

---

# **Valum Documentation**

*Release 0.3.0-dev*

**Antono Vasiljev, Guillaume Poirier-Morency**

October 09, 2016



<b>1</b>	<b>Installation</b>	<b>3</b>
1.1	Packages . . . . .	3
1.2	Dependencies . . . . .	3
1.3	Download the sources . . . . .	4
1.4	Build . . . . .	4
1.5	Install . . . . .	4
1.6	Run the tests . . . . .	5
1.7	Run the sample application . . . . .	5
<b>2</b>	<b>Quickstart</b>	<b>7</b>
2.1	Simple ‘Hello world!’ application . . . . .	7
2.2	Building with valac . . . . .	7
2.3	Building with waf . . . . .	8
2.4	Building with Meson . . . . .	8
2.5	Running the example . . . . .	8
<b>3</b>	<b>Application</b>	<b>11</b>
3.1	Creating an application . . . . .	11
3.2	Binding a route . . . . .	11
3.3	Serving the application . . . . .	12
<b>4</b>	<b>VSGI</b>	<b>13</b>
4.1	HTTP authentication . . . . .	13
4.2	Connection . . . . .	14
4.3	Request . . . . .	14
4.4	Response . . . . .	16
4.5	Cookies . . . . .	19
4.6	Converters . . . . .	21
4.7	Server . . . . .	22
4.8	Entry point . . . . .	28
<b>5</b>	<b>Router</b>	<b>31</b>
5.1	Route . . . . .	31
5.2	Method . . . . .	31
5.3	Named route . . . . .	33
5.4	Use . . . . .	33
5.5	Asterisk . . . . .	34
5.6	Rule . . . . .	34

5.7	Regular expression . . . . .	35
5.8	Matcher callback . . . . .	36
5.9	Scoping . . . . .	36
5.10	Context . . . . .	36
5.11	Next . . . . .	37
5.12	Error handling . . . . .	37
5.13	Subrouting . . . . .	38
5.14	Cleaning up route logic . . . . .	38
<b>6</b>	<b>Redirection and Error</b>	<b>41</b>
6.1	Default handling . . . . .	41
6.2	Informational (1xx) . . . . .	42
6.3	Success (2xx) . . . . .	42
6.4	Redirection (3xx) . . . . .	43
6.5	Client (4xx) and server (5xx) error . . . . .	43
6.6	Errors in next . . . . .	43
<b>7</b>	<b>Middlewares</b>	<b>45</b>
7.1	Authenticate . . . . .	45
7.2	Basepath . . . . .	45
7.3	Basic . . . . .	46
7.4	Content Negotiation . . . . .	47
7.5	Decode . . . . .	49
7.6	Safely . . . . .	49
7.7	Sequence . . . . .	50
7.8	Server-Sent Events . . . . .	50
7.9	Static Resource Delivery . . . . .	51
7.10	Status . . . . .	52
7.11	Subdomain . . . . .	53
7.12	Forward . . . . .	55
<b>8</b>	<b>Recipes</b>	<b>57</b>
8.1	Configuration . . . . .	57
8.2	JSON . . . . .	59
8.3	Persistence . . . . .	61
8.4	Resources . . . . .	62
8.5	Scripting . . . . .	63
8.6	Templating . . . . .	64
<b>9</b>	<b>Hacking</b>	<b>67</b>
9.1	Building . . . . .	67
9.2	Code conventions . . . . .	67
9.3	General strategies . . . . .	67
9.4	Tricky stuff . . . . .	68
9.5	Coverage . . . . .	68
9.6	Tests . . . . .	69
9.7	Version bump . . . . .	69
9.8	Translations . . . . .	69
<b>10</b>	<b>GNU Lesser General Public License</b>	<b>71</b>
10.1	0. Additional Definitions . . . . .	71
10.2	1. Exception to Section 3 of the GNU GPL . . . . .	71
10.3	2. Conveying Modified Versions . . . . .	72
10.4	4. Combined Works . . . . .	72
10.5	5. Combined Libraries . . . . .	73

10.6 6. Revised Versions of the GNU Lesser General Public License . . . . . 73



Valum is a web micro-framework written in Vala and licensed under the LGPLv3. Its source code and releases are available on GitHub: [valum-framework/valum](https://github.com/valum-framework/valum).

This user documentation aims to be as complete as possible and covers topics that are not directly related to the framework, but essential for web development. If you think that this document could be improved, [open a ticket on GitHub](#) to let us know.





---

## Installation

---

This document describes the compilation and installation process. Most of that work is automated with [Meson](#), a build tool written in Python.

### 1.1 Packages

Packages for RPM and Debian based Linux distributions will be provided for stable releases so that the framework can easily be installed in a container or production environment.

#### 1.1.1 Fedora

RPM packages for Fedora (21, 22 and rawhide) are available from the [arteymix/valum-framework](#) Copr repository.

```
dnf copr enable arteymix/valum-framework
```

The *valum* package contains the shared library and *valum-devel* contains all that is necessary to build an application.

```
dnf install valum valum-devel
```

#### 1.1.2 Nix

```
nix-shell -p valum
```

### 1.2 Dependencies

The following dependencies are minimal to build the framework under Ubuntu 12.04 LTS and should be satisfied by most recent Linux distributions.

Package	Version
vala	>=0.26
python	>=3.4
meson	>=0.31
ninja	>=1.5.1
glib-2.0	>=2.32
gio-2.0	>=2.32
gio-unix-2.0	>=2.32
gthread-2.0	>=2.32
libsoup-2.4	>=2.38

Recent dependencies will enable more advanced features:

Package	Version	Feature
gio-2.0	>=2.34	CGI server uses the command line stdin which can be provided by DBus
gio-2.0	>=2.40	CLI arguments parsing
gio-2.0	>=2.44	<code>write_head_async</code> in <code>Response</code>
libsoup-2.4	>=2.48	new server API
libsoup-2.4	>=2.50	uses <code>Soup.ClientContext.steal_connection</code> directly

You can also install additional dependencies to build the examples, you will have to specify the `--enable-examples` flag during the configure step.

Package	Description
json-glib-1.0	JSON library
libmemcached	client for memcached cache storage
libluajit	embed a Lua VM

## 1.3 Download the sources

You may either clone the whole git repository or download one of our [releases from GitHub](#):

```
git clone git://github.com/valum-framework/valum.git && cd valum
```

The `master` branch is a development trunk and is not guaranteed to be very stable. It is always a better idea to checkout the latest tagged release.

## 1.4 Build

```
mkdir build && cd build
meson ..
ninja # or 'ninja-build' on some distribution
```

## 1.5 Install

Installing the build files is optional and if you omit that step, make sure that `LD_LIBRARY_PATH` points to the `build` folder where the shared library has been generated.

```
sudo ninja install
```

The installation is usually prefixed by `/usr/local`, which is generally not in the dynamic library path. You have to export the `LD_LIBRARY_PATH` environment variable for it to work.

```
export LD_LIBRARY_PATH=/usr/local/lib64 # just lib on 32-bit systems
```

## 1.6 Run the tests

```
ninja test
```

If any of them fail, please [open an issue on GitHub](#) so that we can tackle the bug.

## 1.7 Run the sample application

You can run the sample application from the `build` folder if you called `meson` with the `-D enable_examples=true` flag, it uses the `HTTP`.

```
./build/example/app/app
```



---

## Quickstart

---

Assuming that Valum is built and installed correctly (view [Installation](#) for more details), you are ready to create your first application!

Unless you have installed Valum with `--prefix=/usr` or obtained it from your distribution, you might have to export both `PKG_CONFIG_PATH` and `LD_LIBRARY_PATH` environment variables:

```
export LD_LIBRARY_PATH=/usr/local/lib
export PKG_CONFIG_PATH=/usr/local/lib/pkgconfig
```

Some distributions store 64-bit libraries in a separate folder, typically `lib64`.

### 2.1 Simple ‘Hello world!’ application

You can use this sample application and project structure as a basis. The full `valum-framework/example` is available on GitHub and is kept up-to-date with the latest changes in the framework.

```
using Valum;
using VSGI;

var app = new Router ();

app.get ("/", (req, res) => {
    return res.expand_utf8 ("Hello world!");
});

Server.new_with_application ("http", "org.valum.example.App", app.handle).run ({"app", "--port", "3000"}
```

Typically, the `run` function contains CLI argument to make runtime the parametrizable.

It is suggested to use the following structure for your project, but you can do pretty much what you think is the best for your needs.

```
build/
src/
  app.vala
```

### 2.2 Building with valac

Simple applications can be built directly with `valac`:

```
valac --pkg=valum-0.3 -o build/app src/app.vala
```

The `vala` program will build and run the produced binary, which is convenient for testing:

```
vala --pkg=valum-0.3 src/app.vala
```

## 2.3 Building with waf

It is preferable to use a build system like `waf` to automate all this process. Get a release of `waf` and copy this file under the name `wscript` at the root of your project.

```
def options(cfg):
    cfg.load('compiler_c')

def configure(cfg):
    cfg.load('compiler_c vala')
    cfg.check_cfg(package='valum-0.3', uselib_store='VALUM', args='--libs --cflags')

def build(bld):
    bld.load('compiler_c vala')
    bld.program(
        packages = 'valum-0.3',
        target    = 'app',
        source    = 'src/app.vala',
        use       = 'VALUM')
```

You should now be able to build by issuing the following commands:

```
./waf configure
./waf build
```

## 2.4 Building with Meson

`Meson` is highly-recommended for its simplicity and expressiveness. It's not as flexible as `waf`, but it will handle most projects very well.

```
project('example', 'c', 'vala')

valum = dependency('valum-0.3')

executable('app', sources: ['src/app.vala'], dependencies: valum)
```

```
meson . build
ninja -C build
```

## 2.5 Running the example

VSGI produces process-based applications that are either self-hosted or able to communicate with a HTTP server according to a standardized protocol.

The `HTTP` implementation is self-hosting, so you just have to run it and point your browser at <http://127.0.0.1:3003> to see the result.

```
./build/app
```





---

## Application

---

This document explains step-by-step the sample presented in the [Quickstart](#) document.

Many implementations are provided and documented in [Server](#).

### 3.1 Creating an application

An application is defined by a function that respects the `VSGI.ApplicationCallback` delegate. The `Router` provides `handle` for that purpose along with powerful routing facilities for client requests.

```
var app = new Router ();
```

### 3.2 Binding a route

An application constitute of a list of routes matching and handling user requests. The router provides helpers to declare routes which internally use `Route` instances.

```
app.get ("/", (req, res, next, context) => {  
    return res.expand_utf8 ("Hello world!", null);  
});
```

Every route declaration has a callback associated that does the request processing. The callback, named handler, receives four arguments:

- a `Request` that describes a resource being requested
- a `Response` that correspond to that resource
- a `next` continuation to *keep routing*
- a routing `context` to retrieve and store states from previous and for following handlers

---

**Note:** For an alternative, more structured approach to route binding, see *Cleaning up route logic*

---

### 3.3 Serving the application

This part is pretty straightforward: you create a server that will serve your application at port 3003 and since http was specified, it will be served with HTTP.

```
Server.new_with_application ("http", "org.valum.example.App", app.handle).run ({"app", "--port", "3003"});
```

`Server` takes an application identifier and an `ApplicationCallback`, which is respected by the `handle` function.

Minimal application can be defined using a simple lambda function taking a `Request` and `Response`.

```
Server.new_with_application ("http", "org.valum.example.App", (req, res) => {
  res.status = 200;
  return res.expand ("Hello world!", null);
}).run ({"app", "--port", "3003"});
```

Usually, you would only pass the CLI arguments to `run`, so that your runtime can be parametrized easily, but in this case we just want our application to run with fixed parameters. Options are documented per implementation.

```
public static void main (string[] args) {
  var app = new Router ();

  // assume some route declarations...

  Server.new_with_application ("http", "org.valum.example.App", app.handle).run (args);
}
```

VSGI is a middleware that interfaces different web server technologies under a common and simple set of abstractions. For the moment, it is developed along with Valum to target the needs of a web framework, but it will eventually be extracted and distributed as a shared library.

## 4.1 HTTP authentication

VSGI provide implementations of both basic and digest authentication schemes respectively defined in [RFC 7617](#) and [RFC 7616](#).

Both `Authentication` and `Authorization` objects are provided to produce and interpret their corresponding HTTP headers. The typical authentication pattern is highlighted in the following example:

```
using VSGI;

Server.new_for_application ("http", (req, res) => {
  var authentication = BasicAuthentication ("realm");

  var authorization_header = req.headers.get_one ("Authorization");

  if (authorization_header != null) {
    if (authentication.parse_authorization_header (authorization_header,
                                                    out authorization)) {
      var user = User.from_username (authorization.username);
      if (authorization.challenge (user.password)) {
        return res.expand_utf8 ("Authentication successful!");
      }
    }
  }

  res.headers.replace ("WWW-Authenticate", authentication.to_authenticate_header ());

  return res.end ();
}).run ();
```

### 4.1.1 Basic

The `Basic` authentication scheme is the simplest one and expect the user agent to provide username and password in plain text. It should be used exclusively on a secured transport (e.g. HTTPS).

## 4.2 Connection

All resources necessary to process a [Request](#) and produce a [Response](#) are bound to the lifecycle of a connection instance.

**Warning:** It is not recommended to use this directly as it will most likely result in corrupted operations with no regard to the transfer encoding or message format.

The connection can be accessed from the [Request](#) `connection` property. It is a simple `GLib.IOStream` that provides native access to the input and output stream of the used technology.

The following example shows how to bypass processing with higher-level abstractions. It will only work on [HTTP](#), as CGI-like protocols require the status to be part of the response headers.

```
using VSgi;  
  
Server.new_with_application ("http", "org.vsgui.App", (req, res) => {  
    var message = req.connection.output_stream;  
    message.write_all ("200 Success HTTP/1.1\r\n".data, null);  
    message.write_all ("Connection: close\r\n");  
    message.write_all ("Content-Type: text/plain\r\n");  
    message.write_all ("\r\n".data);  
    message.write_all ("Hello world!".data);  
    return true;  
});
```

## 4.3 Request

Requests are representing incoming demands from user agents to resources served by an application.

### 4.3.1 Method

The `Request` class provides constants for the following HTTP methods:

Deprecated since version 0.3: `libsoup-2.4` provide an enumeration of valid HTTP methods and this will be removed once exposed in their Vala API.

- `OPTIONS`
- `GET`
- `HEAD`
- `POST`
- `PUT`
- `DELETE`
- `TRACE`
- `CONNECT`
- `PATCH`

Additionally, an array of supported HTTP methods is provided by `Request.METHODS`.

```

if (req.method == Request.GET) {
    return res.expand_utf8 ("Hello world!");
}

if (req.method == Request.POST) {
    return res.body.splice (req.body, OutputStreamSpliceFlags.NONE);
}

if (req.method in Request.METHODS) {
    // handle a standard HTTP method...
}

```

### 4.3.2 Headers

Request headers are implemented with `Soup.MessageHeaders` and can be accessed from the `headers` property.

```

Server.new_with_application ("http", "org.vsgit.App", (req) => {
    var accept = req.headers.get_one ("Accept");
    return true;
});

```

libsoup-2.4 provides a very extensive set of utilities to process the information contained in headers.

```

SList<string> unacceptable;
Soup.header_parse_quality_list (req.headers.get_list ("Accept"), out unacceptable);

```

### Cookies

Cookies can also be retrieved from the request headers.

### 4.3.3 Query

The HTTP query is provided in various way:

- parsed as a `HashTable<string, string?>` through the `Request.query` property
- raw with `Request.uri.get_query`

If the query is not provided (e.g. no `?` in the URI), then the `Request.query` property will take the `null` value.

---

**Note:** If the query is not encoded according to `application/x-www-form-urlencoded`, it has to be parsed explicitly.

---

To safely obtain a value from the HTTP query, use `Request.lookup_query` with the null-coalescing operator `??`.

```

req.lookup_query ("key") ?? "default value";

```

### 4.3.4 Body

The body is provided as a `GLib.InputStream` by the `body` property. The stream is transparently decoded from any applied transfer encodings.

Implementation will typically consume the status line, headers and newline that separates the headers from the body in the base stream at construct time. It also guarantee that the body has been decoded if any transfer encoding were applied for the transport.

If the content is encoded with the `Content-Encoding` header, it is the responsibility of your application to decode it properly. VSGI provides common [Converters](#) to simplify the task.

## Flatten

New in version 0.2.4.

In some cases, it is practical to flatten the whole request body in a buffer in order to process it as a whole.

The `flatten`, `flatten_bytes` and `flatten_utf8` functions accumulate the request body into a buffer (a [GLib.MemoryOutputStream](#)) and return the corresponding `uint8[]` data buffer.

The request body is always fixed-size since the HTTP specification requires any request to provide a `Content-Length` header. However, the environment should be configured with a hard limit on payload size.

When you are done, it is generally a good thing to close the request body and depending on the used implementation, this could have great benefits such as freeing a file resource.

```
Server.new_with_application ("org.vsgui.App", (req, res) => {
  var payload = req.flatten ();
  return true;
});
```

## Form

`Soup.Form` can be used to parse `application/x-www-form-urlencoded` format, which is submitted by web browsers.

```
Server.new_with_application ("org.vsgui.App", (req, res) => {
  var data = Soup.Form.decode (req.flatten_utf8 (out bytes_read));
  return true;
});
```

## Multipart body

Multipart body support is planned in a future minor release, more information on [issue #81](#). The implementation will be similar to [Soup.MultipartInputStream](#) and provide part access with a filter approach.

## 4.4 Response

Responses are representing resources requested by a user agent. They are actively streamed across the network, preferably using non-blocking asynchronous I/O.

### 4.4.1 Status

The response status can be set with the `status` property. `libsoup-2.4` provides an enumeration in [Soup.Status](#) for that purpose.

The `status` property will default to `200 OK`.

The status code will be written in the response with `write_head` or `write_head_async` if invoked manually or during the first access to its body.

```
Server.new_with_application ("http", "org.vsgui.App", (req, res) => {
  res.status = Soup.Status.MALFORMED;
  return true;
});
```

## 4.4.2 Reason phrase

New in version 0.3.

The reason phrase provide a textual description for the status code. If `null`, which is the default, it will be generated using `Soup.Status.get_phrase`.

```
Server.new_with_application ("http", "org.vsgui.App", (req, res) => {
  res.status = Soup.Status.OK;
  res.reason_phrase = "Everything Went Well"
  return true;
});
```

To obtain final status line sent to the user agent, use the `wrote_status_line` signal.

```
res.wrote_status_line.connect ((http_version, status, reason_phrase) => {
  if (200 <= status < 300) {
    // assuming a success
  }
});
```

## 4.4.3 Headers

The response headers can be accessed as a `Soup.MessageHeaders` from the `headers` property.

```
Server.new_with_application ("http", "org.vsgui.App", (req, res) => {
  res.status = Soup.Status.OK;
  res.headers.set_content_type ("text/plain", null);
  return res.body.write_all ("Hello world!".data, null);
});
```

Headers can be written in the response by invoking `write_head` or its asynchronous version `write_head_async`. The synchronous version is called automatically when the body is accessed for the first time.

```
res.write_head_async.begin (Priority.DEFAULT, null, () => {
  // produce the body...
});
```

**Warning:** Once written, any modification to the `headers` object will be ignored.

The `head_written` property can be tested to see if it's already the case, even though a well written application should assume that already.

```
if (!res.head_written) {
  res.headers.set_content_type ("text/html", null);
}
```

Since headers can still be modified once written, the `wrote_headers` signal can be used to obtain definitive values.

```
res.wrote_headers (() => {
  foreach (var cookie in res.cookies) {
    message (cookie.to_set_cookie_header ());
  }
});
```

#### 4.4.4 Body

The body of a response is accessed through the `body` property. It inherits from `GLib.OutputStream` and provides synchronous and asynchronous streaming capabilities.

The response body is automatically closed following a RAII pattern whenever the `Response` object is disposed.

Note that a reference to the body is not sufficient to maintain the inner `Connection` alive: a reference to either the `Request` or response be maintained.

You can still close the body early as it can provide multiple advantages:

- avoid further and undesired read or write operation
- indicate to the user agent that the body has been fully sent

#### Expand

New in version 0.3.

To deal with fixed-size body, `expand`, `expand_bytes` and `expand_utf8` utilities as well as their respective asynchronous versions are provided.

It will automatically set the `Content-Length` header to the size of the provided buffer, write the response head and pipe the buffer into the body stream and close it properly.

```
Server.new_with_application ("http", "org.vsgl.App", (req, res) => {
  res.expand_utf8 ("Hello world!");
});
```

#### Filtering

One common operation related to stream is filtering. `GLib.FilterOutputStream` and `GLib.ConverterOutputStream` provide, by composition, many filters that can be used for:

- compression and decompression (gzip, deflate, compress, ...)
- charset conversion
- buffering
- writing data

VSGI also provides its own set of `Converters` which cover parts of the HTTP/1.1 specifications such as chunked encoding.

```
var body = new ConverterOutputStream (res.body,
                                     new CharsetConverter (res.body, "iso-8859-1", "utf-8"));
return body.write_all ("Omelette du fromage!", null);
```



Additionally, some filters are applied automatically if the `Transfer-Encoding` header is set. The obtained `GLib.OutputStream` will be wrapped appropriately so that the application can transparently produce its output.

```
res.headers.append ("Transfer-Encoding", "chunked");
return res.body.write_all ("Hello world!".data, null);
```

## Conversion

New in version 0.3.

The body may be converted, see [Converters](#) for more details.

### 4.4.5 Tee

New in version 0.3.

The response body can be splitted pretty much like how the `tee` UNIX utility works. All further write operations will be performed as well on the passed stream, making it possible to process the payload sent to the user agent.

The typical use case would be to implement a file-based cache that would tee the produced response body into a key-based storage.

```
var cache_key = Checksum.compute_for_string (ChecksumType.SHA256, req.uri.to_string ());
var cache_entry = File.new_for_path ("cache/%s".printf (cache_key));

if (cache_entry.query_exists ()) {
  return res.body.splice (cache_entry.read ());
} else {
  res.tee (cache_entry.create (FileCreateFlags.PRIVATE));
}

res.exand_utf8 ("Hello world!");
```

### 4.4.6 End

New in version 0.3.

To properly close the response, writing headers if missing, `end` is provided:

```
Server.new_with_application ("http", "org.vsgj.App", (req, res, next) => {
  res.status = Soup.Status.NO_CONTENT;
  return res.end () && next ();
}).then ((req, res) => {
  // perform blocking operation here...
});
```

To produce a message before closing, favour `extend` utilities.

## 4.5 Cookies

Cookies are stored in [Request](#) and [Response](#) headers as part of the HTTP protocol.

Utilities are provided to perform basic operations based on [Soup.Cookie](#) as those provided by `libsoup-2.4` requires a [Soup.Message](#), which is not common to all implementations.

- extract cookies from request headers
- find a cookie by its name
- marshal cookies for request or response headers (provided by libsoup-2.4)

### 4.5.1 Extract cookies

Cookies can be extracted as a singly-linked list from a [Request](#) or [Response](#) their order of appearance (see [Soup.MessageHeaders.get\\_list](#) for more details).

The `Request.cookies` property will extract cookies from the `Cookie` headers. Only the `name` and `value` fields will be filled as it is the sole information sent by the client.

```
var cookies = req.cookies;
```

The equivalent property exist for [Response](#) and will extract the `Set-Cookie` headers instead. The corresponding [Request URI](#) will be used for the cookies origin.

```
var cookies = res.cookies;
```

The extracted cookies can be manipulated with common [SList](#) operations. However, they must be written back into the [Response](#) for the changes to be effective.

**Warning:** Cookies will be in their order of appearance and [SList.reverse](#) should be used prior to perform a lookup that respects precedence.

```
cookies.reverse ();

for (var cookie in cookies)
    if (cookie.name == "session")
        return cookie;
```

### 4.5.2 Lookup a cookie

You can lookup a cookie by its name from a [Request](#) using `lookup_cookie`, `null` is returned if no such cookies can be found.

**Warning:** Although this is not formally specified, cookies name are considered as being case-sensitive by [CookieUtils](#) utilities.

If it's signed (recommended for sessions), the equivalent `lookup_signed_cookie` exists.

```
string? session_id;
var session = req.lookup_signed_cookie ("session", ChecksumType.SHA512, "secret".data, out session_id);
```

### 4.5.3 Marshall a cookie

libsoup-2.4 provides a complete implementation with the [Soup.Cookie](#) class to represent and marshal cookies for both request and response headers.

The newly created cookie can be sent by adding a `Set-Cookie` header in the [Response](#).

```
var cookie = new Cookie ("name", "value", "0.0.0.0", "/", 60);
res.headers.append ("Set-Cookie", cookie.to_set_cookie_header ());
```

## 4.5.4 Sign and verify

Considering that cookies are persisted by the user agent, it might be necessary to sign to prevent forgery. `CookieUtils.sign` and `CookieUtils.verify` functions are provided for the purposes of signing and verifying cookies.

**Warning:** Be careful when you choose and store the secret key. Also, changing it will break any previously signed cookies, which may still be submitted by user agents.

It's up to you to choose what hashing algorithm and secret: SHA512 is generally recommended.

The `CookieUtils.sign` utility will sign the cookie in-place. It can then be verified using `CookieUtils.verify`.

The value will be stored in the output parameter if the verification process is successful.

```
CookieUtils.sign (cookie, ChecksumType.SHA512, "secret".data);

string value;
if (CookieUtils.verify (cookie, ChecksumType.SHA512, "secret.data", out value)) {
    // cookie's okay and the original value is stored in value
}
```

The signature is computed in a way it guarantees that:

- we have produced the value
- we have produced the name and associated it to the value

The algorithm is the following:

$$\text{HMAC}(\text{checksum\_type}, \text{key}, \text{HMAC}(\text{checksum\_type}, \text{key}, \text{value}) + \text{name}) + \text{value}$$

The verification process does not handle special cases like values smaller than the hashing: cookies are either signed or not, even if their values are incorrectly formed.

If well-formed, cookies are verified in constant-time to prevent time-based attacks.

## 4.6 Converters

VSGI provide stream utilities named converters to convert data according to modern web standards.

These are particularly useful to encode and decode request and response bodies in VSGI implementations.

GLib provide default converters for charset conversion and zlib compression. These can be used to compress the message bodies and convert the string encoding transparently.

- `GLib.CharsetConverter`
- `GLib.ZLibCompressor`
- `GLib.ZLibDecompressor`

Converters can be applied on both the `Request` and `Response` object using the `convert` method.

```
Server.new_with_application ("http", "org.vsgi.App", (req, res) => {
    res.headers.append ("Content-Encoding", "gzip");
    res.convert (new ZlibCompressor (ZlibCompressorFormat.GZIP));
    return res.expand_utf8 ("Hello world!");
});
```

**Warning:** The `Content-Encoding` header must be adapted to reflect the current set of encodings applied (or unapplied) on the payload.

Since conversion typically affect the resulting size of the payload, the `Content-Length` header must be set appropriately. To ease that, the new value can be specified as second argument. Note that `-1` is used to describe an undetermined length.

```
res.convert (new CharsetConverter ("UTF-8", "ascii"), res.headers.get_content_length ());
```

The default, which apply in most cases, is to remove the `Content-Length` header and thus describe an undetermined length.

## 4.7 Server

Server provide HTTP technologies integrations under a common interface. They inherit from `GLib.Application`, providing an optimal integration with the host environment.

### 4.7.1 HTTP

libsoup-2.4 provides a built-in HTTP server that you can use to test your application or spawn workers in production.

```
using Valum;
using VSgi.HTTP;

new Server ("org.vsgl.HTTP", () => {
    res.status = Soup.Status.OK;
    return res.body.write_all ("Hello world!".data, null);
}).run ({"app", "--port", "3003"});
```

### Options

The implementation provides most options provided by `Soup.Server` through command-line options. The available options may vary and can be asserted with the `--help` flag.

Option	Default	Description
<code>--port</code>	3003	port the server is listening on
<code>--all</code>	local	listen on all interfaces
<code>--ipv4-only</code>	disabled	only listen to IPv4 interfaces
<code>--ipv6-only</code>	disabled	only listen on IPv6 interfaces
<code>--file-descriptor</code>	none	listen to the provided file descriptor
<code>--https</code>	disabled	listen for https connections rather than plain http
<code>--ssl-cert-file</code>	none	path to a file containing a PEM-encoded certificate
<code>--ssl-key-file</code>	none	path to a file containing a PEM-encoded private key
<code>--server-header</code>	Valum/0.2	value to use for the “Server” header on Messages processed by this server.
<code>--raw-paths</code>	disabled	percent-encoding in the Request-URI path will not be automatically decoded

### Notes

- if `--all` is not supplied, the server will only listen to local interfaces

- `--all` can be combined with `--ipv4-only` or `--ipv6-only` to listen on all IPv4 or IPv6 interfaces
- if `--https` is specified, you must provide a SSL or TLS certificate along with a private key

## 4.7.2 CGI

CGI is a very simple process-based protocol that uses commonly available process resources:

- environment variables
- standard input stream for the [Request](#)
- standard output stream for the [Response](#)

**Warning:** The CGI protocol expects the response to be written in the standard output: writing there will most surely corrupt the response.

The `VSGI.CGI` implementation provides a basis for its derivatives protocols such as [FastCGI](#) and [SCGI](#) and can be used along with any HTTP server.

The interpretation of the environment prioritize the [CGI/1.1](#) specification while providing fallbacks with values we typically found like `REQUEST_URI`.

Since a process is spawned per request and exits when the latter finishes, scheduled asynchronous tasks might not be processed. To overcome this issue, `hold` and `release` should be used to keep the server alive as long as necessary.

If your task involve the [Request](#) or [Response](#) in its callback, the connection will be kept alive until both are freed.

```
using VSGI.CGI;

Server? server = null;

ApplicationCallback app = (req, res) => {
    Idle.add () => {
        message ("Hello world!");
        server.release ();
    };
    server.hold ();

    // no need to hold & release here, the reference on the request ensures
    // it already
    Idle.add () => {
        req.body.write_all ("Hello world!".data, null);
    };

    return true;
};

server = new Server ("org.vsgi.CGI", app);

server.run ();
```

## lighttpd

There is an example in `examples/cgi` providing a sample `lighttpd` configuration file. Once launched, the application can be accessed at the following address: <http://127.0.0.1:3003/cgi-bin/app/>.

```
lighttpd -D -f examples/cgi/lighttpd.conf
```

### 4.7.3 FastCGI

FastCGI is a binary protocol that multiplexes requests over a single connection.

VSGI uses [FastCGI bindings](#) under the hood to provide a compliant implementation. See [Installation](#) for more information about the framework dependencies.

The whole request cycle is processed in a thread and dispatched in the main context, so it's absolutely safe to use shared states.

#### Options

Option	Default	Description
<code>--port</code>	none	listen on a TCP port from local interface
<code>--socket</code>	none	listen on a UNIX socket path
<code>--file-descriptor</code>	0	listen to the provided file descriptor
<code>--backlog</code>	10	connection queue depth in the <code>listen</code> call

Only one option from `--port`, `--socket` and `--file-descriptor` can be specified.

By default, the FastCGI implementation listens on the file descriptor 0, which is conventionally the case when the process is spawned by an HTTP server.

#### lighttpd

`lighttpd` can be used to develop and potentially deploy your application. An [example of configuration file](#) is available in the `fastcgi` example folder.

You can run the FastCGI example with `lighttpd`:

```
./waf configure --enable-examples && ./waf build  
lighttpd -D -f examples/fastcgi/lighttpd.conf
```

#### Apache

Under Apache, there are two mods available: `mod_fcgid` is more likely to be available as it is part of Apache and `mod_fastcgi` is developed by those who did the FastCGI specifications.

- [mod\\_fcgid](#)
- [mod\\_fastcgi](#)

#### Nginx

Nginx expect a process to be already spawned and will communicate with it using a TCP port or a socket path. Read more about [ngx\\_http\\_fastcgi\\_module](#).

You can spawn a process with `spawn-fcgi`, an utility part of `lighttpd`.

## 4.7.4 SCGI

SCGI (Simple Common Gateway Interface) is a stream-based protocol that is particularly simple to implement.

**Note:** SCGI is the recommended implementation and should be used when available as it takes the best out of GIO asynchronous API.

The implementation uses a [GLib.SocketService](#) and processes multiple requests using non-blocking I/O.

### Options

Option	Default	Description
<code>--any</code>	none	listen on any open TCP port
<code>--port</code>	none	listen on a TCP port from local interface
<code>--file-descriptor</code>	0	listen to the provided file descriptor
<code>--backlog</code>	10	connection queue depth in the <code>listen</code> call

## 4.7.5 General

Basically, you have access to a [DBusConnection](#) to communicate with other process and a [GLib.MainLoop](#) to process events and asynchronous work.

- an application id to identify primary instance
- startup signal emitted right after the registration
- shutdown signal just before the server exits
- a resource base path
- ability to handle CLI arguments

The server can be gracefully terminated by sending a *SIGTERM* signal to the process.

## 4.7.6 Load an implementation

Server implementations are dynamically loaded using [GLib.Module](#). It makes it possible to define its own implementation if necessary.

The shared library name must conform to `libvsgi-<name>` with the appropriate extension. For instance, on GNU/Linux, the `CGI` module is stored in `$(LIBDIR)/vsgi/servers/libvsgi-cgi.so`.

To load an implementation, use the `Server.new` factory, which can receive GObject-style arguments as well.

```
var cgi_server = Server.new ("cgi", "application-id", "org.valum.example.CGI");

if (cgi_server == null) {
  assert_not_reached ();
}

cgi_server.set_application_callback ((req, res) => {
  return res.expand_utf8 ("Hello world!");
});
```

For typical case, use `Server.new_with_application` to initialize the instance with an application identifier and callback:

```
var cgi_server = Server.new_with_application ("cgi", "org.example.CGI", (req, res) => {
  return true;
});
```

For more flexibility, the `ServerModule` class allow a more fine-grained control for loading a server implementation. If non-null, the `directory` property will be used to retrieve the implementation from the given path instead of standard locations.

The computed path of the shared library is available from `path` property, which can be used for debugging purposes.

```
var directory = "/usr/lib64/vsgi/servers";
var cgi_module = new ServerModule (directory, "cgi");

if (!cgi_module.load ()) {
  error ("could not load 'cgi' from '%s'", cgi_module.path);
}

var server = Object.new (cgi_module.server_type);
```

Unloading a module is not necessary: once initially loaded, a use count is kept so that it can be loaded on need or unloaded if not used.

**Warning:** Since a `ServerModule` cannot be disposed (see [GLib.TypeModule](#)), one must be careful of how its reference is being handled. For instance, `Server.new` keeps track of requested implementations and persist them forever.

Mixing direct usages of `ServerModule` and `Server.@new` (and the likes) is not recommended and will result in undefined behaviours if an implementation is loaded more than once.

## 4.7.7 DBus connection

`GLib.Application` will automatically register to the session DBus bus, making IPC (Inter-Process Communication) an easy thing.

It can be used to expose runtime information such as a database connection details or the amount of processing requests. See this [example of DBus server](#) for code examples.

This can be used to request services, communicate between your workers and interact with the runtime.

```
var connection = server.get_dbus_connection ()

connection.call ()
```

## 4.7.8 Options

Each server implementation can optionally take arguments that parametrize its runtime.

If you build your application in a main block, it will not be possible to obtain the CLI arguments to parametrize the runtime. Instead, the code can be written in a usual main function.

```
public static int main (string[] args) {
  Server.new ("http", "org.vsgi.App", (req, res) => {
    res.status = Soup.Status.OK;
    return res.body.write_all ("Hello world!".data, null);
  });
}
```



```
}).run (args);
}
```

If you specify the `--help` flag, you can get more information on the available options which vary from an implementation to another.

```
build/examples/fastcgi --help
```

```
Usage:
  fastcgi [OPTION...]

Help Options:
  -h, --help                Show help options
  --help-all               Show all help options
  --help-gapplication       Show GApplication options

Application Options:
  --forks=0                 Number of fork to create
  -s, --socket              Listen to the provided UNIX domain socket (or named pipe for WinNT)
  -p, --port                Listen to the provided TCP port
  -f, --file-descriptor=0   Listen to the provided file descriptor
  -b, --backlog=10         Listen queue depth used in the listen() call
```

## 4.7.9 Forking

To achieve optimal performances on a multi-core architecture, VSGI support forking at the server level.

**Warning:** Keep in mind that the `fork` system call will actually copy the whole process: no resources (e.g. lock, memory) can be shared unless inter-process communication is used.

The `--forks` option will spawn the requested amount of workers, which should optimally default to the number of available CPUs.

```
server.run ("app", {"--forks=4"});
```

It's also possible to fork manually via the `fork` call.

```
using VSGI.HTTP;

var server = new Server ();

server.listen (options);
server.fork ();

new MainLoop ().run ();
```

It is recommended to fork only through that call since implementations such as `CGI` are not guaranteed to support it.

## 4.7.10 Listen on distinct interfaces

Typically, `fork` is called after `listen` so that all processes share the same file descriptors and interfaces. However, it might be useful to listen to multiple ports (e.g. HTTP and HTTPS).

```
using VSGI.HTTP;
```

```

var server = new Server ();

var parent_options = new VariantDict ();
var child_options = new VariantDict ();

// parent serve HTTP
parent_options.insert_value ("port", new Variant.int32 (80));

// child serve HTTPS
child_options.insert_value ("https");
child_options.insert_value ("port", new Variant.int32 (443));

if (server.fork () > 0) {
    server.listen (parent_options);
} else {
    server.listen (child_options);
}

new MainLoop ().run ();

```

VSGI produces process-based applications that are able to communicate with various HTTP servers using standardized protocols.

## 4.8 Entry point

The entry point of a VSGI application is type-compatible with the `ApplicationCallback` delegate. It is a function of two arguments: a `Request` and a `Response` that return a boolean indicating if the request has been or will be processed.

```

using VSGI;

Server.new_with_application ("http", "org.vsgui.App", (req, res) => {
    // process the request and produce the response...
    return true;
}).run ();

```

If an application indicate that the request has not been processed, it's up to the server implementation to decide what will happen.

### 4.8.1 Error handling

New in version 0.3.

At any moment, an error can be raised and handled by the server implementation which will in turn teardown the connection appropriately.

```

Server.new_with_application ("http", "org.vsgui.App", (req, res) => {
    throw new IOError.FAILED ("some I/O failed");
});

```

### 4.8.2 Asynchronous processing

The asynchronous processing model follows the [RAII pattern](#) and wraps all resources in a connection that inherits from `GLib.IOStream`. It is therefore important that the said connection is kept alive as long as the streams are being

used.

The [Request](#) holds a reference to the said connection and the [Response](#) indirectly does as it holds a reference to the request. Generally speaking, holding a reference on any of these two instances is sufficient to keep the streams usable.

**Warning:** As VSGI relies on reference counting to free the resources underlying a request, you must keep a reference to either the [Request](#) or [Response](#) during the processing, including in asynchronous callbacks.

It is important that the connection persist until all streams operations are done as the following example demonstrates:

```
res.body.write_async.begin ("Hello world!",
                            Priority.DEFAULT,
                            null,
                            (body, result) => {
    // the response reference will make the connection persist
    var written = res.body.write_async.end (result);
});
```



---

## Router

---

Router is the core component of Valum. It dispatches request to the right handler and processes certain error conditions described in [Redirection and Error](#).

The router is constituted of a sequence of `Route` objects which may or may not match incoming requests and perform the process described in their handlers.

### 5.1 Route

The most basic and explicit way of attaching a handler is `Router.route`, which attach the provided `Route` object to the sequence.

```
app.route (new RuleRoute (Method.GET, "/", null, () => {}));
```

Route are simple objects which combine a matching and handling processes. The following sections implicitly treat of route objects such such as `RuleRoute` and `RegexRoute`.

### 5.2 Method

New in version 0.3.

The `Method` flag provide a list of HTTP methods and some useful masks used into route definitions.

Flag	Description
<code>Method.SAFE</code>	safe methods
<code>Method.IDEMPOTENT</code>	idempotent methods (e.g. <code>SAFE</code> and <code>PUT</code> )
<code>Method.CACHEABLE</code>	cacheable methods (e.g. <code>HEAD</code> , <code>GET</code> and <code>POST</code> )
<code>Method.ALL</code>	all standard HTTP methods
<code>Method.OTHER</code>	any non-standard HTTP methods
<code>Method.ANY</code>	anything, including non-standard methods
<code>Method.PROVIDED</code>	indicate that the route provide its methods
<code>Method.META</code>	mask for all meta flags like <code>Method.PROVIDED</code>

---

**Note:** Safe, idempotent and cacheable methods are defined in section 4.2 of [RFC 7231](#).

---

Using a flag makes it really convenient to capture multiple methods with the `|` binary operator.

```
app.rule (Method.GET | Method.POST, "/", (req, res) => {
  // matches GET and POST
});
```

`Method.GET` is defined as `Method.ONLY_GET | Method.HEAD` such that defining the former will also provide a `HEAD` implementation. In general, it's recommended to check the method in order to skip a body that won't be considered by the user agent.

```
app.get ("/", () => {
  res.headers.set_content_type ("text/plain", null);
  if (req.method == Request.HEAD) {
    return res.end (); // skip unnecessary I/O
  }
  return res.expand_utf8 ("Hello world!");
});
```

To provide only the `GET` part, use `Method.ONLY_GET`.

```
app.rule (Method.ONLY_GET, "/", () => {
  res.headers.set_content_type ("text/plain", null);
  return res.expand_utf8 ("Hello world!");
});
```

Per definition, `POST` is considered cacheable, but if it's not desirable, it may be removed from the mask with the unary `~` operator.

```
app.rule (Method.CACHEABLE & ~Method.POST, "/", () => {
  res.headers.set_content_type ("text/plain", null);
  return res.expand_utf8 ("Hello world!");
});
```

## 5.2.1 Non-standard method

To handle non-standard HTTP method, use the `Method.OTHER` along with an explicit check.

```
app.method (Method.OTHER, "/rule", (req, res) => {
  if (req.method != "CUSTOM")
    return next ();
});
```

## 5.2.2 Reverse

New in version 0.3.

Some route implementations can be reversed into URLs by calling `Route.to_url` or the alternative `Route.to_urlv` and `Route.to_url_from_hash`. It may optionally take parameters which, in the case of the rule-based route, correspond to the named captures.

## 5.2.3 Introspection

The router introspect the route sequence to determine what methods are allowed for a given URI and thus produce a nice `Allow` header. To mark a method as *provided*, the `Method.PROVIDED` flag has to be used. This is automatically done for the helpers and the `Router.rule` function described below.

Additionally, the `OPTIONS` and `TRACE` are automatically handled if not specified for a path. The `OPTIONS` will produce a `Allow` header and `TRACE` will feedback the request into the response payload.

## 5.3 Named route

New in version 0.3.

Few of the helpers provided by the router also accept an additional parameter to name the created route object. This can then be used to generate reverse URLs with `Router.url_for`.

---

**Note:** This feature is only support for the rule-based and path-based route implementations.

---

```
var app = new Router ();

app.get ("/", (req, res) => {
    return res.expand_utf8 ("Hello world! %s".printf (app.url_for ("home")));
}, "home");
```

Likewise to `to_url`, it's possible to pass additional parameters as varidic arguments. The following example show how one can serve relocatable static resources and generate URLs in a `Compose` template.

```
using Compose.HTML5;
using Valum;
using Valum.Static;

var app = new Router ();

app.get ("/", () => {
    return res.expand_utf8 (
        html (
            head (
                title ("Hello world!"),
                link ("stylesheet",
                    app.url_for ("static",
                        "path", "bootstrap/dist/css/bootstrap.min.css")),
            body ());
});

app.get ("/static/<path>", serve_from_path ("static"), "static");
```

Other helpers are provided to pass a `GLib.HashTable` via `Router.url_for_hash` or explicit varidic arguments via `Router.url_for_valist`.

## 5.4 Use

New in version 0.3.

The simplest way to attach a handler is `Router.use`, which unconditionally apply the route on the request.

```
app.use ((req, res, next) => {
    var params = new HashTable<string, string> (str_hash, str_equal);
    params["charset"] = "iso-8859-1";
    res.headers.set_content_type ("text/xhtml+xml", params);
    return next ();
});
```

It is typically used to mount a `Middlewares` on the router.

## 5.5 Asterisk

New in version 0.3.

The special `*` URI is handled by the `Router.asterisk` helper. It is typically used along with the `OPTIONS` method to provide a self-description of the web service or application.

```
app.asterisk (Method.OPTIONS, () => {
  return true;
});
```

## 5.6 Rule

Changed in version 0.3: Rule helpers (e.g. `get`, `post`, `rule`) must explicitly be provided with a leading slash.

The rule syntax has been greatly improved to support groups, optionals and wildcards.

The *de facto* way of attaching handler callbacks is based on the rule system. The `Router.rule` as well as all HTTP method helpers use it.

```
app.rule (Method.ALL, "/rule" (req, res) => {
  return true;
});
```

The syntax for rules is given by the following EBNF grammar:

```
rule      = piece | parameter | group | optional | wildcard, [ rule ];
group     = '(', rule, ')';
optional  = (piece | parameter | group), '?';
wildcard  = '*';
parameter = '<', [ type, ':' ], name, '>'; (* considered as a terminal *)
type      = ? any sequence of word character ?;
name      = ? any sequence of word character ?;
piece     = ? any sequence of URL-encoded character ?;
```

### 5.6.1 Remarks

- a piece is a single character, so `/users/?` only indicates that the `/` is optional
- the wildcard `*` matches anything, just like the `.*` regular expression

The following table show valid rules and their corresponding regular expressions. Note that rules are matching the whole path as they are automatically anchored.

Rule	Regular expression
<code>/user</code>	<code>^/user\$</code>
<code>/user/&lt;id&gt;</code>	<code>^/user/(?&lt;id&gt;\w+)\$</code>
<code>/user/&lt;int:id&gt;</code>	<code>^/user/(?&lt;id&gt;\d+)\$</code>
<code>/user(/&lt;int:id&gt;)?</code>	<code>^/user(?:/(?&lt;id&gt;\d+))?\$</code>

### 5.6.2 Types

Valum provides built-in types initialized in the `Router` constructor. The following table details these types and what they match.



Type	Regex	Description
int	\d+	matches non-negative integers like a database primary key
string	\w+	matches any word character
path	(?:\.\.?[\w/-\s/])+	matches a piece of route including slashes, but not ..

Undeclared types default to `string`, which matches any word characters.

It is possible to specify or overwrite types using the `types` map in [Router](#). This example will define the `path` type matching words and slashes using a regular expression literal.

```
app.register_type ("path", new Regex ("[\w/]+", RegexOptions.OPTIMIZE));
```

If you would like `int` to match negatives integer, you may just do:

```
app.register_type ("int", new Regex ("(?:-)?\d+", RegexOptions.OPTIMIZE));
```

Rule parameters are available from the routing context by their name.

```
app.get ("/<controller>/<action>", (req, res, next, context) => {
    var controller = context["controller"].get_string ();
    var action      = context["action"].get_string ();
});
```

### 5.6.3 Helpers

Helpers for the methods defined in the HTTP/1.1 protocol and the extra TRACE methods are included. The path is matched according to the rule system defined previously.

```
app.get ("/", (req, res) => {
    return res.expand_utf8 ("Hello world!");
});
```

The following example deal with a POST request providing using [Soup.Form](#) to decode the payload.

```
app.post ("/login", (req, res) => {
    var data = Soup.Form.decode (req.flatten_utf8 ());

    var username = data["username"];
    var password = data["password"];

    // assuming you have a session implementation in your app
    var session = new Session.authenticated_by (username, password);

    return true;
});
```

## 5.7 Regular expression

Changed in version 0.3: The regex helper must be provided with an explicit leading slash.

If the rule system does not suit your needs, it is always possible to use regular expression. Regular expression will be automatically scoped, anchored and optimized.

```
app.regex (Method.GET, new Regex ("/home/?", RegexOptions.OPTIMIZE), (req, res) => {
    return res.body.write_all ("Matched using a regular expression.".data, true);
});
```

Named captures are registered on the routing context.

```
app.regex (new Regex ("/(?<word>\w+) ", RegexOptions.OPTIMIZE), (req, res, next, ctx) => {
    var word = ctx["word"].get_string ();
});
```

## 5.8 Matcher callback

Request can be matched by a simple callback typed by the `MatcherCallback` delegate.

```
app.matcher (Method.GET, (req) => { return req.uri.get_path () == "/home"; }, (req, res) => {
    // matches /home
});
```

## 5.9 Scoping

Changed in version 0.3: The scope feature does not include a slash, instead you should scope with a leading slash like shown in the following examples.

Scoping is a powerful prefixing mechanism for rules and regular expressions. Route declarations within a scope will be prefixed by `<scope>`.

The `Router` maintains a scope stack so that when the program flow enter a scope, it pushes the fragment on top of that stack and pops it when it exits.

```
app.scope ("/admin", (admin) => {
    // admin is a scoped Router
    app.get ("/users", (req, res) => {
        // matches /admin/users
    });
});

app.get ("/users", (req, res) => {
    // matches /users
});
```

To literally mount an application on a prefix, see the [Basepath](#) middleware.

## 5.10 Context

New in version 0.3.

During the routing, states can be obtained from a previous handler or passed to the next one using the routing context.

Keys are resolved recursively in the tree of context by looking at the parent context if it's missing.

```
app.get ("/", (req, res, next, context) => {
    context["some key"] = "some value";
    return next ();
});

app.get ("/", (req, res, next, context) => {
    var some_value = context["some key"]; // or context.parent["some key"]
});
```

```

    return return res.body.write_all (some_value.data, null);
});

```

## 5.11 Next

Changed in version 0.3: The `next` continuation does not take the request and response objects as parameter. To perform transformation, see [Converters](#) and [Middlewares](#).

The handler takes a callback as an optional third argument. This callback is a continuation that will continue the routing process to the next matching route.

```

app.get ("/", (req, res, next) => {
    return next (); // keep routing
});

app.get ("/", (req, res) => {
    // this is invoked!
});

```

### 5.11.1 Sequence

New in version 0.3.

The [Sequence](#) middleware should be used to chain handling callbacks.

```

app.get ("/", sequence ((req, res, next) => {
    return next ();
}, (req, res) => {
    return res.expand_utf8 ("Hello world!");
}));

```

## 5.12 Error handling

New in version 0.2.1: Prior to this release, any unhandled error would crash the main loop iteration.

Changed in version 0.3: Error and status codes are now handled with a `catch` block or using the [Status](#) middleware.

Changed in version 0.3: The default handling is not ensured by the [Basic](#) middleware.

Changed in version 0.3: Thrown errors are forwarded to VSGL, which process them essentially the same way. See [VSGL](#) for more details.

Similarly to status codes, errors are propagated in the `HandlerCallback` and `NextCallback` delegate signatures and can be handled in a `catch` block.

```

app.use (() => {
    try {
        return next ();
    } catch (IOError err) {
        res.status = 500;
        return res.expand_utf8 (err.message);
    }
});

```

```
app.get ("/", (req, res) => {
    throw new IOError.FAILED ("I/O failed some some reason.");
});
```

Thrown status code can also be caught this way, but it's much more convenient to use the [Status](#) middleware.

## 5.13 Subrouting

Since `VSGI.ApplicationCallback` is type compatible with `HandlerCallback`, it is possible to delegate request handling to another VSGI-compliant application.

In particular, it is possible to treat `Router.handle` like any handling callback.

---

**Note:** This feature is a key design of the router and is intended to be used for a maximum inter-operability with other frameworks based on VSGI.

---

The following example delegates all GET requests to another router which will process in isolation with its own routing context.

```
var app = new Router ();
var api = new Router ();

// delegate all GET requests to api router
app.get ("*", api.handle);
```

One common pattern with subrouting is to attempt another router and fallback on next.

```
var app = new Router ();
var api = new Router ();

app.get ("/some-resource", (req, res) => {
    return api.handle (req, res) || next ();
});
```

## 5.14 Cleaning up route logic

Performing a lot of route bindings can get messy, particularly if you want to split an application several reusable modules. Encapsulation can be achieved by subclassing `Router` and performing initialization in a construct block:

```
public class AdminRouter : Router {

    construct {
        rule (Method.GET, "/admin/user", view);
        rule (Method.GET | Method.POST, "/admin/user/<int:id>", edit);
    }

    public bool view (Request req, Response res) {
        return render_template ("users", Users.all ());
    }

    public bool edit (Request req, Response res) {
        var user = User.find (ctx["id"]);
    }
}
```

```
    if (req.method == "POST") {
        user.values (Soup.Form.decode (req.flatten_utf8 ()));
        user.update ();
    }
    return render_template ("user", user);
}
```

Using subrouting, it can be assembled to a parent router given a rule (or any matching process described in this document). This way, incoming request having the `/admin/` path prefix will be delegated to the `admin` router.

```
var app = new Router ();
app.rule (Method.ALL, "/admin/*", new AdminRouter ().handle);
```

The `Basepath` middleware provide very handy path isolation so that the router can be simply written upon the leading `/` and rebased on any basepath. In that case, we can strip the leading `/admin` in router's rules.

```
var app = new Router ();

// captures '/admin/users' and '/admin/user/<int:id>'
app.use (basepath ("/admin", new AdminRouter ().handle));
```



---

## Redirection and Error

---

Redirection, client and server errors are handled via a simple [exception](#) mechanism.

In a [HandlerCallback](#), you may throw any of [Informational](#), [Success](#), [Redirection](#), [ClientError](#) and [ServerError](#) predefined error domains rather than setting the status and returning from the function.

It is possible to register a handler on the [Router](#) to handle a specific status code.

```
app.use ((req, res, next) => {
  try {
    return next ();
  } catch (Redirection.PERMANENT red) {
    // handle a redirection...
  }
});
```

### 6.1 Default handling

Changed in version 0.3: Default handling is not assured by the [Basic](#) middleware.

The [Router](#) can be configured to handle raised status by setting the response status code and headers appropriately.

```
app.use (basic ());

app.get ("/", () => {
  throw new ClientError.NOT_FOUND ("The request URI '/' was not found.");
});
```

To handle status more elegantly, see the [Status](#) middleware.

```
app.use (status (Status.NOT_FOUND, (req, res, next, ctx, err) => {
  // handle 'err' properly...
}));
```

The error message may be used to fill a specific [Response](#) headers or the response body. The following table describe how the router deal with these cases.

Status	Header	Description
Informational.SWITCHING_PROTOCOLS	Upgrade	Identifier of the protocol to use
Success.CREATED	Location	URL to the newly created resource
Success.PARTIAL_CONTENT	Range	Range of the delivered resource in bytes
Redirection.MOVED_PERMANENTLY	Location	URL to perform the redirection
Redirection.FOUND	Location	URL of the found resource
Redirection.SEE_OTHER	Location	URL of the alternative resource
Redirection.USE_PROXY	Location	URL of the proxy
Redirection.TEMPORARY_REDIRECT	Location	URL to perform the redirection
ClientError.UNAUTHORIZED	WWW-Authenticate	Challenge for authentication
ClientError.METHOD_NOT_ALLOWED	Allow	Comma-separated list of allowed methods
ClientError.UPGRADE_REQUIRED	Upgrade	Identifier of the protocol to use

The following errors does not produce any payload:

- `Information.SWITCHING_PROTOCOLS`
- `Success.NO_CONTENT`
- `Success.RESET_CONTENT`
- `Success.NOT_MODIFIED`

For all other domains, the message will be used as a `text/plain` payload encoded with UTF-8.

The approach taken by Valum is to support at least all status defined by `libsoup-2.4` and those defined in RFC documents. If anything is missing, you can add it and submit us a pull request.

## 6.2 Informational (1xx)

Informational status are used to provide a in-between response for the requested resource. The [Response](#) body must remain empty.

Informational status are enumerated in `Informational` error domain.

## 6.3 Success (2xx)

Success status tells the client that the request went well and provide additional information about the resource. An example would be to throw a `Success.CREATED` error to provide the location of the newly created resource.

Successes are enumerated in `Success` error domain.

```
app.get ("/document/<int:id>", (req, res) => {
  // serve the document by its identifier...
});

app.put ("/document", (req, res) => {
  // create the document described by the request
  throw new Success.CREATED ("/document/%u".printf (id));
});
```



## 6.4 Redirection (3xx)

To perform a redirection, you have to throw a `Redirection` error and use the message as a redirect URL. The `Router` will automatically set the `Location` header accordingly.

Redirections are enumerated in `Redirection` error domain.

```
app.get ("/user/<id>/save", (req, res) => {
  var user = User (req.params["id"]);

  if (user.save ())
    throw new Redirection.MOVED_TEMPORAIRLY ("/user/%u".printf (user.id));
});
```

## 6.5 Client (4xx) and server (5xx) error

Like for redirections, client and server errors are thrown. Errors are predefined in `ClientError` and `ServerError` error domains.

```
app.get ("/not-found", (req, res) => {
  throw new ClientError.NOT_FOUND ("The requested URI was not found.");
});
```

## 6.6 Errors in next

The next continuation is designed to throw these specific errors so that the `Router` can handle them properly.

```
app.use ((req, res, next) => {
  try {
    return next ();
  } catch (ClientError.NOT_FOUND err) {
    // handle a 404...
  }
});

app.get ("/", (req, res, next) => {
  return next (); // will throw a 404
});

app.get ("/", (req, res) => {
  throw new ClientError.NOT_FOUND ("");
});
```



---

## Middlewares

---

Middlewares are reusable pieces of processing that can perform various work from authentication to the delivery of a static resource.

### 7.1 Authenticate

The `authenticate` middleware allow one to perform HTTP basic and digest authentication.

It takes three parameters:

- an `Authentication` object described in [HTTP authentication](#)
- a callback to challenge a user-provided `Authorization` header
- a forward callback invoked on success with the corresponding authorization object

If the authentication fails, a 401 `Unauthorized` status is raised with a `WWW-Authenticate` header.

```
app.use (authenticate (new BasicAuthentication ("realm")), (authorization) => {
    return authorization.challenge ("some password");
}, (req, res, next, ctx, username) => {
    return res.expand_utf8 ("Hello %s".printf (username));
});
```

To perform custom password comparison, it is best to cast the `authorization` parameter and access the password directly.

```
public bool authenticate_user (string username, string password) {
    // authenticate the user against the database...
}

app.use (authenticate (new BasicAuthentication ("realm")), (authorization) => {
    var basic_authorization = authorization as BasicAuthorization;
    return authenticate_user (basic_authorization.username, basic_authorization.password);
});
```

### 7.2 Basepath

The `basepath` middleware allow a better isolation when composing routers by stripping a prefix on the `Request` URI.

The middleware strips and forwards requests which match the provided base path. If the resulting path is empty, it fallbacks to a root `/`.

Error which use their message as a `Location` header are automatically prefixed by the base path.

```
var user = new Router ();

user.get ("/<int:id>", (req, res) => {
  // ...
});

user.post ("/", (req, res) => {
  throw new Success.CREATED ("/5");
});

app.use (basepath ("/user", user.handle));

app.status (Soup.Status.CREATED, (req, res) => {
  assert ("/user/5" == context["message"]);
});
```

If `next` is called while forwarding or an error is thrown, the original path is restored.

```
user.get ("/<int:id>", (req, res, next) => {
  return next (); // path is '/5'
});

app.use (basepath ("/user", user.handle));

app.use ((req, res) => {
  // path is '/user/5'
});
```

One common pattern is to provide a path-based fallback when using the [Subdomain](#) middleware.

```
app.use (subdomain ("api", api.handle));
app.use (basepath ("/api", api.handle));
```

## 7.3 Basic

New in version 0.3.

Previously know under the name of *default handling*, the `basic` middleware provide a conforming handling of raised status codes as described in the [Redirection and Error](#) document.

It aims at providing sane defaults for a top-level middleware.

```
app.use (basic ());

app.get ("/", () => {
  throw new Success.CREATED ("/resource/id");
});
```

If an error is caught, it will perform the following tasks:

1. assign an appropriate status code (500 for other errors)
2. setup required headers (eg. *Location* for a redirection)
3. produce a payload based on the message if required and not already used for a header

The payload will have the `text/plain` content type encoded with UTF-8.

For privacy and security reason, non-status errors (eg. `IOError`) will not be used for the payload. To enable that for specific errors, it's possible to convert them into into a raised status, preferably a 500 Internal Server Error.

```
app.use (() => {
  try {
    return next ();
  } catch (IOError err) {
    throw new ServerError.INTERNAL_SERVER_ERROR (err.message);
  }
})
```

## 7.4 Content Negotiation

Negotiating the resource representation is an essential part of the HTTP protocol.

The negotiation process is simple: expectations are provided for a specific header, if they are met, the processing is forwarded with the highest quality value, otherwise a 406 Not Acceptable status is raised.

```
using Valum.ContentNegotiation;

app.get ("/", negotiate ("Accept", "text/html, text/html+xml",
  (req, res, next, stack, content_type) => {
    // produce a response based on 'content_type'
  }));
```

Or directly by using the default forward callback:

```
app.use (negotiate ("Accept", "text/html"));

// all route declaration may assume that the user agent accept 'text/html'
```

### 7.4.1 Preference and quality

Additionally, the server can state the quality of each expectation. The middleware will maximize the product of quality and user agent preference with respect to the order of declaration and user agent preferences if it happens to be equal.

If, for instance, you would serve a XML document that is just poorly converted from a JSON source, you could state it by giving it a low `q` value. If the user agent as a strong preference the former and a low preference for the latter (eg. `Accept: text/xml; application/json; q=0.1`), it will be served the version with the highest product (eg.  $0.3 * 1 > 1 * 0.3$ ).

```
app.get ("/", negotiate ("Accept", "application/json; text/xml; q=0.3",
  (req, res, next, stack, content_type) => {
    // produce a response based on 'content_type'
  }));
```

### 7.4.2 Error handling

The `Status` middleware may be used to handle the possible 406 Not Acceptable error raised if no expectation can be satisfied.

```
app.use (status (Soup.Status.NOT_ACCEPTABLE, () => {
    // handle '406 Not Acceptable' here
}));

app.use (negotiate ("Accept", "text/xhtml; text/html", () => {
    // produce appropriate resource
}));
```

### 7.4.3 Custom comparison

A custom comparison function can be provided to `negotiate` in order to handle wildcards and other edge cases. The user agent pattern is the first argument and the expectation is the second.

**Warning:** Most of the HTTP/1.1 specification about headers is case-insensitive, use `Soup.str_case_equal` to perform comparisons.

```
app.use (negotiate ("Accept",
    "text/xhtml",
    () => { return true; },
    (a, b) => {
        return a == "*" || Soup.str_case_equal (a, b);
    }
));
```

### 7.4.4 Helpers

For convenience, helpers are provided to handle common headers:

Middleware	Header	Edge cases
<code>accept</code>	Content-Type	<code>*/*, type/*</code> and <code>type/subtype1+subtype2</code>
<code>accept_charset</code>	Content-Type	<code>*</code>
<code>accept_encoding</code>	Content-Encoding	<code>*</code>
<code>accept_language</code>	Content-Language	missing language type
<code>accept_ranges</code>	Content-Ranges	none

The `accept` middleware will assign the media type and preserve all other parameters.

If multiple subtypes are specified (e.g. `application/vnd.api+json`), the middleware will check if the subtypes accepted by the user agent form a subset. This is useful if you serve a specified JSON document format to a client which only state to accept JSON and does not care about the specification itself.

```
accept ("text/html; text/xhtml", (req, res, next, ctx, content_type) => {
    switch (content_type) {
        case "text/html":
            return produce_html ();
        case "text/xhtml":
            return produce_xhtml ();
    }
});
```

The `accept_encoding` middleware will convert the `Response` if it's either `gzip` or `deflate`.

```
accept ("gzip; deflate", (req, res, next, ctx, encoding) => {
    res.expand_utf8 ("Hello world! (compressed with %s)".printf (encoding));
});
```

The `accept_charset` middleware will set the `charset` parameter of the `Content-Type` header, defaulting to `application/octet-stream` if undefined.

## 7.5 Decode

The `decode` middleware is used to unapply various content codings.

```
app.use (decode ());

app.post ("/", (req, res) => {
  var posted_data = req.flatten_utf8 ();
});
```

It is typically put at the top of an application.

Encoding	Action
deflate	<a href="#">Glib.ZlibDecompressor</a>
gzip and x-gzip	<a href="#">Glib.ZlibDecompressor</a>
identity	nothing

If an encoding is not supported, a `501 Not Implemented` is raised and remaining encodings are *reapplied* on the request.

To prevent this behavior, the `DecodeFlags.FORWARD_REMAINING_ENCODINGS` flag can be passed to forward unsupported content codings.

```
app.use (decode (DecodeFlags.FORWARD_REMAINING_ENCODINGS));

app.use (() => {
  if (req.headers.get_one ("Content-Encoding") == "br") {
    req.headers.remove ("Content-Encoding");
    req.convert (new BrotliDecompressor ());
  }
  return next ();
});

app.post ("/", (req, res) => {
  var posted_data = req.flatten_utf8 ();
});
```

## 7.6 Safely

Yet very simple, the `safely` middleware provide a powerful way of discovering possible error conditions and handle them locally.

Only status defined in [Redirection and Error](#) are leaked: the compiler will warn for all other unhandled errors.

```
app.get ("/", safely ((req, res, next, ctx) => {
  try {
    res.expand_utf8 ("Hello world!");
  } catch (IOError err) {
    critical (err.message);
    return false;
  }
});
```

## 7.7 Sequence

New in version 0.3.

The sequence middleware provide a handy way of chaining middlewares.

```
app.post ("/", sequence (decode (), (req, res) => {  
    // handle decoded payload  
}));
```

To chain more than two middlewares, one can chain a middleware with a sequence.

```
app.get ("/admin", sequence ((req, res, next) => {  
    // authenticate user...  
    return next ();  
}, sequence ((req, res, next) => {  
    // produce sensitive data...  
    return next ();  
}, (req, res) => {  
    // produce the response  
}));
```

Vala does not support varidic delegate arguments, which would be much more convenient to describe a sequence.

## 7.8 Server-Sent Events

Valum provides a middleware for the HTML5 [Server-Sent Events](#) protocol to stream notifications over a persistent connection.

The `ServerSentEvents.stream_events` function creates a handling middleware and provide a `send` call-back to transmit the actual events.

```
using Valum;  
using Valum.ServerSentEvents;  
  
app.get ("sse", stream_events ((req, send) => {  
    send (null, "some data");  
}));
```

```
var eventSource = new EventSource ("/sse");  
  
eventSource.onmessage = function(message) {  
    console.log (message.data); // displays 'some data'  
};
```

### 7.8.1 Multi-line messages

Multi-line messages are handled correctly by splitting the data into into multiple `data: chunks`.

```
send (null, "some\nndata");
```

```
data: some  
data: data
```



## 7.9 Static Resource Delivery

Middlewares in the `Valum.Static` namespace ensure delivery of static resources.

```
using Valum.Static;
```

As of convention, all middleware use the `path` context key to resolve the resource to be served. This can easily be specified using a rule parameter with the `path` type.

For more flexibility, one can compute the `path` value and pass the control with `next`. The following example obtain the key from the HTTP query:

```
app.get ("/", sequence ((req, res, next, ctx) => {
    ctx["path"] = req.lookup_query ("path") ?? "index.html";
    return next ();
}), serve_from_file (File.new_for_uri ("resource://")));
```

If a resource is not available (eg. the file does not exist), the control will be forwarded to the next route. The `Sequence` middleware should be used to turn that behaviour into a 404 Not Found.

```
app.get ("/", sequence (serve_from_file (File.new_for_uri ("resource://")),
    (req, res, next, ctx) => {
        throw new ClientError.NOT_FOUND ("The static resource '%s' were not found.", ctx["path"]);
    }));
```

### 7.9.1 GFile

The `serve_from_file` middleware will serve resources relative to a `GLib.File` instance.

```
app.get ("/static/<path:path>", serve_from_file (File.new_for_path ("static")));
```

To deliver from the global resources, use the `resource://` scheme.

```
app.get ("/static/<path:path>", serve_from_file (File.new_for_uri ("resource://static")));
```

Before being served, each file is forwarded to make it possible to modify headers more specifically or raise a last-minute error.

Once done, invoke the `next` continuation to send over the content.

```
app.get ("/static/<path:path>", serve_from_file (File.new_for_path ("static"),
    ServeFlags.NONE,
    (req, res, next, ctx, file) => {
        var user = ctx["user"] as User;
        if (!user.can_access (file)) {
            throw new ClientError.FORBIDDEN ("You cannot access this file.")
        }
        return next ();
    }));
```

The `ServeFlags.X_SENDFILE` will only work for locally available files, meaning that `GLib.File.get_path` is non-null.

### 7.9.2 Resource bundle

The `serve_from_resource` middleware is provided to serve a resource bundle (see `GLib.Resource`) from a given prefix.

```
app.get ("/static/<path:path>", Static.serve_from_resource (Resource.load ("resource"),
                                                         "/static/"));
```

The `ServeFlags.ENABLE_LAST_MODIFIED` is not supported since the necessary information cannot be determined.

The `ServeFlags.X_SENDFILE` is not supported since contained resources are not represented on the filesystem.

## 7.9.3 Options

### ETag

If the `ServeFlags.ENABLE_ETAG` is specified, a checksum of the resource will be generated in the ETag header.

If set and available, it will have precedence over `ServeFlags.ENABLE_LAST_MODIFIED` described below.

### Last-Modified

Unlike ETag, this caching feature is time-based and will indicate the last modification on the resource. This is only available for some GFile backend and will fallback to ETag if enabled as well.

Specify the `ServeFlags.ENABLE_LAST_MODIFIED` to enable this feature.

### X-Sendfile

If the application run under a HTTP server, it might be preferable to let it serve static resources directly.

### Public caching

The `ServeFlags.ENABLE_PUBLIC` let intermediate HTTP servers cache the payload.

### Expose missing permissions

The `ServeFlags.FORBID_ON_MISSING_RIGHTS` will trigger a 403 Forbidden if rights are missing to read a file. This is not a default as it may expose information about the existence of certain files.

## 7.10 Status

Thrown status codes (see [Redirection and Error](#)) can be handled with the `status` middleware.

The received [Request](#) and [Response](#) object are in the same state they were when the status was thrown. An additional parameter provide access to the actual `GLib.Error` object.

```
app.use (status (Soup.Status.NOT_FOUND, (req, res, next, context, err) => {
  // produce a 404 page...
  var message = err.message;
});
```

To jump to the next status handler found upstream in the routing queue, just throw the error. If the error can be resolved, you might want to try `next` once more.

```

app.status (Soup.Status.NOT_FOUND, (req, res) => {
  res.status = 404;
  return res.expand_utf8 ("Not found!");
});

app.status (Soup.Status.NOT_FOUND, (req, res, next, ctx, err) => {
  return next (); // try to route again or jump upstream
});

app.use (() => {
  throw new ClientError.NOT_FOUND ("");
});

```

If an error is not handled, it will eventually be caught by the default status handler, which produce a minimal response.

```

// turns any 404 into a permanent redirection
app.status (Soup.Status.NOT_FOUND, (req, res) => {
  throw new Redirection.PERMANENT ("http://example.com");
});

```

## 7.11 Subdomain

The `subdomain` middleware matches [Request](#) which subdomain is conform to expectations.

---

**Note:** Domains are interpreted in their semantical right-to-left order and matched as suffix.

---

The pattern is specified as the first argument. It may contain asterisk `*` which specify that any supplied label satisfy that position.

```

app.use (subdomain ("api", (req, res) => {
  // match domains like 'api.example.com' and 'v1.api.example.com'
}));

app.use (subdomain ("*.user", (req, res) => {
  // match at least two labels: the first can be anything and the second
  // is exactly 'user'
}));

```

The matched subdomain labels are extracted and passed by parameter.

```

app.use (subdomain ("api", (req, res, next, ctx, subdomains) => {
  // 'subdomains' could be 'api' or 'v1.api'
}));

```

This middleware can be used along with subrouting to mount any [Router](#) on a specific domain pattern.

```

var app = new Router ();
var api = new Router ();

app.use (subdomain ("api", api.handle));

```

### 7.11.1 Strict

There is two matching mode: loose and strict. The loose mode only expect the request to be performed on a suffix-compatible hostname. For instance, `api` would match `api.example.com` and `v1.api.example.com` as well.

To prevent this and perform a `_strict_` match, simply specify the second argument. The domain of the request will have to supply exactly the same amount of labels matching the expectations.

```
// match every request exactly from 'api.*.*'  
app.use (subdomain ("api", api.handle, SubdomainFlags.STRICT));
```

### 7.11.2 Skip labels

By default, the two first labels are ignored since web applications are typically served under two domain levels (eg. `example.com`). If it's not the case, the number of skipped labels can be set to any desirable value.

```
// match exactly 'api.example.com'  
app.use (subdomain ("api.example.com", api.handle, SubdomainFlags.STRICT, 0));
```

The typical way of declaring them involve closures. It is parametrized and returned to perform a specific task:

```
public HandlerCallback middleware (/* parameters here */) {  
    return (req, res, next, ctx) => {  
        var referer = req.headers.get_one ("Referer");  
        ctx["referer"] = new Soup.URI (referer);  
        return next ();  
    };  
}
```

The following example shows a middleware that provide a compressed stream over the [Response](#) body.

```
app.use ((req, res, next) => {  
    res.headers.append ("Content-Encoding", "gzip");  
    res.convert (new ZLibCompressor (ZlibCompressorFormat.GZIP));  
    return next ();  
});  
  
app.get ("/home", (req, res) => {  
    return res.expand_utf8 ("Hello world!"); // transparently compress the output  
});
```

If this is wrapped in a function, which is typically the case, it can even be used directly from the handler.

```
HandlerCallback compress = (req, res, next) => {  
    res.headers.append ("Content-Encoding", "gzip");  
    res.convert (new ZLibCompressor (ZlibCompressorFormat.GZIP));  
    return next ();  
};  
  
app.get ("/home", compress);  
  
app.get ("/home", (req, res) => {  
    return res.expand_utf8 ("Hello world!");  
});
```

Alternatively, a middleware can be used directly instead of being attached to a `Route`, the processing will happen in a `NextCallback`.

```
app.get ("/home", (req, res, next, context) => {
    return compress (req, res, (req, res) => {
        return res.expand_utf8 ("Hello world!");
    }, new Context.with_parent (context));
});
```

## 7.12 Forward

New in version 0.3.

One typical middleware pattern is to take a continuation that is forwarded on success (or any other event) with a single value like it's the case for the [Content Negotiation](#) middlewares.

This can be easily done with `ForwardCallback<T>`. The generic parameter specify the type of the forwarded value.

```
public HandlerCallback accept (string content_types, ForwardCallback<string> forward) {
    return (req, res, next, ctx) => {
        // perform content negotiation and determine 'chosen_content_type'...
        return forward (req, res, next, ctx, chosen_content_type);
    };
}

app.get ("/", accept ("text/xml; application/json", (req, res, next, ctx, content_type) => {
    // produce a response according to 'content_type'...
}));
```

Often, one would simply call the `next` continuation, so a `forward` definition is provided to do that. It is used as a default value for various middlewares such that all the following examples are equivalent:

```
app.use (accept ("text/html" () => {
    return next ();
}));

app.use (accept ("text/html", forward));

app.use (accept ("text/html"));
```

To pass multiple values, it is preferable to explicitly declare them using a delegate.

```
public delegate bool ComplexForwardCallback (Request req,
                                             Response res,
                                             NextCallback next,
                                             Context ctx,
                                             int a,
                                             int b) throws Error;
```



---

## Recipes

---

Recipes are documents providing approaches to common web development tasks and their potential integration with Valum.

### 8.1 Configuration

There exist various way of providing a runtime configuration.

If you need to pass secrets, take a look at the [Libsecret](#) project. It allows one to securely store and retrieve secrets: just unlock the keyring and start your service.

#### 8.1.1 Key file

GLib provide a very handy way of reading and parsing [key files](#), which are widely used across freedesktop specifications.

It should be privileged if the configuration is mostly edited by humans.

```
[app]
public-dir=public

[database]
provider=mysql
connection=
auth=
```

```
using GLib;
using Valum;

var config = new KeyFile ();

config.parse_path ("app.conf");

var app = new Router ();

app.get ("/public/<path:path>",
        Static.serve_from_path (config.get_string ("app", "public-dir")));
```

## 8.1.2 JSON

The `JSON-GLib` project provide a really convenient JSON parser and generator.

```
{
  "app": {
    "publicDir": "public"
  },
  "database": {
    "provider": "mysql",
    "connection": "",
    "auth": ""
  }
}
```

```
using Json;
using Valum;

var parser = new Parser ();
parser.parse_from_file ("config.json");

var config = parser.get_root ();

var app = new Router ();

app.get ("/public/<path:path>",
        Static.serve_from_path (config.get_object ("app").get_string_member ("publicDir")));
```

## 8.1.3 YAML

There is a `GLib wrapper` around `libyaml` that makes it more convenient to use. `YAML` in itself can be seen as a human-readable JSON format.

```
app:
  publicDir: public
database:
  provider: mysql
  connection:
  auth:
```

```
using Valum;
using Yaml;

var config = new Document.from_path ("config.yml").root as Node.Mapping;

var app = new Router ();

app.get ("/public/<path:path>",
        Static.serve_from_path (config.get_mapping ("app").get_scalar ("publicDir").value));
```

## 8.1.4 Other approaches

The following approaches are a bit more complex to setup but can solve more specific use cases:

- `GXml` or `libxml2`
- `GSettings` for a remote (via `DBus`) and monitorable configuration



- environment variables via `GLib.Environment` utilities
- CLI options (see `VSGI.Server.add_main_option` and `VSGI.Server.handle_local_options`)

## 8.2 JSON

JSON is a popular data format for web services and `json-glib` provide a complete implementation that integrates with the `GObject` type system.

The following features will be covered in this document with code examples:

- serialize a `GObject`
- unserialize a `GObject`
- parse an `GLib.InputStream` of JSON like a `Request` body
- generate JSON in a `GLib.OutputStream` like a `Response` body

### 8.2.1 Produce and stream JSON

Using a `Json.Generator`, you can conveniently produce an JSON object and stream synchronously it in the `Response` body.

```
app.get ("/user/<username>", (req, res) => {
  var user      = new Json.Builder ();
  var generator = new Json.Generator ();

  user.set_member_name ("username");
  user.add_string_value (req.params["username"]);

  generator.root  = user.get_root ();
  generator.pretty = false;

  return generator.to_stream (res.body);
});
```

### 8.2.2 Serialize GObject

Your project is likely to have a model abstraction and serialization of `GObject` with `Json.gobject_serialize` is a handy feature. It will recursively build a JSON object from the encountered properties.

```
public class User : Object {
  public string username { construct; get; }

  public User.from_username (string username) {
    // populate the model from the data storage...
  }

  public void update () {
    // persist the model in data storage...
  }
}
```

```

app.get ("/user/<username>", (req, res) => {
    var user      = new User.from_username (req.params["username"]);
    var generator = new Json.Generator ();

    generator.root  = Json.gobject_serialize (user);
    generator.pretty = false;

    return generator.to_stream (res.body);
});

```

With middlewares, you can split the process in multiple reusable steps to avoid code duplication. They are described in the [Router](#) document.

- fetch a model from a data storage
- process the model with data obtained from a `Json.Parser`
- produce a JSON response with `Json.gobject_serialize`

```

app.scope ("/user", (user) => {
    // fetch the user
    app.rule (Method.GET | Method.POST, "<username>", (req, res, next, context) => {
        var user = new User.from_username (context["username"].get_string ());

        if (!user.exists ()) {
            throw new ClientError.NOT_FOUND ("no such user '%s'", context["username"]);
        }

        context["user"] = user;
        return next ();
    });

    // update model data
    app.post ("/<username>", (req, res, next, context) => {
        var username = context["username"].get_string ();
        var user      = context["user"] as User;
        var parser    = new Json.Parser ();

        // whitelist for allowed properties
        string[] allowed = {"username"};

        // update the model when members are read
        parser.object_member.connect ((obj, member) => {
            if (member in allowed)
                user.set_property (member,
                                   obj.get_member (member).get_value ());
        });

        if (!parser.load_from_stream (req.body))
            throw new ClientError.BAD_REQUEST ("unable to parse the request body");

        // persist the changes
        user.update ();

        if (user.username != username) {
            // model location has changed, so we throw a 201 CREATED status
            throw new Success.CREATED ("/user/%s".printf (user.username));
        }

        context["user"] = user;
    });
});

```

```

    return next ();
});

// serialize to JSON any provided GObject
app.rule (Method.GET, "*", (req, res, next, context) => {
    var generator = new Json.Generator ();

    generator.root    = Json.gobject_serialize (context["user"].get_object ());
    generator.pretty  = false;

    res.headers.set_content_type ("application/json", null);

    return generator.to_stream (res.body);
});
});

```

It is also possible to use `Json.Parser.load_from_stream_async` and invoke `next` in the callback with `Router` `invoke` function if you are expecting a considerable user input.

```

parser.load_from_stream_async.begin (req.body, null, (obj, result) => {
    var success = parser.load_from_stream_async.end (result);

    user.update ();

    context["user"] = user;

    // execute 'next' in app context
    return app.invoke (req, res, next);
});

```

## 8.3 Persistence

Multiple persistence solutions have bindings in Vala and can be used by Valum.

- `libgda` for relational databases and more
- `memcached`
- `redis-glib`
- `mongodb-glib`
- `couchdb-glib` which is supported by the Ubuntu team

One good general approach is to use a per-process connection pool since handlers are executing in asynchronous context, your application will greatly benefit from multiple connections.

### 8.3.1 Memcached

You can use `libmemcached.vapi` to access a Memcached cache storage, it is maintained in `nemequ/vala-extra-vapis` GitHub repository.

```

using Valum;
using VSGI;

var app      = new Router ();

```

```

var memcached = new Memcached.Context ();

app.get ("/<key>", (req, res) => {
    var key = req.params["key"];

    int32 flags;
    Memcached.ReturnCode error;
    var value = memcached.get ("hello", out flags, out error);

    return res.expand (value, null);
});

app.post ("/<key>", (req, res) => {
    var key = req.params["key"];
    var buffer = new MemoryStream.resizable ();

    // fill the buffer with the request body
    buffer.splice (req);

    int32 flags;
    Memcached.ReturnCode error;
    var value = memcached.get ("hello", out flags, out error);

    return res.expand (value, null);
});

Server.new_with_application ("http", "org.valum.example.Memcached", app.handle);

```

## 8.4 Resources

GLib provides a powerful [resource api](#) for bundling static resources and optionally link them in the executable.

It has a few advantages:

- resources can be compiled in the text segment of the executable, providing lightning fast loading time
- resource api is simpler than file api and avoids IOError handling
- application do not have to deal with its resource location or minimally if a separate bundle is used

This only applies to small and static resources as it will grow the size of the executable. Also, if the resources are compiled in your executable, changing them will require a recompilation.

Middlewares are provided for that purpose, see [./middlewares/static](#) for more details.

### 8.4.1 Integration

Let's say your project has a few resources:

- CTPL templates in a `templates` folder
- CSS, JavaScript files in `static` folder

Setup a `app.gresource.xml` file that defines what resources will to be bundled.

```

<?xml version="1.0" encoding="UTF-8"?>
<gresources>
  <gresource>

```

```

<file>templates/home.html</file>
<file>templates/404.html</file>
<file>static/css/bootstrap.min.css</file>
</gresource>
</gresources>

```

You can test your setup with:

```
glib-compile-resource app.gresource.xml
```

Latest version of `waf` automatically link `*.gresource.xml` if you load the `glib2` plugin and add the file to your sources.

```

bld.load('glib2')

bld.program(
    packages = ['valum-0.1'],
    target   = 'app',
    source   = bld.path.ant_glob('**/*.vala') + ['app.gresource.xml'],
    uselib   = ['VALUM'])

```

The `app` example serves its static resources this way if you need a code reference.

## 8.5 Scripting

Through `Vala VAPI bindings`, application written with Valum can embed multiple interpreters and JIT to provide facilities for computation and templating.

### 8.5.1 Lua

`luajit` ships with a VAPI you can use to access a Lua VM, just add `--pkg lua` to `valac`.

```
valac --pkg valum-0.1 --pkg lua app.vala
```

```

require 'markdown'
return markdown('## Hello from lua.eval!')

```

```

using Valum;
using VSGI;
using Lua;

var app = new Router ();
var lua = new LuaVM ();

// GET /lua
app.get ("/lua", (req, res) => {
    // evaluate a string containing Lua code
    res.expand_utf8 (some_lua_code, null);

    // evaluate a file containing Lua code
    return res.expand_utf8 (lua.do_file ("scripts/hello.lua"));
});

Server.new_with_application ("http", "org.valum.example.Lua", app.handle).run ();

```

The sample Lua script contains:

```
require 'markdown'  
return markdown("# Hello from Lua!!!")  
-- returned value will be appended to response body
```

Resulting response

```
<h1>Hello from Lua!!!</h1>
```

## 8.5.2 Scheme (TODO)

Scheme can be used to produce template or facilitate computation.

```
app.get ("/hello.scm", (req, res) => {  
  return res.expand_utf8 (scm.run ("scripts/hello.scm"));  
});
```

Scheme code:

```
;; VALUM_ROOT/scripts/hello.scm  
(+ 1 2 3)  
;; returned value will be casted to string  
;; and appended to response body
```

## 8.6 Templating

Template engines are very important tools to craft Web applications and a few libraries exist to handle that tedious work.

### 8.6.1 Compose

For HTML5, [Compose](#) is quite appropriate.

```
app.get ("/", (req, res) => {  
  return res.expand_utf8 (  
    html ({,  
      head ({,  
        title (),  
      },  
      body ({,  
        section (  
          h1 ({, "Section Title"}))));  
    });
```

It comes with two utilities: `take` and `when` to iterate and perform conditional evaluation.

```
var users = Users.all ();  
  
take<User> (() => { return users.next (); },  
  (user) => { return user.username; });  
  
when (User.current ().is_admin,  
  () => { return p ({, "admin" } ),  
  () => { return p ({, "user" } )});
```

Strings are not escaped by default due to the design of the library. Instead, all unsafe value must be escaped properly. For HTML, `e` is provided.

```
e (user.biography);
```

Templates and fragments can be store in Vala source files to separate concerns. In this case, arguments would be used to pass the environment.

```
using Compose.HTML5;

namespace Project.Templates
{
    public string page (string title, string content)
    {
        return
            div ({"id=%s".printf (title)},
                h2 ({}, e (title)),
                content);
    }
}
```





This document addresses hackers who wants to get involved in the framework development.

## 9.1 Building

You can avoid installing the library if you export `LD_LIBRARY_PATH` to the `build` folder.

```
export LD_LIBRARY_PATH=build
./waf configure
./waf build
./build/tests/tests
```

## 9.2 Code conventions

Valum uses the [Vala compiler coding style](#) and these rules are specifically highlighted:

- tabs for indentation
- spaces for alignment
- 80 characters for comment block and 120 for code
- always align blocks of assignation around = sign
- remember that little space between a function name and its arguments
- doclets should be aligned, grouped and ordered alphabetically

## 9.3 General strategies

Produce minimal headers, especially if the response has an empty body as every byte will count.

Since `GET` handle `HEAD` as well, verifying the request method to prevent spending time on producing a body that won't be considered is important.

```
res.headers.set_content_type ("text/html", null);

if (req.method == "HEAD") {
    size_t bytes_written;
    return res.write_head (out bytes_written);
}

return res.expand_utf8 ("<!DOCTYPE html><html></html>");
```

Use the `construct` block to perform post-initialization work. It will be called independently of how the object is constructed.

## 9.4 Tricky stuff

Many parts of the HTTP/1.1 specification is case-insensitive, in these cases, `Soup.str_case_equal` must be used to perform comparisons.

Try to stay by the book and read carefully the specification to ensure that the framework is semantically correct. In particular, the following points:

- choice of a status code
- method is case-sensitive
- URI and query are automatically decoded by `Soup.URI`
- headers and their parameters are case-insensitive
- `\r\n` are used as newlines
- do not handle `Transfer-Encoding`, except for the `libsoup-2.4` implementation with `steal_connection`: at this level, it's up to the HTTP server to perform the transformation

The framework should rely as much as possible upon `libsoup-2.4` to ensure consistent and correct behaviours.

## 9.5 Coverage

`gcov` is used to measure coverage of the tests on the generated C code. The results are automatically uploaded to `Codecov` on a successful build.

You can build Valum with `gcov` if you specify the appropriate `CFLAGS` and `VALAFLAGS` during the configuration step.

```
./waf configure CFLAGS='-fprovide-arcs -ftest-coverage' VALAFLAGS='--debug'
./waf build
./waf --alltests
```

During the execution of the tests, some files will be generated and can be inspected with the `gcov` utility.

```
cd build
gcov src/router.c.1.gcda
```

Would output something like:

```
File 'src/router.c'
Executed lines: 57.83% of 792
Creating 'router.c.gcov'
```

The generated `router.c.gcov` will contain detailed coverage information structured in three columns separated by a `:` character:

- number of executions
- line number
- corresponding line of code

The number of executions can take the following values:

- a `-` symbol means that the line is irrelevant (eg. comment)
- `#####` means that the line is uncovered
- a positive integer indicates how many time the line has executed

Once you have identified an uncovered region, you can supply a test that covers that particular case and submit us a [pull request on GitHub](#).

## 9.6 Tests

Valum is thoroughly tested for regression with the `GLib.Test` framework. Test cases are annotated with `@since` to track when a behaviour was introduced and guarantee its backward compatibility.

You can refer an issue from GitHub by calling `Test.bug` with the issue number.

```
Test.bug ("123");
```

## 9.7 Version bump

Most of the version substitutions is handled during the build, but some places in the code have to be updated manually:

- `VERSION` and `API_VERSION` in `wscript`
- `GIR` version annotations for all declared namespaces
- `version` and `release` in `docs/conf.py`

## 9.8 Translations

To generate POT files, simply launch the following command:

```
ninja docs/pot
```

Then, move to the `docs` folder to update the translations for the language `<lang>`:

```
sphinx-intl update -p ../build/docs/pot -l <lang>
```

This should update the files located in `po/<lang>/LC_MESSAGES`. The docs can then be regenerated with the translations.

```
ninja docs/<lang>
```



---

## GNU Lesser General Public License

---

*Version 3, 29 June 2007 Copyright © 2007 Free Software Foundation, Inc <<http://fsf.org/>>*

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

This version of the GNU Lesser General Public License incorporates the terms and conditions of version 3 of the GNU General Public License, supplemented by the additional permissions listed below.

### 10.1 0. Additional Definitions

As used herein, “this License” refers to version 3 of the GNU Lesser General Public License, and the “GNU GPL” refers to version 3 of the GNU General Public License.

“The Library” refers to a covered work governed by this License, other than an Application or a Combined Work as defined below.

An “Application” is any work that makes use of an interface provided by the Library, but which is not otherwise based on the Library. Defining a subclass of a class defined by the Library is deemed a mode of using an interface provided by the Library.

A “Combined Work” is a work produced by combining or linking an Application with the Library. The particular version of the Library with which the Combined Work was made is also called the “Linked Version”.

The “Minimal Corresponding Source” for a Combined Work means the Corresponding Source for the Combined Work, excluding any source code for portions of the Combined Work that, considered in isolation, are based on the Application, and not on the Linked Version.

The “Corresponding Application Code” for a Combined Work means the object code and/or source code for the Application, including any data and utility programs needed for reproducing the Combined Work from the Application, but excluding the System Libraries of the Combined Work.

### 10.2 1. Exception to Section 3 of the GNU GPL

You may convey a covered work under sections 3 and 4 of this License without being bound by section 3 of the GNU GPL.

## 10.3 2. Conveying Modified Versions

If you modify a copy of the Library, and, in your modifications, a facility refers to a function or data to be supplied by an Application that uses the facility (other than as an argument passed when the facility is invoked), then you may convey a copy of the modified version:

- **a) under this License, provided that you make a good faith effort to** ensure that, in the event an Application does not supply the function or data, the facility still operates, and performs whatever part of its purpose remains meaningful, or
- **b) under the GNU GPL, with none of the additional permissions of** this License applicable to that copy.

### ### 3. Object Code Incorporating Material from Library Header Files

The object code form of an Application may incorporate material from a header file that is part of the Library. You may convey such object code under terms of your choice, provided that, if the incorporated material is not limited to numerical parameters, data structure layouts and accessors, or small macros, inline functions and templates (ten or fewer lines in length), you do both of the following:

- **a) Give prominent notice with each copy of the object code that the** Library is used in it and that the Library and its use are covered by this License.
- **b) Accompany the object code with a copy of the GNU GPL and this license** document.

## 10.4 4. Combined Works

You may convey a Combined Work under terms of your choice that, taken together, effectively do not restrict modification of the portions of the Library contained in the Combined Work and reverse engineering for debugging such modifications, if you also do each of the following:

- **a)** Give prominent notice with each copy of the Combined Work that the Library is used in it and that the Library and its use are covered by this License.
- **b)** Accompany the Combined Work with a copy of the GNU GPL and this license document.
- **c)** For a Combined Work that displays copyright notices during execution, include the copyright notice for the Library among these notices, as well as a reference directing the user to the copies of the GNU GPL and this license document.
- **d)** Do one of the following:
  - **0)** Convey the Minimal Corresponding Source under the terms of this License, and the Corresponding Application Code in a form suitable for, and under terms that permit, the user to recombine or relink the Application with a modified version of the Linked Version to produce a modified Combined Work, in the manner specified by section 6 of the GNU GPL for conveying Corresponding Source.
  - **1)** Use a suitable shared library mechanism for linking with the Library. A suitable mechanism is one that **(a)** uses at run time a copy of the Library already present on the user's computer system, and **(b)** will operate properly with a modified version of the Library that is interface-compatible with the Linked Version.
- **e)** Provide Installation Information, but only if you would otherwise be required to provide such information under section 6 of the GNU GPL, and only to the extent that such information is necessary to install and execute a modified version of the Combined Work produced by recombining or relinking the Application with a modified version of the Linked Version. (If you use option **4d0**, the Installation Information must accompany the Minimal Corresponding Source and Corresponding Application Code. If you use option **4d1**, you must provide the Installation Information in the manner specified by section 6 of the GNU GPL for conveying Corresponding Source.)

## 10.5 5. Combined Libraries

You may place library facilities that are a work based on the Library side by side in a single library together with other library facilities that are not Applications and are not covered by this License, and convey such a combined library under terms of your choice, if you do both of the following:

- **a)** Accompany the combined library with a copy of the same work based on the Library, uncombined with any other library facilities, conveyed under the terms of this License.
- **b)** Give prominent notice with the combined library that part of it is a work based on the Library, and explaining where to find the accompanying uncombined form of the same work.

## 10.6 6. Revised Versions of the GNU Lesser General Public License

The Free Software Foundation may publish revised and/or new versions of the GNU Lesser General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Library as you received it specifies that a certain numbered version of the GNU Lesser General Public License “or any later version” applies to it, you have the option of following the terms and conditions either of that published version or of any later version published by the Free Software Foundation. If the Library as you received it does not specify a version number of the GNU Lesser General Public License, you may choose any version of the GNU Lesser General Public License ever published by the Free Software Foundation.

If the Library as you received it specifies that a proxy can decide whether future versions of the GNU Lesser General Public License shall apply, that proxy’s public statement of acceptance of any version is permanent authorization for you to choose that version for the Library.