be **encoded** as a low-level value of type `L`.

Mappings can be chained using one of the `map` functions.

4.7.2 Codecs

A `Codec[L, H, CF]` is a `Mapping[L, H]`, with additional meta-data: a schema and the format of the low-level value (more on that below).

There are built-in codecs for most common types such as `String`, `Int` etc. Codecs are usually defined as implicit values and resolved implicitly when they are referenced. However, they can also be provided explicitly as needed.

For example, a `query[Int] ("quantity")` specifies an input parameter which corresponds to the `quantity` query parameter and will be mapped as an `Int`. There's an implicit `Codec[List[String], Int, TextPlain]` value that is referenced by the `query` method (which is defined in the `sttp.tapir` package).

In this example, the low-level value is a `List[String]`, as a given query parameter can be absent, have a single or many values. The high-level value is an `Int`. The codec will verify that there's a single query parameter with the given name, and parse it as an `int`. If any of this fails, a failure will be reported.

In a server setting, if the value cannot be parsed as an `int`, a decoding failure is reported, and the endpoint won't match the request, or a `400 Bad Request` response is returned (depending on configuration).

As each codec is also a `Mapping`, codecs can be used to map endpoint inputs and outputs. In such cases, the additional meta-data (schema & format) isn't taken into account.

4.7.3 Optional and multiple parameters

Some inputs/outputs allow optional, or multiple parameters:

- path segments are always required
- query and header values can be optional or multiple (repeated query parameters/headers)
- bodies can be optional, but not multiple

In general, optional parameters are represented as `Option` values, and multiple parameters as `List` values. For example, `header[Option[String]] ("X-Auth-Token")` describes an optional header. An input described as `query[List[String]] ("color")` allows multiple occurrences of the `color` query parameter, with all values gathered into a list.

4.7.4 Schemas

A codec contains a schema, which describes the high-level type. The schema is used when generating documentation and enforcing validation rules.

Schema consists of:

- the schema type, which is one of the values defined in `SchemaType`, such as `SString`, `SBinary`, `SArray` or `SProduct/SCoproduct` (for ADTs). This is the shape of the encoded value - as it is sent over the network
- meta-data: value optionality, description and low-level format name
- validation rules

For primitive types, the schema values are built-in, and defined in the `Schema` companion object.

The schema is left unchanged when mapping a codec, or an input/output, as the underlying representation of the value doesn't change. However, schemas can be changed for individual inputs/outputs using the `.schema (Schema)` method.

When codecs are derived for complex types, e.g. for json mapping, schemas are looked up through implicit `Schema[T]` values. See [custom types](#) for more details.

4.7.5 Codec format

Codecs contain an additional type parameter, which specifies the codec format. Each format corresponds to a media type, which describes the low-level format of the raw value (to which the codec encodes). Some built-in formats include `text/plain`, `application/json` and `multipart/form-data`. Custom formats can be added by creating an implementation of the `sttp.tapir.CodecFormat` trait.

Thanks to codecs being parametrised by codec formats, it is possible to have a `Codec[String, MyCaseClass, TextPlain]` which specifies how to serialize a case class to plain text, and a different `Codec[String, MyCaseClass, Json]`, which specifies how to serialize a case class to json. Both can be implicitly available without implicit resolution conflicts.

Different codec formats can be used in different contexts. When defining a path, query or header parameter, only a codec with the `TextPlain` media type can be used. However, for bodies, any media types is allowed. For example, the input/output described by `jsonBody[T]` requires a json codec.

4.7.6 Next

Read on about [custom types](#).

4.8 Custom types

To support a custom type, you'll need to provide an implicit `Codec` for that type.

This can be done by writing a codec from scratch, mapping over an existing codec, or automatically deriving one. Which of these approaches can be taken, depends on the context in which the codec will be used.

4.8.1 Providing an implicit codec

To create a custom codec, you can either directly implement the `Codec` trait, which requires to provide the following information:

- `encode` and `rawDecode` methods
- schema (for documentation and validation)
- codec format (`text/plain`, `application/json` etc.)

This might be quite a lot of work, that's why it's usually easier to map over an existing codec. To do that, you'll need to provide two mappings:

- a `decode` method which decodes the lower-level type into the custom type, optionally reporting decode failures (the return type is a `DecodeResult`)
- an `encode` method which encodes the custom type into the lower-level type

For example, to support a custom id type:

```
import scala.util._

class MyId private (id: String) {
  override def toString(): String = id
}

object MyId {
  def parse(id: String): Try[MyId] = {
    Success(new MyId(id))
  }
}
```

```
import sttp.tapir._
import sttp.tapir.CodecFormat.TextPlain
```

(continues on next page)

(continued from previous page)

```
def decode(s: String): DecodeResult[MyId] = MyId.parse(s) match {  
  case Success(v) => DecodeResult.Value(v)  
  case Failure(f) => DecodeResult.Error(s, f)  
}  
def encode(id: MyId): String = id.toString  
  
implicit val myIdCodec: Codec[String, MyId, TextPlain] =  
  Codec.string.mapDecode(decode)(encode)
```

Or, using the type alias for codecs in the TextPlain format and String as the raw value:

```
import sttp.tapir.Codec.PlainCodec  
  
implicit val myIdCodec: PlainCodec[MyId] = Codec.string.mapDecode(decode)(encode)
```

Note: Note that inputs/outputs can also be mapped over. In some cases, it's enough to create an input/output corresponding to one of the existing types, and then map over them. However, if you have a type that's used multiple times, it's usually better to define a codec for that type.

Then, you can use the new codec e.g. to obtain an id from a query parameter, or a path segment:

```
endpoint.in(query[MyId]("myId"))  
// or  
endpoint.in(path[MyId])
```

4.8.2 Automatically deriving codecs

In some cases, codecs can be automatically derived:

- for supported [json](#) libraries
- for urlencoded and multipart [forms](#)

Automatic codec derivation usually requires other implicits, such as:

- json encoders/decoders from the [json](#) library
- codecs for individual form fields
- schema of the custom type, through the `Schema[T]` implicit

Note the derivation of e.g. circe json encoders/decoders and tapir schemas are separate processes, and must be configured separately.

4.8.3 Schema derivation

For case classes types, `Schema[_]` values can be derived automatically using [Magnolia](#), given that schemas are defined for all the case class's fields.

Two policies enable to choose the way custom types are derived:

- automatic derivation
- semi automatic derivation

Automatic derivation

Case classes, traits and their children are recursively derived by Magnolia.

Importing `sttp.tapir.generic.auto._` (or extending the `SchemaDerivation` trait) enables fully automatic derivation for `Schema`:

```
import sttp.tapir.Schema
import sttp.tapir.generic.auto._

case class Parent(child: Child)
case class Child(value: String)

// implicit schema used by codecs
implicitly[Schema[Parent]]
```

If you have a case class which contains some non-standard types (other than strings, number, other case classes, collections), you only need to provide schemas for them. Using these, the rest will be derived automatically.

Note that when using `datatypes` integrations, respective codecs must also be imported to enable the derivation, e.g. for `newtype` you'll have to add `import sttp.tapir.codec.newtype._` or extend `TapirCodecNewType`.

Semi-automatic derivation

Semi-automatic derivation can be done using `Schema.derived[T]`.

It only derives selected type `T`. However, derivation is not recursive: schemas must be explicitly defined for every child type.

This mode is easier to debug and helps to avoid issues encountered by automatic mode (wrong schemas for value classes or custom types):

```
import sttp.tapir.Schema

case class Parent(child: Child)
case class Child(value: String)

implicit lazy val sChild: Schema[Child] = Schema.derived
implicit lazy val sParent: Schema[Parent] = Schema.derived
```

Note that while schemas for regular types can be safely defined as `vals`, in case of recursive values, the schema values must be `lazy vals`.

Configuring derivation

It is possible to configure Magnolia's automatic derivation to use `snake_case`, `kebab-case` or a custom field naming policy, by providing an implicit `sttp.tapir.generic.Configuration` value. This influences how the low-level representation is described in documentation:

```
import sttp.tapir.generic.Configuration

implicit val customConfiguration: Configuration =
  Configuration.default.withSnakeCaseMemberNames
```

Alternatively, `Schema[_]` values can be defined by hand, either for whole case classes, or only for some of its fields. For example, here we state that the schema for `MyCustomType` is a `String`:

```
import sttp.tapir._

case class MyCustomType()
implicit val schemaForMyCustomType: Schema[MyCustomType] = Schema.string
// or, if the low-level representation is e.g. a number
implicit val anotherSchemaForMyCustomType: Schema[MyCustomType] = Schema(SchemaType.
  ↳ SInteger)
```

Sealed traits / coproducts

Schema derivation for coproduct types (sealed trait hierarchies) is supported as well. By default, such hierarchies will be represented as a coproduct which contains a list of child schemas, without any discriminator field.

A discriminator field can be specified for coproducts by providing it in the configuration; this will be only used during automatic derivation:

```
import sttp.tapir.generic.Configuration

implicit val customConfiguration: Configuration =
  Configuration.default.withDiscriminator("who_am_i")
```

Alternatively, derived schemas can be customised (see below), and a discriminator can be added by calling the `SchemaType.SCoproduct.addDiscriminatorField(name, schema, mappingOverride)` method.

Finally, if the discriminator is a field that's defined on the base trait (and hence in each implementation), the schemas can be specified using `Schema.oneOfUsingField`, for example (this will also generate the appropriate mapping overrides):

```
sealed trait Entity {
  def kind: String
}
case class Person(firstName:String, lastName:String) extends Entity {
  def kind: String = "person"
}
case class Organization(name: String) extends Entity {
  def kind: String = "org"
}

import sttp.tapir._

val sPerson = Schema.derived[Person]
val sOrganization = Schema.derived[Organization]
implicit val sEntity: Schema[Entity] =
  Schema.oneOfUsingField[Entity, String](_.kind, _.toString) ("person" -> sPerson,
  ↳ "org" -> sOrganization)
```

4.8.4 Customising derived schemas

In some cases, it might be desirable to customise the derived schemas, e.g. to add a description to a particular field of a case class. One way the automatic derivation can be customised is using annotations:

- `@encodedName` sets name for case class's field which is used in the encoded form (and also in documentation)
- `@description` sets description for the whole case class or its field

- `@default` sets default value for a case class field
- `@encodedExample` sets example value for a case class field which is used in the documentation in the encoded form
- `@format` sets the format for a case class field
- `@deprecated` marks a case class's field as deprecated

If the target type isn't accessible or can't be modified, schemas can be customized by looking up an implicit instance of the `Derived[Schema[T]]` type, modifying the value, and assigning it to an implicit schema.

When such an implicit `Schema[T]` is in scope will have higher priority than the built-in low-priority conversion from `Derived[Schema[T]]` to `Schema[T]`.

Schemas for products/coproducts (case classes and case class families) can be traversed and modified using `.modify` method. To traverse collections, use `.each`.

For example:

```
import sttp.tapir._
import sttp.tapir.generic.auto._
import sttp.tapir.generic.Derived

case class Basket(fruits: List[FruitAmount])
case class FruitAmount(fruit: String, amount: Int)
implicit val customBasketSchema: Schema[Basket] = implicitly[Derived[Schema[Basket]]].
  ↪value
  .modify(_.fruits.each.amount) (_.description("How many fruits?"))
```

There is also an unsafe variant of this method, but it should be avoided in most cases. The “unsafe” prefix comes from the fact that the method takes a list of strings, which represent fields, and the correctness of this specification is not checked.

Non-standard collections can be unwrapped in the modification path by providing an implicit value of `ModifyFunctor`.

4.8.5 Next

Read on about [validation](#).

4.9 Validation

Tapir supports validation for primitive/base types. Validation of composite values, whole data structures, business rules enforcement etc. should be done as part of the *server logic* of the endpoint, using the dedicated error output (the `E` in `Endpoint[I, E, O, S]`) to report errors.

4.9.1 Single type validation

Validation rules are part of the `Schema` for a given type, and can be added either directly to the schema, or via the `Codec` or `EndpointInput/EndpointOutput`. For example, when defining a codec for a type, we have the `.validate()` method:

```
import sttp.tapir._
import sttp.tapir.CodecFormat.TextPlain

case class MyId(id: String)

implicit val myIdCodec: Codec[String, MyId, TextPlain] = Codec.string
  .map(MyId(_))(_.id)
  .validate(Validator.pattern("[A-Z].*").contramap(_.id))
```

Validators can also be added to individual inputs/outputs, in addition to whatever the codec provides:

```
import sttp.tapir._

val e = endpoint.in(
  query[Int]("amount")
    .validate(Validator.min(0))
    .validate(Validator.max(100)))
```

Validation rules added using the built-in validators are translated to [OpenAPI](#) documentation.

4.9.2 Validation rules and automatic codec derivation

As validators are parts of schemas, they are looked up as part of the with implicit `Schema[T]` values.

While they can be manually defined, tapir provides tools to derive automatically schemas for custom types (traits and case classes).

4.9.3 Custom type validation

Note that to validate a nested member of a case class, it needs to have a unique type (that is, not an `Int`, as providing an implicit `Validator[Int]` would validate all ints in the hierarchy), as schema lookup is type-driven.

To introduce unique types for primitive values, you can use value classes or [type tagging](#).

For example, to support an integer wrapped in a value type in a json body, we need to provide Circe encoders and decoders (if that's the json library that we are using), schema information with validator:

```
import sttp.tapir._
import sttp.tapir.generic.auto._
import sttp.tapir.json.circe._
import io.circe.{Encoder, Decoder}
import io.circe.generic.semiauto._

case class Amount(v: Int) extends AnyVal
case class FruitAmount(fruit: String, amount: Amount)

implicit val amountSchema: Schema[Amount] = Schema(SchemaType.SInteger)
  .validate(Validator.min(1).contramap(_.v))
implicit val amountEncoder: Encoder[Amount] = Encoder.encodeInt.contramap(_.v)
implicit val amountDecoder: Decoder[Amount] = Decoder.decodeInt.map(Amount.apply)

implicit val decoder: Decoder[FruitAmount] = deriveDecoder[FruitAmount]
implicit val encoder: Encoder[FruitAmount] = deriveEncoder[FruitAmount]

val e: Endpoint[FruitAmount, Unit, Unit, Nothing] =
  endpoint.in(jsonBody[FruitAmount])
```

4.9.4 Decode failures

Codecs support reporting decoding failures, by returning a `DecodeResult` from the `Codec.decode` method. However, this is meant for input/output values which are in an incorrect low-level format, when parsing a “raw value” fails. In other words, decoding failures should be reported for format failures, not business validation errors.

To customise error messages that are returned upon validation/decode failures by the server, see [error handling](#).

4.9.5 Enum validators

Validators for enumerations can be created using the `Validator.enum` method, which either:

- takes a type parameter, which should be an abstract, sealed base type, and using a macro determines the possible implementations
- takes the list of possible values

To properly represent possible values in documentation, the enum validator additionally needs an `encode` method, which converts the enum value to a raw type (typically a string). This method is inferred *only* if the validator is directly added to a codec (without any mapping etc.), for example:

```
import sttp.tapir._
import sttp.tapir.Codec.PlainCodec

sealed trait Color
case object Blue extends Color
case object Red extends Color

implicit def plainCodecForColor: PlainCodec[Color] = {
  Codec.string
    .map[Color]((_: String) match {
      case "red" => Red
      case "blue" => Blue
    }) (_._.toString.toLowerCase)
    .validate(Validator.enum)
}
```

If the enum is nested within an object, regardless of whether the codec for that object is defined by hand or derived, we need to specify the encode function by hand:

```
implicit def colorSchema: Schema[Color] = Schema.string.validate(Validator.enum,
  ↪ encode(_._.toString.toLowerCase))
```

Like other validators/schemas, enum schemas need to be added to a codec manually or through an implicit value, if the codec and validator is automatically derived.

4.9.6 Next

Read on about [json support](#).

4.10 Working with JSON

Json values are supported through codecs, which encode/decode values to json strings. Most often, you’ll be using a third-party library to perform the actual json parsing/printing. Currently, [Circe](#), [µPickle](#), [Spray JSON](#) and [Play JSON](#) are supported.

All of the integrations, when imported into scope, define a `jsonBody[T]` method. This method depends on library-specific implicits being in scope, and derives from them a json codec. The derivation also requires implicit `Schema[T]` and `Validator[T]` instances, which should be automatically derived. For more details see documentation on supporting [custom types](#).

If you have a custom, implicit `Codec[String, T, Json]` instance, you should use the `anyJsonBody[T]` method instead. This description of endpoint input/output, instead of deriving a codec basing on other library-specific implicits, uses the json codec that is in scope.

4.10.1 Circe

To use Circe add the following dependency to your project:

```
"com.softwaremill.sttp.tapir" %% "tapir-json-circe" % "0.17.15"
```

Next, import the package (or extend the `TapirJsonCirce` trait, see [MyTapir](#)):

```
import sttp.tapir.json.circe._
```

This will allow automatically deriving Codecs which, given an in-scope circe `Encoder/Decoder` and a `Schema`, will create a codec using the json media type. Circe includes a couple of approaches to generating encoders/decoders (manual, semi-auto and auto), so you may choose whatever suits you.

Note that when using Circe's auto derivation, any encoders/decoders for custom types must be in scope as well.

Additionally, the above import brings into scope the `jsonBody[T]` body input/output description, which uses the above codec.

For example, to automatically generate a JSON codec for a case class:

```
import sttp.tapir._
import sttp.tapir.json.circe._
import sttp.tapir.generic.auto._
import io.circe.generic.auto._

case class Book(author: String, title: String, year: Int)

val bookInput: EndpointIO[Book] = jsonBody[Book]
```

Circe lets you select an instance of `io.circe.Printer` to configure the way JSON objects are rendered. By default Tapir uses `Printer.nospaces`, which would render:

```
import io.circe._

Json.obj(
  "key1" -> Json.fromString("present"),
  "key2" -> Json.Null
)
```

as

```
{"key1": "present", "key2": null}
```

Suppose we would instead want to omit null-values from the object and pretty-print it. You can configure this by overriding the `jsonPrinter` in `tapir.circe.json.TapirJsonCirce`:

```
import sttp.tapir.json.circe._
import io.circe.Printer

object MyTapirJsonCirce extends TapirJsonCirce {
  override def jsonPrinter: Printer = Printer.spaces2.copy(dropNullValues = true)
}

import MyTapirJsonCirce._
```

Now the above JSON object will render as

```
{"key1": "present"}
```

4.10.2 μ Pickle

To use μ Pickle add the following dependency to your project:

```
"com.softwaremill.sttp.tapir" %% "tapir-json-upickle" % "0.17.15"
```

Next, import the package (or extend the `TapirJsonuPickle` trait, see *MyTapir* and add `TapirJsonuPickle` not `TapirCirceJson`):

```
import sttp.tapir.json.upickle._
```

μ Pickle requires a `ReadWriter` in scope for each type you want to serialize. In order to provide one use the `macroRW` macro in the companion object as follows:

```
import sttp.tapir._
import sttp.tapir.generic.auto._
import upickle.default._
import sttp.tapir.json.upickle._

case class Book(author: String, title: String, year: Int)

object Book {
  implicit val rw: ReadWriter[Book] = macroRW
}

val bookInput: EndpointIO[Book] = jsonBody[Book]
```

Like Circe, μ Pickle allows you to control the rendered json output. Please see the [Custom Configuration](#) of the manual for details.

For more examples, including making a custom encoder/decoder, see [TapirJsonuPickleTests.scala](#)

4.10.3 Play JSON

To use Play JSON add the following dependency to your project:

```
"com.softwaremill.sttp.tapir" %% "tapir-json-play" % "0.17.15"
```

Next, import the package (or extend the `TapirJsonPlay` trait, see *MyTapir* and add `TapirJsonPlay` not `TapirCirceJson`):

```
import sttp.tapir.json.play._
```

Play JSON requires `Reads` and `Writes` implicit values in scope for each type you want to serialize.

4.10.4 Spray JSON

To use Spray JSON add the following dependency to your project:

```
"com.softwaremill.sttp.tapir" %% "tapir-json-spray" % "0.17.15"
```

Next, import the package (or extend the `TapirJsonSpray` trait, see *MyTapir* and add `TapirJsonSpray` not `TapirCirceJson`):

```
import sttp.tapir.json.spray._
```

Spray JSON requires a `JsonFormat` implicit value in scope for each type you want to serialize.

4.10.5 Tethys JSON

To use Tethys JSON add the following dependency to your project:

```
"com.softwaremill.sttp.tapir" %% "tapir-json-tethys" % "0.17.15"
```

Next, import the package (or extend the `TapirJsonTethys` trait, see *MyTapir* and add `TapirJsonTethys` not `TapirCirceJson`):

```
import sttp.tapir.json.tethysjson._
```

Tethys JSON requires `JsonReader` and `JsonWriter` implicit values in scope for each type you want to serialize.

4.10.6 Jsoniter Scala

To use `Jsoniter-scala` add the following dependency to your project:

```
"com.softwaremill.sttp.tapir" %% "tapir-jsoniter-scala" % "0.17.15"
```

Next, import the package (or extend the `TapirJsonJsoniter` trait, see *MyTapir* and add `TapirJsonJsoniter` not `TapirCirceJson`):

```
import sttp.tapir.json.jsoniter._
```

Jsoniter Scala requires `JsonValueCodec` implicit value in scope for each type you want to serialize.

4.10.7 Json4s

To use `json4s` add the following dependencies to your project:

```
"com.softwaremill.sttp.tapir" %% "tapir-json4s" % "0.17.15"
```

And one of the implementations:

```
"org.json4s" %% "json4s-native" % "3.6.11"
// Or
"org.json4s" %% "json4s-jackson" % "3.6.11"
```

Next, import the package (or extend the `TapirJson4s` trait, see [MyTapir](#) and add `TapirJson4s` instead of `TapirCirceJson`):

```
import sttp.tapir.json.json4s._
```

`Json4s` requires `Serialization` and `Formats` implicit values in scope, for example:

```
import org.json4s._
// ...
implicit val serialization: Serialization = org.json4s.jackson.Serialization
implicit val formats: Formats = org.json4s.jackson.Serialization.formats(NoTypeHints)
```

4.10.8 Other JSON libraries

To add support for additional JSON libraries, see the [sources](#) for the Circe codec (which is just a couple of lines of code).

4.10.9 Schemas

To derive json codecs automatically, not only implicits from the base library are needed (e.g. a circe `Encoder/Decoder`), but also an implicit `Schema[T]` value, which provides a mapping between a type `T` and its schema. A schema-for value contains a single schema: `Schema field`.

See [custom types](#) for details.

4.10.10 Next

Read on about [working with forms](#).

4.11 Forms

4.11.1 URL-encoded forms

An URL-encoded form input/output can be specified in two ways. First, it is possible to map all form fields as a `Seq[(String, String)]`, or `Map[String, String]` (which is more convenient if fields can't have multiple values):

```
import sttp.tapir._

formBody[Seq[(String, String)]]: EndpointIO.Body[String, Seq[(String, String)]]
formBody[Map[String, String]]: EndpointIO.Body[String, Map[String, String]]
```

Second, form data can be mapped to a case class. The codec for the case class is automatically derived using a macro at compile-time. The fields of the case class should have types, for which there is a plain text codec. For example:

```
import sttp.tapir._
import sttp.tapir.generic.auto._

case class RegistrationForm(name: String, age: Int, news: Boolean, city: Option[String])

formBody[RegistrationForm]: EndpointIO.Body[String, RegistrationForm]
```

Each form-field is named the same as the case-class-field. The names can be transformed to snake or kebab case by providing an implicit `tapir.generic.Configuration`, or customised using the `@encodedName` annotation.

4.11.2 Multipart forms

Similarly as above, multipart form input/outputs can be specified in two ways. To map to all parts of a multipart body, use:

```
import sttp.tapir._
import sttp.model.Part

multipartBody: EndpointIO.Body[Seq[RawPart], Seq[Part[Array[Byte]]]]
```

`Part` is a case class containing the name of the part, disposition parameters, headers, and the body. The bodies will be mapped as byte arrays (`Array[Byte]`). Custom multipart codecs can be defined with the `Codec.multipartCodec` method, and then used with `multipartBody[T]`.

As with URL-encoded forms, multipart bodies can be mapped directly to case classes, however without the restriction on codecs for individual fields. Given a field of type `T`, first a plain text codec is looked up, and if one isn't found, any codec for any media type (e.g. JSON) is searched for.

Each part is named the same as the case-class-field. The names can be transformed to snake or kebab case by providing an implicit `sttp.tapir.generic.Configuration`, or customised using the `@encodedName` annotation.

Additionally, the case class to which the multipart body is mapped can contain both normal fields, and fields of type `Part[T]`. This is useful, if part metadata (e.g. the filename) is relevant.

For example:

```
import sttp.tapir._
import sttp.model.Part
import java.io.File
import sttp.tapir.generic.auto._

case class RegistrationForm(userData: User, photo: Part[File], news: Boolean)
case class User(email: String)

multipartBody[RegistrationForm]: EndpointIO.Body[Seq[RawPart], RegistrationForm]
```

4.11.3 Next

Read on about [authentication](#).

4.12 Authentication

Inputs which carry authentication data wrap another input can be marked as such by declaring them using members of the `auth` object. Apart from predefined codecs for some authentication methods, such inputs will be treated differently when generating documentation. Otherwise, they behave as normal inputs which map to the given type.

Currently, the following authentication inputs are available (assuming `import sttp.tapir._`):

- `auth.apiKey(anotherInput)`: wraps any other input and designates it as an api key. The input is typically a header, cookie or a query parameter
- `auth.basic[T]`: reads data from the `Authorization` header, removing the `Basic` prefix. To parse the data as a base64-encoded username/password combination, use: `basic[UsernamePassword]`.
- `auth.bearer[T]`: reads data from the `Authorization` header, removing the `Bearer` prefix. To get the string as a token, use: `bearer[String]`.
- `auth.oauth2.authorizationCode(authorizationUrl, tokenUrl, scopes, refreshUrl): EndpointInput[String]`: creates an OAuth2 authorization using authorization code - sign in using an auth service (for documentation, requires defining also the `oauth2-redirect.html`, see [Generating OpenAPI documentation](#))

Multiple authentication inputs indicate that all of the given authentication values should be provided. Specifying alternative authentication methods (where only one value out of many needs to be provided) is currently not supported. However, optional authentication can be described by mapping to optional types, e.g. `bearer[Option[String]]`.

When interpreting a route as a server, it is useful to define the authentication input first, to be able to share the authentication logic among multiple endpoints easily. See [server logic](#) for more details.

For each `auth` scheme, one can define `WWW-Authenticate` headers that should be returned by the server in case input is not provided. Default behavior is to return 401 status with the headers needed for authentication.

For example if you define `endpoint.get.in("path").in(auth.basic[UsernamePassword]())` then your browser will show you a password prompt.

4.12.1 Next

Read on about [streaming support](#).

4.13 Streaming support

Both input and output bodies can be mapped to a stream, by using `streamBody(streams)`. The parameter `streams` must implement the `Streams[S]` capability, and determines the precise type of the binary stream supported by the given non-blocking streams implementation. The interpreter must then support the given capability. Refer to the documentation of server/client interpreters for more information.

Note: Here, streams refer to asynchronous, non-blocking, “reactive” stream implementations, such as [akka-streams](#), [fs2](#) or [zio-streams](#). If you’d like to use blocking streams (such as `InputStream`), these are available through e.g. `inputStreamBody` without any additional requirements on the interpreter.

Adding a stream body input/output influences both the type of the input/output, as well as the 4th type parameter of `Endpoint`, which specifies the requirements regarding supported stream types for interpreters.

When using a stream body, the schema (needed for documentation) and format (media type) of the body must be provided by hand, as they cannot be inferred from the raw stream type. For example, to specify that the output is an akka-stream, which is a (presumably large) serialised list of json objects mapping to the `Person` class:

```
import sttp.tapir._
import sttp.tapir.generic.auto._
import sttp.capabilities.akka.AkkaStreams
import akka.stream.scaladsl._
import akka.util.ByteString

case class Person(name: String)

// copying the derived json schema type
endpoint.out(streamBody(AkkaStreams)(Schema(Schema.derived[List[Person]].schemaType),
↳ CodecFormat.Json()))
```

See also the *runnable streaming example*.

4.13.1 Next

Read on about [web sockets](#).

4.14 Web sockets

Web sockets are supported through stream pipes, converting a stream of incoming messages to a stream of outgoing messages. That's why web socket endpoints require both the `Streams` and `WebSocket` capabilities (see [streaming support](#) for more information on streams).

4.14.1 Typed web sockets

Web sockets outputs can be used in two variants. In the first, both requests and responses are handled by `codecs`. Typically, a codec handles either text or binary messages, signalling decode failure (which closes the web socket), if an unsupported frame is passed to decoding.

For example, here's an endpoint where the requests are strings (hence only text frames are handled), and the responses are parsed/formatted as json:

```
import sttp.tapir._
import sttp.capabilities.akka.AkkaStreams
import sttp.tapir.json.circe._
import sttp.tapir.generic.auto._
import io.circe.generic.auto._

case class Response(msg: String, count: Int)
endpoint.out(
  websocketBody[String, CodecFormat.TextPlain, Response, CodecFormat.
↳ Json](AkkaStreams)
```

When creating a `websocketBody`, we need to provide the following parameters:

- the type of requests, along with its codec format (which is used to lookup the appropriate codec, as well as determines the media type in documentation)
- the type of responses, along with its codec format

- the `Streams` implementation, which determines the pipe type

By default, ping-pong frames are handled automatically, fragmented frames are combined, and close frames aren't decoded, but this can be customised through methods on `WebSocketBody`.

4.14.2 Raw web sockets

Alternatively, it's possible to obtain a raw pipe transforming `WebSocketFrames`:

```
import akka.stream.scaladsl.Flow
import sttp.tapir._
import sttp.capabilities.akka.AkkaStreams
import sttp.capabilities.WebSockets
import sttp.ws.WebSocketFrame

endpoint.out(webSocketBodyRaw(AkkaStreams)) : Endpoint [
  Unit,
  Unit,
  Flow[WebSocketFrame, WebSocketFrame, Any],
  AkkaStreams with WebSockets]
```

Such a pipe by default doesn't handle ping-pong frames automatically, doesn't concatenate fragmented flames, and passes close frames to the pipe as well. As before, this can be customised by methods on the returned output.

Request/response schemas can be customised through `.requestsSchema` and `.responsesSchema`.

4.14.3 Interpreting as a sever

When interpreting a web socket endpoint as a server, the *server logic* needs to provide a streaming-specific pipe from requests to responses. E.g. in akka's case, this will be `Flow[REQ, RESP, Any]`.

Refer to the documentation of interpreters for more details, as not all interpreters support all settings.

4.14.4 Interpreting as a client

When interpreting a web socket endpoint as a client, after applying the input parameters, the result is a pipe representing message processing as it happens on the server.

Refer to the documentation of interpreters for more details, as there are interpreter-specific additional requirements.

4.14.5 Interpreting as documentation

Web socket endpoints can be interpreted into *AsyncAPI documentation*.

4.14.6 Determining if the request is a web socket upgrade

The `isWebSocket` endpoint input can be used to determine if the request contains the web socket upgrade headers. The input only impacts server interpreters, doesn't affect documentation and its value is discarded by client interpreters.

4.14.7 Next

Read on about [datatypes integrations](#).

4.15 Datatypes integrations

4.15.1 Cats datatypes integration

The `tapir-cats` module contains additional instances for some `cats` datatypes as well as additional syntax:

```
"com.softwaremill.sttp.tapir" %% "tapir-cats" % "0.17.15"
```

- `import sttp.tapir.integ.cats.codec._` - brings schema, validator and codec instances
- `import sttp.tapir.integ.cats.syntax._` - brings additional syntax for `tapir` types

4.15.2 Refined integration

If you use `refined`, the `tapir-refined` module will provide implicit codecs and validators for `T Refined P` as long as a codec for `T` already exists:

```
"com.softwaremill.sttp.tapir" %% "tapir-refined" % "0.17.15"
```

You'll need to extend the `sttp.tapir.codec.refined.TapirCodecRefined` trait or import `sttp.tapir.codec.refined._` to bring the implicit values into scope.

The refined codecs contain a validator which wrap/unwrap the value from/to its refined equivalent.

Some predicates will bind correctly to the vanilla `tapir Validator`, while others will bind to a custom validator that might not be very clear when reading the generated documentation. Correctly bound predicates can be found in `integration/refined/src/main/scala/sttp/tapir/codec/refined/TapirCodecRefined.scala`. If you are not satisfied with the validator generated by `tapir-refined`, you can provide an implicit `ValidatorForPredicate[T, P]` in scope using `'ValidatorForPredicate.fromPrimitiveValidator'` to build it (do not hesitate to contribute your work!).

4.15.3 Enumeratum integration

The `tapir-enumeratum` module provides schemas, validators and codecs for `Enumeratum` enumerations. To use, add the following dependency:

```
"com.softwaremill.sttp.tapir" %% "tapir-enumeratum" % "0.17.15"
```

Then, import `sttp.tapir.codec.enumeratum`, or extends the `sttp.tapir.codec.enumeratum.TapirCodecEnumeratum` trait.

This will bring into scope implicit values for values extending `*EnumEntry`.

4.15.4 Enumeration integration

There is no library for the use of the build in scala `Enumeration`, but it can be implemented by hand.

The example code below will generate `enums` to the open-api documentation.

```
import sttp.tapir._

trait EnumHelper { e: Enumeration =>
  import io.circe._
```

(continues on next page)

(continued from previous page)

```

implicit val enumDecoder: Decoder[e.Value] = Decoder.decodeEnumeration(e)
implicit val enumEncoder: Encoder[e.Value] = Encoder.encodeEnumeration(e)

// needs to be a def or lazy val so that the enumeration values are available!
implicit def schemaForEnum: Schema[e.Value] = Schema.string.validate(Validator.
  ↪enum(e.values.toList, v => Option(v)))
}
object Color extends Enumeration with EnumHelper {
  type Color = Value
  val Blue = Value("blue")
  val Red   = Value("red")
}

```

4.15.5 NewType integration

If you use `scala-newtype`, the `tapir-newtype` module will provide implicit codecs and schemas for a types with a `@newtype` and `@newsubtype` annotations as long as a codec and schema for its underlying value already exists:

```
"com.softwaremill.sttp.tapir" %% "tapir-newtype" % "0.17.15"
```

Then, import `sttp.tapir.codec.newtype._`, or extend the `sttp.tapir.codec.enumeratum.TapirCodecNewType` trait to bring the implicit values into scope.

4.15.6 Derevo integration

The `tapir-derevo` module provides a way to derive schema for your type using `@derive` annotation. For details refer to [derevo documentation](#). To use, add the following dependency:

```
"com.softwaremill.sttp.tapir" %% "tapir-derevo" % "0.17.15"
```

Then you can derive schema for your ADT along with other typeclasses besides ADT declaration itself:

```

import derevo.derive
import sttp.tapir.derevo.schema

@derive(schema)
case class Person(name: String, age: Int)

//or with custom description

@derive(schema("Type of currency in the country"))
sealed trait Currency
object Currency {
  case object CommunisticCurrency extends Currency
  case class USD(amount: Long) extends Currency
}

```

The annotation will simply generate a `Schema[T]` for your type `T` and put it into companion object. Generation rules are the same as in `Schema.derived[T]`.

4.16 ZIO integration

The `tapir-zio` module defines type aliases and extension methods which make it more ergonomic to work with `ZIO` and `tapir`. Moreover, the `tapir-zio-http4s-server` contains an interpreter useful when exposing the endpoints using the `http4s` server.

You'll need the following dependencies:

```
"com.softwaremill.sttp.tapir" %% "tapir-zio" % "0.17.15"
"com.softwaremill.sttp.tapir" %% "tapir-zio-http4s-server" % "0.17.15"
```

Next, instead of the usual `import sttp.tapir._`, you should import:

```
import sttp.tapir.ztapir._
```

This brings into scope all of the `basic` input/output descriptions, which can be used to define an endpoint. Additionally, it defines the `ZEndpoint` type alias, which should be used instead of `Endpoint`.

Note: You should have only one of these imports in your source file. Otherwise, you'll get naming conflicts. The `import sttp.tapir.ztapir._` import is meant as a complete replacement of `import sttp.tapir._`.

4.16.1 Server logic

When defining the business logic for an endpoint, the following methods are available, which replace the *standard ones*:

- `def zServerLogic(logic: I => ZIO[R, E, O]): ZServerEndpoint[R, I, E, O]`
- `def zServerLogicPart(logicPart: T => ZIO[R, E, U])`
- `def zServerLogicForCurrent(logicPart: I => ZIO[R, E, U])`

The first defines complete server logic, while the second and third allow defining server logic in parts.

4.16.2 Exposing endpoints using the http4s server

To interpret a `ZServerEndpoint` as a `http4s` server, add the following import:

```
import sttp.tapir.server.http4s.ztapir.ZHttp4sServerInterpreter
```

This adds the following method on `ZEndpoint`:

- `def toRoutes[R](logic: I => ZIO[R, E, O]): HttpRoutes[ZIO[R with Clock, Throwable, *]]`

And the following methods on `ZServerEndpoint` or `List[ZServerEndpoint]`:

- `def toRoutes[R]: HttpRoutes[ZIO[R with Clock, Throwable, *]]`

Note that the resulting `HttpRoutes` always require a clock in their environment.

If you have multiple endpoints with different environmental requirements, the environment must be first widened so that it is uniform across all endpoints, using the `.widen` method:

```
import org.http4s.HttpRoutes
import sttp.tapir.ztapir._
import sttp.tapir.server.http4s.ztapir.ZHttp4sServerInterpreter
import zio.{Has, RIO, ZIO}
import zio.clock.Clock
import zio.interop.catz._

trait Component1
trait Component2
type Service1 = Has[Component1]
type Service2 = Has[Component2]

val serverEndpoint1: ZServerEndpoint[Service1, Unit, Unit, Unit] = ???
val serverEndpoint2: ZServerEndpoint[Service2, Unit, Unit, Unit] = ???

type Env = Service1 with Service2
val routes: HttpRoutes[RIO[Env with Clock, *]] =
  ZHttp4sServerInterpreter.from(List(
    serverEndpoint1.widen[Env],
    serverEndpoint2.widen[Env]
  )).toRoutes // this is where zio-cats interop is needed
```

4.16.3 Example

Three examples of using the ZIO integration are available. The first two showcase basic functionality, while the third shows how to use partial server logic methods:

- ZIO basic example
- ZIO environment example
- ZIO partial server logic example

4.17 Running as an akka-http server

To expose an endpoint as an akka-http server, first add the following dependency:

```
"com.softwaremill.sttp.tapir" %% "tapir-akka-http-server" % "0.17.15"
```

This will transitively pull some Akka modules in version 2.6. If you want to force your own Akka version (for example 2.5), use sbt exclusion. Mind the Scala version in artifact name:

```
"com.softwaremill.sttp.tapir" %% "tapir-akka-http-server" % "0.17.15" exclude("com.
↳ typesafe.akka", "akka-stream_2.12")
```

Now import the object:

```
import sttp.tapir.server.akkahttp.AkkaHttpServerInterpreter
```

The `AkkaHttpServerInterpreter` object contains methods such as: `toRoute`, `toRouteRecoverErrors` and `toDirective`.

4.17.1 Using `toRoute` and `toRouteRecoverErrors`

The `toRoute` method requires the logic of the endpoint to be given as a function of type:

```
I => Future[Either[E, O]]
```

The `toRouteRecoverErrors` method recovers errors from failed futures, and hence requires that `E` is a subclass of `Throwable` (an exception); it expects a function of type `I => Future[O]`.

For example:

```
import sttp.tapir._
import sttp.tapir.server.akkahttp.AkkaHttpServerInterpreter
import scala.concurrent.Future
import akka.http.scaladsl.server.Route

def countCharacters(s: String): Future[Either[Unit, Int]] =
  Future.successful(Right[Unit, Int](s.length))

val countCharactersEndpoint: Endpoint[String, Unit, Int, Any] =
  endpoint.in(stringBody).out(plainBody[Int])

val countCharactersRoute: Route = AkkaHttpServerInterpreter.
  ↪toRoute(countCharactersEndpoint)(countCharacters)
```

Note that the second argument to `toRoute` is a function with one argument, a tuple of type `I`. This means that functions which take multiple arguments need to be converted to a function using a single argument using `.tupled`:

```
import sttp.tapir._
import sttp.tapir.server.akkahttp._
import scala.concurrent.Future
import akka.http.scaladsl.server.Route

def logic(s: String, i: Int): Future[Either[Unit, String]] = ???
val anEndpoint: Endpoint[(String, Int), Unit, String, Any] = ???
val aRoute: Route = AkkaHttpServerInterpreter.toRoute(anEndpoint)((logic _).tupled)
```

4.17.2 Using `toDirective`

The `toDirective` method splits parsing the input and encoding the output. The directive provides the input parameters, type `I`, and a function that can be used to encode the output.

For example:

```
import sttp.tapir._
import sttp.tapir.server.akkahttp.AkkaHttpServerInterpreter
import scala.concurrent.Future
import akka.http.scaladsl.server.Route

def countCharacters(s: String): Future[Either[Unit, Int]] =
  Future.successful(Right[Unit, Int](s.length))

val countCharactersEndpoint: Endpoint[String, Unit, Int, Any] =
  endpoint.in(stringBody).out(plainBody[Int])

val countCharactersRoute: Route = AkkaHttpServerInterpreter.
  ↪toDirective(countCharactersEndpoint).tapply {
```

(continues on next page)

(continued from previous page)

```

    case (input, completion) => completion(countCharacters(input))
  }

```

4.17.3 Combining directives

The tapir-generated `Route/Directive` captures from the request only what is described by the endpoint. Combine with other akka-http directives to add additional behavior, or get more information from the request.

For example, wrap the tapir-generated route in a metrics route, or nest a security directive in the tapir-generated directive.

Edge-case endpoints, which require special logic not expressible using tapir, can be implemented directly using akka-http. For example:

```

import sttp.tapir._
import sttp.tapir.server.akkahttp.AkkaHttpServerInterpreter
import akka.http.scaladsl.server._

case class User(email: String)
def metricsDirective: Directive0 = ???
def securityDirective: Directive1[User] = ???
val tapirEndpoint: Endpoint[String, Unit, Unit, Any] = endpoint.in(path[String]("input"
  ↪))

val myRoute: Route = metricsDirective {
  securityDirective { user =>
    AkkaHttpServerInterpreter.toRoute(tapirEndpoint) { input =>
      ???
      /* here we can use both `user` and `input` values */
    }
  }
}

```

Note that `Routes` can only be nested within other directives. `Directives` can nest in other directives and can also contain nested directives. For example:

```

import akka.http.scaladsl.server.Directive0
import akka.http.scaladsl.server.Directives.Authenticator
import akka.http.scaladsl.server.Directives.authenticateBasic
import akka.http.scaladsl.server.Route
import sttp.tapir._
import sttp.tapir.server.akkahttp.AkkaHttpServerInterpreter

import scala.concurrent.Future

def countCharacters(s: String): Future[Either[Unit, Int]] =
  Future.successful(Right[Unit, Int](s.length))

val countCharactersEndpoint: Endpoint[String, Unit, Int, Any] =
  endpoint.in(stringBody).out(plainBody[Int])

case class User(email: String)
val authenticator: Authenticator[User] = ???
def authorizationDirective(user: User, input: String): Directive0 = ???

```

(continues on next page)

(continued from previous page)

```
val countCharactersRoute: Route =
  authenticateBasic("realm", authenticator) { user =>
    AkkaHttpServerInterpreter.toDirective(countCharactersEndpoint).tapply {
      case (input, completion) =>
        authorizationDirective(user, input) {
          completion(countCharacters(input))
        }
    }
  }
```

4.17.4 Streaming

The akka-http interpreter accepts streaming bodies of type `Source[ByteString, Any]`, as described by the `AkkaStreams` capability. Both response bodies and request bodies can be streamed. Usage: `streamBody(AkkaStreams)(schema, format)`.

The capability can be added to the classpath independently of the interpreter through the `"com.softwaremill.sttp.shared" %% "akka"` dependency.

4.17.5 Web sockets

The interpreter supports web sockets, with pipes of type `Flow[REQ, RESP, Any]`. See [web sockets](#) for more details.

akka-http does not expose control frames (Ping, Pong and Close), so any setting regarding them are discarded, and ping/pong frames which are sent explicitly are ignored. [Automatic pings](#) can be instead enabled through configuration.

4.17.6 Configuration

The interpreter can be configured by providing an implicit `AkkaHttpServerOptions` value and status mappers, see [server options](#) for details.

4.17.7 Defining an endpoint together with the server logic

It's also possible to define an endpoint together with the server logic in a single, more concise step. See [server logic](#) for details.

4.18 Running as an http4s server

To expose an endpoint as an [http4s](#) server, first add the following dependency:

```
"com.softwaremill.sttp.tapir" %% "tapir-http4s-server" % "0.17.15"
```

and import the object:

```
import sttp.tapir.server.http4s.Http4sServerInterpreter
```

This objects contains the `toRoutes` and `toRoutesRecoverErrors` methods. This first requires the logic of the endpoint to be given as a function of type:

```
I => F[Either[E, O]]
```

where `F[_]` is the chosen effect type. The second recovers errors from failed effects, and hence requires that `E` is a subclass of `Throwable` (an exception); it expects a function of type `I => F[O]`. For example:

```
import sttp.tapir._
import sttp.tapir.server.http4s.Http4sServerInterpreter
import cats.effect.IO
import org.http4s.HttpRoutes
import cats.effect.{ContextShift, Timer}

// will probably come from somewhere else
implicit val cs: ContextShift[IO] =
  IO.contextShift(scala.concurrent.ExecutionContext.global)
implicit val t: Timer[IO] =
  IO.timer(scala.concurrent.ExecutionContext.global)

def countCharacters(s: String): IO[Either[Unit, Int]] =
  IO.pure(Right[Unit, Int](s.length))

val countCharactersEndpoint: Endpoint[String, Unit, Int, Any] =
  endpoint.in(stringBody).out(plainBody[Int])
val countCharactersRoutes: HttpRoutes[IO] =
  Http4sServerInterpreter.toRoutes(countCharactersEndpoint)(countCharacters _)
```

Note that the second argument to `toRoute` is a function with one argument, a tuple of type `I`. This means that functions which take multiple arguments need to be converted to a function using a single argument using `.tupled`:

```
import sttp.tapir._
import sttp.tapir.server.http4s.Http4sServerInterpreter
import cats.effect.IO
import org.http4s.HttpRoutes
import cats.effect.{ContextShift, Timer}

// will probably come from somewhere else
implicit val cs: ContextShift[IO] =
  IO.contextShift(scala.concurrent.ExecutionContext.global)
implicit val t: Timer[IO] =
  IO.timer(scala.concurrent.ExecutionContext.global)

def logic(s: String, i: Int): IO[Either[Unit, String]] = ???
val anEndpoint: Endpoint[(String, Int), Unit, String, Any] = ???
val routes: HttpRoutes[IO] = Http4sServerInterpreter.toRoutes(anEndpoint)((logic _).
  ↪tupled)
```

The created `HttpRoutes` are the usual `http4s` Kleisli-based transformation of a Request to a Response, and can be further composed using `http4s` middlewares or request-transforming functions. The tapir-generated `HttpRoutes` captures from the request only what is described by the endpoint.

It's completely feasible that some part of the input is read using a `http4s` wrapper function, which is then composed with the tapir endpoint descriptions. Moreover, “edge-case endpoints”, which require some special logic not expressible using tapir, can be always implemented directly using `http4s`.

4.18.1 Streaming

The `http4s` interpreter accepts streaming bodies of type `Stream[F, Byte]`, as described by the `Fs2Streams` capability. Both response bodies and request bodies can be streamed. Usage:

```
streamBody(Fs2Streams[F])(schema, format).
```

The capability can be added to the classpath independently of the interpreter through the "com.softwaremill.sttp.shared" %% "http4s" dependency.

4.18.2 Web sockets

The interpreter supports web sockets, with pipes of type `Pipe[F, REQ, RESP]`. See [web sockets](#) for more details.

4.18.3 Configuration

The interpreter can be configured by providing an implicit `Http4sServerOptions` value and status mappers, see [server options](#) for details.

The http4s options also includes configuration for the blocking execution context to use, and the io chunk size.

4.18.4 Defining an endpoint together with the server logic

It's also possible to define an endpoint together with the server logic in a single, more concise step. See [server logic](#) for details.

4.19 Running as a Finatra server

To expose an endpoint as an [finatra](#) server, first add the following dependency:

```
"com.softwaremill.sttp.tapir" %% "tapir-finatra-server" % "0.17.15"
```

and import the object:

```
import sttp.tapir.server.finatra.FinatraServerInterpreter
```

or if you would like to use cats-effect project, you can add the following dependency:

```
"com.softwaremill.sttp.tapir" %% "tapir-finatra-server-cats" % "0.17.15"
```

and import the object:

```
import sttp.tapir.server.finatra.cats.FinatraCatsServerInterpreter
```

The interpreter objects contain the `toRoute` and `toRouteRecoverErrors` method. The first one requires the logic of the endpoint to be given as a function of type (note this is a `Twitter Future` or a cats-effect project's `Effect` type):

```
I => Future[Either[E, O]]
```

or giben an effect type that supports cats-effect's `Effect[F]` type, for example, project [Catbird's Rerunnbale](#):

```
I => Rerunnable[Either[E, O]]
```

The second recovers errors from failed futures, and hence requires that `E` is a subclass of `Throwable` (an exception); it expects a function of type `I => Future[O]` or type `I => F[O]` if `F` supports cat-effect's `Effect` type.

For example:

```
import sttp.tapir._
import sttp.tapir.server.finatra.{ FinatraServerInterpreter, FinatraRoute }
import com.twitter.util.Future

def countCharacters(s: String): Future[Either[Unit, Int]] =
  Future.value(Right[Unit, Int](s.length))

val countCharactersEndpoint: Endpoint[String, Unit, Int, Any] =
  endpoint.in(stringBody).out(plainBody[Int])

val countCharactersRoute: FinatraRoute = FinatraServerInterpreter.
  ↪toRoute(countCharactersEndpoint)(countCharacters)
```

or a cats-effect's example:

```
import cats.effect.IO
import sttp.tapir._
import sttp.tapir.server.finatra.FinatraRoute
import sttp.tapir.server.finatra.cats.FinatraCatsServerInterpreter

def countCharacters(s: String): IO[Either[Unit, Int]] =
  IO.pure(Right[Unit, Int](s.length))

val countCharactersEndpoint: Endpoint[String, Unit, Int, Any] =
  endpoint.in(stringBody).out(plainBody[Int])

val countCharactersRoute: FinatraRoute = FinatraCatsServerInterpreter.
  ↪toRoute(countCharactersEndpoint)(countCharacters)
```

Note that the second argument to `toRoute` is a function with one argument, a tuple of type `I`. This means that functions which take multiple arguments need to be converted to a function using a single argument using `.tupled`:

```
import sttp.tapir._
import sttp.tapir.server.finatra.{ FinatraServerInterpreter, FinatraRoute }
import com.twitter.util.Future

def logic(s: String, i: Int): Future[Either[Unit, String]] = ???
val anEndpoint: Endpoint[(String, Int), Unit, String, Any] = ???
val aRoute: FinatraRoute = FinatraServerInterpreter.toRoute(anEndpoint)((logic _).
  ↪tupled)
```

Now that you've created the `FinatraRoute`, add `TapirController` as a trait to your `Controller`. You can then add the created route with `addTapirRoute`.

```
import sttp.tapir.server.finatra._
import com.twitter.finatra.http.Controller

val aRoute: FinatraRoute = ???
class MyController extends Controller with TapirController {
  addTapirRoute(aRoute)
}
```

4.19.1 Defining an endpoint together with the server logic

It's also possible to define an endpoint together with the server logic in a single, more concise step. See [server logic](#) for details.

4.20 Running as a Play server

To expose endpoint as a `play-server` first add the following dependencies:

```
"com.softwaremill.sttp.tapir" %% "tapir-play-server" % "0.17.15"
```

and (if you don't already depend on Play)

```
"com.typesafe.play" %% "play-akka-http-server" % "2.8.7"
```

or

```
"com.typesafe.play" %% "play-netty-server" % "2.8.7"
```

depending on whether you want to use netty or akka based http-server under the hood.

Then import the object:

```
import sttp.tapir.server.play.PlayServerInterpreter
```

This object contains the `toRoute` and `toRoutesRecoverError` methods. This first requires the logic of the endpoint to be given as a function of type:

```
I => Future[Either[E, O]]
```

The second recovers errors from failed effects, and hence requires that `E` is a subclass of `Throwable` (an exception); it expects a function of type `I => Future[O]`. For example:

```
import sttp.tapir._
import sttp.tapir.server.play.PlayServerInterpreter
import scala.concurrent.ExecutionContext.Implicits.global
import scala.concurrent.Future
import akka.stream.Materializer
import play.api.routing.Router.Routes

implicit val materializer: Materializer = ???

def countCharacters(s: String): Future[Either[Unit, Int]] =
  Future(Right[Unit, Int](s.length))

val countCharactersEndpoint: Endpoint[String, Unit, Int, Any] =
  endpoint.in(stringBody).out(plainBody[Int])
val countCharactersRoutes: Routes =
  PlayServerInterpreter.toRoute(countCharactersEndpoint)(countCharacters _)
```

Note that the second argument to `toRoute` is a function with one argument, a tuple of type `I`. This means that functions which take multiple arguments need to be converted to a function using a single argument using `.tupled`:

```
import sttp.tapir._
import sttp.tapir.server.play._
import scala.concurrent.ExecutionContext.Implicits.global
import scala.concurrent.Future
import akka.stream.Materializer
import play.api.routing.Router.Routes

implicit val materializer: Materializer = ???
```

(continues on next page)

(continued from previous page)

```
def logic(s: String, i: Int): Future[Either[Unit, String]] = ???
val anEndpoint: Endpoint[(String, Int), Unit, String, Any] = ???
val aRoute: Routes = PlayServerInterpreter.toRoute(anEndpoint)((logic _).tupled)
```

4.20.1 Bind the routes

Creating the HTTP server manually

An HTTP server can then be started as in the following example:

```
import play.core.server._
import play.api.routing.Router.Routes

val aRoute: Routes = ???

object Main {
  // JVM entry point that starts the HTTP server
  def main(args: Array[String]): Unit = {
    val playConfig = ServerConfig(port =
      sys.props.get("http.port").map(_.toInt).orElse(Some(9000))
    )
    NettyServer.fromRouterWithComponents(playConfig) { components =>
      aRoute
    }
  }
}
```

As part of an existing Play application

Or, if you already have an existing Play application, you can create a Router class and bind it to the application.

First, add a line like following in the routes files:

```
-> /api api.ApiRouter
```

Then create a class like this:

```
class ApiRouter @Inject() () extends SimpleRouter {
  override def routes: Routes = {
    // Routes generated by Tapir (see above)
    aRoute
    .orElse(anotherRoute)
  }
}
```

Find more details about how to bind a Router to your application in the [Play framework documentation](#).

4.20.2 Configuration

The interpreter can be configured by providing an implicit PlayServerOptions value and status mappers, see [server options](#) for details.

4.20.3 Defining an endpoint together with the server logic

It's also possible to define an endpoint together with the server logic in a single, more concise step. See [server logic](#) for details.

4.21 Running as a Vert.X server

Endpoints can be mounted as Vert.x Routes on top of a Vert.x Router.

Vert.x interpreter can be used with different effect systems (cats-effect, ZIO) as well as Scala's standard Future.

4.21.1 Scala's standard Future

Add the following dependency

```
"com.softwaremill.sttp.tapir" %% "tapir-vertx-server" % "0.17.15"
```

to use this interpreter with Future.

Then import the object:

```
import sttp.tapir.server.vertx.VertxFutureServerInterpreter._
```

This object contains the following methods:

- `route(e: Endpoint[I, E, O, Any])(logic: I => Future[Either[E, O]]):` returns a function `Router => Route` that will create a route matching the endpoint definition, and attach your logic as an handler. Errors will be recovered automatically (but generically)
- `routeRecoverErrors(e: Endpoint[I, E, O, Any])(logic: I => Future[O]):` returns a function `Router => Route` that will create a route matching the endpoint definition, and attach your logic as an handler. You're providing your own way to deal with errors happening in the logic function.
- `blockingRoute(e: Endpoint[I, E, O, Any])(logic: I => Future[Either[E, O]]):` returns a function `Router => Route` that will create a route matching the endpoint definition, and attach your logic as an blocking handler. Errors will be recovered automatically (but generically)
- `blockingRouteRecoverErrors(e: Endpoint[I, E, O, Any])(logic: I => Future[O]):` returns a function `Router => Route` that will create a route matching the endpoint definition, and attach your logic as a blocking handler. You're providing your own way to deal with errors happening in the logic function.

The methods recovering errors from failed effects, require E to be a subclass of `Throwable` (an exception); and expect a function of type `I => Future[O]`.

Note that the second argument to `route` etc. is a function with one argument, a tuple of type `I`. This means that functions which take multiple arguments need to be converted to a function using a single argument using `.tupled`:

```
import sttp.tapir._
import sttp.tapir.server.vertx.{ VertxFutureEndpointOptions,
  ↪VertxFutureServerInterpreter }
import io.vertx.ext.web._
import scala.concurrent.Future

implicit val options: VertxFutureEndpointOptions = ???
def logic(s: String, i: Int): Future[Either[Unit, String]] = ???
```

(continues on next page)

(continued from previous page)

```
val anEndpoint: Endpoint[(String, Int), Unit, String, Any] = ???
val aRoute: Router => Route = VortexFutureServerInterpreter.route(anEndpoint)((logic _
↳).tupled)
```

In practice, routes will be mounted on a router, this router can then be used as a request handler for your http server. An HTTP server can then be started as in the following example:

```
import sttp.tapir._
import sttp.tapir.server.vertx.VertxFutureEndpointOptions
import sttp.tapir.server.vertx.VertxFutureServerInterpreter._
import io.vertx.core.Vertx
import io.vertx.ext.web._
import scala.concurrent.{Await, Future}
import scala.concurrent.duration._

object Main {
  // JVM entry point that starts the HTTP server
  def main(args: Array[String]): Unit = {
    implicit val options: VertxFutureEndpointOptions = ???
    val vertx = Vertx.vertx()
    val server = vertx.createHttpServer()
    val router = Router.router(vertx)
    val anEndpoint: Endpoint[(String, Int), Unit, String, Any] = ??? // your_
↳definition here
    def logic(s: String, i: Int): Future[Either[Unit, String]] = ??? // your logic_
↳here
    val attach = route(anEndpoint)((logic _).tupled)
    attach(router) // your endpoint is now attached to the router, and the route has_
↳been created
    Await.result(server.requestHandler(router).listen(9000).asScala, Duration.Inf)
  }
}
```

4.21.2 Configuration

Every endpoint can be configured by providing an implicit `VertxFutureEndpointOptions`, see [server options](#) for details. You can also provide your own `ExecutionContext` to execute the logic.

4.21.3 Defining an endpoint together with the server logic

It's also possible to define an endpoint together with the server logic in a single, more concise step. See [server logic](#) for details.

4.21.4 Cats Effect typeclasses

Add the following dependency

```
"com.softwaremill.sttp.tapir" %% "tapir-vertx-server" % "0.17.15"
"com.softwaremill.sttp.shared" %% "fs2" % "LatestVersion"
```

to use this interpreter with Cats Effect typeclasses.

Then import the object:

```
import sttp.tapir.server.vertx.VertxCatsServerInterpreter._
```

This object contains the following methods:

- `route[F[_], I, E, O](e: Endpoint[I, E, O, Fs2Streams[F]])(logic: I => F[Either[E, O]])` returns a function `Router => Route` that will create a route matching the endpoint definition, and attach your logic as an handler. Errors will be recovered automatically.
- `routeRecoverErrors[F[_], I, E, O](e: Endpoint[I, E, O, Fs2Streams[F]])(logic: I => F[O])` returns a function `Router => Route` that will create a route matching the endpoint definition, and attach your logic as an handler. You're providing your own way to deal with errors happening in the `logic` function.

Here is simple example which starts HTTP server with one route:

```
import cats.effect._
import cats.syntax.flatMap._
import io.vertx.core.Vertx
import io.vertx.ext.web.Router
import sttp.tapir._
import sttp.tapir.server.vertx.VertxCatsServerInterpreter._
import sttp.tapir.server.vertx.VertxEffectfulEndpointOptions

object App extends IOApp {

  implicit val opts = VertxEffectfulEndpointOptions()

  val responseEndpoint =
    endpoint
      .in("response")
      .in(query[String]("key"))
      .out(plainBody[String])

  def handler(req: String): IO[Either[Unit, String]] =
    IO.pure(Right(req))

  val attach = route(responseEndpoint)(handler)

  override def run(args: List[String]): IO[ExitCode] =
    Resource.make(IO.delay{
      val vertx = Vertx.vertx()
      val server = vertx.createHttpServer()
      val router = Router.router(vertx)
      attach(router)
      server.requestHandler(router).listen(8080)
    } >>= (_.liftF[IO]))({ server =>
      IO.delay(server.close) >>= (_.liftF[IO].void)
    }).use(_ => IO.never)
}
```

This interpreter also supports streaming using FS2 streams:

```
import cats.effect._
import fs2._
import sttp.capabilities.fs2.Fs2Streams
import sttp.tapir._
import sttp.tapir.server.vertx.VertxCatsServerInterpreter._
import sttp.tapir.server.vertx.VertxEffectfulEndpointOptions
```

(continues on next page)

(continued from previous page)

```

implicit val effect: ConcurrentEffect[IO] = ???

implicit val opts = VertxEffectfulEndpointOptions()

val streamedResponse =
  endpoint
    .in("stream")
    .in(query[Int]("key"))
    .out(streamBody(Fs2Streams[IO])(Schema.string, CodecFormat.TextPlain()))

val attach = route(streamedResponse) { key =>
  IO.pure(Right(Stream.chunk(Chunk.array("Hello world!".getBytes)).repeatN(key)))
}

```

4.21.5 ZIO

Add the following dependency

```

"com.softwaremill.sttp.tapir" %% "tapir-vertx-server" % "0.17.15"
"com.softwaremill.sttp.shared" %% "zio" % "LatestVersion"

```

to use this interpreter with ZIO.

Then import the object:

```

import sttp.tapir.server.vertx.VertxZioServerInterpreter._

```

This object contains method `route[R, I, E, O](e: Endpoint[I, E, O, ZioStreams])(logic: I => ZIO[R, E, O])` which returns a function `Router => Route` that will create a route matching the endpoint definition, and attach your logic as an handler.

Here is simple example which starts HTTP server with one route:

```

import io.vertx.core.Vertx
import io.vertx.ext.web.Router
import sttp.tapir._
import sttp.tapir.server.vertx.VertxZioServerInterpreter._
import sttp.tapir.server.vertx.VertxEffectfulEndpointOptions

import zio._

object Short extends zio.App {

  implicit val opts = VertxEffectfulEndpointOptions()

  implicit val runtime = Runtime.default

  val responseEndpoint =
    endpoint
      .in("response")
      .in(query[String]("key"))
      .out(plainBody[String])

  val attach = route(responseEndpoint) { key => UIO.succeed(key) }

```

(continues on next page)

(continued from previous page)

```

override def run(args: List[String]): URIO[ZEnv, ExitCode] =
  ZManaged.make(ZIO.effect {
    val vertx = Vertx.vertx()
    val server = vertx.createHttpServer()
    val router = Router.router(vertx)
    attach(router)
    server.requestHandler(router).listen(8080)
  } >=> (_.asTask)) ({ server =>
    ZIO.effect(server.close()).flatMap(_.asTask).orDie
  }).useForever.as(ExitCode.success).orDie
}

```

This interpreter supports streaming using ZStreams.

4.22 Server options

Each interpreter accepts an implicit options value, which contains configuration values for:

- how to create a file (when receiving a response that is mapped to a file, or when reading a file-mapped multipart part)
- how to handle decode failures (see also [error handling](#))
- debug-logging of request handling

To customise the server options, define an implicit value, which will be visible when converting an endpoint or multiple endpoints to a route/routes. For example, for AkkaHttpServerOptions:

```

import sttp.tapir.server.akkahttp.AkkaHttpServerOptions

implicit val customServerOptions: AkkaHttpServerOptions =
  AkkaHttpServerOptions.default.copy(decodeFailureHandler = ???)

```

4.23 Server logic

The logic that should be run when an endpoint is invoked can be passed when interpreting the endpoint as a server, and converting to a server-implementation-specific route. However, there's also the option to define an endpoint coupled with the server logic - either given entirely or gradually, in parts. This might make it easier to work with endpoints and their collections in a server setting.

4.23.1 Defining an endpoint together with the server logic

It's possible to combine an endpoint description with the server logic in a single object, `ServerEndpoint[I, E, O, S, F]`. Such an endpoint contains not only an endpoint of type `Endpoint[I, E, O, S]`, but also a logic function `I => F[Either[E, O]]`, for some effect `F`.

The book example can be more concisely written as follows:

```

import sttp.tapir._
import sttp.tapir.server._
import sttp.tapir.server.akkahttp.AkkaHttpServerInterpreter

```

(continues on next page)

(continued from previous page)

```
import scala.concurrent.Future
import akka.http.scaladsl.server.Route

val countCharactersServerEndpoint: ServerEndpoint[String, Unit, Int, Any, Future] =
  endpoint.in(stringBody).out(plainBody[Int]).serverLogic { s =>
    Future.successful[Either[Unit, Int]](Right(s.length))
  }

val countCharactersRoute: Route = AkkaHttpServerInterpreter.
  ↪toRoute(countCharactersServerEndpoint)
```

A `ServerEndpoint` can then be converted to a route using `.toRoute/``.toRoutes` methods (without any additional parameters; the exact method name depends on the server interpreter), or to documentation.

Moreover, a list of server endpoints can be converted to routes or documentation as well:

```
import sttp.tapir._
import sttp.tapir.server.akkahttp.AkkaHttpServerInterpreter
import scala.concurrent.Future
import akka.http.scaladsl.server.Route

val endpoint1 = endpoint.in("hello").out(stringBody)
  .serverLogic { _ => Future.successful[Either[Unit, String]](Right("world")) }

val endpoint2 = endpoint.in("ping").out(stringBody)
  .serverLogic { _ => Future.successful[Either[Unit, String]](Right("pong")) }

val route: Route = AkkaHttpServerInterpreter.toRoute(List(endpoint1, endpoint2))
```

Note that when dealing with endpoints which have multiple input parameters, the server logic function is a function of a *single* argument, which is a tuple; hence you'll need to pattern-match using `case` to extract the parameters:

```
import sttp.tapir._
import scala.concurrent.Future

val echoEndpoint = endpoint
  .in(query[Int]("count"))
  .in(stringBody)
  .out(stringBody)
  .serverLogic { case (count, body) =>
    Future.successful[Either[Unit, String]](Right(body * count))
  }
```

Recovering errors from failed effects

If your error type is an exception (extends `Throwable`), and if errors that occur in the server logic are represented as failed effects, you can use a variant of the methods above, which extract the error from the failed effect, and respond with the error output appropriately.

This can be done with the `serverLogicRecoverErrors(f: E => F[O])` method. Note that the `E` type parameter isn't directly present here; however, the method also contains a requirement that `E` is an exception, and will only recover errors which are subtypes of `E`. Any others will be propagated without changes.

For example:

```
import sttp.tapir._
import scala.concurrent.Future

case class MyError(msg: String) extends Exception
val testEndpoint = endpoint
  .in(query[Boolean] ("fail"))
  .errorOut(stringBody.map(MyError) (_.msg))
  .out(stringBody)
  .serverLogicRecoverErrors { fail =>
    if (fail) {
      Future.successful("OK") // note: no Right() wrapper
    } else {
      Future.failed(new MyError("Not OK")) // no Left() wrapper, a failed future
    }
  }
}
```

4.23.2 Extracting common route logic

Quite often, especially for *authentication*, some part of the route logic is shared among multiple endpoints. However, these functions don't compose in a straightforward way, as authentication usually operates on a single input, which is only a part of the whole logic's input. That's why there are two options for providing the server logic in parts.

Defining an extendable, base endpoint with partial server logic

First, you might want to define a base endpoint, with some part of the server logic already defined, and later extend it with additional inputs/outputs, and successive logic parts.

When using this method, note that the error type must be fully defined upfront, and cannot be extended later. That's because all of the partial logic functions can return either an error (of the given type), or a partial result.

To define partial logic for the inputs defined so far, you should use the `serverLogicForCurrent` method on an endpoint. This accepts a method, `f: I => F[Either[E, U]]`, which given the entire (current) input defined so far, returns either an error, or a partial result.

For example, we can create a partial server endpoint given an authentication function, and an endpoint describing the authentication inputs:

```
import sttp.tapir._
import sttp.tapir.server._
import scala.concurrent.Future

implicit val ec = scala.concurrent.ExecutionContext.global

case class User(name: String)
def auth(token: String): Future[Either[Int, User]] = Future {
  if (token == "secret") Right(User("Spock"))
  else Left(1001) // error code
}

val secureEndpoint: PartialServerEndpoint[User, Unit, Int, Unit, Any, Future] =
  ↪ endpoint
    .in(header[String] ("X-AUTH-TOKEN"))
    .errorOut(plainBody[Int])
    .serverLogicForCurrent(auth)
```

The result is a value of type `PartialServerEndpoint`, which can be extended with further inputs and outputs, just as a normal endpoint (except for error outputs, which are fixed).

Successive logic parts (again consuming the entire input defined so far, and producing another partial result) can be given by calling `serverLogicForCurrent` again. The logic can be completed - given a tuple of partial results, and unconsumed inputs - using the `serverLogic` method:

```
val secureHelloWorld1WithLogic = secureEndpoint.get
    .in("hello1")
    .in(query[String]("salutation"))
    .out(stringBody)
    .serverLogic { case (user, salutation) =>
        Future.successful[Either[Int, String]] (Right(s"${salutation}, ${user.name}!"))
    }
```

Once the server logic is complete, a `ServerEndpoint` is returned, which can be then interpreted as a server.

The methods mentioned also have variants which recover errors from failed effects (as described above), using `serverLogicForCurrentRecoverErrors` and `serverLogicRecoverErrors`.

Providing server logic in parts, for an already defined endpoint

Secondly, you might already have the entire endpoint defined upfront (e.g. in a separate module, which doesn't know anything about the server logic), and would like to provide the server logic in parts.

This can be done using the `serverLogicPart` method, which takes a function of type `f: T => F[Either[E, U]]` as a parameter. Here `T` is some prefix of the endpoint's input `I`. The function then consumes some part of the input, and produces an error, or a partial result.

Unlike previously, here the endpoint is considered complete, and cannot be later extended: no additional inputs or outputs can be defined.

Note: Note that the partial logic functions should be fully typed, to help with inference. It might not be possible for the compiler to infer, which part of the input should be consumed.

For example, if we have an endpoint:

```
import sttp.tapir._

val secureHelloWorld2: Endpoint[(String, String), Int, String, Any] = endpoint
    .in(header[String]("X-AUTH-TOKEN"))
    .errorOut(plainBody[Int])
    .get
    .in("hello2")
    .in(query[String]("salutation"))
    .out(stringBody)
```

We can provide the server logic in parts (using the same `auth` method as above):

```
import scala.concurrent.Future

case class User(name: String)
def auth(token: String): Future[Either[Int, User]] = ???

val secureHelloWorld2WithLogic = secureHelloWorld2
    .serverLogicPart(auth)
```

(continues on next page)

(continued from previous page)

```
.andThen { case (user, salutation) =>
  Future.successful[Either[Int, String]] (Right(s"$salutation, ${user.name}!"))
}
```

Here, we first define a single part using `serverLogicPart`, and then complete the logic (consuming the remaining inputs) using `andThen`. The result is, like previously, a `ServerEndpoint`, which can be interpreted as a server.

Multiple parts can be provided using `andThenPart` invocations (consuming successive input parts). There are also variants of the methods, which recover errors from failed effects: `serverLogicPartRecoverErrors`, `andThenRecoverErrors` and `andThenPartRecoverErrors`.

4.23.3 Status codes

By default, successful responses are returned with the 200 OK status code, and errors with 400 Bad Request. However, this can be customised by specifying how an *output maps to the status code*.

4.24 Error handling

4.24.1 Exception handling

There's no exception handling built into tapir. However, tapir contains a more general error handling mechanism, as the endpoints can contain dedicated error outputs.

If the logic function, which is passed to the server interpreter, fails (i.e. throws an exception, which results in a failed `Future` or `IO/Task`), this is propagated to the library (akka-http or http4s).

However, any exceptions can be recovered from and mapped to an error value. For example:

```
import sttp.tapir._
import sttp.tapir.server.akkahttp.AkkaHttpServerInterpreter
import scala.concurrent.Future
import scala.util._

implicit val ec = scala.concurrent.ExecutionContext.global
type ErrorInfo = String
def logic(s: String): Future[Int] = ???

def handleErrors[T](f: Future[T]): Future[Either[ErrorInfo, T]] =
  f.transform {
    case Success(v) => Success(Right(v))
    case Failure(e) =>
      println(s"Exception when running endpoint logic: $e")
      Success(Left(e.getMessage))
  }

AkkaHttpServerInterpreter.toRoute(
  endpoint
    .errorOut(plainBody[ErrorInfo])
    .out(plainBody[Int])
    .in(query[String]("name")) {
    (logic _).andThen(handleErrors)
  }
)
```

In the above example, errors are represented as `Strings` (aliased to `ErrorInfo` for readability). When the logic completes successfully an `Int` is returned. Any exceptions that are raised are logged, and represented as a value of type `ErrorInfo`.

Following the convention, the left side of the `Either[ErrorInfo, T]` represents an error, and the right side success.

Alternatively, errors can be recovered from failed effects and mapped to the error output - provided that the `E` type in the endpoint description is itself a subclass of exception. This can be done using the `toRouteRecoverErrors` method.

4.24.2 Handling decode failures

Quite often user input will be malformed and decoding will fail. Should the request be completed with a `400 Bad Request` response, or should the request be forwarded to another endpoint? By default, tapir follows OpenAPI conventions, that an endpoint is uniquely identified by the method and served path. That's why:

- an “endpoint doesn't match” result is returned if the request method or path doesn't match. The http library should attempt to serve this request with the next endpoint. The path doesn't match if a path segment is missing, there's a constant value mismatch or a decoding error (e.g. parsing a segment to an `Int` fails)
- otherwise, we assume that this is the correct endpoint to serve the request, but the parameters are somehow malformed. A `400 Bad Request` response is returned if a query parameter, header or body causes any decode failure, or if the decoding a path capture causes a validation error.

This can be customised by providing an implicit instance of `sttp.tapir.server.DecodeFailureHandler`, which basing on the request, failing input and failure description can decide, whether to return a “no match” or a specific response.

Only the first failure is passed to the `DecodeFailureHandler`. Inputs are decoded in the following order: method, path, query, header, body.

Note that the decode failure handler is used **only** for failures that occur during decoding of path, query, body and header parameters - while invoking `Codec.decode`. It does not handle any failures or exceptions that occur when invoking the logic of the endpoint.

Default failure handler

The default decode failure handler is a case class, consisting of functions which decide whether to respond with an error or return a “no match”, create error messages and create the response.

To reuse the existing default logic, parts of the default behavior can be swapped, e.g. to return responses in a different format (other than textual):

```
import sttp.tapir._
import sttp.tapir.server._
import sttp.tapir.server.akkahttp.{ AkkaHttpServerInterpreter, AkkaHttpServerOptions }
import sttp.tapir.generic.auto._
import sttp.tapir.json.circe._
import sttp.model.StatusCode
import io.circe.generic.auto._

implicit val ec = scala.concurrent.ExecutionContext.global
case class MyFailure(msg: String)
def myFailureResponse(response: DefaultDecodeFailureResponse, message: String): _
  ↳ DecodeFailureHandling =
  DecodeFailureHandling.response(ServerDefaults.failureOutput(jsonBody[MyFailure])) (
```

(continues on next page)

(continued from previous page)

```

    (response, MyFailure(message))
  )

  val myDecodeFailureHandler = ServerDefaults.decodeFailureHandler.copy(
    response = myFailureResponse
  )

  implicit val myServerOptions: AkkaHttpServerOptions = AkkaHttpServerOptions.default
    .copy(
      decodeFailureHandler = myDecodeFailureHandler
    )

```

Note that when specifying that a response should be returned upon a failure, we need to provide the endpoint output which should be used to create the response, as well as a value for this output.

The default decode failure handler also has the option to return a 400 Bad Request, instead of a no-match (ultimately leading to a 404 Not Found), when the “shape” of the path matches (that is, the number of segments in the request and endpoint’s paths are the same), but when decoding some part of the path ends in an error. See the `badRequestOnPathErrorIfPathShapeMatches` in `ServerDefaults`.

Finally, you can provide custom error messages for validation errors (which optionally describe what failed) and failure errors (which describe the source of the error).

A completely custom implementation of the `DecodeFailureHandler` function can also be used.

4.25 Debugging

When dealing with multiple endpoints, how to find out which endpoint handled a request, or why an endpoint didn’t handle a request?

For this purpose, tapir provides optional logging. The logging options (and messages) can be customised by changing the default `LogRequestHandling` class, which is part of `server options`.

The following can be customised:

1. `logWhenHandled`: when a request is handled by an endpoint, or when the inputs can’t be decoded, and the decode failure maps to a response (default: `true`, `DEBUG` log)
2. `logAllDecodeFailures`: when the inputs can’t be decoded, and the decode failure doesn’t map to a response (the next endpoint will be tried; default: `false`, `DEBUG` log)
3. `logLogicExceptions`: when there’s an exception during evaluation of the server logic (default: `true`, `ERROR` log)

Logging all decode failures (2) might be helpful when debugging, but can also produce a large amount of logs, hence it’s disabled by default.

Even if logging for a particular category (as described above) is set to `true`, normal logger rules apply - if you don’t see the logs, please verify your logging levels for the appropriate packages.

4.26 Using as an sttp client

Add the dependency:

```
"com.softwaremill.sttp.tapir" %% "tapir-sttp-client" % "0.17.15"
```

To make requests using an endpoint definition using the `sttp client`, import:

```
import sttp.tapir.client.sttp.SttpClientInterpreter
```

This object contains two methods:

- `toRequestUnsafe(Endpoint, Uri)`: given the base URI returns a function, which might throw an exception if decoding of the result fails

```
I => Request[Either[E, O], Any]
```

- `toRequest(Endpoint, Uri)`: given the base URI returns a function, which represents decoding errors as the `DecodeResult` class

```
I => Request[DecodeResult[Either[E, O]], Any]
```

Note that the returned functions have one argument: the input values of the endpoint. This might be a single type, a tuple, or a case class, depending on the endpoint description.

After providing the input parameters, a description of the request to be made is returned, with the input value encoded as appropriate request parameters: path, query, headers and body. This can be further customised and sent using any `sttp` backend. The response will then contain the decoded error or success values (note that this can be the body enriched with data from headers/status code).

See the [runnable example](#) for example usage.

4.26.1 Web sockets

To interpret a web socket endpoint, an additional streams-specific import is needed, so that the interpreter can convert `sttp`'s `WebSocket` instance into a pipe. This logic is looked up via the `WebSocketToPipe` implicit.

The required imports are as follows:

```
import sttp.tapir.client.sttp.ws.akkahttp._ // for akka-streams // for akka-streams
import sttp.tapir.client.sttp.ws.fs2._      // for fs2
```

No additional dependencies are needed, as both of the above implementations are included in the main interpreter, with dependencies on `akka-streams` and `fs2` being marked as optional (hence these are not transitive).

4.26.2 Scala.JS

The client interpreter also supports `Scala.JS`, the request must then be sent using the `sttp client Scala.JS Fetch backend`.

4.27 Using as a Play client

Add the dependency:

```
"com.softwaremill.sttp.tapir" %% "tapir-play-client" % "0.17.15"
```

To make requests using an endpoint definition using the `play client`, import:

```
import sttp.tapir.client.play.PlayClientInterpreter
```

This objects contains two methods:

- `toRequestUnsafe(Endpoint, String)`: given the base URI returns a function, which will generate a request and a response parser which might throw an exception when decoding of the result fails

```
I => (StandaloneWSRequest, StandaloneWSResponse => Either[E, O])
```

- `toRequest(Endpoint, String)`: given the base URI returns a function, which will generate a request and a response parser which represents decoding errors as the `DecodeResult` class

```
I => (StandaloneWSRequest, StandaloneWSResponse => DecodeResult[Either[E, O]])
```

Note that the returned functions have one-argument: the input values of end endpoint. This might be a single type, a tuple, or a case class, depending on the endpoint description.

After providing the input parameters, the two following are returned:

- a description of the request to be made, with the input value encoded as appropriate request parameters: path, query, headers and body. This can be further customised and sent using regular Play methods.
- a response parser to be applied to the response got after executing the request. The result will then contain the decoded error or success values (note that this can be the body enriched with data from headers/status code).

Example:

```
import sttp.tapir._
import sttp.tapir.client.play.PlayClientInterpreter
import sttp.capabilities.akka.AkkaStreams

import scala.concurrent.ExecutionContext.Implicits.global
import scala.concurrent.Future

import play.api.libs.ws.StandaloneWSClient

def example[I, E, O, R >: AkkaStreams](implicit wsClient: StandaloneWSClient) {
  val e: Endpoint[I, E, O, R] = ???
  val inputArgs: I = ???

  val (req, responseParser) = PlayClientInterpreter
    .toRequestUnsafe(e, s"http://localhost:9000")
    .apply(inputArgs)

  val result: Future[Either[E, O]] = req
    .execute()
    .map(responseParser)
}
```

4.27.1 Limitations

Multipart requests are not supported.

Streaming capabilities:

- only `AkkaStreams` is supported

4.28 Generating OpenAPI documentation

To use, add the following dependencies:

```
"com.softwaremill.sttp.tapir" %% "tapir-openapi-docs" % "0.17.15"
"com.softwaremill.sttp.tapir" %% "tapir-openapi-circe-yaml" % "0.17.15"
```

Tapir contains a case class-based model of the openapi data structures in the `openapi/openapi-model` subproject (the model is independent from all other tapir modules and can be used stand-alone).

An endpoint can be converted to an instance of the model by importing the `sttp.tapir.docs.openapi.OpenAPIDocsInterpreter` object:

```
import sttp.tapir._
import sttp.tapir.openapi.OpenAPI
import sttp.tapir.docs.openapi.OpenAPIDocsInterpreter

val booksListing = endpoint.in(path[String] ("bookId"))

val docs: OpenAPI = OpenAPIDocsInterpreter.toOpenAPI(booksListing, "My Bookshop", "1.0"
↪)
```

Such a model can then be refined, by adding details which are not auto-generated. Working with a deeply nested case class structure such as the `OpenAPI` one can be made easier by using a lens library, e.g. [Quicklens](#).

The documentation is generated in a large part basing on [schemas](#). Schemas can be [automatically derived and customised](#).

Quite often, you'll need to define the servers, through which the API can be reached. To do this, you can modify the returned `OpenAPI` case class either directly or by using a helper method:

```
import sttp.tapir.openapi.Server

val docsWithServers: OpenAPI = OpenAPIDocsInterpreter.toOpenAPI(booksListing, "My_
↪Bookshop", "1.0")
  .servers(List(Server("https://api.example.com/v1").description("Production server
↪")))
```

Multiple endpoints can be converted to an `OpenAPI` instance by calling the method on a list of endpoints:

```
OpenAPIDocsInterpreter.toOpenAPI(List(addBook, booksListing, booksListingByGenre),
↪ "My Bookshop", "1.0")
```

The `openapi` case classes can then be serialised to YAML using [Circe](#):

```
import sttp.tapir.openapi.circe.yaml._

println(docs.toYaml)
```

Or to JSON:

```
import io.circe.Printer
import io.circe.syntax._
import sttp.tapir.openapi.circe._

println(Printer.spaces2.print(docs.asJson))
```

4.28.1 Options

Options can be customised by providing an implicit instance of `OpenAPIDocsOptions`, when calling `.toOpenAPI`.

- `operationIdGenerator`: each endpoint corresponds to an operation in the OpenAPI format and should have a unique operation id. By default, the name of endpoint is used as the operation id, and if this is not available, the operation id is auto-generated by concatenating (using camel-case) the request method and path.

4.28.2 Exposing OpenAPI documentation

Exposing the OpenAPI documentation can be very application-specific. However, tapir contains modules which contain akka-http/http4s routes for exposing documentation using [Swagger UI](#) or [Redoc](#):

```
// Akka HTTP
"com.softwaremill.sttp.tapir" %% "tapir-swagger-ui-akka-http" % "0.17.15"
"com.softwaremill.sttp.tapir" %% "tapir-redoc-akka-http" % "0.17.15"

// Finatra
"com.softwaremill.sttp.tapir" %% "tapir-swagger-ui-finatra" % "0.17.15"

// HTTP4S
"com.softwaremill.sttp.tapir" %% "tapir-swagger-ui-http4s" % "0.17.15"
"com.softwaremill.sttp.tapir" %% "tapir-redoc-http4s" % "0.17.15"

// Play
"com.softwaremill.sttp.tapir" %% "tapir-swagger-ui-play" % "0.17.15"
"com.softwaremill.sttp.tapir" %% "tapir-redoc-play" % "0.17.15"

// Vert.x
"com.softwaremill.sttp.tapir" %% "tapir-swagger-ui-vertx" % "0.17.15"
```

Note: `tapir-swagger-ui-akka-http` transitively pulls some Akka modules in version 2.6. If you want to force your own Akka version (for example 2.5), use `sbt exclusion`. Mind the Scala version in artifact name:

```
"com.softwaremill.sttp.tapir" %% "tapir-swagger-ui-akka-http" % "0.17.15" exclude(
  ↪ "com.typesafe.akka", "akka-stream_2.12")
```

Usage example for akka-http:

```
import sttp.tapir._
import sttp.tapir.docs.openapi.OpenAPIDocsInterpreter
import sttp.tapir.openapi.circe.yaml._
import sttp.tapir.swagger.akkahttp.SwaggerAkka

val myEndpoints: Seq[Endpoint[_, _, _, _]] = ???
val docsAsYaml: String = OpenAPIDocsInterpreter.toOpenAPI(myEndpoints, "My App", "1.0")
  ↪).toYaml
// add to your akka routes
new SwaggerAkka(docsAsYaml).routes
```

For redoc, use `RedocAkkaHttp`.

For http4s, use the `SwaggerHttp4s` or `RedocHttp4s` classes.

For Play, use `SwaggerPlay` or `RedocPlay` classes.

For Vert.x, use `SwaggerVertx` class.

Using with sbt-assembly

The `tapir-swagger-ui-*` modules rely on a file in the `META-INF` directory tree, to determine the version of the Swagger UI. You need to take additional measures if you package your application with `sbt-assembly` because the default merge strategy of the `assembly` task discards most artifacts in that directory. To avoid a `NullPointerException`, you need to include the following file explicitly:

```
assemblyMergeStrategy in assembly := {
  case PathList("META-INF", "maven", "org.webjars", "swagger-ui", "pom.properties") =>
    MergeStrategy.singleOnError
  case x =>
    val oldStrategy = (assemblyMergeStrategy in assembly).value
    oldStrategy(x)
}
```

4.29 Generating AsyncAPI documentation

To use, add the following dependencies:

```
"com.softwaremill.sttp.tapir" %% "tapir-asyncapi-docs" % "0.17.15"
"com.softwaremill.sttp.tapir" %% "tapir-asyncapi-circe-yaml" % "0.17.15"
```

Tapir contains a case class-based model of the asyncapi data structures in the `asyncapi/asyncapi-model` sub-project (the model is independent from all other tapir modules and can be used stand-alone).

An endpoint can be converted to an instance of the model by using the `sttp.tapir.docs.asyncapi.AsyncAPIInterpreter` object:

```
import sttp.capabilities.akka.AkkaStreams
import sttp.tapir._
import sttp.tapir.asyncapi.AsyncAPI
import sttp.tapir.docs.asyncapi.AsyncAPIInterpreter
import sttp.tapir.generic.auto._
import sttp.tapir.json.circe._
import io.circe.generic.auto._

case class Response(msg: String, count: Int)
val echoWS = endpoint.out(
  websocketBody[String, CodecFormat.TextPlain, Response, CodecFormat.
    ↪Json](AkkaStreams))

val docs: AsyncAPI = AsyncAPIInterpreter.toAsyncAPI(echoWS, "Echo web socket", "1.0")
```

Such a model can then be refined, by adding details which are not auto-generated. Working with a deeply nested case class structure such as the `AsyncAPI` one can be made easier by using a lens library, e.g. [Quicklens](#).

The documentation is generated in a large part basing on [schemas](#). Schemas can be [automatically derived](#) and [customised](#).

Quite often, you'll need to define the servers, through which the API can be reached. Any servers provided to the `.toAsyncAPI` invocation will be supplemented with security requirements, as specified by the endpoints:

```
import sttp.tapir.asyncapi.Server

val docsWithServers: AsyncAPI = AsyncAPIInterpreter.toAsyncAPI(
```

(continues on next page)

(continued from previous page)

```
echoWS,  
  "Echo web socket",  
  "1.0",  
  List("production" -> Server("api.example.com", "wss"))  
)
```

Servers can also be later added through methods on the `AsyncAPI` object.

Multiple endpoints can be converted to an `AsyncAPI` instance by calling the method using a list of endpoints.

The `asyncapi` case classes can then be serialised, either to JSON or YAML using [Circe](#):

```
import sttp.tapir.asyncapi.circe.yaml._  
  
println(docs.toYaml)
```

4.29.1 Options

Options can be customised by providing an implicit instance of `AsyncAPIDocsOptions`, when calling `.toAsyncAPI`.

- `subscribeOperationId`: basing on the endpoint's path and the entire endpoint, determines the id of the subscribe operation. This can be later used by code generators as the name of the method to receive messages from the socket.
- `publishOperationId`: as above, but for publishing (sending messages to the web socket).

4.29.2 Exposing AsyncAPI documentation

`AsyncAPI` documentation can be exposed through the [AsyncAPI playground](#).

4.30 Testing

4.30.1 White box testing

If you are unit testing your application you should stub all external services.

If you are using `sttp` client to send HTTP requests, and if the external APIs, which your application consumes, are described using `tapir`, you can create a stub of the service by converting endpoints to `SttpBackendStub` (see the [sttp documentation](#) for details on how the stub works).

Add the dependency:

```
"com.softwaremill.sttp.tapir" %% "tapir-sttp-stub-server" % "0.17.15"
```

And the following imports:

```
import sttp.client3.testing.SttpBackendStub  
import sttp.tapir.server.stub._
```

Given the following endpoint:

```
import sttp.tapir._
import sttp.tapir.generic.auto._
import sttp.tapir.json.circe._
import io.circe.generic.auto._

case class ResponseWrapper(value: Double)

val endpoint = sttp.tapir.endpoint
  .in("api" / "sometest4")
  .in(query[Int]("amount"))
  .post
  .out(jsonBody[ResponseWrapper])
```

Convert any endpoint to `SttpBackendStub`:

```
import sttp.client3.monad.IdMonad

implicit val backend = SttpBackendStub
  .apply(IdMonad)
  .whenRequestMatchesEndpoint(endpoint)
  .thenSuccess(ResponseWrapper(1.0))
```

4.30.2 Black box testing

When testing application as a whole component, running for example in docker, you might want to stub external services with which you application interacts, to steer their behaviors.

To do that you might want to use well-known solutions like e.g. [wiremock](#) or [mock-server](#), but if their api is described using tapir you might want to use [livestub](#), which combines nicely with the rest of the sttp ecosystem.

4.31 Generate endpoint definitions from an OpenAPI YAML

Note: This is a really early alpha implementation.

4.31.1 Installation steps

Add the sbt plugin to the project/plugins.sbt:

```
addSbtPlugin("com.softwaremill.sttp.tapir" % "sbt-openapi-codegen" % "0.17.15")
```

Enable the plugin for your project in the build.sbt:

```
enablePlugins(OpenapiCodegenPlugin)
```

Add your OpenApi file to the project, and override the `openapiSwaggerFile` setting in the build.sbt:

```
openapiSwaggerFile := baseDirectory.value / "swagger.yaml"
```

At this point your compile step will try to generate the endpoint definitions to the `sttp.tapir.generated.TapirGeneratedEndpoints` object, where you can access the defined case-classes and endpoint definitions.

4.31.2 Usage and options

The generator currently supports these settings, you can override them in the `build.sbt`;

setting	default value	description
<code>openapiSwaggerFile</code>	<code>baseDirectory.value / "swagger.yaml"</code>	The swagger file with the api definitions.
<code>openapiPackage</code>	<code>sttp.tapir.generated</code>	The name for the generated package.
<code>openapiObject</code>	<code>TapirGeneratedEndpoints</code>	The name for the generated object.

The general usage is;

```
import sttp.tapir.generated._
import sttp.tapir.docs.openapi._
import sttp.tapir.openapi.circe.yaml._

val docs = TapirGeneratedEndpoints.generatedEndpoints.toOpenAPI("My Bookshop", "1.0")
```

Limitations

Currently, the generated code depends on `"io.circe" %% "circe-generic"`. In the future probably we will make the encoder/decoder json lib configurable (PRs welcome).

We currently miss a lot of OpenApi features like:

- tags
- enums/ADTs
- missing model types and meta descriptions (like date, minLength)
- file handling

4.32 Other interpreters

At its core, tapir creates a data structure describing the HTTP endpoints. This data structure can be freely interpreted also by code not included in the library. Below is a list of projects, which provide tapir interpreters.

4.32.1 GraphQL

[Caliban](#) allows you to easily turn your Tapir endpoints into a GraphQL API. More details in the [documentation](#).

4.32.2 tapir-gen

[tapir-gen](#) extends tapir to do client code generation. The goal is to auto-generate clients in multiple-languages with multiple libraries.

[scala-opentracing](#) contains a module which provides a small integration layer that allows you to create traced http endpoints from tapir Endpoint definitions.

4.33 Creating your own Tapir

Tapir uses a number of packages which contain either the data classes for describing endpoints or interpreters of this data (turning endpoints into a server or a client). Importing these packages every time you want to use Tapir may be tedious, that's why each package object inherits all of its functionality from a trait.

Hence, it is possible to create your own object which combines all of the required functionalities and provides a single-import whenever you want to use tapir. For example:

```
object MyTapir extends Tapir
  with AkkaHttpServerInterpreter
  with SttpClientInterpreter
  with OpenAPIDocsInterpreter
  with SchemaDerivation
  with TapirJsonCirce
  with TapirOpenAPICirceYaml
  with TapirAliases
```

Then, a single import `MyTapir._` and all Tapir data types and interpreter methods will be in scope!

4.34 Troubleshooting

4.34.1 StackOverflowException during compilation

Sidenote for scala 2.12.4 and higher: if you encounter an issue with compiling your project because of a `StackOverflowException` related to [this scala bug](#), please increase your stack memory. Example:

```
sbt -J-Xss4M clean compile
```

4.34.2 Logging of generated macros code

For some cases, it may be helpful to examine how generated macros code looks like. To do that, just set an environmental variable and check compilation logs for details.

```
export TAPIR_LOG_GENERATED_CODE=true
```

4.35 Contributing

Tapir is an early stage project. Everything might change. All suggestions welcome :)

See the list of [issues](#) and pick one! Or report your own.

If you are having doubts on the *why* or *how* something works, don't hesitate to ask a question on [gitter](#) or via [github](#). This probably means that the documentation, scaladocs or code is unclear and can be improved for the benefit of all.

4.35.1 Acknowledgments

Tuple-concatenating code is copied from [akka-http](#)

Generic derivation configuration is copied from [circe](#)