



# Functional Description version 1.9.1

Web Server, LLC



## Contents

| 1   | Ann  | otation  |                                  | 1   |
|-----|------|----------|----------------------------------|-----|
| 2   | Gene | eral Inf | formation                        | 2   |
| 3   | Func | ctional  | Characteristics                  | 4   |
|     | 3.1  | Run-Ti   | ime Control                      | 4   |
|     |      | 3.1.1    | Using Signals                    | 5   |
|     |      | 3.1.2    | Changing Configuration           | 5   |
|     |      | 3.1.3    | Rotating Log Files               | 6   |
|     | 3.2  | Connec   | ctions, Sessions, Requests, Logs | 7   |
|     |      | 3.2.1    | Connection processing mechanisms | 7   |
|     |      | 3.2.2    | HTTP request processing          | 8   |
|     |      | 3.2.3    | TCP/UDP session processing       | 8   |
|     |      | 3.2.4    | Processing requests              | 8   |
|     |      | 3.2.5    | Proxying and Load Balancing      | 12  |
|     |      | 3.2.6    | Logging                          | 15  |
|     | 3.3  | Built-ii | n Modules                        | 15  |
|     |      | 3.3.1    | Core Module                      | 15  |
|     |      | 3.3.2    | HTTP Module                      | 25  |
|     |      | 3.3.3    | Stream Module                    | 317 |
|     |      | 3.3.4    | Mail Module                      | 384 |
|     |      | 3.3.5    | Google PerfTools                 | 404 |
|     |      | 3.3.6    | WASM Module                      | 404 |
|     |      | 3.3.7    | Core Module                      | 408 |
|     |      | 3.3.8    | HTTP Modules                     | 408 |
|     |      | 3.3.9    | Streaming Modules                | 409 |
|     |      | 3.3.10   | Mail Modules                     | 410 |
|     |      | 3.3.11   | Google PerfTools Module          | 410 |
|     |      | 3.3.12   | WASM Modules                     | 410 |
| 4   | Inte | llectual | Property Rights                  | 411 |
| Inc | dex  |          |                                  | 412 |



|         | -1 |
|---------|----|
| CHAPTER |    |

Annotation

This document contains a description of the functional features of Angie PRO. Angie PRO is an efficient, powerful, and scalable web server.



## CHAPTER 2

## General Information

Angie PRO is the only commercial web server developed and localized in Russia.

A web server is a class of software that provides access to network resources via the HTTP protocol to end users. Angie PRO, for example, can be used to operate websites, mobile applications, self-service kiosks in the subway, and multimedia systems on long-distance trains. Every time a user opens a website, uses a mobile application, interacts with a self-service kiosk in the subway, or even with a multimedia system on the "Sapsan" train, the user's request can be processed by Angie PRO.

#### Angie PRO is:

- A general-purpose web server. Written in C.
- An L4-L7 load balancer. Allows load balancing between servers for both TCP/UDP protocols and HTTP.
- A proxy and caching server. Enables faster operation of web services through a flexible caching mechanism.
- Available on all popular platforms. Compiled and tested on Alpine, Debian, Oracle, RED OS, Rocky, and Ubuntu.
- High performance. One of the most efficient web servers in the world.

#### Why choose Angie PRO:

- Compatibility with NGINX OSS. Angie PRO is fully compatible with Nginx, allowing any existing Nginx user to transition to Angie PRO without significant costs or service downtime.
- Enhanced statistics and real-time monitoring. Angie PRO offers complete real-time server load monitoring, enabling dynamic configuration management based on load profiles and ensuring full service availability.
- Dynamic configuration of proxied server groups. Allows management of proxied server group settings through a convenient REST interface without service interruption.
- Cache element removal. Provides the ability to remove cache elements via a user-friendly API without service downtime.
- Active health checks for proxied servers. Checks for "liveness" and proxies only to those groups of proxied servers that respond according to a specified algorithm.
- Dynamic key-value storage. Enables dynamic management of Angie PRO configuration variables via HTTP API.



- Dynamic DNS updates.
- Session-affinity proxying.
- Repository with dynamic third-party modules. Angie PRO supports most NGINX third-party modules and allows for seamless installation, guaranteeing functionality and support.
- Shared memory zone synchronization. Capability to use cache zones, limit\_req, etc., in the Angie PRO cluster.
- Hiding or personal branding of the server name in response headers. Ability to change or hide the name and version of the web server from users.

A list of foreign software with similar functional characteristics to Angie PRO includes Nginx, Nginx Plus, Apache, Envoy, products utilizing NGINX solutions (OpenResty, Tengine, Cloudflare), and Yandex's cloud solutions.



## CHAPTER 3

**Functional Characteristics** 

### 3.1 Run-Time Control

To start Angie, use systemd with the following command:

\$ sudo service angie start

It is recommended to check the configuration syntax beforehand. Here is how:

\$ sudo angie -t && sudo service angie start

To reload the configuration:

\$ sudo angie -t && sudo service angie reload

To stop Angie:

\$ sudo service angie stop

After installation, run the following command to ensure that Angie is up and running:

\$ curl localhost:80

#### 1 Note

The methods for running the open-source version of Angie may vary depending on the installation method.

Angie has one master process and several worker processes. The master process is responsible for reading and evaluating the configuration and maintaining the worker processes. Worker processes handle the actual request processing. Angie uses an event-based model and OS-dependent mechanisms to efficiently distribute requests among the worker processes. The number of worker processes is defined in the configuration file and may be either fixed for a given configuration or automatically adjusted based on the number of available CPU cores (see worker processes).

When configured, Angie will also flush certain shared memory zones (currently, the keys\_zone in proxy\_cache\_path) to the disk before exiting, so a newly started master process can restore them with



improved performance. If the restore fails due to a change in zone size, binary version incompatibility, or other reasons, Angie will log an alert (failed to restore zone at address) and will not use the zone restore mechanism.

### 3.1.1 Using Signals

Angie can also be controlled using signals. By default, the process ID of the master process is written to the file /run/angie.pid. This filename can be changed at configuration time or in angie.conf using the *pid* directive. The master process supports the following signals:

| TERM, INT | Fast shutdown  |
|-----------|--|
| QUIT      | Graceful shutdown  |
| HUP       | Reload configuration, update time zone (only for FreeBSD and Linux), start new worker processes with the updated configuration, <i>gracefully</i> shut down old worker processes |
| USR1      | Reopen log files   |
| USR2      | Upgrade the executable file  |
| WINCH     | Graceful shutdown of worker processes  |

You can send signals using kill:

```
$ sudo kill -QUIT $(cat /run/angie.pid)
```

Individual worker processes can also be controlled using signals, although this is optional. The supported signals are:

| TERM, INT | Fast shutdown  |
|-----------|--|
| QUIT      | Graceful shutdown  |
| USR1      | Reopen log files   |
| WINCH     | Abnormal termination for debugging (requires debug_points to be enabled) |

#### 3.1.2 Changing Configuration

In order for Angie to re-read the configuration file, a HUP signal should be sent to the master process. The master process first checks the syntax validity and then attempts to apply the new configuration, which includes opening new log files and listen sockets. If applying the new configuration fails, the master process rolls back the changes and continues operating with the old configuration. If the application succeeds, the master process starts new worker processes and sends messages to the old worker processes, requesting them to shut down *gracefully*. The old worker processes close their listen sockets and continue to service existing clients. After all clients have been served, the old worker processes are shut down.

Angie tracks configuration changes for each process. Generation numbers start at 1 when the server is first started. These numbers are incremented with each configuration reload and are visible in the process titles:

```
$ sudo angie
$ ps aux | grep angie
angie: master process v1.9.1 #1 [angie]
angie: worker process #1
```

After a successful configuration reload (regardless of whether there are actual changes), Angie increments the generation number for processes that received the new configuration:

```
$ sudo kill -HUP $(cat /run/angie.pid)
$ ps aux | grep angie
```



```
angie: master process v1.9.1 #2 [angie]
angie: worker process #2
```

If any worker processes from previous generations continue to operate, they will become immediately visible:

```
$ ps aux | grep angie
angie: worker process #1
angie: worker process #2
```

#### 1 Note

Do not confuse the configuration generation number with a 'process number'; Angie does not use continuous process numbering for practical purposes.

#### 3.1.3 Rotating Log Files

To rotate log files, first rename the files. Then, send a USR1 signal to the master process. The master process will re-open all currently open log files and assign them to an unprivileged user under which the worker processes are running. After successfully re-opening the files, the master process closes all open files and notifies the worker processes to re-open their log files. Worker processes will also open the new files and close the old ones immediately. As a result, the old files become available for post-processing, such as compression, almost immediately.

#### On-the-fly Executable Upgrade

To upgrade the server executable, first replace the old executable file with the new one. Then, send a USR2 signal to the master process. The master process will rename its current file with the process ID to a new file with the .oldbin suffix, e.g., /usr/local/angie/logs/angie.pid.oldbin, and then start the new executable, which in turn starts new worker processes.

Note that the old master process does not close its listen sockets and can be managed to restart its worker processes if necessary. If the new executable does not perform as expected, you can take one of the following actions:

- Send the HUP signal to the old master process. This will start new worker processes without rereading the configuration. You can then shut down all new processes *gracefully* by sending the QUIT signal to the new master process.
- Send the TERM signal to the new master process. It will send a message to its worker processes, requesting them to exit immediately. If any processes do not exit, send the KILL signal to force them to exit. When the new master process exits, the old master process will automatically start new worker processes.

If the new master process exits, the old master process will remove the <code>.oldbin</code> suffix from the file name with the process ID.

If the upgrade is successful, send the QUIT signal to the old master process, and only the new processes will remain.

When configured, Angie will also flush certain shared memory zones (currently, the keys\_zone in proxy\_cache\_path) to the disk before upgrading, so a newly started master process can restore them with improved performance. If the restore fails due to a change in zone size, binary version incompatibility, or other reasons, Angie will log an alert (failed to restore zone at address) and will not use the zone restore mechanism.



#### **Command-Line Options**

| -h, -?           | Display help for command-line parameters, then exit.                                    |
|------------------|---|
| ,                | 1 / 1   |
| build-env        | Display auxiliary information about the build environment, then exit.                   |
| -c file          | Use file as the configuration file instead of the default file.                         |
| -е file          | Use file as the error log file instead of the default file. The special value stderr    |
|                  | specifies the standard error output.  |
| -g directives    | Apply additional global configuration directives, for example: angle -g "pid /          |
|                  | <pre>var/run/angie.pid; worker_processes `sysctl -n hw.ncpu`;".</pre>                   |
| -m, -M           | Display a list of built-in (-m) or built-in and loaded (-M) modules, then exit.         |
| -р <i>prefix</i> | Specify the <i>prefix</i> path for angle (the directory where server files are located; |
|                  | the default is /usr/local/angie/).  |
| -q               | Only display error messages if -t or -T is set; otherwise, do nothing.                  |
| -s $signal$      | Send a signal to the master process, such as stop, quit, reopen, reload, and so         |
|                  | on.   |
| -t               | Test the configuration file, then exit. Angie checks the configuration syntax and       |
|                  | recursively includes any files mentioned in it.   |
| -T               | Same as -t, but also outputs the summary configuration to standard output after         |
|                  | recursively including all files mentioned in the configuration.                         |
| -V               | Display the Angie version, then exit.   |
| -                | - · ·   |
| -V               | Display the Angie version, the compiler version, build time and the build param-        |
|                  | eters used, then exit.  |

## 3.2 Connections, Sessions, Requests, Logs

## 3.2.1 Connection processing mechanisms

Angie supports various connection processing methods. The availability of a specific method depends on the platform being used. On platforms that support multiple methods, Angie typically selects the most efficient method automatically. However, if necessary, a connection processing method can be explicitly chosen using the *use* directive.

The following connection processing methods are available:

| Method    | Description   |
|-----------|---|
| select    | A standard method. The supporting module is built automatically on platforms that do not have more efficient methods. Thewith-select_module andwithout-select_module build options can be used to forcibly enable or disable the building of this module. |
| poll      | A standard method. The supporting module is built automatically on platforms that do not have more efficient methods. Thewith-poll_module andwithout-poll_module build options can be used to forcibly enable or disable the building of this module.     |
| kqueue    | An efficient method available on FreeBSD 4.1+, OpenBSD 2.9+, NetBSD 2.0, and macOS.   |
| epoll     | An efficient method available on Linux 2.6+.  |
| /dev/poll | An efficient method available on Solaris 7 $11/99+$ , HP/UX $11.22+$ (eventport), IRIX $6.5.15+$ , and Tru64 UNIX $5.1A+$ .   |
| eventport | The event ports method is available on Solaris 10+. (Due to known issues, using the /dev/poll method is recommended instead.)   |



## 3.2.2 HTTP request processing

An HTTP request goes through a series of phases, where a specific type of processing is performed at each phase.

| _              |   |
|----------------|---|
| Post-read      | The initial phase. The <i>RealIP</i> module is invoked during this phase.                       |
| Server-rewrite | The phase where directives from the <i>Rewrite</i> module, defined in a server block            |
|                | (but outside a location block), are processed.  |
| Find-config    | A special phase where a <i>location</i> is selected based on the request URI.                   |
| Rewrite        | Similar to the Server-rewrite phase, but it applies to rewrite rules defined                    |
|                | within the location block selected in the previous phase.                                       |
| Post-rewrite   | A special phase where the request is redirected to a new location, as in the                    |
|                | Find-config phase, if its URI was modified during the Rewrite phase.                            |
| Preaccess      | During this phase, standard Angie modules like <i>Limit Req</i> register their handlers.        |
| Access         | The phase where the client's authorization to make the request is verified, typi-               |
|                | cally by invoking standard Angie modules such as Auth Basic.                                    |
| Post-access    | A special phase where the <i>satisfy any</i> directive is processed.                            |
| Precontent     | Standard module directives, such as try_files and mirror, register their handlers               |
|                | during this phase.  |
| Content        | The phase where the response is usually generated. Multiple standard Angie                      |
|                | modules register their handlers at this stage, including <i>Index</i> . The <i>proxy_pass</i> , |
|                | fastcgi_pass, uwsgi_pass, scgi_pass and grpc_pass directives are also handled                   |
|                | here.   |
|                | Handlers are called sequentially until one of them produces the output.                         |
| Log            | The final phase, where request logging is performed. Currently, only the Log                    |
|                | module registers its handler at this stage for access logging.                                  |
|                |   |

### 3.2.3 TCP/UDP session processing

A TCP/UDP session from a client goes through a series of phases, where a specific type of processing is performed at each phase:

| Post-accept | The initial phase after accepting a client connection. The <i>RealIP</i> module is invoked at this phase.  |
|-------------|--|
| Pre-access  | A preliminary phase for checking access. The <i>Set</i> modules are invoked during this phase.   |
| Access      | The phase for limiting client access before actual data processing. The <i>Access</i> module is invoked at this stage.   |
| SSL         | The phase where TLS/SSL termination occurs. The <i>SSL</i> module is invoked during this phase.  |
| Preread     | The phase for reading initial bytes of data into the <i>preread buffer</i> to allow modules such as <i>SSL Preread</i> to analyze the data before processing.                                |
| Content     | A mandatory phase where the data is actually processed, typically involving the <i>Return</i> module to send a response to the client. The <i>proxy_pass</i> directive is also handled here. |
| Log         | The final phase where the outcome of client session processing is recorded. The $Log$ module is invoked at this phase.   |

## 3.2.4 Processing requests

#### Virtual server selection

Initially, a connection is created within the context of a default server. The server name can then be determined in the following stages of request processing, each of which is involved in the selection of server configuration:

• During the SSL handshake, in advance, according to the SNI.



- After processing the request line.
- After processing the Host header field.

If the server name is not determined after processing the request line or the Host header field, Angie will use an empty name as the server name.

At each of these stages, different server configurations may be applied. Therefore, certain directives should be specified with caution:

- In the case of the *ssl\_protocols* directive, the protocol list is set by the OpenSSL library before the server configuration is applied according to the name requested through SNI. As a result, protocols should only be specified for the default server.
- The *client\_header\_buffer\_size* and *merge\_slashes* directives are applied before reading the request line. Therefore, these directives use either the default server configuration or the server configuration chosen by SNI.
- In the case of the <code>ignore\_invalid\_headers</code>, <code>large\_client\_header\_buffers</code>, and <code>underscores\_in\_headers</code> directives, which are involved in processing request header fields, the server configuration additionally depends on whether it was updated according to the request line or the <code>Host</code> header field.
- An error response is handled using the *error\_page* directive in the server that is currently processing the request.

#### Name-based virtual servers

Angie first determines which server should handle the request. Consider a simple configuration where all three virtual servers listen on port 80:

```
server {
    listen 80;
    server_name example.org www.example.org;
    # ...
}
server {
    listen 80;
    server_name example.net www.example.net;
    # ...
}
server {
    listen 80;
    server_name example.com www.example.com;
    # ...
}
```

In this configuration, Angie determines which server should handle the request based solely on the Host header field. If the value of this header does not match any server name or if the request does not contain this header field, Angie will route the request to the default server for this port. In the configuration above, the default server is the first one — which is Angie's standard default behavior. It can also be explicitly specified which server should be the default using the default\_server parameter in the listen directive:

```
server {
   listen 80 default_server;
```



```
server_name example.net www.example.net;
# ...
}
```

#### 1 Note

Note that the default server is a property of the listen socket, not of the server name.

#### Internationalized names

Internationalized domain names (IDNs) should be specified using an ASCII (Punycode) representation in the server name directive:

```
server {
    listen 80;
    server_name xn--elafmkfd.xn--80akhbyknj4f; # пример.испытание
    # ...
}
```

#### Preventing requests with undefined server names

If requests without the Host header field should not be allowed, a server that simply drops such requests can be defined:

```
server {
    listen 80;
    server_name "";
    return 444;
}
```

In this configuration, the server name is set to an empty string, which matches requests without the Host header field. A special non-standard code 444 is then returned, which closes the connection.

## Combining name-based and IP-based virtual servers

Let's examine a more complex configuration where some virtual servers listen on different addresses:

```
server {
    listen 192.168.1.1:80;
    server_name example.org www.example.org;
    # ...
}
server {
    listen 192.168.1.1:80;
    server_name example.net www.example.net;
    # ...
}
server {
    listen 192.168.1.2:80;
    server_name example.com www.example.com;
```



```
# ...
}
```

In this configuration, Angie first tests the IP address and port of the request against the *listen* directives of the *server* blocks. It then tests the Host header field of the request against the *server\_name* entries of the *server* blocks that matched the IP address and port. If the server name is not found, the request will be processed by the default server. For example, a request for www.example.com received on port 192.168.1.1:80 will be handled by the default server for that port — i.e., by the first server — since www.example.com is not defined for this port.

As previously mentioned, a default server is a property of the listen port, and different default servers may be defined for different ports:

```
server {
    listen 192.168.1.1:80;
    server_name example.org www.example.org;
    # ...
}

server {
    listen 192.168.1.1:80 default_server;
    server_name example.net www.example.net;
    # ...
}

server {
    listen 192.168.1.2:80 default_server;
    server_name example.com www.example.com;
    # ...
}
```

### **Choosing locations**

Consider a simple PHP website configuration:

```
server {
    listen 80;
    server_name example.org www.example.org;
    root /data/www;

    location / {
        index index.html index.php;
    }

    location ~* \.(gif|jpg|png)$ {
        expires 30d;
    }

    location ~ \.php$ {
        fastcgi_pass localhost:9000;
        fastcgi_param SCRIPT_FILENAME
```



```
$document_root$fastcgi_script_name;
include fastcgi_params;
}
```

Angie first searches for the most specific prefix location given by literal strings, regardless of their listed order. In the configuration above, the only prefix location is location /, which matches any request and will be used as a last resort. Angie then checks locations defined by regular expressions in the order they appear in the configuration file. The first matching expression stops the search, and Angie will use that location. If no regular expression matches a request, Angie will use the most specific prefix location found earlier.

#### 1 Note

Locations of all types test only the URI part of the request line, excluding arguments. This is because arguments in the query string can be specified in various ways, for example:

- /index.php?user=john&page=1
- /index.php?page=1&user=john

Additionally, query strings may contain any number of parameters:

• /index.php?page=1&something+else&user=john

Now let's look at how requests would be processed in the configuration above:

- The request /logo.gif is first matched by the prefix location / and then by the regular expression . (gif|jpg|png)\$. Therefore, it is handled by the latter location. Using the directive root /data/www, the request is mapped to the file /data/www/logo.gif, and the file is sent to the client.
- The request /index.php is also initially matched by the prefix location / and then by the regular expression .(php)\$. Consequently, it is handled by the latter location, and the request is passed to a FastCGI server listening on localhost:9000. The <code>fastcgi\_param</code> directive sets the FastCGI parameter SCRIPT\_FILENAME to /data/www/index.php, and the FastCGI server executes the file. The variable <code>\$fastcgi script name</code> is set to the request URI, i.e., /index.php.
- The request /about.html is matched only by the prefix location /, so it is handled in this location. Using the directive root /data/www, the request is mapped to the file /data/www/about.html, and the file is sent to the client.

Handling the request / is more complex. It is matched only by the prefix location /, so it is handled by this location. The index directive then tests for the existence of index files according to its parameters and the root /data/www directive. If the file /data/www/index.html does not exist but the file /data/www/index.php does, the directive performs an internal redirect to /index.php, and Angie searches the locations again as if the request had been sent by a client. As previously mentioned, the redirected request will eventually be handled by the FastCGI server.

#### 3.2.5 Proxying and Load Balancing

One common use of Angie is to set it up as a proxy server. In this role, Angie receives requests, forwards them to the proxied servers, retrieves responses from those servers, and sends the responses back to the clients.

A simple proxy server:

```
server {
  location / {
```



```
proxy_pass http://backend:8080;
}
```

The *proxy\_pass* directive instructs Angie to pass client requests to the backend backend:8080 (the proxied server). There are many additional *directives* available for further configuring a proxy connection.

#### FastCGI Proxying

Angie can be used to route requests to FastCGI servers that run applications built with various frameworks and programming languages, such as PHP.

The most basic Angie configuration for working with a FastCGI server involves using the <code>fastcgi\_pass</code> directive instead of the <code>proxy\_pass</code> directive, along with <code>fastcgi\_param</code> directives to set parameters passed to the FastCGI server. Suppose the FastCGI server is accessible on <code>localhost:9000</code>. In PHP, the <code>SCRIPT\_FILENAME</code> parameter is used to determine the script name, and the <code>QUERY\_STRING</code> parameter is used to pass request parameters. The resulting configuration would be:

```
server {
    location / {
        fastcgi_pass localhost:9000;
        fastcgi_param SCRIPT_FILENAME $document_root$fastcgi_script_name;
        fastcgi_param QUERY_STRING $query_string;
}

location ~ \.(gif|jpg|png)$ {
        root /data/images;
    }
}
```

This configuration sets up a server that routes all requests, except those for static images, to the proxied server operating on localhost:9000 via the FastCGI protocol.

#### WebSocket Proxying

To upgrade a connection from  $\mathrm{HTTP}/1.1$  to WebSocket, the protocol switch mechanism available in  $\mathrm{HTTP}/1.1$  is used.

However, there is a subtlety: since the Upgrade header is a hop-by-hop header, it is not passed from the client to the proxied server. With forward proxying, clients may use the CONNECT method to circumvent this issue. This approach does not work with reverse proxying, as clients are unaware of any proxy servers, and special processing on the proxy server is required.

Angie implements a special mode of operation that allows setting up a tunnel between a client and a proxied server if the proxied server returns a response with code 101 (Switching Protocols), and the client requests a protocol switch via the Upgrade header in the request.

As mentioned, hop-by-hop headers, including Upgrade and Connection, are not passed from the client to the proxied server. Therefore, for the proxied server to be aware of the client's intention to switch to the WebSocket protocol, these headers must be explicitly passed:

```
location /chat/ {
    proxy_pass http://backend;
    proxy_http_version 1.1;
    proxy_set_header Upgrade $http_upgrade;
    proxy_set_header Connection "upgrade";
}
```



A more sophisticated example demonstrates how the value of the Connection header field in a request to the proxied server depends on the presence of the Upgrade field in the client request header:

```
http {
    map $http_upgrade $connection_upgrade {
        default upgrade;
        '' close;
}

server {
        ...
        location /chat/ {
            proxy_pass http://backend;
            proxy_http_version 1.1;
            proxy_set_header Upgrade $http_upgrade;
            proxy_set_header Connection $connection_upgrade;
        }
    }
}
```

By default, the connection will be closed if the proxied server does not transmit any data within 60 seconds. This timeout can be increased using the  $proxy\_read\_timeout$  directive. Alternatively, the proxied server can be configured to periodically send WebSocket ping frames to reset the timeout and check if the connection is still active.

#### **Load Balancing**

Load balancing across multiple application instances is a widely used technique to optimize resource utilization, maximize throughput, reduce latency, and ensure fault-tolerant configurations.

Angie can be used as a highly efficient HTTP load balancer to distribute traffic to multiple application servers, thereby enhancing the performance, scalability, and reliability of web applications.

The simplest configuration for load balancing with Angie might look like this:

```
http {
    upstream myapp1 {
        server srv1.example.com;
        server srv2.example.com;
        server srv3.example.com;
}

server {
        listen 80;
        location / {
            proxy_pass http://myapp1;
        }
    }
}
```



In the example above, three instances of the same application are running on srv1 through srv3. When a load balancing method is not explicitly configured, it defaults to round-robin. Other supported load balancing mechanisms include: weight, least\_conn, and ip\_hash. The reverse proxy implementation in Angie also supports in-band (or passive) server health checks. These are configured using the max fails and fail timeout directives within the server block in the upstream context.

#### 3.2.6 Logging



#### Note

In addition to the options listed here, you can also enable the debugging log.

#### Syslog

The error\_log and access\_log directives support logging to syslog. The following parameters are used to configure logging to syslog:

| server=address  | Specifies the address of a syslog server. The address can be a domain name or an IP address, with an optional port, or a UNIX domain socket path specified after the "unix:" prefix. If the port is not specified, UDP port 514 is used. If a domain name resolves to multiple IP addresses, the first resolved address is used.  |
|-----------------|---|
| facility=string | Sets the facility for syslog messages, as defined in RFC 3164. Possible facilities include: "kern", "user", "mail", "daemon", "auth", "intern", "lpr", "news", "uucp", "clock", "authpriv", "ftp", "ntp", "audit", "alert", "cron", "local0""local7". The default is "local7".  |
| severity=string | Defines the severity level of syslog messages for access_log, as specified in RFC 3164. Possible values are the same as those for the second parameter (level) of the error_log directive. The default is "info". The severity of error messages is determined by Angie, so this parameter is ignored in the error_log directive. |
| tag=string      | Sets the tag for syslog messages. The default tag is "angie".   |
| nohostname      | Disables the addition of the :sa,p:hostname field in the syslog message header.   |

#### Example syslog configuration:

```
error_log syslog:server=192.168.1.1 debug;
access_log syslog:server=unix:/var/log/angie.sock,nohostname;
access_log syslog:server=[2001:db8::1]:12345,facility=local7,tag=angie,severity=infou
```



#### 1 Note

Syslog entries are reported no more than once per second to prevent flooding.

#### 3.3 Built-in Modules

This guide describes the built-in modules of Angie, provides configuration examples, lists their directives and parameters, as well as built-in variables.

#### 3.3.1 Core Module

The module provides essential functionality and configuration directives necessary for the basic operation of the server, and handles critical tasks such as managing worker processes, configuring event-driven models, and processing incoming connections and requests. It includes key directives for setting up the main process, error logging, and controlling the behavior of the server at a low level.



#### **Configuration Example**

```
user www www;
worker_processes 2;
error_log /var/log/error.log info;
events {
    use kqueue; worker_connections 2048;
```

#### **Directives**

#### accept mutex

| Syntax  | accept_mutex on   off;       |
|---------|------------------------------|
| Default | <pre>accept_mutex off;</pre> |
| Context | events                       |

When accept\_mutex is enabled, worker processes will accept new connections in turn. Without this setting, all worker processes are notified of new connections, which can lead to inefficient use of system resources if the volume of new connections is low.

#### 1 Note

There is no need to enable accept\_mutex on systems that support the EPOLLEXCLUSIVE flag or when using the reuseport directive.

#### accept \_ mutex \_ delay

| Syntax  | accept_mutex_delay $time;$           |
|---------|--------------------------------------|
| Default | <pre>accept_mutex_delay 500ms;</pre> |
| Context | events                               |

If accept mutex is enabled, this directive specifies the maximum time a worker process will wait to continue accepting new connections while another worker process is already handling new connections.

#### daemon

| Syntax  | daemon on   off; |
|---------|------------------|
| Default | daemon on;       |
| Context | main             |

Determines whether Angie should run as a daemon. This is primarily used during development.

#### debug\_connection

| Syntax  | debug_connection address   CIDR   unix:; |
|---------|--|
| Default | _  |
| Context | events                                   |



Enables debugging logs for specific client connections. Other connections will use the logging level set by the *error\_log* directive. You can specify connections by IPv4 or IPv6 address, network, or hostname. For connections using UNIX domain sockets, use the unix: parameter to enable debugging logs.

```
events {
    debug_connection 127.0.0.1;
    debug_connection localhost;
    debug_connection 192.0.2.0/24;
    debug_connection ::1;
    debug_connection 2001:0db8::/32;
    debug_connection unix:;
    # ...
}
```

#### Important

For this directive to work, Angie must be built with debugging log enabled.

#### debug points

| Syntax  | debug_points abort   stop; |
|---------|----------------------------|
| Default | _                          |
| Context | main                       |

This directive is used for debugging.

When an internal error occurs, such as a socket leak during worker process restarts, enabling debug\_points will either create a core file (abort) or stop the process (stop) for further analysis with a system debugger.

#### env

| Syntax  | env variable[=value]; |
|---------|-----------------------|
| Default | env TZ;               |
| Context | main                  |

By default, Angie removes all environment variables inherited from its parent process except for the TZ variable. This directive allows you to preserve some inherited variables, modify their values, or create new environment variables.

These variables are then:

- ullet Inherited during a live upgrade of an executable file
- Used by the *Perl* module
- Available to worker processes

Note that controlling system libraries in this way may not always be effective, as libraries often check variables only during initialization, which occurs before this directive takes effect. The TZ variable is always inherited and accessible to the *Perl* module unless explicitly configured otherwise.

Example:



```
env MALLOC_OPTIONS;
env PERL5LIB=/data/site/modules;
env OPENSSL_ALLOW_PROXY_CERTS=1;
```

## 1 Note

The ANGIE environment variable is used internally by Angie and should not be set directly by the user.

## error\_log

| Syntax  | error_log file [level];  |
|---------|--|
| Default | error_log logs/error.log error; (the path depends on theerror-log-path build option) |
| Context | main, http, mail, stream, server, location   |

Configures logging, allowing multiple logs to be specified at the same configuration level. If a log file is not explicitly defined at the main configuration level, the default file will be used.

The first parameter specifies the file to store the log. The special value **stderr** selects the standard error stream. To configure logging to *syslog*, use the "syslog:" prefix. To log to a cyclic memory buffer, use the "memory:" prefix followed by the buffer size; this is typically used for debugging.

The second parameter sets the logging level, which can be one of the following: debug, info, notice, warn, error, crit, alert, or emerg. These levels are listed in order of increasing severity. Setting a log level will capture messages of equal and higher severity:

| Setting | Levels Captured                                      |
|---------|--|
| debug   | debug, info, notice, warn, error, crit, alert, emerg |
| info    | info, notice, warn, error, crit, alert, emerg        |
| notice  | notice, warn, error, crit, alert, emerg              |
| warn    | warn, error, crit, alert, emerg                      |
| error   | error, crit, alert, emerg                            |
| crit    | crit, alert, emerg                                   |
| alert   | alert, emerg   |
| emerg   | emerg  |

If this parameter is omitted, error is used as the default logging level.

## Important

For the debug logging level to work, Angie must be built with debugging log enabled.

#### events

| Syntax  | events { }; |
|---------|-------------|
| Default | _           |
| Context | main        |

Provides the configuration file context for directives that affect connection processing.



#### include

| Syntax  | $	ext{include } file \mid mask;$ |
|---------|----------------------------------|
| Default | _                                |
| Context | any                              |

Includes another file, or files that match the specified mask, into the configuration. The included files must contain syntactically correct directives and blocks.

#### Example:

```
include mime.types;
include vhosts/*.conf;
```

### load module

| Syntax  | ${\tt load\_module}\; file;$ |
|---------|------------------------------|
| Default | _                            |
| Context | main                         |

Loads a dynamic module from the specified file. If a relative path is provided, it is interpreted based on the --prefix build option. To verify the path:

```
$ sudo angie -V
```

#### Example:

```
load_module modules/ngx_mail_module.so;
```

#### lock\_file

| Syntax  | ${	t lock\_file};$  |
|---------|---|
| Default | <pre>lock_file logs/angie.lock; (the path depends on thelock-path build option)</pre> |
| Context | main  |

Angie uses a locking mechanism to implement  $accept\_mutex$  and serialize access to shared memory. On most systems, locks are managed using atomic operations, making this directive unnecessary. On certain systems, however, an alternative lock file mechanism is used. This directive sets a prefix for lock file names.

#### master process

| Syntax  | master_process on   off; |
|---------|--------------------------|
| Default | master_process on;       |
| Context | main                     |

Determines whether worker processes are started. This directive is intended for Angie developers.



#### multi\_accept

| Syntax  | multi_accept on   off;       |
|---------|------------------------------|
| Default | <pre>multi_accept off;</pre> |
| Context | events                       |

| on  | A worker process will accept all new connections simultaneously. |
|-----|--|
| off | A worker process will accept one new connection at a time.       |

#### 1 Note

This directive is ignored if the kqueue connection processing method is used, as it provides the number of new connections ready to be accepted.

#### pcre\_jit

| Syntax  | <pre>pcre_jit on   off;</pre> |
|---------|-------------------------------|
| Default | <pre>pcre_jit off;</pre>      |
| Context | main                          |

Enables or disables "just-in-time compilation" (PCRE JIT) for regular expressions known at the time of configuration parsing.

PCRE JIT can significantly accelerate regular expression processing.

#### Important

JIT is available in PCRE libraries from version 8.20, provided they are built with the --enable-jit configuration option. When Angie is built with the PCRE library (--with-pcre=), JIT support is enabled using the --with-pcre-jit option.

#### pid

| Syntax  | $\verb pid  file   \verb off ;$   |
|---------|---|
| Default | <pre>pid logs/angie.pid; (the path depends on thepid-path build option)</pre> |
| Context | main  |

Specifies the *file* that will store the ID of the Angie main process. The file is created atomically, which ensures its contents are always correct. The off setting disables the creation of this file.

#### 1 Note

If the *file* setting is modified during reconfiguration but points to a symlink of the previous PID file, the file will not be recreated.



#### ssl\_engine

| Syntax  | ssl_engine device; |
|---------|--------------------|
| Default | _                  |
| Context | main               |

Specifies the name of the hardware SSL accelerator.

#### ssl object cache inheritable

| Syntax  | ssl_object_cache_inheritable on   off; |
|---------|--|
| Default | ssl_object_cache_inheritable on;       |
| Context | main                                   |

If enabled, SSL objects (SSL certificates, secret keys, trusted CA certificates, CRL lists) are inherited across configuration reloads.

SSL objects loaded from files are inherited if their modification time and file index have not changed since the previous configuration load. Secret keys specified as engine:name:id are never inherited, while secret keys specified as data:value are always inherited.

SSL objects loaded from variables cannot be inherited.

Example:

```
ssl_object_cache_inheritable on;
http {
    server {
        ssl_certificate example.com.crt;
        ssl_certificate_key example.com.key;
    }
}
```

#### thread pool

| Syntax  | thread_pool name threads=number [max_queue=number];        |
|---------|--|
| Default | <pre>thread_pool default threads=32 max_queue=65536;</pre> |
| Context | main   |

Defines the name and parameters of a thread pool used for multi-threaded file reading and sending,  $without\ blocking$  worker processes.

The threads parameter specifies the number of threads in the pool.

If all threads are busy, new tasks will wait in the queue. The max\_queue parameter limits the number of tasks that can wait in the queue. By default, the queue can hold up to 65536 tasks. When the queue overflows, new tasks are completed with an error.

#### timer\_resolution

| Syntax  | timer_resolution interval; |
|---------|----------------------------|
| Default | _                          |
| Context | main                       |



Reduces timer resolution in worker processes, thereby decreasing the frequency of gettimeofday() system calls. By default, gettimeofday() is called each time kernel events are queried. With reduced resolution, it is called only once per the specified interval.

#### Example:

```
timer_resolution 100ms;
```

The internal implementation of the interval depends on the method used:

- The EVFILT\_TIMER filter if kqueue is used.
- The timer\_create() function if eventport is used.
- The setitimer() function otherwise.

#### use

| Syntax  | use $method;$ |
|---------|---------------|
| Default | _             |
| Context | events        |

Specifies the *connection processing method* to use. Typically, there is no need to specify it explicitly, as Angie defaults to the most efficient method.

#### user

| Syntax  | user user [group];  |
|---------|---|
| Default | user <user build="" option=""> <group build="" option="">;</group></user> |
| Context | main  |

Defines the user and group credentials for worker processes (see also the build options). If only the user is set, the specified user name is also used for the group.

#### worker aio requests

| Syntax  | worker_aio_requests number; |
|---------|-----------------------------|
| Default | worker_aio_requests 32;     |
| Context | events                      |

When using aio with the epoll connection processing method, sets the maximum number of outstanding asynchronous I/O operations for a single worker process.

## worker\_connections

| Syntax  | worker_connections number; |
|---------|----------------------------|
| Default | worker_connections 512;    |
| Context | events                     |

Sets the maximum number of simultaneous connections a worker process can open.

Note that this number includes all connections, such as those with proxied servers, not just client connections. Additionally, the actual number of simultaneous connections cannot exceed the system's limit on open files, which can be adjusted using *worker rlimit nofile*.



#### worker\_cpu\_affinity

| Syntax  | worker_cpu_affinity $cpumask$ ; worker_cpu_affinity auto [ $cpumask$ ]; |
|---------|---|
| Default | _   |
| Context | main  |

Binds worker processes to specific sets of CPUs. Each CPU set is represented by a bitmask indicating allowed CPUs. A separate set should be defined for each worker process. By default, worker processes are not bound to any specific CPUs.

For example:

```
worker_processes 4;
worker_cpu_affinity 0001 0010 0100 1000;
```

This configuration binds each worker process to a separate CPU.

Alternatively:

```
worker_processes 2;
worker_cpu_affinity 0101 1010;
```

This binds the first worker process to CPU0 and CPU2, and the second worker process to CPU1 and CPU3. This setup is suitable for hyper-threading.

The special value auto allows automatic binding of worker processes to available CPUs:

```
worker_processes auto;
worker_cpu_affinity auto;
```

The optional mask parameter can be used to limit the CPUs available for automatic binding:

```
worker_cpu_affinity auto 01010101;
```

## Important

This directive is only available on FreeBSD and Linux.

#### worker\_priority

| Syntax  | worker_priority number;       |
|---------|-------------------------------|
| Default | <pre>worker_priority 0;</pre> |
| Context | main                          |

Defines the scheduling priority for worker processes, similar to the **nice** command: a negative *number* indicates higher priority. The allowed range typically varies from -20 to 20.

Example:

```
worker_priority -10;
```



### worker\_processes

| Syntax  | worker_processes number   auto; |
|---------|---------------------------------|
| Default | worker_processes 1;             |
| Context | main                            |

Defines the number of worker processes.

The optimal value depends on various factors, including the number of CPU cores, the number of hard disk drives, and the load pattern. If unsure, starting with the number of available CPU cores is recommended. The value auto attempts to autodetect the optimal worker process number.

#### worker rlimit core

| Syntax  | worker_rlimit_core size; |
|---------|--------------------------|
| Default | _                        |
| Context | main                     |

Changes the limit on the maximum size of a core file (RLIMIT\_CORE) for worker processes. This allows increasing the limit without restarting the main process.

#### worker rlimit nofile

| Syntax  | worker_rlimit_nofile $number;$ |
|---------|--------------------------------|
| Default | _                              |
| Context | main                           |

Changes the limit on the maximum number of open files (RLIMIT\_NOFILE) for worker processes. This allows increasing the limit without restarting the main process.

## $worker\_shutdown\_timeout$

| Syntax  | $worker\_shutdown\_timeout\ time;$ |
|---------|------------------------------------|
| Default | _                                  |
| Context | main                               |

Configures a timeout for the graceful shutdown of worker processes. Once the *time* expires, Angie will attempt to close all active connections to complete the shutdown process.

A graceful shutdown is initiated by sending a *QUIT signal* to the master process, which tells the worker processes to stop accepting new connections while allowing existing connections to complete. Workers continue to process active requests until they finish, after which they exit cleanly. If the connections remain open beyond worker\_shutdown\_timeout, Angie will forcefully close those connections to finalize the shutdown. Also, keepalive connections from clients are closed only if they remain inactive for at least the duration of *lingering\_timeout*.

#### working\_directory

| Syntax  | working_directory directory; |
|---------|------------------------------|
| Default | _                            |
| Context | main                         |



Defines the current working directory for a worker process. This is primarily used for writing core files, so the worker process must have write permissions for the specified directory.

#### 3.3.2 HTTP Module

#### Access

The module controls access to server resources based on client IP addresses or networks. It allows to permit or block specific IPs, IP ranges, or UNIX domain sockets to enhance security by restricting access to sensitive areas of a website or application.

Access can also be restricted by using a password with the *Auth Basic* module or based on the result of a subrequest with the *Auth Request* module. To apply both address and password restrictions at the same time, use the *satisfy* directive.

#### **Configuration Example**

```
location / {
    deny 192.168.1.1;
    allow 192.168.1.0/24;
    allow 10.1.1.0/16;
    allow 2001:0db8::/32;
    deny all;
}
```

Rules are evaluated sequentially until a match is found. In this example, access is allowed only for the IPv4 networks 10.1.1.0/16 and 192.168.1.0/24, excluding the specific address 192.168.1.1, and for the IPv6 network 2001:0db8::/32. When there are many rules, it is preferable to use variables from the Geo module.

#### **Directives**

#### allow

| Syntax  | allow address   CIDR   unix:   all;  |
|---------|--------------------------------------|
| Default | _                                    |
| Context | http, server, location, limit_except |

Allows access for a specified network or address. The special value all means all client IPs.

Added in version 1.5.1: The special value unix: allows access for any UNIX domain sockets.

#### deny

| Syntax  | deny address   CIDR   unix:   all;   |
|---------|--------------------------------------|
| Default | _                                    |
| Context | http, server, location, limit_except |

Denies access for a specified network or address. The special value all means all client IPs.

Added in version 1.5.1: The special value unix: denies access for any UNIX domain sockets.



#### **ACME**

Provides automatic certificate retrieval using the ACME protocol.

When building from the source code, this module isn't built by default; it should be enabled with the --with-http\_acme\_module build option.

In packages and images from our repos, the module is included in the build.

#### **Configuration Example**

Examples of configuration and instructions for setup can be found in the acme config section.

#### **Directives**

#### acme

| Syntax  | acme $name$ ; |
|---------|---------------|
| Default | _             |
| Context | server        |

For all domains specified in the <code>server\_name</code> directives of all <code>server</code> blocks that refer to the <code>acme\_client</code> called <code>name</code>, a single certificate will be obtained; if the <code>server\_name</code> configuration changes, the certificate is renewed to reflect the updates.

Each time Angie starts, certificates are requested for all domains that are missing a valid certificate. Possible reasons include expired certificates; missing or unreadable files; configuration changes for a certificate.

#### 1 Note

Currently, domains specified with regular expressions are not supported and will be ignored.

Wildcard domains are supported only with challenge=dns setting in acme\_client.

This directive can be specified multiple times to load certificates of different types, for example, RSA and ECDSA:

```
listen 443 ssl;
server_name example.com www.example.com;

ssl_certificate $acme_cert_rsa;
ssl_certificate_key $acme_cert_key_rsa;

ssl_certificate $acme_cert_ecdsa;
ssl_certificate_key $acme_cert_key_ecdsa;
acme_rsa;
acme_ecdsa;
}
```



#### acme\_client

| Syntax  | $acme\_client \ name \ uri \ [enabled=on \   \ off] \ [key\_type=type] \ [key\_bits=number]$               |
|---------|--|
|         | $[email=email]$ $[max\_cert\_size=number]$ $[renew\_before\_expiry=time]$                                  |
|         | $[	exttt{renew_on_load}]  [	exttt{retry_after_error=} off   time]  [	exttt{challenge=dns}     	ext{http}]$ |
|         | $[\mathtt{account\_key} = file];$  |
| Default |  |
| Context | http   |

Defines an ACME client under a globally unique *name*. It must be valid for a file directory, is a string with variables and will be treated as case insensitive.



The client name defined here identifies it in the Angie configuration, allowing you to match acme\_client, acme directives, and module variables that use this name; don't confuse it with your domain or server name.

The second mandatory parameter is the *uri* of the ACME directory. For example, the URI of Let's Encrypt ACME directory is listed as https://acme-v02.api.letsencrypt.org/directory.

For this directive to work, a resolver must be configured in the same context.

## 1 Note

For testing purposes, certificate authorities usually provide separate staging environments. For example, Let's Encrypt staging environment is currently https://acme-staging-v02.api.letsencrypt.org/directory.



| enabled          | Enables or disables certificate renewal for the client; this is useful, for example, for temporarily suspending without removing the client from the configuration. Default: on.  |
|------------------|---|
| key_type         | The type of private key algorithm for the certificate. Valid values: rsa, ecdsa. Default: ecdsa.  |
| key_bits         | Number of bits in the certificate key. Default: 256 for ecdsa, 2048 for rsa.  |
| email            | Optional email address for feedback; used when creating an account on the CA server.  |
| max_cert_size    | Specifies the maximum allowed size of a new certificate file in bytes to reserve space for a new certificate in shared memory; the more domains the certificate is requested for, the more space is required.  If a certificate already exists at startup but its size exceeds the value of max_cert_size, the max_cert_size value is dynamically increased to match the size of the existing certificate file.  If the size of a renewed certificate exceeds max_cert_size, the renewal will fail with an error.  Default: 8192. |
| renew_before_exp | Time before the certificate expires when certificate renewal should start.<br>Default: 30d.   |
| renew_on_load    | Specifies that the certificate should be forcibly renewed each time the configuration is loaded.  |
| retry_after_erro | Time to retry when a certificate couldn't be obtained. If set to off, the client won't retry to obtain the certificate after a failure.  Default: 2h.   |
| challenge        | Specifies the verification type for the ACME client. Allowed values: dns, http. Default: http.  |
| account_key      | Specifies the full path to a file containing a key in PEM format. This is useful if you want to use an existing account key instead of generating one automatically, or if you need to use one key for multiple ACME clients.   |
|                  | Supported key types:  • RSA keys with lengths that are multiples of 8, ranging from 2048 to 8192 bits.  |
|                  | • ECDSA keys with lengths of 256, 384, or 521 bits.  When specifying the account_key parameter, make sure the key file actually exists. If the file is missing, Angie will attempt to create it at the specified path. Note that keys for ACME clients are created in the order the corresponding clients are mentioned in the configuration in acme_client, acme, or acme_hook directives. Therefore, if one client should use a key created for another, that other client must appear earlier in the configuration.            |
|                  | Additionally, keys are only created for clients that have the enabled=on parameter set.   |

## $acme\_client\_path$

| Syntax  | acme_client_path path; |
|---------|------------------------|
| Default | _                      |
| Context | http                   |

Overrides the *path* to the directory for storing certificates and keys, specified at build time with the --http-acme-client-path build option.



#### acme\_dns\_port

| Syntax  | $acme\_dns\_port   ip:port   [ip6]:port;$ |
|---------|---|
| Default | <pre>acme_dns_port 53;</pre>              |
| Context | http                                      |

Specifies the port that the module uses to handle DNS queries from the ACME server over UDP. The port number must be between 1 and 65535.

Specifying an IP address along with the optional port is also supported. Both IPv4 addresses in the form ip:port and IPv6 addresses in the form [ip6]:port can be used:

```
acme_dns_port 8053;
acme_dns_port 127.0.0.1;
acme_dns_port [::1];
```

To use a port number of 1024 or lower, Angie must run with *superuser* privileges.

#### acme hook

| Syntax  | acme_hook name [uri]; |
|---------|-----------------------|
| Default | _                     |
| Context | location              |

The directive links the server to the specified ACME client. Handler (hook) calls implemented by an external service are made through the location context where it is defined.

| name | Specifies the associated ACME client.                                  |
|------|--|
| uri  | A string with variables; defines the request string for handler calls. |
|      | Default is /.  |

For example, the following configuration passes the values of *hook variables* to a FastCGI application through the request string:

```
acme_hook example uri=/acme_hook/$acme_hook_name?domain=$acme_hook_domain&key=$acme_

hook_keyauth;
fastcgi_param REQUEST_URI $request_uri;
fastcgi_pass ...;
```

#### **Built-in Variables**

```
$acme_cert_<name>
```

Contents of the last certificate file (if any) obtained by the client with this name.

```
$acme_cert_key_<name>
```

Contents of the certificate key file used by the client with this name.

#### Important

The certificate file is available only if the ACME client has received at least one certificate, whereas the key file is available immediately after startup.



#### \$acme\_hook\_challenge

The verification type. Possible values: dns, http.

#### \$acme\_hook\_client

The name of the ACME client initiating the request.

#### \$acme\_hook\_domain

The domain being verified. If it is a wildcard domain, it will be passed without the \*. prefix.

#### \$acme\_hook\_keyauth

The authorization string:

- For DNS verification, it is used as the value of the TXT record, whose name is formed as \_acme-challenge. + \$acme hook domain + ..
- For HTTP verification, this string must be used as the content of the response requested by the ACME server.

#### \$acme\_hook\_name

The name of the hook. For different verification types, it may have different values and meanings:

| Value                      | Meaning for DNS verification   | Meaning for HTTP verification   |
|----------------------------|--|---|
| add (adding the hook)      | The corresponding TXT record<br>must be added to the DNS con-<br>figuration. | The appropriate response to the HTTP request must be prepared.  |
| remove (removing the hook) | The TXT record can be removed from the DNS configuration.                    | This HTTP request is no longer<br>relevant; a previously created<br>file with the authorization string<br>can be removed. |

#### \$acme\_hook\_token

The verification token. For HTTP verification, it is used as the name of the requested file:  $/.well-known/acme-challenge/+ acme_hook_token$ .

#### Addition

The module is a filter that adds text before and after a response.

When building from the source code, this module isn't built by default; it should be enabled with the --with-http\_addition\_module build option.

In packages and images from our repos, the module is included in the build.

#### **Configuration Example**

```
location / {
   add_before_body /before_action;
   add_after_body /after_action;
}
```



#### **Directives**

#### add\_before\_body

| Syntax  | add_before_body uri;   |
|---------|------------------------|
| Default | _                      |
| Context | http, server, location |

Adds the text returned as a result of processing a given subrequest before the response body. An empty string ("") as a parameter cancels addition inherited from the previous configuration level.

#### add\_after\_body

| Syntax  | add_after_body uri;    |
|---------|------------------------|
| Default | _                      |
| Context | http, server, location |

Adds the text returned as a result of processing a given subrequest after the response body. An empty string ("") as a parameter cancels addition inherited from the previous configuration level.

#### addition types

| Syntax  | addition_types mime-type; |
|---------|---------------------------|
| Default | addition_types text/html; |
| Context | http, server, location    |

Allows adding text in responses with the specified MIME types, in addition to "text/html". The special value "\*" matches any MIME type.

#### **API**

The module provides HTTP RESTful interface for accessing in JSON format basic information about a web server instance, as well as metrics of client connections, shared memory zones, DNS queries, HTTP requests, HTTP responses cache, TCP/UDP sessions of *Stream Module*, and zones of *Limit Conn*, *Limit Conn*, *Limit Reg* and *Upstream* modules.

The API accepts GET and HEAD HTTP requests; other request methods cause an error:

```
{
    "error": "MethodNotAllowed",
    "description": "The POST method is not allowed for the requested API element \"/\
    \"."
}
```

The Angie PRO API has a *dynamic configuration* that allows updating the settings without reloading the configuration or restarting Angie itself; currently, it enables configuring individual peers in an upstream.

#### **Directives**

#### api

| Syntax  | api $path;$ |
|---------|-------------|
| Default | _           |
| Context | location    |



Enables HTTP RESTful interface in location.

path parameter is mandatory and works similar to the alias directive, but operates over API tree, rather than filesystem hierarchy.

When specified in a prefix location:

```
location /stats/ {
    api /status/http/server_zones/;
}
```

the part of request URI matching particular prefix location /stats/ will be replaced by the path / status/http/server\_zones/ in the directive parameter. For example, on a request of /stats/foo/ the /status/http/server\_zones/foo/ API element will be accessed.

Also it's possible to use variables:  $api/status/\$module/server\_zones/\$name/$  and the directive can also be specified inside a regexp location:

```
location ~^/api/([^/]+)/(.*)$ {
    api /status/http/$1_zones/$2;
}
```

here, similar to *alias*, parameter defines the whole path to the API element. E.g., from /api/location/data/ the following positional variables will be populated:

```
$1 = "location"
$2 = "data/"
```

which results after interpolation in /status/http/location\_zones/data/.

## Note

In Angie PRO, you can decouple the  $dynamic \ configuration \ API$  from the immutable  $metrics \ API$  that reflects current state:

```
location /config/ {
    api /config/;
}
location /status/ {
    api /status/;
}
```

Overall, it serves for precise configuration of API access rights, e.g.:

```
location /status/ {
    api /status/;
    allow 127.0.0.1;
    deny all;
}
```

Or:



#### api\_config\_files

```
Syntax api_config_files on | off;
Default off
Context location
```

Enables or disables the config\_files object, which enumerates the contents of all Angie config files that are currently loaded by the server instance, in the /status/angie/ API section. For example, with this configuration:

```
location /status/ {
    api /status/;
    api_config_files on;
}
```

A query to /status/angie/ returns approximately this:

By default, the object is disabled because the config files can contain extra sensitive, confidential details.

#### Metrics

Angie exposes usage metrics in the /status/ section of the API; you can make it accessible by defining a respective location. Full access:

```
location /status/ {
    api /status/;
}
```

Subtree access, already discussed earlier:

```
location /stats/ {
    api /status/http/server_zones/;
}
```

#### **Example configuration**

Configured with location /status/, resolver, http upstream, http server, location, cache, limit\_conn in http and limit\_req zones:

```
http {
    resolver 127.0.0.53 status_zone=resolver_zone;
    proxy_cache_path /var/cache/angie/cache keys_zone=cache_zone:2m;
    limit_conn_zone $binary_remote_addr zone=limit_conn_zone:10m;
    limit_req_zone $binary_remote_addr zone=limit_req_zone:10m rate=1r/s;
}
```



```
upstream upstream {
        zone upstream 256k;
        server backend.example.com service=_example._tcp resolve max_conns=5;
    server {
        server_name www.example.com;
        listen 443 ssl;
        status_zone http_server_zone;
        proxy_cache cache_zone;
        access_log /var/log/access.log main;
        location / {
            root /usr/share/angie/html;
            status_zone location_zone;
            limit_conn limit_conn_zone 1;
            limit_req zone=limit_req_zone burst=5;
        }
        location /status/ {
            api /status/;
            allow 127.0.0.1;
            deny all;
        }
   }
}
```

Responds to curl https://www.example.com/status/ with the following JSON:

#### JSON tree

```
"angie": {
    "version":"1.9.1",
    "address": "192.168.16.5",
    "generation":1,
    "load_time":"2025-05-29T12:58:39.789Z"
},
"connections": {
    "accepted":2257,
    "dropped":0,
    "active":3,
    "idle":1
},
"slabs": {
    "cache_zone": {
        "pages": {
             "used":2,
            "free":506
        },
```



```
"slots": {
        "64": {
            "used":1,
            "free":63,
            "reqs":1,
            "fails":0
        },
        "512": {
            "used":1,
            "free":7,
            "reqs":1,
            "fails":0
        }
    }
},
"limit_conn_zone": {
    "pages": {
        "used":2,
        "free":2542
    },
    "slots": {
        "64": {
            "used":1,
            "free":63,
            "reqs":74,
            "fails":0
        },
        "128": {
            "used":1,
            "free":31,
            "reqs":1,
            "fails":0
        }
   }
},
"limit_req_zone": {
    "pages": {
        "used":2,
        "free":2542
    },
    "slots": {
        "64": {
            "used":1,
            "free":63,
            "reqs":1,
            "fails":0
        },
        "128": {
            "used":2,
            "free":30,
```



```
"reqs":3,
                 "fails":0
            }
        }
    }
},
"http": {
    "server_zones": {
        "http_server_zone": {
            "ssl": {
                "handshaked":4174,
                 "reuses":0,
                 "timedout":0,
                 "failed":0
            },
            "requests": {
                "total":4327,
                 "processing":0,
                 "discarded":8
            },
            "responses": {
                 "200":4305,
                 "302":12,
                 "404":4
            },
            "data": {
                 "received":733955,
                 "sent":59207757
            }
        }
    },
    "location_zones": {
        "location_zone": {
            "requests": {
                 "total":4158,
                 "discarded":0
            },
            "responses": {
                 "200":4157,
                 "304":1
            },
            "data": {
                 "received":538200,
                 "sent":177606236
            }
        }
    },
    "caches": {
        "cache_zone": {
            "size":0,
```



```
"cold":false,
        "hit": {
            "responses":0,
            "bytes":0
        },
        "stale": {
            "responses":0,
            "bytes":0
        },
        "updating": {
            "responses":0,
            "bytes":0
        },
        "revalidated": {
            "responses":0,
            "bytes":0
        },
        "miss": {
            "responses":0,
            "bytes":0,
            "responses_written":0,
            "bytes_written":0
        },
        "expired": {
            "responses":0,
            "bytes":0,
            "responses_written":0,
            "bytes_written":0
        },
        "bypass": {
            "responses":0,
            "bytes":0,
            "responses_written":0,
            "bytes_written":0
        }
   }
},
"limit_conns": {
    "limit_conn_zone": {
        "passed":73,
        "skipped":0,
        "rejected":0,
        "exhausted":0
    }
},
"limit_reqs": {
    "limit_req_zone": {
        "passed":54816,
        "skipped":0,
```



```
"delayed":65,
            "rejected":26,
            "exhausted":0
        }
    },
    "upstreams": {
        "upstream": {
             "peers": {
                 "192.168.16.4:80": {
                     "server": "backend.example.com",
                     "service": "_example._tcp",
                     "backup":false,
                     "weight":5,
                     "state":"up",
                     "selected": {
                         "current":2,
                         "total":232
                     },
                     "max_conns":5,
                     "responses": {
                         "200":222,
                         "302":12
                     },
                     "data": {
                         "sent":543866,
                         "received":27349934
                     },
                     "health": {
                         "fails":0,
                         "unavailable":0,
                         "downtime":0
                     },
                     "sid":"<server_id>"
                }
            },
            "keepalive":2
        }
    }
},
"resolvers": {
    "resolver_zone": {
        "queries": {
            "name":442,
            "srv":2,
            "addr":0
        },
        "responses": {
            "success":440,
             "timedout":1,
```



Each JSON branch can be requested separately with the request constructed accordingly, e.g.:

```
$ curl https://www.example.com/status/angie
$ curl https://www.example.com/status/connections
$ curl https://www.example.com/status/slabs
$ curl https://www.example.com/status/slabs/<zone>/slots
$ curl https://www.example.com/status/slabs/<zone>/slots/64
$ curl https://www.example.com/status/http/
$ curl https://www.example.com/status/http/server_zones
$ curl https://www.example.com/status/http/server_zones/<http_server_zone>/ssl
```

## 1 Note

By default, the module uses ISO 8601 strings for date values; to use the integer epoch format instead, add the date=epoch parameter to the query string:

```
$ curl https://www.example.com/status/angie/load_time

"2024-04-01T00:59:59+01:00"

$ curl https://www.example.com/status/angie/load_time?date=epoch
1711929599
```

#### Server status

/status/angie



| version      | String; version of the running Angie web server  |
|--------------|--|
| build        | String; particular build name when it specified during compilation   |
| build_time   | String; the build time of the Angie executable in the date format  |
| address      | String; the address of the server that accepted API request  |
| generation   | Number; total number of configuration reloads since last start   |
| load_time    | String or number; time of the last configuration reload, formatted as a <i>date</i> ; strings have millisecond resolution  |
| config_files | Object; its members are absolute pathnames of all Angie configuration files that are currently loaded by the server instance, and their values are string representations of the files' contents, for example:  {     "/etc/angie/angie.conf": "server {\n listen 80;\n #\     ¬\n}\n" } |
|              | * Caution  |
|              | The config_files object is available in /status/angie/ only if the api_config_files directive is enabled.  |

## Connections global metrics

/status/connections

```
{
   "accepted": 2257,
   "dropped": 0,
   "active": 3,
   "idle": 1
}
```

| accepted | Number; the total number of accepted client connections |
|----------|---|
| dropped  | Number; the total number of dropped client connections  |
| active   | Number; the current number of active client connections |
| idle     | Number; the current number of idle client connections   |

## Slab allocator metrics of shared memory zones

/status/slabs/<zone>

Usage statistics of shared memory zones that utilize slab allocation, such as  $limit\_conn$ ,  $limit\_req$ , and  $HTTP\ cache$ :

```
limit_conn_zone $binary_remote_addr zone=limit_conn_zone:10m;
limit_req_zone $binary_remote_addr zone=limit_req_zone:10m rate=1r/s;
proxy_cache cache_zone;
```

The specified shared memory zone will collect the following statistics:



| pages | Object; memory pages statistics  |
|-------|--|
| used  | Number; the number of currently used memory pages  |
| free  | Number; the number of currently free memory pages  |
| slots | Object; memory slots statistics for each slot size. The slots object contains fields for requested memory slot sizes (e.g. 8, 16, 32, etc., up to half of the page size) |
| used  | Number; the number of currently used memory slots of specified size  |
| free  | Number; the number of currently free memory slots of specified size  |
| reqs  | Number; the total number of attempts to allocate memory slot of specified size   |
| fails | Number; the number of unsuccessful attempts to allocate memory slot of specified size  |

### Example:

```
{
    "pages": {
        "used": 2,
        "free": 506
},

    "slots": {
        "64": {
            "used": 1,
            "free": 63,
            "reqs": 1,
            "fails": 0
}
}
```

## **Resolver DNS queries**

# /status/resolvers/<zone>

To collect resolver statistics, the resolver directive must set the  $status\_zone$  parameter (HTTP and Stream):

```
resolver 127.0.0.53 status_zone=resolver_zone;
```

The specified shared memory zone will collect the following statistics:



| queries        | Object; queries statistics   |
|----------------|--|
| name           | Number; the number of queries to resolve names to addresses (A and AAAA queries) |
| srv            | Number; the number of queries to resolve services to addresses (SRV queries)     |
| addr           | Number; the number of queries to resolve addresses to names (PTR queries)        |
| responses      | Object; responses statistics   |
| success        | Number; the number of successful responses                                       |
| timedout       | Number; the number of timed out queries  |
| format_error   | Number; the number responses with code 1 (Format Error)                          |
| server_failure | Number; the number responses with code 2 (Server Failure)                        |
| not_found      | Number; the number responses with code 3 (Name Error)                            |
| unimplemented  | Number; the number responses with code 4 (Not Implemented)                       |
| refused        | Number; the number responses with code 5 (Refused)                               |
| other          | Number; the number of queries completed with other non-zero code                 |
| sent           | Object; sent DNS queries statistics  |
| a              | Number; the number of A type queries   |
| aaaa           | Number; the number of AAAA type queries  |
| ptr            | Number; the number of PTR type queries   |
| srv            | Number; the number of SRV type queries   |

The response codes are described in RFC 1035, section 4.1.1.

Various DNS record types are detailed in RFC 1035, RFC 2782, and RFC 3596.

Example:

```
"queries": {
    "name": 442,
    "srv": 2,
    "addr": 0
  },
  "responses": {
    "success": 440,
    "timedout": 1,
    "format_error": 0,
    "server_failure": 1,
    "not_found": 1,
    "unimplemented": 0,
    "refused": 1,
    "other": 0
  },
  "sent": {
    "a": 185,
    "aaaa": 245,
    "srv": 2,
    "ptr": 12
  }
}
```

## **HTTP** server and location

/status/http/server\_zones/<zone>

To collect the server metrics, set the *status\_zone* directive in the *server* context:



```
server {
    ...
    status_zone server_zone;
}
```

To group the metrics by a custom value, use the alternative syntax. Here, the metrics are aggregated by \$host, with each group reported as a standalone zone:

```
status_zone $host zone=server_zone:5;
```

The specified shared memory zone will collect the following statistics:

| ssl           | Object; SSL statistics. Present if server sets listen ssl;                 |
|---------------|--|
| handshaked    | Number; the total number of successful SSL handshakes                      |
| reuses        | Number; the total number of session reuses during SSL handshake            |
| timedout      | Number; the total number of timed out SSL handshakes                       |
| failed        | Number; the total number of failed SSL handshakes                          |
| requests      | Object; requests statistics  |
| total         | Number; the total number of client requests                                |
| processing    | Number; the number of currently being processed client requests            |
| discarded     | Number; the total number of client requests completed without sending a    |
|               | response   |
| responses     | Object; responses statistics   |
| <code></code> | Number; a non-zero number of responses with status <code> (100-599)</code> |
| xxx           | Number; a non-zero number of responses with other status codes             |
| data          | Object; data statistics  |
| received      | Number; the total number of bytes received from clients                    |
| sent          | Number; the total number of bytes sent to clients                          |

### Example:

```
{
    "ssl":{
        "handshaked":4174,
        "reuses":0,
        "timedout":0,
        "failed":0
    },
    "requests":{
        "total":4327,
        "processing":0,
        "discarded":0
    },
    "responses":{
        "200":4305,
        "302":6,
        "304":12,
        "404":4
   },
    "data":{
        "received":733955,
        "sent":59207757
    }
}
```



/status/http/location\_zones/<zone>

To collect the location metrics, set the status zone directive in the context of location or if in location:

```
location / {
    root /usr/share/angie/html;
    status_zone location_zone;

if ($request_uri ~* "^/condition") {
    # ...
    status_zone if_location_zone;
  }
}
```

To group the metrics by a custom value, use the alternative syntax. Here, the metrics are aggregated by \$host, with each group reported as a standalone zone:

```
status_zone $host zone=server_zone:5;
```

The specified shared memory zone will collect the following statistics:

| requests      | Object; requests statistics  |
|---------------|--|
| total         | Number; the total number of client requests                                |
| discarded     | Number; the total number of client requests completed without sending a    |
|               | response   |
| responses     | Object; responses statistics   |
| <code></code> | Number; a non-zero number of responses with status <code> (100-599)</code> |
| XXX           | Number; a non-zero number of responses with other status codes             |
| data          | Object; data statistics  |
| received      | Number; the total number of bytes received from clients                    |
| sent          | Number; the total number of bytes sent to clients                          |

### Example:

```
{
    "requests": {
        "total": 4158,
        "discarded": 0
},

"responses": {
        "200": 4157,
        "304": 1
},

"data": {
        "received": 538200,
        "sent": 177606236
}
}
```

## Stream server

/status/stream/server\_zones/<zone>

To collect the server metrics, set the status\_zone directive in the server context:



```
server {
    ...
    status_zone server_zone;
}
```

To group the metrics by a custom value, use the alternative syntax. Here, the metrics are aggregated by \$host, with each group reported as a standalone zone:

```
status_zone $host zone=server_zone:5;
```

The specified shared memory zone will collect the following statistics:

| Object; SSL statistics. Present if server sets listen ssl;  |
|---|
| Number; the total number of successful SSL handshakes   |
| Number; the total number of session reuses during SSL handshake   |
| Number; the total number of timed out SSL handshakes  |
| Number; the total number of failed SSL handshakes   |
| Object; connections statistics  |
| Number; the total number of client connections  |
| Number; the number of currently being processed client connections  |
| Number; the total number of client connections completed without creating a session   |
| Number; the total number of client connections relayed to another listening port with pass directives   |
| Object; sessions statistics   |
| Number; the number of sessions completed with code 200, which means successful completion   |
| Number; the number of sessions completed with code 400, which happens when client data could not be parsed, e.g. the PROXY protocol header          |
| Number; the number of sessions completed with code 403, when access was forbidden, for example, when access is limited for certain client addresses |
| Number; the number of sessions completed with code 500, the internal server error   |
| Number; the number of sessions completed with code 502, bad gateway, for example, if an upstream server could not be selected or reached            |
| Number; the number of sessions completed with code 503, service unavailable, for example, when access is limited by the number of connections       |
| Object; data statistics   |
| Number; the total number of bytes received from clients   |
| Number; the total number of bytes sent to clients   |
|   |

# Example:

```
{
    "ssl": {
        "handshaked": 24,
        "reuses": 0,
        "timedout": 0,
        "failed": 0
},

"connections": {
        "total": 24,
        "processing": 1,
        "discarded": 0,
        "passed": 2
},
```



```
"sessions": {
    "success": 24,
    "invalid": 0,
    "forbidden": 0,
    "internal_error": 0,
    "bad_gateway": 0,
    "service_unavailable": 0
},

"data": {
    "received": 2762947,
    "sent": 53495723
}
}
```

#### **HTTP** caches

```
proxy_cache cache_zone;
```

/status/http/caches/<cache>

For each zone configured with *proxy* cache, the following data is stored:

```
{
 "name_zone": {
    "size": 0,
    "cold": false,
    "hit": {
      "responses": 0,
      "bytes": 0
    },
    "stale": {
      "responses": 0,
      "bytes": 0
    },
    "updating": {
      "responses": 0,
      "bytes": 0
    },
    "revalidated": {
      "responses": 0,
      "bytes": 0
    },
    "miss": {
      "responses": 0,
      "bytes": 0,
      "responses_written": 0,
      "bytes_written": 0
    },
    "expired": {
```



```
"responses": 0,
   "bytes": 0,
   "responses_written": 0,
   "bytes_written": 0
},

"bypass": {
   "responses": 0,
   "bytes": 0,
   "responses_written": 0,
   "bytes_written": 0
}
}
```

| size              | Number; the current size of the cache  |
|-------------------|--|
| max_size          | Number; configured limit on the maximum size of the cache                    |
| cold              | Boolean; true while the cache loader loads data from disk                    |
| hit               | Object; statistics of valid cached responses (proxy_cache_valid)             |
| responses         | Number; the total number of responses read from the cache                    |
| bytes             | Number; the total number of bytes read from the cache                        |
| stale             | Object; statistics of expired responses taken from the cache                 |
|                   | $(proxy\_cache\_use\_stale)$   |
| responses         | Number; the total number of responses read from the cache                    |
| bytes             | Number; the total number of bytes read from the cache                        |
| updating          | Object; statistics of expired responses taken from the cache while responses |
|                   | were being updated (proxy cache use stale updating)                          |
| responses         | Number; the total number of responses read from the cache                    |
| bytes             | Number; the total number of bytes read from the cache                        |
| revalidated       | Object; statistics of expired and revalidated responses taken from the cache |
|                   | $(proxy\_cache\_revalidate)$   |
| responses         | Number; the total number of responses read from the cache                    |
| bytes             | Number; the total number of bytes read from the cache                        |
| miss              | Object; statistics of responses not found in the cache                       |
| responses         | Number; the total number of corresponding responses                          |
| bytes             | Number; the total number of bytes read from the proxied server               |
| responses_written | Number; the total number of responses written to the cache                   |
| bytes_written     | Number; the total number of bytes written to the cache                       |
| expired           | Object; statistics of expired responses not taken from the cache             |
| responses         | Number; the total number of corresponding responses                          |
| bytes             | Number; the total number of bytes read from the proxied server               |
| responses_written | Number; the total number of responses written to the cache                   |
| bytes_written     | Number; the total number of bytes written to the cache                       |
| bypass            | Object; statistics of responses not looked up in the cache                   |
|                   | $(proxy\_cache\_bypass)$   |
| responses         | Number; the total number of corresponding responses                          |
| bytes             | Number; the total number of bytes read from the proxied server               |
| responses_written | Number; the total number of responses written to the cache                   |
| bytes_written     | Number; the total number of bytes written to the cache                       |
| <u> </u>          | , ,  |

### Added in version 1.2.0: PRO

In Angie PRO, if *cache sharding* is enabled with *proxy\_cache\_path* directives, individual shards are exposed as object members of a **shards** object:



| shards          | Object; lists individual shards as members                          |
|-----------------|---|
| <shard></shard> | Object; represents an individual shard with its cache path for name |
| size            | Number; the shard's current size                                    |
| max_size        | Number; maximum shard size, if configured                           |
| cold            | Boolean; true while the cache loader loads data from disk           |

### limit conn

```
limit_conn_zone $binary_remote_addr zone=limit_conn_zone:10m;
```

/status/http/limit\_conns/<zone>, /status/stream/limit\_conns/<zone>

Objects for each configured  $limit\_conn\ in\ http$  or  $limit\_conn\ in\ stream$  contexts with the following fields:

```
{
   "passed": 73,
   "skipped": 0,
   "rejected": 0,
   "exhausted": 0
}
```

| passed    | Number; the total number of passed connections  |
|-----------|---|
| skipped   | Number; the total number of connections passed with zero-length key, or key exceeding 255 bytes |
| rejected  | Number; the total number of connections exceeding the configured limit                          |
| exhausted | Number; the total number of connections rejected due to exhaustion of zone storage              |

### limit\_req

```
limit_req_zone $binary_remote_addr zone=limit_req_zone:10m rate=1r/s;
```

/status/http/limit\_reqs/<zone>

Objects for each configured *limit req* with the following fields:

```
{
    "passed": 54816,
```



```
"skipped": 0,
"delayed": 65,
"rejected": 26,
"exhausted": 0
}
```

| passed    | Number; the total number of passed connections   |
|-----------|--|
| skipped   | Number; the total number of requests passed with zero-length key, or key exceeding 65535 bytes |
| delayed   | Number; the total number of delayed requests   |
| rejected  | Number; the total number of rejected requests  |
| exhausted | Number; the total number of requests rejected due to exhaustion of zone storage                |

## HTTP upstream

Added in version 1.1.0.

To enable collection of the following metrics, set the zone directive in the upstream context, for instance:

```
upstream upstream {
    zone upstream 256k;
    server backend.example.com service=_example._tcp resolve max_conns=5;
    keepalive 4;
}
```

/status/http/upstreams/<upstream>

where <upstream> is the name of any upstream specified with the zone directive

```
{
    "peers": {
        "192.168.16.4:80": {
            "server": "backend.example.com",
            "service": "_example._tcp",
            "backup": false,
            "weight": 5,
            "state": "up",
            "selected": {
                "current": 2,
                "total": 232
            },
            "max_conns": 5,
            "responses": {
                "200": 222,
                "302": 12
            },
            "data": {
                "sent": 543866,
                "received": 27349934
            },
            "health": {
                 "fails": 0,
```



```
"unavailable": 0,
                "downtime": 0
            },
            "sid": "<server_id>"
        }
   },
    "keepalive": 2
}
```



| peers  | Object; contains the metrics of the upstream's peers as subobjects whose names are canonical representations of the peers' addresses. Members of each subobject: |
|--|--|
| server   | String; the parameter of the server directive  |
| service  | String; name of service as it's specified in <i>server</i> directive, if configured  |
| slow_start   | Number; the specified <i>slow start</i> value for the server, expressed in seconds.  |
| (PRO 1.4.0+)   | When setting the value via the respective subsection of the dynamic configu-   |
| (1100 1.110   )  | ration API, you can specify either a number or a time value with millisecond   |
|  | precision.   |
| backup   | Boolean; true for backup servers   |
| weight   | Number; configured weight  |
| state  | String; the current state of the peer and what requests are sent to it:  |
|  | • busy: indicates that the number of requests to the server has reached  |
|  | the limit set by <i>max conns</i> , and no new requests are sent   |
|  | • down: manually disabled, no requests are sent  |
|  | • recovering: recovering after a failure according to <i>slow start</i> , more   |
|  | and more requests are sent over time   |
|  | • unavailable: reached the max fails limit, only trial client requests   |
|  | are sent at intervals defined by fail_timeout;   |
|  | • up: operational, requests are sent as usual  |
|  | Additional states in Angie PRO:  |
|  | • checking: configured as essential and being checked, only probe  |
|  | requests are sent  |
|  | • draining: similar to down, but requests from previously bound ses-   |
|  | sions (via <i>sticky</i> ) are still sent  |
|  | • unhealthy: non-operational, only <i>probe requests</i> are sent  |
| selected   | Object; peer selection statistics  |
| current  | Number; the current number of connections to peer  |
| total  | Number; total number of requests forwarded to peer   |
| last   | String or number; time when peer was last selected, formatted as a <i>date</i>   |
| max_conns  | Number; the configured <i>maximum</i> number of simultaneous connections, if   |
|  | specified  |
| responses  | Object; responses statistics   |
| <code></code>  | Number; a non-zero number of responses with status < code > (100-599)  |
| xxx  | Number; a non-zero number of responses with other status codes   |
| data   | Object; data statistics  |
| received   | Number; the total number of bytes received from peer   |
| sent   | Number; the total number of bytes sent to peer   |
| health   | Object; health statistics  |
| fails  | Number; the total number of unsuccessful attempts to communicate with  |
|  | the peer   |
| unavailable  | Number; how many times peer became unavailable due to reaching the   |
|  | max_fails limit  |
| downtime   | Number; the total time (in milliseconds) when peer was unavailable for   |
| 1  | selection  |
| downstart  | String or number; time when peer became unavailable, formatted as a  |
| hooden tier-   | date Number; average time (in milliseconds) to receive the response headers from   |
| $\begin{array}{c} \texttt{header\_time} \\ (\text{PRO } 1.3.0+) \end{array}$ | Number; average time (in miniseconds) to receive the response headers from the peer; see response time factor $(PRO)$  |
| response_time  | Number; average time (in milliseconds) to receive the entire peer response;  |
| (PRO 1.3.0+)   | see response time factor (PRO)   |
| (1100 1.3.0+)  | String; configured id of the server in upstream group  |
| keepalive  | Number; the number of currently cached connections   |
| backup_switch  | Object; contains the current state of the active backup logic, present if  |
| 240114P_DWI0011  | backup switch (PRO) is configured for the upstream   |
| active   | Number; active group identifier, if any  |
| timeout  | Number; time to expire in milliseconds, after which the balancer will re-  |
|  | check the groups for healthy peers; does not appear for the primary group  |
| 2 2 2 1 1 1 1  | 1  |



## health/probes (PRO)

#### Changed in version 1.2.0: PRO

If the upstream has  $upstream\_probe\ (PRO)$  probes configured, the health object also has a probes subobject that stores the peer's health probe counters, while the peer's state can also be checking and unhealthy, apart from the values listed in the table above:

```
{
    "192.168.16.4:80": {
        "state": "unhealthy",
        "...": "...",
        "health": {
            "count": 10,
            "fails": 10,
            "last": "2025-05-29T09:56:07Z"
            }
        }
    }
}
```

The checking value of state isn't counted as downtime and means that the peer, which has a probe configured as essential, hasn't been checked yet; the unhealthy value means that the peer is malfunctioning. Both states also imply that the peer isn't included in load balancing. For details of health probes, see *upstream\_probe*.

## Counters in probes:

| count | Number; total probes for this peer                     |
|-------|--|
| fails | Number; total failed probes                            |
| last  | String or number; last probe time, formatted as a date |

#### queue

## Changed in version 1.4.0.

If a request queue is configured for the upstream, the upstream object also contains a nested queue object, which holds counters for requests in the queue:

```
{
    "queue": {
        "queued": 20112,
        "waiting": 1011,
        "dropped": 6031,
        "timedout": 560,
        "overflows": 13
    }
}
```

The counter values are aggregated across all worker processes:

| queued    | Number; total count of requests that entered the queue  |
|-----------|---|
| waiting   | Number; current count of requests in the queue  |
| dropped   | Number; total count of requests removed from the queue due to the client prematurely closing the connection |
| timedout  | Number; total count of requests removed from the queue due to timeout                                       |
| overflows | Number; total count of queue overflow occurrences   |



### Stream upstream

To enable collection of the following metrics, set the zone directive in the upstream context, for instance:

```
upstream upstream {
    zone upstream 256k;
    server backend.example.com service=_example._tcp resolve max_conns=5;
    keepalive 4;
}
```

/status/stream/upstreams/<upstream>

Here,  $\langle upstream \rangle$  is the name of an upstream that is configured with a zone directive.

```
"peers": {
        "192.168.16.4:1935": {
            "server": "backend.example.com",
            "service": "_example._tcp",
            "backup": false,
            "weight": 5,
            "state": "up",
            "selected": {
                 "current": 2,
                "total": 232
            },
            "max_conns": 5,
            "data": {
                "sent": 543866,
                "received": 27349934
            },
            "health": {
                "fails": 0,
                "unavailable": 0,
                "downtime": 0
            }
        }
    }
}
```



| peers   | Object; contains the metrics of the upstream's peers as subobjects whose names are canonical representations of the peers' addresses. Members of each subobject:  |
|---|---|
| server  | String; address set by the <i>server</i> directive  |
| service   | String; service name, if set by <i>server</i> directive   |
| slow_start  | Number; the specified <i>slow</i> start value for the server, expressed in seconds.   |
| (PRO 1.4.0+)  | When setting the value via the <i>respective subsection</i> of the dynamic configuration API, you can specify either a number or a time value with millisecond precision.   |
| backup  | Boolean; true for backup server   |
| weight  | Number; the weight of the peer  |
| state   | <ul> <li>String; the current state of the peer and what requests are sent to it:</li> <li>busy: indicates that the number of requests to the server has reached the limit set by max_conns, and no new requests are sent</li> <li>down: manually disabled, no requests are sent</li> <li>recovering: recovering after a failure according to slow_start, more and more requests are sent over time</li> <li>unavailable: reached the max_fails limit, only trial client requests are sent at intervals defined by fail_timeout;</li> <li>up: operational, requests are sent as usual</li> <li>Additional states in Angie PRO:</li> <li>checking: configured as essential and being checked, only probe requests are sent</li> <li>draining: similar to down, but requests from previously bound sessions (via sticky) are still sent</li> <li>unhealthy: non-operational, only probe requests are sent</li> </ul> |
| selected  | Object; the peer's selection metrics  |
|   |   |
| current<br>total  | Number; current connections to the peer Number; total connections forwarded to the peer   |
|   |   |
| last<br>max_conns   | String or number; time when the peer was last selected, formatted as a <i>date</i> Number; <i>maximum</i> number of simultaneous active connections to the peer, if set   |
| data  | Object; data transfer metrics   |
| received  | Number; total bytes received from the peer  |
| sent  | Number; total bytes sent to the peer  |
| health  | Object; peer health metrics   |
| fails   | Number; total failed attempts to reach the peer   |
| unavailable   | Number; times the peer became unavailable due to reaching the max fails   |
| downtime  | Number; total time (in milliseconds) that the peer was unavailable for selection  |
| downstart   | String or number; time when the peer last became ${\tt unavailable}$ , formatted as a ${\it date}$  |
| $\begin{array}{c} \texttt{connect\_time} \\ (\text{PRO } 1.4.0+) \end{array}$     | Number; average time (in milliseconds) taken to establish a connection with the peer; see the <i>response_time_factor</i> ( <i>PRO</i> ) directive.   |
| $\begin{array}{c} \texttt{first\_byte\_time} \\ (\text{PRO } 1.4.0+) \end{array}$ | Number; average time (in milliseconds) to receive the first byte of the response from the peer; see the $response\_time\_factor~(PRO)$ directive.   |
| last_byte_time (PRO 1.4.0+)   | Number; average time (in milliseconds) to receive the complete response from the peer; see the <i>response_time_factor</i> ( <i>PRO</i> ) directive.  |

## Changed in version 1.4.0: PRO

In Angie PRO, if the upstream has  $upstream\_probe~(PRO)$  probes configured, the health object also has a probes subobject that stores the peer's health probe counters, while the peer's state can also be checking and unhealthy, apart from the values listed in the table above:



The checking value of state isn't counted as downtime and means that the peer, which has a probe configured as essential, hasn't been checked yet; the unhealthy value means that the peer is malfunctioning. Both states also imply that the peer isn't included in load balancing. For details of health probes, see *upstream probe*.

#### Counters in probes:

| count | Number; total probes for this peer                     |
|-------|--|
| fails | Number; total failed probes                            |
| last  | String or number; last probe time, formatted as a date |

## Dynamic Configuration API (PRO only)

Added in version 1.2.0.

The API includes a /config section that enables dynamic updates to Angie's configuration in JSON with PUT, PATCH, and DELETE HTTP requests. All updates are atomic; new settings are applied as a whole, or none are applied at all. On error, Angie reports the reason.

### Subsections of /config

Currently, configuration of individual servers within upstreams is available in the /config section for the *HTTP* and *stream* modules; the number of settings eligible for dynamic configuration is steadily increasing.

#### /config/http/upstreams/<upstream>/servers/<name>

Enables configuring individual upstream peers, including deleting existing peers or adding new ones. URI path parameters:

| <upstream></upstream> | Name of the upstream; to be configurable via /config, it must have a zone directive configured, defining a shared memory zone.   |
|-----------------------|--|
| <name></name>         | <ul> <li>The peer's name within the upstream, defined as <service>@<host>, where:</host></service></li> <li><service>@ is an optional service name, used for SRV record resolution.</service></li> <li><host> is the domain name of the service (if resolve is present) or its IP; an optional port can be defined here.</host></li> </ul> |

For example, the following configuration:



```
upstream backend {
    server backend.example.com service=_http._tcp resolve;
    server 127.0.0.1;
    zone backend 1m;
}
```

Allows the following peer names:

This API subsection enables setting the weight, max\_conns, max\_fails, fail\_timeout, backup, down and sid parameters, as described in *server*.

```
There is no separate drain option here; to enable drain, set down to the string value drain:

$ curl -X PUT -d \"drain\" \
http://127.0.0.1/config/http/upstreams/backend/servers/backend.example.com/down
```

#### Example:

```
curl http://127.0.0.1/config/http/upstreams/backend/servers/backend.example.com?
```

```
{
    "weight": 1,
    "max_conns": 0,
    "max_fails": 1,
    "fail_timeout": 10,
    "backup": true,
    "down": false,
    "sid": ""
}
```

Actually available parameters are limited to the ones supported by the current load balancing method of the *upstream*. So, if the upstream is configured as random:

```
upstream backend {
   zone backend 256k;
   server backend.example.com resolve max_conns=5;
   random;
}
```

You will be unable to add a new peer that defines backup:

```
$ curl -X PUT -d '{ "backup": true }' \
http://127.0.0.1/config/http/upstreams/backend/servers/backend1.example.com
```

```
{
    "error": "FormatError",
    "description": "The \"backup\" field is unknown."
}
```



## 1 Note

Even with a compatible load balancing method, the backup parameter can only be set at new peer creation.

#### /config/stream/upstreams/<upstream>/servers/<name>

Enables configuring individual upstream peers, including deleting existing peers or adding new ones. URI path parameters:

| <upstream></upstream> | Name of the upstream; to be configurable via /config, it must have a zone directive configured, defining a shared memory zone.   |
|-----------------------|--|
| <name></name>         | <ul> <li>The peer's name within the upstream, defined as <service>@<host>, where:</host></service></li> <li><service>@ is an optional service name, used for SRV record resolution.</service></li> <li><host> is the domain name of the service (if resolve is present) or its IP; an optional port can be defined here.</host></li> </ul> |

For example, the following configuration:

```
upstream backend {
    server backend.example.com:8080 service=_example._tcp resolve;
    server 127.0.0.1:12345;
    zone backend 1m;
}
```

Allows the following peer names:

This API subsection enables setting the weight, max\_conns, max\_fails, fail\_timeout, backup and down parameters, as described in server.

# 1 Note

There is no separate drain option here; to enable drain, set down to the string value drain:

```
$ curl -X PUT -d \"drain\" \
http://127.0.0.1/config/stream/upstreams/backend/servers/backend.example.com/down
```

Example:

```
curl http://127.0.0.1/config/stream/upstreams/backend/servers/backend.example.com?

defaults=on
```

```
"weight": 1,
    "max_conns": 0,
    "max_fails": 1,
    "fail_timeout": 10,
    "backup": true,
```



```
"down": false,
}
```

Actually available parameters are limited to the ones supported by the current load balancing method of the *upstream*. So, if the upstream is configured as random:

```
upstream backend {
    zone backend 256k;
    server backend.example.com resolve max_conns=5;
    random;
}
```

You will be unable to add a new peer that defines backup:

```
$ curl -X PUT -d '{ "backup": true }' \
http://127.0.0.1/config/stream/upstreams/backend/servers/backend1.example.com
```

```
{
    "error": "FormatError",
    "description": "The \"backup\" field is unknown."
}
```

## 1 Note

Even with a compatible load balancing method, the backup parameter can only be set at new peer creation.

When deleting servers, you can set the connection\_drop=<value> argument (PRO) to override the proxy\_connection\_drop settings:

## **HTTP Methods**

Let's consider the semantics of all HTTP methods applicable to this section, given this upstream configuration:

```
http {
    # ...

upstream backend {
    zone upstream 256k;
    server backend.example.com resolve max_conns=5;
    # ...
}
```



```
server {
    # ...

    location /config/ {
        api /config/;

        allow 127.0.0.1;
        deny all;
    }
}
```

#### **GET**

The GET HTTP method queries an entity at any existing path within /config, just as it does for other API sections.

For example, the /config/http/upstreams/backend/servers/ upstream server branch enables these queries:

You can obtain default parameter values with defaults=on:

```
$ curl http://127.0.0.1/config/http/upstreams/backend/servers?defaults=on
```

```
{
    "backend.example.com": {
        "weight": 1,
        "max_conns": 5,
        "max_fails": 1,
        "fail_timeout": 10,
        "backup": false,
        "down": false,
        "sid": ""
    }
}
```

#### **PUT**

The PUT HTTP method creates a new JSON entity at the specified path or *entirely* replaces an existing one.

For example, to set the max\_fails parameter, not specified earlier, of the backend.example.com server within the backend upstream:

```
$ curl -X PUT -d '2' \
   http://127.0.0.1/config/http/upstreams/backend/servers/backend.example.com/max_
   →fails
```

```
{
    "success": "Updated",
```



```
"description": "Existing configuration API entity \"/config/http/upstreams/

⇒backend/servers/backend.example.com/max_fails\" was updated with replacing."
}
```

Verify the changes:

```
$ curl http://127.0.0.1/config/http/upstreams/backend/servers/backend.example.com
```

```
{
    "max_conns": 5,
    "max_fails": 2
}
```

#### DELETE

The DELETE HTTP method deletes previously defined settings at the specified path; at doing that, it returns to the default values if there are any.

For example, to delete the previously set max\_fails parameter of the backend.example.com server within the backend upstream:

```
$ curl -X DELETE \
http://127.0.0.1/config/http/upstreams/backend/servers/backend.example.com/max_
fails
```

```
{
    "success": "Reset",
    "description": "Configuration API entity \"/config/http/upstreams/backend/servers/
    ⇒backend.example.com/max_fails\" was reset to default."
}
```

Verify the changes using defaults=on:

```
$ curl http://127.0.0.1/config/http/upstreams/backend/servers/backend.example.com?

defaults=on
```

```
{
    "weight": 1,
    "max_conns": 5,
    "max_fails": 1,
    "fail_timeout": 10,
    "backup": false,
    "down": false,
    "sid": ""
}
```

The max\_fails setting is back to its default value.

When deleting servers, you can set the connection\_drop=<value> argument (PRO) to override the proxy\_connection\_drop, grpc\_connection\_drop, fastcgi\_connection\_drop, scgi\_connection\_drop, and uwsgi\_connection\_drop settings:

```
$ curl -X DELETE \
   http://127.0.0.1/config/http/upstreams/backend/servers/backend1.example.com?
   →connection_drop=off
$ curl -X DELETE \
   http://127.0.0.1/config/http/upstreams/backend/servers/backend2.example.com?
```



```
$ curl -X DELETE \
   http://127.0.0.1/config/http/upstreams/backend/servers/backend3.example.com?
   →connection_drop=1000
```

#### **PATCH**

The PATCH HTTP method creates a new entity at the specified path or partially replaces or complements an existing one (RFC 7386) by supplying a JSON definition in its payload.

The method operates as follows: if the entities from the new definition exist in the configuration, they are overwritten; otherwise, they are added.

For example, to change the down setting of the backend.example.com server within the backend upstream, leaving the rest intact:

```
$ curl -X PATCH -d '{ "down": true }' \
http://127.0.0.1/config/http/upstreams/backend/servers/backend.example.com
```

```
{
    "success": "Updated",
    "description": "Existing configuration API entity \"/config/http/upstreams/
    →backend/servers/backend.example.com\" was updated with merging."
}
```

Verify the changes:

```
$ curl http://127.0.0.1/config/http/upstreams/backend/servers/backend.example.com
```

```
{
    "max_conns": 5,
    "down": true
}
```

The JSON object supplied with the PATCH request was merged with the existing one instead of overwriting it, as would be the case with PUT.

The null values are a corner case; they are used to delete specific configuration items during such merge.

# 1 Note

This deletion is identical to DELETE; in particular, it reinstates the default values.

For example, to delete the down setting added earlier and simultaneously update max\_conns:

```
$ curl -X PATCH -d '{ "down": null, "max_conns": 6 }' \
   http://127.0.0.1/config/http/upstreams/backend/servers/backend.example.com
```

Verify the changes:



\$ curl http://127.0.0.1/config/http/upstreams/backend/servers/backend.example.com

```
{
    "max_conns": 6
}
```

The down parameter, for which a null was supplied, was deleted; max\_conns was updated.

#### **Auth Basic**

Allows limiting access to resources by validating the user name and password using the "HTTP Basic Authentication" protocol.

Access can also be limited by *address* or by the *result of subrequest*. Simultaneous limitation of access by address and by password is controlled by the *satisfy* directive.

### **Configuration Example**

#### **Directives**

### auth basic

| Syntax  | auth_basic string   off;             |
|---------|--------------------------------------|
| Default | <pre>auth_basic off;</pre>           |
| Context | http, server, location, limit_except |

Enables validation of user name and password using the "HTTP Basic Authentication" protocol. The specified parameter is used as a *realm*. Parameter value can contain variables.

| off | cancels the effect of the auth_basic directive inherited from the previous config- |
|-----|--|
|     | uration level  |

### auth basic user file

```
Syntax auth_basic_user_file file;
Default —
Context http, server, location, limit_except
```

Specifies a file that keeps user names and passwords, in the following format:

```
# comment
name1:password1
name2:password2:comment
name3:password3
```

The file name can contain variables.

The following password types are supported:



- encrypted with the crypt() function; can be generated using the htpasswd utility from the Apache HTTP Server distribution or the "openssl passwd" command;
- hashed with the Apache variant of the MD5-based password algorithm (apr1); can be generated with the same tools;
- specified by the "{scheme}data" syntax as described in RFC 2307; currently implemented schemes include PLAIN (an example one, should not be used), SHA (plain SHA-1 hashing, should not be used) and SSHA (salted SHA-1 hashing, used by some software packages, notably OpenLDAP and Dovecot).

## \* Caution

Support for SHA scheme was added only to aid in migration from other web servers. It should not be used for new passwords, since unsalted SHA-1 hashing that it employs is vulnerable to rainbow table attacks.

#### **Auth Request**

Implements client authorization based on the result of a subrequest. If the subrequest returns a 2xx response code, the access is allowed. If it returns 401 or 403, the access is denied with the corresponding error code. Any other response code returned by the subrequest is considered an error.

For the 401 error, the client also receives the "WWW-Authenticate" header from the subrequest response.

When building from the source code, this module isn't built by default; it should be enabled with the --with-http\_auth\_request\_module build option.

In packages and images from our repos, the module is included in the build.

The module may be combined with other access modules, such as *Access* and *Auth Basic*, via the *satisfy* directive.

#### **Configuration Example**

```
location /private/ {
    auth_request /auth;
# ...
}

location = /auth {
    proxy_pass ...
    proxy_pass_request_body off;
    proxy_set_header Content-Length "";
    proxy_set_header X-Original-URI $request_uri;
}
```

# Directives

#### auth request

```
Syntax auth_request uri | off;
Default auth_request off;
Context http, server, location
```

Enables authorization based on the result of a subrequest and sets the URI to which the subrequest will be sent.



## auth\_request\_set

| Syntax  | auth_request_set \$variable value; |
|---------|------------------------------------|
| Default | _                                  |
| Context | http, server, location             |

Sets the request variable to the given value after the authorization request completes. The value may contain variables from the authorization request, such as **\$upstream\_http\_\***.

#### **AutoIndex**

The module processes requests ending with the slash character (/) and produces a directory listing. Usually, a request is passed to the  $http\_autoindex$  module when the  $http\_index$  module cannot find an index file.

## **Configuration Example**

```
location / {
   autoindex on;
}
```

#### **Directives**

#### autoindex

| Syntax  | autoindex on   off;    |
|---------|------------------------|
| Default | autoindex off;         |
| Context | http, server, location |

Enables or disables the directory listing output.

#### autoindex exact size

| Syntax  | autoindex_exact_size on   off;      |
|---------|-------------------------------------|
| Default | <pre>autoindex_exact_size on;</pre> |
| Context | http, server, location              |

For the HTML *format*, specifies whether exact file sizes should be output in the directory listing, or rather rounded to kilobytes, megabytes, and gigabytes.

## autoindex\_format

| Syntax  | autoindex_format html   xml   json   jsonp; |
|---------|---|
| Default | <pre>autoindex_format html;</pre>           |
| Context | http, server, location                      |

Sets the format of a directory listing.

When the JSONP format is used, the name of a callback function is set with the *callback* request argument. If the argument is missing or has an empty value, then the JSON format is used.

The XML output can be transformed using the http xslt module.



## autoindex\_localtime

| Syntax  | <pre>autoindex_localtime on   off;</pre> |
|---------|--|
| Default | <pre>autoindex_localtime off;</pre>      |
| Context | http, server, location                   |

For the HTML format, specifies whether times in the directory listing should be output in the local time zone or UTC.

#### **Browser**

The module creates variables whose values depend on the value of the "User-Agent" request header field.

#### **Variables**

#### \$modern\_browser

equals the value set by the modern browser value directive, if a browser was identified as modern;

#### \$ancient\_browser

equals the value set by the ancient browser value directive, if a browser was identified as ancient;

#### \$msie

equals "1" if a browser was identified as MSIE of any version.

## **Configuration Example**

## Choosing an index file:

```
modern_browser_value "modern.";
modern_browser msie 5.5;
modern_browser gecko 1.0.0;
modern_browser opera 9.0;
modern_browser safari 413;
modern_browser konqueror 3.0;
index index.${modern_browser}html index.html;
```

#### Redirection for old browsers:

```
modern_browser msie 5.0;
modern_browser gecko 0.9.1;
modern_browser opera 8.0;
modern_browser safari 413;
modern_browser konqueror 3.0;

modern_browser unlisted;
ancient_browser Links Lynx netscape4;
if ($ancient_browser) {
   rewrite ^ /ancient.html;
}
```



#### **Directives**

### ancient\_browser

| Syntax  | ancient_browser $string \;$ |
|---------|-----------------------------|
| Default | _                           |
| Context | http, server, location      |

If any of the specified substrings is found in the "User-Agent" request header field, the browser will be considered ancient. The special string "netscape4" corresponds to the regular expression " $^Mozilla/[1-4]$ ".

### ancient browser value

| Syntax  | ancient_browser_value $string;$     |
|---------|-------------------------------------|
| Default | <pre>ancient_browser_value 1;</pre> |
| Context | http, server, location              |

Sets a value for the *\$ancient browser* variables.

### modern browser

| Syntax  | modern_browser browser version; modern_browser unlisted; |
|---------|--|
| Default | _  |
| Context | http, server, location                                   |

Specifies a version starting from which a browser is considered modern. A browser can be any one of the following: msie, gecko (browsers based on Mozilla), opera, safari, or konqueror.

Versions can be specified in the following formats: X, X.X, X.X.X, or X.X.X.X. The maximum values for each of the format are 4000, 4000.99, 4000.99.99, and 4000.99.99, respectively.

The special value unlisted specifies to consider a browser as modern if it was not listed by the modern\_browser and ancient\_browser directives. Otherwise such a browser is considered ancient. If a request does not provide the "User-Agent" field in the header, the browser is treated as not being listed.

## modern\_browser\_value

| Syntax  | $modern\_browser\_value\ string;$  |
|---------|------------------------------------|
| Default | <pre>modern_browser_value 1;</pre> |
| Context | http, server, location             |

Sets a value for the \$modern browser variables.

#### Charset

The module adds the specified charset to the "Content-Type" response header field. In addition, the module can convert data from one charset to another, with some limitations:

- conversion is performed one way from server to client,
- only single-byte charsets can be converted
- or single-byte charsets to/from UTF-8.



#### **Configuration Example**

```
include conf/koi-win;
charset windows-1251;
source_charset koi8-r;
```

#### **Directives**

#### charset

| Syntax  | charset charset   off;                 |
|---------|--|
| Default | charset off;                           |
| Context | http, server, location, if in location |

Adds the specified charset to the "Content-Type" response header field. If this charset is different from the charset specified in the *source charset* directive, a conversion is performed.

The parameter off cancels the addition of charset to the "Content-Type" response header field.

A charset can be defined with a variable:

```
charset $charset;
```

In such a case, all possible values of a variable need to be present in the configuration at least once in the form of the <code>charset\_map</code>, <code>charset</code>, or <code>source\_charset</code> directives. For <code>utf-8</code>, <code>windows-1251</code>, and <code>koi8-r</code> charsets, it is sufficient to include the files <code>conf/koi-win</code>, <code>conf/koi-utf</code>, and <code>conf/win-utf</code> into configuration. For other charsets, simply making a fictitious conversion table works, for example:

```
charset_map iso-8859-5 _ { }
```

In addition, a charset can be set in the "X-Accel-Charset" response header field. This capability can be disabled using the <code>proxy\_ignore\_headers</code>, <code>fastcgi\_ignore\_headers</code>, <code>uwsgi\_ignore\_headers</code>, <code>scgi\_ignore\_headers</code>, and <code>grpc\_ignore\_headers</code> directives.

#### charset map

```
Syntax charset_map charset1 charset2 { ... }

Default —

Context http
```

Describes the conversion table from one charset to another. A reverse conversion table is built using the same data. Character codes are given in hexadecimal. Missing characters in the range 80-FF are replaced with "?". When converting from UTF-8, characters missing in a one-byte charset are replaced with " $\mathcal{E}\#XXXX$ ;".

Example:

```
charset_map koi8-r windows-1251 {
    C0 FE ; # small yu
    C1 E0 ; # small a
    C2 E1 ; # small b
    C3 F6 ; # small ts
}
```

When describing a conversion table to UTF-8, codes for the UTF-8 charset should be given in the second column, for example:



```
charset_map koi8-r utf-8 {
   CO D18E ; # small yu
   C1 D0BO ; # small a
   C2 D0B1 ; # small b
   C3 D186 ; # small ts
}
```

Full conversion tables from koi8-r to windows-1251, and from koi8-r and windows-1251 to utf-8 are provided in the distribution files conf/koi-win, conf/koi-utf, and conf/win-utf.

## charset\_types

| Syntax  | charset_types mime-type;                                     |
|---------|--|
| Default | charset_types text/html text/xml text/plain text/vnd.wap.wml |
|         | application/javascript application/rss+xml;                  |
| Context | http, server, location                                       |

Enables module processing in responses with the specified MIME types in addition to text/html. The special value \* matches any MIME type.

## override charset

| Syntax  | override_charset on   off;             |
|---------|--|
| Default | <pre>override_charset off;</pre>       |
| Context | http, server, location, if in location |

Determines whether a conversion should be performed for answers received from a proxied or a FastCGI/uwsgi/SCGI/gRPC server when the answers already carry a charset in the "Content-Type" response header field. If conversion is enabled, a charset specified in the received response is used as a source charset.

# 1 Note

If a response is received in a subrequest then the conversion from the response charset to the main request charset is always performed, regardless of the *override charset* directive setting.

### source\_charset

| Syntax  | source_charset encoding;               |
|---------|--|
| Default | _                                      |
| Context | http, server, location, if in location |

Defines the source charset of a response. If this charset is different from the charset specified in the *charset* directive, a conversion is performed.

#### **DAV**

The module is intended for file management automation via the WebDAV protocol. The module processes HTTP and WebDAV methods PUT, DELETE, MKCOL, COPY, and MOVE.

When building from the source code, this module isn't built by default; it should be enabled with the --with-http\_dav\_module build option.



In packages and images from our repos, the module is included in the build.

# Important

WebDAV clients that require additional WebDAV methods to operate will not work with this module.

## **Configuration Example**

#### **Directives**

# create\_full\_put\_path

```
Syntax create_full_put_path on | off;
Default create_full_put_path off;
Context http, server, location
```

The WebDAV specification only allows creating files in already existing directories. This directive allows creating all needed intermediate directories.

### dav\_access

| Syntax  | dav_access users:permissions;  |
|---------|--------------------------------|
| Default | <pre>dav_access user:rw;</pre> |
| Context | http, server, location         |

Sets access permissions for newly created files and directories, e.g.:

```
dav_access user:rw group:rw all:r;
```

If any group or all access permissions are specified then user permissions may be omitted:

```
dav_access group:rw all:r;
```



## dav\_methods

| Syntax  | ${\tt dav\_methods} \ {\tt off} \   \ method \;$ |
|---------|--|
| Default | <pre>dav_methods off;</pre>                      |
| Context | http, server, location                           |

Allows the specified HTTP and WebDAV methods. The parameter off denies all methods processed by this module. The following methods are supported: PUT, DELETE, MKCOL, COPY, and MOVE.

A file uploaded with the PUT method is first written to a temporary file, and then the file is renamed. Starting from version 0.8.9, temporary files and the persistent store can be put on different file systems. However, be aware that in this case a file is copied across two file systems instead of the cheap renaming operation. It is thus recommended that for any given location both saved files and a directory holding temporary files, set by the <code>client\_body\_temp\_path</code> directive, are put on the same file system.

When creating a file with the PUT method, it is possible to specify the modification date by passing it in the "Date" header field.

## min\_delete\_depth

| Syntax  | min_delete_depth number;       |
|---------|--------------------------------|
| Default | <pre>min_delete_depth 0;</pre> |
| Context | http, server, location         |

Allows the DELETE method to remove files provided that the number of elements in a request path is not less than the specified number. For example, the directive

```
min_delete_depth 4;
```

allows removing files on requests

```
/users/00/00/name
/users/00/00/name/pic.jpg
/users/00/00/page.html
```

and denies the removal of

```
/users/00/00
```

## **Empty GIF**

The module emits a single-pixel transparent GIF.

#### **Configuration Example**

```
location = /_.gif {
    empty_gif;
}
```

#### **Directives**

# empty\_gif

| Syntax  | empty_gif; |
|---------|------------|
| Default | _          |
| Context | location   |



Turns on module processing in the surrounding location.

#### **FastCGI**

The module allows passing requests to a FastCGI server.

### **Configuration Example**

#### **Directives**

### fastcgi bind

| Syntax  | fastcgi_bind address [transparent]   off; |
|---------|---|
| Default | _   |
| Context | http, server, location                    |

Makes outgoing connections to a FastCGI server originate from the specified local IP address with an optional port. Parameter value can contain variables. The special value off cancels the effect of the fastcgi\_bind directive inherited from the previous configuration level, which allows the system to auto-assign the local IP address and port.

The transparent parameter allows outgoing connections to a FastCGI server originate from a non-local IP address, for example, from a real IP address of a client:

```
fastcgi_bind $remote_addr transparent;
```

For this parameter to work, Angie worker processes usually need to run with superuser privileges. On Linux, this is not required: if the transparent parameter is specified, worker processes inherit the  $CAP\ NET\ RAW$  capability from the master process.

# Important

The kernel routing table should also be configured to intercept network traffic from the FastCGI server.

### fastcgi\_buffer\_size

| Syntax  | fastcgi_buffer_size $size;$           |
|---------|---------------------------------------|
| Default | <pre>fastcgi_buffer_size 4k 8k;</pre> |
| Context | http, server, location                |

Sets the size of the buffer used for reading the first part of the response received from the FastCGI server. This part usually contains a small response header. By default, the buffer size is equal to one memory page. This is either 4K or 8K, depending on a platform. It can be made smaller, however.



## fastcgi\_buffering

| Syntax  | fastcgi_buffering on   off; |
|---------|-----------------------------|
| Default | fastcgi_buffering on;       |
| Context | http, server, location      |

Enables or disables buffering of responses from the FastCGI server.

| off | Angie receives a response from the FastCGI server as soon as possible, saving       |
|-----|---|
|     | it into the buffers set by the fastcgi_buffer_size and fastcgi_buffers directives.  |
|     | If the whole response does not fit into memory, a part of it can be saved to        |
|     | a temporary file on the disk. Writing to temporary files is controlled by the       |
|     | $fastcgi\_max\_temp\_file\_size$ and $fastcgi\_temp\_file\_write\_size$ directives. |
| on  | the response is passed to a client synchronously, immediately as it is received.    |
|     | Angie will not try to read the whole response from the FastCGI server. The          |
|     | maximum size of the data that Angie can receive from the server at a time is set    |
|     | by the $fastcgi\_buffer\_size$ directive.   |

Buffering can also be enabled or disabled by passing "yes" or "no" in the "X-Accel-Buffering" response header field. This capability can be disabled using the *fastcgi ignore headers* directive.

### fastcgi buffers

| Syntax  | fastcgi_buffers number size; |
|---------|------------------------------|
| Default | fastcgi_buffers 8 4k 8k;     |
| Context | http, server, location       |

Sets the number and size of the buffers used for reading a response from the FastCGI server, for a single connection.

By default, the buffer size is equal to one memory page. This is either 4K or 8K, depending on a platform.

### fastcgi busy buffers size

| Syntax  | fastcgi_busy_buffers_size $size;$              |
|---------|--|
| Default | <pre>fastcgi_busy_buffers_size 8k   16k;</pre> |
| Context | http, server, location                         |

When buffering of responses from the FastCGI server is enabled, limits the total size of buffers that can be busy sending a response to the client while the response is not yet fully read. In the meantime, the rest of the buffers can be used for reading the response and, if needed, buffering part of the response to a temporary file. By default, size is limited by the size of two buffers set by the fastcgi\_buffer\_size and fastcgi\_buffers directives.

## fastcgi\_cache

| Syntax  | fastcgi_cache zone   off;     |
|---------|-------------------------------|
| Default | <pre>fastcgi_cache off;</pre> |
| Context | http, server, location        |

Defines a shared memory zone used for caching. The same zone can be used in several places. Parameter value can contain variables. The off parameter disables caching inherited from the previous configuration level.



### fastcgi\_cache\_background\_update

| Syntax  | <pre>fastcgi_cache_background_update on   off;</pre> |
|---------|--|
| Default | <pre>fastcgi_cache_background_update off;</pre>      |
| Context | http, server, location                               |

Allows starting a background subrequest to update an expired cache item, while a stale cached response is returned to the client. Note that it is necessary to allow the usage of a stale cached response when it is being updated.

### fastcgi cache bypass

| Syntax  | fastcgi_cache_bypass $string;$ |
|---------|--------------------------------|
| Default | _                              |
| Context | http, server, location         |

Defines conditions under which the response will not be taken from a cache. If at least one value of the string parameters is not empty and is not equal to "0" then the response will not be taken from the cache:

```
fastcgi_cache_bypass $cookie_nocache $arg_nocache$arg_comment;
fastcgi_cache_bypass $http_pragma $http_authorization;
```

Can be used along with the fastcgi no cache directive.

#### fastcgi cache key

| Syntax  | fastcgi_cache_key $string;$ |
|---------|-----------------------------|
| Default | _                           |
| Context | http, server, location      |

Defines a key for caching, for example

```
fastcgi_cache_key localhost:9000$request_uri;
```

### fastcgi cache lock

| Syntax  | fastcgi_cache_lock on   off;       |
|---------|------------------------------------|
| Default | <pre>fastcgi_cache_lock off;</pre> |
| Context | http, server, location             |

When enabled, only one request at a time will be allowed to populate a new cache element identified according to the  $fastcgi\_cache\_key$  directive by passing a request to a FastCGI server. Other requests of the same cache element will either wait for a response to appear in the cache or the cache lock for this element to be released, up to the time set by the  $fastcgi\_cache\_lock\_timeout$  directive.

### fastcgi\_cache\_lock\_age

| Syntax  | <pre>fastcgi_cache_lock_age time;</pre> |
|---------|---|
| Default | <pre>fastcgi_cache_lock_age 5s;</pre>   |
| Context | http, server, location                  |



If the last request sent to the FastCGI server to fill a new cache entry has not completed in the specified time, another request may be sent to the FastCGI server.

# fastcgi\_cache\_lock\_timeout

| Syntax  | <pre>fastcgi_cache_lock_timeout time;</pre> |
|---------|---|
| Default | <pre>fastcgi_cache_lock_timeout 5s;</pre>   |
| Context | http, server, location                      |

Sets a timeout for *fastcgi\_cache\_lock*. When the time expires, the request will be passed to the FastCGI server, however, the response will not be cached.

# fastcgi\_cache\_max\_range\_offset

| Syntax  | $fastcgi\_cache\_max\_range\_offset \ number;$ |
|---------|--|
| Default | _  |
| Context | http, server, location                         |

Sets an offset in bytes for byte-range requests. If the range is beyond the offset, the range request will be passed to the FastCGI server and the response will not be cached.

# fastcgi cache methods

| Syntax  | fastcgi_cache_methods GET   HEAD   POST;   |
|---------|--|
| Default | <pre>fastcgi_cache_methods GET HEAD;</pre> |
| Context | http, server, location                     |

If the client request method is listed in this directive then the response will be cached. "GET" and "HEAD" methods are always added to the list, though it is recommended to specify them explicitly. See also the *fastcgi no cache* directive.

### fastcgi\_cache\_min\_uses

| Syntax  | fastcgi_cache_min_uses $number;$     |
|---------|--------------------------------------|
| Default | <pre>fastcgi_cache_min_uses 1;</pre> |
| Context | http, server, location               |

Sets the number of requests after which the response will be cached.

#### fastcgi cache path

| Syntax  | $ \begin{array}{cccccccccccccccccccccccccccccccccccc$ |
|---------|---|
| Default | _   |
| Context | http, server, location                                |

Sets the path and other parameters of a cache. Cache data are stored in files. Both the key and file name in a cache are a result of applying the MD5 function to the proxied URL.



The levels parameter defines hierarchy levels of a cache: from 1 to 3, each level accepts values 1 or 2. For example, in the following configuration

```
fastcgi_cache_path /data/angie/cache levels=1:2 keys_zone=one:10m;
```

file names in a cache will look like this:

```
/data/angie/cache/c/29/b7f54b2df7773722d382f4809d65029c
```

A cached response is first written to a temporary file, and then the file is renamed. Temporary files and the cache can be put on different file systems. However, be aware that in this case a file is copied across two file systems instead of the cheap renaming operation. It is thus recommended that for any given location both cache and a directory holding temporary files are put on the same file system.

A directory for temporary files is set based on the use\_temp\_path parameter.

| on  | If this parameter is omitted or set to the value on, the directory set by the fastcgi_temp_path directive for the given location will be used. |
|-----|--|
| off | temporary files will be put directly in the cache directory.   |

In addition, all active keys and information about data are stored in a shared memory zone, whose name and size are configured by the keys\_zone parameter. One megabyte zone can store about 8 thousand keys.

Cached data that are not accessed during the time specified by the inactive parameter get removed from the cache regardless of their freshness.

By default, inactive is set to 10 minutes.

A special **cache manager** process monitors the maximum cache size, and the minimum amount of free space on the file system with cache. When the size is exceeded or there is not enough free space, it removes the least recently used data. The data is removed in iterations.

| max_size         | maximum cache size  |
|------------------|---|
| min_free         | minimum amount of free space on the file system with cache                        |
| manager_files    | limits the number of items to be deleted during one iteration<br>By default, 100. |
| manager_threshol | limits the duration of one iteration By default, 200 milliseconds                 |
| manager_sleep    | configures a pause between interactions By default, 50 milliseconds               |

A minute after Angie starts, the special **cache loader** process is activated. It loads information about previously cached data stored on file system into a cache zone. The loading is also done in iterations.

| loader_files     | limits the number of items to load during one iteration<br>By default, 100 |
|------------------|--|
| loader_threshold | limits the duration of one iteration By default, 200 milliseconds          |
| loader_sleep     | configures a pause between interactions By default, 50 milliseconds        |

### fastcgi cache revalidate

| Syntax  | fastcgi_cache_revalidate on   off;       |
|---------|--|
| Default | <pre>fastcgi_cache_revalidate off;</pre> |
| Context | http, server, location                   |



Enables revalidation of expired cache items using conditional requests with the "If-Modified-Since" and "If-None-Match" header fields.

#### fastcgi cache use stale

| Syntax  | fastcgi_cache_use_stale error   timeout   invalid_header   updating |
|---------|---|
|         | http_500   http_503   http_403   http_429   off;                    |
| Default | <pre>fastcgi_cache_use_stale off;</pre>                             |
| Context | http, server, location  |

Determines in which cases a stale cached response can be used when an error occurs during communication with the FastCGI server. The directive's parameters match the parameters of the fastcgi next upstream directive.

| error    | permits using a stale cached response if a FastCGI server to process a request cannot be selected.  |
|----------|---|
| updating | additional parameter, permits using a stale cached response if it is currently being updated. This allows minimizing the number of accesses to FastCGI servers when updating cached data. |

Using a stale cached response can also be enabled directly in the response header for a specified number of seconds after the response became stale.

- The stale-while-revalidate extension of the "Cache-Control" header field permits using a stale cached response if it is currently being updated.
- The stale-if-error extension of the "Cache-Control" header field permits using a stale cached response in case of an error.

# 1 Note

This has lower priority than using the directive parameters.

To minimize the number of accesses to FastCGI servers when populating a new cache element, the <code>fastcgi\_cache\_lock</code> directive can be used.

# fastcgi cache valid

| Syntax  | fastcgi_cache_valid [code] time; |
|---------|----------------------------------|
| Default | _                                |
| Context | http, server, location           |

Sets caching time for different response codes. For example, the following directives

```
fastcgi_cache_valid 200 302 10m;
fastcgi_cache_valid 404 1m;
```

set 10 minutes of caching for responses with codes 200 and 302 and 1 minute for responses with code 404.

If only caching time is specified

```
fastcgi_cache_valid 5m;
```



then only 200, 301, and 302 responses are cached.

In addition, the any parameter can be specified to cache any responses:

```
fastcgi_cache_valid 200 302 10m;
fastcgi_cache_valid 301 1h;
fastcgi_cache_valid any 1m;
```

# 1 Note

Parameters of caching can also be set directly in the response header. This has higher priority than setting of caching time using the directive.

- The "X-Accel-Expires" header field sets caching time of a response in seconds. The zero value disables caching for a response. If the value starts with the @ prefix, it sets an absolute time in seconds since Epoch, up to which the response may be cached.
- If the header does not include the "X-Accel-Expires" field, parameters of caching may be set in the header fields "Expires" or "Cache-Control".
- If the header includes the "Set-Cookie" field, such a response will not be cached.
- If the header includes the "Vary" field with the special value "\*", such a response will not be cached. If the header includes the "Vary" field with another value, such a response will be cached taking into account the corresponding request header fields.

Processing of one or more of these response header fields can be disabled using the  $fastcgi\_ignore\_headers$  directive.

# fastcgi catch stderr

| Syntax  | fastcgi_catch_stderr $string;$ |
|---------|--------------------------------|
| Default | _                              |
| Context | http, server, location         |

Sets a string to search for in the error stream of a response received from a FastCGI server. If the string is found then it is considered that the FastCGI server has returned an *invalid response*. This allows handling application errors in Angie, for example:

```
location /php/ {
    fastcgi_pass backend:9000;
    ...
    fastcgi_catch_stderr "PHP Fatal error";
    fastcgi_next_upstream error timeout invalid_header;
}
```

#### fastcgi\_connect\_timeout

| Syntax  | fastcgi_connect_timeout $time;$         |
|---------|---|
| Default | <pre>fastcgi_connect_timeout 60s;</pre> |
| Context | http, server, location                  |

Defines a timeout for establishing a connection with a FastCGI server. It should be noted that this timeout cannot usually exceed 75 seconds.



### fastcgi\_connection\_drop

| Syntax  | ${\tt fastcgi\_connection\_drop}\ time\  \ {\tt on}\  \ {\tt off};$ |
|---------|---|
| Default | <pre>fastcgi_connection_drop off;</pre>                             |
| Context | http, server, location  |

Enables termination of all connections to the proxied server after it has been removed from the group or marked as permanently unavailable by a reresolve process or the API command DELETE.

A connection is terminated when the next read or write event is processed for either the client or the proxied server.

Setting *time* enables a connection termination timeout; with on set, connections are dropped immediately.

## fastcgi\_force\_ranges

| Syntax  | fastcgi_force_ranges on   off;       |
|---------|--------------------------------------|
| Default | <pre>fastcgi_force_ranges off;</pre> |
| Context | http, server, location               |

Enables byte-range support for both cached and uncached responses from the FastCGI server regardless of the "Accept-Ranges" field in these responses.

#### fastcgi hide header

| Syntax  | fastcgi_hide_header field; |
|---------|----------------------------|
| Default | _                          |
| Context | http, server, location     |

By default, Angie does not pass the header fields "Status" and "X-Accel-..." from the response of a FastCGI server to a client. The <code>fastcgi\_hide\_header</code> directive sets additional fields that will not be passed. If, on the contrary, the passing of fields needs to be permitted, the <code>fastcgi\_pass\_header</code> directive can be used.

# fastcgi\_ignore\_client\_abort

| Syntax  | fastcgi_ignore_client_abort on   off;       |
|---------|---|
| Default | <pre>fastcgi_ignore_client_abort off;</pre> |
| Context | http, server, location                      |

Determines whether the connection with a FastCGI server should be closed when a client closes the connection without waiting for a response.

### fastcgi ignore headers

| Syntax  | ${	t fastcgi\_ignore\_headers} \ field;$ |
|---------|--|
| Default | _  |
| Context | http, server, location                   |

Disables processing of certain response header fields from the FastCGI server. The following fields can be ignored: "X-Accel-Redirect", "X-Accel-Expires", "X-Accel-Limit-Rate", "X-Accel-Buffering", "X-Accel-Expires", "Expires", "Cache-Control", "Set-Cookie", and "Vary".



If not disabled, processing of these header fields has the following effect:

- "X-Accel-Expires", "Expires", "Cache-Control", "Set-Cookie", and "Vary" set the *parameters* of response caching;
- "X-Accel-Redirect" performs an *internal* redirect to the specified URI;
- $\bullet$  "X-Accel-Limit-Rate" sets the  $\it rate\ limit$  for transmission of a response to a client;
- "X-Accel-Buffering" enables or disables buffering of a response;
- "X-Accel-Charset" sets the desired *charset* of a response.

### fastcgi index

| Syntax  | $fastcgi\_index \ name;$ |
|---------|--------------------------|
| Default | _                        |
| Context | http, server, location   |

Sets a file name that will be appended after a URI that ends with a slash, in the value of the \$fastcgi\_script\_name variable. For example, with these settings

```
fastcgi_index index.php;
fastcgi_param SCRIPT_FILENAME /home/www/scripts/php$fastcgi_script_name;
```

and the /page.php request, the SCRIPT\_FILENAME parameter will be equal to /home/www/scripts/php/page.php,

and with the / request it will be equal to /home/www/scripts/php/index.php.

#### fastcgi intercept errors

| Syntax  | fastcgi_intercept_errors on   off;       |
|---------|--|
| Default | <pre>fastcgi_intercept_errors off;</pre> |
| Context | http, server, location                   |

Determines whether FastCGI server responses with codes greater than or equal to 300 should be passed to a client or be intercepted and redirected to Angie for processing with the *error page* directive.

### fastcgi keep conn

| Syntax  | fastcgi_keep_conn on   off;       |
|---------|-----------------------------------|
| Default | <pre>fastcgi_keep_conn off;</pre> |
| Context | http, server, location            |

By default, a FastCGI server will close a connection right after sending the response. However, when this directive is set to the value on, Angie will instruct a FastCGI server to keep connections open. This is necessary, in particular, for *keepalive* connections to FastCGI servers to function.

# fastcgi limit rate

| Syntax  | fastcgi_limit_rate rate;         |
|---------|----------------------------------|
| Default | <pre>fastcgi_limit_rate 0;</pre> |
| Context | http, server, location           |



Limits the speed of reading the response from the FastCGI server. The *rate* is specified in bytes per second and can contain variables.

0 disables rate limiting

# 1 Note

The limit is set per a request, and so if Angie simultaneously opens two connections to the FastCGI server, the overall rate will be twice as much as the specified limit. The limitation works only if buffering of responses from the FastCGI server is enabled.

### fastcgi max temp file size

| Syntax  | fastcgi_max_temp_file_size $size;$           |
|---------|--|
| Default | <pre>fastcgi_max_temp_file_size 1024m;</pre> |
| Context | http, server, location                       |

When buffering of responses from the FastCGI server is enabled, and the whole response does not fit into the buffers set by the fastcgi\_buffer\_size and fastcgi\_buffers directives, a part of the response can be saved to a temporary file. This directive sets the maximum size of the temporary file. The size of data written to the temporary file at a time is set by the fastcgi\_temp\_file\_write\_size directive.

0 disables buffering of responses to temporary files

#### 1 Note

This restriction does not apply to responses that will be *cached* or stored on *disk*.

### fastcgi\_next\_upstream

| Syntax  | fastcgi_next_upstream error   timeout   invalid_header   http_500   http_503   http_403   http_404   http_429   non_idempotent   off; |
|---------|---|
| Default | <pre>fastcgi_next_upstream error timeout;</pre>   |
| Context | http, server, location  |

Specifies in which cases a request should be passed to the next server:



| error          | an error occurred while establishing a connection with the server, passing a request to it, or reading the response header;  |
|----------------|--|
| timeout        | a timeout has occurred while establishing a connection with the server, passing a request to it, or reading the response header;   |
| invalid_header | a server returned an empty or invalid response;  |
| http_500       | a server returned a response with the code 500;  |
| http_503       | a server returned a response with the code 503;  |
| http_403       | a server returned a response with the code 403;  |
| http_404       | a server returned a response with the code 404;  |
| http_429       | a server returned a response with the code 429;  |
| non_idempotent | normally, requests with a non-idempotent method (POST, LOCK, PATCH) are not passed to the next server if a request has been sent to an upstream server; enabling this option explicitly allows retrying such requests; |
| off            | disables passing a request to the next server.   |

# 1 Note

One should bear in mind that passing a request to the next server is only possible if nothing has been sent to a client yet. That is, if an error or timeout occurs in the middle of the transferring of a response, fixing this is impossible.

The directive also defines what is considered an unsuccessful attempt of communication with a server.

| error<br>timeout<br>invalid_header | always considered unsuccessful attempts, even if they are not specified in the directive |
|------------------------------------|--|
| http_500<br>http_503<br>http_429   | considered unsuccessful attempts only if they are specified in the directive             |
| http_403<br>http_404               | never considered unsuccessful attempts   |

Passing a request to the next server can be limited by the *number of tries* and by time.

### fastcgi next upstream timeout

| Syntax  | ${\tt fastcgi\_next\_upstream\_timeout}\ time;$ |
|---------|---|
| Default | <pre>fastcgi_next_upstream_timeout 0;</pre>     |
| Context | http, server, location                          |

Limits the time during which a request can be passed to the *next server*.

| 0 turns off this limitation |
|-----------------------------|
|-----------------------------|

## fastcgi\_next\_upstream\_tries

| Syntax  | fastcgi_next_upstream_tries number;       |
|---------|---|
| Default | <pre>fastcgi_next_upstream_tries 0;</pre> |
| Context | http, server, location                    |



Limits the number of possible tries for passing a request to the *next server*.

| 0 | turns off this limitation |  |
|---|---------------------------|--|
|---|---------------------------|--|

## fastcgi\_no\_cache

| Syntax  | fastcgi_no_cache $string;$ |
|---------|----------------------------|
| Default | _                          |
| Context | http, server, location     |

Defines conditions under which the response will not be saved to a cache. If at least one value of the string parameters is not empty and is not equal to "0" then the response will not be saved:

```
fastcgi_no_cache $cookie_nocache $arg_nocache$arg_comment;
fastcgi_no_cache $http_pragma $http_authorization;
```

Can be used along with the fastcgi cache bypass directive.

## fastcgi param

| Syntax  | <pre>fastcgi_param parameter value [if_not_empty];</pre> |
|---------|--|
| Default | _  |
| Context | http, server, location                                   |

Sets a parameter that should be passed to the FastCGI server. The value can contain text, variables, and their combination. These directives are inherited from the previous configuration level if and only if there are no <code>fastcgi\_param</code> directives defined on the current level.

The following example shows the minimum required settings for PHP:

The SCRIPT\_FILENAME parameter is used in PHP for determining the script name, and the QUERY STRING parameter is used to pass request parameters.

For scripts that process POST requests, the following three parameters are also required:

If PHP was built with the --enable-force-cgi-redirect configuration parameter, the REDIRECT\_STATUS parameter should also be passed with the value "200":

```
fastcgi_param REDIRECT_STATUS 200;
```

If the directive is specified with if\_not\_empty then such a parameter will be passed to the server only if its value is not empty:



### fastcgi\_pass

| Syntax  | <pre>fastcgi_pass address;</pre> |
|---------|----------------------------------|
| Default | _                                |
| Context | location, if in location         |

Sets the address of a FastCGI server. The address can be specified as a domain name or IP address, and a port:

```
fastcgi_pass localhost:9000;
```

or as a UNIX domain socket path:

```
fastcgi_pass unix:/tmp/fastcgi.socket;
```

If a domain name resolves to several addresses, all of them will be used in a round-robin fashion. In addition, an address can be specified as a *server group*. If a group is used, you cannot specify the port with it; instead, specify the port for each server within the group individually.

Parameter value can contain variables. In this case, if an address is specified as a domain name, the name is searched among the described *server groups*, and, if not found, is determined using a *resolver*.

### fastcgi pass header

| Syntax  | ${	t fastcgi\_pass\_header} \ field;$ |
|---------|---------------------------------------|
| Default | _                                     |
| Context | http, server, location                |

Permits passing otherwise disabled header fields from a FastCGI server to a client.

# fastcgi pass request body

| Syntax  | fastcgi_pass_request_body on   off; |
|---------|-------------------------------------|
| Default | _                                   |
| Context | http, server, location              |

Indicates whether the original request body is passed to the FastCGI server. See also the fastcgi\_pass\_request\_headers directive.

# fastcgi\_pass\_request\_headers

| Syntax  | fastcgi_pass_request_headers on   off; |
|---------|--|
| Default | _                                      |
| Context | http, server, location                 |

Indicates whether the header fields of the original request are passed to the FastCGI server. See also the  $fastcgi\_pass\_request\_body$  directive.

### fastcgi read timeout

| Syntax  | ${\tt fastcgi\_read\_timeout}\ time;$ |
|---------|---------------------------------------|
| Default | <pre>fastcgi_read_timeout 60s;</pre>  |
| Context | http, server, location                |



Defines a timeout for reading a response from the FastCGI server. The timeout is set only between two successive read operations, not for the transmission of the whole response. If the FastCGI server does not transmit anything within this time, the connection is closed.

# fastcgi request buffering

| Syntax  | <pre>fastcgi_request_buffering on   off;</pre> |
|---------|--|
| Default | <pre>fastcgi_request_buffering on;</pre>       |
| Context | http, server, location                         |

Enables or disables buffering of a client request body.

| on  | the entire request body is $read$ from the client before sending the request to a FastCGI server.   |
|-----|---|
| off | the request body is sent to the FastCGI server immediately as it is received. In this case, the request cannot be passed to the <i>next server</i> , if Angie already started sending the request body. |

### fastcgi send lowat

| Syntax  | ${\tt fastcgi\_send\_lowat}$ $size;$ |
|---------|--------------------------------------|
| Default | <pre>fastcgi_send_lowat 0;</pre>     |
| Context | http, server, location               |

If the directive is set to a non-zero value, Angie will try to minimize the number of send operations on outgoing connections to a FastCGI server by using either  $NOTE\_LOWAT$  flag of the kqueue method, or the  $SO\_SNDLOWAT$  socket option, with the specified size.



### fastcgi send timeout

| Syntax  | fastcgi_send_timeout $time$ ;        |
|---------|--------------------------------------|
| Default | <pre>fastcgi_send_timeout 60s;</pre> |
| Context | http, server, location               |

Sets a timeout for transmitting a request to the FastCGI server. The timeout is set only between two successive write operations, not for the transmission of the whole request. If the FastCGI server does not receive anything within this time, the connection is closed.

### fastcgi socket keepalive

| Syntax  | fastcgi_socket_keepalive on   off;       |
|---------|--|
| Default | <pre>fastcgi_socket_keepalive off;</pre> |
| Context | http, server, location                   |

Configures the "TCP keepalive" behavior for outgoing connections to a FastCGI server.



| 11 11 | By default, the operating system's settings are in effect for the socket. |
|-------|---|
| on    | the SO_KEEPALIVE socket option is turned on for the socket.               |

### fastcgi\_split\_path\_info

| Syntax  | fastcgi_split_path_info regex; |
|---------|--------------------------------|
| Default | _                              |
| Context | location                       |

Defines a regular expression that captures a value for the  $fastcgi\_path\_info$  variable. The regular expression should have two captures: the first becomes a value of the  $fastcgi\_script\_name$  variable, the second becomes a value of the  $fastcgi\_path$  info variable. For example, with these settings

and the /show.php/article/0001 request, the SCRIPT\_FILENAME parameter will be equal to /path/to/php/show.php, and the PATH\_INFO parameter will be equal to /article/0001.

### fastcgi store

| Syntax  | fastcgi_store on   off   string; |
|---------|----------------------------------|
| Default | <pre>fastcgi_store off;</pre>    |
| Context | http, server, location           |

Enables saving of files to a disk.

| on  | saves files with paths corresponding to the directives alias or root. |
|-----|---|
| off | disables saving of files  |

In addition, the file name can be set explicitly using the string with variables:

```
fastcgi_store /data/www$original_uri;
```

The modification time of files is set according to the received "Last-Modified" response header field. The response is first written to a temporary file, and then the file is renamed. Temporary files and the persistent store can be put on different file systems. However, be aware that in this case a file is copied across two file systems instead of the cheap renaming operation. It is thus recommended that for any given location both saved files and a directory holding temporary files, set by the  $fastcgi\_temp\_path$  directive, are put on the same file system.

This directive can be used to create local copies of static unchangeable files, e.g.:



```
fastcgi_pass backend:9000;
...

fastcgi_store on;
fastcgi_store_access user:rw group:rw all:r;
fastcgi_temp_path /data/temp;

alias /data/www/;
}
```

### fastcgi store access

| Syntax  | fastcgi_store_access users:permissions;  |
|---------|--|
| Default | <pre>fastcgi_store_access user:rw;</pre> |
| Context | http, server, location                   |

Sets access permissions for newly created files and directories, e.g.:

```
fastcgi_store_access user:rw group:rw all:r;
```

If any group or all access permissions are specified then user permissions may be omitted:

```
fastcgi_store_access group:rw all:r;
```

### fastcgi temp file write size

| Syntax  | fastcgi_temp_file_write_size size;              |
|---------|---|
| Default | <pre>fastcgi_temp_file_write_size 8k 16k;</pre> |
| Context | http, server, location                          |

Limits the size of data written to a temporary file at a time, when buffering of responses from the FastCGI server to temporary files is enabled. By default, size is limited by two buffers set by the fastcgi\_buffer\_size and fastcgi\_buffers directives. The maximum size of a temporary file is set by the fastcgi\_max\_temp\_file\_size directive.

### fastcgi\_temp\_path

| Syntax  | <pre>fastcgi_temp_path path [level1 [level2 [level3]]]`;</pre>  |
|---------|---|
| Default | <pre>fastcgi_temp_path fastcgi_temp; (the path depends on thehttp-fastcgi-temp-path build option)</pre> |
| Context | http, server, location  |

Defines a directory for storing temporary files with data received from FastCGI servers. Up to three-level subdirectory hierarchy can be used underneath the specified directory. For example, in the following configuration

```
fastcgi_temp_path /spool/angie/fastcgi_temp 1 2;
```

a temporary file might look like this:

```
/spool/angie/fastcgi_temp/7/45/00000123457
```

See also the use temp path parameter of the fastcgi cache path directive.



#### Parameters Passed to a FastCGI Server

HTTP request header fields are passed to a FastCGI server as parameters. In applications and scripts running as FastCGI servers, these parameters are usually made available as environment variables. For example, the "User-Agent" header field is passed as the HTTP\_USER\_AGENT parameter. In addition to HTTP request header fields, it is possible to pass arbitrary parameters using the <code>fastcgi\_param</code> directive.

#### **Built-in Variables**

The http\_fastcgi module supports built-in variables that can be used to set parameters using the fastcgi param directive:

#### \$fastcgi\_script\_name

request URI or, if a URI ends with a slash, request URI with an index file name configured by the <code>fastcgi\_index</code> directive appended to it. This variable can be used to set the SCRIPT\_FILENAME and PATH\_TRANSLATED parameters that determine the script name in PHP. For example, for the <code>/info/request</code> with the following directives

```
fastcgi_index index.php;
fastcgi_param SCRIPT_FILENAME /home/www/scripts/php$fastcgi_script_name;
```

the SCRIPT FILENAME parameter will be equal to /home/www/scripts/php/info/index.php.

When using the fastcgi\_split\_path\_info directive, the \$fastcgi\_script\_name variable equals the value of the first capture set by the directive.

```
$fastcgi_path_info
```

the value of the second capture set by the  $fastcgi\_split\_path\_info$  directive. This variable can be used to set the PATH\_INFO parameter.

#### **FLV**

The module provides pseudo-streaming server-side support for Flash Video (FLV) files.

It handles requests with the start argument in the request URI's query string specially, by sending back the contents of a file starting from the requested byte offset and with the prepended FLV header.

When building from the source code, this module isn't built by default; it should be enabled with the --with-http\_flv\_module build option.

In packages and images from our repos, the module is included in the build.

# **Configuration Example**

```
location ~ \.flv$ {
    flv;
}
```

#### **Directives**

#### fΙν

| Syntax  | flv;     |
|---------|----------|
| Default | _        |
| Context | location |



Turns on module processing in a surrounding location.

#### Geo

The module creates variables with values depending on the client IP address.

### **Configuration Example**

```
geo $geo {
    default 0;

    127.0.0.1 2;
    192.168.1.0/24 1;
    10.1.0.0/16 1;

    ::1 2;
    2001:0db8::/32 1;
}
```

#### **Directives**

#### geo

```
Syntax geo [$address] $variable { ... }

Default —

Context http
```

Describes the dependency of values of the specified variable on the client IP address. By default, the address is taken from the  $\$remote\_addr$  variable, but it can also be taken from another variable, for example:

```
geo $arg_remote_addr $geo {
    ...;
}
```

### 1 Note

Since variables are evaluated only when used, the mere existence of even a large number of declared geo variables does not cause any extra costs for request processing.

If the value of a variable does not represent a valid IP address then the "255.255.255.255" address is used.

Addresses are specified either as prefixes in CIDR notation (including individual addresses) or as ranges.

The following special parameters are also supported:



| delete          | deletes the specified network  |
|-----------------|--|
| default         | the value set to the variable if the client address does not match any of the specified addresses. When addresses are specified in CIDR notation, 0.0.0.0/0 and ::/0 can be used instead of default. When default is not specified, the default value will be an empty string  |
| include         | includes a file with addresses and values. There can be several inclusions.  |
| proxy           | defines trusted addresses. When a request comes from a trusted address, an address from the X-Forwarded-For request header field will be used instead. In contrast to the regular addresses, trusted addresses are checked sequentially.   |
| proxy_recursive | enables recursive address search. If recursive search is disabled then instead of the original client address that matches one of the trusted addresses, the last address sent in X-Forwarded-For will be used. If recursive search is enabled then instead of the original client address that matches one of the trusted addresses, the last non-trusted address sent in X-Forwarded-For will be used. |
| ranges          | indicates that addresses are specified as ranges. This parameter should be the first. To speed up loading of a geo base, addresses should be put in ascending order.   |

### Example:

```
geo $country {
    default
                   ZZ;
    include
                   conf/geo.conf;
                   127.0.0.0/16;
    delete
                   192.168.100.0/24;
    proxy
    proxy
                   2001:0db8::/32;
    127.0.0.0/24
                   US;
                   RU;
    127.0.0.1/32
    10.1.0.0/16
                   RU;
    192.168.1.0/24 UK;
}
```

The conf/geo.conf file could contain the following lines:

```
10.2.0.0/16 RU;
192.168.2.0/24 RU;
```

The value of the most specific match is used. For example, for the 127.0.0.1 address, the value RU will be chosen, not US.

Sample range description:



#### **GeoIP**

Creates variables with values depending on the client IP address, using the precompiled MaxMind databases or their counterparts.

When using the databases with IPv6 support, IPv4 addresses are looked up as IPv4-mapped IPv6 addresses.

When building from the source code, this module isn't built by default; it should be enabled with the --with-http\_geoip\_module build option.

# • Important

This module requires the MaxMind GeoIP database or a counterpart such as MaxMind GeoLite2.

#### **Configuration Example**

#### **Directives**

#### geoip\_country

| Syntax  | ${	t geoip\_country} \ file;$ |
|---------|-------------------------------|
| Default | _                             |
| Context | http                          |

Specifies a database used to determine the country depending on the client IP address. The following variables are available when using this database:

```
$geoip_country_c two-letter country code, for example, "RU", "US".
$geoip_country_c three-letter country code, for example, "RUS", "USA".
$geoip_country_n country name, for example, "Russian Federation", "United States".
```

## geoip\_city

| Syntax  | geoip_city file; |
|---------|------------------|
| Default | _                |
| Context | http             |

Specifies a database used to determine the country, region, and city depending on the client IP address. The following variables are available when using this database:



| <pre>\$geoip_city_cont</pre> | two-letter continent code, for example, "EU", "NA".  |
|------------------------------|--|
| \$geoip_city_coun            | two-letter country code, for example, "RU", "US".  |
| \$geoip_city_coun            | three-letter country code, for example, "RUS", "USA".  |
| \$geoip_city_coun            | country name, for example, "Russian Federation", "United States".  |
| <pre>\$geoip_dma_code</pre>  | DMA region code in US (also known as "metro code"), according to the geotargeting in Google AdWords API.                                 |
| \$geoip_latitude             | latitude.  |
| <pre>\$geoip_longitude</pre> | longitude.   |
| <pre>\$geoip_region</pre>    | two-symbol country region code (region, territory, state, province, federal land and the like), for example, "48", "DC".                 |
| \$geoip_region_na            | country region name (region, territory, state, province, federal land and the like), for example, "Moscow City", "District of Columbia". |
| <pre>\$geoip_city</pre>      | city name, for example, "Moscow", "Washington".  |
| \$geoip_postal_co            | postal code.   |

# geoip\_org

| Syntax  | geoip_org file; |
|---------|-----------------|
| Default | _               |
| Context | http            |

Specifies a database used to determine the organization depending on the client IP address. The following variable is available when using this database:

| \$geoip_org | organization name, for example, "The University of Melbourne". |  |
|-------------|--|--|
|-------------|--|--|

# geoip\_proxy

| Syntax  | geoip_proxy file; |
|---------|-------------------|
| Default | _                 |
| Context | http              |

Defines trusted addresses. When a request comes from a trusted address, an address from the X-Forwarded-For request header field will be used instead.

### geoip\_proxy\_recursive

| Syntax  | geoip_proxy_recursive on   off;       |
|---------|---------------------------------------|
| Default | <pre>geoip_proxy_recursive off;</pre> |
| Context | http                                  |

If recursive search is disabled then instead of the original client address that matches one of the trusted addresses, the last address sent in X-Forwarded-For will be used. If recursive search is enabled then instead of the original client address that matches one of the trusted addresses, the last non-trusted address sent in X-Forwarded-For will be used.



#### **gRPC**

Allows passing requests to a gRPC server. The module requires the HTTP/2.

#### **Configuration Example**

```
server {
    listen 9000;

    http2 on;

    location / {
        grpc_pass 127.0.0.1:9000;
    }
}
```

#### **Directives**

### grpc bind

| Syntax  | <pre>grpc_bind address [transparent]   off;</pre> |
|---------|---|
| Default | _   |
| Context | http, server, location                            |

Makes outgoing connections to a gRPC server originate from the specified local IP address with an optional port. Parameter value can contain variables. The special value off cancels the effect of the  $grpc\_bind$  directive inherited from the previous configuration level, which allows the system to autoassign the local IP address and port.

The transparent parameter allows outgoing connections to a gRPC server originate from a non-local IP address, for example, from a real IP address of a client:

```
grpc_bind $remote_addr transparent;
```

In order for this parameter to work, it is usually necessary to run Angie worker processes with the *superuser* privileges. On Linux it is not required as if the **transparent** parameter is specified, worker processes inherit the *CAP NET RAW* capability from the master process.

# Important

It is necessary to configure kernel routing table to intercept network traffic from the gRPC server.

# grpc\_buffer\_size

| Syntax  | <pre>grpc_buffer_size size;</pre>  |
|---------|------------------------------------|
| Default | <pre>grpc_buffer_size 4k 8k;</pre> |
| Context | http, server, location             |

Sets the size of the buffer used for reading the first part of the response received from the gRPC server. The response is passed to the client synchronously, as soon as it is received.



### grpc\_connect\_timeout

| Syntax  | <pre>grpc_connect_timeout time;</pre> |
|---------|---------------------------------------|
| Default | <pre>grpc_connect_timeout 60s;</pre>  |
| Context | http, server, location                |

Defines a timeout for establishing a connection with a gRPC server. It should be noted that this timeout cannot usually exceed 75 seconds.

### grpc connection drop

| Syntax  | <pre>grpc_connection_drop time   on   off;</pre> |
|---------|--|
| Default | <pre>grpc_connection_drop off;</pre>             |
| Context | http, server, location                           |

Enables termination of all connections to the proxied server after it has been removed from the group or marked as permanently unavailable by a *reresolve* process or the *API command* DELETE.

A connection is terminated when the next read or write event is processed for either the client or the proxied server.

Setting *time* enables a connection termination timeout; with on set, connections are dropped immediately.

#### grpc hide header

| Syntax  | ${	t grpc\_hide\_header} \ field;$ |
|---------|------------------------------------|
| Default | _                                  |
| Context | http, server, location             |

By default, Angie does not pass the header fields "Date", "Server", and "X-Accel-..." from the response of a gRPC server to a client. The  $grpc\_hide\_header$  directive sets additional fields that will not be passed. If, on the contrary, the passing of fields needs to be permitted, the  $grpc\_pass\_header$  directive can be used.

# grpc ignore headers

| Syntax  | grpc_ignore_headers field; |
|---------|----------------------------|
| Default | _                          |
| Context | http, server, location     |

Disables processing of certain response header fields from the gRPC server. The following fields can be ignored: "X-Accel-Redirect" and "X-Accel-Charset".

If not disabled, processing of these header fields has the following effect:

- "X-Accel-Redirect" performs an internal redirect to the specified URI;
- "X-Accel-Charset" sets the desired *charset* of a response.

### grpc intercept errors

| Syntax  | <pre>grpc_intercept_errors on   off;</pre> |
|---------|--|
| Default | <pre>grpc_intercept_errors off;</pre>      |
| Context | http, server, location                     |



Determines whether gRPC responses with codes greater than or equal to 300 should be passed to a client or be intercepted and redirected to Angie for processing with the *error* page directive.

### grpc next upstream

| Syntax  | <pre>grpc_next_upstream error   timeout   invalid_header   http_500   http_502  </pre> |
|---------|--|
|         | http_503   http_504   http_403   http_404   http_429   non_idempotent   off;           |
| Default | <pre>grpc_next_upstream error timeout;</pre>   |
| Context | http, server, location   |

Specifies in which cases a request should be passed to the next server in the upstream pool:

| error          | an error occurred while establishing a connection with the server, passing a request to it, or reading the response header;  |
|----------------|--|
| timeout        | a timeout has occurred while establishing a connection with the server, passing  |
| invalid_header | a request to it, or reading the response header;<br>a server returned an empty or invalid response;  |
| http_500       | a server returned a response with the code 500;  |
| http_502       | a server returned a response with the code 502;  |
| http_503       | a server returned a response with the code 503;  |
| http_504       | a server returned a response with the code 504;  |
| http_403       | a server returned a response with the code 403;  |
| http_404       | a server returned a response with the code 404;  |
| http_429       | a server returned a response with the code 429;  |
| non_idempotent | normally, requests with a non-idempotent method (POST, LOCK, PATCH) are not passed to the next server if a request has been sent to an upstream server; enabling this option explicitly allows retrying such requests; |
| off            | disables passing a request to the next server.   |

### 1 Note

One should bear in mind that passing a request to the next server is only possible if nothing has been sent to a client yet. That is, if an error or timeout occurs in the middle of the transferring of a response, fixing this is impossible.

The directive also defines what is considered an unsuccessful attempt of communication with a server.

| error, timeout, invalid_header                               | always considered unsuccessful attempts, even if they are not specified in the directive |
|--|--|
| http_500,<br>http_502,<br>http_503,<br>http_504,<br>http_429 | considered unsuccessful attempts only if they are specified in the directive             |
| http_403,<br>http_404  | never considered unsuccessful attempts   |

Passing a request to the next server can be limited by the number of tries and by time.



### grpc\_next\_upstream\_timeout

| Syntax  | ${\tt grpc\_next\_upstream\_timeout}\ time;$ |
|---------|--|
| Default | <pre>grpc_next_upstream_timeout 0;</pre>     |
| Context | http, server, location                       |

Limits the time during which a request can be passed to the *next* server.

### grpc next upstream tries

| Syntax  | ${\tt grpc\_next\_upstream\_tries}\ number;$ |
|---------|--|
| Default | <pre>grpc_next_upstream_tries 0;</pre>       |
| Context | http, server, location                       |

Limits the number of possible tries for passing a request to the *next* server.

| 0 turns off this limitation |  |
|-----------------------------|--|
|-----------------------------|--|

### grpc\_pass

| Syntax  | <pre>grpc_pass address;</pre> |
|---------|-------------------------------|
| Default | _                             |
| Context | location, if in location      |

Sets gRPC server address. The address can be specified as a domain name or IP address, and a port:

```
grpc_pass localhost:9000;
```

or as a UNIX domain socket path:

```
grpc_pass unix:/tmp/grpc.socket;
```

Alternatively, the grpc:// scheme can be used:

```
grpc_pass grpc://127.0.0.1:9000;
```

To use gRPC over SSL, the grpcs:// scheme should be used:

```
grpc_pass grpcs://127.0.0.1:443;
```

If a domain name resolves to several addresses, all of them will be used in a round-robin fashion. In addition, an address can be specified as a  $server\ group$ .

Parameter value can contain variables. In this case, if an address is specified as a domain name, the name is searched among the described server groups, and, if not found, is determined using a *resolver*.



### grpc\_pass\_header

| Syntax  | ${	t grpc\_pass\_header} \ field;$ |
|---------|------------------------------------|
| Default | _                                  |
| Context | http, server, location             |

Permits passing otherwise disabled header fields from a gRPC server to a client.

#### grpc read timeout

| Syntax  | <pre>grpc_read_timeout time;</pre> |
|---------|------------------------------------|
| Default | <pre>grpc_read_timeout 60s;</pre>  |
| Context | http, server, location             |

Defines a timeout for reading a response from the gRPC server. The timeout is set only between two successive read operations, not for the transmission of the whole response. If the gRPC server does not transmit anything within this time, the connection is closed.

#### grpc send timeout

| Syntax  | <pre>grpc_send_timeout time;</pre> |
|---------|------------------------------------|
| Default | <pre>grpc_send_timeout 60s;</pre>  |
| Context | http, server, location             |

Sets a timeout for transmitting a request to the gRPC server. The timeout is set only between two successive write operations, not for the transmission of the whole request. If the gRPC server does not receive anything within this time, the connection is closed.

### grpc set header

| Syntax  | <pre>grpc_set_header field value;</pre>                     |
|---------|---|
| Default | <pre>grpc_set_header Content-Length \$content_length;</pre> |
| Context | http, server, location                                      |

Allows redefining or appending fields to the request header *passed* to the gRPC server. The value can contain text, variables, and their combinations. These directives are inherited from the previous configuration level if and only if there are no grpc set header directives defined on the current level.

If the value of a header field is an empty string then this field will not be passed to a gRPC server:

```
grpc_set_header Accept-Encoding "";
```

### grpc\_socket\_keepalive

| Syntax  | <pre>grpc_socket_keepalive on   off;</pre> |
|---------|--|
| Default | <pre>grpc_socket_keepalive off;</pre>      |
| Context | http, server, location                     |

Configures the "TCP keepalive" behavior for outgoing connections to a gRPC server.



| 11 11 | By default, the operating system's settings are in effect for the socket. |
|-------|---|
| on    | The SO_KEEPALIVE socket option is turned on for the socket.               |

### grpc\_ssl\_certificate

| Syntax  | <pre>grpc_ssl_certificate file;</pre> |
|---------|---------------------------------------|
| Default | _                                     |
| Context | http, server, location                |

Specifies a file with the certificate in the PEM format used for authentication to a gRPC SSL server. Variables can be used in the file name.

### grpc ssl certificate cache

| Syntax  | <pre>grpc_ssl_certificate_cache off; grpc_ssl_certificate_cache max=N [inactive=time] [valid=time];</pre> |
|---------|---|
| Default | <pre>grpc_ssl_certificate_cache off;</pre>  |
| Context | http, server, location  |

Defines a cache that stores SSL certificates and secret keys specified using variables.

The directive supports the following parameters:

- max sets the maximum number of elements in the cache. When the cache overflows, the least recently used (LRU) elements are removed.
- inactive defines the time after which an element is removed if it has not been accessed. The default is 10 seconds.
- valid defines the time during which a cached element is considered valid and can be reused. The default is 60 seconds. After this period, certificates are reloaded or revalidated.
- off disables the cache.

#### Example:

# grpc\_ssl\_certificate\_key

| Syntax  | ${	t grpc\_ssl\_certificate\_key} \ file;$ |
|---------|--|
| Default | _  |
| Context | http, server, location                     |

Specifies a file with the secret key in the PEM format used for authentication to a gRPC SSL server.

The value "engine: name: id" can be specified instead of the file, which loads a secret key with a specified id from the OpenSSL engine name. Variables can be used in the file name.



### grpc\_ssl\_ciphers

| Syntax  | <pre>grpc_ssl_ciphers ciphers;</pre> |
|---------|--------------------------------------|
| Default | <pre>grpc_ssl_ciphers DEFAULT;</pre> |
| Context | http, server, location               |

Specifies the enabled ciphers for requests to a gRPC SSL server. The ciphers are specified in the format understood by the OpenSSL library.

The list of ciphers depends on the version of OpenSSL installed. The full list can be viewed using the openssl ciphers command.

#### **A** Attention

The grpc\_ssl\_ciphers directive does not configure ciphers for TLS 1.3 when using OpenSSL. To tune TLS 1.3 ciphers with OpenSSL, use the grpc\_ssl\_conf\_command directive, which was added to support advanced SSL configuration.

- In LibreSSL, TLS 1.3 ciphers can be configured using grpc\_ssl\_ciphers.
- In BoringSSL, TLS 1.3 ciphers cannot be configured at all.

## grpc\_ssl\_conf\_command

| Syntax  | <pre>grpc_ssl_conf_command name value;</pre> |
|---------|--|
| Default | _  |
| Context | http, server, location                       |

Sets arbitrary OpenSSL configuration commands when establishing a connection with the gRPC SSL server.

### Important

The directive is supported when using OpenSSL 1.0.2 or higher. To configure TLS 1.3 ciphers with OpenSSL, use the ciphersuites command.

Several  $grpc\_ssl\_conf\_command$  directives can be specified on the same level. These directives are inherited from the previous configuration level if and only if there are no  $grpc\_ssl\_conf\_command$  directives defined on the current level.

### \* Caution

Note that configuring OpenSSL directly might result in unexpected behavior.

## grpc\_ssl\_crl

| Syntax  | grpc_ssl_crl file;     |
|---------|------------------------|
| Default | _                      |
| Context | http, server, location |

Specifies a file with revoked certificates (CRL) in the PEM format used to *verify* the certificate of the gRPC SSL server.



### grpc\_ssl\_name

| Syntax  | <pre>grpc_ssl_name name;</pre>                  |
|---------|---|
| Default | <pre>grpc_ssl_name `host from grpc_pass;`</pre> |
| Context | http, server, location                          |

Allows overriding the server name used to verify the certificate of the gRPC SSL server and to be passed through SNI when establishing a connection with the gRPC SSL server.

By default, the host part of the grpc pass URL is used.

# grpc\_ssl\_password\_file

| Syntax  | ${	t grpc\_ssl\_password\_file}$ $file;$ |
|---------|--|
| Default | _  |
| Context | http, server, location                   |

Specifies a file with passphrases for *secret keys* where each passphrase is specified on a separate line. Passphrases are tried in turn when loading the key.

## grpc\_ssl\_protocols

| Syntax  | grpc_ssl_protocols [SSLv2] [SSLv3] [TLSv1] [TLSv1.1] [TLSv1.2] [TLSv1.3]; |
|---------|---|
| Default | <pre>grpc_ssl_protocols TLSv1.2 TLSv1.3;</pre>                            |
| Context | http, server, location  |

Changed in version 1.2.0: TLSv1.3 parameter added to default set.

Enables the specified protocols for requests to a gRPC HTTPS server.

## grpc\_ssl\_server\_name

| Syntax  | <pre>grpc_ssl_server_name on   off;</pre> |
|---------|---|
| Default | <pre>grpc_ssl_server_name off;</pre>      |
| Context | http, server, location                    |

Enables or disables passing the server name set by the  $grpc\_ssl\_name$  directive via the Server Name Indication TLS extension (SNI, RFC 6066) while establishing a connection with the gRPC SSL server.

### grpc\_ssl\_session\_reuse

| Syntax  | <pre>grpc_ssl_session_reuse on   off;</pre> |
|---------|---|
| Default | <pre>grpc_ssl_session_reuse on;</pre>       |
| Context | http, server, location                      |

Determines whether SSL sessions can be reused when working with the gRPC server. If the errors "SSL3 GET FINISHED:digest check failed" appear in the logs, try disabling session reuse.



# grpc\_ssl\_trusted\_certificate

| Syntax  | ${	t grpc\_ssl\_trusted\_certificate} \ file;$ |
|---------|--|
| Default | _  |
| Context | http, server, location                         |

Specifies a file with trusted CA certificates in the PEM format used to *verify* the certificate of the gRPC SSL server.

### grpc ssl verify

| Syntax  | <pre>grpc_ssl_verify on   off;</pre> |
|---------|--------------------------------------|
| Default | <pre>grpc_ssl_verify off;</pre>      |
| Context | http, server, location               |

Enables or disables verification of the gRPC SSL server certificate.

# grpc ssl verify depth

| Syntax  | <pre>grpc_ssl_verify_depth number;</pre> |
|---------|--|
| Default | <pre>grpc_ssl_verify_depth 1;</pre>      |
| Context | http, server, location                   |

Sets the verification depth in the gRPC SSL server certificates chain.

### **GunZIP**

The module is a filter that decompresses responses with "Content-Encoding: gzip" for clients that do not support "gzip" encoding method. The module will be useful when it is desirable to store data compressed to save space and reduce I/O costs.

When building from the source code, this module isn't built by default; it should be enabled with the --with-http\_gunzip\_module build option.

In packages and images from our repos, the module is included in the build.

### **Configuration Example**

```
location /storage/ {
    gunzip on;
# ...
}
```

#### **Directives**

### gunzip

| Syntax  | gunzip on   off;       |
|---------|------------------------|
| Default | <pre>gunzip off;</pre> |
| Context | http, server, location |

Enables or disables decompression of gzipped responses for clients that lack gzip support. If enabled, the following directives are also taken into account when determining if clients support gzip: gzip\_http\_version, gzip\_proxied and gzip\_disable. See also the gzip\_vary directive.



# gunzip\_buffers

| Syntax  | <pre>gunzip_buffers number size;</pre>   |
|---------|--|
| Default | <pre>gunzip_buffers 32 4k   16 8k;</pre> |
| Context | http, server, location                   |

Sets the number and size of buffers used to decompress a response. By default, the buffer size is equal to one memory page. This is either 4K or 8K, depending on a platform.

#### **GZip**

The module is a filter that compresses responses using the "gzip" method. This often helps to reduce the size of transmitted data by half or even more.

# **Caution**

When using the SSL/TLS protocol, compressed responses may be subject to BREACH attacks.

## **Configuration Example**

```
gzip
gzip_min_length 1000;
gzip_proxied expired no-cache no-store private auth;
gzip_types text/plain application/xml;
```

The \$gzip ratio variable can be used to log the achieved compression ratio.

#### **Directives**

#### gzip

| Syntax  | gzip on   off;                         |
|---------|--|
| Default | gzip off;                              |
| Context | http, server, location, if in location |

Enables or disables gzipping of responses.

# gzip\_buffers

| Syntax  | gzip_buffers number size;   |
|---------|-----------------------------|
| Default | gzip_buffers 32 4k   16 8k; |
| Context | http, server, location      |

Sets the number and size of buffers used to compress a response. By default, the buffer size is equal to one memory page. This is either 4K or 8K, depending on a platform.

## gzip\_comp\_level

| Syntax  | <pre>gzip_comp_level level;</pre> |
|---------|-----------------------------------|
| Default | <pre>gzip_comp_level 1;</pre>     |
| Context | http, server, location            |



Sets a gzip compression level of a response. Acceptable values are in the range from 1 to 9.

## gzip disable

| Syntax  | gzip_disable regex;    |
|---------|------------------------|
| Default | _                      |
| Context | http, server, location |

Disables gzipping of responses for requests with "User-Agent" header fields matching any of the specified regular expressions.

The special mask msie6 corresponds to the regular expression "MSIE [4-6].", but works faster. "MSIE 6.0; ... SV1" is excluded from this mask.

## gzip\_http\_version

| Syntax  | gzip_http_version 1.0   1.1;      |
|---------|-----------------------------------|
| Default | <pre>gzip_http_version 1.1;</pre> |
| Context | http, server, location            |

Sets the minimum HTTP version of a request required to compress a response.

### gzip min length

| Syntax  | gzip_min_length length;        |
|---------|--------------------------------|
| Default | <pre>gzip_min_length 20;</pre> |
| Context | http, server, location         |

Sets the minimum length of a response that will be gzipped. The length is determined only from the "Content-Length" response header field.

## gzip\_proxied

| Syntax  | <pre>gzip_proxied off   expired   no-cache   no-store   private   no_last_modified</pre> |
|---------|--|
|         | no_etag   auth   any;  |
| Default | <pre>gzip_proxied off;</pre>   |
| Context | http, server, location   |

Enables or disables gzipping of responses for proxied requests depending on the request and response. The fact that the request is proxied is determined by the presence of the "Via" request header field. The directive accepts multiple parameters:



| off              | disables compression for all proxied requests, ignoring other parameters;          |
|------------------|--|
| expired          | enables compression if a response header includes the "Expires" field with a value |
|                  | that disables caching;   |
| no-cache         | enables compression if a response header includes the "Cache-Control" field with   |
|                  | the "no-cache" parameter;  |
| no-store         | enables compression if a response header includes the "Cache-Control" field with   |
|                  | the "no-store" parameter;  |
| private          | enables compression if a response header includes the "Cache-Control" field with   |
|                  | the "private" parameter;   |
| no_last_modified | enables compression if a response header does not include the "Last-Modified"      |
|                  | field;   |
| no_etag          | enables compression if a response header does not include the "ETag" field;        |
| auth             | enables compression if a request header includes the "Authorization" field;        |
| any              | enables compression for all proxied requests.                                      |

### gzip\_types

| Syntax  | gzip_types mime-type;            |
|---------|----------------------------------|
| Default | <pre>gzip_types text/html;</pre> |
| Context | http, server, location           |

Enables gzipping of responses for the specified MIME types in addition to text/html. The special value "\*" matches any MIME type. Responses with the text/html type are always compressed.

## gzip\_vary

| Syntax  | gzip_vary on   off;       |
|---------|---------------------------|
| Default | <pre>gzip_vary off;</pre> |
| Context | http, server, location    |

Enables or disables inserting the "Vary: Accept-Encoding" response header field if the directives gzip, gzip static or gunzip are active.

# **Built-in Variables**

### \$gzip\_ratio

achieved compression ratio, computed as the ratio between the original and compressed response sizes.

#### **GZip Static**

Allows sending precompressed files with the ".gz" filename extension instead of regular files.

When building from the source code, this module isn't built by default; it should be enabled with the --with-http\_gzip\_static\_module build option.

In packages and images from our repos, the module is included in the build.

### **Configuration Example**

```
gzip_static on;
gzip_proxied expired no-cache no-store private auth;
```



#### **Directives**

### gzip\_static

| Syntax  | gzip_static on   off   always; |
|---------|--------------------------------|
| Default | <pre>gzip_static off;</pre>    |
| Context | http, server, location         |

Enables (on) or disables (off) checking the existence of precompressed files. The following directives are also taken into account:  $gzip\_http\_version$ ,  $gzip\_proxied$ ,  $gzip\_disable$  and  $gzip\_vary$ .

With always, gzipped files are used in all cases, without checking if the client supports it. This is useful if there are no uncompressed files on the disk anyway or the *GunZIP* module is used.

The files can be compressed using the gzip command, or any other compatible one. It is recommended that the modification date and time of original and compressed files be the same.

#### **Headers**

Allows adding the "Expires" and "Cache-Control" header fields, and arbitrary fields, to a response header.

#### **Configuration Example**

```
expires 24h;
expires modified +24h;
expires @24h;
expires 0;
expires -1;
expires epoch;
expires $expires;
add_header Cache-Control private;
```

# Directives

#### add header

| Syntax  | add_header name value [always];        |
|---------|--|
| Default | _                                      |
| Context | http, server, location, if in location |

Adds the specified field to a response header provided that the response code equals 200, 201, 204, 206, 301, 302, 303, 304, 307, or 308. Parameter value can contain variables.

There could be several add\_header directives. These directives are inherited from the previous configuration level if and only if there are no add\_header directives defined on the current level.

If the always parameter is specified, the header field will be added regardless of the response code.

# add\_trailer

| Syntax  | add_trailer name value [always];       |
|---------|--|
| Default | _                                      |
| Context | http, server, location, if in location |



Adds the specified field to the end of a response provided that the response code equals 200, 201, 206, 301, 302, 303, 307, or 308. Parameter value can contain variables.

There could be several add\_trailer directives. These directives are inherited from the previous configuration level if and only if there are no add\_trailer directives defined on the current level.

If the always parameter is specified the specified field will be added regardless of the response code.

#### expires

| Syntax  | expires [modified] $time$ ;<br>expires epoch   max   off; |
|---------|---|
| Default | expires off;  |
| Context | http, server, location, if in location                    |

Enables or disables adding or modifying the "Expires" and "Cache-Control" response header fields provided that the response code equals 200, 201, 204, 206, 301, 302, 303, 304, 307, or 308. The parameter can be a positive or negative time.

The time in the "Expires" field is computed as a sum of the current time and time specified in the directive. If the modified parameter is used then the time is computed as a sum of the file's modification time and the time specified in the directive.

In addition, it is possible to specify a time of day using the "@" prefix:

```
expires @15h30m;
```

The contents of the "Cache-Control" field depends on the sign of the specified time:

- time is negative "Cache-Control: no-cache".
- time is positive or zero "Cache-Control: max-age=`t`", where t is a time specified in the directive, in seconds.

| epoch | sets "Expires" to the value "Thu, 01 Jan 1970 00:00:01 GMT", and "Cache-Control" to "no-cache". |
|-------|---|
| max   | sets "Expires" to the value "Thu, 31 Dec 2037 23:55:55 GMT", and "Cache-Control" to 10 years.   |
| off   | disables adding or modifying the "Expires" and "Cache-Control" response header fields.          |

The last parameter value can contain variables:

```
map $sent_http_content_type $expires {
    default off;
    application/pdf 42d;
    ~image/ max;
}
expires $expires;
```

### **Image Filter**

The module is a filter that transforms images in JPEG, GIF, PNG, and WebP formats.

When building from the source code, this module isn't built by default; it should be enabled with the --with-http\_image\_filter\_module build option.

In our repositories, the module is built dynamically and is available as a separate package named angie-module-image-filter or angie-pro-module-image-filter.



### Important

This module utilizes the libgd library. It is recommended to use the latest available version of the library.

To transform images in WebP format, the libgd library must be compiled with WebP support.

### **Configuration Example**

```
location /img/ {
    proxy_pass    http://backend;
    image_filter resize 150 100;
    image_filter rotate 90;
    error_page    415 = /empty;
}
location = /empty {
    empty_gif;
}
```

#### **Directives**

# image filter

```
Syntax

• image_filter off;
• image_filter test;
• image_filter size;
• image_filter rotate 90 | 180 | 270;
• image_filter resize width height;
• image_filter crop width height;

Default image_filter off;
Context location
```

Sets the type of transformation to perform on images:

| off                    | turns off module processing in a surrounding location.   |
|------------------------|--|
| test                   | ensures that responses are images in either JPEG, GIF, PNG, or WebP format. Otherwise, the 415 (Unsupported Media Type) error is returned.   |
| size                   | outputs information about images in a JSON format, e.g.: "img": { "width": 100, "height": 100, "type": "gif" } In case of an error, the output is as follows: {}   |
| rotate<br>90 180 270   | rotates images counter-clockwise by the specified number of degrees. Parameter value can contain variables. This mode can be used either alone or along with the resize and crop transformations.  |
| resize width<br>height | proportionally reduces an image to the specified sizes. To reduce by only one dimension, another dimension can be specified as "-". In case of an error, the server will return code 415 (Unsupported Media Type). Parameter values can contain variables. When used along with the <i>rotate</i> parameter, the rotation happens <b>after</b> reduction.  |
| crop width height      | proportionally reduces an image to the larger side size and crops extraneous edges by another side. To reduce by only one dimension, another dimension can be specified as "-". In case of an error, the server will return code 415 (Unsupported Media Type). Parameter values can contain variables. When used along with the <i>rotate</i> parameter, the rotation happens <b>before</b> reduction. |



## image\_filter\_buffer

| Syntax  | $image\_filter\_buffer$ $size;$    |
|---------|------------------------------------|
| Default | <pre>image_filter_buffer 1M;</pre> |
| Context | http, server, location             |

Sets the maximum size of the buffer used for reading images. When the size is exceeded the server returns error 415 (Unsupported Media Type).

# image\_filter\_interlace

| Syntax  | <pre>image_filter_interlace on   off;</pre> |
|---------|---|
| Default | <pre>image_filter_interlace off;</pre>      |
| Context | http, server, location                      |

If enabled, final images will be interlaced. For JPEG, final images will be in "progressive JPEG" format.

# image filter jpeg quality

| Syntax  | <pre>image_filter_jpeg_quality quality;</pre> |
|---------|---|
| Default | <pre>image_filter_jpeg_quality 75;</pre>      |
| Context | http, server, location                        |

Sets the desired quality of the transformed JPEG images. Acceptable values are in the range from 1 to 100. Lesser values usually imply both lower image quality and less data to transfer. The maximum recommended value is 95. Parameter value can contain variables.

# image\_filter\_sharpen

| Syntax  | <pre>image_filter_sharpen percent;</pre> |
|---------|--|
| Default | <pre>image_filter_sharpen 0;</pre>       |
| Context | http, server, location                   |

Increases sharpness of the final image. The sharpness percentage can exceed 100. The 0 value disables sharpening. Parameter value can contain variables.

# $image\_filter\_transparency$

| Syntax  | <pre>image_filter_transparency on   off;</pre> |
|---------|--|
| Default | <pre>image_filter_transparency on;</pre>       |
| Context | http, server, location                         |

Defines whether transparency should be preserved when transforming GIF images or PNG images with colors specified by a palette. The loss of transparency results in images of a better quality. The alpha channel transparency in PNG is always preserved.

## image\_filter\_webp\_quality

| Syntax  | <pre>image_filter_webp_quality quality;</pre> |
|---------|---|
| Default | <pre>image_filter_webp_quality 80;</pre>      |
| Context | http, server, location                        |



Sets the desired quality of the transformed WebP images. Acceptable values are in the range from 1 to 100. Lesser values usually imply both lower image quality and less data to transfer. Parameter value can contain variables.

#### Index

The module processes requests ending with the slash character (/). Such requests can also be processed by the http autoindex and http random index modules.

#### **Configuration Example**

```
location / {
   index index.$geo.html index.html;
}
```

#### **Directives**

#### index

```
Syntax index file ...;
Default index index.html;
Context http, server, location
```

Defines files that will be used as an index. The file name can contain variables. Files are checked in the specified order. The last element of the list can be a file with an absolute path. Example:

```
index index.$geo.html index.0.html /index.html;
```

It should be noted that using an index file causes an internal redirect, and the request can be processed in a different location. For example, with the following configuration:

```
location = / {
   index index.html;
}
location / {
# ...
}
```

A "" request will actually be processed in the second location as /index.html.

## JS

The module is used to implement handlers in njs - a subset of the JavaScript language.

In our repositories, the module is built dynamically and is available as a separate package named angie-module-njs or angie-pro-module-njs.

## **Configuration Example**

```
http {
    js_import http.js;

    js_set $foo http.foo;
    js_set $summary http.summary;
    js_set $hash http.hash;
```



```
resolver 127.0.0.53;
    server {
       listen 8000;
       location / {
           add_header X-Foo $foo;
           js_content http.baz;
       location = /summary {
           return 200 $summary;
        location = /hello {
           js_content http.hello;
       location = /fetch {
                                        http.fetch;
           js_content
            js_fetch_trusted_certificate /path/to/ISRG_Root_X1.pem;
        }
        location = /crypto {
           add_header Hash $hash;
           return 200;
       }
   }
}
```

The http.js file:

```
function foo(r) {
   r.log("hello from foo() handler");
   return "foo";
function summary(r) {
   var a, s, h;
   s = "JS summary\n\n";
   s += "Method: " + r.method + "\n";
   s += "HTTP version: " + r.httpVersion + "\n";
   s += "Host: " + r.headersIn.host + "\n";
   s += "Remote Address: " + r.remoteAddress + "\n";
   s += "URI: " + r.uri + "\n";
   s += "Headers:\n";
   for (h in r.headersIn) {
       s += "header'" + h + "'is'" + r.headersIn[h] + "'\n";
   s += "Args:\n";
   for (a in r.args) {
       s += " arg '" + a + "' is '" + r.args[a] + "'\n";
```



```
return s;
}
function baz(r) {
   r.status = 200;
    r.headersOut.foo = 1234;
    r.headersOut['Content-Type'] = "text/plain; charset=utf-8";
    r.headersOut['Content-Length'] = 15;
    r.sendHeader();
    r.send("nginx");
    r.send("java");
    r.send("script");
    r.finish();
}
function hello(r) {
    r.return(200, "Hello world!");
}
async function fetch(r) {
    let results = await Promise.all([ngx.fetch('https://google.com/'),
                                      ngx.fetch('https://google.ru/')]);
    r.return(200, JSON.stringify(results, undefined, 4));
}
async function hash(r) {
    let hash = await crypto.subtle.digest('SHA-512', r.headersIn.host);
    r.setReturnValue(Buffer.from(hash).toString('hex'));
export default {foo, summary, baz, hello, fetch, hash};
```

## **Directives**

# js body filter

| Syntax  | <pre>js_body_filter function   module.function [buffer_type=string   buffer];</pre> |
|---------|---|
| Default | _   |
| Context | location, if in location, limit_except  |

Sets an njs function as a response body filter. The filter function is called for each data chunk of a response body with the following arguments:

| r     | the HTTP request object  |
|-------|--|
| data  | the incoming data chunk, may be a string or Buffer depending on the buffer_type value, by default is a string. |
| flags | an object with the following properties: - $last - a$ boolean value, true if data is the last buffer           |

The filter function can pass its own modified version of the input data chunk to the next body filter by calling r.sendBuffer(). For example, to transform all the lowercase letters in the response body:



```
function filter(r, data, flags) {
   r.sendBuffer(data.toLowerCase(), flags);
}
```

To stop filtering (following data chunks will be passed to client without calling  $js\_body\_filter$ ), r.done() can be used.

If the filter function changes the length of the response body, then it is required to clear out the "Content-Length" response header (if any) in  $js\_header\_filter$  to enforce chunked transfer encoding.

# 1 Note

As the  $js\_body\_filter$  handler returns its result immediately, it supports only synchronous operations. Thus, asynchronous operations such as r.subrequest() or setTimeout() are not supported.

## js content

| Syntax  | <pre>js_content function   module.function;</pre> |
|---------|---|
| Default | _   |
| Context | location, if in location, limit_except            |

Sets an nis function as a location content handler. Module functions can be referenced.

## js\_fetch\_buffer\_size

| Syntax  | js_fetch_buffer_size size;           |
|---------|--------------------------------------|
| Default | <pre>js_fetch_buffer_size 16k;</pre> |
| Context | http, server, location               |

Sets the size of the buffer used for reading and writing with Fetch API.

#### js fetch ciphers

| Syntax  | <pre>js_fetch_ciphers ciphers;</pre>          |
|---------|---|
| Default | <pre>js_fetch_ciphers HIGH:!aNULL:!MD5;</pre> |
| Context | http, server, location                        |

Specifies the enabled ciphers for HTTPS connections with Fetch API. The ciphers are specified in the format understood by the OpenSSL library.

The list of ciphers depends on the version of OpenSSL installed. The full list can be viewed using the openssl ciphers command.

# js\_fetch\_max\_response\_buffer\_size

| Syntax  | <pre>js_fetch_max_response_buffer_size size;</pre> |
|---------|--|
| Default | <pre>js_fetch_max_response_buffer_size 1m;</pre>   |
| Context | http, server, location                             |

Sets the maximum size of the response received with Fetch API.



## js\_fetch\_protocols

| Syntax  | <pre>js_fetch_protocols [TLSv1] [TLSv1.1] [TLSv1.2] [TLSv1.3];</pre> |
|---------|--|
| Default | <pre>js_fetch_protocols TLSv1 TLSv1.1 TLSv1.2;</pre>                 |
| Context | http, server, location   |

Enables the specified protocols for HTTPS connections with Fetch API.

#### js fetch timeout

| Syntax  | <pre>js_fetch_timeout time;</pre> |
|---------|-----------------------------------|
| Default | <pre>js_fetch_timeout 60s;</pre>  |
| Context | http, server, location            |

Defines a timeout for reading and writing for Fetch API. The timeout is set only between two successive read/write operations, not for the whole response. If no data is transmitted within this time, the connection is closed.

# js\_fetch\_trusted\_certificate

| Syntax  | ${\tt js\_fetch\_trusted\_certificate}$ $file;$ |
|---------|---|
| Default | _   |
| Context | http, server, location                          |

Specifies a file with trusted CA certificates in the PEM format used to verify the HTTPS certificate with Fetch API.

# js fetch verify

| Syntax  | js_fetch_verify on   off;      |
|---------|--------------------------------|
| Default | <pre>js_fetch_verify on;</pre> |
| Context | http, server, location         |

Enables or disables verification of the HTTPS server certificate with Fetch API.

# js\_fetch\_verify\_depth

| Syntax  | <pre>js_fetch_verify_depth number;</pre> |
|---------|--|
| Default | <pre>js_fetch_verify_depth 100;</pre>    |
| Context | http, server, location                   |

Sets the verification depth in the HTTPS server certificates chain with Fetch API.

# js header filter

| Syntax  | ${\tt js\_header\_filter}\ function \mid module.function;$ |
|---------|--|
| Default | _  |
| Context | location, if in location, limit_except                     |



Sets an njs function as a response header filter. The directive allows changing arbitrary header fields of a response header.

#### 1 Note

As the js header filter handler returns its result immediately, it supports only synchronous operations. Thus, asynchronous operations such as r.subrequest() or setTimeout() are not supported.

## js import

| Syntax  | <pre>js_import module.js   export_name from module.js;</pre> |
|---------|--|
| Default | _  |
| Context | http, server, location                                       |

Imports a module that implements location and variable handlers in njs. The export name is used as a namespace to access module functions. If the export name is not specified, the module name will be used as a namespace.

```
js_import http.js;
```

Here, the module name http is used as a namespace while accessing exports. If the imported module exports foo(), http.foo is used to refer to it.

Several js\_import directives can be specified.

## js\_path

| Syntax  | <pre>js_path path;</pre> |
|---------|--------------------------|
| Default | _                        |
| Context | http, server, location   |

Sets an additional path for njs modules.

## js preload object

| Syntax  | <pre>js_preload_object name.json   name from file.json;</pre> |
|---------|---|
| Default | _   |
| Context | http, server, location  |

Preloads an immutable object at configure time. The name is used as a name of the global variable though which the object is available in njs code. If the name is not specified, the file name will be used instead.

```
js_preload_object map.json;
```

Here, the map is used as a name while accessing the preloaded object.

Several js preload object directives can be specified.



## js\_set

| Syntax  | ${\tt js\_set}$ ${\tt \$variable}$ function $ $ module.function; |
|---------|--|
| Default | _  |
| Context | http, server, locatio  |

Sets an njs function for the specified variable. Module functions can be referenced.

The function is called when the variable is referenced for the first time for a given request. The exact moment depends on a *phase* at which the variable is referenced. This can be used to perform some logic not related to variable evaluation. For example, if the variable is referenced only in the *log\_format* directive, its handler will not be executed until the log phase. This handler can be used to do some cleanup right before the request is freed.

# 1 Note

As the  $js\_set$  handler returns its result immediately, it supports only synchronous callbacks. Thus, asynchronous callbacks such as r.subrequest() or setTimeout() are not supported.

# js\_shared\_dict\_zone

| Syntax  | <pre>js_shared_dict_zone z [evict];</pre> | zone=name:size | $[\mathtt{timeout} = time]$ | [type=string | number] |
|---------|---|----------------|-----------------------------|--------------|---------|
| Default | _   |                |                             |              |         |
| Context | http                                      |                |                             |              |         |

Sets the name and size of the shared memory zone that keeps the key-value dictionary shared between worker processes.

| type    | optional parameter, allows redefining the value type to number; by default, the shared dictionary uses string for keys and values |
|---------|---|
| timeout | optional parameter, sets the time after which all shared dictionary entries are removed from the zone                             |
| evict   | optional parameter, removes the oldest key-value pair when the zone storage is exhausted  |

## Examples:

```
example.conf:
    # Creates a 1Mb dictionary with string values,
    # removes key-value pairs after 60 seconds of inactivity:
    js_shared_dict_zone zone=foo:1M timeout=60s;

# Creates a 512Kb dictionary with string values,
    # forcibly removes oldest key-value pairs when the zone is exhausted:
    js_shared_dict_zone zone=bar:512K timeout=30s evict;

# Creates a 32Kb permanent dictionary with number values:
    js_shared_dict_zone zone=num:32k type=number;
```

```
example.js:
   function get(r) {
     r.return(200, ngx.shared.foo.get(r.args.key));
```



```
function set(r) {
    r.return(200, ngx.shared.foo.set(r.args.key, r.args.value));
}

function delete(r) {
    r.return(200, ngx.shared.bar.delete(r.args.key));
}

function increment(r) {
    r.return(200, ngx.shared.num.incr(r.args.key, 2));
}
```

# js\_var

| Syntax  | js_var \$variable [value]; |
|---------|----------------------------|
| Default | _                          |
| Context | stream, server             |

Declares a writable variable. The value can contain text, variables, and their combination. The variable is not overwritten after a redirect unlike variables created with the *set* directive.

## Request Argument

Each HTTP njs handler receives one argument, a request object.

## Limit Conn

The module is used to limit the number of connections per the defined key, in particular, the number of connections from a single IP address.

Not all connections are counted. A connection is counted only if it has a request being processed by the server and the whole request header has already been read.

#### **Configuration Example**

```
http {
    limit_conn_zone $binary_remote_addr zone=addr:10m;
    ...
    server {
        ...
        location /download/ {
            limit_conn addr 1;
        }
}
```

#### **Directives**



## limit\_conn

| Syntax  | <pre>limit_conn zone number;</pre> |
|---------|------------------------------------|
| Default | _                                  |
| Context | http, server, location             |

Sets the shared memory zone and the maximum allowed number of connections for a given key value. When this limit is exceeded, the server will return the *error* in reply to a request. For example, the directives

```
limit_conn_zone $binary_remote_addr zone=addr:10m;
server {
    location /download/ {
        limit_conn addr 1;
    }
```

allow only one connection per IP address at a time.

## Note

In HTTP/2 and SPDY, each concurrent request is considered a separate connection.

There could be several limit\_conn directives. For example, the following configuration will limit the number of connections to the server per client IP and, at the same time, the total number of connections to the virtual server:

```
limit_conn_zone $binary_remote_addr zone=perip:10m;
limit_conn_zone $server_name zone=perserver:10m;
server {
    ...
    limit_conn perip 10;
    limit_conn perserver 100;
}
```

These directives are inherited from the previous configuration level if and only if there are no limit\_conn directives defined on the current level.

## limit \_conn \_dry \_run

| Syntax  | <pre>limit_conn_dry_run on   off;</pre> |
|---------|---|
| Default | <pre>limit_conn_dry_run off;</pre>      |
| Context | http, server, location                  |

Enables the dry run mode. In this mode, the number of connections is not limited, however, in the *shared memory zone*, the number of excessive connections is accounted as usual.

#### limit conn log level

| Syntax  | <pre>limit_conn_log_level info   notice   warn   error;</pre> |
|---------|---|
| Default | <pre>limit_conn_log_level error;</pre>                        |
| Context | http, server, location  |



Sets the desired logging level for cases when the server limits the number of connections.

## limit conn status

| Syntax  | limit_conn_status code;           |
|---------|-----------------------------------|
| Default | <pre>limit_conn_status 503;</pre> |
| Context | http, server, location            |

Sets the status code to return in response to rejected requests.

#### limit conn zone

| Syntax  | ${	t limit\_conn\_zone}\ key\ {	t zone} = name:size;$ |
|---------|---|
| Default | _   |
| Context | http  |

Sets parameters for a shared memory zone that will keep states for various keys. In particular, the state includes the current number of connections. The key can contain text, variables, and their combination. Requests with an empty key value are not accounted.

Usage example:

```
limit_conn_zone $binary_remote_addr zone=addr:10m;
```

Here, a client IP address serves as a key. Note that instead of \$remote\_addr, the \$binary\_remote\_addr variable is used here.

The **\$remote\_addr** variable's size can vary from 7 to 15 bytes. The stored state occupies either 32 or 64 bytes of memory on 32-bit platforms and always 64 bytes on 64-bit platforms.

The \$binary\_remote\_addr variable's size is always 4 bytes for IPv4 addresses or 16 bytes for IPv6 addresses. The stored state always occupies 32 or 64 bytes on 32-bit platforms and 64 bytes on 64-bit platforms.

One megabyte zone can keep about 32 thousand 32-byte states or about 16 thousand 64-byte states. If the zone storage is exhausted, the server will return the *error* to all further requests.

# **Built-in Variables**

#### \$limit\_conn\_status

keeps the result of limiting the number of connections: PASSED, REJECTED or REJECTED\_DRY\_RUN

#### Limit Req

The module is used to limit the request processing rate per a defined key, in particular, the processing rate of requests coming from a single IP address. The limitation is done using the "leaky bucket" method.

## Configuration Example

```
http {
    limit_req_zone $binary_remote_addr zone=one:10m rate=1r/s;
    ...
    server {
```



```
location /search/ {
    limit_req zone=one burst=5;
}
```

#### **Directives**

#### limit req

| Syntax  | <pre>limit_req zone number [burst=number] [nodelay   delay=number];</pre> |
|---------|---|
| Default | _   |
| Context | http, server, location  |

Sets the shared memory zone and the maximum burst size of requests. If the requests rate exceeds the rate configured for a zone, their processing is delayed such that requests are processed at a defined rate. Excessive requests are delayed until their number exceeds the maximum burst size in which case the request is terminated with an *error*. By default, the maximum burst size is equal to zero. For example, the directives

```
limit_req_zone $binary_remote_addr zone=one:10m rate=1r/s;
server {
    location /search/ {
        limit_req zone=one burst=5;
    }
```

allow not more than 1 request per second at an average, with bursts not exceeding 5 requests.

If delaying of excessive requests while requests are being limited is not desired, the parameter nodelay should be used:

```
limit_req zone=one burst=5 nodelay;
```

The delay parameter specifies a limit at which excessive requests become delayed. Default value is zero, i.e. all excessive requests are delayed.

There could be several limit\_req directives. For example, the following configuration will limit the processing rate of requests coming from a single IP address and, at the same time, the request processing rate by the virtual server:

```
limit_req_zone $binary_remote_addr zone=perip:10m rate=1r/s;
limit_req_zone $server_name zone=perserver:10m rate=10r/s;
server {
    ...
    limit_req zone=perip burst=5 nodelay;
    limit_req zone=perserver burst=10;
}
```

These directives are inherited from the previous configuration level if and only if there are no limit\_req directives defined on the current level.



## limit\_req\_dry\_run

| Syntax  | <pre>limit_req_dry_run on   off;</pre> |
|---------|--|
| Default | <pre>limit_req_dry_run off;</pre>      |
| Context | http, server, location                 |

Enables the dry run mode. In this mode, requests processing rate is not limited, however, in the *shared memory zone*, the number of excessive requests is accounted as usual.

## limit req log level

| Syntax  | limit_req_log_level info   notice   warn   error; |
|---------|---|
| Default | <pre>limit_req_log_level error;</pre>             |
| Context | http, server, location                            |

Sets the desired logging level for cases when the server refuses to process requests due to rate exceeding, or delays request processing. Logging level for delays is one point less than for refusals; for example, if limit\_req\_log\_level notice is specified, delays are logged with the info level.

#### limit\_req\_status

| Syntax  | limit_req_status $code;$         |
|---------|----------------------------------|
| Default | <pre>limit_req_status 503;</pre> |
| Context | http, server, location           |

Sets the status code to return in response to rejected requests.

# limit\_req\_zone

| Syntax  | limit_req_zone key zone=name:size rate=rate; |
|---------|--|
| Default | _  |
| Context | http   |

Sets parameters for a shared memory zone that will keep states for various keys. In particular, the state stores the current number of excessive requests. The key can contain text, variables, and their combination. Requests with an empty key value are not accounted.

Usage example:

```
limit_req_zone $binary_remote_addr zone=one:10m rate=1r/s;
```

Here, the states are kept in a 10 megabyte zone one, and an average request processing rate for this zone cannot exceed 1 request per second.

A client IP address serves as a key. Note that instead of **\$remote\_addr**, the **\$binary\_remote\_addr** variable is used here.

The \$binary\_remote\_addr variable's size is always 4 bytes for IPv4 addresses or 16 bytes for IPv6 addresses. The stored state always occupies 64 bytes on 32-bit platforms and 128 bytes on 64-bit platforms.

One megabyte zone can keep about 16 thousand 64-byte states or about 8 thousand 128-byte states.

If the zone storage is exhausted, the least recently used state is removed. If even after that a new state cannot be created, the request is terminated with an *error*.



The rate is specified in requests per second (r/s). If a rate of less than one request per second is desired, it is specified in request per minute (r/m). For example, half-request per second is 30r/m.

#### **Built-in Variables**

```
$limit_req_status
```

keeps the result of limiting the request processing rate: PASSED, DELAYED, REJECTED, DELAYED\_DRY\_RUN, or REJECTED\_DRY\_RUN

#### Log

The module writes request logs in the specified format.

Requests are logged in the context of a location where processing ends. It may be different from the original location, if an *internal redirect* happens during request processing.

#### **Configuration Example**

#### **Directives**

#### access log

| Syntax  | <pre>access_log path [format [buffer=size] [gzip[=level]] [flush=time] [if=condition]]; access_log off;</pre> |
|---------|---|
| Default | access_log logs/access.log combined; (the path depends on thehttp-log-path build option)                      |
| Context | http, server, location, if in location, limit_except  |

Sets the path, format, and configuration for a buffered log write. Several logs can be specified on the same configuration level. Logging to syslog can be configured by specifying the "syslog:" prefix in the first parameter. The special value off cancels all access\_log directives on the current level. If the format is not specified then the predefined "combined" format is used.

If either the buffer or gzip parameter is used, writes to log will be buffered.

## \* Caution

The buffer size must not exceed the size of an atomic write to a disk file. For FreeBSD this size is unlimited.

When buffering is enabled, the data will be written to the file:

- if the next log line does not fit into the buffer;
- if the buffered data is older than specified by the flush parameter;
- when a worker process is *re-opening log files* or is shutting down.

If the gzip parameter is used, then the buffered data will be compressed before writing to the file. The compression level can be set between 1 (fastest, less compression) and 9 (slowest, best compression). By default, the buffer size is equal to 64K bytes, and the compression level is set to 1. Since the data is compressed in atomic blocks, the log file can be decompressed or read by "zcat" at any time.



#### Example:

```
access_log /path/to/log.gz combined gzip flush=5m;
```

# Important

For gzip compression to work, Angie must be built with the zlib library.

The file path can contain variables, but such logs have some constraints:

- the *user* whose credentials are used by worker processes should have permissions to create files in a directory with such logs;
- buffered writes do not work;
- the file is opened and closed for each log write. However, since the descriptors of frequently used files can be stored in a cache, writing to the old file can continue during the time specified by the open log file cache directive's valid parameter
- during each log write the existence of the request's root directory is checked, and if it does not exist the log is not created. It is thus a good idea to specify both *root* and *access\_log* on the same configuration level:

```
server {
   root /spool/vhost/data/$host;
   access_log /spool/vhost/logs/$host;
   ...
```

The if parameter enables conditional logging. A request will not be logged if the condition evaluates to "0" or an empty string. In the following example, the requests with response codes 2xx and 3xx will not be logged:

```
map $status $loggable {
    ~^[23] 0;
    default 1;
}
access_log /path/to/access.log combined if=$loggable;
```

## log format

| Syntax  | log_format name [escape=default   json   none] string; |
|---------|--|
| Default | <pre>log_format combined "";</pre>                     |
| Context | http   |

Specifies log format.

The escape parameter allows setting json or default characters escaping in variables, by default, default escaping is used. The none value disables escaping.

For default escaping, characters """, "\", and other characters with values less than 32 or above 126 are escaped as " $\xspace{\times}$ xXX". If the variable value is not found, a hyphen "-" will be logged.

For json escaping, all characters not allowed in JSON strings will be escaped: characters """ and "\" are escaped as "\"" and "\\", characters with values less than 32 are escaped as "\n", "\r", "\t", "\b", "\f", or "\u00XX".

Header lines sent to a client have the prefix sent\_http\_, for example, \$sent\_http\_content\_range.

The configuration always includes the predefined combined format:



# open log file cache

| Syntax  | $\label{eq:condition} \begin{split} & \texttt{open\_log\_file\_cache\ max} = N\ [\texttt{inactive} = time]\ [\texttt{min\_uses} = N]\ [\texttt{valid} = time]; \\ & \texttt{open\_log\_file\_cache\ off}; \end{split}$ |
|---------|--|
| Default | <pre>open_log_file_cache off;</pre>  |
| Context | http, server, location   |

Defines a cache that stores the file descriptors of frequently used logs whose names contain variables. The directive has the following parameters:

| max      | sets the maximum number of descriptors in a cache; if the cache becomes full the least recently used (LRU) descriptors are closed                |
|----------|--|
| inactive | sets the time after which the cached descriptor is closed if there were no access during this time; by default, 10 seconds                       |
| min_uses | sets the minimum number of file uses during the time defined by the inactive parameter to let the descriptor stay open in a cache; by default, 1 |
| valid    | sets the time after which it should be checked that the file still exists with the same name; by default, 60 seconds                             |
| off      | disables caching   |

#### Usage example:

```
open_log_file_cache max=1000 inactive=20s valid=1m min_uses=2;
```

#### Map

The module creates variables whose values depend on values of other variables.

#### **Configuration Example**

```
map $http_host $name {
    hostnames;
    default
                  0;
    example.com
                 1;
    *.example.com 1;
    example.org
    *.example.org 2;
    .example.net 3;
    wap.*
                  4;
}
map $http_user_agent $mobile {
    default 0;
    "~Opera Mini" 1;
```



#### **Directives**

#### map

```
Syntax map string $variable { ... }

Default —
Context http
```

Creates a new variable. Its value depends on the first parameter, specified as a string with variables, for example:

```
set $var1 "foo";
set $var2 "bar";

map $var1$var2 $new_variable {
    default "foobar_value";
}
```

Here, the variable **\$new\_variable** will have a value composed of the two variables **\$var1** and **\$var2**, or a default value if these variables are not defined.

# 1 Note

Since variables are evaluated only when they are used, the mere declaration even of a large number of "map" variables does not add any extra costs to request processing.

Parameters inside the map block specify a mapping between source and resulting values.

Source values are specified as strings or regular expressions.

Strings are matched ignoring the case.

A regular expression should either start with a "~" symbol for a case-sensitive matching, or with the "~\*" symbols for case-insensitive matching. A regular expression can contain named and positional captures that can later be used in other directives along with the resulting variable.

If a source value matches one of the names of special parameters described below, it should be prefixed with the "" symbol.

The resulting value can contain text, variable and their combination.

The following special parameters are also supported:

| default $value$ | sets the resulting value if the source value matches none of the specified variants. When <i>default</i> is not specified, the default resulting value will be an empty string. |
|-----------------|---|
| hostnames       | indicates that source values can be hostnames with a prefix or suffix mask. This parameter should be specified before the list of values.                                       |

For example,

```
*.example.com 1; example.* 1;
```

The following two records

```
example.com 1;
*.example.com 1;
```



can be combined:

```
.example.com 1;
```

| $\verb"include" \mathit{file}$ | includes a file with values. There can be several inclusions. |
|--------------------------------|---|
| volatile                       | indicates that the variable is not cacheable.                 |

If the source value matches more than one of the specified variants, e.g. both a mask and a regular expression match, the first matching variant will be chosen, in the following order of priority:

- 1. String value without a mask
- 2. Longest string value with a prefix mask, e.g. \*.example.com
- 3. Longest string value with a suffix mask, e.g. mail.\*
- 4. First matching regular expression (in order of appearance in a configuration file)
- 5. Default value (default)

## map hash bucket size

| Syntax  | map_hash_bucket_size $size$ ;              |
|---------|--|
| Default | <pre>map_hash_bucket_size 32 64 128;</pre> |
| Context | http                                       |

Sets the bucket size for the map variables hash tables. Default value depends on the processor's cache line size. The details of setting up hash tables are provided separately.

## map hash max size

| Syntax  | map_hash_max_size $size;$          |
|---------|------------------------------------|
| Default | <pre>map_hash_max_size 2048;</pre> |
| Context | http                               |

Sets the maximum size of the map variables hash tables. The details of setting up hash tables are provided separately.

## Memcached

The module is used to obtain responses from a memcached server. The key is set in the *\$memcached\_key* variable. A response should be put in memcached in advance by means external to Angie.

# **Configuration Example**



#### **Directives**

## memcached\_bind

| Syntax  | memcached_bind address [transparent]   off; |
|---------|---|
| Default | _   |
| Context | http, server, location                      |

Makes outgoing connections to a memcached server originate from the specified local IP address with an optional port. Parameter value can contain variables. The special value off cancels the effect of the memcached\_bind directive inherited from the previous configuration level, which allows the system to auto-assign the local IP address and port.

The transparent parameter allows outgoing connections to a memcached server originate from a non-local IP address, for example, from a real IP address of a client:

```
memcached_bind $remote_addr transparent;
```

In order for this parameter to work, it is usually necessary to run Angie worker processes with the superuser privileges. On Linux it is not required as if the transparent parameter is specified, worker processes inherit the  $CAP\_NET\_RAW$  capability from the master process.

## Important

It is necessary to configure kernel routing table to intercept network traffic from the memcached server.

# memcached\_buffer\_size

| Syntax  | memcached_buffer_size size;             |
|---------|---|
| Default | <pre>memcached_buffer_size 4k 8k;</pre> |
| Context | http, server, location                  |

Sets the size of the buffer used for reading the first part of the response received from the memcached server. The response is passed to the client synchronously, as soon as it is received.

#### memcached connect timeout

| Syntax  | ${\tt memcached\_connect\_timeout}\ time;$ |
|---------|--|
| Default | <pre>memcached_connect_timeout 60s;</pre>  |
| Context | http, server, location                     |

Defines a timeout for establishing a connection with a memcached server. It should be noted that this timeout cannot usually exceed 75 seconds.

# memcached\_gzip\_flag

| Syntax  | ${\tt memcached\_gzip\_flag} \ flag;$ |
|---------|---------------------------------------|
| Default | _                                     |
| Context | http, server, location                |

Enables the test for the flag presence in the memcached server response and sets the "Content-Encoding" response header field to "gzip" if the flag is set.



## $memcached\_next\_upstream$

| Syntax  | ${\tt memcached\_next\_upstream\;error} \mid {\tt timeout} \mid {\tt invalid\_response} \mid {\tt not\_found} \mid {\tt off}$ |
|---------|---|
|         | ;   |
| Default | memcached_next_upstream error timeout;  |
| Context | http, server, location  |

Specifies in which cases a request should be passed to the next server in the upstream pool:

| error            | an error occurred while establishing a connection with the server, passing a request to it, or reading the response header;      |
|------------------|--|
| timeout          | a timeout has occurred while establishing a connection with the server, passing a request to it, or reading the response header; |
| invalid_response | a server returned an empty or invalid response;  |
| not_found        | a response was not found on the server;  |
| off              | disables passing a request to the next server.   |

## 1 Note

One should bear in mind that passing a request to the next server is only possible if nothing has been sent to a client yet. That is, if an error or timeout occurs in the middle of the transferring of a response, fixing this is impossible.

The directive also defines what is considered an unsuccessful attempt of communication with a server.

| error, timeout,  | always considered unsuccessful attempts, even if they are not specified in the |
|------------------|--|
| invalid_response | directive  |
| not_found        | never considered unsuccessful attempts   |

Passing a request to the next server can be limited by the *number of tries* and by time.

# memcached next upstream timeout

| Syntax  | $memcached_next\_upstream\_timeout\ time;$    |
|---------|---|
| Default | <pre>memcached_next_upstream_timeout 0;</pre> |
| Context | http, server, location                        |

Limits the time during which a request can be passed to the *next* server.

| 0 | turns off this limitation |  |
|---|---------------------------|--|
|---|---------------------------|--|

# memcached\_next\_upstream\_tries

| Syntax  | memcached_next_upstream_tries number;       |
|---------|---|
| Default | <pre>memcached_next_upstream_tries 0;</pre> |
| Context | http, server, location                      |

Limits the number of possible tries for passing a request to the *next* server.



| 0 | turns off this limitation |  |
|---|---------------------------|--|
|---|---------------------------|--|

## memcached pass

| Syntax  | $memcached_pass \ uri;$  |
|---------|--------------------------|
| Default | _                        |
| Context | location, if in location |

Sets the memcached server address. The address can be specified as a domain name or IP address, and a port:

```
memcached_pass localhost:11211;
```

or as a UNIX domain socket path:

```
memcached_pass unix:/tmp/memcached.socket;
```

If a domain name resolves to several addresses, all of them will be used in a round-robin fashion. In addition, an address can be specified as a *server group*. If a group is used, you cannot specify the port with it; instead, specify the port for each server within the group individually.

# $memcached\_read\_timeout$

| Syntax  | ${\tt memcached\_read\_timeout}\ time;$ |
|---------|---|
| Default | <pre>memcached_read_timeout 60s;</pre>  |
| Context | http, server, location                  |

Defines a timeout for reading a response from the memcached server. The timeout is set only between two successive read operations, not for the transmission of the whole response. If the memcached server does not transmit anything within this time, the connection is closed.

## memcached send timeout

| Syntax  | ${\tt memcached\_send\_timeout}\ time;$ |
|---------|---|
| Default | <pre>memcached_send_timeout 60s;</pre>  |
| Context | http, server, location                  |

Sets a timeout for transmitting a request to the memcached server. The timeout is set only between two successive write operations, not for the transmission of the whole request. If the memcached server does not receive anything within this time, the connection is closed.

# memcached\_socket\_keepalive

| Syntax  | memcached_socket_keepalive on   off;       |
|---------|--|
| Default | <pre>memcached_socket_keepalive off;</pre> |
| Context | http, server, location                     |

Configures the "TCP keepalive" behavior for outgoing connections to a memcached server.

| 11 11 | By default, the operating system's settings are in effect for the socket. |
|-------|---|
| on    | The SO_KEEPALIVE socket option is turned on for the socket.               |



#### **Built-in Variables**

#### \$memcached\_key

Defines a key for obtaining response from a memcached server.

#### Mirror

The module implements mirroring of an original request by creating background mirror subrequests. Responses to mirror subrequests are ignored.

#### **Configuration Example**

```
location / {
    mirror /mirror;
    proxy_pass http://backend;
}
location = /mirror {
    internal;
    proxy_pass http://test_backend$request_uri;
}
```

#### **Directives**

#### mirror

```
Syntax mirror uri \mid off;
Default mirror off;
Context http, server, location
```

Sets the URI to which an original request will be mirrored. Several mirrors can be specified on the same configuration level.

#### mirror\_request\_body

```
Syntax mirror_request_body on | off;
Default mirror_request_body on;
Context http, server, location
```

Indicates whether the client request body is mirrored. When enabled, the client request body will be read prior to creating mirror subrequests. In this case, unbuffered client request body proxying set by the <code>proxy\_request\_buffering</code>, <code>fastcgi\_request\_buffering</code>, <code>scgi\_request\_buffering</code> and <code>uwsgi\_request\_buffering</code> directives will be disabled.

```
location / {
    mirror /mirror;
    mirror_request_body off;
    proxy_pass http://backend;
}

location = /mirror {
    internal;
    proxy_pass http://log_backend;
    proxy_pass_request_body off;
```



```
proxy_set_header Content-Length "";
proxy_set_header X-Original-URI $request_uri;
}
```

#### MP4

The module provides pseudo-streaming server-side support for MP4 files. Such files typically have the .mp4, .m4v, or .m4a filename extensions.

In packages and images from our repos, the module is included in the build.

Pseudo-streaming works in alliance with a compatible media player. The player sends an HTTP request to the server with the start time specified in the query string argument (named simply start and specified in seconds), and the server responds with the stream such that its start position corresponds to the requested time, for example:

```
http://example.com/elephants_dream.mp4?start=238.88
```

This allows performing a random seeking at any time, or starting playback in the middle of the timeline.

To support seeking, H.264-based formats store metadata in a so-called "moov atom". It is a part of the file that holds the index information for the whole file.

To start playback, the player first needs to read metadata. This is done by sending a special request with the start=0 argument. A lot of encoding software insert the metadata at the end of the file. This is suboptimal for pseudo-streaming, because the player has to download the entire file before starting playback. If the metadata are located at the beginning of the file, it is enough for Angie to simply start sending back the file contents. If the metadata are located at the end of the file, Angie must read the entire file and prepare a new stream so that the metadata come before the media data. This involves some CPU, memory, and disk I/O overhead, so it is a good idea to prepare an original file for pseudo-streaming in advance, rather than having Angie do this on every such request.

The module also supports the end argument of an HTTP request which sets the end point of playback. The end argument can be specified with the start argument or separately:

```
http://example.com/elephants_dream.mp4?start=238.88&end=555.55
```

For a matching request with a non-zero start or end argument, Angie will read the metadata from the file, prepare the stream with the requested time range, and send it to the client. This has the same overhead as described above.

If the start argument points to a non-key video frame, the beginning of such video will be broken. To fix this issue, the video *can* be prepended with the key frame before start point and with all intermediate frames between them. These frames will be hidden from playback using an edit list.

If a matching request does not include the **start** and **end** arguments, there is no overhead, and the file is sent simply as a static resource. Some players also support byte-range requests, and thus do not require this module.

When building from the source code, this module isn't built by default; it should be enabled with the --with-http\_mp4\_module build option.

```
Caution

If a third-party mp4 module was previously used, it should be disabled.
```

A similar pseudo-streaming support for FLV files is provided by the FLV.



#### **Configuration Example**

```
set_real_ip_from 192.168.1.0/24;
set_real_ip_from 192.168.2.1;
set_real_ip_from 2001:0db8::/32;
real_ip_header X-Forwarded-For;
real_ip_recursive on;
```

#### **Directives**

#### mp4

| Syntax  | mp4;     |
|---------|----------|
| Default | _        |
| Context | location |

Turns on module processing in a surrounding location.

## mp4 buffer size

| Syntax  | mp4_buffer_size $size;$          |
|---------|----------------------------------|
| Default | <pre>mp4_buffer_size 512K;</pre> |
| Context | http, server, location           |

Sets the initial size of the buffer used for processing MP4 files.

## mp4\_max\_buffer\_size

| Syntax  | mp4_max_buffer_size $size;$         |
|---------|-------------------------------------|
| Default | <pre>mp4_max_buffer_size 10M;</pre> |
| Context | http, server, location              |

During metadata processing, a larger buffer may become necessary. Its size cannot exceed the specified size, or else Angie will return the 500 (Internal Server Error) server error, and log the following message:

"/some/movie/file.mp4"mp4 moov atom is too large: 12583268, you may want to increase mp4\_max\_buffer\_size

## mp4\_limit\_rate

```
Syntax mp4_limit_rate on | off | factor;
Default mp4_limit_rate off;
Context http, server, location
```

Rate-limits the transfer of the requested MP4 file to the client. To calculate the limit, the *factor* is multiplied by the average bitrate of the file.

- The off value disables rate limiting.
- ullet The on value sets a factor of 1.1.
- The limit is applied after reaching the value set by  $mp4\_limit\_rate\_after$ .

The requests are rate limited individually: if the client opens two connections, the resulting rate doubles. In this regard, consider using  $limit\_conn$  and accompanying directives.



## mp4\_limit\_rate\_after

| Syntax  | ${\tt mp4\_limit\_rate\_after}\ time;$ |
|---------|--|
| Default | <pre>mp4_limit_rate_after 60s;</pre>   |
| Context | http, server, location                 |

Sets (in terms of playback time) the amount of mediadata transferred that triggers the rate limit set by mp4 limit rate.

## mp4\_start\_key\_frame

| Syntax  | <pre>mp4_start_key_frame on   off;</pre> |
|---------|--|
| Default | <pre>mp4_start_key_frame off;</pre>      |
| Context | http, server, location                   |

Forces output video to always start with a key video frame. If the start argument does not point to a key frame, initial frames are hidden using an mp4 edit list. Edit lists are supported by major players and browsers such as Chrome, Safari, QuickTime and ffmpeg, partially supported by Firefox.

#### Perl

The module is used to implement location and variable handlers in Perl and insert Perl calls into SSI.

When building from the source code, this module isn't built by default; it should be enabled with the --with-http\_perl\_module build option.

In our repositories, the module is built dynamically and is available as a separate package named angie-module-perl or angie-pro-module-perl.

# Important

This module requires Perl version 5.6.1 or higher. The C compiler should be compatible with the one used to build Perl.

#### **Known Issues**

The module is experimental, caveat emptor applies.

In order for Perl to recompile the modified modules during reconfiguration, it should be built with the -Dusemultiplicity=yes or -Dusethreads=yes parameters. Also, to make Perl leak less memory at run time, it should be built with the -Dusemymalloc=no parameter. To check the values of these parameters in an already built Perl (preferred values are specified in the example), run:

```
$ perl -V:usemultiplicity -V:usemymalloc
usemultiplicity='define';
usemymalloc='n';
```

Note that after rebuilding Perl with the new -Dusemultiplicity=yes or -Dusethreads=yes parameters, all binary Perl modules will have to be rebuilt as well — they will just stop working with the new Perl.

There is a possibility that the main process and then worker processes will grow in size after every reconfiguration. If the main process grows to an unacceptable size, the *live upgrade* procedure can be applied without changing the executable file.

While the Perl module is performing a long-running operation, such as resolving a domain name, connecting to another server, or querying a database, other requests assigned to the current worker process will not be processed. It is thus recommended to perform only such operations that have predictable and short execution time, such as accessing the local file system.



#### **Configuration Example**

```
http {
    perl_modules perl/lib;
    perl_require hello.pm;
   perl_set $msie6 '
        sub {
            my $r = shift;
            my $ua = $r->header_in("User-Agent");
            return "" if $ua =~ /Opera/;
            return "1" if $ua =~ / MSIE [6-9]\.\d+/;
            return "";
        }
    ١;
    server {
        location / {
            perl hello::handler;
        }
    }
```

The perl/lib/hello.pm module:

```
package hello;
use nginx;
sub handler {
    my $r = shift;

    $r->send_http_header("text/html");
    return OK if $r->header_only;

    $r->print("hello!\n<br/>");

    if (-f $r->filename or -d _) {
        $r->print($r->uri, " exists!\n");
    }

    return OK;
}

1;
__END__
```

#### **Directives**

perl

```
Syntax perl module :: function | 'sub { ... }';

Default —
Context location, limit_except
```



Sets a Perl handler for the given location.

## perl modules

| Syntax  | perl_modules path; |
|---------|--------------------|
| Default | _                  |
| Context | http               |

Sets an additional path for Perl modules.

#### perl require

| Syntax  | perl_require $module;$ |
|---------|------------------------|
| Default | _                      |
| Context | http                   |

Defines the name of a module that will be loaded during each reconfiguration. Several *perl\_require* directives can be present.

#### perl set

| Syntax  | perl_set \$variable module :: function   'sub { }'; |
|---------|---|
| Default | _   |
| Context | http  |

Installs a Perl handler for the specified variable.

#### Calling Perl from SSI

An SSI command calling Perl has the following format:

```
<!--# perl sub="module::function" arg="parameter1" arg="parameter2" ...
-->
```

## The \$r Request Object Methods

#### \$r->args

Returns request arguments.

## \$r->filename

Returns a filename corresponding to the request URI.

```
$r->has_request_body (handler)
```

Returns 0 if there is no body in a request. If there is a body, the specified handler is set for the request and 1 is returned. After reading the request body, Angie will call the specified handler. Note that the handler function should be passed by reference. Example:

```
package hello;
use nginx;
```



```
sub handler {
    my $r = shift;
    if ($r->request_method ne "POST") {
        return DECLINED;
    if ($r->has_request_body(\&post)) {
        return OK;
    }
    return HTTP_BAD_REQUEST;
}
sub post {
    my $r = shift;
    $r->send_http_header;
    $r->print("request_body: \"", $r->request_body, "\"<br/>");
    $r->print("request_body_file: \"", $r->request_body_file, "\"<br/>\n");
    return OK;
}
1;
__END__
```

#### \$r->allow\_ranges

Enables the use of byte ranges when sending responses.

```
$r->discard_request_body
```

Instructs Angie to discard the request body.

```
$r->header_in (field)
```

Returns the value of the specified client request header field.

```
$r->header_only
```

Determines whether the whole response or only its header should be sent to the client.

```
$r->header_out (field, value)
```

Sets a value for the specified response header field.

```
$r->internal_redirect (uri)
```

Does an internal redirect to the specified uri. An actual redirect happens after the Perl handler execution is completed. Method accepts escaped URIs and supports redirections to *named locations*.



## \$r->log\_error (errno, message)

Writes the specified message into the *error\_log*. If *errno* is non-zero, an error code and its description will be appended to the message.

```
$r->print (text, ...)
```

Passes data to a client.

#### \$r->request\_body

Returns the client request body if it has not been written to a temporary file. To ensure that the client request body is in memory, its size should be limited by  $client\_max\_body\_size$ , and a sufficient buffer size should be set using  $client\_body\_buffer\_size$ .

#### \$r->request\_body\_file

Returns the name of the file with the client request body. After the processing, the file should be removed. To always write a request body to a file, *client body in file only* should be enabled.

#### \$r->request\_method

Returns the client request HTTP method.

#### \$r->remote\_addr

Returns the client IP address.

#### \$r->flush

Immediately sends data to the client.

```
$r->sendfile (name [, offset [, length ]])
```

Sends the specified file content to the client. Optional parameters specify the initial offset and length of the data to be transmitted. The actual data transmission happens after the Perl handler has completed.

```
$r->send_http_header ([type])
```

Sends the response header to the client. The optional type parameter sets the value of the "Content-Type" response header field. If the value is an empty string, the "Content-Type" header field will not be sent.

```
$r->status (code)
```

Sets a response code.

## \$r->sleep (milliseconds, handler)

Sets the specified handler and stops request processing for the specified time. In the meantime, Angie continues to process other requests. After the specified time has elapsed, Angie will call the installed handler. Note that the handler function should be passed by reference. In order to pass data between handlers, \$r->variable() should be used. Example:

```
package hello;
use nginx;
```



```
sub handler {
    my $r = shift;

    $r->discard_request_body;
    $r->variable("var", "OK");
    $r->sleep(1000, \&next);

    return OK;
}

sub next {
    my $r = shift;

    $r->send_http_header;
    $r->print($r->variable("var"));

    return OK;
}

1;
__END__
```

## \$r->unescape (text)

Decodes a text encoded in the "%XX" form.

#### \$r->uri

Returns a request URI.

#### \$r->variable (name [, value ])

Returns or sets the value of the specified variable. Variables are local to each request.

# Prometheus

Collects the Angie metrics, using configuration-defined templates, and returns the template-generated metrics in the Prometheus format.

# ▲ Attention

To collect these metrics, enable a shared memory zone in appropriate contexts using:

- the zone directive in an http upstream or a stream upstream;
- the status zone directive;
- the status\_zone parameter in the resolver directive.

#### **Configuration Example**

Three metrics that collect request statistics for a server's shared memory zones, combined into the custom template and exposed at /p8s:

```
http {
    prometheus_template custom {
```



```
'angie_http_server_zones_requests_total{zone="$1"}' $p8s_value
            path=~^/http/server_zones/([^/]+)/requests/total$
            type=counter;
        'angie_http_server_zones_requests_processing{zone="$1"}' $p8s_value
            path=~^/http/server_zones/([^/]+)/requests/processing$
            type=gauge;
        'angie_http_server_zones_requests_discarded{zone="$1"}' $p8s_value
            path=~^/http/server_zones/([^/]+)/requests/discarded$
            type=counter;
   }
    # ...
    server {
        listen 80;
        location =/p8s {
            prometheus custom;
        # ...
   }
}
```

Angie has a supplementary prometheus\_all.conf file that defines a number of common metrics to be used as the all template:

# File Contents (Angie)

```
prometheus_template all {
angie_connections_accepted $p8s_value
    path=/connections/accepted
    type=counter
    'help=The total number of accepted client connections.';
angie_connections_dropped $p8s_value
   path=/connections/dropped
    type=counter
    'help=The total number of dropped client connections.';
angie_connections_active $p8s_value
   path=/connections/active
    type=gauge
    'help=The current number of active client connections.';
angie_connections_idle $p8s_value
    path=/connections/idle
    type=gauge
    'help=The current number of idle client connections.';
'angie_slabs_pages_used{zone="$1"}' $p8s_value
```



```
path=~^/slabs/([^/]+)/pages/used$
   type=gauge
   'help=The number of currently used memory pages in a slab zone.';
'angie_slabs_pages_free{zone="$1"}' $p8s_value
   path=~^/slabs/([^/]+)/pages/free$
   type=gauge
   'help=The number of currently free memory pages in a slab zone.';
'angie_slabs_pages_slots_used{zone="$1",size="$2"}' $p8s_value
   path=~^/slabs/([^/]+)/slots/([^/]+)/used$
   type=gauge
   'help=The number of currently used memory slots of a specific size in a slab zone.
'angie_slabs_pages_slots_free{zone="$1",size="$2"}' $p8s_value
   path=~^/slabs/([^/]+)/slots/([^/]+)/free$
   type=gauge
   'help=The number of currently free memory slots of a specific size in a slab zone.
';
'angie_slabs_pages_slots_reqs{zone="$1",size="$2"}' $p8s_value
   path=~^/slabs/([^/]+)/slots/([^/]+)/reqs$
   type=counter
   'help=The total number of attempts to allocate a memory slot of a specific sizeu
→in a slab zone.';
'angie_slabs_pages_slots_fails{zone="$1",size="$2"}' $p8s_value
   path=~^/slabs/([^/]+)/slots/([^/]+)/fails$
   'help=The number of unsuccessful attempts to allocate a memory slot of a specific
⇒size in a slab zone.';
'angie_resolvers_queries{zone="$1",type="$2"}' $p8s_value
   type=counter
   'help=The number of queries of a specific type to resolve in a resolver zone.';
'angie_resolvers_sent{zone="$1",type="$2"}' $p8s_value
   path=~^/resolvers/([^/]+)/sent/([^/]+)$
   type=counter
   'help=The number of sent DNS queries of a specific type to resolve in a resolver
⇒zone.';
'angie_resolvers_responses{zone="$1",status="$2"}' $p8s_value
   path=~^/resolvers/([^/]+)/responses/([^/]+)$
   type=counter
   'help=The number of resolution results with a specific status in a resolver zone.
'angie_http_server_zones_ssl_handshaked{zone="$1"}' $p8s_value
   path=~^/http/server_zones/([^/]+)/ssl/handshaked$
   type=counter
   'help=The total number of successful SSL handshakes in an HTTP server zone.';
```



```
'angie_http_server_zones_ssl_reuses{zone="$1"}' $p8s_value
   path=~^/http/server_zones/([^/]+)/ssl/reuses$
   type=counter
   'help=The total number of session reuses during SSL handshakes in an HTTP server
⇒zone.';
'angie_http_server_zones_ssl_timedout{zone="$1"}' $p8s_value
   path=~^/http/server_zones/([^/]+)/ssl/timedout$
   type=counter
   'help=The total number of timed-out SSL handshakes in an HTTP server zone.';
'angie_http_server_zones_ssl_failed{zone="$1"}' $p8s_value
   path=~^/http/server_zones/([^/]+)/ssl/failed$
   type=counter
   'help=The total number of failed SSL handshakes in an HTTP server zone.';
'angie_http_server_zones_requests_total{zone="$1"}' $p8s_value
   path=~^/http/server_zones/([^/]+)/requests/total$
   type=counter
   'help=The total number of client requests received in an HTTP server zone.';
'angie_http_server_zones_requests_processing{zone="$1"}' $p8s_value
   path=~^/http/server_zones/([^/]+)/requests/processing$
   type=gauge
   'help=The number of client requests currently being processed in an HTTP server
⇒zone.';
'angie_http_server_zones_requests_discarded{zone="$1"}' $p8s_value
   path=~^/http/server_zones/([^/]+)/requests/discarded$
   type=counter
   'help=The total number of client requests completed in an HTTP server zoneu
→without sending a response.';
'angie_http_server_zones_responses{zone="$1",code="$2"}' $p8s_value
   path=~^/http/server_zones/([^/]+)/responses/([^/]+)$
   type=counter
   'help=The number of responses with a specific status in an HTTP server zone.';
'angie_http_server_zones_data_received{zone="$1"}' $p8s_value
   path=~^/http/server_zones/([^/]+)/data/received$
   type=counter
   'help=The total number of bytes received from clients in an HTTP server zone.';
'angie_http_server_zones_data_sent{zone="$1"}' $p8s_value
   path=~^/http/server_zones/([^/]+)/data/sent$
   type=counter
   'help=The total number of bytes sent to clients in an HTTP server zone.';
'angie_http_location_zones_requests_total{zone="$1"}' $p8s_value
   path=~^/http/location_zones/([^/]+)/requests/total$
   type=counter
   'help=The total number of client requests in an HTTP location zone.';
```



```
'angie_http_location_zones_requests_discarded{zone="$1"}' $p8s_value
   path=~^/http/location_zones/([^/]+)/requests/discarded$
   type=counter
   'help=The total number of client requests completed in an HTTP location zone
⇒without sending a response.';
'angie_http_location_zones_responses{zone="$1",code="$2"}' $p8s_value
   path=~^/http/location_zones/([^/]+)/responses/([^/]+)$
   type=counter
   'help=The number of responses with a specific status in an HTTP location zone.';
'angie_http_location_zones_data_received{zone="$1"}' $p8s_value
   path=~^/http/location_zones/([^/]+)/data/received$
   type=counter
   'help=The total number of bytes received from clients in an HTTP location zone.';
'angie_http_location_zones_data_sent{zone="$1"}' $p8s_value
   path=~^/http/location_zones/([^/]+)/data/sent$
   type=counter
   'help=The total number of bytes sent to clients in an HTTP location zone.';
'angie_http_upstreams_peers_state{upstream="$1",peer="$2"}' $p8st_all_ups_state
   path=~^/http/upstreams/([^/]+)/peers/([^/]+)/state$
   type=gauge
   'help=The current state of an upstream peer in "HTTP": 1 - up, 2 - down, 3 -⊔
→unavailable, or 4 - recovering.';
'angie_http_upstreams_peers_selected_current{upstream="$1",peer="$2"}' $p8s_value
   path=~^/http/upstreams/([^/]+)/peers/([^/]+)/selected/current$
   'help=The number of requests currently being processed by an upstream peer in
→"HTTP".';
'angie_http_upstreams_peers_selected_total{upstream="$1",peer="$2"}' $p8s_value
   path=^{-}/http/upstreams/([^/]+)/peers/([^/]+)/selected/total
   type=counter
   'help=The total number of attempts to use an upstream peer in "HTTP".';
'angie_http_upstreams_peers_responses{upstream="$1",peer="$2",code="$3"} ' $p8s_value
   path=~^/http/upstreams/([^/]+)/peers/([^/]+)/responses/([^/]+)$
   type=counter
   'help=The number of responses with a specific status received from an upstreamu
→peer in "HTTP".';
'angie_http_upstreams_peers_data_sent{upstream="$1",peer="$2"}' $p8s_value
   path=~^/http/upstreams/([^/]+)/peers/([^/]+)/data/sent$
   type=counter
   'help=The total number of bytes sent to an upstream peer in "HTTP".';
'angie_http_upstreams_peers_data_received{upstream="$1",peer="$2"}' $p8s_value
```



```
path=~^/http/upstreams/([^/]+)/peers/([^/]+)/data/received$
   type=counter
   'help=The total number of bytes received from an upstream peer in "HTTP".';
'angie_http_upstreams_peers_health_fails{upstream="$1",peer="$2"}' $p8s_value
   path=~^/http/upstreams/([^/]+)/peers/([^/]+)/health/fails$
   type=counter
   'help=The total number of unsuccessful attempts to communicate with an upstreamu
→peer in "HTTP".';
'angie_http_upstreams_peers_health_unavailable{upstream="$1",peer="$2"}' $p8s_value
   path=~^/http/upstreams/([^/]+)/peers/([^/]+)/health/unavailable$
   type=counter
   'help=The number of times when an upstream peer in "HTTP" became "unavailable"
→due to reaching the max_fails limit.';
'angie_http_upstreams_peers_health_downtime{upstream="$1",peer="$2"}' $p8s_value
   path=~^/http/upstreams/([^/]+)/peers/([^/]+)/health/downtime$
   type=counter
   'help=The total time (in milliseconds) that an upstream peer in "HTTP" was
'angie_http_upstreams_keepalive{upstream="$1"}' $p8s_value
   path=~^/http/upstreams/([^/]+)/keepalive$
   'help=The number of currently cached keepalive connections for an HTTP upstream.';
'angie_http_caches_responses{zone="$1",status="$2"}' $p8s_value
   path=~^/http/caches/([^/]+)/([^/]+)/responses$
   type=counter
   'help=The total number of responses processed in an HTTP cache zone with a
→specific cache status.';
'angie_http_caches_bytes{zone="$1",status="$2"}' $p8s_value
   path = ^{\sim}/http/caches/([^{/}]+)/([^{/}]+)/bytes 
   type=counter
   'help=The total number of bytes processed in an HTTP cache zone with a specificu
→cache status.';
path=~^/http/caches/([^/]+)/([^/]+)/responses_written$
   type=counter
   'help=The total number of responses written to an HTTP cache zone with a specificu
→cache status.';
'angie_http_caches_bytes_written{zone="$1",status="$2"}' $p8s_value
   path=~^/http/caches/([^/]+)/([^/]+)/bytes_written$
   type=counter
   'help=The total number of bytes written to an HTTP cache zone with a specific
→cache status.';
'angie_http_caches_size{zone="$1"}' $p8s_value
   path=~^/http/caches/([^/]+)/size$
```



```
type=gauge
   'help=The current size (in bytes) of cached responses in an HTTP cache zone.';
'angie_http_caches_shards_size{zone="$1",path="$2"}' $p8s_value
   path=~^/http/caches/([^/]+)/shards/([^/]+)/size$
   type=gauge
   'help=The current size (in bytes) of cached responses in a shard path of an HTTPu
→cache zone.';
'angie_http_limit_conns{zone="$1",status="$2"}' $p8s_value
   path=~^/http/limit_conns/([^/]+)/([^/]+)$
   type=counter
   'help=The number of requests processed by an HTTP limit_conn zone with a specificu
→result.';
'angie_http_limit_reqs{zone="$1",status="$2"}' $p8s_value
   path=~^/http/limit_reqs/([^/]+)/([^/]+)$
   type=counter
   'help=The number of requests processed by an HTTP limit_reqs zone with a specific⊔
→result.';
'angie_stream_server_zones_ssl_handshaked{zone="$1"}' $p8s_value
   path=~^/stream/server_zones/([^/]+)/ssl/handshaked$
   type=counter
   'help=The total number of successful SSL handshakes in a stream server zone.';
'angie_stream_server_zones_ssl_reuses{zone="$1"}' $p8s_value
   path=~^/stream/server_zones/([^/]+)/ssl/reuses$
   type=counter
   'help=The total number of session reuses during SSL handshakes in a stream server
⇒zone.';
'angie_stream_server_zones_ssl_timedout{zone="$1"}' $p8s_value
   path=~^/stream/server_zones/([^/]+)/ssl/timedout$
   type=counter
   'help=The total number of timed-out SSL handshakes in a stream server zone.';
'angie_stream_server_zones_ssl_failed{zone="$1"}' $p8s_value
   path=~^/stream/server_zones/([^/]+)/ssl/failed$
   type=counter
   'help=The total number of failed SSL handshakes in a stream server zone.';
'angie_stream_server_zones_connections_total{zone="$1"}' $p8s_value
   path=~^/stream/server_zones/([^/]+)/connections/total$
   type=counter
   'help=The total number of client connections received in a stream server zone.';
'angie_stream_server_zones_connections_processing{zone="$1"}' $p8s_value
   path=~^/stream/server_zones/([^/]+)/connections/processing$
   'help=The number of client connections currently being processed in a stream_
⇒server zone.';
```



```
path=~^/stream/server_zones/([^/]+)/connections/discarded$
   type=counter
   'help=The total number of client connections completed in a stream server zone
⇒without establishing a session.';
'angie_stream_server_zones_connections_passed{zone="$1"}' $p8s_value
   path=~^/stream/server_zones/([^/]+)/connections/passed$
   type=counter
   'help=The total number of client connections in a stream server zone passed for
→handling to a different listening socket.';
'angie_stream_server_zones_sessions{zone="$1",status="$2"}' $p8s_value
   path=~^/stream/server_zones/([^/]+)/sessions/([^/]+)$
   type=counter
   'help=The number of sessions finished with a specific status in a stream server
⇔zone.';
'angie_stream_server_zones_data_received{zone="$1"}' $p8s_value
   path=~^/stream/server_zones/([^/]+)/data/received$
   type=counter
   'help=The total number of bytes received from clients in a stream server zone.';
'angie_stream_server_zones_data_sent{zone="$1"}' $p8s_value
   path=~^/stream/server_zones/([^/]+)/data/sent$
   type=counter
   'help=The total number of bytes sent to clients in a stream server zone.';
'angie_stream_upstreams_peers_state{upstream="$1",peer="$2"}' $p8st_all_ups_state
   path=~^/stream/upstreams/([^/]+)/peers/([^/]+)/state$
   type=gauge
   'help=The current state of an upstream peer in "stream": 1 - up, 2 - down, 3 -
→unavailable, or 4 - recovering.';
'angie_stream_upstreams_peers_selected_current{upstream="$1",peer="$2"} ' $p8s_value
   path=~^/stream/upstreams/([^/]+)/peers/([^/]+)/selected/current$
   type=gauge
   'help=The number of sessions currently being processed by an upstream peer in
→"stream".';
'angie_stream_upstreams_peers_selected_total{upstream="$1",peer="$2"}' $p8s_value
   path=~^/stream/upstreams/([^/]+)/peers/([^/]+)/selected/total$
   type=counter
   'help=The total number of attempts to use an upstream peer in "stream".';
'angie_stream_upstreams_peers_data_sent{upstream="$1",peer="$2"}' $p8s_value
   path=~^/stream/upstreams/([^/]+)/peers/([^/]+)/data/sent$
   type=counter
   'help=The total number of bytes sent to an upstream peer in "stream".';
'angie_stream_upstreams_peers_data_received{upstream="$1",peer="$2"}' $p8s_value
   path=~^/stream/upstreams/([^/]+)/peers/([^/]+)/data/received$
```



```
type=counter
    'help=The total number of bytes received from an upstream peer in "stream".';
'angie_stream_upstreams_peers_health_fails{upstream="$1",peer="$2"}' $p8s_value
   path=~^/stream/upstreams/([^/]+)/peers/([^/]+)/health/fails$
    type=counter
    'help=The total number of unsuccessful attempts to communicate with an upstreamu
→peer in "stream".';
'angie_stream_upstreams_peers_health_unavailable{upstream="$1",peer="$2"}' $p8s_value
   path=~^/stream/upstreams/([^/]+)/peers/([^/]+)/health/unavailable$
   type=counter
    'help=The number of times when an upstream peer in "stream" became "unavailable"
→due to reaching the max_fails limit.';
'angie_stream_upstreams_peers_health_downtime{upstream="$1",peer="$2"}' $p8s_value
   path=~^/stream/upstreams/([^/]+)/peers/([^/]+)/health/downtime$
   type=counter
    'help=The total time (in milliseconds) that an upstream peer in "stream" was
}
map $p8s_value $p8st_all_ups_state {
   volatile;
   "up"
                  1;
   "down"
                  2;
   "unavailable" 3;
   "recovering"
#
   "unhealthy"
                 5;
#
    "checking"
    "draining"
#
                  7;
   "busy"
                  8;
   default
                 0;
}
```

#### File Contents (Angie PRO)

```
prometheus_template all {
    angie_connections_accepted $p8s_value
        path=/connections/accepted
        type=counter
        'help=The total number of accepted client connections.';

angie_connections_dropped $p8s_value
        path=/connections/dropped
        type=counter
        'help=The total number of dropped client connections.';

angie_connections_active $p8s_value
        path=/connections/active
        type=gauge
        'help=The current number of active client connections.';

angie_connections_idle $p8s_value
        path=/connections/idle
```



```
type=gauge
   'help=The current number of idle client connections.';
'angie_slabs_pages_used{zone="$1"}' $p8s_value
   path=~^/slabs/([^/]+)/pages/used$
   type=gauge
   'help=The number of currently used memory pages in a slab zone.';
'angie_slabs_pages_free{zone="$1"}' $p8s_value
   path=~^/slabs/([^/]+)/pages/free$
   type=gauge
   'help=The number of currently free memory pages in a slab zone.';
'angie_slabs_pages_slots_used{zone="$1",size="$2"}' $p8s_value
   path=~^/slabs/([^/]+)/slots/([^/]+)/used$
   type=gauge
   'help=The number of currently used memory slots of a specific size in a slab zone.
'angie_slabs_pages_slots_free{zone="$1",size="$2"}' $p8s_value
   path=~^/slabs/([^/]+)/slots/([^/]+)/free$
   type=gauge
   'help=The number of currently free memory slots of a specific size in a slab zone.
'angie_slabs_pages_slots_reqs{zone="$1",size="$2"}' $p8s_value
   path=~^/slabs/([^/]+)/slots/([^/]+)/reqs$
   type=counter
   'help=The total number of attempts to allocate a memory slot of a specific sizeu
→in a slab zone.';
'angie_slabs_pages_slots_fails{zone="$1",size="$2"}' $p8s_value
   path=~^/slabs/([^/]+)/slots/([^/]+)/fails$
   type=counter
   \hbox{'help=The number of unsuccessful attempts to allocate a memory slot of a specific}_{\hbox{\it u}}
⇒size in a slab zone.';
'angie_resolvers_queries{zone="$1",type="$2"}' $p8s_value
   path=~^/resolvers/([^/]+)/queries/([^/]+)$
   type=counter
   'help=The number of queries of a specific type to resolve in a resolver zone.';
'angie_resolvers_sent{zone="$1",type="$2"}' $p8s_value
   path=~^/resolvers/([^/]+)/sent/([^/]+)$
   type=counter
   'help=The number of sent DNS queries of a specific type to resolve in a resolver
⇒zone.';
'angie_resolvers_responses{zone="$1",status="$2"}' $p8s_value
   path=~^/resolvers/([^/]+)/responses/([^/]+)$
   type=counter
   'help=The number of resolution results with a specific status in a resolver zone.
';
```



```
'angie_http_server_zones_ssl_handshaked{zone="$1"}' $p8s_value
   path=~^/http/server_zones/([^/]+)/ssl/handshaked$
   type=counter
   'help=The total number of successful SSL handshakes in an HTTP server zone.';
'angie_http_server_zones_ssl_reuses{zone="$1"}' $p8s_value
   path=~^/http/server_zones/([^/]+)/ssl/reuses$
   type=counter
   'help=The total number of session reuses during SSL handshakes in an HTTP server
⇒zone.';
'angie_http_server_zones_ssl_timedout{zone="$1"}' $p8s_value
   path=~^/http/server_zones/([^/]+)/ssl/timedout$
   type=counter
   'help=The total number of timed-out SSL handshakes in an HTTP server zone.';
'angie_http_server_zones_ssl_failed{zone="$1"}' $p8s_value
   path=~^/http/server_zones/([^/]+)/ssl/failed$
   type=counter
   'help=The total number of failed SSL handshakes in an HTTP server zone.';
'angie_http_server_zones_requests_total{zone="$1"}' $p8s_value
   path=~^/http/server_zones/([^/]+)/requests/total$
   type=counter
   'help=The total number of client requests received in an HTTP server zone.';
'angie_http_server_zones_requests_processing{zone="$1"}' $p8s_value
   path=~^/http/server_zones/([^/]+)/requests/processing$
   'help=The number of client requests currently being processed in an HTTP server
⇒zone.';
'angie_http_server_zones_requests_discarded{zone="$1"}' $p8s_value
   path=~^/http/server_zones/([^/]+)/requests/discarded$
   type=counter
   \verb|'help=The total number of client requests completed in an HTTP server zone_{\sqcup}
→without sending a response.';
'angie_http_server_zones_responses{zone="$1",code="$2"}' $p8s_value
   path=~^/http/server_zones/([^/]+)/responses/([^/]+)$
   tvpe=counter
   'help=The number of responses with a specific status in an HTTP server zone.';
'angie_http_server_zones_data_received{zone="$1"}' $p8s_value
   path=~^/http/server_zones/([^/]+)/data/received$
   type=counter
   'help=The total number of bytes received from clients in an HTTP server zone.';
'angie_http_server_zones_data_sent{zone="$1"}' $p8s_value
   path=~^/http/server_zones/([^/]+)/data/sent$
   type=counter
   'help=The total number of bytes sent to clients in an HTTP server zone.';
```



```
'angie_http_location_zones_requests_total{zone="$1"}' $p8s_value
   path=~^/http/location_zones/([^/]+)/requests/total$
   type=counter
   'help=The total number of client requests in an HTTP location zone.';
'angie_http_location_zones_requests_discarded{zone="$1"}' $p8s_value
   path=~^/http/location_zones/([^/]+)/requests/discarded$
   type=counter
   'help=The total number of client requests completed in an HTTP location zoneu
⇒without sending a response.';
'angie_http_location_zones_responses{zone="$1",code="$2"}' $p8s_value
   path=~^/http/location_zones/([^/]+)/responses/([^/]+)$
   type=counter
   'help=The number of responses with a specific status in an HTTP location zone.';
'angie_http_location_zones_data_received{zone="$1"}' $p8s_value
   path=~^/http/location_zones/([^/]+)/data/received$
   type=counter
   'help=The total number of bytes received from clients in an HTTP location zone.';
'angie_http_location_zones_data_sent{zone="$1"}' $p8s_value
   path=~^/http/location_zones/([^/]+)/data/sent$
   type=counter
   'help=The total number of bytes sent to clients in an HTTP location zone.';
'angie_http_upstreams_peers_backup{upstream="$1",peer="$2"}' $p8st_all_ups_backup
   path=~^/http/upstreams/([^/]+)/peers/([^/]+)/backup$
   type=gauge
   'help=The HTTP upstream peer backup group level.';
'angie_http_upstreams_peers_state{upstream="$1",peer="$2"}' $p8st_all_ups_state
   path=~^/http/upstreams/([^/]+)/peers/([^/]+)/state$
   type=gauge
   'help=The current state of an upstream peer in "HTTP": 1 - up, 2 - down, 3 -u
→unavailable, 4 - recovering, 5 - unhealthy, 6 - checking, or 7 - draining.';
'angie_http_upstreams_peers_selected_current{upstream="$1",peer="$2"}' $p8s_value
   path=~^/http/upstreams/([^/]+)/peers/([^/]+)/selected/current$
   'help=The number of requests currently being processed by an upstream peer in
→"HTTP".';
'angie_http_upstreams_peers_selected_total{upstream="$1",peer="$2"}' $p8s_value
   path=~^/http/upstreams/([^/]+)/peers/([^/]+)/selected/total$
   type=counter
   'help=The total number of attempts to use an upstream peer in "HTTP".';
path=^{\hfill}/http/upstreams/([^/]+)/peers/([^/]+)/responses/([^/]+)
```



```
type=counter
   'help=The number of responses with a specific status received from an upstream_
→peer in "HTTP".';
'angie_http_upstreams_peers_data_sent{upstream="$1",peer="$2"}' $p8s_value
   path=~^/http/upstreams/([^/]+)/peers/([^/]+)/data/sent$
   type=counter
   'help=The total number of bytes sent to an upstream peer in "HTTP".';
'angie_http_upstreams_peers_data_received{upstream="$1",peer="$2"}' $p8s_value
   path=~^/http/upstreams/([^/]+)/peers/([^/]+)/data/received$
   type=counter
   'help=The total number of bytes received from an upstream peer in "HTTP".';
'angie_http_upstreams_peers_health_fails{upstream="$1",peer="$2"}' $p8s_value
   path=~^/http/upstreams/([^/]+)/peers/([^/]+)/health/fails$
   type=counter
   'help=The total number of unsuccessful attempts to communicate with an upstreamu
→peer in "HTTP".';
'angie_http_upstreams_peers_health_unavailable{upstream="$1",peer="$2"}' $p8s_value
   path=~^/http/upstreams/([^/]+)/peers/([^/]+)/health/unavailable$
   tvpe=counter
   'help=The number of times when an upstream peer in "HTTP" became "unavailable" u
→due to reaching the max_fails limit.';
'angie_http_upstreams_peers_health_downtime{upstream="$1",peer="$2"}' $p8s_value
   path=~^/http/upstreams/([^/]+)/peers/([^/]+)/health/downtime$
   type=counter
   'help=The total time (in milliseconds) that an upstream peer in "HTTP" was
→"unavailable".';
'angie_http_upstreams_peers_health_header_time{upstream="$1",peer="$2"}' $p8s_value
   path=~^/http/upstreams/([^/]+)/peers/([^/]+)/health/header_time$
   type=gauge
   'help=Average time (in milliseconds) to receive the response headers from anu
→upstream peer in "HTTP".';
'angie_http_upstreams_peers_health_response_time{upstream="$1",peer="$2"}' $p8s_value
   path=~^/http/upstreams/([^/]+)/peers/([^/]+)/health/response_time$
   type=gauge
   'help=Average time (in milliseconds) to receive the complete response from an
→upstream peer in "HTTP".';
'angie_http_upstreams_peers_health_probes_count{upstream="$1",peer="$2"}' $p8s_value
   path=~^/http/upstreams/([^/]+)/peers/([^/]+)/health/probes/count$
   type=counter
   'help=The total number of probes for this peer.';
path=~^/http/upstreams/([^/]+)/peers/([^/]+)/health/probes/fails$
   type=counter
   'help=The total number of failed probes for this peer.';
```



```
'angie_http_upstreams_keepalive{upstream="$1"}' $p8s_value
   path=~^/http/upstreams/([^/]+)/keepalive$
   type=gauge
   'help=The number of currently cached keepalive connections for an HTTP upstream.';
'angie_http_upstreams_backup_switch_active{upstream="$1"}' $p8s_value
   path=~^/http/upstreams/([^/]+)/backup_switch/active$
   type=gauge
   'help=The currently active HTTP upstream servers backup group level.';
'angie_http_upstreams_queue_queued{upstream="$1"}' $p8s_value
   path=~^/http/upstreams/([^/]+)/queue/queued$
   type=counter
   'help=The total number of queued requests for an HTTP upstream.';
'angie_http_upstreams_queue_waiting{upstream="$1"}' $p8s_value
   path=~^/http/upstreams/([^/]+)/queue/waiting$
   type=gauge
   'help=The number of requests currently waiting in an HTTP upstream queue.';
'angie_http_upstreams_queue_dropped{upstream="$1"}' $p8s_value
   path=~^/http/upstreams/([^/]+)/queue/dropped$
   type=counter
   'help=The total number of requests dropped from an HTTP upstream queue becauseu
→the client had prematurely closed the connection.';
'angie_http_upstreams_queue_timedout{upstream="$1"}' $p8s_value
   path=~^/http/upstreams/([^/]+)/queue/timedout$
   type=counter
   'help=The total number of requests timed out from an HTTP upstream queue.';
'angie_http_upstreams_queue_overflows{upstream="$1"}' $p8s_value
   path=~^/http/upstreams/([^/]+)/queue/overflows$
   type=counter
   'help=The total number of requests rejected by an HTTP upstream queue because theu
⇒size limit had been reached.';
'angie_http_caches_responses{zone="$1",status="$2"}' $p8s_value
   path=~^/http/caches/([^/]+)/([^/]+)/responses$
   type=counter
   'help=The total number of responses processed in an HTTP cache zone with a
→specific cache status.';
'angie_http_caches_bytes{zone="$1",status="$2"}' $p8s_value
   path=^{-}/http/caches/([^/]+)/([^/]+)/bytes
   type=counter
   'help=The total number of bytes processed in an HTTP cache zone with a specificu
→cache status.';
'angie_http_caches_responses_written{zone="$1",status="$2"}' $p8s_value
   path=~^/http/caches/([^/]+)/([^/]+)/responses_written$
   type=counter
   \hbox{'help=The total number of responses written to an HTTP cache zone with a specific}_{\hbox{\it u}}
```



```
'angie_http_caches_bytes_written{zone="$1",status="$2"}' $p8s_value
   type=counter
   'help=The total number of bytes written to an HTTP cache zone with a specific⊔
'angie_http_caches_size{zone="$1"}' $p8s_value
   path=~^/http/caches/([^/]+)/size$
   type=gauge
   'help=The current size (in bytes) of cached responses in an HTTP cache zone.';
'angie_http_caches_shards_size{zone="$1",path="$2"}' $p8s_value
   path=~^/http/caches/([^/]+)/shards/([^/]+)/size$
   type=gauge
   'help=The current size (in bytes) of cached responses in a shard path of an HTTPu
'angie_http_limit_conns{zone="$1",status="$2"}' $p8s_value
   path=~^/http/limit_conns/([^/]+)/([^/]+)$
   type=counter
   'help=The number of requests processed by an HTTP limit_conn zone with a specific_
⇔result.';
'angie_http_limit_reqs{zone="$1",status="$2"}' $p8s_value
   path=~^/http/limit_reqs/([^/]+)/([^/]+)$
   type=counter
   'help=The number of requests processed by an HTTP limit_reqs zone with a specificu
⇔result.';
'angie_stream_server_zones_ssl_handshaked{zone="$1"}' $p8s_value
   path=~^/stream/server_zones/([^/]+)/ssl/handshaked$
   type=counter
   'help=The total number of successful SSL handshakes in a stream server zone.';
'angie_stream_server_zones_ssl_reuses{zone="$1"}' $p8s_value
   path=~^/stream/server_zones/([^/]+)/ssl/reuses$
   type=counter
   'help=The total number of session reuses during SSL handshakes in a stream server
⇒zone.';
'angie_stream_server_zones_ssl_timedout{zone="$1"}' $p8s_value
   path=~^/stream/server_zones/([^/]+)/ssl/timedout$
   type=counter
   'help=The total number of timed-out SSL handshakes in a stream server zone.';
'angie_stream_server_zones_ssl_failed{zone="$1"}' $p8s_value
   path=~^/stream/server_zones/([^/]+)/ssl/failed$
   type=counter
   'help=The total number of failed SSL handshakes in a stream server zone.';
'angie_stream_server_zones_connections_total{zone="$1"}' $p8s_value
```



```
path=~^/stream/server_zones/([^/]+)/connections/total$
   type=counter
   'help=The total number of client connections received in a stream server zone.';
'angie_stream_server_zones_connections_processing{zone="$1"}' $p8s_value
   path=~^/stream/server_zones/([^/]+)/connections/processing$
   type=gauge
   'help=The number of client connections currently being processed in a streamu
→server zone.';
'angie_stream_server_zones_connections_discarded{zone="$1"}' $p8s_value
   path=~^/stream/server_zones/([^/]+)/connections/discarded$
   type=counter
   'help=The total number of client connections completed in a stream server zoneu
→without establishing a session.';
'angie_stream_server_zones_connections_passed{zone="$1"}' $p8s_value
   path=~^/stream/server_zones/([^/]+)/connections/passed$
   type=counter
   'help=The total number of client connections in a stream server zone passed for \Box
→handling to a different listening socket.';
'angie_stream_server_zones_sessions{zone="$1",status="$2"}' $p8s_value
   path=~^/stream/server_zones/([^/]+)/sessions/([^/]+)$
   type=counter
   'help=The number of sessions finished with a specific status in a stream server
⇒zone.';
'angie_stream_server_zones_data_received{zone="$1"}' $p8s_value
   path=~^/stream/server_zones/([^/]+)/data/received$
   type=counter
   'help=The total number of bytes received from clients in a stream server zone.';
'angie_stream_server_zones_data_sent{zone="$1"}' $p8s_value
   path=~^/stream/server_zones/([^/]+)/data/sent$
   type=counter
   'help=The total number of bytes sent to clients in a stream server zone.';
'angie_stream_upstreams_peers_backup{upstream="$1",peer="$2"}' $p8st_all_ups_backup
   path=~^/stream/upstreams/([^/]+)/peers/([^/]+)/backup$
   'help=The "stream" upstream peer backup group level.';
'angie_stream_upstreams_peers_state{upstream="$1",peer="$2"}' $p8st_all_ups_state
   path=~^/stream/upstreams/([^/]+)/peers/([^/]+)/state$
   type=gauge
   'help=The current state of an upstream peer in "stream": 1 - up, 2 - down, 3 -u
→unavailable, 4 - recovering, 5 - unhealthy, 6 - checking, or 7 - draining.';
'angie_stream_upstreams_peers_selected_current{upstream="$1",peer="$2"} ' $p8s_value
   path=~^/stream/upstreams/([^/]+)/peers/([^/]+)/selected/current$
   type=gauge
```



```
'help=The number of sessions currently being processed by an upstream peer in
→"stream".';
'angie_stream_upstreams_peers_selected_total{upstream="$1",peer="$2"}' $p8s_value
   path=~^/stream/upstreams/([^/]+)/peers/([^/]+)/selected/total$
   type=counter
   'help=The total number of attempts to use an upstream peer in "stream".';
'angie_stream_upstreams_peers_data_sent{upstream="$1",peer="$2"}' $p8s_value
   path=~^/stream/upstreams/([^/]+)/peers/([^/]+)/data/sent$
   type=counter
   'help=The total number of bytes sent to an upstream peer in "stream".';
'angie_stream_upstreams_peers_data_received{upstream="$1",peer="$2"}' $p8s_value
   path=~^/stream/upstreams/([^/]+)/peers/([^/]+)/data/received$
   type=counter
   'help=The total number of bytes received from an upstream peer in "stream".';
'angie_stream_upstreams_peers_data_pkt_sent{upstream="$1",peer="$2"}' $p8s_value
   path=~^/stream/upstreams/([^/]+)/peers/([^/]+)/data/pkt_sent$
   type=counter
   'help=The total number of packets sent to an upstream peer in "stream".';
'angie_stream_upstreams_peers_data_pkt_received{upstream="$1",peer="$2"}' $p8s_value
   path=~^/stream/upstreams/([^/]+)/peers/([^/]+)/data/pkt_received$
   type=counter
   'help=The total number of packets received from an upstream peer in "stream".';
'angie_stream_upstreams_peers_health_fails{upstream="$1",peer="$2"}' $p8s_value
   path=~^/stream/upstreams/([^/]+)/peers/([^/]+)/health/fails$
   type=counter
   'help=The total number of unsuccessful attempts to communicate with an upstream,
→peer in "stream".';
'angie_stream_upstreams_peers_health_unavailable{upstream="$1",peer="$2"}' $p8s_value
   path=~^/stream/upstreams/([^/]+)/peers/([^/]+)/health/unavailable$
   type=counter
   'help=The number of times when an upstream peer in "stream" became "unavailable" u
→due to reaching the max_fails limit.';
'angie_stream_upstreams_peers_health_downtime{upstream="$1",peer="$2"}' $p8s_value
   path=~^/stream/upstreams/([^/]+)/peers/([^/]+)/health/downtime$
   type=counter
   'help=The total time (in milliseconds) that an upstream peer in "stream" was
→"unavailable".';
'angie_stream_upstreams_peers_health_connect_time{upstream="$1",peer="$2"} ' $p8s_value
   path=~^/stream/upstreams/([^/]+)/peers/([^/]+)/health/connect_time$
   type=gauge
   'help=Average time (in milliseconds) to connect to an upstream peer in "stream".';
'angie_stream_upstreams_peers_health_first_byte_time{upstream="$1",peer="$2"}' $p8s_
   path=~^/stream/upstreams/([^/]+)/peers/([^/]+)/health/first_byte_time$
   type=gauge
```



```
'help=Average time (in milliseconds) to receive the first byte from an upstreamu
→peer in "stream".';
angie_stream_upstreams_peers_health_last_byte_time{upstream="$1",peer="$2"}' $p8s_
   path=~^/stream/upstreams/([^/]+)/peers/([^/]+)/health/last_byte_time$
    type=gauge
    'help=Average time (in milliseconds) of the whole communication session with anu
→upstream peer in "stream".';
'angie_stream_upstreams_peers_health_probes_count{upstream="$1",peer="$2"}' $p8s_value
   path=~^/stream/upstreams/([^/]+)/peers/([^/]+)/health/probes/count$
    type=counter
    'help=The total number of probes for this peer.';
'angie_stream_upstreams_peers_health_probes_fails{upstream="$1",peer="$2"}' $p8s_value
   path=~^/stream/upstreams/([^/]+)/peers/([^/]+)/health/probes/fails$
    type=counter
    'help=The total number of failed probes for this peer.';
}
map $p8s_value $p8st_all_ups_state {
   volatile;
    "up"
                  1;
   "down"
   "unavailable" 3;
   "recovering"
                5;
    "unhealthy"
    "checking"
                  6;
    "draining"
                  7;
    "busy"
                  8;
   default
                 0;
}
map $p8s_value $p8st_all_ups_backup {
   volatile;
    "false"
                   0;
    "true"
                  1;
   default
                  $p8s_value;
}
```

Usage:

```
http {
   include prometheus_all.conf;

#...
server {
   listen 80;
   location =/p8s {
      prometheus all;
   }
}
```



```
# ...
}
}
```

```
$ curl localhost/p8s

# Angie Prometheus template "all"
...
```

#### **Directives**

#### prometheus

| Syntax  | prometheus template; |
|---------|----------------------|
| Default | _                    |
| Context | location             |

Sets the handler template for the location context; the template should be defined by a prometheus\_template directive. Upon a request, this location evaluates and returns the template metrics in the Prometheus format.

```
location =/p8s {
    prometheus custom;
}
```

```
$ curl localhost/p8s

# Angie Prometheus template "custom"
...
```

# $prometheus\_template$

| Syntax  | $\verb prometheus_template  template { }$ |
|---------|---|
| Default | _   |
| Context | http                                      |

Defines a named template of metrics to be collected and exported by Angie; the template should be used with a *prometheus* directive.

#### 1 Note

Angie also has a predefined *all* template that includes a number of common metrics.

The template can contain any number of metrics definitions, each structured as follows: <metric> <variable> [path=<match>] [type=<type>] [help=<help>];



| metric   | Sets the name of the metric to be added to the response in the Prometheus format. Can contain an optional label definition $(\ldots)$ , for example:                               |
|----------|--|
|          | http_requests_total{method="\$1",code="\$2"}   |
|          | Label values can use Angie variables; if match is a regex, you can also use its  |
|          | capture groups with labels. Such variables and groups are evaluated when the metric's <i>variable</i> itself is evaluated.   |
| variable | Sets the name of the variable to be evaluated and added to the response as the metric value. If there's no such variable or the resulting value is empty (""), no metric is added. |

The metric's value is evaluated from the *variable*; upon a successful evaluation, the metric is added to the response, for example:

```
'angie_time{version="$angie_version"}' $msec;
```

```
$ curl localhost/p8s
angie_time{version="1.9.1"} 1695119820.562
```

| path=match | Is matched against all end-to-end metric paths in the /status section of Angie |
|------------|--|
|            | API, allowing to add several metric instances to the response at once.         |

During matching, the paths include the starting slash but not the ending one, for example: /angie/generation; matching is case-insensitive. There are two matching modes:

```
path=exact_match Matches strings character by character.

path="regex_matc" PCRE-based match; can define capture groups to be used with the metric name labels.
```

If match matches a path, the value of the Angie metric at this path is stored in the \$p8s\_value variable, which can be used as variable when path= is set.

When regex match is used, multiple matches can occur; in this case, a metric is added to the response for *each* matching path. With capture groups, this enables series of metrics that share one name but have different labels, for example:

```
'angie_slabs_slots_free{zone="$1",size="$2"}' $p8s_value
   path=~^/slabs/([^/]+)/slots/([^/]+)/free$;
```

This adds a metric for each zone and size combination the configuration contains:

```
angie_slabs_slots_free{zone="one",size="8"} 502
angie_slabs_slots_free{zone="one",size="16"} 249
angie_slabs_slots_free{zone="one",size="32"} 122
angie_slabs_slots_free{zone="one",size="128"} 22
angie_slabs_slots_free{zone="one",size="512"} 4
angie_slabs_slots_free{zone="two",size="8"} 311
...
```

If there are no matches (regardless of the mode), no metric is added.

#### 1 Note

The path= option is available only when Angie is built with the API module.



| both are added to the response with the metric itself without any transformation or validation. | type=type,<br>help=help | Set the metric's type and help string, respectively, using the Prometheus format; both are added to the response with the metric itself without any transformation or validation. |
|---|-------------------------|---|
|---|-------------------------|---|

#### **Built-in Variables**

The http\_prometheus module has a built-in variable that receives its value when the paths from the /status section of Angie API are matched against the match option of a metric defined by the prometheus template directive.

#### \$p8s\_value

If the *match* of a metric defined by *prometheus\_template* matches a path, the value of the Angie metric at this path is stored in the \$p8s\_value variable. It's intended to be used as the *variable* in path=-based metric definitions.

The values of Angie metrics stored in \$p8s\_value may deviate from the requirements of the Prometheus format. In this case, the *map* directive can help transform strings into numbers:

If the Angie metric is Boolean, namely, true or false, the variable is set to "1" or "0", respectively; if the metric is null, the variable is set to "(null)". For date values, the integer epoch format is used.

#### **Proxy**

Allows passing requests to another (proxied) server.

### **Configuration Example**

#### **Directives**

# proxy\_bind

| Syntax  | proxy_bind address [transparent]   off; |
|---------|---|
| Default | _                                       |
| Context | http, server, location                  |

Makes outgoing connections to a proxied server originate from the specified local IP address with an optional port. Parameter value can contain variables. The special value off cancels the effect of the



proxy\_bind directive inherited from the previous configuration level, which allows the system to auto-assign the local IP address and port.

The transparent parameter allows outgoing connections to a proxied server originate from a non-local IP address, for example, from a real IP address of a client:

```
proxy_bind $remote_addr transparent;
```

In order for this parameter to work, it is usually necessary to run Angie worker processes with the superuser privileges. On Linux it is not required as if the transparent parameter is specified, worker processes inherit the  $CAP\_NET\_RAW$  capability from the master process.

# Important

It is necessary to configure kernel routing table to intercept network traffic from the proxied server.

## proxy buffer size

| Syntax  | proxy_buffer_size $size$ ;          |
|---------|-------------------------------------|
| Default | <pre>proxy_buffer_size 4k 8k;</pre> |
| Context | http, server, location              |

Sets the size of the buffer used for reading the first part of the response received from the proxied server. This part usually contains a small response header. By default, the buffer size is equal to one memory page. This is either 4K or 8K, depending on a platform. It can be made smaller, however.

### proxy\_buffering

| Syntax  | proxy_buffering on   off;      |
|---------|--------------------------------|
| Default | <pre>proxy_buffering on;</pre> |
| Context | http, server, location         |

Enables or disables buffering of responses from the proxied server.

| on  | Angie receives a response from the proxied server as soon as possible, saving it into the buffers set by the proxy_buffer_size and proxy_buffers directives. If the whole response does not fit into memory, a part of it can be saved to a temporary file on the disk. Writing to temporary files is controlled by the proxy_max_temp_file_size and proxy_temp_file_write_size directives. |
|-----|---|
| off | The response is passed to a client synchronously, immediately as it is received. Angie will not try to read the whole response from the proxied server. The maximum size of the data that Angie can receive from the server at a time is set by the <code>proxy_buffer_size</code> directive.   |

Buffering can also be enabled or disabled by passing "yes" or "no" in the "X-Accel-Buffering" response header field. This capability can be disabled using the *proxy ignore headers* directive.

### proxy\_buffers

| Syntax  | proxy_buffers number size; |
|---------|----------------------------|
| Default | proxy_buffers 8 4k   8k;   |
| Context | http, server, location     |



Sets the number and size of the buffers used for reading a response from the proxied server, for a single connection.

By default, the buffer size is equal to one memory page. This is either 4K or 8K, depending on a platform.

### proxy\_busy\_buffers\_size

| Syntax  | proxy_busy_buffers_size $size$ ;             |
|---------|--|
| Default | <pre>proxy_busy_buffers_size 8k   16k;</pre> |
| Context | http, server, location                       |

When *buffering* of responses from the proxied server is enabled, limits the total size of buffers that can be busy sending a response to the client while the response is not yet fully read. In the meantime, the rest of the buffers can be used for reading the response and, if needed, buffering part of the response to a temporary file.

By default, size is limited by the size of two buffers set by the proxy\_buffer\_size and proxy\_buffers directives.

#### proxy cache

| Syntax  | proxy_cache zone   off [path=path]; |
|---------|-------------------------------------|
| Default | <pre>proxy_cache off;</pre>         |
| Context | http, server, location              |

Defines a shared memory zone for caching. A zone can be used in the configuration multiple times. The parameter's value allows variables.

Added in version 1.2.0: PRO

In Angie PRO, you can specify multiple  $proxy\_cache\_path$  directives that share the same keys\_zone value to implement  $cache\ sharding$ . If you do, set the path parameter of the  $proxy\_cache$  directive that references this keys\_zone:

| path=path | The value is determined when the backend's response is <i>cached</i> , which implies that variables are involved, including those that store some information from the               |
|-----------|--|
|           | response.  |
|           | If the response is obtained from the cache, <i>path</i> isn't reevaluated; thus, a response from the cache will preserve its original <i>path</i> until it's deleted from the cache. |

This allows choosing between cache paths by applying map directives or scripts to responses from the backend. A Content-Type example:

```
proxy_cache_path /cache/one keys_zone=zone:10m;
proxy_cache_path /cache/two keys_zone=zone;

map $upstream_http_content_type $cache {
    ~^text/ one;
    default two;
}

server {
    ...
```



```
location / {
    proxy_pass http://backend;
    proxy_cache zone path=/cache/$cache;
}
```

This adds two cache paths and a variable mapping to choose between them. If Content-Type starts with text/, the first path is used; otherwise, the second.

# proxy\_cache\_background\_update

| Syntax  | <pre>proxy_cache_background_update on   off;</pre> |
|---------|--|
| Default | <pre>proxy_cache_background_update off;</pre>      |
| Context | http, server, location                             |

Allows starting a background subrequest to update an expired cache item, while a stale cached response is returned to the client.

### **A**ttention

Note that it is necessary to allow the usage of a stale cached response when it is being updated.

# proxy\_cache\_bypass

| Syntax  | proxy_cache_bypass;    |
|---------|------------------------|
| Default | _                      |
| Context | http, server, location |

Defines conditions under which the response will not be taken from a cache. If at least one value of the string parameters is not empty and is not equal to "0" then the response will not be taken from the cache:

```
proxy_cache_bypass $cookie_nocache $arg_nocache$arg_comment;
proxy_cache_bypass $http_pragma $http_authorization;
```

Can be used along with the proxy no cache directive.

### proxy cache convert head

| Syntax  | proxy_cache_convert_head on   off;      |
|---------|---|
| Default | <pre>proxy_cache_convert_head on;</pre> |
| Context | http, server, location                  |

Enables or disables the conversion of the "HEAD" method to "GET" for caching. When the conversion is disabled, the *cache key* should be configured to include the *\$request\_method*.

# proxy\_cache\_key

| Syntax  | proxy_cache_key $string;$                                     |
|---------|---|
| Default | <pre>proxy_cache_key \$scheme\$proxy_host\$request_uri;</pre> |
| Context | http, server, location  |



Defines a key for caching, for example

```
proxy_cache_key "$host$request_uri $cookie_user";
```

By default, the directive's value is close to the string

```
proxy_cache_key $scheme$proxy_host$uri$is_args$args;
```

### proxy\_cache\_lock

| Syntax  | proxy_cache_lock on   off;       |
|---------|----------------------------------|
| Default | <pre>proxy_cache_lock off;</pre> |
| Context | http, server, location           |

When enabled, only one request at a time will be allowed to populate a new cache element identified according to the  $proxy\_cache\_key$  directive by passing a request to a proxied server. Other requests of the same cache element will either wait for a response to appear in the cache or the cache lock for this element to be released, up to the time set by the  $proxy\_cache\_lock\_timeout$  directive.

### proxy cache lock age

| Syntax  | proxy_cache_lock_age time;          |
|---------|-------------------------------------|
| Default | <pre>proxy_cache_lock_age 5s;</pre> |
| Context | http, server, location              |

If the last request passed to the proxied server for populating a new cache element has not completed for the specified time, one more request may be passed to the proxied server.

# proxy\_cache\_lock\_timeout

| Syntax  | ${\tt proxy\_cache\_lock\_timeout}\ time;$ |
|---------|--|
| Default | <pre>proxy_cache_lock_timeout 5s;</pre>    |
| Context | http, server, location                     |

Sets a timeout for  $proxy\_cache\_lock$ . When the time expires, the request will be passed to the proxied server, however, the response will not be cached.

### proxy cache max range offset

| Syntax  | $\verb proxy_cache_max_range_offset  number; \\$ |
|---------|--|
| Default | _  |
| Context | http, server, location                           |

Sets an offset in bytes for byte-range requests. If the range is beyond the offset, the range request will be passed to the proxied server and the response will not be cached.

## proxy cache methods

| Syntax  | proxy_cache_methods GET   HEAD   POST;   |
|---------|--|
| Default | <pre>proxy_cache_methods GET HEAD;</pre> |
| Context | http, server, location                   |



If the client request method is listed in this directive then the response will be cached. "GET" and "HEAD" methods are always added to the list, though it is recommended to specify them explicitly. See also the  $proxy\_no\_cache$  directive.

### proxy cache min uses

| Syntax  | proxy_cache_min_uses number;       |
|---------|------------------------------------|
| Default | <pre>proxy_cache_min_uses 1;</pre> |
| Context | http, server, location             |

Sets the number of requests after which the response will be cached.

## proxy\_cache\_path

| Syntax  | proxy_cache_path $path$ [levels= $levels$ ] [use_temp_path=on   off] keys_zone= $name:size$ [:file= $file$ ] [inactive= $time$ ] [max_size= $size$ ] [min_free= $size$ ] [manager_files= $number$ ] [manager_sleep= $time$ ] [manager_threshold= $time$ ] [loader_files= $number$ ] [loader_sleep= $time$ ] [loader_threshold= $time$ ]; |
|---------|--|
| Default |  |
| Context | http   |

Sets the *path* and other parameters of a cache. Cache data are stored in files. The file name in a cache is a result of applying the MD5 function to the *cache key*.

| levels | defines hierarchy levels of a cache: from 1 to 3, each level accepts values 1 or 2. |
|--------|---|
|        | J I I I I I I I I I I I I I I I I I I I   |

For example, in the following configuration:

```
proxy_cache_path /data/angie/cache levels=1:2 keys_zone=one:10m;
```

file names in a cache will look like this:

```
/data/angie/cache/c/29/b7f54b2df7773722d382f4809d65029c
```

A cached response is first written to a temporary file, and then the file is renamed. Temporary files and the cache can be put on different file systems. However, be aware that in this case a file is copied across two file systems instead of the cheap renaming operation. It is thus recommended that for any given location both cache and a directory holding temporary files are put on the same file system.



| use_temp_path=on | Sets the directory for temporary files  |
|------------------|---|
| on               | If this parameter is omitted or set to the value on, the directory set by the   |
|                  | proxy_temp_path directive for the given location will be used.  |
| off              | Temporary files will be put directly in the cache directory.  |
| keys_zone        | Configures the name and size for a shared memory zone to store all active keys and information about data.  One megabyte zone can store about 8,000 keys.   |
|                  | When the optional <i>file</i> parameter is used with keys_zone, Angie flushes the contents of this zone to the disk on master process exit and attempts to restore it   |
|                  | at the same memory address at next <i>startup</i> or after a <i>binary upgrade</i> ; to achieve more robust persistence and improve cache loading time.   |
|                  | If the zone cannot be restored due to a change in size, binary version incompat-  |
|                  | ibility, or other reasons, Angie will log an alert (failed to restore zone at address) and will not use the zone restore mechanism. Instead, the incompatible file will be renamed as .old; you can either delete it, or restore its name and revert Angie to the configuration and version where it was created in the first |
|                  | place.  |
|                  | Attention   |
|                  | Ensure that the <i>file</i> path is valid and has the correct permissions for Angie to use it and prevent unauthorized access at the same time; relative paths are prefix-based.  |
| inactive         | Cached data that are not accessed during the time specified by this parameter get removed from the cache regardless of their freshness. By default, it is set to 10 minutes.  |

# 1 Note

Added in version 1.2.0: PRO

In Angie PRO, multiple  $proxy\_cache\_path$  directives that share the same keys\_zone value are allowed. Only the first such directive may set the shared memory zone size. The choice between such directives is made by the path parameter of the relevant  $proxy\_cache$  directive.

A special **cache manager** process monitors the maximum cache size and the minimum amount of free space on the file system with cache and when the size is exceeded or there is not enough free space, it removes the least recently used data. The data is removed in iterations.

| max_size         | maximum cache size   |
|------------------|--|
| min_free         | minimum amount of free space on the file system with cache                       |
| manager_files    | limits the number of items to be deleted during one iteration<br>By default, 100 |
| manager_threshol | limits the duration of one iteration<br>By default, 200 milliseconds             |
| manager_sleep    | configures a pause between interactions By default, 50 milliseconds              |

A minute after Angie starts, the special **cache loader** process is activated. It scans the file system for previously cached data and loads that information into the cache zone. This process is carried out in iterations; each iteration processes a limited number of items set by <code>loader\_files</code>, ensures it does not exceed the <code>loader\_threshold</code>, then pauses for a short interval set by <code>loader\_sleep</code> before proceeding to the next batch. These iterations continue until the loader has processed all existing cache entries on disk:



| loader_files     | limits the number of items to load during one iteration<br>By default, 100 |
|------------------|--|
| loador throabold | limits the duration of one iteration                                       |
| Toader_threshord |  |
|                  | By default, 200 milliseconds   |
| loader_sleep     | configures a pause between iterations                                      |
|                  | By default, 50 milliseconds  |

## 1 Note

Setting the file path for the keys\_zone parameter doesn't interfere with the cache loader behavior.

#### proxy cache revalidate

| Syntax  | <pre>proxy_cache_revalidate on   off;</pre> |
|---------|---|
| Default | <pre>proxy_cache_revalidate off;</pre>      |
| Context | http, server, location                      |

Enables revalidation of expired cache items using conditional requests with the "If-Modified-Since" and "If-None-Match" header fields.

## proxy cache use stale

| Syntax  | proxy_cache_use_stale error   timeout   invalid_header   updating   http_500   http_502   http_503   http_504   http_403   http_404   http_429   off; |
|---------|---|
| Default | <pre>proxy_cache_use_stale off;</pre>   |
| Context | http, server, location  |

Determines in which cases a stale cached response can be used during communication with the proxied server. The directive's parameters match the parameters of the proxy\_next\_upstream directive.

| error    | permits using a stale cached response if a proxied server to process a request cannot be selected.  |
|----------|---|
| updating | additional parameter, permits using a stale cached response if it is currently being updated. This allows minimizing the number of accesses to proxied servers when updating cached data. |

Using a stale cached response can also be enabled directly in the response header for a specified number of seconds after the response became stale:

- The stale-while-revalidate extension of the "Cache-Control" header field permits using a stale cached response if it is currently being updated.
- The stale-if-error extension of the "Cache-Control" header field permits using a stale cached response in case of an error.

## 1 Note

This has lower priority than using the directive parameters.

To minimize the number of accesses to proxied servers when populating a new cache element, the  $proxy\_cache\_lock$  directive can be used.



### proxy\_cache\_valid

| Syntax  | proxy_cache_valid [code] time; |
|---------|--------------------------------|
| Default | _                              |
| Context | http, server, location         |

Sets caching time for different response codes. For example, the following directives

```
proxy_cache_valid 200 302 10m;
proxy_cache_valid 404 1m;
```

set 10 minutes of caching for responses with codes 200 and 302 and 1 minute for responses with code 404.

If only caching time is specified

```
proxy_cache_valid 5m;
```

then only 200, 301, and 302 responses are cached.

In addition, the any parameter can be specified to cache any responses:

### 1 Note

Parameters of caching can also be set directly in the response header. This has higher priority than setting of caching time using the directive

- The "X-Accel-Expires" header field sets caching time of a response in seconds. The zero value disables caching for a response. If the value starts with the @ prefix, it sets an absolute time in seconds since Epoch, up to which the response may be cached.
- If the header does not include the "X-Accel-Expires" field, parameters of caching may be set in the header fields "Expires" or "Cache-Control".
- If the header includes the "Set-Cookie" field, such a response will not be cached.
- If the header includes the "Vary" field with the special value "\*", such a response will not be cached. If the header includes the "Vary" field with another value, such a response will be cached taking into account the corresponding request header fields.

Processing of one or more of these response header fields can be disabled using the  $proxy\_ignore\_headers$  directive.

# proxy\_connect\_timeout

| Syntax  | proxy_connect_timeout $time$ ;        |
|---------|---------------------------------------|
| Default | <pre>proxy_connect_timeout 60s;</pre> |
| Context | http, server, location                |

Defines a timeout for establishing a connection with a proxied server. It should be noted that this timeout cannot usually exceed 75 seconds.



# proxy\_connection drop

| Syntax  | proxy_connection_drop time   on   off; |
|---------|--|
| Default | <pre>proxy_connection_drop off;</pre>  |
| Context | http, server, location                 |

Enables termination of all connections to the proxied server after it has been removed from the group or marked as permanently unavailable by a reresolve process or the API command DELETE.

A connection is terminated when the next read or write event is processed for either the client or the proxied server.

Setting *time* enables a connection termination timeout; with on set, connections are dropped immediately.

## proxy\_cookie\_domain

| Syntax  | <pre>proxy_cookie_domain off;<br/>proxy_cookie_domain domain replacement;</pre> |
|---------|---|
| Default | <pre>proxy_cookie_domain off;</pre>   |
| Context | http, server, location  |

Sets a text that should be changed in the domain attribute of the "Set-Cookie" header fields of a proxied server response. Suppose a proxied server returned the "Set-Cookie" header field with the attribute "domain=localhost". The directive

```
proxy_cookie_domain localhost example.org;
```

will rewrite this attribute to "domain=example.org".

A dot at the beginning of the *domain* and *replacement* strings and the domain attribute is ignored. Matching is case-insensitive.

The domain and replacement strings can contain variables:

```
proxy_cookie_domain www.$host $host;
```

The directive can also be specified using regular expressions. In this case, domain should start with a " $\sim$ " symbol. A regular expression can contain named and positional captures, and replacement can reference them:

```
proxy_cookie_domain ~\.(?P<sl_domain>[-0-9a-z]+\.[a-z]+)$ $sl_domain;
```

Multiple proxy\_cookie\_domain directives may be specified at the same level:

```
proxy_cookie_domain localhost example.org;
proxy_cookie_domain ~\.([a-z]+\.[a-z]+)$ $1;
```

If several directives can be applied to the cookie, the first matching directive will be chosen.

The off parameter cancels the effect of the  $proxy\_cookie\_domain$  directives inherited from the previous configuration level.

#### proxy cookie flags

| Syntax  | proxy_cookie_flags off   cookie [flag]; |
|---------|---|
| Default | <pre>proxy_cookie_flags off;</pre>      |
| Context | http, server, location                  |



Sets one or more flags for the cookie. The cookie can contain text, variables, and their combinations. The flag can contain text, variables, and their combinations.

The secure, httponly, samesite=strict, samesite=lax, samesite=none parameters add the corresponding flags.

The nosecure, nohttponly, nosamesite parameters remove the corresponding flags.

The cookie can also be specified using regular expressions. In this case, cookie should start with a "~" symbol.

Several proxy\_cookie\_flags directives can be specified on the same configuration level:

```
proxy_cookie_flags one httponly;
proxy_cookie_flags ~ nosecure samesite=strict;
```

If several directives can be applied to the cookie, the first matching directive will be chosen. In the example, the *httponly* flag is added to the cookie *one*, for all other cookies the *samesite=strict* flag is added and the *secure* flag is deleted.

The off parameter cancels the effect of the  $proxy\_cookie\_flags$  directives inherited from the previous configuration level.

#### proxy cookie path

| Syntax  | <pre>proxy_cookie_path off;<br/>proxy_cookie_path path replacement;</pre> |
|---------|---|
| Default | <pre>proxy_cookie_path off;</pre>   |
| Context | http, server, location  |

Sets a text that should be changed in the path attribute of the Set-Cookie header fields of a proxied server response. Suppose a proxied server returned the "Set-Cookie" header field with the attribute "path=/two/some/uri/". The directive

```
proxy_cookie_path /two/ /;
```

will rewrite this attribute to "path=/some/uri/".

The path and replacement strings can contain variables:

```
proxy_cookie_path $uri /some$uri;
```

The directive can also be specified using regular expressions. In this case, path should either start with a " $^{-}$ " symbol for a case-sensitive matching, or with the " $^{-}$ \*" symbols for case-insensitive matching. The regular expression can contain named and positional captures, and replacement can reference them:

```
proxy_cookie_path ~*^/user/([^/]+) /u/$1;
```

Several proxy cookie path directives can be specified on the same level:

```
proxy_cookie_path /one/ /;
proxy_cookie_path / /two/;
```

If several directives can be applied to the cookie, the first matching directive will be chosen.

The off parameter cancels the effect of the proxy\_cookie\_path directives inherited from the previous configuration level.



### proxy\_force\_ranges

| Syntax  | <pre>proxy_force_ranges off;</pre> |
|---------|------------------------------------|
| Default | <pre>proxy_force_ranges off;</pre> |
| Context | http, server, location             |

Enables byte-range support for both cached and uncached responses from the proxied server regardless of the "Accept-Ranges" field in these responses.

### proxy headers hash bucket size

| Syntax  | $\verb proxy_headers_hash_bucket_size  size; \\$ |
|---------|--|
| Default | <pre>proxy_headers_hash_bucket_size 64;</pre>    |
| Context | http, server, location                           |

Sets the bucket size for hash tables used by the *proxy\_hide\_header* and *proxy\_set\_header* directives. The details of setting up hash tables are provided separately.

### proxy\_headers\_hash\_max\_size

| Syntax  | proxy_headers_hash_max_size $size;$         |
|---------|---|
| Default | <pre>proxy_headers_hash_max_size 512;</pre> |
| Context | http, server, location                      |

Sets the maximum size of hash tables used by the  $proxy\_hide\_header$  and  $proxy\_set\_header$  directives. The details of setting up hash tables are provided separately.

## proxy hide header

| Syntax  | proxy_hide_header field; |
|---------|--------------------------|
| Default | _                        |
| Context | http, server, location   |

By default, Angie does not pass the header fields "Date", "Server", "X-Pad", and "X-Accel-..." from the response of a proxied server to a client. The  $proxy\_hide\_header$  directive sets additional fields that will not be passed. If, on the contrary, the passing of fields needs to be permitted, the  $proxy\_pass\_header$  directive can be used.

### proxy\_http\_version

| Syntax  | proxy_http_version 1.0   1.1   3;                    |
|---------|--|
| Default | <pre>proxy_http_version 1.0;</pre>                   |
| Context | http, server, location, if in location, limit_except |

Sets the HTTP protocol version for proxying. By default, version 1.0 is used. Version 1.1 or higher is recommended for use with  $keepalive\ connections$ .



## proxy\_http3\_hq

| Syntax  | <pre>proxy_http3_hq on   off;</pre> |
|---------|-------------------------------------|
| Default | <pre>proxy_http3_hq off;</pre>      |
| Context | http, server                        |

Toggles the special hq-interop negotiation mode, which is used for QUIC interop tests that Angie relies on.

### Attention

Enable this mode only to run specialized tests that explicitly require it.

### proxy http3 max concurrent streams

| Syntax  | proxy_http3_max_concurrent_streams number;         |
|---------|--|
| Default | <pre>proxy_http3_max_concurrent_streams 128;</pre> |
| Context | http, server                                       |

Initializes HTTP/3 and QUIC settings and sets the maximum number of concurrent HTTP/3 request streams in a connection. Requires enabling  $keepalive\ connections$ .

## proxy\_http3\_max\_table\_capacity

| Syntax  | proxy_http3_max_table_capacity number;          |
|---------|---|
| Default | <pre>proxy_http3_max_table_capacity 4096;</pre> |
| Context | http, server, location                          |

Sets the dynamic table capacity for proxy connections.

# 1 Note

A similar  $http3\_max\_table\_capacity$  directive does this for server connections. To avoid errors, dynamic table usage is disabled when proxying with caching is enabled.

## proxy\_http3\_stream\_buffer\_size

| Syntax  | <pre>proxy_http3_stream_buffer_size size;</pre> |
|---------|---|
| Default | <pre>proxy_http3_stream_buffer_size 64k;</pre>  |
| Context | http, server                                    |

Sets the size of the read-write buffer used with QUIC streams.

# proxy\_ignore\_client\_abort

| Syntax  | <pre>proxy_ignore_client_abort on   off;</pre> |
|---------|--|
| Default | <pre>proxy_ignore_client_abort off;</pre>      |
| Context | http, server, location                         |



Determines whether the connection with a proxied server should be closed when a client closes the connection without waiting for a response.

#### proxy ignore headers

| Syntax  | proxy_ignore_headers field; |
|---------|-----------------------------|
| Default | _                           |
| Context | http, server, location      |

Disables processing of certain response header fields from the proxied server. The following fields can be ignored: "X-Accel-Redirect", "X-Accel-Expires", "X-Accel-Limit-Rate", "X-Accel-Buffering", "X-Accel-Charset", "Expires", "Cache-Control", "Set-Cookie", and "Vary".

If not disabled, processing of these header fields has the following effect:

- "X-Accel-Expires", "Expires", "Cache-Control", "Set-Cookie" and "Vary" set the parameters of response *caching*;
- "X-Accel-Redirect" performs an *internal* redirect to the specified URI;
- "X-Accel-Limit-Rate" sets the rate limit for transmission of a response to a client;
- "X-Accel-Buffering" enables or disables buffering of a response;
- "X-Accel-Charset" sets the desired *charset* of a response.

#### proxy\_intercept\_errors

| Syntax  | proxy_intercept_errors on   off;       |
|---------|--|
| Default | <pre>proxy_intercept_errors off;</pre> |
| Context | http, server, location                 |

Determines whether proxied responses with codes greater than or equal to 300 should be passed to a client or be intercepted and redirected to Angie for processing with the *error* page directive.

## proxy\_limit\_rate

| Syntax  | proxy_limit_rate rate;         |
|---------|--------------------------------|
| Default | <pre>proxy_limit_rate 0;</pre> |
| Context | http, server, location         |

Limits the speed of reading the response from the proxied server. The *rate* is specified in bytes per second and can contain variables.

0 disables rate limiting



The limit is set per a request, and so if Angie simultaneously opens two connections to the proxied server, the overall rate will be twice as much as the specified limit. The limitation works only if *buffering* of responses from the proxied server is enabled.



## proxy\_max\_temp\_file\_size

| Syntax  | <pre>proxy_max_temp_file_size size;</pre>  |
|---------|--|
| Default | <pre>proxy_max_temp_file_size 1024m;</pre> |
| Context | http, server, location                     |

When buffering of responses from the proxied server is enabled, and the whole response does not fit into the buffers set by the proxy\_buffer\_size and proxy\_buffers directives, a part of the response can be saved to a temporary file. This directive sets the maximum size of the temporary file. The size of data written to the temporary file at a time is set by the proxy\_temp\_file\_write\_size directive.

| 0 | disables buffering of responses to temporary files |  |
|---|--|--|
|---|--|--|

# 1 Note

This restriction does not apply to responses that will be cached or *stored on disk*.

### proxy method

| Syntax  | proxy_method method;   |
|---------|------------------------|
| Default | _                      |
| Context | http, server, location |

Specifies the HTTP method to use in requests forwarded to the proxied server instead of the method from the client request. Parameter value can contain variables.

## proxy\_next\_upstream

| Syntax  | proxy_next_upstream error   timeout   invalid_header   http_500   http_502   http_503   http_504   http_403   http_404   http_429   non_idempotent   off; |
|---------|---|
| Default | <pre>proxy_next_upstream error timeout;</pre>   |
| Context | http, server, location  |

Specifies in which cases a request should be passed to the next server in the upstream pool:

| error          | an error occurred while establishing a connection with the server, passing a request to it, or reading the response header; |
|----------------|---|
| timeout        | a timeout has occurred while establishing a connection with the server, passing   |
|                | a request to it, or reading the response header;  |
| invalid_header | a server returned an empty or invalid response;   |
| http_500       | a server returned a response with the code 500;   |
| http_502       | a server returned a response with the code 502;   |
| http_503       | a server returned a response with the code 503;   |
| http_504       | a server returned a response with the code 504;   |
| http_403       | a server returned a response with the code 403;   |
| http_404       | a server returned a response with the code 404;   |
| http_429       | a server returned a response with the code 429;   |
| non_idempotent | normally, requests with a non-idempotent method (POST, LOCK, PATCH) are not   |
|                | passed to the next server if a request has been sent to an upstream server; enabling  |
|                | this option explicitly allows retrying such requests;   |
| off            | disables passing a request to the next server.  |





One should bear in mind that passing a request to the next server is only possible if nothing has been sent to a client yet. That is, if an error or timeout occurs in the middle of the transferring of a response, fixing this is impossible.

The directive also defines what is considered an unsuccessful attempt of communication with a server.

| error<br>timeout<br>invalid_header                       | always considered unsuccessful attempts, even if they are not specified in the directive |
|--|--|
| http_500<br>http_502<br>http_503<br>http_504<br>http_429 | considered unsuccessful attempts only if they are specified in the directive             |
| http_403<br>http_404                                     | never considered unsuccessful attempts   |

Passing a request to the next server can be limited by the *number of tries* and by time.

### proxy next upstream timeout

| Syntax  | proxy_next_upstream_timeout time;         |
|---------|---|
| Default | <pre>proxy_next_upstream_timeout 0;</pre> |
| Context | http, server, location                    |

Limits the time during which a request can be passed to the *next* server.

| 0 | turns off this limitation |  |
|---|---------------------------|--|
|---|---------------------------|--|

### proxy\_next\_upstream\_tries

| Syntax  | <pre>proxy_next_upstream_tries number;</pre> |
|---------|--|
| Default | <pre>proxy_next_upstream_tries 0;</pre>      |
| Context | http, server, location                       |

Limits the number of possible tries for passing a request to the *next* server.

| 0 turns off this limitation |
|-----------------------------|
|-----------------------------|

### proxy\_no\_cache

| Syntax  | proxy_no_cache string; |
|---------|------------------------|
| Default | _                      |
| Context | http, server, location |

Defines conditions under which the response will not be saved to a cache. If at least one value of the string parameters is not empty and is not equal to "0" then the response will not be saved:



```
proxy_no_cache $cookie_nocache $arg_nocache$arg_comment;
proxy_no_cache $http_pragma $http_authorization;
```

Can be used along with the proxy\_cache\_bypass directive.

#### proxy\_pass

| Syntax  | $	exttt{proxy_pass } uri;$             |
|---------|--|
| Default | _                                      |
| Context | location, if in location, limit_except |

Sets the protocol and address of a proxied server and an optional URI to which a location should be mapped. As a protocol, http or https can be specified. The address can be specified as a domain name or IP address, and an optional port:

```
proxy_pass http://localhost:8000/uri/;
```

or as a UNIX domain socket path specified after the word unix and enclosed in colons:

```
proxy_pass http://unix:/tmp/backend.socket:/uri/;
```

If a domain name resolves to several addresses, all of them will be used in a round-robin fashion. In addition, an address can be specified as a *server group*.

Parameter value can contain variables. In this case, if an address is specified as a domain name, the name is searched among the described server groups, and, if not found, is determined using a *resolver*.

A request URI is passed to the server as follows:

• If the proxy\_pass directive is specified with a URI, then when a request is passed to the server, the part of a *normalized* request URI matching the location is replaced by a URI specified in the directive:

```
location /name/ {
    proxy_pass http://127.0.0.1/remote/;
}
```

• If proxy\_pass is specified without a URI, the request URI is passed to the server in the same form as sent by a client when the original request is processed, or the full normalized request URI is passed when processing the changed URI:

```
location /some/path/ {
    proxy_pass http://127.0.0.1;
}
```

In some cases, the part of a request URI to be replaced cannot be determined:

• When location is specified using a regular expression, and also inside named *locations* 

In these cases, proxy\_pass should be specified without a URI.

• When the URI is changed inside a proxied location using the *rewrite* directive, and this same configuration will be used to process a request (*break*):

```
location /name/ {
    rewrite     /name/([^/]+) /users?name=$1 break;
    proxy_pass http://127.0.0.1;
}
```



In this case, the URI specified in the directive is ignored and the full changed request URI is passed to the server.

• When variables are used in proxy\_pass:

```
location /name/ {
    proxy_pass http://127.0.0.1$request_uri;
}
```

In this case, if URI is specified in the directive, it is passed to the server as is, replacing the original request URI.

WebSocket proxying requires special configuration.

## proxy\_pass\_header

```
Syntax proxy_pass_header field ...;

Default —
Context http, server, location
```

Permits passing otherwise disabled header fields from a proxied server to a client.

### proxy pass request body

| Syntax  | proxy_pass_request_body on   off;      |
|---------|--|
| Default | <pre>proxy_pass_request_body on;</pre> |
| Context | http, server, location                 |

Indicates whether the original request body is passed to the proxied server.

```
location /x-accel-redirect-here/ {
    proxy_method GET;
    proxy_pass_request_body off;
    proxy_set_header Content-Length "";

    proxy_pass ...;
}
```

See also the proxy set header and proxy pass request headers directives.

### proxy pass request headers

```
Syntax proxy_pass_request_headers on | off;
Default proxy_pass_request_headers on;
Context http, server, location
```

Indicates whether the header fields of the original request are passed to the proxied server.

```
location /x-accel-redirect-here/ {
    proxy_method GET;
    proxy_pass_request_headers off;
    proxy_pass_request_body off;

    proxy_pass ...;
}
```



See also the  $proxy\_set\_header$  and  $proxy\_pass\_request\_body$  directives.

### proxy pass trailers

| Syntax  | proxy_pass_trailers on   off;       |
|---------|-------------------------------------|
| Default | <pre>proxy_pass_trailers off;</pre> |
| Context | http, server, location              |

Allows passing trailer fields from a proxied server to a client.

A trailer section in HTTP/1.1 is explicitly enabled.

```
location / {
    proxy_http_version 1.1;
    proxy_set_header Connection "te";
    proxy_set_header TE "trailers";
    proxy_pass_trailers on;
    proxy_pass ...;
}
```

## proxy\_quic\_active\_connection\_id\_limit

| Syntax  | <pre>proxy_quic_active_connection_id_limit number;</pre> |
|---------|--|
| Default | <pre>proxy_quic_active_connection_id_limit 2;</pre>      |
| Context | http, server   |

Sets the QUIC active\_connection\_id\_limit transport parameter value. This is the maximum number of active connection IDs that can be maintained per server.

### proxy\_quic\_gso

| Syntax  | proxy_quic_gso on   off; |
|---------|--------------------------|
| Default | proxy_quic_gso off;      |
| Context | http, server             |

Toggles sending data in QUIC-optimized batch mode using (generic segmentation offload).

## proxy\_quic\_host\_key

| Syntax  | proxy_quic_host_key file; |
|---------|---------------------------|
| Default | _                         |
| Context | http, server              |

Sets a file with the secret key used with QUIC to encrypt Stateless Reset and Address Validation tokens. By default, a random key is generated at each restart. Tokens generated with old keys are not accepted.

# $proxy\_read\_timeout$

| Syntax  | proxy_read_timeout $time$ ;        |
|---------|------------------------------------|
| Default | <pre>proxy_read_timeout 60s;</pre> |
| Context | http, server, location             |



Defines a timeout for reading a response from the proxied server. The timeout is set only between two successive read operations, not for the transmission of the whole response. If the proxied server does not transmit anything within this time, the connection is closed.

### proxy redirect

| Syntax  | <pre>proxy_redirect default; proxy_redirect off; proxy_redirect redirect replacement;</pre> |
|---------|---|
| Default | <pre>proxy_redirect default;</pre>  |
| Context | http, server, location  |

Sets the text that should be changed in the "Location" and "Refresh" header fields of a proxied server response.

Suppose a proxied server returned the header field:

```
Location: http://localhost:8000/two/some/uri/
```

The directive

```
proxy_redirect http://localhost:8000/two/ http://frontend/one/;
```

will rewrite this string to:

```
Location: http://frontend/one/some/uri/
```

A server name may be omitted in the *replacement* string:

```
proxy_redirect http://localhost:8000/two/ /;
```

then the primary server's name and port, if different from 80, will be inserted.

The default replacement specified by the default parameter uses the parameters of the *location* and *proxy pass* directives. Hence, the two configurations below are equivalent:

```
location /one/ {
    proxy_pass     http://upstream:port/two/;
    proxy_redirect default;
```

```
location /one/ {
    proxy_pass     http://upstream:port/two/;
    proxy_redirect http://upstream:port/two/ /one/;
```

### \* Caution

The default parameter is not permitted if proxy\_pass is specified using variables.

A replacement string can contain variables:

```
proxy_redirect http://localhost:8000/ http://$host:$server_port/;
```

A redirect can also contain variables:

```
proxy_redirect http://$proxy_host:8000/ /;
```



The directive can be specified using regular expressions. In this case, redirect should either start with the " $\sim$ " symbol for a case-sensitive matching, or with the " $\sim$ " symbols for case-insensitive matching. The regular expression can contain named and positional captures, and replacement can reference them:

```
proxy_redirect ~^(http://[^:]+):\d+(/.+)$ $1$2;
proxy_redirect ~*/user/([^/]+)/(.+)$ http://$1.example.com/$2;
```

Several proxy redirect directives can be specified on the same level:

```
proxy_redirect default;
proxy_redirect http://localhost:8000/ /;
proxy_redirect http://www.example.com/ /;
```

If several directives can be applied to the header fields of a proxied server response, the first matching directive will be chosen.

The off parameter cancels the effect of the  $proxy\_redirect$  directives inherited from the previous configuration level.

Using this directive, it is also possible to add host names to relative redirects issued by a proxied server:

```
proxy_redirect / /;
```

### proxy request buffering

| Syntax  | <pre>proxy_request_buffering on   off;</pre> |
|---------|--|
| Default | <pre>proxy_request_buffering on;</pre>       |
| Context | http, server, location                       |

Enables or disables buffering of a client request body.

| on  | the entire request body is <i>read</i> from the client before sending the request to a proxied server.  |
|-----|---|
| off | the request body is sent to the proxied server immediately as it is received. In this case, the request cannot be passed to the <i>next server</i> if Angie already started sending the request body. |

When HTTP/1.1 chunked transfer encoding is used to send the original request body, the request body will be buffered regardless of the directive value unless HTTP/1.1 is enabled for proxying.

# proxy\_send\_lowat

| Syntax  | <pre>proxy_send_lowat size;</pre> |
|---------|-----------------------------------|
| Default | <pre>proxy_send_lowat 0;</pre>    |
| Context | http, server, location            |

If the directive is set to a non-zero value, Angie will try to minimize the number of send operations on outgoing connections to a proxied server by using either  $NOTE\_LOWAT$  flag of the kqueue method, or the  $SO\_SNDLOWAT$  socket option, with the specified size.

```
    Note

This directive is ignored on Linux, Solaris, and Windows.
```



### proxy\_send\_timeout

| Syntax  | $\verb"proxy_send_time" ime;$      |
|---------|------------------------------------|
| Default | <pre>proxy_send_timeout 60s;</pre> |
| Context | http, server, location             |

Sets a timeout for transmitting a request to the proxied server. The timeout is set only between two successive write operations, not for the transmission of the whole request. If the proxied server does not receive anything within this time, the connection is closed.

### proxy set body

| Syntax  | proxy_set_body $value;$ |
|---------|-------------------------|
| Default | _                       |
| Context | http, server, location  |

Allows redefining the request body passed to the proxied server. The value can contain text, variables, and their combination.

### proxy\_set\_header

| Syntax  | proxy_set_header field value;                  |
|---------|--|
| Default | <pre>proxy_set_header Host \$proxy_host;</pre> |
| Context | http, server, location                         |

Allows redefining or appending fields to the request header *passed* to the proxied server. The *value* can contain text, variables, and their combinations. These directives are inherited from the previous configuration level if and only if there are no *proxy\_set\_header* directives defined on the current level. By default, only two fields are redefined:

If caching is enabled, the header fields "If-Modified-Since", "If-Unmodified-Since", "If-None-Match", "If-Match", "Range", and "If-Range" from the original request are not passed to the proxied server.

An unchanged "Host" request header field can be passed like this:

```
proxy_set_header Host $http_host;
```

However, if this field is not present in a client request header then nothing will be passed. In such a case it is better to use the *\$host* variable - its value equals the server name in the "Host" request header field or the primary server name if this field is not present:

```
proxy_set_header Host $host;
```

In addition, the server name can be passed together with the port of the proxied server:

```
proxy_set_header Host $host:$proxy_port;
```

If the value of a header field is an empty string then this field will not be passed to a proxied server:

```
proxy_set_header Accept-Encoding "";
```



### proxy\_socket\_keepalive

| Syntax  | <pre>proxy_socket_keepalive on   off;</pre> |
|---------|---|
| Default | <pre>proxy_socket_keepalive off;</pre>      |
| Context | http, server, location                      |

Configures the "TCP keepalive" behavior for outgoing connections to a proxied server.

| 11 11 | By default, the operating system's settings are in effect for the socket. |
|-------|---|
| on    | The SO_KEEPALIVE socket option is turned on for the socket.               |

### proxy ssl certificate

| Syntax  | proxy_ssl_certificate file [file]; |
|---------|------------------------------------|
| Default | _                                  |
| Context | http, server, location             |

Specifies a file with the certificate in the PEM format used for authentication to a proxied HTTPS server. Variables can be used in the file name.

Added in version 1.2.0.

When  $proxy\_ssl\_ntls$  enabled, directive accepts two arguments instead of one, sign and encryption parts of certificate:

```
location /proxy {
    proxy_ssl_ntls on;

proxy_ssl_certificate sign.crt enc.crt;
    proxy_ssl_certificate_key sign.key enc.key;

proxy_ssl_ciphers "ECC-SM2-WITH-SM4-SM3:ECDHE-SM2-WITH-SM4-SM3:RSA";

proxy_pass https://backend:443;
}
```

#### proxy ssl certificate cache

| Syntax  | $\label{lem:proxy_ssl_certificate_cache off;} $$proxy_ssl_certificate\_cache max=N [inactive=time] [valid=time];}$ |
|---------|--|
| Default | <pre>proxy_ssl_certificate_cache off;</pre>  |
| Context | http, server, location   |

Defines a cache that stores SSL certificates and secret keys specified using variables.

The directive supports the following parameters:

- max sets the maximum number of elements in the cache. When the cache overflows, the least recently used (LRU) elements are removed.
- inactive defines the time after which an element is removed if it has not been accessed. The default is 10 seconds.
- valid defines the time during which a cached element is considered valid and can be reused. The default is 60 seconds. After this period, certificates are reloaded or revalidated.



• off — disables the cache.

### Example:

## proxy ssl certificate key

| Syntax  | proxy_ssl_certificate_key file [file]; |
|---------|--|
| Default | _                                      |
| Context | http, server, location                 |

Specifies a file with the secret key in the PEM format used for authentication to a proxied HTTPS server.

The value "engine: name: id" can be specified instead of the file, which loads a secret key with a specified id from the OpenSSL engine name. Variables can be used in the file name.

Added in version 1.2.0.

When proxy\_ssl\_ntls enabled, directive accepts two arguments instead of one: sign and encryption parts of key:

```
location /proxy {
    proxy_ssl_ntls on;

proxy_ssl_certificate sign.crt enc.crt;
    proxy_ssl_certificate_key sign.key enc.key;

proxy_ssl_ciphers "ECC-SM2-WITH-SM4-SM3:ECDHE-SM2-WITH-SM4-SM3:RSA";

proxy_pass https://backend:443;
}
```

## proxy ssl ciphers

| Syntax  | proxy_ssl_ciphers ciphers;            |
|---------|---------------------------------------|
| Default | <pre>proxy_ssl_ciphers DEFAULT;</pre> |
| Context | http, server, location                |

Specifies the enabled ciphers for requests to a proxied HTTPS server. The ciphers are specified in the format understood by the OpenSSL library.

The list of ciphers depends on the version of OpenSSL installed. The full list can be viewed using the openssl ciphers command.

### **A** Attention

The proxy\_ssl\_ciphers directive does not configure ciphers for TLS 1.3 when using OpenSSL. To tune TLS 1.3 ciphers with OpenSSL, use the proxy\_ssl\_conf\_command directive, which was added to support advanced SSL configuration.

- In LibreSSL, TLS 1.3 ciphers can be configured using proxy\_ssl\_ciphers.
- In BoringSSL, TLS 1.3 ciphers cannot be configured at all.



# proxy\_ssl\_conf\_command

| Syntax  | <pre>proxy_ssl_conf_command name value;</pre> |
|---------|---|
| Default | _   |
| Context | http, server, location                        |

Sets arbitrary OpenSSL configuration commands when establishing a connection with the proxied HTTPS server.

# Important

The directive is supported when using OpenSSL 1.0.2 or higher. To configure TLS 1.3 ciphers with OpenSSL, use the ciphersuites command.

Several proxy\_ssl\_conf\_command directives can be specified on the same level. These directives are inherited from the previous configuration level if and only if there are no proxy\_ssl\_conf\_command directives defined on the current level.

## \* Caution

Note that configuring OpenSSL directly might result in unexpected behavior.

# proxy\_ssl\_crl

| Syntax  | proxy_ssl_crl file;    |
|---------|------------------------|
| Default | _                      |
| Context | http, server, location |

Specifies a file with revoked certificates (CRL) in the PEM format used to *verify* the certificate of the proxied HTTPS server.

# proxy\_ssl\_name

| Syntax  | proxy_ssl_name name;                    |
|---------|---|
| Default | <pre>proxy_ssl_name \$proxy_host;</pre> |
| Context | http, server, location                  |

Allows overriding the server name used to *verify* the certificate of the proxied HTTPS server and to be *passed through SNI* when establishing a connection with the proxied HTTPS server.

By default, the host part of the proxy pass URL is used.

## proxy\_ssl\_ntls

Added in version 1.2.0.

| Syntax  | proxy_ssl_ntls on   off;       |
|---------|--------------------------------|
| Default | <pre>proxy_ssl_ntls off;</pre> |
| Context | http, server                   |

Enables client-side support for NTLS using TongSuo library.



```
location /proxy {
    proxy_ssl_ntls on;

proxy_ssl_certificate sign.crt enc.crt;
    proxy_ssl_certificate_key sign.key enc.key;

proxy_ssl_ciphers "ECC-SM2-WITH-SM4-SM3:ECDHE-SM2-WITH-SM4-SM3:RSA";

proxy_pass https://backend:443;
}
```

# Important

Build Angie using the --with-ntls build option and link with NTLS-enabled SSL library

## proxy ssl password file

| Syntax  | $	exttt{proxy\_ssl\_password\_file}$ $file;$ |
|---------|--|
| Default | _  |
| Context | http, server, location                       |

Specifies a file with passphrases for *secret keys* where each passphrase is specified on a separate line. Passphrases are tried in turn when loading the key.

## proxy\_ssl\_protocols

| Syntax  | proxy_ssl_protocols [SSLv2] [SSLv3] [TLSv1] [TLSv1.1] [TLSv1.2] [TLSv1.3]; |
|---------|--|
| Default | <pre>proxy_ssl_protocols TLSv1.2 TLSv1.3;</pre>                            |
| Context | http, server, location   |

Changed in version 1.2.0: TLSv1.3 parameter added to default set.

Enables the specified protocols for requests to a proxied HTTPS server.

# proxy\_ssl\_server\_name

| Syntax  | <pre>proxy_ssl_server_name on   off;</pre> |
|---------|--|
| Default | <pre>proxy_ssl_server_name off;</pre>      |
| Context | http, server, location                     |

Enables or disables passing the server name set by the *proxy\_ssl\_name* directive via the Server Name Indication TLS extension (SNI, RFC 6066) while establishing a connection with the proxied HTTPS server.



# proxy\_ssl\_session\_reuse

| Syntax  | <pre>proxy_ssl_session_reuse on   off;</pre> |
|---------|--|
| Default | <pre>proxy_ssl_session_reuse on;</pre>       |
| Context | http, server, location                       |

Determines whether SSL sessions can be reused when working with the proxied server. If the errors " $SSL3\_GET\_FINISHED:digest\ check\ failed$ " appear in the logs, try disabling session reuse.

## proxy ssl trusted certificate

| Syntax  | $	ext{proxy\_ssl\_trusted\_certificate} \ file;$ |
|---------|--|
| Default | _  |
| Context | http, server, location                           |

Specifies a file with trusted CA certificates in the PEM format used to *verify* the certificate of the proxied HTTPS server.

# proxy\_ssl\_verify

| Syntax  | proxy_ssl_verify on   off;       |
|---------|----------------------------------|
| Default | <pre>proxy_ssl_verify off;</pre> |
| Context | http, server, location           |

Enables or disables verification of the proxied HTTPS server certificate.

# proxy ssl verify depth

| Syntax  | proxy_ssl_verify_depth $number;$     |
|---------|--------------------------------------|
| Default | <pre>proxy_ssl_verify_depth 1;</pre> |
| Context | http, server, location               |

Sets the verification depth in the proxied HTTPS server certificates chain.

# proxy\_store

| Syntax  | proxy_store on   off   string; |
|---------|--------------------------------|
| Default | <pre>proxy_store off;</pre>    |
| Context | http, server, location         |

Enables saving of files to a disk.

| on  | saves files with paths corresponding to the directives alias or root |
|-----|--|
| off | disables saving of files   |

The file name can be set explicitly using the string with variables:

```
proxy_store /data/www$original_uri;
```



The modification time of files is set according to the received "Last-Modified" response header field. The response is first written to a temporary file, and then the file is renamed. Temporary files and the persistent store can be put on different file systems. However, be aware that in this case a file is copied across two file systems instead of the cheap renaming operation. It is thus recommended that for any given location both saved files and a directory holding temporary files, set by the proxy\_temp\_path directive, are put on the same file system.

This directive can be used to create local copies of static unchangeable files, e.g.:

```
location /images/ {
    root
                        /data/www;
                       404 = /fetch$uri;
    error_page
}
location /fetch/ {
    internal;
    proxy_pass
                       http://backend/;
    proxy_store
                       on;
    proxy_store_access user:rw group:rw all:r;
    proxy_temp_path
                       /data/temp;
    alias
                       /data/www/;
}
```

or like this:

```
location /images/ {
    root
                        /data/www;
                       404 = 0fetch;
    error_page
location @fetch {
    internal;
                       http://backend;
    proxy_pass
    proxy_store
                       on;
    proxy_store_access user:rw group:rw all:r;
    proxy_temp_path
                       /data/temp;
                       /data/www;
    root
}
```

### proxy\_store\_access

```
Syntax proxy_store_access users:permissions ...;

Default proxy_store_access user:rw;

Context http, server, location
```

Sets access permissions for newly created files and directories, e.g.:

```
proxy_store_access user:rw group:rw all:r;
```

If any group or all access permissions are specified then user permissions may be omitted:

```
proxy_store_access group:rw all:r;
```



## proxy\_temp\_file\_write\_size

| Syntax  | proxy_temp_file_write_size size;              |
|---------|---|
| Default | <pre>proxy_temp_file_write_size 8k 16k;</pre> |
| Context | http, server, location                        |

Limits the size of data written to a temporary file at a time, when buffering of responses from the proxied server to temporary files is enabled. By default, size is limited by two buffers set by the proxy\_buffer\_size and proxy\_buffers directives. The maximum size of a temporary file is set by the proxy max temp file size directive.

### proxy temp path

| Syntax  | <pre>proxy_temp_path path [level1 [level2 [level3]]]`;</pre>    |
|---------|---|
| Default | <pre>proxy_temp_path proxy_temp; (the path depends on the</pre> |
|         | http-proxy-temp-path build option)                              |
| Context | http, server, location  |

Defines a directory for storing temporary files with data received from proxied servers. Up to three-level subdirectory hierarchy can be used underneath the specified directory. For example, in the following configuration

```
proxy_temp_path /spool/angie/proxy_temp 1 2;
```

a temporary file might look like this:

```
/spool/angie/proxy_temp/7/45/00000123457
```

See also the use temp path parameter of the proxy cache path directive.

# **Built-in Variables**

The  $http\_proxy$  module supports built-in variables that can be used to compose headers using the proxy set header directive:

### \$proxy\_host

name and port of a proxied server as specified in the proxy\_pass directive;

## \$proxy\_port

port of a proxied server as specified in the proxy pass directive, or the protocol's default port;

## \$proxy\_add\_x\_forwarded\_for

the "X-Forwarded-For" client request header field with the  $\$remote\_addr$  variable appended to it, separated by a comma. If the "X-Forwarded-For" field is not present in the client request header, the \$proxy=add=x for variable is equal to the \$remote=addr variable.

#### Random Index

The module processes requests ending with the slash character (/) and picks a random file in a directory to serve as an index file. The module is processed before the  $http\_index$  module.

When building from the source code, this module isn't built by default; it should be enabled with the --with-http\_random\_index\_module build option.

In packages and images from our repos, the module is included in the build.



### **Configuration Example**

```
location / {
    random_index on;
}
```

#### **Directives**

### random index

| Syntax  | random_index on   off; |
|---------|------------------------|
| Default | random_index off;      |
| Context | location               |

Enables or disables module processing in a surrounding location.

### **RealIP**

The module is used to change the client address and optional port to those sent in the specified header field.

When building from the source code, this module isn't built by default; it should be enabled with the --with-http\_realip\_module build option.

In packages and images from our repos, the module is included in the build.

## **Configuration Example**

```
set_real_ip_from 192.168.1.0/24;
set_real_ip_from 192.168.2.1;
set_real_ip_from 2001:0db8::/32;
real_ip_header X-Forwarded-For;
real_ip_recursive on;
```

## **Directives**

# set real ip from

```
Syntax set_real_ip_from address | CIDR | unix:;

Default —
Context http, server, location
```

Defines trusted addresses that are known to send correct replacement addresses. If the special value unix: is specified, all UNIX domain sockets will be trusted. Trusted addresses may also be specified using a hostname.

## real ip header

| Syntax  | $\verb real_ip_header  field \mid \verb X-Real-IP  \mid \verb X-Forwarded-For  proxy_protocol ;$ |
|---------|--|
| Default | real_ip_header X-Real-IP;  |
| Context | http, server, location   |

Defines the request header field whose value will be used to replace the client address.



The request header field value that contains an optional port is also used to replace the client port. The address and port should be specified according to RFC 3986.

The proxy\_protocol parameter changes the client address to the one from the PROXY protocol header. The PROXY protocol must be previously enabled by setting the proxy\_protocol parameter in the listen directive.

## real ip recursive

| Syntax  | real_ip_recursive on   off; |
|---------|-----------------------------|
| Default | real_ip_recursive off;      |
| Context | http, server, location      |

If recursive search is disabled, the original client address that matches one of the trusted addresses is replaced by the last address sent in the request header field defined by the  $real\_ip\_header$  directive. If recursive search is enabled, the original client address that matches one of the trusted addresses is replaced by the last non-trusted address sent in the request header field.

#### **Built-in Variables**

```
$realip_remote_addr
```

keeps the original client address

```
$realip_remote_port
```

keeps the original client port

### Referer

The module is used to block access to a site for requests with invalid values in the "Referer" header field. It should be kept in mind that fabricating a request with an appropriate "Referer" field value is quite easy, and so the intended purpose of this module is not to block such requests thoroughly but to block the mass flow of requests sent by regular browsers. It should also be taken into consideration that regular browsers may not send the "Referer" field even for valid requests.

### **Configuration Example**

#### **Directives**

referer hash bucket size

```
Syntax referer_hash_bucket_size size;
Default referer_hash_bucket_size 64;
Context server, location
```

Sets the bucket size for the valid referers hash tables. The details of setting up hash tables are provided in a separate document.



## referer hash max size

| Syntax  | referer_hash_max_size $size;$ |
|---------|-------------------------------|
| Default | referer_hash_max_size 2048;   |
| Context | server, location              |

Sets the maximum size of the valid referers hash tables. The details of setting up hash tables are provided in a separate document.

## valid referers

| Syntax  | valid_referers none   blocked   server_names   string; |
|---------|--|
| Default | _  |
| Context | server, location                                       |

Specifies the "Referer" request header field values that will cause the built-in *\$invalid\_referer* variable to be set to an empty string. Otherwise, the variable will be set to "1". Search for a match is case-insensitive.

Parameters can be as follows:

| none                  | the "Referer" field is missing in the request header;   |
|-----------------------|---|
| blocked               | the "Referer" field is present in the request header, but its value has been deleted by a firewall or proxy server; such values are strings that do not start with http:/ |
|                       | / or https://;  |
| server_names          | the "Referer" request header field contains one of the server names;  |
| arbitrary             | defines a server name and an optional URI prefix. A server name can have an "*"   |
| string                | at the beginning or end. During the checking, the server's port in the "Referer" field is ignored;  |
| regular<br>expression | the first symbol should be a "~". It should be noted that an expression will be matched against the text starting after the http:// or https://.                          |

## Example:

#### **Built-in Variables**

#### \$invalid\_referer

Empty string, if the "Referer" request header field value is considered valid, otherwise "1".

## Rewrite

The module is used to change request URI using PCRE regular expressions, return redirects, and conditionally select configurations.

The break, if, return, rewrite and set directives are processed in the following order:

- the directives of this module specified on the *server* level are executed sequentially;
- repeatedly:
  - a location is searched based on a request URI;
  - the directives of this module specified inside the found location are executed sequentially;



- the loop is repeated if a request URI was rewritten, but not more than 10 times.

### **Directives**

#### break

| Syntax  | break;               |
|---------|----------------------|
| Default | _                    |
| Context | server, location, if |

Stops processing the current set of http rewrite directives.

If a directive is specified inside the location, further processing of the request continues in this location.

### Example:

```
if ($slow) {
    limit_rate 10k;
    break;
}
```

#### if

```
Syntax if (condition) { ... }

Default —

Context server, location
```

The specified condition is evaluated. If true, this module directives specified inside the braces are executed, and the request is assigned the configuration inside the *if* directive. Configurations inside the *if* directives are inherited from the previous configuration level.

A condition may be any of the following:

- a variable name; false if the value of a variable is an empty string or "0";
- comparison of a variable with a string using the "=" and "!=" operators;
- matching of a variable against a regular expression using the "~" (for case-sensitive matching) and "~\*" (for case-insensitive matching) operators. Regular expressions can contain captures that are made available for later reuse in the \$1..\$9 variables. Negative operators "!~" and "!~\*" are also available. If a regular expression includes the "}" or ";" characters, the whole expressions should be enclosed in single or double quotes.
- checking of a file existence with the "-f" and "!-f" operators;
- checking of a directory existence with the "-d" and "!-d" operators;
- checking of a file, directory, or symbolic link existence with the "-e" and "!-e" operators;
- checking for an executable file with the "-x" and "!-x" operators.

## Examples:

```
if ($http_user_agent ~ MSIE) {
    rewrite ^(.*)$ /msie/$1 break;
}
if ($http_cookie ~* "id=([^;]+)(?:;|$)") {
    set $id $1;
}
```



```
if ($request_method = POST) {
    return 405;
}

if ($slow) {
    limit_rate 10k;
}

if ($invalid_referer) {
    return 403;
}
```

## 1 Note

A value of the *\$invalid referer Built-in variable* is set by the *valid referers* directive.

#### return

| Syntax  | $egin{array}{ll} { m return} \ code \ [text]; \ { m return} \ code \ URL; \ { m return} \ URL; \end{array}$ |
|---------|---|
| Default | _   |
| Context | server, location, if  |

Stops processing and returns the specified code to a client. The non-standard code 444 closes a connection without sending a response header.

It is possible to specify either a redirect URL (for codes 301, 302, 303, 307, and 308) or the response body text (for other codes). A response body text and redirect URL can contain variables. As a special case, a redirect URL can be specified as a URI local to this server, in which case the full redirect URL is formed according to the request scheme (\$scheme) and the \$server\_name\_in\_redirect and \$port\_in\_redirect directives.

In addition, a URL for temporary redirect with the code 302 can be specified as the sole parameter. Such a parameter should start with the http://, https://, or "\$scheme" string. A URL can contain variables.

See also the *error* page directive.

## rewrite

| Syntax  | rewrite regex replacement [flag]; |
|---------|-----------------------------------|
| Default | _                                 |
| Context | server, location, if              |

If the specified regular expression matches a request URI, URI is changed as specified in the replacement string. The *rewrite* directives are executed sequentially in order of their appearance in the configuration file. It is possible to terminate further processing of the directives using flags. If a replacement string starts with http://, https://, or "\$scheme", the processing stops and the redirect is returned to a client.

An optional flag parameter can be one of:



| last      | stops processing the current set of <i>http_rewrite</i> , directives and starts a search for a new <i>location</i> matching the changed URI; |
|-----------|--|
| break     | stops processing the current set of <i>http_rewrite</i> directives as with the <i>break</i> directive;                                       |
| redirect  | returns a temporary redirect with the 302 code; used if a replacement string does not start with http://, https:// or "\$scheme";            |
| permanent | returns a permanent redirect with the 301 code.  |

The full redirect URL is formed according to the request scheme (\$scheme) and the server\_name\_in\_redirect and port\_in\_redirect directives.

### Example:

```
server {
# ...
    rewrite ^(/download/.*)/media/(.*)\..*$ $1/mp3/$2.mp3 last;
    rewrite ^(/download/.*)/audio/(.*)\..*$ $1/mp3/$2.ra last;
    return 403;
# ...
}
```

But if these directives are put inside the "/download/" location, the last flag should be replaced by break, or otherwise Angie will make 10 cycles and return the 500 error:

```
location /download/ {
    rewrite ^(/download/.*)/media/(.*)\..*$ $1/mp3/$2.mp3 break;
    rewrite ^(/download/.*)/audio/(.*)\..*$ $1/mp3/$2.ra break;
    return 403;
}
```

If a replacement string includes the new request arguments, the previous request arguments are appended after them. If this is undesired, putting a question mark at the end of a replacement string avoids having them appended, for example:

```
rewrite ^/users/(.*)$ /show?user=$1? last;
```

If a regular expression includes the "}" or ";" characters, the whole expressions should be enclosed in single or double quotes.

### rewrite log

| Syntax  | rewrite_log on   off;      |
|---------|----------------------------|
| Default | rewrite_log off;           |
| Context | http, server, location, if |

Enables or disables logging of  $http\_rewrite$  module directives processing results into the  $error\_log$  at the notice level.

#### set

| Syntax  | set \$variable value; |
|---------|-----------------------|
| Default | _                     |
| Context | server, location, if  |

Sets a value for the specified variable. The value can contain text, variables, and their combination.



## uninitialized variable warn

| Syntax  | uninitialized_variable_warn on   off; |
|---------|---------------------------------------|
| Default | uninitialized_variable_warn on;       |
| Context | http, server, location, if            |

Controls whether warnings about uninitialized variables are logged.

## **Internal Implementation**

The http\_rewrite module directives are compiled at the configuration stage into internal instructions that are interpreted during request processing. An interpreter is a simple virtual stack machine.

For example, the directives

```
location /download/ {
    if ($forbidden) {
        return 403;
    }

    if ($slow) {
        limit_rate 10k;
    }

    rewrite ^/(download/.*)/media/(.*)\..*$ /$1/mp3/$2.mp3 break;
}
```

will be translated into these instructions:

```
variable $forbidden
check against zero
    return 403
    end of code
variable $slow
check against zero
match of regular expression
copy "/"
copy $1
copy $1
copy "/mp3/"
copy $2
copy ".mp3"
end of regular expression
end of code
```

Note that there are no instructions for the  $limit\_rate$  directive above as it is unrelated to the  $http\_rewrite$  module. A separate configuration is created for the if block. If the condition holds true, a request is assigned this configuration where  $limit\_rate$  equals to 10k.

The directive

```
rewrite ^/(download/.*)/media/(.*)\..*$ /$1/mp3/$2.mp3 break;
```

can be made smaller by one instruction if the first slash in the regular expression is put inside the parentheses:

```
rewrite ^(/download/.*)/media/(.*)\..*$ $1/mp3/$2.mp3 break;
```

The corresponding instructions will then look like this:



```
match of regular expression
copy $1
copy "/mp3/"
copy $2
copy ".mp3"
end of regular expression
end of code
```

### **SCGI**

Allows passing requests to SCGI server.

### **Configuration Example**

```
location / {
   include   scgi_params;
   scgi_pass localhost:9000;
}
```

#### **Directives**

## scgi bind

| Syntax  | scgi_bind address [transparent]   off; |
|---------|--|
| Default | _                                      |
| Context | http, server, location                 |

Makes outgoing connections to a SCGI server originate from the specified local IP address with an optional port. Parameter value can contain variables. The special value off cancels the effect of the  $scgi\_bind$  directive inherited from the previous configuration level, which allows the system to auto-assign the local IP address and port.

The transparent parameter allows outgoing connections to a SCGI server originate from a non-local IP address, for example, from a real IP address of a client:

```
scgi_bind $remote_addr transparent;
```

In order for this parameter to work, it is usually necessary to run Angie worker processes with the superuser privileges. On Linux it is not required as if the transparent parameter is specified, worker processes inherit the CAP NET RAW capability from the master process.

# Important

It is necessary to configure kernel routing table to intercept network traffic from the SCGI server.

## scgi\_buffer\_size

| Syntax  | scgi_buffer_size size;  |
|---------|-------------------------|
| Default | scgi_buffer_size 4k 8k; |
| Context | http, server, location  |

Sets the size of the buffer used for reading the first part of the response received from the SCGI server. This part usually contains a small response header. By default, the buffer size is equal to one memory page. This is either 4K or 8K, depending on a platform. It can be made smaller, however.



# scgi\_buffering

| Syntax  | scgi_buffering $size;$ |
|---------|------------------------|
| Default | scgi_buffering on;     |
| Context | http, server, location |

Enables or disables buffering of responses from the SCGI server.

| on  | Angie receives a response from the SCGI server as soon as possible, saving             |
|-----|--|
|     | it into the buffers set by the $scgi\_buffer\_size$ and $scgi\_buffers$ directives. If |
|     | the whole response does not fit into memory, a part of it can be saved to              |
|     | a temporary file on the disk. Writing to temporary files is controlled by the          |
|     | $scgi\_max\_temp\_file\_size$ and $scgi\_temp\_file\_write\_size$ directives.          |
| off | The response is passed to a client synchronously, immediately as it is received.       |
|     | Angie will not try to read the whole response from the SCGI server. The maxi-          |
|     | mum size of the data that Angie can receive from the server at a time is set by        |
|     | the $scgi\_buffer\_size$ directive.  |

Buffering can also be enabled or disabled by passing "yes" or "no" in the "X-Accel-Buffering" response header field. This capability can be disabled using the <code>scgi\_ignore\_headers</code> directive.

## scgi buffers

| Syntax  | scgi_buffers number size; |
|---------|---------------------------|
| Default | scgi_buffers 8 4k   8k;   |
| Context | http, server, location    |

Sets the number and size of the buffers used for reading a response from the SCGI server, for a single connection.

By default, the buffer size is equal to one memory page. This is either 4K or 8K, depending on a platform.

## scgi busy buffers size

| Syntax  | scgi_busy_buffers_size size;                |
|---------|---|
| Default | <pre>scgi_busy_buffers_size 8k   16k;</pre> |
| Context | http, server, location                      |

When *buffering* of responses from the SCGI server is enabled, limits the total size of buffers that can be busy sending a response to the client while the response is not yet fully read. In the meantime, the rest of the buffers can be used for reading the response and, if needed, buffering part of the response to a temporary file.

By default, size is limited by the size of two buffers set by the scgi buffer size and scgi buffers directives.

# scgi\_cache

| Syntax  | scgi_cache zone   off; |
|---------|------------------------|
| Default | scgi_cache off;        |
| Context | http, server, location |

Defines a shared memory zone used for caching. The same zone can be used in several places. Parameter value can contain variables.



| off disables caching inherited from the previous configuration level. |  |
|---|--|
|---|--|

# scgi\_cache\_background\_update

| Syntax  | scgi_cache_background_update on   off; |
|---------|--|
| Default | scgi_cache_background_update off;      |
| Context | http, server, location                 |

Allows starting a background subrequest to update an expired cache item, while a stale cached response is returned to the client.

# **A** Attention

Note that it is necessary to allow the usage of a stale cached response when it is being updated.

## scgi cache bypass

| Syntax  | scgi_cache_bypass;     |
|---------|------------------------|
| Default | _                      |
| Context | http, server, location |

Defines conditions under which the response will not be taken from a cache. If at least one value of the string parameters is not empty and is not equal to "0" then the response will not be taken from the cache:

Can be used along with the scgi no cache directive.

## scgi cache key

| Syntax  | scgi_cache_key $string;$ |
|---------|--------------------------|
| Default | _                        |
| Context | http, server, location   |

Defines a key for caching, for example

```
scgi_cache_key localhost:9000$request_uri;
```

## scgi\_cache\_lock

| Syntax  | scgi_cache_lock on   off;       |
|---------|---------------------------------|
| Default | <pre>scgi_cache_lock off;</pre> |
| Context | http, server, location          |

When enabled, only one request at a time will be allowed to populate a new cache element identified according to the  $scgi\_cache\_key$  directive by passing a request to a SCGI server. Other requests of the same cache element will either wait for a response to appear in the cache or the cache lock for this element to be released, up to the time set by the  $scgi\_cache\_lock\_timeout$  directive.



## scgi\_cache\_lock\_age

| Syntax  | <pre>scgi_cache_lock_age time;</pre> |
|---------|--------------------------------------|
| Default | scgi_cache_lock_age 5s;              |
| Context | http, server, location               |

If the last request passed to the SCGI server for populating a new cache element has not completed for the specified time, one more request may be passed to the SCGI server.

## scgi cache lock timeout

| Syntax  | scgi_cache_lock_timeout $time$ ; |
|---------|----------------------------------|
| Default | scgi_cache_lock_timeout 5s;      |
| Context | http, server, location           |

Sets a timeout for  $scgi\_cache\_lock$ . When the time expires, the request will be passed to the SCGI server, however, the response will not be cached.

## scgi\_cache\_max\_range\_offset

| Syntax  | scgi_cache_max_range_offset number; |
|---------|-------------------------------------|
| Default | _                                   |
| Context | http, server, location              |

Sets an offset in bytes for byte-range requests. If the range is beyond the offset, the range request will be passed to the SCGI server and the response will not be cached.

# scgi cache methods

| Syntax  | scgi_cache_methods GET   HEAD   POST; |
|---------|---------------------------------------|
| Default | scgi_cache_methods GET HEAD;          |
| Context | http, server, location                |

If the client request method is listed in this directive then the response will be cached. "GET" and "HEAD" methods are always added to the list, though it is recommended to specify them explicitly. See also the  $scgi\_no\_cache$  directive.

# scgi\_cache\_min\_uses

| Syntax  | scgi_cache_min_uses number; |
|---------|-----------------------------|
| Default | scgi_cache_min_uses 1;      |
| Context | http, server, location      |

Sets the number of requests after which the response will be cached.



## scgi\_cache\_path

| Syntax  | ${\tt scgi\_cache\_path}  path  [{\tt levels=} levels]  [{\tt use\_temp\_path=on}     {\tt off}]$  |
|---------|--|
|         | $\texttt{keys\_zone} = name: size \qquad [\texttt{inactive} = time] \qquad [\texttt{max\_size} = size] \qquad [\texttt{min\_free} = size]$ |
|         | $[\texttt{manager\_files} = number]  [\texttt{manager\_sleep} = time]  [\texttt{manager\_threshold} = time]$                               |
|         | $[\texttt{loader\_files} = number] \ [\texttt{loader\_sleep} = time] \ [\texttt{loader\_threshold} = time];$                               |
| Default | _  |
| Context | http   |

Sets the path and other parameters of a cache. Cache data are stored in files. The file name in a cache is a result of applying the MD5 function to the *cache key*.

The levels parameter defines hierarchy levels of a cache: from 1 to 3, each level accepts values 1 or 2. For example, in the following configuration:

```
scgi_cache_path /data/angie/cache levels=1:2 keys_zone=one:10m;
```

file names in a cache will look like this:

```
/data/angie/cache/c/29/b7f54b2df7773722d382f4809d65029c
```

A cached response is first written to a temporary file, and then the file is renamed. Temporary files and the cache can be put on different file systems. However, be aware that in this case a file is copied across two file systems instead of the cheap renaming operation. It is thus recommended that for any given location both cache and a directory holding temporary files are put on the same file system.

The directory for temporary files is set based on the use\_temp\_path parameter.

| on  | If this parameter is omitted or set to the value on, the directory set by the $scgi\_temp\_path$ directive for the given $location$ will be used. |
|-----|---|
| off | Temporary files will be put directly in the cache directory.  |

In addition, all active keys and information about data are stored in a shared memory zone, whose name and size are configured by the keys\_zone parameter. One megabyte zone can store about 8 thousand keys.

Cached data that are not accessed during the time specified by the inactive parameter get removed from the cache regardless of their freshness.

By default, *inactive* is set to 10 minutes.

A special **cache manager** process monitors the maximum cache size and the minimum amount of free space on the file system with cache and when the size is exceeded or there is not enough free space, it removes the least recently used data. The data is removed in iterations.

| max_size         | maximum cache size   |
|------------------|--|
| min_free         | minimum amount of free space on the file system with cache                       |
| manager_files    | limits the number of items to be deleted during one iteration<br>By default, 100 |
| manager_threshol | limits the duration of one iteration By default, 200 milliseconds                |
| manager_sleep    | configures a pause between iterations By default, 50 milliseconds                |

A minute after Angie starts, the special **cache loader** process is activated. It loads information about previously cached data stored on file system into a cache zone. The loading is also done in iterations.



| loader_files     | limits the number of items to load during one iteration<br>By default, 100 |
|------------------|--|
| loader_threshold | limits the duration of one iteration By default, 200 milliseconds          |
| loader_sleep     | configures a pause between iterations By default, 50 milliseconds          |

# scgi\_cache\_revalidate

| Syntax  | scgi_cache_revalidate on   off; |
|---------|---------------------------------|
| Default | scgi_cache_revalidate off;      |
| Context | http, server, location          |

Enables revalidation of expired cache items using conditional requests with the "If-Modified-Since" and "If-None-Match" header fields.

# scgi\_cache\_use\_stale

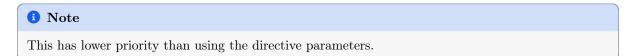
| Syntax  | ${\tt scgi\_cache\_use\_stale~error} \mid {\tt timeout} \mid {\tt invalid\_header} \mid {\tt updating} \mid {\tt http\_500}$ |
|---------|--|
|         | http_502   http_503   http_504   http_403   http_404   http_429   off;   |
| Default | scgi_cache_use_stale off;  |
| Context | http, server, location   |

Determines in which cases a stale cached response can be used during communication with the SCGI server. The directive's parameters match the parameters of the scgi next upstream directive.

| error    | permits using a stale cached response if a SCGI server to process a request cannot be selected.  |
|----------|--|
| updating | additional parameter, permits using a stale cached response if it is currently being updated. This allows minimizing the number of accesses to SCGI servers when updating cached data. |

Using a stale cached response can also be enabled directly in the response header for a specified number of seconds after the response became stale:

- The stale-while-revalidate extension of the "Cache-Control" header field permits using a stale cached response if it is currently being updated.
- The stale-if-error extension of the "Cache-Control" header field permits using a stale cached response in case of an error.



To minimize the number of accesses to SCGI servers when populating a new cache element, the scgi cache lock directive can be used.



## scgi\_cache\_valid

| Syntax  | $scgi\_cache\_valid [code] time;$ |
|---------|-----------------------------------|
| Default | _                                 |
| Context | http, server, location            |

Sets caching time for different response codes. For example, the following directives

set 10 minutes of caching for responses with codes 200 and 302 and 1 minute for responses with code 404

If only caching time is specified

```
scgi_cache_valid 5m;
```

then only 200, 301, and 302 responses are cached.

In addition, the any parameter can be specified to cache any responses:

## 1 Note

Parameters of caching can also be set directly in the response header. This has higher priority than setting of caching time using the directive

- The "X-Accel-Expires" header field sets caching time of a response in seconds. The zero value disables caching for a response. If the value starts with the @ prefix, it sets an absolute time in seconds since Epoch, up to which the response may be cached.
- If the header does not include the "X-Accel-Expires" field, parameters of caching may be set in the header fields "Expires" or "Cache-Control".
- If the header includes the "Set-Cookie" field, such a response will not be cached.
- If the header includes the "Vary" field with the special value "\*", such a response will not be cached. If the header includes the "Vary" field with another value, such a response will be cached taking into account the corresponding request header fields.

Processing of one or more of these response header fields can be disabled using the scgi\_ignore\_headers directive.

## scgi connect timeout

| Syntax  | scgi_connect_timeout time;           |
|---------|--------------------------------------|
| Default | <pre>scgi_connect_timeout 60s;</pre> |
| Context | http, server, location               |

Defines a timeout for establishing a connection with a SCGI server. It should be noted that this timeout cannot usually exceed 75 seconds.



## scgi\_connection\_drop

| Syntax  | scgi_connection_drop time   on   off; |
|---------|---------------------------------------|
| Default | scgi_connection_drop off;             |
| Context | http, server, location                |

Enables termination of all connections to the proxied server after it has been removed from the group or marked as permanently unavailable by a *reresolve* process or the *API command* DELETE.

A connection is terminated when the next read or write event is processed for either the client or the proxied server.

Setting time enables a connection termination timeout; with on set, connections are dropped immediately.

# scgi force ranges

| Syntax  | scgi_force_ranges off;            |
|---------|-----------------------------------|
| Default | <pre>scgi_force_ranges off;</pre> |
| Context | http, server, location            |

Enables byte-range support for both cached and uncached responses from the SCGI server regardless of the "Accept-Ranges" field in these responses.

# scgi\_hide\_header

| Syntax  | ${	t scgi_hide_header} \ field;$ |
|---------|----------------------------------|
| Default | _                                |
| Context | http, server, location           |

By default, Angie does not pass the header fields "Status" and "X-Accel-..." from the response of a SCGI server to a client. The  $scgi\_hide\_header$  directive sets additional fields that will not be passed. If, on the contrary, the passing of fields needs to be permitted, the  $scgi\_pass\_header$  directive can be used.

### scgi ignore client abort

| Syntax  | <pre>scgi_ignore_client_abort on   off;</pre> |
|---------|---|
| Default | <pre>scgi_ignore_client_abort off;</pre>      |
| Context | http, server, location                        |

Determines whether the connection with a SCGI server should be closed when a client closes the connection without waiting for a response.

# scgi\_ignore\_headers

| Syntax  | scgi_ignore_headers field; |
|---------|----------------------------|
| Default | _                          |
| Context | http, server, location     |

Disables processing of certain response header fields from the SCGI server. The following fields can be ignored: "X-Accel-Redirect", "X-Accel-Expires", "X-Accel-Limit-Rate", "X-Accel-Buffering", "X-Accel-Expires", "Expires", "Cache-Control", "Set-Cookie", and "Vary".

If not disabled, processing of these header fields has the following effect:



- "X-Accel-Expires", "Expires", "Cache-Control", "Set-Cookie" and "Vary" set the parameters of response *caching*;
- "X-Accel-Redirect" performs an *internal* redirect to the specified URI;
- "X-Accel-Limit-Rate" sets the rate limit for transmission of a response to a client;
- "X-Accel-Buffering" enables or disables buffering of a response;
- $\bullet$  "X-Accel-Charset" sets the desired charset of a response.

## scgi\_intercept\_errors

| Syntax  | scgi_intercept_errors on   off;       |
|---------|---------------------------------------|
| Default | <pre>scgi_intercept_errors off;</pre> |
| Context | http, server, location                |

Determines whether SCGI-responses with codes greater than or equal to 300 should be passed to a client or be intercepted and redirected to Angie for processing with the *error page* directive.

## scgi limit rate

| Syntax  | scgi_limit_rate rate;         |
|---------|-------------------------------|
| Default | <pre>scgi_limit_rate 0;</pre> |
| Context | http, server, location        |

Limits the speed of reading the response from the SCGI server. The *rate* is specified in bytes per second and can contain variables.

0 disables rate limiting

## 1 Note

The limit is set per a request, and so if Angie simultaneously opens two connections to the SCGI server, the overall rate will be twice as much as the specified limit. The limitation works only if buffering of responses from the SCGI server is enabled.

## scgi max temp file size

| Syntax  | $scgi_max_temp_file_size \ size;$         |
|---------|---|
| Default | <pre>scgi_max_temp_file_size 1024m;</pre> |
| Context | http, server, location                    |

When buffering of responses from the SCGI server is enabled, and the whole response does not fit into the buffers set by the  $scgi\_buffer\_size$  and  $scgi\_buffers$  directives, a part of the response can be saved to a temporary file. This directive sets the maximum size of the temporary file. The size of data written to the temporary file at a time is set by the  $scgi\_temp\_file\_write\_size$  directive.

0 disables buffering of responses to temporary files





This restriction does not apply to responses that will be cached or *stored on disk*.

# scgi\_next\_upstream

| Syntax  | scgi_next_upstream error   timeout   invalid_header   http_500   http_503 |
|---------|---|
|         | http_403   http_404   http_429   non_idempotent   off;                    |
| Default | <pre>scgi_next_upstream error timeout;</pre>                              |
| Context | http, server, location  |

Specifies in which cases a request should be passed to the next server in the upstream pool:

| error          | an error occurred while establishing a connection with the server, passing a request to it, or reading the response header;  |
|----------------|--|
| timeout        | a timeout has occurred while establishing a connection with the server, passing a request to it, or reading the response header;   |
| invalid_header | a server returned an empty or invalid response;  |
| http_500       | a server returned a response with the code 500;  |
| http_503       | a server returned a response with the code 503;  |
| http_403       | a server returned a response with the code 403;  |
| http_404       | a server returned a response with the code 404;  |
| http_429       | a server returned a response with the code 429;  |
| non_idempotent | normally, requests with a non-idempotent method (POST, LOCK, PATCH) are not passed to the next server if a request has been sent to an upstream server; enabling this option explicitly allows retrying such requests; |
| off            | disables passing a request to the next server.   |

# 1 Note

One should bear in mind that passing a request to the next server is only possible if nothing has been sent to a client yet. That is, if an error or timeout occurs in the middle of the transferring of a response, fixing this is impossible.

The directive also defines what is considered an unsuccessful attempt of communication with a server.

| error<br>timeout<br>invalid_header | always considered unsuccessful attempts, even if they are not specified in the directive |
|------------------------------------|--|
| http_500<br>http_503<br>http_429   | considered unsuccessful attempts only if they are specified in the directive             |
| http_403<br>http_404               | never considered unsuccessful attempts   |

Passing a request to the next server can be limited by the *number of tries* and by time.



## scgi\_next\_upstream\_timeout

| Syntax  | ${\tt scgi\_next\_upstream\_timeout}\ time;$ |
|---------|--|
| Default | <pre>scgi_next_upstream_timeout 0;</pre>     |
| Context | http, server, location                       |

Limits the time during which a request can be passed to the *next* server.

# scgi\_next\_upstream tries

| Syntax  | scgi_next_upstream_tries $number;$     |
|---------|--|
| Default | <pre>scgi_next_upstream_tries 0;</pre> |
| Context | http, server, location                 |

Limits the number of possible tries for passing a request to the *next* server.

| 0 turns off this limitation |  |
|-----------------------------|--|
|-----------------------------|--|

## scgi\_no\_cache

| Syntax  | scgi_no_cache string;  |
|---------|------------------------|
| Default | _                      |
| Context | http, server, location |

Defines conditions under which the response will not be saved to a cache. If at least one value of the string parameters is not empty and is not equal to "0" then the response will not be saved:

Can be used along with the  $scgi\_cache\_bypass$  directive.

# scgi param

| Syntax  | <pre>scgi_param parameter value [if_not_empty];</pre> |
|---------|---|
| Default | _   |
| Context | http, server, location                                |

Sets a parameter that should be passed to the SCGI server. The value can contain text, variables, and their combination. These directives are inherited from the previous configuration level if and only if there are no scgi param directives defined on the current level.

Standard CGI environment variables should be provided as SCGI headers, see the scgi\_params file provided in the distribution:

```
location / {
   include scgi_params;
# ...
}
```



If the directive is specified with if\_not\_empty then such a parameter will be passed to the server only if its value is not empty:

```
scgi_param HTTPS $https if_not_empty;
```

## scgi pass

| Syntax  | $\mathtt{scgi\_pass}\;uri;$            |
|---------|--|
| Default | _                                      |
| Context | location, if in location, limit_except |

Sets the address of an SCGI server. The address can be specified as a domain name or IP address, and an optional port:

```
scgi_pass localhost:9000;
```

or as a UNIX domain socket path specified after the word unix:

```
scgi_pass unix:/tmp/scgi.socket;
```

If a domain name resolves to several addresses, all of them will be used in a round-robin fashion. In addition, an address can be specified as a *server group*. If a group is used, you cannot specify the port with it; instead, specify the port for each server within the group individually.

Parameter value can contain variables. In this case, if an address is specified as a domain name, the name is searched among the described server groups, and, if not found, is determined using a *resolver*.

# scgi\_pass\_header

| Syntax  | scgi_pass_header field; |
|---------|-------------------------|
| Default | _                       |
| Context | http, server, location  |

Permits passing otherwise disabled header fields from a SCGI server to a client.

## scgi\_pass\_request\_body

| Syntax  | scgi_pass_request_body on   off;      |
|---------|---------------------------------------|
| Default | <pre>scgi_pass_request_body on;</pre> |
| Context | http, server, location                |

Indicates whether the original request body is passed to the SCGI server. See also the scgi pass request headers directive.

## scgi pass request headers

| Syntax  | scgi_pass_request_headers on   off;      |
|---------|--|
| Default | <pre>scgi_pass_request_headers on;</pre> |
| Context | http, server, location                   |

Indicates whether the header fields of the original request are passed to the SCGI server. See also the  $scgi\_pass\_request\_body$  directive.



# scgi\_read\_timeout

| Syntax  | ${\tt scgi\_read\_timeout}\ time;$ |
|---------|------------------------------------|
| Default | <pre>scgi_read_timeout 60s;</pre>  |
| Context | http, server, location             |

Defines a timeout for reading a response from the SCGI server. The timeout is set only between two successive read operations, not for the transmission of the whole response. If the SCGI server does not transmit anything within this time, the connection is closed.

# scgi request buffering

| Syntax  | scgi_request_buffering on   off; |
|---------|----------------------------------|
| Default | scgi_request_buffering on;       |
| Context | http, server, location           |

Enables or disables buffering of a client request body.

| on  | the entire request body is $read$ from the client before sending the request to a SCGI server.   |
|-----|--|
| off | the request body is sent to the SCGI server immediately as it is received. In this case, the request cannot be passed to the <i>next server</i> if Angie already started sending the request body. |

When HTTP/1.1 chunked transfer encoding is used to send the original request body, the request body will be buffered regardless of the directive value.

# scgi\_send\_timeout

| Syntax  | ${	t scgi\_send\_timeout} \ time;$ |
|---------|------------------------------------|
| Default | <pre>scgi_send_timeout 60s;</pre>  |
| Context | http, server, location             |

Sets a timeout for transmitting a request to the SCGI server. The timeout is set only between two successive write operations, not for the transmission of the whole request. If the SCGI server does not receive anything within this time, the connection is closed.

# scgi\_socket\_keepalive

| Syntax  | scgi_socket_keepalive on   off;       |
|---------|---------------------------------------|
| Default | <pre>scgi_socket_keepalive off;</pre> |
| Context | http, server, location                |

Configures the "TCP keepalive" behavior for outgoing connections to a SCGI server.

| 11 11 | By default, the operating system's settings are in effect for the socket. |
|-------|---|
| on    | The SO_KEEPALIVE socket option is turned on for the socket.               |



## scgi\_store

| Syntax  | $scgi\_store on \mid off \mid string;$ |
|---------|--|
| Default | scgi_store off;                        |
| Context | http, server, location                 |

Enables saving of files to a disk.

| on  | saves files with paths corresponding to the directives alias or root |
|-----|--|
| off | disables saving of files   |

The file name can be set explicitly using the string with variables:

```
scgi_store /data/www$original_uri;
```

The modification time of files is set according to the received "Last-Modified" response header field. The response is first written to a temporary file, and then the file is renamed. Temporary files and the persistent store can be put on different file systems. However, be aware that in this case a file is copied across two file systems instead of the cheap renaming operation. It is thus recommended that for any given location both saved files and a directory holding temporary files, set by the  $scgi\_temp\_path$  directive, are put on the same file system.

This directive can be used to create local copies of static unchangeable files, e.g.:

```
location /images/ {
                       /data/www;
    root
                      404 = /fetch\$uri;
    error_page
}
location /fetch/ {
    internal;
                     backend:9000;
    scgi_pass
    . . .
    scgi_store
                      on;
    scgi_store_access user:rw group:rw all:r;
    scgi_temp_path
                      /data/temp;
    alias
                       /data/www/;
}
```

## scgi\_store\_access

```
Syntax scgi_store_access users:permissions ...;
Default scgi_store_access user:rw;
Context http, server, location
```

Sets access permissions for newly created files and directories, e.g.:

```
scgi_store_access user:rw group:rw all:r;
```

If any group or all access permissions are specified then user permissions may be omitted:



```
scgi_store_access group:rw all:r;
```

## scgi temp file write size

| Syntax  | <pre>scgi_temp_file_write_size size;</pre>   |
|---------|--|
| Default | <pre>scgi_temp_file_write_size 8k 16k;</pre> |
| Context | http, server, location                       |

Limits the size of data written to a temporary file at a time, when buffering of responses from the SCGI server to temporary files is enabled. By default, size is limited by two buffers set by the  $scgi\_buffer\_size$  and  $scgi\_buffers$  directives. The maximum size of a temporary file is set by the  $scgi\_max\_temp\_file\_size$  directive.

## scgi\_temp\_path

| Syntax  | scgi_temp_path path [level1 [level2 [level3]]]`;   |
|---------|--|
| Default | <pre>scgi_temp_path scgi_temp; (the path depends on thehttp-scgi-temp-path build option)</pre> |
| Context | http, server, location   |

Defines a directory for storing temporary files with data received from SCGI servers. Up to three-level subdirectory hierarchy can be used underneath the specified directory. For example, in the following configuration

```
scgi_temp_path /spool/angie/scgi_temp 1 2;
```

a temporary file might look like this:

```
/spool/angie/scgi_temp/7/45/00000123457
```

See also the use\_temp\_path parameter of the scgi\_cache\_path directive.

## Secure Link

The module is used to check authenticity of requested links, protect resources from unauthorized access, and limit link lifetime.

The authenticity of a requested link is verified by comparing the checksum value passed in a request with the value computed for the request. If a link has a limited lifetime and the time has expired, the link is considered outdated. The status of these checks is made available in the \$secure\_link variable.

The module provides two alternative operation modes. The first mode is enabled by the secure\_link\_secret directive and is used to check authenticity of requested links as well as protect resources from unauthorized access. The second mode is enabled by the secure\_link and secure\_link\_md5 directives and is also used to limit lifetime of links.

When building from the source code, this module isn't built by default; it should be enabled with the --with-http\_secure\_link\_module build option.

In packages and images from our repos, the module is included in the build.

### **Directives**



# secure\_link

| Syntax  | secure_link expression; |
|---------|-------------------------|
| Default | _                       |
| Context | http, server, location  |

Defines a string with variables from which the checksum value and lifetime of a link will be extracted.

Variables used in an expression are usually associated with a request; see example below.

The checksum value extracted from the string is compared with the MD5 hash value of the expression defined by the  $secure\ link\ md5$  directive.

If the checksums are different, the  $\$secure\_link$  variable is set to an empty string. If the checksums are the same, the link lifetime is checked.

If the link has a limited lifetime and the time has expired, the *\$secure\_link* variable is set to 0. Otherwise, it is set to 1. The MD5 hash value passed in a request is encoded in base64url.

If a link has a limited lifetime, the expiration time is set in seconds since Epoch (Thu, 01 Jan 1970 00:00:00 GMT). The value is specified in the expression after the MD5 hash, and is separated by a comma. The expiration time passed in a request is available through the  $\$secure\_link\_expires$  variable for a use in the  $secure\_link\_md5$  directive. If the expiration time is not specified, a link has the unlimited lifetime.

## secure link md5

| Syntax  | secure_link_md5 expression; |
|---------|-----------------------------|
| Default | _                           |
| Context | http, server, location      |

Defines an expression for which the MD5 hash value will be computed and compared with the value passed in a request.

The expression should contain the secured part of a link (resource) and a secret ingredient. If the link has a limited lifetime, the expression should also contain \$secure link expires.

To prevent unauthorized access, the expression may contain some information about the client, such as its address and browser version.

# Example:

```
location /s/ {
    secure_link $arg_md5,$arg_expires;
    secure_link_md5 "$secure_link_expires$uri$remote_addr secret";

if ($secure_link = "") {
    return 403;
}

if ($secure_link = "0") {
    return 410;
}
```

The "/s/link?md5=\_e4Nc3iduzkWRm01TBBNYw&expires=2147483647" link restricts access to "/s/link" for the client with the IP address 127.0.0.1. The link also has the limited lifetime until January 19, 2038 (GMT).



On UNIX, the md5 request argument value can be obtained as:

```
echo -n '2147483647/s/link127.0.0.1 secret' | \
openssl md5 -binary | openssl base64 | tr +/ -_ | tr -d =
```

## secure link secret

| Syntax  | secure_link_secret word; |
|---------|--------------------------|
| Default | _                        |
| Context | location                 |

Defines a secret word used to check authenticity of requested links.

The full URI of a requested link looks as follows:

```
/prefix/hash/link
```

where hash is a hexadecimal representation of the MD5 hash computed for the concatenation of the link and secret word, and prefix is an arbitrary string without slashes.

If the requested link passes the authenticity check, the  $\$secure\_link$  variable is set to the link extracted from the request URI. Otherwise, the  $\$secure\_link$  variable is set to an empty string.

Example:

```
location /p/ {
    secure_link_secret secret;

if ($secure_link = "") {
    return 403;
  }

rewrite ^ /secure/$secure_link;
}

location /secure/ {
    internal;
}
```

A request of "/p/5e814704a28d9bc1914ff19fa0c4a00a/link" will be internally redirected to "/secure/link".

On UNIX, the hash value for this example can be obtained as:

```
echo -n 'linksecret' | openssl md5 -hex
```

## **Built-in Variables**

## \$secure\_link

The status of a link check. The specific value depends on the selected operation mode.

### \$secure\_link\_expires

The lifetime of a link passed in a request; intended to be used only in the secure link md5 directive.



#### **Slice**

The module is a filter that splits a request into subrequests, each returning a certain range of response. The filter provides more effective caching of big responses.

When building from the source code, this module isn't built by default; it should be enabled with the --with-http\_slice\_module build option.

In packages and images from our repos, the module is included in the build.

### **Configuration Example**

In this example, the response is split into 1-megabyte cacheable slices.

#### **Directives**

#### slice

```
Syntax slice size;
Default slice 0;
Context http, server, location
```

Sets the size of the slice. The zero value disables splitting responses into slices.

# **A** Warning

Note that a too low value may result in excessive memory usage and opening a large number of files.

In order for a subrequest to return the required range, the *\$slice\_range* variable should be passed to the proxied server as the "Range" request header field. If *caching* is enabled, *\$slice\_range* should be added to the *cache key* and caching of responses with 206 status code should be *enabled*.

### **Built-in Variables**

```
$slice_range
```

The current slice range in HTTP byte range format, for example, bytes=0-1048575.

## **Split Clients**

The module emits variables for A/B testing, canary releases, or other scenarios that require routing a percentage of clients to one server or configuration while directing the rest elsewhere.

# Configuration Example



#### **Directives**

# split\_clients

| Syntax  | $	exttt{split_clients} \ string \ 	exttt{$\$ variable} \ 	exttt{ } \ldots \ 	exttt{$\}}$ |
|---------|--|
| Default | _  |
| Context | http   |

Creates a *\$variable* by hashing the *string*; the variables in *string* are substituted, the result is hashed, then the hash is mapped to the *\$variable*'s string value.

The hash function uses MurmurHash2 (32-bit), and its entire value range (0 to 4294967295) is mapped to buckets in order of appearance; the percentages determine the size of the buckets. A wildcard (\*) may occur last; hashes that fall outside other buckets are mapped to its assigned value.

An example:

Here, the hashed values of the interpolated \$remote\_addrAAA string are distributed as follows:

- $\bullet$  values 0 to 21474835 (a sample of 0.5%) yield .one
- values 21474836 to 107374180 (a sample of 2%) yield .two
- values 107374181 to 4294967295 (all other values) yield "" (an empty string)

# SSI

The module is a filter that processes SSI (Server Side Includes) commands in responses passing through it.

# **Configuration Example**

```
location / {
    ssi on;
# ...
}
```

### **Directives**

ssi

| Syntax  | ssi on   off;                          |
|---------|--|
| Default | ssi off;                               |
| Context | http, server, location, if in location |



Enables or disables processing of SSI commands in responses.

## ssi last modified

| Syntax  | ssi_last_modified on   off; |
|---------|-----------------------------|
| Default | ssi_last_modified off;      |
| Context | http, server, location      |

Allows preserving the "Last-Modified" header field from the original response during SSI processing to facilitate response caching.

By default, the header field is removed as contents of the response are modified during processing and may contain dynamically generated elements or parts that are changed independently of the original response.

# ssi\_min\_file\_chunk

| Syntax  | ssi_min_file_chunk $size;$ |
|---------|----------------------------|
| Default | ssi_min_file_chunk 1k;     |
| Context | http, server, location     |

Sets the minimum size for parts of a response stored on disk, starting from which it makes sense to send them using sendfile.

# ssi\_silent\_errors

| Syntax  | ssi_silent_errors on   off; |
|---------|-----------------------------|
| Default | ssi_silent_errors off;      |
| Context | http, server, location      |

If enabled, suppresses the output of the "[an error occurred while processing the directive]" string if an error occurred during SSI processing.

## ssi types

| Syntax  | ssi_types mime-type;   |
|---------|------------------------|
| Default | ssi_types text/html;   |
| Context | http, server, location |

Enables processing of SSI commands in responses with the specified MIME types in addition to text/html. The special value "\*" matches any MIME type.

# ssi\_value\_length

| Syntax  | $	exttt{ssi_value_length} \ length;$ |
|---------|--------------------------------------|
| Default | ssi_value_length 256;                |
| Context | http, server, location               |

Sets the maximum length of parameter values in SSI commands.



#### SSI Commands

SSI commands have the following generic format:

```
<!--# command parameter1=value1 parameter2=value2 ... -->
```

The following commands are supported:

#### block

Defines a block that can be used as a stub in the *include* command. The block can contain other SSI commands. The command has the following parameter:

#### name

block name.

Example:

```
<!--# block name="one" -->
stub
<!--# endblock -->
```

### config

Sets some parameters used during SSI processing, namely:

#### errmsg

a string that is output if an error occurs during SSI processing. By default, the following string is output:

[an error occurred while processing the directive]

## timefmt

a format string passed to the strftime() function used to output date and time. By default, the following format is used:

```
"%A, %d-%b-%Y %H:%M:%S %Z"
```

The "%s" format is suitable to output time in seconds.

## echo

Outputs the value of a variable. The command has the following parameters:

### var

the variable name.

### encoding

the encoding method. Possible values include none, url, and entity.

By default, entity is used.

## default

a non-standard parameter that sets a string to be output if a variable is undefined. By default, (none) is output. The command



```
<!--# echo var="name" default="no" -->
```

replaces the following sequence of commands:

```
<!--# if expr="$name" --> <!--# echo var="name" --> <!--# else --> no <!--# endif -->
```

if

Performs a conditional inclusion. The following commands are supported:

```
<!--# if expr="..." -->
...
<!--# elif expr="..." -->
...
<!--# else -->
...
<!--# endif -->
```

Only one level of nesting is currently supported. The command has the following parameter:

expr

expression. An expression can be:

• variable existence check:

```
<!--# if expr="$name" -->
```

• comparison of a variable with a text:

```
<!--# if expr="$name = text" -->
<!--# if expr="$name != text" -->
```

• comparison of a variable with a regular expression:

```
<!--# if expr="$name = /text/" -->
<!--# if expr="$name != /text/" -->
```

If a *text* contains variables, their values are substituted. A regular expression can contain positional and named captures that can later be used through variables, for example:

include

Includes the result of another request into a response. The command has the following parameters:

file

specifies an included file, for example:

```
<!--# include file="footer.html" -->
```



#### virtual

specifies an included request, for example:

```
<!--# include virtual="/remote/body.php?argument=value" -->
```

Several requests specified on one page and processed by proxied or FastCGI/uwsgi/SCGI/gRPC servers run in parallel. If sequential processing is desired, the wait parameter should be used.

#### stub

a non-standard parameter that names the block whose content will be output if the included request results in an empty body or if an error occurs during the request processing, for example:

```
<!--# block name="one" -->&nbsp;<!--# endblock -->
<!--# include virtual="/remote/body.php?argument=value" stub="one" -->
```

The replacement block content is processed in the included request context.

#### wait

a non-standard parameter that instructs to wait for a request to fully complete before continuing with SSI processing, for example:

```
<!--# include virtual="/remote/body.php?argument=value" wait="yes" -->
```

#### set

a non-standard parameter that instructs to write a successful result of request processing to the specified variable, for example:

```
<!--# include virtual="/remote/body.php?argument=value" set="one" -->
```

The maximum size of the response is set by the subrequest output buffer size directive:

```
location /remote/ {
    subrequest_output_buffer_size 64k;
# ...
}
```

#### set

Sets a value of a variable. The command has the following parameters:

#### var

the variable name.

## value

the variable value. If an assigned value contains variables, their values are substituted.

### **Built-in Variables**

#### \$date\_local

current time in the local time zone. The format is set by the *config* command with the *timefmt* parameter.



#### \$date\_gmt

current time in GMT. The format is set by the *config* command with the *timefmt* parameter.

#### **SSL**

Provides the necessary support for HTTPS.

When building from the source code, this module isn't built by default; it should be enabled with the --with-http\_ssl\_module build option.

In packages and images from our repos, the module is included in the build.

## Important

This module requires the OpenSSL library.

## **Configuration Example**

To reduce the processor load it is recommended to

- set the number of worker processes equal to the number of processors,
- enable *keep-alive* connections,
- $\bullet$  enable the *shared* session cache,
- disable the *built-in* session cache,
- and possibly increase the session *lifetime* (by default, 5 minutes):

```
worker_processes auto;
http {
    # ...
    server {
        listen
                            443 ssl;
        keepalive_timeout 70;
        ssl_protocols
                           TLSv1.2 TLSv1.3;
                            AES128-SHA: AES256-SHA: RC4-SHA: DES-CBC3-SHA: RC4-MD5;
        ssl_ciphers
        ssl_certificate
                          /usr/local/angie/conf/cert.pem;
        ssl_certificate_key /usr/local/angie/conf/cert.key;
        ssl_session_cache shared:SSL:10m;
        ssl_session_timeout 10m;
   }
```

#### **Directives**

## ssl buffer size

```
Syntax ssl_buffer_size size;
Default ssl_buffer_size 16k;
Context http, server
```

Sets the size of the buffer used for sending data.



By default, the buffer size is 16k, which corresponds to minimal overhead when sending big responses. To minimize Time To First Byte it may be beneficial to use smaller values, for example:

```
ssl_buffer_size 4k;
```

## ssl\_certificate

```
Syntax ssl_certificate file;
Default —
Context http, server
```

Specifies a file with the certificate in the PEM format for the given virtual server. If intermediate certificates should be specified in addition to a primary certificate, they should be specified in the same file in the following order: the primary certificate comes first, then the intermediate certificates. A secret key in the PEM format may be placed in the same file.

This directive can be specified multiple times to load certificates of different types, for example, RSA and ECDSA:

Only OpenSSL 1.0.2 or higher supports separate certificate chains for different certificates. With older versions, only one certificate chain can be used.

Note that using variables implies that a certificate will be loaded for each SSL handshake, and this may have a negative impact on performance.

The value data: \$variable can be specified instead of the file, which loads a certificate from a variable without using intermediate files. Note that inappropriate use of this syntax may have its security implications, such as writing secret key data to error log.

#### Important

It should be kept in mind that due to the HTTPS protocol limitations for maximum interoperability virtual servers should listen on different IP addresses.

Added in version 1.2.0: If  $ssl\_ntls$  is enabled, the directive can accept two arguments (the signature and the encryption parts of the key) instead of one:



## ssl certificate cache

| Syntax  | $ssl\_certificate\_cache off;$ $ssl\_certificate\_cache max=N [inactive=time] [valid=time];$ |
|---------|--|
| Default | ssl_certificate_cache off;   |
| Context | http, server   |

Defines a cache that stores SSL certificates and secret keys specified using variables.

The directive supports the following parameters:

- max sets the maximum number of elements in the cache. When the cache overflows, the least recently used (LRU) elements are removed.
- inactive defines the time after which an element is removed if it has not been accessed. The default is 10 seconds.
- valid defines the time during which a cached element is considered valid and can be reused. The default is 60 seconds. After this period, certificates are reloaded or revalidated.
- off disables the cache.

#### Example:

```
ssl_certificate $ssl_server_name.crt;
ssl_certificate_key $ssl_server_name.key;
ssl_certificate_cache max=1000 inactive=20s valid=1m;
```

## ssl certificate key

```
Syntax ssl_certificate_key file;
Default —
Context http, server
```

Specifies a file with the secret key in the PEM format for the given virtual server.

# Uniables can be used in the file name when using OpenSSL 1.0.2 or higher.

The value engine:name:id can be specified instead of the file, which loads a secret key with a specified id from the OpenSSL engine name.



The value data: \$variable can be specified instead of the file, which loads a secret key from a variable without using intermediate files. Note that inappropriate use of this syntax may have its security implications, such as writing secret key data to error log.

Added in version 1.2.0: If ssl\_ntls is enabled, the directive can accept two arguments (the signature and the encryption parts of the key) instead of one:

## ssl ciphers

```
Syntax ssl_ciphers ciphers;
Default ssl_ciphers HIGH:!aNULL:!MD5;
Context http, server
```

Specifies the enabled ciphers. The ciphers are specified in the format understood by the OpenSSL library, for example:

```
ssl_ciphers ALL:!aNULL:!EXPORT56:RC4+RSA:+HIGH:+MEDIUM:+LOW:+SSLv2:+EXP;
```

The list of ciphers depends on the version of OpenSSL installed. The full list can be viewed using the openssl ciphers command.

## **A** Attention

The ssl\_ciphers directive does not configure ciphers for TLS 1.3 when using OpenSSL. To tune TLS 1.3 ciphers with OpenSSL, use the ssl\_conf\_command directive, which was added to support advanced SSL configuration.

- In LibreSSL, TLS 1.3 ciphers can be configured using ssl\_ciphers.
- In BoringSSL, TLS 1.3 ciphers cannot be configured at all.

## ssl client certificate

| Syntax  | ssl_client_certificate file; |
|---------|------------------------------|
| Default | _                            |
| Context | http, server                 |

Specifies a file with trusted CA certificates in the PEM format used to *verify* client certificates and OCSP responses if *ssl\_stapling* is enabled.

The list of certificates will be sent to clients. If this is not desired, the  $ssl\_trusted\_certificate$  directive can be used.



## ssl\_conf\_command

| Syntax  | ssl_conf_command name value; |
|---------|------------------------------|
| Default | _                            |
| Context | http, server                 |

Sets arbitrary OpenSSL configuration commands.

## Important

The directive is supported when using OpenSSL 1.0.2 or higher. To configure TLS 1.3 ciphers with OpenSSL, use the ciphersuites command.

Several  $ssl\_conf\_command$  directives can be specified on the same level:

```
ssl_conf_command Options PrioritizeChaCha; ssl_conf_command Ciphersuites TLS_CHACHA20_POLY1305_SHA256;
```

These directives are inherited from the previous configuration level if and only if there are no  $ssl\_conf\_command$  directives defined on the current level.

## 🛠 Caution

Configuring OpenSSL directly might result in unexpected behavior.

## ssl\_crl

| Syntax  | ssl_crl file; |
|---------|---------------|
| Default | _             |
| Context | http, server  |

Specifies a file with revoked certificates (CRL) in the PEM format used to verify client certificates.

#### ssl dhparam

| Syntax  | ${	t ssl\_dhparam} \ file;$ |
|---------|-----------------------------|
| Default | _                           |
| Context | http, server                |

Specifies a file with DH parameters for DHE ciphers.

By default no parameters are set, and therefore DHE ciphers will not be used.

## ssl\_early\_data

| Syntax  | ssl_early_data on   off; |
|---------|--------------------------|
| Default | ssl_early_data off;      |
| Context | http, server             |

Enables or disables TLS 1.3 early data.



Requests sent within early data are subject to replay attacks. To protect against such attacks at the application layer, the \$ssl\_early\_data variable should be used.

```
proxy_set_header Early-Data $ssl_early_data;
```

## Important

The directive is supported when using OpenSSL 1.1.1 or higher and BoringSSL.

## ssl\_ecdh\_curve

| Syntax  | ssl_ecdh_curve curve; |
|---------|-----------------------|
| Default | ssl_ecdh_curve auto;  |
| Context | http, server          |

Specifies a curve for ECDHE ciphers.

## Important

When using OpenSSL 1.0.2 or higher, it is possible to specify multiple curves, for example:

```
ssl_ecdh_curve prime256v1:secp384r1;
```

The special value auto instructs Angie to use a list built into the OpenSSL library when using OpenSSL 1.0.2 or higher, or *prime256v1* with older versions.

## Important

When using OpenSSL 1.0.2 or higher, this directive sets the list of curves supported by the server. Thus, in order for ECDSA certificates to work, it is important to include the curves used in the certificates.

## ssl ntls

Added in version 1.2.0.

```
Syntax ssl_ntls on | off;
Default ssl_ntls off;
Context http, server
```

Enables server-side support for NTLS using TongSuo library.

```
listen ... ssl;
ssl_ntls on;
```

#### Important

Build Angie using the --with-ntls build option and link with NTLS-enabled SSL library



## ssl ocsp

| Syntax  | ssl_ocsp on   off   leaf; |
|---------|---------------------------|
| Default | ssl_ocsp off;             |
| Context | http, server              |

Enables OCSP validation of the client certificate chain. The leaf parameter enables validation of the client certificate only.

For the OCSP validation to work, the ssl\_verify\_client directive should be set to on or optional.

To resolve the OCSP responder hostname, the resolver directive should also be specified.

#### Example:

## ssl ocsp cache

| Syntax  | ssl_ocsp_cache off   [shared:name:size]; |
|---------|--|
| Default | ssl_ocsp_cache off;                      |
| Context | http, server                             |

Sets name and size of the cache that stores client certificates status for OCSP validation. The cache is shared between all worker processes. A cache with the same name can be used in several virtual servers.

The off parameter prohibits the use of the cache.

#### ssl ocsp responder

| Syntax  | ssl_ocsp_responder uri; |
|---------|-------------------------|
| Default | _                       |
| Context | http, server            |

Overrides the URI of the OCSP responder specified in the "Authority Information Access" certificate extension for validation of client certificates.

Only http:// OCSP responders are supported:

```
ssl_ocsp_responder http://ocsp.example.com/;
```

## ssl password file

| Syntax  | ${	t ssl\_password\_file}$ ; |
|---------|------------------------------|
| Default | _                            |
| Context | http, server                 |

Specifies a file with passphrases for *secret keys* where each passphrase is specified on a separate line. Passphrases are tried in turn when loading the key.



## Example:

```
http {
    ssl_password_file /etc/keys/global.pass;
    ...

server {
    server_name www1.example.com;
    ssl_certificate_key /etc/keys/first.key;
}

server {
    server_name www2.example.com;

# named pipe can also be used instead of a file
    ssl_password_file /etc/keys/fifo;
    ssl_certificate_key /etc/keys/second.key;
}
```

## ssl prefer server ciphers

| Syntax  | ssl_prefer_server_ciphers on   off; |
|---------|-------------------------------------|
| Default | ssl_prefer_server_ciphers off;      |
| Context | http, server                        |

Specifies that server ciphers should be preferred over client ciphers when using the SSLv3 and TLS protocols.

## ssl protocols

| Syntax  | ssl_protocols [SSLv2] [SSLv3] [TLSv1.1] [TLSv1.2] [TLSv1.3]; |
|---------|--|
| Default | ssl_protocols TLSv1.2 TLSv1.3;                               |
| Context | http, server   |

Changed in version 1.2.0: TLSv1.3 parameter added to default set.

Enables the specified protocols.

## Important

The TLSv1.1 and TLSv1.2 parameters work only when OpenSSL 1.0.1 or higher is used.

The TLSv1.3 parameter works only when OpenSSL 1.1.1 or higher is used.

#### ssl reject handshake

| Syntax  | ssl_reject_handshake on   off; |
|---------|--------------------------------|
| Default | ssl_reject_handshake off;      |
| Context | http, server                   |

If enabled, SSL handshakes in the server block will be rejected.



For example, in the following configuration, SSL handshakes with server names other than *example.com* are rejected:

## ssl session cache

| Syntax  | ssl_session_cache off   none   [builtin[:size]] [shared:name:size]; |
|---------|---|
| Default | ssl_session_cache none;   |
| Context | http, server  |

Sets the types and sizes of caches that store session parameters. A cache can be of any of the following types:

| off     | the use of a session cache is strictly prohibited: Angie explicitly tells a client that sessions may not be reused.  |
|---------|--|
| none    | the use of a session cache is gently disallowed: Angie tells a client that sessions may be reused, but does not actually store session parameters in the cache.  |
| builtin | a cache built in OpenSSL; used by one worker process only. The cache size is specified in sessions. If size is not given, it is equal to 20480 sessions. Use of the built-in cache can cause memory fragmentation.   |
| shared  | a cache shared between all worker processes. The cache size is specified in bytes; one megabyte can store about 4000 sessions. Each shared cache should have an arbitrary name. A cache with the same name can be used in several virtual servers. It is also used to automatically generate, store, and periodically rotate TLS session ticket keys unless configured explicitly using the $ssl\_session\_ticket\_key$ directive. |

Both cache types can be used simultaneously, for example:

```
ssl_session_cache builtin:1000 shared:SSL:10m;
```

but using only shared cache without the built-in cache should be more efficient.

## ssl\_session\_ticket\_key

| Syntax  | ssl_session_ticket_key file; |
|---------|------------------------------|
| Default | _                            |
| Context | http, server                 |

Sets a file with the secret key used to encrypt and decrypt TLS session tickets. The directive is necessary if the same key has to be shared between multiple servers. By default, a randomly generated key is used.

If several keys are specified, only the first key is used to encrypt TLS session tickets. This allows configuring key rotation, for example:



```
ssl_session_ticket_key current.key;
ssl_session_ticket_key previous.key;
```

The file must contain 80 or 48 bytes of random data and can be created using the following command:

```
openssl rand 80 > ticket.key
```

Depending on the file size either AES256 (for 80-byte keys) or AES128 (for 48-byte keys) is used for encryption.

## ssl session tickets

| Syntax  | ssl_session_tickets on   off; |
|---------|-------------------------------|
| Default | ssl_session_tickets on;       |
| Context | http, server                  |

Enables or disables session resumption through TLS session tickets.

## ssl session timeout

| Syntax  | ${\tt ssl\_session\_timeout}\ time;$ |
|---------|--------------------------------------|
| Default | ssl_session_timeout 5m;              |
| Context | http, server                         |

Specifies a time during which a client may reuse the session parameters.

## ssl stapling

| Syntax  | ssl_stapling on   off; |
|---------|------------------------|
| Default | ssl_stapling off;      |
| Context | http, server           |

Enables or disables stapling of OCSP responses by the server. Example:

```
ssl_stapling on;
resolver 127.0.0.53;
```

For the OCSP stapling to work, the certificate of the server certificate issuer should be known. If the  $ssl\_certificate$  file does not contain intermediate certificates, the certificate of the server certificate issuer should be present in the  $ssl\_trusted\_certificate$  file.

#### **A**ttention

For a resolution of the OCSP responder hostname, the resolver directive should also be specified.

## $ssl\_stapling\_file$

| Syntax  | ssl_stapling_file file; |
|---------|-------------------------|
| Default | _                       |
| Context | http, server            |



When set, the stapled OCSP response will be taken from the specified file instead of querying the OCSP responder specified in the server certificate.

The file should be in the DER format as produced by the openss1 ocsp command.

## ssl stapling responder

| Syntax  | ssl_stapling_responder uri; |
|---------|-----------------------------|
| Default | _                           |
| Context | http, server                |

Overrides the URI of the OCSP responder specified in the "Authority Information Access". certificate extension.

Only http:// OCSP responders are supported:

```
ssl_stapling_responder http://ocsp.example.com/;
```

## ssl stapling verify

| Syntax  | ssl_stapling_verify on   off; |
|---------|-------------------------------|
| Default | ssl_stapling_verify off;      |
| Context | http, server                  |

Enables or disables verification of OCSP responses by the server.

For verification to work, the certificate of the server certificate issuer, the root certificate, and all intermediate certificates should be configured as trusted using the *ssl trusted certificate* directive.

## ssl\_trusted\_certificate

| Syntax  | ${	t ssl\_trusted\_certificate}$ $file;$ |
|---------|--|
| Default | _  |
| Context | http, server                             |

Specifies a file with trusted CA certificates in the PEM format used to *verify* client certificates and OCSP responses if *ssl stapling* is enabled.

In contrast to the certificate set by  $ssl\_client\_certificate$ , the list of these certificates will not be sent to clients.

## ssl\_verify\_client

| Syntax  | ssl_verify_client on   off   optional   optional_no_ca; |
|---------|---|
| Default | ssl_verify_client off;                                  |
| Context | http, server  |

Enables verification of client certificates. The verification result is stored in the  $\$ssl\_client\_verify$  variable.

| optional       | requests the client certificate and verifies it if the certificate is present.       |  |  |
|----------------|--|--|--|
| optional_no_ca | requests the client certificate but does not require it to be signed by a trusted CA |  |  |
|                | certificate. This is intended for the use in cases when a service that is external   |  |  |
|                | to Angie performs the actual certificate verification.                               |  |  |



## ssl\_verify\_depth

| Syntax  | $	exttt{ssl\_verify\_depth } number;$ |
|---------|---------------------------------------|
| Default | ssl_verify_depth 1;                   |
| Context | http, server                          |

Sets the verification depth in the client certificates chain.

#### **Error Processing**

The http\_ssl module supports several non-standard error codes that can be used for redirects using the error page directive:

| 495 | an error has occurred during the client certificate verification; |
|-----|---|
| 496 | a client has not presented the required certificate;              |
| 497 | a regular request has been sent to the HTTPS port.                |

The redirection happens after the request is fully parsed and the variables, such as \$request\_uri, \$uri, \$args and others, are available.

#### **Built-in Variables**

The http ssl module supports built-in variables:

#### \$ssl\_alpn\_protocol

returns the protocol selected by ALPN during the SSL handshake, or an empty string otherwise.

#### \$ssl\_cipher

returns the name of the cipher used for an established SSL connection.

#### \$ssl\_ciphers

returns the list of ciphers supported by the client. Known ciphers are listed by names, unknown are shown in hexadecimal, for example:

AES128-SHA:AES256-SHA:0x00ff

# Important

The variable is fully supported only when using OpenSSL version 1.0.2 or higher. With older versions, the variable is available only for new sessions and lists only known ciphers.

#### \$ssl\_client\_escaped\_cert

returns the client certificate in the PEM format (urlencoded) for an established SSL connection;

#### \$ssl\_client\_fingerprint

returns the SHA1 fingerprint of the client certificate for an established SSL connection;



#### \$ssl\_client\_i\_dn

returns the "issuer DN" string of the client certificate for an established SSL connection according to RFC 2253;

#### \$ssl\_client\_i\_dn\_legacy

returns the "issuer DN" string of the client certificate for an established SSL connection.

#### \$ssl\_client\_raw\_cert

returns the client certificate in the PEM format for an established SSL connection.

## \$ssl\_client\_s\_dn

returns the "subject DN" string of the client certificate for an established SSL connection according to RFC 2253.

#### \$ssl\_client\_s\_dn\_legacy

returns the "subject DN" string of the client certificate for an established SSL connection.

#### \$ssl\_client\_serial

returns the serial number of the client certificate for an established SSL connection.

#### \$ssl\_client\_v\_end

returns the end date of the client certificate.

#### \$ssl\_client\_v\_remain

returns the number of days until the client certificate expires.

## \$ssl\_client\_v\_start

returns the start date of the client certificate.

#### \$ssl\_client\_verify

returns the result of client certificate verification: SUCCESS, FAILED:reason, and NONE if a certificate was not present.

#### \$ssl\_curve

returns the negotiated curve used for SSL handshake key exchange process. Known curves are listed by names, unknown are shown in hexadecimal, for example:

prime256v1

## Important

The variable is supported only when using OpenSSL version 3.0 or higher. With older versions, the variable value will be an empty string.



#### \$ssl\_curves

returns the list of curves supported by the client. Known curves are listed by names, unknown are shown in hexadecimal, for example:

0x001d:prime256v1:secp521r1:secp384r1

## Important

The variable is supported only when using OpenSSL version 1.0.2 or higher. With older versions, the variable value will be an empty string.

The variable is available only for new sessions.

#### \$ssl\_early\_data

returns "1" if TLS 1.3 early data is used and the handshake is not complete, otherwise "".

#### \$ssl\_protocol

returns the protocol of an established SSL connection.

```
$ssl_server_cert_type
```

takes the values RSA, DSA, ECDSA, ED448, ED25519, SM2, RSA-PSS, or unknown depending on the type of server certificate and key.

#### \$ssl\_server\_name

returns the server name requested through SNI;

```
$ssl_session_id
```

returns the session identifier of an established SSL connection.

#### \$ssl\_session\_reused

returns  ${\tt r}$  if an SSL session was reused, or "." otherwise.

#### Stub Status

The module provides access to basic status information.

When building from the source code, this module isn't built by default; it should be enabled with the --with-http\_stub\_status\_module build option.

In packages and images from our repos, the module is included in the build.

## **Configuration Example**

```
location = /basic_status {
    stub_status;
}
```

This configuration creates a simple web page with basic status data which may look like as follows:

```
Active connections: 291
server accepts handled requests
```



16630948 16630948 31070465

Reading: 6 Writing: 179 Waiting: 106

#### **Directives**

## stub\_status

| Syntax  | stub_status;     |
|---------|------------------|
| Default | _                |
| Context | server, location |

The basic status information will be accessible from the surrounding location.

#### Data

The following status information is provided:

#### Active connections

The current number of active client connections including Waiting connections.

#### accepts

The total number of accepted client connections.

## handled

The total number of handled connections. Generally, the parameter value is the same as accepts unless some resource limits have been reached (for example, the *worker\_connections* limit).

#### requests

The total number of client requests.

#### Reading

The current number of connections where Angie is reading the request header.

## Writing

The current number of connections where Angie is writing the response back to the client.

#### Waiting

The current number of idle client connections waiting for a request.

#### **Built-in Variables**

#### \$connections\_active

Same as the  $Active\ connections$  value.



#### \$connections\_reading

Same as the *Reading* value.

\$connections\_writing

Same as the Writing value.

\$connections\_waiting

Same as the Waiting value.

#### Sub

The module is a filter that modifies a response by replacing one specified string with another.

When building from the source code, this module isn't built by default; it should be enabled with the --with-http\_sub\_module build option.

In packages and images from our repos, the module is included in the build.

#### **Configuration Example**

```
location / {
    sub_filter '<a href="http://127.0.0.1:8080/' '<a href="https://$host/';
    sub_filter '<img src="http://127.0.0.1:8080/' '<img src="https://$host/';
    sub_filter_once on;
}</pre>
```

## **Directives**

## sub\_filter

| Syntax  | <pre>sub_filter string replacement;</pre> |
|---------|---|
| Default | _   |
| Context | http, server, location                    |

Sets a string to replace and a replacement string. The string to replace is matched ignoring the case. The string to replace and replacement string can contain variables. Several sub\_filter directives can be specified on the same configuration level. These directives are inherited from the previous configuration level if and only if there are no sub\_filter directives defined on the current level.

#### sub filter last modified

| Syntax  | <pre>sub_filter_last_modified on   off;</pre> |
|---------|---|
| Default | <pre>sub_filter_last_modified off;</pre>      |
| Context | http, server, location                        |

Allows preserving the "Last-Modified" header field from the original response during replacement to facilitate response caching.

By default, the header field is removed as contents of the response are modified during processing.



## sub\_filter\_once

| Syntax  | <pre>sub_filter_once on   off;</pre> |
|---------|--------------------------------------|
| Default | <pre>sub_filter_once on;</pre>       |
| Context | http, server, location               |

Indicates whether to look for each string to replace once or repeatedly.

#### sub filter types

| Syntax  | sub_filter_types mime-type;            |
|---------|--|
| Default | <pre>sub_filter_types text/html;</pre> |
| Context | http, server, location                 |

Enables string replacement in responses with the specified MIME types in addition to text/html. The special value "\*" matches any MIME type.

## **Upstream**

The module is used to define groups of servers that can be referenced by the proxy\_pass, fastcgi\_pass, uwsgi\_pass, scgi\_pass, memcached\_pass and grpc\_pass directives.

## **Configuration Example**

```
upstream backend {
    zone backend 1m;
    server backend1.example.com
                                       weight=5;
    server backend2.example.com:8080;
    server backend3.example.com
                                       service=_example._tcp resolve;
    server unix:/tmp/backend3;
    server backup1.example.com:8080
                                       backup;
    server backup2.example.com:8080
                                       backup;
}
resolver 127.0.0.53 status_zone=resolver;
server {
    location / {
        proxy_pass http://backend;
}
```

#### **Directives**

## backup\_switch (PRO)

Added in version 1.9.0: PRO

| Syntax  | backup_switch permanent`[=`time]; |
|---------|-----------------------------------|
| Default | _                                 |
| Context | upstream                          |

The directive enables active backups for the upstream where it occurs. Once a request fails to select



a server in the primary group and resorts to the backup group, that group will become *active* if the directive is defined. Subsequent requests are handled by first looking for servers in the active group.

When permanent is defined without a *time* value, the group remains active once selected, and no automatic reevaluation occurs. If the *time* limit is set, the active status times out after the specified interval. and the balancer reevaluates the primary group, reverting to it if the servers are healthy.

#### Example:

```
upstream my_backend {
    server primary1.example.com;
    server primary2.example.com;

server backup1.example.com backup;
    server backup2.example.com backup;

backup_switch permanent=2m;
}
```

If the balancer fails over from the primary servers to the backup group, all subsequent requests are served by that backup group for 2 minutes. Once 2 minutes elapse, the balancer reevaluates the primary servers and makes them active again if they're found to be healthy.

## bind conn (PRO)

```
Syntax bind_conn value;
Default —
Context upstream
```

Enables binding the server connection to the client when the *value*, which is set as a string of variables, becomes anything other than "" and "0".

## **A** Attention

The bind\_conn directive must be used after all directives that set the load balancing method; otherwise, it won't work. If *sticky* is also used, bind\_conn should appear after sticky.

#### **A** Attention

When using the directive, configure the  $http\_proxy$  module to allow keep alive connections, for example:

```
proxy_http_version 1.1;
proxy_set_header Connection "";
```

A typical use case for the directive is proxying NTLM-authenticated connections, where the client should be bound to the server when the negotiation starts:

```
map $http_authorization $ntlm {
    ~*^N(?:TLM|egotiate) 1;
}
upstream ntlm_backend {
    server 127.0.0.1:8080;
    bind_conn $ntlm;
```



```
server {
    # ...
    location / {
        proxy_pass http://ntlm_backend;
        proxy_http_version 1.1;
        proxy_set_header Connection "";
        # ...
}
```

## feedback (PRO)

Added in version 1.6.0: PRO

| Syntax  | <pre>feedback variable [last_byte];</pre> | [inverse] | [factor=number] | $[{\tt account} = condition\_variable]$ |
|---------|---|-----------|-----------------|---|
| Default | _   |           |                 |   |
| Context | upstream                                  |           |                 |   |

Enables a feedback-based load balancing mechanism for the upstream. It adjusts the load balancing decisions dynamically, multiplying each peer's weight by its average feedback value that is affected by the value of a *variable* over time and is subject to an optional condition.

The following parameters are accepted:

| variable  | The variable from which the feedback value is taken. It should represent a performance or health metric, and is intended to be supplied by the peer in header fields or otherwise.  The value is assessed at each response from the peer and factored into the rolling average according to inverse and factor settings.  |
|-----------|---|
| inverse   | If set, the feedback value is interpreted inversely, meaning lower values indicate better performance.  |
| factor    | The factor by which the feedback value is weighted when calculating the average. Valid values are integers between 0 and 99. By default — 90.  The average feedback is calculated using the exponential moving average formula. The larger is the factor, the less is the average affected by new values; if the factor is set to 90, the result has 90% of the previous value and only 10% of the new value. |
| account   | Specifies a condition variable that controls which responses should be included in the calculation. The average is updated with the feedback value only if the condition variable for the response isn't "" or "0".  Note  By default, responses from probes aren't included in the calculation; combining the \$upstream_probe variable with account allows to include these                                 |
|           | responses or even exclude everything else.  |
| last_byte | Allows processing feedback from the upstream server after the full response has been received, instead of just after the header.  |

Example:



```
upstream backend {
    zone backend 1m;
    feedback $feedback_value factor=80 account=$condition_value;
    server backend1.example.com;
    server backend2.example.com;
}
map $upstream_http_custom_score $feedback_value {
                                 100:
    "medium"
                                 75;
    "low"
                                 50;
    default
                                 10;
}
map $upstream_probe $condition_value {
    "high_priority" "1";
    "low_priority"
                    "0";
    default
                     "1";
}
```

This categorizes server responses into different feedback levels based on specific scores obtained from response header fields, and also adds a condition mapped from  $supstream\_probe$  to account only for the responses from the high\_priority probe or responses to regular client requests.

#### hash

| Syntax  | key [consistent]; |
|---------|-------------------|
| Default | _                 |
| Context | upstream          |

Specifies a load balancing method for a server group where the client-server mapping is based on the hashed key value. The key can contain text, variables, and their combinations. Note that adding or removing a server from the group may result in remapping most of the keys to different servers. The method is compatible with the Cache::Memcached Perl library.

If the consistent parameter is specified, the ketama consistent hashing method will be used instead. The method ensures that only a few keys will be remapped to different servers when a server is added to or removed from the group. This helps to achieve a higher cache hit ratio for caching servers. The method is compatible with the Cache::Memcached::Fast Perl library with the ketama\_points parameter set to 160.

## ip hash

| Syntax  | ip_hash; |
|---------|----------|
| Default | _        |
| Context | upstream |

Specifies that a group should use a load balancing method where requests are distributed between servers based on client IP addresses. The first three octets of the client IPv4 address, or the entire IPv6 address, are used as a hashing key. The method ensures that requests from the same client will always be passed to the same server except when this server is unavailable. In the latter case client requests will be passed to another server. Most probably, it will always be the same server as well.



If one of the servers needs to be temporarily removed, it should be marked with the down parameter in order to preserve the current hashing of client IP addresses.

```
upstream backend {
   ip_hash;

   server backend1.example.com;
   server backend2.example.com;
   server backend3.example.com down;
   server backend4.example.com;
}
```

#### keepalive

| Syntax  | keepalive connections; |
|---------|------------------------|
| Default | _                      |
| Context | upstream               |

Activates the cache for connections to upstream servers.

The connections parameter sets the maximum number of idle keepalive connections to upstream servers that are preserved in the cache of each worker process. When this number is exceeded, the least recently used connections are closed.

## 1 Note

It should be particularly noted that the *keepalive* directive does not limit the total number of connections to upstream servers that an Angie worker process can open. The connections parameter should be set to a number small enough to let upstream servers process new incoming connections as well.

## **A** Attention

The keepalive directive must be used after all directives that set the load balancing method; otherwise, it won't work.

Example configuration of memcached upstream with keepalive connections:

```
upstream memcached_backend {
    server 127.0.0.1:11211;
    server 10.0.0.2:11211;

    keepalive 32;
}
server {
    #...

    location /memcached/ {
        set $memcached_key $uri;
        memcached_pass memcached;
    }
}
```



For HTTP, the *proxy\_http\_version* directive should be set to "1.1" and the "Connection" header field should be cleared:

```
upstream http_backend {
    server 127.0.0.1:8080;

    keepalive 16;
}
server {
    #...

    location /http/ {
        proxy_pass http://http_backend;
        proxy_http_version 1.1;
        proxy_set_header Connection "";
    # ...
    }
}
```

## 1 Note

Alternatively,  $\mathrm{HTTP}/1.0$  persistent connections can be used by passing the "Connection: Keep-Alive" header field to an upstream server, though this method is not recommended.

For FastCGI servers, it is required to set <code>fastcgi\_keep\_conn</code> for keepalive connections to work:

```
upstream fastcgi_backend {
    server 127.0.0.1:9000;

    keepalive 8;
}
server {
    #...

    location /fastcgi/ {
        fastcgi_pass fastcgi_backend;
        fastcgi_keep_conn on;
    # ...
    }
}
```

#### 1 Note

SCGI and uwsgi protocols do not have a notion of keepalive connections.

## keepalive requests

| Syntax  | $keepalive\_requests\ number;$ |
|---------|--------------------------------|
| Default | keepalive_requests 1000;       |
| Context | upstream                       |

Sets the maximum number of requests that can be served through one keepalive connection. After the



maximum number of requests is made, the connection is closed.

Closing connections periodically is necessary to free per-connection memory allocations. Therefore, using too high maximum number of requests could result in excessive memory usage and not recommended.

## keepalive\_time

| Syntax  | $\verb keepalive_time  time;$ |
|---------|-------------------------------|
| Default | keepalive_time 1h;            |
| Context | upstream                      |

Limits the maximum time during which requests can be processed through one keepalive connection. After this time is reached, the connection is closed following the subsequent request processing.

## keepalive\_timeout

| Syntax  | $\verb keepalive_time  ime;$ |
|---------|------------------------------|
| Default | keepalive_timeout 60s;       |
| Context | upstream                     |

Sets a timeout during which an idle keepalive connection to an upstream server will stay open.

## least conn

| Syntax  | least_conn; |
|---------|-------------|
| Default | _           |
| Context | upstream    |

Specifies that a group should use a load balancing method where a request is passed to the server with the least number of active connections, taking into account weights of servers. If there are several such servers, they are tried in turn using a weighted round-robin balancing method.

## least time (PRO)

| Syntax  | <pre>least_time header   last_byte [factor=number] [account=condition_variable];</pre> |
|---------|--|
| Default | _  |
| Context | upstream   |

Specifies that the group should use a load balancing method where an active server's chance of receiving the request is inversely proportional to its average response time; the less it is, the more requests a server gets.

| header    | The directive accounts for response headers only.                   |
|-----------|---|
| last_byte | The directive uses the average time to receive the entire response. |

Added in version 1.7.0: PRO



| factor  | Serves the same purpose as $response\_time\_factor~(PRO)$ and overrides it if set.  |
|---------|---|
| account | Specifies a condition variable that controls which responses should be included in the calculation. The average is updated only if the condition variable for the response isn't "" or "0".                 |
|         | 1 Note  |
|         | By default, responses from <i>probes</i> aren't included in the calculation; combining the <i>\$upstream_probe</i> variable with account allows to include these responses or even exclude everything else. |

The respective moving averages, adjusted for factor and account, are also presented as header\_time and response\_time in the server's health object among the upstream metrics in the API.

#### queue (PRO)

Added in version 1.4.0: PRO

| Syntax  | queue $number$ [timeout= $time$ ]; |
|---------|------------------------------------|
| Default | _                                  |
| Context | upstream                           |

If it is not possible to assign a proxied server to a request on the first attempt (for example, during a brief service interruption or when there is a surge in load reaching the max conns limit), the request is not rejected; instead, Angie attempts to enqueue it for processing.

The number in the directive sets the maximum number of requests in the queue for a worker process. If the queue is full, a 502 (Bad Gateway) error is returned to the client.



## 1 Note

The logic of the proxy next upstream directive also applies to queued requests. Specifically, if a server was selected for a request but it cannot be handed over to it, the request may be returned to the queue.

If a server is not selected to process a queued request within the time set by timeout (default is 60 seconds), a 502 (Bad Gateway) error is returned to the client. Requests from clients that prematurely close the connection are also removed from the queue; there are counters for requests passing through the queue in the API.



#### Attention

The queue directive must be used after all directives that set the load balancing method; otherwise, it won't work.

## random

| Syntax  | random [two]; |
|---------|---------------|
| Default | _             |
| Context | upstream      |



Specifies that a group should use a load balancing method where a request is passed to a randomly selected server, taking into account weights of servers.

The optional two parameter instructs Angie to randomly select two servers and then choose a server using the specified method. The default method is *least\_conn* which passes a request to a server with the least number of active connections.

## response time factor (PRO)

| Syntax  | $response\_time\_factor\ number;$ |
|---------|-----------------------------------|
| Default | response_time_factor 90;          |
| Context | upstream                          |

If the *least\_time* (*PRO*) load balancing method is used, sets the smoothing factor for the **previous** value when average response time is calculated using the exponential moving average formula.

The larger is the number, the less is the average affected by new values; if the number is set to 90, the result has 90% of the previous value and only 10% of the new value. The allowed range is 0 to 99, inclusive.

The respective moving averages are presented as header\_time (headers only) and response\_time (entire responses) in the server's health object among the *upstream metrics* in the API.

## 1 Note

The calculation accounts for successful reponses only; what is considered an unsuccessful response is defined by the <code>proxy\_next\_upstream</code>, <code>fastcgi\_next\_upstream</code>, <code>uwsgi\_next\_upstream</code>, <code>scgi\_next\_upstream</code>, <code>memcached\_next\_upstream</code>, and <code>grpc\_next\_upstream</code> directives. Besides, <code>header\_time</code> is updated only if all headers are received and processed, and <code>response\_time</code> is updated only if the entire reponse is received.

#### server

| Syntax  | server address [parameters]; |
|---------|------------------------------|
| Default | _                            |
| Context | upstream                     |

Defines the address and other parameters of a server. The address can be specified as a domain name or IP address, with an optional port, or as a UNIX domain socket path specified after the unix: prefix. If a port is not specified, the port 80 is used. A domain name that resolves to several IP addresses defines multiple servers at once.

The following parameters can be defined:

| weight=number               | sets the weight of the server; by default, 1.                                       |
|-----------------------------|---|
| ${\tt max\_conns} = number$ | limits the maximum number of simultaneous active connections to the proxied         |
|                             | server. Default value is 0, meaning there is no limit. If the server group does not |
|                             | reside in the <i>shared memory</i> , the limitation works per each worker process.  |

## 1 Note

If idle keepalive connections, multiple workers, and the  $shared\ memory$  are enabled, the total number of active and idle connections to the proxied server may exceed the  $max\_conns$  value.



 $max_fails=number$  — sets the number of unsuccessful attempts to communicate with the server that should happen in the duration set by fail timeout to consider the server unavailable; it is then retried after the same duration.

What is considered an unsuccessful attempt is defined by the proxy next upstream,  $fastcgi\_next\_upstream$ ,  $uwsgi\_next\_upstream$ ,  $scgi\_next\_upstream$ , memcached next upstream, and grpc next upstream directives.

When max\_fails is reached, the peer is also considered unhealthy by the upstream probe (PRO) probes; it won't receive client requests until the probes consider it healthy again.

## 1 Note

If a server in an upstream resolves into multiple peers, its max\_fails setting applies to each peer individually.

If an upstream contains only one peer after all its server directives are resolved, the max\_fails setting has no effect and will be ignored.

| max_fails=1 | the default number of unsuccessful attempts |
|-------------|---|
| max_fails=0 | disables the accounting of attempts         |

fail\_timeout=time — sets the period of time during which a number of unsuccessful attempts to communicate with the server (max\_fails) should happen to consider the server unavailable. The server then becomes unavailable for the same amount of time before it is retried.

By default, this is set to 10 seconds.

## 1 Note

If a server in an upstream resolves into multiple peers, its fail\_timeout setting applies to each peer individually.

If an upstream contains only one peer after all its server directives are resolved, the fail\_timeout setting has no effect and will be ignored.

| backup | marks the server as a backup server. It will be passed requests when the primary servers are unavailable.  |
|--------|--|
|        | If backup_switch (PRO) is configured, its active backup logic is also applied.   |
| down   | marks the server as permanently unavailable.   |
| drain  | sets the server to draining; this means it receives only requests from the sessions that were bound earlier with <i>sticky</i> . Otherwise it behaves similarly to down. |

## Caution

The parameter backup cannot be used along with the hash, ip hash, and random load balancing methods.

The down and drain options are mutually exclusive.

Added in version 1.1.0.



| resolve      | enables monitoring changes to the list of IP addresses that corresponds to a domain name, updating it without a configuration reload. The group should be stored in a <i>shared memory zone</i> ; also, you need to define a <i>resolver</i> . |
|--------------|--|
| service=name | enables resolving DNS SRV records and sets the service name. For this parameter to work, specify the <i>resolve</i> server parameter, providing a hostname without a port number.  |
|              | If there are no dots in the service name, the name is formed according to the  |
|              | RFC standard: the service name is prefixed with _, then _tcp is added after a  |
|              | dot. Thus, the service name http will result in _httptcp.  |
|              | Angie resolves the SRV records by combining the normalized service name and  |
|              | the hostname and obtaining the list of servers for the combination via DNS, along  |
|              | with their priorities and weights.   |
|              | • Top-priority SRV records (ones that share the minimum priority value)  |
|              | resolve into primary servers, and other records become backup servers. If  |
|              | backup is set with server, top-priority SRV records resolve into backup  |
|              | servers, and other records are ignored.  |
|              | • Weight influences the selection of servers by the assigned capacity: higher  |
|              | weights receive more requests. If set by both the server directive and the   |
|              | SRV record, the weight set by server is used.  |

This example will look up the \_http.\_tcp.backend.example.com record:

server backend.example.com service=http resolve;

Added in version 1.2.0: Angie

Added in version 1.1.0-P1: Angie PRO

| $\mathtt{sid} \texttt{=} id$ | sets the server ID within the group. If the parameter is omitted, the ID is set to |
|------------------------------|--|
|                              | the hexadecimal MD5 hash value of either the IP address and port or the UNIX       |
|                              | domain socket path.  |

Added in version 1.4.0.

| slow_start=time | sets the time to recover the weight for a server that goes back online, if load        |
|-----------------|--|
|                 | balancing uses the round-robin or least_conn method.                                   |
|                 | If the value is set and the server is again considered available and healthy as        |
|                 | defined by $max\_fails$ and $upstream\_probe~(PRO)$ , the server will steadily recover |
|                 | its designated weight within the allocated timeframe.                                  |
|                 | If the value isn't set, the server in a similar situation will recover its designated  |
|                 | weight immediately.  |



If there's only one server in an upstream, slow\_start has no effect and will be ignored.

## state (PRO)

Added in version 1.2.0: PRO



| Syntax  | state $\mathit{file}$ ; |
|---------|-------------------------|
| Default | _                       |
| Context | upstream                |

Specifies the *file* where the upstream's server list is persisted. When installing from our packages, a designated /var/lib/angie/state/ (/var/db/angie/state/ on FreeBSD) directory with appropriate permissions is created to store these files, so you will only need to add the file's basename in the configuration:

```
upstream backend {
   zone backend 1m;
   state /var/lib/angie/state/<FILE NAME>;
}
```

The format of this server list is similar to **server**. The contents of the file change whenever there is any modification to servers in the /config/http/upstreams/ section via the configuration API. The file is read at Angie start or configuration reload.

## **Caution**

For the state directive to be used in an upstream block, the block should have no server directives; instead, it must have a shared memory zone (zone).

## sticky

Added in version 1.2.0: Angie

Added in version 1.1.0-P1: Angie PRO

| Syntax  | sticky cookie name [attr=value]; sticky route \$variable; sticky learn zone=zone create=\$create_var1 lookup=\$lookup_var1 [header] [timeout=time]; sticky learn zone=zone lookup=\$lookup_var1 remote_action=uri remote_result=\$remote var [timeout=time]; |
|---------|--|
|         | Tomodo_results wremove_var [ormeodo varies],   |
| Default | _  |
| Context | upstream   |

Configures the binding of client sessions to proxied servers in the mode specified by the first parameter; to drain requests from servers that have sticky defined, use the drain option in the server block.

## **A** Attention

The sticky directive must be used after all directives that set the load balancing method; otherwise, it won't work. If bind\_conn (PRO) is also used, bind\_conn should appear after sticky.

#### cookie mode

This mode uses cookies to maintain session persistence. It is more suitable for situations where cookies are already used for session management.

Here, a client's request, not yet bound to any server, is sent to a server chosen according to the configured balancing method. Also, Angie sets a cookie with a unique value identifying the server.



The cookie's name (name) is set by the sticky directive, and the value (value) corresponds to the *sid* parameter of the *server* directive. Note that the parameter is additionally hashed if the *sticky\_secret* directive is set.

Subsequent client requests that contain this cookie are forwarded to the server identified by the cookie's value, which is the server with the specified *sid*. If selecting a server fails or the chosen server can't handle the request, another server is selected according to the configured balancing method.

The directive allows assigning attributes to the cookie; the only attribute set by default is path=/. Attribute values are specified as strings with variables. To remove an attribute, set an empty value for it: attr=. Thus, sticky cookie path= creates a cookie without path.

Here, Angie creates a cookie named srv\_id with a one-hour lifespan and a variable-specified domain:

```
upstream backend {
    server backend1.example.com:8080;
    server backend2.example.com:8080;

    sticky cookie srv_id domain=$my_domain max-age=3600;
}
```

#### route mode

This mode uses predefined route identifiers that can be embedded in URLs, cookies, or other request properties. It is less flexible because it relies on predefined values but can suit better if such identifiers are already in place.

Here, when a proxied server receives a request, it can assign a route to the client and return its identifier in a way that both the client and the server are aware of. The value of the *sid* parameter of the *server* directive must be used as the route identifier. Note that the parameter is additionally hashed if the *sticky\_secret* directive is set.

Subsequent requests from clients that wish to use this route must contain the identifier issued by the server in a way that ensures it ends up in Angie variables, for example, in *cookie* or *request arguments*.

The directive lists the specific variables used for routing. To select the server to which the incoming request is forwarded, the first non-empty variable is used; it is then compared with the *sid* parameter of the *server* directive. If selecting a server fails or the chosen server can't handle the request, another server is selected according to the configured balancing method.

Here, Angie looks for the route identifier in the route cookie, and then in the route request argument:

```
upstream backend {
   server backend1.example.com:8080 "sid=server 1";
   server backend2.example.com:8080 "sid=server 2";
   sticky route $cookie_route $arg_route;
}
```

learn mode (PRO 1.4.0+)

This mode uses a dynamically generated key to associate a client with a particular proxied server; it's more flexible because it assigns servers on the go, stores sessions in a shared memory zone, and supports different ways of passing session identifiers.

Here, a session is created based on the response from the proxied server. The create and lookup parameters list variables indicating how new sessions are created and existing sessions are looked up. Both parameters can occur multiple times.

The session identifier is the value of the first non-empty variable specified with create; for example, this could be a *cookie from the proxied server*.

Sessions are stored in a shared memory zone; its name and size are set by the zone parameter. If a session has been inactive for the time set by timeout, it is deleted. The default is 10 minutes.



Subsequent requests from clients that wish to use the session must contain its identifier, ensuring that it ends up in a non-empty variable specified with lookup; its value will then be matched against sessions in shared memory. If selecting a server fails or the chosen server can't handle the request, another server is selected according to the configured balancing method.

The header parameter allows creating a session immediately after receiving headers from the proxied server. Without it, a session is created only after processing the request.

In the example, Angie creates a session, setting a cookie named examplecookie in the response:

```
upstream backend {
    server backend1.example.com:8080;
    server backend2.example.com:8080;

    sticky learn
        lookup=$cookie_examplecookie
        zone=client_sessions:1m;
}
```

learn mode with remote\_action (PRO 1.8.0+)

The remote\_action and remote\_result parameters enable dynamically assigning and managing session IDs via remote session storage. Here, the shared memory acts as a local cache, while the remote storage is the authoritative source. Thus, the create parameter is incompatible with remote\_action because session IDs need to be created remotely. If a session has been inactive for the time set by timeout, it is deleted. The remote\_action setting doesn't affect the timeout. The default is 10 minutes.

The initial session ID always comes from lookup; if it can be found in the local shared memory, Angie proceeds to select the appropriate peer.

If this session ID isn't found locally, Angie sends a synchronous subrequest to remote storage. The remote\_action parameter sets the URI of the remote storage, which should handle session lookup and creation as follows:

- It accepts the session ID from lookup and the locally suggested server ID to be associated with this session via custom header fields or in some other way. On Angie's side, two special variables are provided for this purpose: \$sticky\_sessid and \$sticky\_sid\$, respectively. The sticky\_sid value comes from the sid= setting in the upstream block of the server directive, if it's set; otherwise, it is an MD5 hash of the server's name.
- A 200 response from the remote storage indicates it has accepted the session and saved it with the suggested values for later retrieval.
- A 409 response from the remote storage indicates that this session ID is already populated. In this case, the response should suggest an alternative session ID in the X-Sticky-Sid header field. Angie saves this ID in the variable set by the remote\_result parameter.

In this example, Angie creates a session, uses the \$cookie\_bar variable for the initial session ID, and stores alternative session IDs reported by the remote storage in \$upstream\_http\_x\_sticky\_sid:

```
http {
    upstream u1 {
        server srv1;
        server srv2;

    sticky learn zone=sz:1m
        lookup=$cookie_bar
        remote_action=/remote_session
        remote_result=$upstream_http_x_sticky_sid;

    zone z 1m;
```



```
server {

listen localhost;

location / {

    proxy_pass http://u1/;
}

location /remote_session {

    internal;
    proxy_set_header X-Sticky-Sessid $sticky_sessid;
    proxy_set_header X-Sticky-Sid $sticky_sid;
    proxy_set_header X-Sticky-Last $msec;
    proxy_pass http://remote;
}
}
```

## sticky\_secret

Added in version 1.2.0: Angie

Added in version 1.1.0-P1: Angie PRO

```
Syntax sticky_secret string;
Default —
Context upstream
```

Adds the *string* as the salt value to the MD5 hashing function for the *sticky* directive in cookie and route modes. The *string* may contain variables, for example, *\$remote\_addr*:

```
upstream backend {
    server backend1.example.com:8080;
    server backend2.example.com:8080;

    sticky cookie cookie_name;
    sticky_secret my_secret.$remote_addr;
}
```

Salt is appended to the value being hashed; to verify the hashing mechanism independently:

```
$ echo -n "<VALUE><SALT>" | md5sum
```

## sticky strict

Added in version 1.2.0: Angie

Added in version 1.1.0-P1: Angie PRO

```
Syntax sticky_strict on | off;
Default sticky_strict off;
Context upstream
```



When enabled, causes Angie to return an HTTP 502 error to the client if the desired server is unavailable, rather than using any other available server as it would when no servers in the upstream are available.

#### upstream

| Syntax  | $\mathtt{upstream}\ name\ \{\\ \}$ |
|---------|------------------------------------|
| Default | _                                  |
| Context | http                               |

Defines a group of servers. Servers can listen on different ports. In addition, servers listening on TCP and UNIX domain sockets can be mixed.

#### Example:

By default, requests are distributed between the servers using a weighted round-robin balancing method. In the above example, each 7 requests will be distributed as follows: 5 requests go to backend1.example.com and one request to each of the second and third servers.

If an error occurs during communication with a server, the request will be passed to the next server, and so on until all of the functioning servers will be tried. If a successful response could not be obtained from any of the servers, the client will receive the result of the communication with the last server.

#### zone

| Syntax  | zone name [size]; |
|---------|-------------------|
| Default | _                 |
| Context | upstream          |

Defines the name and size of the shared memory zone that keeps the group's configuration and run-time state that are shared between worker processes. Several groups may share the same zone. In this case, it is enough to specify the size only once.

## **Built-in Variables**

The http\_upstream module supports the following built-in variables:

#### \$sticky\_sessid

Used with remote\_action in *sticky*; stores the initial session ID taken from lookup.

## \$sticky\_sid

Used with remote\_action in *sticky*; stores the server ID tentatively associated with the session.



#### \$upstream\_addr

stores the IP address and port, or the path to the UNIX domain socket of the upstream server. If several servers were contacted during request processing, their addresses are separated by commas, e.g.:

192.168.1.1:80, 192.168.1.2:80, unix:/tmp/sock

If an internal redirect from one server group to another happens, initiated by "X-Accel-Redirect" or *error page*, then the server addresses from different groups are separated by colons, e.g.:

192.168.1.1:80, 192.168.1.2:80, unix:/tmp/sock: 192.168.10.1:80, 192.168.10.2:80

If a server cannot be selected, the variable keeps the name of the server group.

#### \$upstream\_bytes\_received

number of bytes received from an upstream server. Values from several connections are separated by commas and colons like addresses in the *\$upstream* addr variable.

#### \$upstream\_bytes\_sent

number of bytes sent to an upstream server. Values from several connections are separated by commas and colons like addresses in the  $\$upstream\_addr$  variable.

#### \$upstream\_cache\_status

keeps the status of accessing a response cache. The status can be either MISS, BYPASS, EXPIRED, STALE, UPDATING, REVALIDATED or HIT:

- MISS: The response isn't found in the cache, and the request is forwarded to the upstream server.
- BYPASS: The cache is bypassed, and the request is directly forwarded to the upstream server.
- EXPIRED: The cached response is stale, and a new request for the updated content is sent to the upstream server.
- STALE: The cached response is stale, but will be served to the clients until an update has been eventually fetched from the upstream server.
- UPDATING: The cached response is stale, but will be served to the clients until the currently ongoing update from the upstream server has been finished.
- REVALIDATED: The cached response is stale, but is successfully revalidated and doesn't need an update from the upstream server.
- HIT: The response was served from the cache.

If the cache was bypassed entirely without accessing it, the variable isn't set.

#### \$upstream\_connect\_time

keeps time spent on establishing a connection with the upstream server; the time is kept in seconds with millisecond resolution. In case of SSL, includes time spent on handshake. Times of several connections are separated by commas and colons like addresses in the  $\$upstream\_addr$  variable.

#### \$upstream\_cookie\_<name>

cookie with the specified *name* sent by the upstream server in the "Set-Cookie" response header field. Only the cookies from the response of the last server are saved.



#### \$upstream\_header\_time

stores time spent on receiving the response header from the upstream server; the time is kept in seconds with millisecond resolution. Times of several responses are separated by commas and colons like addresses in the  $\$upstream \ addr$  variable.

#### \$upstream\_http\_<name>

stores server response header fields. For example, the "Server" response header field is available through the *\$upstream\_http\_server* variable. The rules of converting header field names to variable names are the same as for the variables that start with the "\$http\_" prefix. Only the header fields from the response of the last server are saved.

#### \$upstream\_queue\_time

stores time the request spent in the queue before a server was selected; the time is kept in seconds with millisecond resolution. Times of several selection attempts are separated by commas and colons, like addresses in the \$upstream addr variable.

#### \$upstream\_response\_length

keeps the length of the response obtained from the upstream server; the length is kept in bytes. Lengths of several responses are separated by commas and colons like addresses in the *\$upstream addr* variable.

## \$upstream\_response\_time

keeps time spent on receiving the response from the upstream server; the time is kept in seconds with millisecond resolution. Times of several responses are separated by commas and colons like addresses in the \$upstream addr variable.

#### \$upstream\_status

keeps status code of the response obtained from the upstream server. Status codes of several responses are separated by commas and colons like addresses in the *\$upstream\_addr* variable. If a server cannot be selected, the variable keeps the 502 (Bad Gateway) status code.

## \$upstream\_sticky\_status

Status of sticky requests.

| 11 11 | Request sent to upstream without sticky enabled.  |
|-------|---|
| NEW   | Request without sticky information.   |
| HIT   | Request with sticky information routed to the desired backend.                                  |
| MISS  | Request with sticky information routed to the backend selected by the load balancing algorithm. |

Values from multiple connections are separated by commas and colons, similar to addresses in the \$upstream addr variable.

#### \$upstream\_trailer\_<name>

stores fields from the end of the response obtained from the upstream server.



#### **Upstream Probe**

The module implements active health probes for http upstream.

## **Configuration Example**

```
server {
    listen ...;
    location /backend {
        proxy_pass http://backend;
        upstream_probe backend_probe
            uri=/probe
            port=10004
            interval=5s
            test=$good
            essential
            fails=3
            passes=3
            max_body=10m
            mode=idle;
    }
}
```

#### **Directives**

#### upstream probe (PRO)

Added in version 1.2.0: PRO

| Syntax  | . , .    | name [uri=address] test=condition] [essent ax_body=number] [mode=a | t <b>=</b> 11 | [fails=number] |
|---------|----------|--|---------------|----------------|
| Default | _        |  |               |                |
| Context | location |  |               |                |

Defines an active health probe for peers within the *upstream* groups that are specified for *proxy\_pass*, *uwsgi\_pass*, and so on in the same location context with the upstream\_probe directive. Subsequently, Angie regularly probes each peer of the upstream group according to the parameters configured here.

A peer's probe is passed if the request to the peer succeeds, considering all parameter settings of the upstream\_probe directive and the settings that control how upstreams are used by the directive's location context. This includes the proxy\_next\_upstream and uwsgi\_next\_upstream directives, etc.; also, proxy\_set\_header and so on.

To make use of the probes, the upstream must have a shared memory zone (zone). One upstream may be configured with several probes.

The following parameters are accepted:



| name       | Mandatory name of the probe.   |  |  |
|------------|--|--|--|
| uri        | Request URI to be added to the argument for $proxy\_pass$ , $uwsgi\_pass$ , etc. By default — $/$ .  |  |  |
| port       | Alternative port number for the probe request.   |  |  |
| interval   | Interval between probes. By default $-5s$ .  |  |  |
| method     | HTTP method of the probe. By default $-$ <b>GET</b> .  |  |  |
| test       | The condition for the probe, defined as a string with variables. If variable substitution yields "" or "0", the probe is not passed.   |  |  |
| essential  | If set, the initial state of the peer is being checked, so the peer doesn't receive client requests until the probe is passed.   |  |  |
| persistent | Setting this parameter requires enabling essential first; persistent peers that were deemed healthy prior to a configuration reload start receiving requests without being required to pass this probe first.  |  |  |
| fails      | Number of subsequent failed probes that renders the peer unhealthy. By default $-$ 1.  |  |  |
| passes     | Number of subsequent passed probes that renders the peer healthy. By default $-1$ .  |  |  |
| max_body   | Maximum amount of memory for the response body. By default $-$ 256k.   |  |  |
| mode       | <ul> <li>Probe mode, depending on the peers' health:</li> <li>always — peers are probed regardless of their state;</li> <li>idle — probes affect unhealthy peers and peers where interval has elapsed since the last client request.</li> <li>onfail — only unhealthy peers are probed.</li> <li>By default — always.</li> </ul> |  |  |

## Example:

```
upstream backend {
   zone backend 1m;
   server backend1.example.com;
    server backend2.example.com;
}
map $upstream_status $good {
          "1";
    200
server {
   listen ...;
    location /backend {
       proxy_pass http://backend;
        upstream_probe backend_probe
           uri=/probe
            port=10004
            interval=5s
            test=$good
            essential
            persistent
            fails=3
            passes=3
           max_body=10m
            mode=idle;
```



}

Details of probe operation:

- Initially, the peer won't receive client requests until it passes all essential probes configured for it, skipping persistent ones if the configuration was reloaded and the peer was deemed healthy prior to that. If there are no such probes, the peer is considered healthy.
- The peer is considered unhealthy and won't receive client requests, if any of the probes configured for it hits fails or the peer reaches max fails.
- For an unhealthy peer to be considered healthy again, all probes configured for it must reach their respective passes; after that, max fails is also considered.

#### **Built-in Variables**

The http\_upstream module supports the following built-in variables:

```
$upstream_probe (PRO)
```

Name of the currently active upstream probe.

```
$upstream_probe_body (PRO)
```

Peer response body, received during an upstream probe; its size is limited by max\_body.

#### UserID

The module sets cookies suitable for client identification. Received and set cookies can be logged using the built-in variables  $\$uid\_got$  and  $\$uid\_set$ . This module is compatible with the mod\_uid module for Apache.

#### **Configuration Example**

```
userid on;
userid_name uid;
userid_domain example.com;
userid_path /;
userid_expires 365d;
userid_p3p 'policyref="/w3c/p3p.xml", CP="CUR ADM OUR NOR STA NID"';
```

### Directives

# userid

| Syntax  | userid on   v1   log   off; |
|---------|-----------------------------|
| Default | userid off;                 |
| Context | http, server, location      |

Enables or disables setting cookies and logging the received cookies:

| on  | enables the setting of version 2 cookies and logging of the received cookies; |
|-----|---|
| v1  | enables the setting of version 1 cookies and logging of the received cookies; |
| log | disables the setting of cookies, but enables logging of the received cookies; |
| off | disables the setting of cookies and logging of the received cookies.          |



# userid\_domain

| Syntax  | $\verb"userid_domain" name \mid \verb"none";$ |
|---------|---|
| Default | userid_domain none;                           |
| Context | http, server, location                        |

Defines a domain for which the cookie is set. The none parameter disables setting of a domain for the cookie.

#### userid expires

| Syntax  | $\verb"userid_expires" time \mid \verb"max" \mid \verb"off";$ |
|---------|---|
| Default | userid_expires off;   |
| Context | http, server, location  |

Sets a time during which a browser should keep the cookie. The parameter max will cause the cookie to expire on "31 Dec 2037 23:55:55 GMT". The parameter off will cause the cookie to expire at the end of a browser session.

#### userid flags

| Syntax  | userid_flags off $\mid$ $flag$ ; |
|---------|----------------------------------|
| Default | userid_flags off;                |
| Context | http, server, location           |

If the parameter is not off, defines one or more additional flags for the cookie: secure, httponly, samesite=strict, samesite=lax, samesite=none.

#### userid mark

| Syntax  | $	exttt{userid_mark } letter \mid digit \mid = \mid 	exttt{off};$ |
|---------|---|
| Default | <pre>userid_mark off;</pre>                                       |
| Context | http, server, location  |

If the parameter is not off, enables the cookie marking mechanism and sets the character used as a mark. This mechanism is used to add or change  $userid\_p3p$  and/or a cookie expiration time while preserving the client identifier. A mark can be any letter of the English alphabet (case-sensitive), digit, or the "=" character.

If the mark is set, it is compared with the first padding symbol in the base64 representation of the client identifier passed in a cookie. If they do not match, the cookie is resent with the specified mark, expiration time, and "P3P" header.

## userid\_name

| Syntax  | userid_name name;      |
|---------|------------------------|
| Default | userid_name uid;       |
| Context | http, server, location |

Sets the cookie name.



# userid\_p3p

| Syntax  | userid_p3p $string \mid none;$ |
|---------|--------------------------------|
| Default | userid_p3p none;               |
| Context | http, server, location         |

Sets a value for the "P3P" header field that will be sent along with the cookie. If the directive is set to the special value none, the "P3P" header will not be sent in a response.

### userid path

| Syntax  | ${\tt userid\_path}\;path;$ |
|---------|-----------------------------|
| Default | userid_path /;              |
| Context | http, server, location      |

Defines a path for which the cookie is set.

# userid\_service

| Syntax  | userid_service number;                   |
|---------|--|
| Default | userid_service IP address of the server; |
| Context | http, server, location                   |

If identifiers are issued by multiple servers (services), each service should be assigned its own number to ensure that client identifiers are unique. For version 1 cookies, the default value is zero. For version 2 cookies, the default value is the number composed from the last four octets of the server's IP address.

#### **Built-in Variables**

### \$uid\_got

The cookie name and received client identifier.

### \$uid\_reset

If the variable is set to a non-empty string that is not 0, the client identifiers are reset. The special value log additionally leads to the output of messages about the reset identifiers to the *error* log.

#### \$uid\_set

The cookie name and sent client identifier.

#### uWSGI

Allows passing requests to a uwsgi server.

### **Configuration Example**

```
location / {
   include   uwsgi_params;
   uwsgi_pass localhost:9000;
}
```



#### **Directives**

### uwsgi\_bind

| Syntax  | ${\tt uwsgi\_bind} \ address \ [{\tt transparent}] \   \ {\tt off};$ |
|---------|--|
| Default | _  |
| Context | http, server, location   |

Makes outgoing connections to a uwsgi server originate from the specified local IP address with an optional port. Parameter value can contain variables. The special value off cancels the effect of the  $uwsgi\_bind$  directive inherited from the previous configuration level, which allows the system to autoassign the local IP address and port.

The transparent parameter allows outgoing connections to a uwsgi server originate from a non-local IP address, for example, from a real IP address of a client:

```
uwsgi_bind $remote_addr transparent;
```

In order for this parameter to work, it is usually necessary to run Angie worker processes with the *superuser* privileges. On Linux it is not required as if the **transparent** parameter is specified, worker processes inherit the  $CAP\_NET\_RAW$  capability from the master process.

# Important

It is necessary to configure kernel routing table to intercept network traffic from the uwsgi server.

# uwsgi\_buffer\_size

| Syntax  | uwsgi_buffer_size $size$ ; |
|---------|----------------------------|
| Default | uwsgi_buffer_size 4k 8k;   |
| Context | http, server, location     |

Sets the size of the buffer used for reading the first part of the response received from the uwsgi server. This part usually contains a small response header. By default, the buffer size is equal to one memory page. This is either 4K or 8K, depending on a platform. It can be made smaller, however.

#### uwsgi buffering

| Syntax  | uwsgi_buffering $size;$ |
|---------|-------------------------|
| Default | uwsgi_buffering on;     |
| Context | http, server, location  |

Enables or disables buffering of responses from the uwsgi server.

| on  | Angie receives a response from the uwsgi server as soon as possible, saving it into the buffers set by the uwsgi_buffer_size and uwsgi_buffers directives. If the whole response does not fit into memory, a part of it can be saved to a temporary file on the disk. Writing to temporary files is controlled by the uwsgi max temp file size and uwsgi temp file write size directives. |
|-----|---|
| off | The response is passed to a client synchronously, immediately as it is received. Angie will not try to read the whole response from the uwsgi server. The maximum size of the data that Angie can receive from the server at a time is set by the <code>uwsgi_buffer_size</code> directive.   |



Buffering can also be enabled or disabled by passing "yes" or "no" in the "X-Accel-Buffering" response header field. This capability can be disabled using the *uwsgi ignore headers* directive.

## uwsgi buffers

| Syntax  | uwsgi_buffers number size; |
|---------|----------------------------|
| Default | uwsgi_buffers 8 4k   8k;   |
| Context | http, server, location     |

Sets the number and size of the buffers used for reading a response from the uwsgi server, for a single connection.

By default, the buffer size is equal to one memory page. This is either 4K or 8K, depending on a platform.

# uwsgi\_busy\_buffers\_size

| Syntax  | uwsgi_busy_buffers_size $size;$   |
|---------|-----------------------------------|
| Default | uwsgi_busy_buffers_size 8k   16k; |
| Context | http, server, location            |

When *buffering* of responses from the uwsgi server is enabled, limits the total size of buffers that can be busy sending a response to the client while the response is not yet fully read. In the meantime, the rest of the buffers can be used for reading the response and, if needed, buffering part of the response to a temporary file.

By default, size is limited by the size of two buffers set by the  $uwsgi\_buffer\_size$  and  $uwsgi\_buffers$  directives.

## uwsgi cache

| Syntax  | uwsgi_cache zone   off; |
|---------|-------------------------|
| Default | uwsgi_cache off;        |
| Context | http, server, location  |

Defines a shared memory zone used for caching. The same zone can be used in several places. Parameter value can contain variables.

| off | disables caching inherited from the previous configuration level. |  |
|-----|---|--|
|-----|---|--|

## uwsgi cache background update

| Syntax  | uwsgi_cache_background_update on   off; |
|---------|---|
| Default | uwsgi_cache_background_update off;      |
| Context | http, server, location                  |

Allows starting a background subrequest to update an expired cache item, while a stale cached response is returned to the client.

### **A** Attention

Note that it is necessary to allow the usage of a stale cached response when it is being updated.



## uwsgi\_cache\_bypass

| Syntax  | uwsgi_cache_bypass;    |
|---------|------------------------|
| Default | _                      |
| Context | http, server, location |

Defines conditions under which the response will not be taken from a cache. If at least one value of the string parameters is not empty and is not equal to "0" then the response will not be taken from the cache:

Can be used along with the uwsgi no cache directive.

## uwsgi cache key

| Syntax  | uwsgi_cache_key $string$ ; |
|---------|----------------------------|
| Default | _                          |
| Context | http, server, location     |

Defines a key for caching, for example

```
uwsgi_cache_key localhost:9000$request_uri;
```

### uwsgi\_cache\_lock

| Syntax  | uwsgi_cache_lock on   off; |
|---------|----------------------------|
| Default | uwsgi_cache_lock off;      |
| Context | http, server, location     |

When enabled, only one request at a time will be allowed to populate a new cache element identified according to the  $uwsgi\_cache\_key$  directive by passing a request to a uwsgi server. Other requests of the same cache element will either wait for a response to appear in the cache or the cache lock for this element to be released, up to the time set by the  $uwsgi\_cache\_lock\_timeout$  directive.

## uwsgi cache lock age

| Syntax  | <pre>uwsgi_cache_lock_age time;</pre> |
|---------|---------------------------------------|
| Default | uwsgi_cache_lock_age 5s;              |
| Context | http, server, location                |

If the last request passed to the uwsgi server for populating a new cache element has not completed for the specified time, one more request may be passed to the uwsgi server.

## uwsgi\_cache\_lock\_timeout

| Syntax  | $\verb"uwsgi_cache_lock_time" out $time$;$ |
|---------|--|
| Default | <pre>uwsgi_cache_lock_timeout 5s;</pre>    |
| Context | http, server, location                     |



Sets a timeout for *uwsgi\_cache\_lock*. When the time expires, the request will be passed to the uwsgi server, however, the response will not be cached.

## uwsgi cache max range offset

| Syntax  | ${\tt uwsgi\_cache\_max\_range\_offset} \ number;$ |
|---------|--|
| Default | _  |
| Context | http, server, location                             |

Sets an offset in bytes for byte-range requests. If the range is beyond the offset, the range request will be passed to the uwsgi server and the response will not be cached.

## uwsgi cache methods

| Syntax  | uwsgi_cache_methods GET   HEAD   POST; |
|---------|--|
| Default | uwsgi_cache_methods GET HEAD;          |
| Context | http, server, location                 |

If the client request method is listed in this directive then the response will be cached. "GET" and "HEAD" methods are always added to the list, though it is recommended to specify them explicitly. See also the  $uwsgi\_no\_cache$  directive.

### uwsgi cache min uses

| Syntax  | uwsgi_cache_min_uses $number;$     |
|---------|------------------------------------|
| Default | <pre>uwsgi_cache_min_uses 1;</pre> |
| Context | http, server, location             |

Sets the number of requests after which the response will be cached.

### uwsgi\_cache\_path

| Syntax  | $\verb  uwsgi_cache_path   path   [levels=levels]   [use_temp_path=on   off] $  |
|---------|---|
|         | $\texttt{keys\_zone} = name : size \qquad [\texttt{inactive} = time] \qquad [\texttt{max\_size} = size] \qquad [\texttt{min\_free} = size]$ |
|         | $[manager\_files=number]$ $[manager\_sleep=time]$ $[manager\_threshold=time]$   |
|         | $[\texttt{loader\_files} = number] \ [\texttt{loader\_sleep} = time] \ [\texttt{loader\_threshold} = time];$                                |
| Default | _   |
| Context | http  |

Sets the path and other parameters of a cache. Cache data are stored in files. The file name in a cache is a result of applying the MD5 function to the *cache key*.

The levels parameter defines hierarchy levels of a cache: from 1 to 3, each level accepts values 1 or 2. For example, in the following configuration:

```
uwsgi_cache_path /data/angie/cache levels=1:2 keys_zone=one:10m;
```

file names in a cache will look like this:

```
/data/angie/cache/c/29/b7f54b2df7773722d382f4809d65029c
```

A cached response is first written to a temporary file, and then the file is renamed. Temporary files and the cache can be put on different file systems. However, be aware that in this case a file is copied across



two file systems instead of the cheap renaming operation. It is thus recommended that for any given location both saved files and a directory holding temporary files are put on the same file system.

The directory for temporary files is set based on the use\_temp\_path parameter.

| on  | If this parameter is omitted or set to the value on, the directory set by the <i>uwsqi temp path</i> directive for the given location will be used. |
|-----|---|
| off | Temporary files will be put directly in the cache directory.  |

In addition, all active keys and information about data are stored in a shared memory zone, whose name and size are configured by the keys\_zone parameter. One megabyte zone can store about 8 thousand keys.

Cached data that are not accessed during the time specified by the inactive parameter get removed from the cache regardless of their freshness.

By default, inactive is set to 10 minutes.

A special **cache manager** process monitors the maximum cache size and the minimum amount of free space on the file system with cache and when the size is exceeded or there is not enough free space, it removes the least recently used data. The data is removed in iterations.

| max_size         | maximum uwsgi_cache size   |
|------------------|--|
| min_free         | minimum amount of free space on the file system with cache                       |
| manager_files    | limits the number of items to be deleted during one iteration<br>By default, 100 |
| manager_threshol | limits the duration of one iteration By default, 200 milliseconds                |
| manager_sleep    | configures a pause between interactions By default, 50 milliseconds              |

A minute after Angie starts, the special **cache loader** process is activated. It loads information about previously cached data stored on file system into a cache zone. The loading is also done in iterations.

| loader_files   | limits the number of items to load during one iteration<br>By default, 100 |
|----------------|--|
| loader_thresho | ld limits the duration of one iteration By default, 200 milliseconds       |
| loader_sleep   | configures a pause between interactions By default, 50 milliseconds        |

### uwsgi cache revalidate

| Syntax  | uwsgi_cache_revalidate on   off;       |
|---------|--|
| Default | <pre>uwsgi_cache_revalidate off;</pre> |
| Context | http, server, location                 |

Enables revalidation of expired cache items using conditional requests with the "If-Modified-Since" and "If-None-Match" header fields.



## uwsgi\_cache\_use\_stale

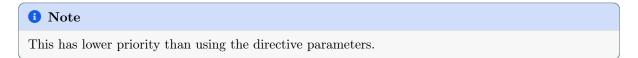
| Syntax  | uwsgi_cache_use_stale error   timeout   invalid_header   updating   http_500 |
|---------|--|
|         | http_503   http_403   http_404   http_429   off;                             |
| Default | <pre>uwsgi_cache_use_stale off;</pre>  |
| Context | http, server, location   |

Determines in which cases a stale cached response can be used during communication with the uwsgi server. The directive's parameters match the parameters of the *uwsgi next upstream* directive.

| error    | permits using a stale cached response if a uwsgi server to process a request cannot be selected.  |
|----------|---|
| updating | additional parameter, permits using a stale cached response if it is currently being updated. This allows minimizing the number of accesses to uwsgi servers when updating cached data. |

Using a stale cached response can also be enabled directly in the response header for a specified number of seconds after the response became stale:

- The stale-while-revalidate extension of the "Cache-Control" header field permits using a stale cached response if it is currently being updated.
- The stale-if-error extension of the "Cache-Control" header field permits using a stale cached response in case of an error.



To minimize the number of accesses to uwsgi servers when populating a new cache element, the  $uwsgi\_cache\_lock$  directive can be used.

# uwsgi\_cache\_valid

| Syntax  | uwsgi_cache_valid [code] time; |
|---------|--------------------------------|
| Default | _                              |
| Context | http, server, location         |

Sets caching time for different response codes. For example, the following directives

set 10 minutes of caching for responses with codes 200 and 302 and 1 minute for responses with code 404.

If only caching time is specified

```
uwsgi_cache_valid 5m;
```

then only 200, 301, and 302 responses are cached.

In addition, the any parameter can be specified to cache any responses:



```
uwsgi_cache_valid 200 302 10m;uwsgi_cache_valid 301 1h;uwsgi_cache_valid any 1m;
```

# 1 Note

Parameters of caching can also be set directly in the response header. This has higher priority than setting of caching time using the directive

- The "X-Accel-Expires" header field sets caching time of a response in seconds. The zero value disables caching for a response. If the value starts with the @ prefix, it sets an absolute time in seconds since Epoch, up to which the response may be cached.
- If the header does not include the "X-Accel-Expires" field, parameters of caching may be set in the header fields "Expires" or "Cache-Control".
- If the header includes the "Set-Cookie" field, such a response will not be cached.
- If the header includes the "Vary" field with the special value "\*", such a response will not be cached. If the header includes the "Vary" field with another value, such a response will be cached taking into account the corresponding request header fields.

Processing of one or more of these response header fields can be disabled using the  $uwsgi\_ignore\_headers$  directive.

# uwsgi connect timeout

| Syntax  | uwsgi_connect_timeout $time$ ;        |
|---------|---------------------------------------|
| Default | <pre>uwsgi_connect_timeout 60s;</pre> |
| Context | http, server, location                |

Defines a timeout for establishing a connection with a uwsgi server. It should be noted that this timeout cannot usually exceed 75 seconds.

## uwsgi connection drop

| Syntax  | uwsgi_connection_drop time   on   off; |
|---------|--|
| Default | uwsgi_connection_drop off;             |
| Context | http, server, location                 |

Enables termination of all connections to the proxied server after it has been removed from the group or marked as permanently unavailable by a *reresolve* process or the *API command* DELETE.

A connection is terminated when the next read or write event is processed for either the client or the proxied server.

Setting *time* enables a connection termination timeout; with on set, connections are dropped immediately.

# uwsgi\_force\_ranges

| Syntax  | uwsgi_force_ranges off;            |
|---------|------------------------------------|
| Default | <pre>uwsgi_force_ranges off;</pre> |
| Context | http, server, location             |

Enables byte-range support for both cached and uncached responses from the uwsgi server regardless of the "Accept-Ranges" field in these responses.



## uwsgi\_hide\_header

| Syntax  | ${\tt uwsgi\_hide\_header}\ field;$ |
|---------|-------------------------------------|
| Default | _                                   |
| Context | http, server, location              |

By default, Angie does not pass the header fields "Status" and "X-Accel-..." from the response of a uwsgi server to a client. The  $uwsgi\_hide\_header$  directive sets additional fields that will not be passed. If, on the contrary, the passing of fields needs to be permitted, the  $uwsgi\_pass\_header$  directive can be used.

# uwsgi ignore client abort

| Syntax  | <pre>uwsgi_ignore_client_abort on   off;</pre> |
|---------|--|
| Default | <pre>uwsgi_ignore_client_abort off;</pre>      |
| Context | http, server, location                         |

Determines whether the connection with a uwsgi server should be closed when a client closes the connection without waiting for a response.

## uwsgi\_ignore\_headers

| Syntax  | uwsgi_ignore_headers field; |
|---------|-----------------------------|
| Default | _                           |
| Context | http, server, location      |

Disables processing of certain response header fields from the uwsgi server. The following fields can be ignored: "X-Accel-Redirect", "X-Accel-Expires", "X-Accel-Limit-Rate", "X-Accel-Buffering", "X-Accel-Expires", "Expires", "Cache-Control", "Set-Cookie", and "Vary".

If not disabled, processing of these header fields has the following effect:

- "X-Accel-Expires", "Expires", "Cache-Control", "Set-Cookie" and "Vary" set the parameters of response *caching*;
- "X-Accel-Redirect" performs an *internal* redirect to the specified URI;
- "X-Accel-Limit-Rate" sets the rate limit for transmission of a response to a client;
- "X-Accel-Buffering" enables or disables buffering of a response;
- "X-Accel-Charset" sets the desired *charset* of a response.

#### uwsgi intercept errors

| Syntax  | <pre>uwsgi_intercept_errors on   off;</pre> |
|---------|---|
| Default | <pre>uwsgi_intercept_errors off;</pre>      |
| Context | http, server, location                      |

Determines whether uwsgi server responses with codes greater than or equal to 300 should be passed to a client or be intercepted and redirected to Angie for processing with the *error\_page* directive.



## uwsgi\_limit\_rate

| Syntax  | uwsgi_limit_rate rate;         |
|---------|--------------------------------|
| Default | <pre>uwsgi_limit_rate 0;</pre> |
| Context | http, server, location         |

Limits the speed of reading the response from the uwsgi server. The *rate* is specified in bytes per second and can contain variables.

0 disables rate limiting

# 1 Note

The limit is set per a request, and so if Angie simultaneously opens two connections to the uwsgi server, the overall rate will be twice as much as the specified limit. The limitation works only if buffering of responses from the uwsgi server is enabled.

# uwsgi\_max\_temp\_file\_size

| Syntax  | uwsgi_max_temp_file_size $size$ ; |
|---------|-----------------------------------|
| Default | uwsgi_max_temp_file_size 1024m;   |
| Context | http, server, location            |

When buffering of responses from the uwsgi server is enabled, and the whole response does not fit into the buffers set by the uwsgi\_buffer\_size and uwsgi\_buffers directives, a part of the response can be saved to a temporary file. This directive sets the maximum size of the temporary file. The size of data written to the temporary file at a time is set by the uwsgi temp file write size directive.

0 disables buffering of responses to temporary files

#### 1 Note

This restriction does not apply to responses that will be cached or *stored on disk*.

# uwsgi\_modifier1

| Syntax  | uwsgi_modifier1 $size$ ;      |
|---------|-------------------------------|
| Default | <pre>uwsgi_modifier1 0;</pre> |
| Context | http, server, location        |

Sets the value of the modifier field in the uwsgi packet header.

### uwsgi modifier2

| Syntax  | uwsgi_modifier2 size;         |
|---------|-------------------------------|
| Default | <pre>uwsgi_modifier2 0;</pre> |
| Context | http, server, location        |



Sets the value of the modifier 2 field in the uwsgi packet header.

# uwsgi next upstream

| Syntax  | uwsgi_next_upstream error   timeout   invalid_header   http_500   http_503   http_403   http_404   http_429   non_idempotent   off; |
|---------|---|
| Default | uwsgi_next_upstream error timeout;  |
| Context | http, server, location  |

Specifies in which cases a request should be passed to the next server in the upstream pool:

| error          | an error occurred while establishing a connection with the server, passing a request to it, or reading the response header;  |
|----------------|--|
| timeout        | a timeout has occurred while establishing a connection with the server, passing a request to it, or reading the response header;   |
| invalid_header | a server returned an empty or invalid response;  |
| http_500       | a server returned a response with the code 500;  |
| http_503       | a server returned a response with the code 503;  |
| http_403       | a server returned a response with the code 403;  |
| http_404       | a server returned a response with the code 404;  |
| http_429       | a server returned a response with the code 429;  |
| non_idempotent | normally, requests with a non-idempotent method (POST, LOCK, PATCH) are not passed to the next server if a request has been sent to an upstream server; enabling this option explicitly allows retrying such requests; |
| off            | disables passing a request to the next server.   |

# 1 Note

One should bear in mind that passing a request to the next server is only possible if nothing has been sent to a client yet. That is, if an error or timeout occurs in the middle of the transferring of a response, fixing this is impossible.

The directive also defines what is considered an unsuccessful attempt of communication with a server.

| error timeout invalid_header     | always considered unsuccessful attempts, even if they are not specified in the directive |
|----------------------------------|--|
| http_500<br>http_503<br>http_429 | considered unsuccessful attempts only if they are specified in the directive             |
| http_403<br>http_404             | never considered unsuccessful attempts   |

Passing a request to the next server can be limited by the *number of tries* and by time.

# $uwsgi\_next\_upstream\_timeout$

| Syntax  | uwsgi_next_upstream_timeout $time$ ;      |
|---------|---|
| Default | <pre>uwsgi_next_upstream_timeout 0;</pre> |
| Context | http, server, location                    |

Limits the time during which a request can be passed to the *next* server.



0 turns off this limitation

### uwsgi next upstream tries

| Syntax  | uwsgi_next_upstream_tries number;       |
|---------|---|
| Default | <pre>uwsgi_next_upstream_tries 0;</pre> |
| Context | http, server, location                  |

Limits the number of possible tries for passing a request to the *next* server.

| 0 turns off this limitation |  |
|-----------------------------|--|
|-----------------------------|--|

### uwsgi\_no\_cache

| Syntax  | uwsgi_no_cache string; |
|---------|------------------------|
| Default | _                      |
| Context | http, server, location |

Defines conditions under which the response will not be saved to a cache. If at least one value of the string parameters is not empty and is not equal to "0" then the response will not be saved:

Can be used along with the uwsgi cache bypass directive.

#### uwsgi param

| Syntax  | <pre>uwsgi_param parameter value [if_not_empty];</pre> |
|---------|--|
| Default | _  |
| Context | http, server, location                                 |

Sets a parameter that should be passed to the uwsgi server. The value can contain text, variables, and their combination. These directives are inherited from the previous configuration level if and only if there are no uwsgi\_param directives defined on the current level.

Standard CGI environment variables should be provided as uwsgi headers, see the *uwsgi\_params* file provided in the distribution:

```
location / {
   include uwsgi_params;
# ...
}
```

If the directive is specified with if\_not\_empty then such a parameter will be passed to the server only if its value is not empty:

```
uwsgi_param HTTPS $https if_not_empty;
```



### uwsgi\_pass

| Syntax  | ${\tt uwsgi\_pass} \ [protocol://] \ address;$ |
|---------|--|
| Default | _  |
| Context | location, if in location                       |

Sets the protocol and address of a uwsgi server and an optional URI to which a location should be mapped. As a protocol, uwsgi or suwsgi (secured uwsgi, uwsgi over SSL) can be specified. The address can be specified as a domain name or IP address, and a port:

```
uwsgi_pass localhost:9000;
uwsgi_pass uwsgi://localhost:9000;
uwsgi_pass suwsgi://[2001:db8::1]:9090;
```

or as a UNIX domain socket path:

```
uwsgi_pass unix:/tmp/uwsgi.socket;
```

If a domain name resolves to several addresses, all of them will be used in a round-robin fashion. In addition, an address can be specified as a *server group*. If a group is used, you cannot specify the port with it; instead, specify the port for each server within the group individually.

Parameter value can contain variables. In this case, if an address is specified as a domain name, the name is searched among the described server groups, and, if not found, is determined using a *resolver*.

## uwsgi pass header

| Syntax  | uwsgi_pass_header field; |
|---------|--------------------------|
| Default | _                        |
| Context | http, server, location   |

Permits passing *otherwise disabled* header fields from a uwsgi server to a client.

### uwsgi pass request body

| Syntax  | <pre>uwsgi_pass_request_body on   off;</pre> |
|---------|--|
| Default | <pre>uwsgi_pass_request_body on;</pre>       |
| Context | http, server, location                       |

Indicates whether the original request body is passed to the uwsgi server. See also the uwsgi pass request headers directive.

#### uwsgi pass request headers

| Syntax  | uwsgi_pass_request_headers on   off;      |
|---------|---|
| Default | <pre>uwsgi_pass_request_headers on;</pre> |
| Context | http, server, location                    |

Indicates whether the header fields of the original request are passed to the uwsgi server. See also the  $uwsgi\_pass\_request\_body$  directives.



## uwsgi\_read\_timeout

| Syntax  | $\verb"uwsgi_read_time" out $time";$ |
|---------|--------------------------------------|
| Default | <pre>uwsgi_read_timeout 60s;</pre>   |
| Context | http, server, location               |

Defines a timeout for reading a response from the uwsgi server. The timeout is set only between two successive read operations, not for the transmission of the whole response. If the uwsgi server does not transmit anything within this time, the connection is closed.

# uwsgi request buffering

| Syntax  | uwsgi_request_buffering on   off;      |
|---------|--|
| Default | <pre>uwsgi_request_buffering on;</pre> |
| Context | http, server, location                 |

Enables or disables buffering of a client request body.

| on  | the entire request body is <i>read</i> from the client before sending the request to a uwsgi server.  |
|-----|---|
| off | the request body is sent to the uwsgi server immediately as it is received. In this case, the request cannot be passed to the <i>next server</i> if Angie already started sending the request body. |

When HTTP/1.1 chunked transfer encoding is used to send the original request body, the request body will be buffered regardless of the directive value.

# uwsgi\_send\_timeout

| Syntax  | uwsgi_send_timeout $time$ ;        |
|---------|------------------------------------|
| Default | <pre>uwsgi_send_timeout 60s;</pre> |
| Context | http, server, location             |

Sets a timeout for transmitting a request to the uwsgi server. The timeout is set only between two successive write operations, not for the transmission of the whole request. If the uwsgi server does not receive anything within this time, the connection is closed.

# uwsgi socket keepalive

| Syntax  | uwsgi_socket_keepalive on   off;       |
|---------|--|
| Default | <pre>uwsgi_socket_keepalive off;</pre> |
| Context | http, server, location                 |

Configures the "TCP keepalive" behavior for outgoing connections to a uwsgi server.

| 11 11 | By default, the operating system's settings are in effect for the socket. |
|-------|---|
| on    | The SO_KEEPALIVE socket option is turned on for the socket.               |



# uwsgi\_ssl\_certificate

| Syntax  | ${\tt uwsgi\_ssl\_certificate}\ file;$ |
|---------|--|
| Default | _                                      |
| Context | http, server, location                 |

Specifies a file with the certificate in the PEM format used for authentication to a secured uwsgi server. Variables can be used in the file name.

## uwsgi ssl certificate cache

| Syntax  | uwsgi_ssl_certificate_cache off;<br>uwsgi_ssl_certificate_cache max= $N$ [inactive= $time$ ] [valid= $time$ ]; |
|---------|--|
| Default | uwsgi_ssl_certificate_cache off;   |
| Context | 9  |
| Context | http, server, location   |

Defines a cache that stores SSL certificates and secret keys specified using variables.

The directive supports the following parameters:

- max sets the maximum number of elements in the cache. When the cache overflows, the least recently used (LRU) elements are removed.
- inactive defines the time after which an element is removed if it has not been accessed. The default is 10 seconds.
- valid defines the time during which a cached element is considered valid and can be reused. The default is 60 seconds. After this period, certificates are reloaded or revalidated.
- off disables the cache.

#### Example:

# uwsgi\_ssl\_certificate\_key

| Syntax  | ${\tt uwsgi\_ssl\_certificate\_key} \ \mathit{file};$ |
|---------|---|
| Default | _   |
| Context | http, server, location                                |

Specifies a file with the secret key in the PEM format used for authentication to a secured uwsgi server.

The value "engine: name: id" can be specified instead of the file, which loads a secret key with a specified id from the OpenSSL engine name. Variables can be used in the file name.

### uwsgi\_ssl\_ciphers

| Syntax  | uwsgi_ssl_ciphers ciphers; |
|---------|----------------------------|
| Default | uwsgi_ssl_ciphers DEFAULT; |
| Context | http, server, location     |

Specifies the enabled ciphers for requests to a secured uwsgi server. The ciphers are specified in the format understood by the OpenSSL library.



The list of ciphers depends on the version of OpenSSL installed. The full list can be viewed using the openssl ciphers command.

# **A** Attention

The uwsgi\_ssl\_ciphers directive does not configure ciphers for TLS 1.3 when using OpenSSL. To tune TLS 1.3 ciphers with OpenSSL, use the uwsgi\_ssl\_conf\_command directive, which was added to support advanced SSL configuration.

- In LibreSSL, TLS 1.3 ciphers can be configured using uwsgi\_ssl\_ciphers.
- In BoringSSL, TLS 1.3 ciphers cannot be configured at all.

#### uwsgi ssl conf command

| Syntax  | uwsgi_ssl_conf_command name value; |
|---------|------------------------------------|
| Default | _                                  |
| Context | http, server, location             |

Sets arbitrary OpenSSL configuration commands when establishing a connection with the secured uwsgi server.

# Important

The directive is supported when using OpenSSL 1.0.2 or higher. To configure TLS 1.3 ciphers with OpenSSL, use the ciphersuites command.

Several  $uwsgi\_ssl\_conf\_command$  directives can be specified on the same level. These directives are inherited from the previous configuration level if and only if there are no  $uwsgi\_ssl\_conf\_command$  directives defined on the current level.

# \* Caution

Note that configuring OpenSSL directly might result in unexpected behavior.

### uwsgi\_ssl\_crl

| Syntax  | uwsgi_ssl_crl file;    |
|---------|------------------------|
| Default | _                      |
| Context | http, server, location |

Specifies a file with revoked certificates (CRL) in the PEM format used to *verify* the certificate of the secured uwsgi server.

#### uwsgi ssl name

| Syntax  | uwsgi_ssl_name name;                              |
|---------|---|
| Default | <pre>uwsgi_ssl_name `host from uwsgi_pass;`</pre> |
| Context | http, server, location                            |



Allows overriding the server name used to verify the certificate of the secured uwsgi server and to be passed through SNI when establishing a connection with the secured uwsgi server.

By default, the host part of the uwsgi pass URL is used.

#### uwsgi ssl password file

| Syntax  | ${\tt uwsgi\_ssl\_password\_file}$ $file;$ |
|---------|--|
| Default | _  |
| Context | http, server, location                     |

Specifies a file with passphrases for *secret keys* where each passphrase is specified on a separate line. Passphrases are tried in turn when loading the key.

### uwsgi ssl protocols

| Syntax  | uwsgi_ssl_protocols [SSLv2] [SSLv3] [TLSv1] [TLSv1.1] [TLSv1.2] [TLSv1.3]; |
|---------|--|
| Default | uwsgi_ssl_protocols TLSv1.2 TLSv1.3;                                       |
| Context | http, server, location   |

Changed in version 1.2.0: TLSv1.3 parameter added to default set.

Enables the specified protocols for requests to a secured uwsgi server.

# uwsgi ssl server name

| Syntax  | uwsgi_ssl_server_name on   off;       |
|---------|---------------------------------------|
| Default | <pre>uwsgi_ssl_server_name off;</pre> |
| Context | http, server, location                |

Enables or disables passing the server name set by the  $uwsgi\_ssl\_name$  directive via the Server Name Indication TLS extension (SNI, RFC 6066) while establishing a connection with the secured uwsgi server.

# uwsgi\_ssl\_session\_reuse

| Syntax  | uwsgi_ssl_session_reuse on   off; |
|---------|-----------------------------------|
| Default | uwsgi_ssl_session_reuse on;       |
| Context | http, server, location            |

Determines whether SSL sessions can be reused when working with the uwsgi server. If the errors "SSL3\_GET\_FINISHED:digest check failed" appear in the logs, try disabling session reuse.

# uwsgi ssl trusted certificate

| Syntax  | ${\tt uwsgi\_ssl\_trusted\_certificate} \ \mathit{file};$ |
|---------|---|
| Default | _   |
| Context | http, server, location                                    |

Specifies a file with trusted CA certificates in the PEM format used to *verify* the certificate of the secured uwsgi server.



## uwsgi\_ssl\_verify

| Syntax  | <pre>uwsgi_ssl_verify on   off;</pre> |
|---------|---------------------------------------|
| Default | <pre>uwsgi_ssl_verify off;</pre>      |
| Context | http, server, location                |

Enables or disables verification of the secured uwsgi server certificate.

## uwsgi ssl verify depth

| Syntax  | <pre>uwsgi_ssl_verify_depth number;</pre> |
|---------|---|
| Default | <pre>uwsgi_ssl_verify_depth 1;</pre>      |
| Context | http, server, location                    |

Sets the verification depth in the secured uwsgi server certificates chain.

#### uwsgi store

| Syntax  | uwsgi_store on   off   string; |
|---------|--------------------------------|
| Default | uwsgi_store off;               |
| Context | http, server, location         |

Enables saving of files to a disk.

| on  | saves files with paths corresponding to the directives alias or root |
|-----|--|
| off | disables saving of files   |

The file name can be set explicitly using the string with variables:

```
uwsgi_store /data/www$original_uri;
```

The modification time of files is set according to the received "Last-Modified" response header field. The response is first written to a temporary file, and then the file is renamed. Temporary files and the persistent store can be put on different file systems. However, be aware that in this case a file is copied across two file systems instead of the cheap renaming operation. It is thus recommended that for any given location both saved files and a directory holding temporary files, set by the <code>uwsgi\_temp\_path</code> directive, are put on the same file system.

This directive can be used to create local copies of static unchangeable files, e.g.:

```
location /images/ {
    root
                        /data/www;
                        404 = /fetch$uri;
    error_page
}
location /fetch/ {
    internal;
    uwsgi_pass
                        backend:9000;
    . . .
    uwsgi_store
                        on;
    uwsgi_store_access user:rw group:rw all:r;
    uwsgi_temp_path
                        /data/temp;
```



```
alias /data/www/;
}
```

### uwsgi store access

| Syntax  | uwsgi_store_access users:permissions;  |
|---------|--|
| Default | <pre>uwsgi_store_access user:rw;</pre> |
| Context | http, server, location                 |

Sets access permissions for newly created files and directories, e.g.:

```
uwsgi_store_access user:rw group:rw all:r;
```

If any group or all access permissions are specified then user permissions may be omitted:

```
uwsgi_store_access group:rw all:r;
```

## uwsgi temp file write size

| Syntax  | uwsgi_temp_file_write_size $size;$            |
|---------|---|
| Default | <pre>uwsgi_temp_file_write_size 8k 16k;</pre> |
| Context | http, server, location                        |

Limits the size of data written to a temporary file at a time, when buffering of responses from the uwsgi server to temporary files is enabled. By default, size is limited by two buffers set by the uwsgi\_buffer\_size and uwsgi\_buffers directives. The maximum size of a temporary file is set by the uwsgi\_max\_temp\_file\_size directive.

### uwsgi temp path

| Syntax  | uwsgi_temp_path path [level1 [level2 [level3]]]`;   |
|---------|---|
| Default | <pre>uwsgi_temp_path uwsgi_temp; (the path depends on thehttp-uwsgi-temp-path build option)</pre> |
| Context | http, server, location  |

Defines a directory for storing temporary files with data received from uwsgi servers. Up to three-level subdirectory hierarchy can be used underneath the specified directory. For example, in the following configuration

```
uwsgi_temp_path /spool/angie/uwsgi_temp 1 2;
```

a temporary file might look like this:

```
/spool/angie/uwsgi_temp/7/45/00000123457
```

See also the use temp path parameter of the uwsgi cache path directive.

#### HTTP/2

Provides support for HTTP/2.

When building from the source code, this module isn't built by default; it should be enabled with the --with-http\_v2\_module build option.



In packages and images from our repos, the module is included in the build.

### **Configuration Example**

```
server {
    listen 443 ssl;

    http2 on;

    ssl_certificate server.crt;
    ssl_certificate_key server.key;
}
```

## Important

Note that accepting  $\mathrm{HTTP}/2$  connections over TLS requires the "Application-Layer Protocol Negotiation" (ALPN) TLS extension support, which is available since OpenSSL version 1.0.2.

If the  $ssl\_prefer\_server\_ciphers$  directive is set to the value "on", the ciphers should be configured to comply with RFC 9113, Appendix A black list and supported by clients.

#### **Directives**

### http2

Added in version 1.2.0.

| Syntax  | http2 on   off; |
|---------|-----------------|
| Default | http2 off;      |
| Context | http, server    |

Enables the HTTP/2 protocol.

# http2 body preread size

| Syntax  | http2_body_preread_size $size;$ |
|---------|---------------------------------|
| Default | _                               |
| Context | http, server                    |

Sets the size of the buffer per each request in which the request body may be saved before it is started to be processed.

### http2\_chunk\_size

| Syntax  | http2_chunk_size size; |
|---------|------------------------|
| Default | http2_chunk_size 8k;   |
| Context | http, server, location |

Sets the maximum size of chunks into which the response body is sliced. A too low value results in higher overhead. A too high value impairs prioritization due to HOL blocking.



## http2\_max\_concurrent\_pushes

Deprecated since version 1.2.0.

| Syntax  | http2_max_concurrent_pushes number; |
|---------|-------------------------------------|
| Default | http2_max_concurrent_pushes 10;     |
| Context | http, server                        |

Limits the maximum number of concurrent *push* requests in a connection.

### http2 max concurrent streams

| Syntax  | http2_max_concurrent_streams number; |
|---------|--------------------------------------|
| Default | http2_max_concurrent_streams 128;    |
| Context | http, server                         |

Sets the maximum number of concurrent HTTP/2 streams in a connection.

# http2\_push

Deprecated since version 1.2.0.

| Syntax  | $\mathtt{http2\_push}\ uri\  \ \mathtt{off};$ |
|---------|---|
| Default | http2_push off;                               |
| Context | http, server, location                        |

Preemptively sends (pushes) a request to the specified uri along with the response to the original request. Only relative URIs with absolute path will be processed, for example:

```
http2_push /static/css/main.css;
```

The uri value can contain variables.

Several  $http2\_push$  directives can be specified on the same configuration level. The off parameter cancels the effect of the  $http2\_push$  directives inherited from the previous configuration level.

## http2 push preload

Deprecated since version 1.2.0.

| Syntax  | http2_push_preload on   off; |
|---------|------------------------------|
| Default | http2_push_preload off;      |
| Context | http, server, location       |

Enables automatic conversion of preload links specified in the "Link" response header fields into push requests.

# http2\_recv\_buffer\_size

| Syntax  | http2_recv_buffer_size $size;$ |
|---------|--------------------------------|
| Default | http2_recv_buffer_size 256k;   |
| Context | http                           |

Sets the size of the per worker input buffer.



#### **Built-in Variables**

The http v2 module supports the following built-in variables:

### \$http2

negotiated protocol identifier:

| h2    | for HTTP/2 over TLS               |
|-------|-----------------------------------|
| h2c   | for $HTTP/2$ over cleartext $TCP$ |
| 11 11 | an empty string otherwise         |

### HTTP/3

Enables HTTP/3 support for client connections, as well as for connections with proxied servers configured using the following  $http\_proxy$  directives:

```
proxy_http3_hq
proxy_http3_max_concurrent_streams
proxy_http3_max_table_capacity
proxy_http3_stream_buffer_size
proxy_http_version
proxy_pass
proxy_quic_active_connection_id_limit
proxy_quic_gso
proxy_quic_host_key
```

When building from the source code, this module isn't built by default; it should be enabled with the --with-http\_v3\_module build option.

In packages and images from our repos, the module is included in the build.

#### **Configuration Example**

```
http {
   log_format quic '$remote_addr - $remote_user [$time_local] '
                    '"$request" $status $body_bytes_sent '
                    '"$http_referer" "$http_user_agent" "$http3"';
   access_log logs/access.log quic;
    server {
        # for better compatibility it's recommended
        # to use the same port for http/3 and https
        listen 8443 quic reuseport;
        listen 8443 ssl;
        ssl_certificate
                          certs/example.com.crt;
        ssl_certificate_key certs/example.com.key;
        location / {
            # used to advertise the availability of HTTP/3
            add_header Alt-Svc 'h3=":8443"; ma=86400';
```



} }

# Important

Note that accepting  $\mathrm{HTTP}/3$  connections over TLS requires the TLSv1.3 protocol support, which is available since OpenSSL version 1.1.1.

#### **Directives**

# http3

| Syntax  | http3 on   off; |
|---------|-----------------|
| Default | http3 on;       |
| Context | http, server    |

Enables HTTP/3 protocol negotiation.

# http3 hq

| Syntax  | http3_hq on   off; |
|---------|--------------------|
| Default | http3_hq off;      |
| Context | http, server       |

Enables  $\mathrm{HTTP}/0.9$  protocol negotiation used in QUIC interoperability tests.

### **A** Attention

Enable this mode only to run specialized tests that explicitly require it.

# http3\_max\_concurrent\_streams

| Syntax  | http3_max_concurrent_streams number; |
|---------|--------------------------------------|
| Default | http3_max_concurrent_streams 128;    |
| Context | http, server                         |

Initializes  $\mathrm{HTTP}/3$  and  $\mathrm{QUIC}$  settings and sets the maximum number of concurrent  $\mathrm{HTTP}/3$  request streams in a connection.

# http3\_max\_table\_capacity

| Syntax  | $\verb http3_max_table_capacity  number;$ |
|---------|---|
| Default | http3_max_table_capacity 4096;            |
| Context | http, server                              |

Sets the dynamic table capacity for server connections.



# 1 Note

A similar  $proxy\_http3\_max\_table\_capacity$  directive does this for proxy connections. To avoid errors, dynamic table usage is disabled when proxying with caching is enabled.

## http3 stream buffer size

| Syntax  | http3_stream_buffer_size $size$ ; |
|---------|-----------------------------------|
| Default | http3_stream_buffer_size 64k;     |
| Context | http, server                      |

Sets the size of the buffer used for reading and writing of the QUIC streams.

## quic active connection id limit

| Syntax  | quic_active_connection_id_limit number;       |
|---------|---|
| Default | <pre>quic_active_connection_id_limit 2;</pre> |
| Context | http, server                                  |

Sets the QUIC active\_connection\_id\_limit transport parameter value. This is the maximum number of client connection IDs which can be stored on the server.

# quic\_bpf

| Syntax  | <pre>quic_bpf on   off;</pre> |
|---------|-------------------------------|
| Default | <pre>quic_bpf off;</pre>      |
| Context | main                          |

Enables routing of QUIC packets using eBPF. When enabled, this allows supporting QUIC connection migration.

### Important

The directive is only supported on Linux 5.7+.

# quic\_gso

| Syntax  | quic_gso on   off;       |
|---------|--------------------------|
| Default | <pre>quic_gso off;</pre> |
| Context | http, server             |

Enables sending in optimized batch mode using segmentation offloading.

# Important

Optimized sending is supported only on Linux featuring UDP SEGMENT.



## quic\_host\_key

| Syntax  | $	exttt{quic_host_key} \ file;$ |
|---------|---------------------------------|
| Default | _                               |
| Context | http, server                    |

Sets a *file* with the secret key used to encrypt stateless reset and address validation tokens. By default, a random key is generated on each reload. Tokens generated with old keys are not accepted.

#### quic retry

| Syntax  | quic_retry on   off;       |
|---------|----------------------------|
| Default | <pre>quic_retry off;</pre> |
| Context | http, server               |

Enables the QUIC Address Validation feature. This includes sending a new token in a *Retry* packet or a *NEW TOKEN* frame and validating a token received in the *Initial* packet.

#### **Built-in Variables**

The http\_v3 module supports the following built-in variables:

#### \$http3

negotiated protocol identifier:

| h3    | for HTTP/3 connections    |
|-------|---------------------------|
| hq    | for hq connections        |
| 11 11 | an empty string otherwise |

## \$quic\_connection

QUIC connection serial number

## **XSLT**

The module is a filter that transforms XML responses using one or more XSLT stylesheets.

When building from the source code, this module isn't built by default; it should be enabled with the --with-http\_xslt\_module build option.

In our repositories, the module is built dynamically and is available as a separate package named angie-module-xslt or angie-pro-module-xslt.

# Important

This module requires the libxml2 and libxslt libraries.

#### **Configuration Example**

```
location / {
    xml_entities /site/dtd/entities.dtd;
    xslt_stylesheet /site/xslt/one.xslt param=value;
```



```
xslt_stylesheet /site/xslt/two.xslt;
}
```

#### **Directives**

#### xml entities

| Syntax  | xml_entities path;     |
|---------|------------------------|
| Default | _                      |
| Context | http, server, location |

Specifies the DTD file that declares character entities. This file is compiled at the configuration stage. For technical reasons, the module is unable to use the external subset declared in the processed XML, so it is ignored and a specially defined file is used instead. This file should not describe the XML structure. It is enough to declare just the required character entities, for example:

```
<!ENTITY nbsp "&#xa0;">
```

### xslt last modified

| Syntax  | xslt_last_modified on   off; |
|---------|------------------------------|
| Default | xslt_last_modified off;      |
| Context | http, server, location       |

Allows preserving the "Last-Modified" header field from the original response during XSLT transformations to facilitate response caching.

By default, the header field is removed as contents of the response are modified during transformations and may contain dynamically generated elements or parts that are changed independently of the original response.

## xslt param

| Syntax  | xslt_param parameter value; |
|---------|-----------------------------|
| Default | _                           |
| Context | http, server, location      |

Defines the parameters for XSLT stylesheets. The value is treated as an XPath expression. The value can contain variables. To pass a string value to a stylesheet, the  $xslt\_string\_param$  directive can be used.

There could be several  $xslt\_param$  directives. These directives are inherited from the previous configuration level if and only if there are no  $xslt\_param$  and  $xslt\_string\_param$  directives defined on the current level.

# xslt\_string\_param

| Syntax  | xslt_string_param parameter value; |
|---------|------------------------------------|
| Default | _                                  |
| Context | http, server, location             |

Defines the string parameters for XSLT stylesheets. XPath expressions in the value are not interpreted. The value can contain variables.



There could be several  $xslt\_string\_param$  directives. These directives are inherited from the previous configuration level if and only if there are no  $xslt\_param$  and  $xslt\_string\_param$  directives defined on the current level.

#### xslt stylesheet

| Syntax  | xslt_stylesheet stylesheet [parameter=value]; |
|---------|---|
| Default | _   |
| Context | location                                      |

Defines the XSLT stylesheet and its optional parameters. A stylesheet is compiled at the configuration stage.

Parameters can either be specified separately, or grouped in a single line using the ":" delimiter. If a parameter includes the ":" character, it should be escaped as "%3A". Also, libxslt requires to enclose parameters that contain non-alphanumeric characters into single or double quotes, for example:

```
param1='http%3A//www.example.com':param2=value2
```

The parameters description can contain variables, for example, the whole line of parameters can be taken from a single variable:

It is possible to specify several stylesheets. They will be applied sequentially in the specified order.

# xslt types

| Syntax  | xslt_types mime-type;           |
|---------|---------------------------------|
| Default | <pre>xslt_types text/xml;</pre> |
| Context | http, server, location          |

Enables transformations in responses with the specified MIME types in addition to text/xml. The special value "\*" matches any MIME type. If the transformation result is an HTML response, its MIME type is changed to text/html.

The core HTTP module implements the basic functionality of an HTTP server: this includes defining server blocks, configuring locations for request routing, serving static files and controlling access, configuring redirects, supporting keep-alive connections, and managing request and response headers.

The other modules in this section extend this functionality, allowing you to flexibly configure and optimize the HTTP server for various scenarios and requirements.

#### **Directives**

#### absolute redirect

| Syntax  | absolute_redirect on   off; |
|---------|-----------------------------|
| Default | absolute_redirect on;       |
| Context | http, server, location      |



If disabled, redirects issued by Angie will be relative.

See also server name in redirect and port in redirect directives.

aio

```
 \begin{array}{lll} {\rm Syntax} & {\rm aio\ on\ |\ off\ |\ threads\ [=pool];} \\ {\rm Default} & {\rm aio\ off;} \\ {\rm Context} & {\rm http,\ server,\ location} \end{array}
```

Enables or disables the use of asynchronous file I/O (AIO) on FreeBSD and Linux:

```
location /video/ {
  aio         on;
  output_buffers 1 64k;
}
```

On FreeBSD, AIO can be used starting from FreeBSD 4.3. Prior to FreeBSD 11.0, AIO can either be linked statically into a kernel:

```
options VFS_AIO
```

or loaded dynamically as a kernel loadable module:

```
kldload aio
```

On Linux, AIO can be used starting from kernel version 2.6.22. Also, it is necessary to enable *directio*, or otherwise reading will be blocking:

```
location /video/ {
  aio          on;
  directio     512;
  output_buffers 1 128k;
}
```

On Linux, *directio* can only be used for reading blocks that are aligned on 512-byte boundaries (or 4K for XFS). File's unaligned end is read in blocking mode. The same holds true for byte range requests and for FLV requests not from the beginning of a file: reading of unaligned data at the beginning and end of a file will be blocking.

When both AIO and *sendfile* are enabled on Linux, AIO is used for files that are larger than or equal to the size specified in the *directio* directive, while *sendfile* is used for files of smaller sizes or when *directio* is disabled:

```
location /video/ {
  sendfile    on;
  aio    on;
  directio    8m;
}
```

Finally, files can be read and sent using multi-threading, without blocking a worker process:

```
location /video/ {
  sendfile    on;
  aio         threads;
}
```

Read and send file operations are offloaded to threads of the specified *pool*. If the pool name is omitted, the pool with the name "default" is used. The pool name can also be set with variables:



```
aio threads=pool$disk;
```

By default, multi-threading is disabled, it should be enabled with the --with-threads configuration parameter. Currently, multi-threading is compatible only with the epoll, kqueue and eventport methods. Multi-threaded sending of files is only supported on Linux.

See also the *sendfile* directive.

### aio\_write

| Syntax  | aio_write on   off;    |
|---------|------------------------|
| Default | aio_write off;         |
| Context | http, server, location |

If *aio* is enabled, specifies whether it is used for writing files. Currently, this only works when using *aio* threads and is limited to writing temporary files with data received from proxied servers.

#### alias

| Syntax  | alias $path$ ; |
|---------|----------------|
| Default | _              |
| Context | location       |

Defines a replacement for the specified location. For example, with the following configuration:

```
location /i/ {
  alias /data/w3/images/;
}
```

on request of /i/top.gif, the file /data/w3/images/top.gif will be sent.

The path value can contain variables, except \$document root and \$realpath root.

If alias is used inside a location defined with a regular expression then such regular expression should contain captures and alias should refer to these captures, for example:

```
location ~ ^/users/(.+\.(?:gif|jpe?g|png))$ {
   alias /data/w3/images/$1;
}
```

When location matches the last part of the directive's value:

```
location /images/ {
  alias /data/w3/images/;
}
```

it is better to use the *root* directive instead:

```
location /images/ {
  root /data/w3;
}
```



### auth\_delay

| Syntax  | $auth\_delay \ time;$     |
|---------|---------------------------|
| Default | <pre>auth_delay 0s;</pre> |
| Context | http, server, location    |

Delays processing of unauthorized requests with 401 response code to prevent timing attacks when access is limited by password or by the result of subrequest.

## auto redirect

| Syntax  | auto_redirect [on   off   default]; |
|---------|-------------------------------------|
| Default | <pre>auto_redirect default;</pre>   |
| Context | http, server, location              |

The directive controls the redirection behavior when a prefix location ends with a slash:

```
location /prefix/ {
    auto_redirect on;
}
```

Here, a request for /prefix causes a redirect to /prefix/.

The value on explicitly enables redirection, while off disables it. When set to default, redirection is enabled only if the location processes requests with api, proxy\_pass, fastcgi\_pass, uwsgi\_pass, segi pass, memcached pass, or grpc pass.

#### chunked \_transfer \_encoding

| Syntax  | chunked_transfer_encoding on   off;      |
|---------|--|
| Default | <pre>chunked_transfer_encoding on;</pre> |
| Context | http, server, location                   |

Allows disabling chunked transfer encoding in  $\mathrm{HTTP}/1.1$ . It may come in handy when using a software failing to support chunked encoding despite the standard's requirement.

### client\_body\_buffer\_size

| Syntax  | <pre>client_body_buffer_size size;</pre>   |
|---------|--|
| Default | <pre>client_body_buffer_size 8k 16k;</pre> |
| Context | http, server, location                     |

Sets buffer size for reading client request body. In case the request body is larger than the buffer, the whole body or only its part is written to a *temporary file*. By default, buffer size is equal to two memory pages. This is 8K on x86, other 32-bit platforms, and x86-64. It is usually 16K on other 64-bit platforms.

## client\_body\_in\_file\_only

| Syntax  | client_body_in_file_only on   clean   off; |
|---------|--|
| Default | <pre>client_body_in_file_only off;</pre>   |
| Context | http, server, location                     |



Determines whether Angie should save the entire client request body into a file. This directive can be used during debugging, or when using the  $request\_body\_file$  variable, or the  $request\_body\_file$  method of the module  $request\_body\_file$  method of the method of the module  $request\_body\_file$  method of the module requ

When set to the value on, temporary files are not removed after request processing.

| on    | temporary files are not removed after request processing                   |
|-------|--|
| clean | will cause the temporary files left after request processing to be removed |

#### client body in single buffer

| Syntax  | <pre>client_body_in_single_buffer on   off;</pre> |
|---------|---|
| Default | <pre>client_body_in_single_buffer off;</pre>      |
| Context | http, server, location                            |

Determines whether Angie should save the entire client request body in a single buffer. The directive is recommended when using the *\$request\_body* variable, to save the number of copy operations involved.

## client body temp path

| Syntax  | client_body_temp_path path [level1 [level2 [level3]]];  |
|---------|---|
| Default | <pre>client_body_temp_path client_body_temp; (the path depends on thehttp-proxy-temp-path build option)</pre> |
| Context | http, server, location  |

Defines a directory for storing temporary files holding client request bodies. Up to three-level subdirectory hierarchy can be used under the specified directory. For example, in the following configuration

```
client_body_temp_path /spool/angie/client_temp 1 2;
```

a path to a temporary file might look like this:

```
/spool/angie/client_temp/7/45/00000123457
```

## client body timeout

| Syntax  | client_body_timeout $time;$         |
|---------|-------------------------------------|
| Default | <pre>client_body_timeout 60s;</pre> |
| Context | http, server, location              |

Defines a timeout for reading client request body. The timeout is set only for a period between two successive read operations, not for the transmission of the whole request body. If a client does not transmit anything within this time, the request is terminated with the 408 (Request Time-out) error.

### client header buffer size

| Syntax  | client_header_buffer_size $size;$        |
|---------|--|
| Default | <pre>client_header_buffer_size 1k;</pre> |
| Context | http, server                             |

Sets buffer size for reading client request header. For most requests, a buffer of 1K bytes is enough. However, if a request includes long cookies, or comes from a WAP client, it may not fit into 1K. If a



request line or a request header field does not fit into this buffer then larger buffers, configured by the large client header buffers directive, are allocated.

If the directive is specified on the *server* level, the value from the default server can be used. Details are provided in the *Virtual server selection* section.

# client\_header\_timeout

| Syntax  | client_header_timeout $time$ ;        |
|---------|---------------------------------------|
| Default | <pre>client_header_timeout 60s;</pre> |
| Context | http, server                          |

Defines a timeout for reading client request header. If a client does not transmit the entire header within this time, the request is terminated with the 408 (Request Time-out) error.

# client max body size

| Syntax  | client_max_body_size $size;$        |
|---------|-------------------------------------|
| Default | <pre>client_max_body_size 1m;</pre> |
| Context | http, server, location              |

Sets the maximum allowed size of the client request body. If the size in a request exceeds the configured value, the 413 (Request Entity Too Large) error is returned to the client. Please be aware that browsers cannot correctly display this error.

| 0 disa | ables checking of client request body size |
|--------|--|
|--------|--|

# connection pool size

| Syntax  | connection_pool_size size;      |
|---------|---------------------------------|
| Default | connection_pool_size 256   512; |
| Context | http, server, location          |

Allows accurate tuning of per-connection memory allocations. This directive has minimal impact on performance and should not generally be used.

By default:

| 256 (bytes) | 32-bit platforms |  |
|-------------|------------------|--|
| 512 (bytes) | 64-bit platforms |  |

# default\_type

| Syntax  | default_type mime-type;             |
|---------|-------------------------------------|
| Default | <pre>default_type text/plain;</pre> |
| Context | http, server, location              |

Defines the default MIME type of a response. Mapping of file name extensions to MIME types can be set with the types directive.



#### directio

| Syntax  | directio $size \mid \texttt{off};$ |
|---------|------------------------------------|
| Default | directio off;                      |
| Context | http, server, location             |

Enables the use of the  $O\_DIRECT$  flag (FreeBSD, Linux), the  $F\_NOCACHE$  flag (macOS), or the directio() function (Solaris), when reading files that are larger than or equal to the specified size. The directive automatically disables the use of sendfile for a given request. It can be useful for serving large files:

```
directio 4m;
```

or when using aio on Linux.

# directio alignment

| Syntax  | directio_alignment $size;$ |
|---------|----------------------------|
| Default | directio_alignment 512;    |
| Context | http, server, location     |

Sets the alignment for *directio*. In most cases, a 512-byte alignment is enough. However, when using XFS under Linux, it needs to be increased to 4K.

# disable\_symlinks

| Syntax  | disable_symlinks off;                           |
|---------|---|
|         | disable_symlinks on   if_not_owner [from=part]; |
| Default | disable_symlinks off;                           |
| Context | http, server, location                          |

Determines how symbolic links should be treated when opening files:

| off                     | Symbolic links in the pathname are allowed and not checked. This is the default behavior.   |
|-------------------------|---|
| on                      | If any component of the pathname is a symbolic link, access to a file is denied.  |
| <pre>if_not_owner</pre> | Access to a file is denied if any component of the pathname is a symbolic link, and the link and object that the link points to have different owners.  |
| ${\tt from} = part$     | When checking symbolic links (parameters on and if_not_owner), all components of the pathname are normally checked. Checking of symbolic links in the initial part of the pathname may be avoided by specifying additionally the from=part parameter. In this case, symbolic links are checked only from the pathname component that follows the specified initial part. If the value is not an initial part of the pathname checked, the whole pathname is checked as if this parameter was not specified at all. If the value matches the whole file name, symbolic links are not checked. The parameter value can contain variables. |

# Example:

```
disable_symlinks on from=$document_root;
```

This directive is only available on systems that have the openat() and fstatat() interfaces. Such systems include modern versions of FreeBSD, Linux, and Solaris.



# **A** Warning

Parameters on and if\_not\_owner add a processing overhead.

On systems that do not support opening of directories only for search, to use these parameters it is required that worker processes have read permissions for all directories being checked.

# 1 Note

The AutoIndex, Random Index and DAV modules currently ignore this directive.

#### error\_page

```
Syntax error_page code ... [=[response]] uri;
Default —
Context http, server, location, if in location
```

Defines the URI that will be shown for the specified errors. A uri value can contain variables.

#### Example:

```
error_page 404 /404.html;
error_page 500 502 503 504 /50x.html;
```

This causes an internal redirect to the specified uri with the client request method changed to "GET" (for all methods other than "GET" and "HEAD").

Furthermore, it is possible to change the response code to another using the **=response** syntax, for example:

```
error_page 404 =200 /empty.gif;
```

If an error response is processed by a proxied server or a FastCGI/uwsgi/SCGI/gRPC server, and the server may return different response codes (e.g., 200, 302, 401 or 404), it is possible to respond with the code it returns:

```
error_page 404 = /404.php;
```

If there is no need to change URI and method during internal redirection it is possible to pass error processing into a named location:

```
location / {
   error_page 404 = @fallback;
}

location @fallback {
   proxy_pass http://backend;
}
```

## Note

If uri processing leads to an error, the status code of the last occurred error is returned to the client.

It is also possible to use URL redirects for error processing:



```
error_page 403         http://example.com/forbidden.html;
error_page 404 =301 http://example.com/notfound.html;
```

In this case, by default, the response code 302 is returned to the client. It can only be changed to one of the redirect status codes (301, 302, 303, 307, and 308).

#### etag

| Syntax  | etag on   off;         |
|---------|------------------------|
| Default | etag on;               |
| Context | http, server, location |

Enables or disables automatic generation of the "ETag" response header field for static resources.

# http

| Syntax  | http { } |
|---------|----------|
| Default | _        |
| Context | main     |

Provides the configuration file context in which the HTTP server directives are specified.

# if modified since

| Syntax  | if_modified_since off   exact   before; |
|---------|---|
| Default | <pre>if_modified_since exact;</pre>     |
| Context | http, server, location                  |

Specifies how to compare modification time of a response with the time in the *If-Modified-Since* request header field:

| off    | the response is always considered modified  |
|--------|---|
| exact  | exact match   |
| before | modification time of the response is less than or equal to the time in the <i>If-Modified-Since</i> request header field. |

# ignore\_invalid\_headers

| Syntax  | ignore_invalid_headers on   off; |
|---------|----------------------------------|
| Default | ignore_invalid_headers on;       |
| Context | http, server                     |

Controls whether header fields with invalid names should be ignored. Valid names are composed of English letters, digits, hyphens, and possibly underscores (as controlled by the ref: underscores\_in\_headers directive).

If the directive is specified on the server level, the value from the default server can be used.



### internal

| Syntax  | internal; |
|---------|-----------|
| Default | _         |
| Context | location  |

Specifies that a given location can only be used for internal requests. For external requests, the client error 404 (Not Found) is returned. Internal requests are the following:

- requests redirected by the error page, index, random index and try files directives;
- ullet requests redirected by the X-Accel-Redirect response header field from an upstream server;
- subrequests formed by the *include virtual* command of the *http\_ssi* module, by the *http\_addition* module directives, and by *auth request* and *mirror* directives;
- requests changed by the *rewrite* directive.

#### Example:

```
error_page 404 /404.html;

location = /404.html {
  internal;
}
```

# 1 Note

There is a limit of 10 internal redirects per request to prevent request processing cycles that can occur in incorrect configurations. If this limit is reached, the error 500 (Internal Server Error) is returned. In such cases, the *rewrite or internal redirection cycle* message can be seen in the error log.

# keepalive\_disable

| Syntax  | keepalive_disable none   browser; |
|---------|-----------------------------------|
| Default | keepalive_disable msie6;          |
| Context | http, server, location            |

Disables keep-alive connections with misbehaving browsers. The browser parameters specify which browsers will be affected.

| none   | enables keep-alive connections with all browsers   |
|--------|--|
| msie6  | disables keep-alive connections with old versions of MSIE, once a POST request is received                     |
| safari | disables keep-alive connections with Safari and Safari-like browsers on macOS and macOS-like operating systems |

### keepalive requests

| Syntax  | keepalive_requests number; |
|---------|----------------------------|
| Default | keepalive_requests 1000;   |
| Context | http, server, location     |



Sets the maximum number of requests that can be served through one keep-alive connection. After the maximum number of requests are made, the connection is closed.

Closing connections periodically is necessary to free per-connection memory allocations. Therefore, using too high maximum number of requests could result in excessive memory usage and not recommended.

# keepalive\_time

| Syntax  | $\verb keepalive_time  time;$ |
|---------|-------------------------------|
| Default | keepalive_time 1h;            |
| Context | http, server, location        |

Limits the maximum time during which requests can be processed through one keep-alive connection. After this time is reached, the connection is closed following the subsequent request processing.

# keepalive\_timeout

| Syntax  | keepalive_timeout timeout [header_timeout]; |
|---------|---|
| Default | keepalive_timeout 75s;                      |
| Context | http, server, location                      |

| timeout | sets a timeout during which a keep-alive client connection will stay open on the server side |
|---------|--|
| 0       | disables keep-alive client connections   |

The *optional* second parameter sets a value in the "Keep-Alive: timeout=time" response header field. Two parameters may differ.

The "Keep-Alive: timeout=time" header field is recognized by Mozilla and Konqueror. MSIE closes keep-alive connections by itself in about 60 seconds.

# large\_client\_header\_buffers

| Syntax  | large_client_header_buffers number size;     |
|---------|--|
| Default | <pre>large_client_header_buffers 4 8k;</pre> |
| Context | http, server                                 |

Sets the maximum number and size of buffers used for reading large client request header. A request line cannot exceed the size of one buffer, or the 414 (Request-URI Too Large) error is returned to the client. A request header field cannot exceed the size of one buffer as well, or the 400 (Bad Request) error is returned to the client. Buffers are allocated only on demand. By default, the buffer size is equal to 8K bytes. If after the end of request processing a connection is transitioned into the keep-alive state, these buffers are released.

If the directive is specified on the server level, the value from the default server can be used.

# limit except

| Syntax  | <pre>limit_except method1 [method2] { };</pre> |
|---------|--|
| Default | _  |
| Context | location                                       |



Limits allowed HTTP methods inside a location. The method can be one of the following: GET, HEAD, POST, PUT, DELETE, MKCOL, COPY, MOVE, OPTIONS, PROPFIND, PROPPATCH, LOCK, UNLOCK or PATCH. Allowing the GET method also enabled the HEAD method. Access to other methods can be limited using the directives from *Access* and *Auth Basic* modules:

```
limit_except GET {
  allow 192.168.1.0/32;
  deny all;
}
```

# 1 Note

This example will limit access to all methods, except GET and HEAD.

# limit rate

```
Syntax limit_rate rate;
Default limit_rate 0;
Context http, server, location, if in location
```

Limits the rate of response transmission to a client. The rate is specified in bytes per second. The zero value disables rate limiting. The limit is set per a request, and so if a client simultaneously opens two connections, the overall rate will be twice as much as the specified limit.

Parameter value can contain variables. It may be useful in cases where rate should be limited depending on a certain condition:

```
map $slow $rate {
   1   4k;
   2   8k;
}
limit_rate $rate;
```

Rate limit can also be set in the *\$limit rate* variable, however, this method is not recommended:

```
server {
  if ($slow) {
    set $limit_rate 4k;
  }
}
```

Rate limit can also be set in the "X-Accel-Limit-Rate" header field of a proxied server response. This capability can be disabled using the <code>proxy\_ignore\_headers</code>, <code>fastcgi\_ignore\_headers</code>, <code>uwsgi\_ignore\_headers</code> and <code>scgi\_ignore\_headers</code> directives.

### limit rate after

| Syntax  | limit_rate_after $size;$               |
|---------|--|
| Default | <pre>limit_rate_after 0;</pre>         |
| Context | http, server, location, if in location |



Sets the initial amount after which the further transmission of a response to a client will be rate limited. Parameter value can contain variables.

### Example:

```
location /flv/ {
  flv;
  limit_rate_after 500k;
  limit_rate 50k;
}
```

# lingering close

| Syntax  | lingering_close on   always   off; |
|---------|------------------------------------|
| Default | <pre>lingering_close on;</pre>     |
| Context | http, server, location             |

Controls how Angie closes client connections.

| on     | instructs Angie to <i>wait</i> for and <i>process</i> additional data from a client before fully closing a connection, but only if heuristics suggests that a client may be sending more data. |
|--------|--|
| always | will cause Angie to unconditionally wait for and process additional client data.   |
| off    | tells Angie to never wait for more data and close the connection immediately. This behavior breaks the protocol and should not be used under normal circumstances.                             |

To control closing HTTP/2 connections, the directive must be specified on the *server* level.

# lingering time

| Syntax  | lingering_time time;           |
|---------|--------------------------------|
| Default | <pre>lingering_time 30s;</pre> |
| Context | http, server, location         |

When *lingering\_close* is in effect, this directive specifies the maximum time during which Angie will process (read and ignore) additional data coming from a client. After that, the connection will be closed, even if there will be more data.

# lingering\_timeout

| Syntax  | ${\tt lingering\_timeout}\ time;$ |
|---------|-----------------------------------|
| Default | lingering_timeout 5s;             |
| Context | http, server, location            |

When  $lingering\_close$  is in effect, the directive specifies the maximum waiting time for more client data to arrive. If data are not received during this time, the connection is closed. Otherwise, the data are read and ignored, and Angie starts waiting for more data again. This "wait-read-ignore" cycle is repeated no longer than specified by the  $lingering\_time$  directive.

During a graceful shutdown, keepalive connections from clients are closed only if they remain inactive for at least the duration of lingering\_timeout.



#### listen

| Syntax  | ${\tt listen} \ \ address[:port] \ \ [{\tt default\_server}] \ \ [{\tt ssl}] \ \ [{\tt http2} \ \   \ \ {\tt quic}] \ \ [{\tt proxy\_protocol}]$   |
|---------|--|
|         | [setfib=number] $[fastopen=number]$ $[backlog=number]$ $[rcvbuf=size]$   |
|         | $[\mathtt{sndbuf} = size]$ $[\mathtt{accept\_filter} = filter]$ $[\mathtt{deferred}]$ $[\mathtt{bind}]$ $[\mathtt{ipv6only=on} \mid \mathtt{off}]$ |
|         | $[\texttt{reuseport}] \ [\texttt{so\_keepalive=} on   off   [keepidle] : [keepintvl] : [keepcnt]];$  |
|         | listen port [default_server] [ssl] [http2   quic] [proxy_protocol]   |
|         | $[\mathtt{setfib} = number] \qquad [\mathtt{fastopen} = number] \qquad [\mathtt{backlog} = number] \qquad [\mathtt{rcvbuf} = size]$                |
|         | $[\mathtt{sndbuf} = size]$ $[\mathtt{accept\_filter} = filter]$ $[\mathtt{deferred}]$ $[\mathtt{bind}]$ $[\mathtt{ipv6only=on} \mid \mathtt{off}]$ |
|         | $[\texttt{reuseport}] \ [\texttt{so\_keepalive=} on   off   [keepidle] : [keepintvl] : [keepcnt]];$  |
|         | listen unix:path [default_server] [ssl] [http2   quic] [proxy_protocol]  |
|         | [backlog=number] [rcvbuf=size] [sndbuf=size] [accept_filter=filter] [deferred]   |
|         | $[bind] [so\_keepalive=on off [keepidle]:[keepintvl]:[keepcnt]];$  |
| Default | listen *:80   *:8000;  |
| Context | server   |

Sets the address and port for listen socket, or the path for a UNIX domain socket on which the server will accept requests. An address may also be a hostname, for example:

```
listen 127.0.0.1:8000;
listen 127.0.0.1;
listen 8000;
listen *:8000;
listen localhost:8000;
```

IPv6 addresses are specified in square brackets:

```
listen [::]:8000;
listen [::1];
```

UNIX domain sockets are specified with the unix: prefix:

```
listen unix:/var/run/angie.sock;
```

Both address and port, or only address or port, can be specified. When some parts are omitted, the behavior varies:

- If only the address is given, port 80 is used.
- If only the port is given, Angie listens on all available IPv4 (and IPv6, if enabled) interfaces. The first defined server block for that port becomes the default for requests with an unmatched Host header.
- If the directive is omitted entirely, Angie uses \*:80 when running with superuser privileges or \*:8000 otherwise.



| default_server | The server with this parameter specified will be the default server for the given address:port pair (together they form a <i>listening socket</i> ). If there are no directives with the default_server parameter, the default server for the listening socket will be the first server in the configuration that serves this socket. |
|----------------|---|
| ssl            | allows specifying that all connections accepted on this port should work in SSL mode. This allows for a more compact configuration for the server that handles both HTTP and HTTPS requests.  |
| http2          | configures the port to accept HTTP/2 connections. Normally, for this to work the ssl parameter should be specified as well, but Angie can also be configured to accept HTTP/2 connections without SSL. Deprecated since version 1.2.0. Use the $http2$ directive instead.   |
| quic           | configures the port to accept QUIC connections. To use this option, Angie must have the <i>HTTP3 module</i> enabled and configured. With quic set, you can also specify reuseport so multiple worker processes can be used.   |
| proxy_protocol | allows specifying that all connections accepted on this port should use the PROXY protocol.   |

The listen directive can have several additional parameters specific to socket-related system calls. These parameters can be specified in any listen directive, but only once for a given address:port pair.

| setfib=`number`  | this parameter sets the associated routing table, FIB (the SO_SETFIB option) for the listening socket. This currently works only on FreeBSD. |
|------------------|--|
| fastopen=`number | enables "TCP Fast Open" for the listening socket and limits the maximum length   |
|                  | for the queue of connections that have not yet completed the three-way hand-   |
|                  | shake.   |

# **Caution**

Do not enable this feature unless the server can handle receiving the same SYN packet with data more than once.



| backlog=`number`                  | sets the backlog parameter in the listen() call that limits the maximum length for the queue of pending connections. By default, backlog is set to -1 on FreeBSD, DragonFly BSD, and macOS, and to 511 on other platforms.   |
|-----------------------------------|--|
| rcvbuf=size                       | sets the receive buffer size (the SO_RCVBUF option) for the listening socket.  |
| $\mathtt{sndbuf} \texttt{=} size$ | sets the send buffer size (the SO_SNDBUF option) for the listening socket.   |
| accept_filter=fi                  | sets the name of accept filter (the SO_ACCEPTFILTER option) for the listening socket that filters incoming connections before passing them to accept(). This works only on FreeBSD and NetBSD 5.0+. Possible values are dataready and httpready.   |
| deferred                          | instructs to use a deferred accept() (the TCP_DEFER_ACCEPT socket option) on Linux.  |
| bind                              | instructs to make a separate bind() call for a given address:port pair. This is useful because if there are several <i>listen</i> directives with the same port but different addresses, and one of the listen directives listens on all addresses for the given port (*:port), Angie will bind() only to *:port. It should be noted that the getsockname() system call will be made in this case to determine the address that accepted the connection. If the setfib, fastopen, backlog, rcvbuf, sndbuf, accept_filter, deferred, ipv6only, reuseport or so_keepalive parameters are used then for a given address:port pair a separate bind() call will always be made. |
| ipv6only=on   off                 | this parameter determines (via the IPV6_V60NLY socket option) whether an IPv6 socket listening on a wildcard address [::] will accept only IPv6 connections or both IPv6 and IPv4 connections. This parameter is turned on by default. It can only be set once on start.   |
| reuseport                         | this parameter instructs to create an individual listening socket for each worker process (using the SO_REUSEPORT socket option on Linux 3.9+ and DragonFly BSD, or SO_REUSEPORT_LB on FreeBSD 12+), allowing a kernel to distribute incoming connections between worker processes. This currently works only on Linux 3.9+, DragonFly BSD, and FreeBSD 12+.   |

# \* Caution

Inappropriate use of this option may have its security implications.

 $so_keepalive=on | off | [`keepidle]: | keepintvl|: | keepintvl| ` | keepintvl|$ 

Configures the "TCP keepalive" behavior for the listening socket.

| 1.1 | if this parameter is omitted then the operating system's settings will be in effect for the socket |
|-----|--|
| on  | the SO_KEEPALIVE option is turned on for the socket  |
| off | the SO_KEEPALIVE option is turned off for the socket   |

Some operating systems support setting of TCP keepalive parameters on a per-socket basis using the TCP\_KEEPIDLE, TCP\_KEEPINTVL, and TCP\_KEEPCNT socket options. On such systems (currently, Linux 2.4+, NetBSD 5+, and FreeBSD 9.0-STABLE), they can be configured using the keepidle, keepintvl, and keepcnt parameters. One or two parameters may be omitted, in which case the system default setting for the corresponding socket option will be in effect. For example,

# so\_keepalive=30m::10

will set the idle timeout (TCP\_KEEPIDLE) to 30 minutes, leave the probe interval (TCP\_KEEPINTVL) at its system default, and set the probes count (TCP\_KEEPCNT) to 10 probes.

Example:



```
listen 127.0.0.1 default_server accept_filter=dataready backlog=1024;
```

#### location

| Syntax  | $\texttt{location} \; ([\; = \;   \; \widehat{\ } \;   \; \widehat{\ } \; \widehat{\ } \; ] \; uri \;   \; @name) + \; \{ \; \; \}$ |
|---------|---|
| Default | _   |
| Context | server, location  |

Sets the configuration depending on whether the request URI matches any of the matching expressions.

The matching is performed against a normalized URI, after decoding the text encoded in the "%XX" form, resolving references to relative path components "." and "..", and possible *compression* of two or more adjacent slashes into a single slash.

A location can either be defined by a prefix string, or by a regular expression.

Regular expressions are specified with the preceding modifier:

| ~* | Case-insensitive matching |
|----|---------------------------|
| ~  | Case-sensitive matching   |

To find a location that matches a request, Angie first checks the locations defined with prefix strings (known as prefix locations). Among them, the location with the longest matching prefix is selected and tentatively stored.

# 1 Note

For case-insensitive operating systems such as macOS, prefix string matching is case insensitive. However, matching is limited to single-byte locales.

Then, regex-based locations are evaluated in order of their appearance in the configuration file. Their evaluation stops at the first match, and the corresponding configuration is used. If no matching regex location is found, Angie uses the configuration of the tentatively stored prefix location.

With some exceptions mentioned below, location blocks can be nested.

Regex locations may define capture groups that can later be used with other directives.

If the matching prefix location uses the ~~ modifier, regex locations aren't checked.

Also, the = modifier enables exact URI matching mode for a location; if an exact match is found, the lookup stops. For example, if / requests are frequent, defining location =/ speeds up their processing because the lookup stops at the exact match. Obviously, such locations can't contain nested locations.

Example:

```
location =/ {
    #configuration A
}

location / {
    #configuration B
}

location /documents/ {
    #configuration C
}
```



```
location ^~/images/ {
    #configuration D
}
location ~*\.(gif|jpg|jpeg)$ {
    #configuration E
}
```

- A / request matches configuration A,
- an /index.html request matches configuration B,
- a /documents/document.html request matches configuration C,
- an /images/1.gif request matches configuration D,
- and a /documents/1.jpg request matches configuration E.

### 1 Note

If a prefix location ends with a slash character and *auto\_redirect* is enabled, the following occurs: When a request arrives with the URI that has no trailing slash but otherwise matches the prefix exactly, a permanent 301 code redirect is returned, pointing to the requested URI with the slash appended.

With an exact URI-matching location, redirection isn't applied:

```
location /user/ {
  proxy_pass http://user.example.com;
}
location =/user {
  proxy_pass http://login.example.com;
}
```

The @ prefix defines a named location. Such locations aren't used for regular request processing, but instead can be used for request redirection. They cannot be nested and cannot contain nested locations.

# **Combined locations**

Several location contexts that define identical configuration blocks can be compacted by listing all their matching expressions in a single location with a single configuration block. That's called a *combined* location.

Suppose that configurations A, D, and E from the previous example define identical configurations; you can combine them into one location:

A named location can also be a part of the combination:

```
location =/
           @named_combined {
    #...
}
```



### Caution

A combined location can't have a space between the matching expression and its modifier. Proper form: location ~\*/match(ing|es|er)\$ ....

# 1 Note

Currently, a combined location cannot **immediately** contain neither *proxy\_pass* and similar directives with URI set, nor api or alias. However, these directives can be used by locations nested inside a combined location.

# log\_not\_found

| Syntax  | <pre>log_not_found on   off;</pre> |
|---------|------------------------------------|
| Default | <pre>log_not_found on;</pre>       |
| Context | http, server, location             |

Enables or disables logging of errors about not found files into error log.

# $log\_subrequest$

| Syntax  | log_subrequest on   off;       |
|---------|--------------------------------|
| Default | <pre>log_subrequest off;</pre> |
| Context | http, server, location         |

Enables or disables logging of subrequests into access\_log.

# max headers

| Syntax  | max_headers number; |
|---------|---------------------|
| Default | max_headers 1000;   |
| Context | http, server        |

Sets the maximum number of client request header fields allowed. If this limit is exceeded, a 400 (Bad Request) error is returned.

When this directive is set at the *server* level, the value from the default server may be applied. For more information, refer to the *Virtual server selection* section.

# max\_ranges

| Syntax  | $max\_ranges number;$  |
|---------|------------------------|
| Default | _                      |
| Context | http, server, location |

Limits the maximum allowed number of ranges in byte-range requests. Requests that exceed the limit are processed as if there were no byte ranges specified. By default, the number of ranges is not limited.

| 0 | disables the byte-range support completely |  |
|---|--|--|
|---|--|--|



# merge\_slashes

| Syntax  | merge_slashes on   off; |
|---------|-------------------------|
| Default | merge_slashes on;       |
| Context | http, server            |

Enables or disables compression of two or more adjacent slashes in a URI into a single slash.

Note that compression is essential for the correct matching of prefix string and regular expression locations. Without it, the //scripts/one.php request would not match

```
[location /scripts/ { }
```

and might be processed as a static file. So it gets converted to /scripts/one.php.

Turning the compression off can become necessary if a URI contains base64-encoded names, since base64 uses the "/" character internally. However, for security considerations, it is better to avoid turning the compression off.

If the directive is specified on the server level, the value from the default server can be used.

### msie padding

| Syntax  | msie_padding on   off;      |
|---------|-----------------------------|
| Default | <pre>msie_padding on;</pre> |
| Context | http, server, location      |

Enables or disables adding comments to responses for MSIE clients with status greater than 400 to increase the response size to 512 bytes.

# msie refresh

| Syntax  | msie_refresh on   off; |
|---------|------------------------|
| Default | msie_refresh off;      |
| Context | http, server, location |

Enables or disables issuing refreshes instead of redirects for MSIE clients.

# open\_file\_cache

| Syntax  | open_file_cache off; open_file_cache max= $N$ [inactive= $time$ ]; |
|---------|--|
| Default | open_file_cache off;   |
| Context | http, server, location   |

Configures a cache that can store:

- open file descriptors, their sizes and modification times;
- information on existence of directories;
- file lookup errors, such as "file not found", "no read permission", and so on.

Caching of errors should be enabled separately by the  $open\_file\_cache\_errors$  directive.



| max      | sets the maximum number of elements in the cache; on cache overflow the least recently used (LRU) elements are removed;                               |
|----------|---|
| inactive | defines a time after which an element is removed from the cache if it has not been accessed during this time;<br>By default, it is set to 60 seconds. |
| off      | disables the cache.   |

# Example:

# open\_file\_cache\_errors

| Syntax  | open_file_cache_errors on   off;       |
|---------|--|
| Default | <pre>open_file_cache_errors off;</pre> |
| Context | http, server, location                 |

Enables or disables caching of file lookup errors by open\_file\_cache.

# open\_file\_cache\_min\_uses

| Syntax  | open_file_cache_min_uses $number;$     |
|---------|--|
| Default | <pre>open_file_cache_min_uses 1;</pre> |
| Context | http, server, location                 |

Sets the minimum`number` of file accesses during the period configured by the inactive parameter of the <code>open\_file\_cache</code> directive, required for a file descriptor to remain open in the cache.

# open file cache valid

| Syntax  | ${\tt open\_file\_cache\_valid}\ time;$ |
|---------|---|
| Default | open_file_cache_valid 60s;              |
| Context | http, server, location                  |

Sets a time after which <code>open\_file\_cache</code> elements should be validated.

# $output\_buffers$

| Syntax  | output_buffers number size; |
|---------|-----------------------------|
| Default | output_buffers 2 32k;       |
| Context | http, server, location      |

Sets the `number ` and size of the buffers used for reading a response from a disk.



# port\_in\_redirect

| Syntax  | <pre>port_in_redirect on   off;</pre> |
|---------|---------------------------------------|
| Default | <pre>port_in_redirect on;</pre>       |
| Context | http, server, location                |

Enables or disables specifying the port in absolute redirects issued by Angie.

The use of the primary server name in redirects is controlled by the server name in redirect directive.

# $postpone\_output$

| Syntax  | postpone_output $size;$ |
|---------|-------------------------|
| Default | postpone_output 1460;   |
| Context | http, server, location  |

If possible, the transmission of client data will be postponed until Angie has at least size bytes of data to send.

| 0 | disables postponing data transmission |  |
|---|---------------------------------------|--|
|---|---------------------------------------|--|

### read ahead

| Syntax  | read_ahead $size;$     |
|---------|------------------------|
| Default | read_ahead 0;          |
| Context | http, server, location |

Sets the amount of pre-reading for the kernel when working with file.

On Linux, the posix\_fadvise(0, 0, 0, POSIX\_FADV\_SEQUENTIAL) system call is used, and so the size parameter is ignored.

# recursive error pages

| Syntax  | recursive_error_pages on   off; |
|---------|---------------------------------|
| Default | recursive_error_pages off;      |
| Context | http, server, location          |

Enables or disables doing several redirects using the  $error\_page$  directive. The number of such redirects is limited.

# request pool size

| Syntax  | request_pool_size $size;$ |
|---------|---------------------------|
| Default | request_pool_size 4k;     |
| Context | http, server              |

Allows accurate tuning of per-request memory allocations. This directive has minimal impact on performance and should not generally be used.



# reset \_ timedout \_ connection

| Syntax  | reset_timedout_connection on   off; |
|---------|-------------------------------------|
| Default | reset_timedout_connection off;      |
| Context | http, server, location              |

Enables or disables resetting timed out connections and connections closed with the non-standard code 444. The reset is performed as follows. Before closing a socket, the SO\_LINGER option is set on it with a timeout value of 0. When the socket is closed, TCP RST is sent to the client, and all memory occupied by this socket is released. This helps avoid keeping an already closed socket with filled buffers in a FIN\_WAIT1 state for a long time.



# 1 Note

timed out keep-alive connections are closed normally.

#### resolver

| Syntax  | resolver address [status_zone=zone]; | $[{\tt valid=} time]$ | [ipv4=on | off] | [ipv6=on | off] |
|---------|--------------------------------------|-----------------------|----------|------|----------|------|
| Default | _                                    |                       |          |      |          |      |
| Context | http, server, location, upstre       | eam                   |          |      |          |      |

Configures name servers used to resolve names of upstream servers into addresses, for example:

```
resolver 127.0.0.53 [::1]:5353;
```

The address can be specified as a domain name or IP address, with an optional port. If port is not specified, the port 53 is used. Name servers are queried in a round-robin fashion.

By default, Angie caches answers using the TTL value of a response.

| valid optional parameter allows overriding cached entry validity |
|--|
|--|

```
resolver 127.0.0.53 [::1]:5353 valid=30s;
```

By default, Angie will look up both IPv4 and IPv6 addresses while resolving.

| ipv4=off | disables looking up of IPv4 addresses |  |
|----------|---------------------------------------|--|
| ipv6=off | disables looking up of IPv6 addresses |  |

optional parameter; enables the collection of DNS server request and response status\_zone metrics (/status/resolvers/<zone>) in the specified zone.

# 🗘 Tip

To prevent DNS spoofing, it is recommended configuring DNS servers in a properly secured trusted local network.



# 🗘 Tip

When running in Docker, use its internal DNS server address such as 127.0.0.11.

# $resolver\_timeout$

| Syntax  | ${\tt resolver\_timeout}\ time;$ |
|---------|----------------------------------|
| Default | resolver_timeout 30s;            |
| Context | http, server, location, upstream |

Sets a timeout for name resolution, for example:

```
resolver_timeout 5s;
```

#### root

| Syntax  | root path;                             |
|---------|--|
| Default | <pre>root html;</pre>                  |
| Context | http, server, location, if in location |

Sets the root directory for requests. For example, with the following configuration

```
location /i/ {
  root /data/w3;
}
```

The /data/w3/i/top.gif file will be sent in response to the /i/top.gif request.

The path value can contain variables, except \$document\_root and \$realpath\_root.

A path to the file is constructed by merely adding a URI to the value of the root directive. If a URI has to be modified, the *alias* directive should be used.

# satisfy

```
Syntax satisfy all | any;
Default satisfy all;
Context http, server, location
```

Allows access if all (all) or at least one (any) of these modules allow access: Access, Auth Basic, or Auth Request.

```
location / {
  satisfy any;

allow 192.168.1.0/32;
  deny all;

auth_basic     "closed site";
  auth_basic_user_file conf/htpasswd;
}
```



# send\_lowat

| Syntax  | ${	t send\_lowat} \ size;$ |
|---------|----------------------------|
| Default | send_lowat 0;              |
| Context | http, server, location     |

If the directive is set to a non-zero value, Angie will try to minimize the number of send operations on client sockets by using either NOTE\_LOWAT flag of the ref:kqueue method or the SO\_SNDLOWAT socket option. In both cases the specified size is used.

# send timeout

| Syntax  | send_timeout $time;$   |
|---------|------------------------|
| Default | send_timeout 60s;      |
| Context | http, server, location |

Sets a timeout for transmitting a response to the client. The timeout is set only between two successive write operations, not for the transmission of the whole response. If the client does not receive anything within this time, the connection is closed.

#### sendfile

| Syntax  | sendfile on   off;                     |
|---------|--|
| Default | sendfile off;                          |
| Context | http, server, location, if in location |

Enables or disables the use of sendfile().

aio can be used to pre-load data for sendfile():

```
location /video/ {
  sendfile    on;
  tcp_nopush    on;
  aio    on;
}
```

In this configuration, sendfile() is called with the SF\_NODISKIO flag which causes it not to block on disk I/O, but, instead, report back that the data are not in memory. Angie then initiates an asynchronous data load by reading one byte. On the first read, the FreeBSD kernel loads the first 128K bytes of a file into memory, although next reads will only load data in 16K chunks. This can be changed using the  $read\_ahead$  directive.

# sendfile max chunk

| Syntax  | $sendfile_max\_chunk \ size;$     |
|---------|-----------------------------------|
| Default | <pre>sendfile_max_chunk 2m;</pre> |
| Context | http, server, location            |

Limits the amount of data that can be transferred in a single sendfile() call. Without the limit, one fast connection may seize the worker process entirely.



#### server

| Syntax  | server { } |
|---------|------------|
| Default | _          |
| Context | http       |

Sets configuration for a virtual server. There is no clear separation between IP-based (based on the IP address) and name-based (based on the "Host" request header field) virtual servers. Instead, the *listen* directives describe all addresses and ports that should accept connections for the server, and the *server name* directive lists all server names.

Example configurations are provided in the How Angie processes a request document.

# server\_name

```
Syntax server_name name ...;

Default server_name "";

Context server
```

Sets names of a virtual server, for example:

```
server {
   server_name example.com www.example.com;
}
```

The first name becomes the primary server name.

Server names can include an asterisk ("\*") replacing the first or last part of a name:

```
server {
   server_name example.com *.example.com www.example.*;
}
```

Such names are called wildcard names.

The first two of the names mentioned above can be combined in one:

```
server {
  server_name .example.com;
}
```

It is also possible to use regular expressions in server names, preceding the name with a tilde ("~"):

```
server {
   server_name ~^www\d+\.example\.com$ www.example.com;
}
```

Regular expressions can contain captures that can later be used in other directives:

```
server {
  server_name ~^(www\.)?(.+)$;

  location / {
    root /sites/$2;
  }
}
```



```
server_name _;
location / {
   root /sites/default;
}
```

Named captures in regular expressions create variables that can later be used in other directives:

```
server {
    server_name ~^(www\.)?(?<domain>.+)$;

    location / {
        root /sites/$domain;
    }
}

server {
    server_name _;
    location / {
        root /sites/default;
    }
}
```

### 1 Note

If the directive is set to \$hostname, the hostname of the web server is used.

You can also specify an empty server name (""):

```
server {
    server_name www.example.com "";
}
```

When searching for a virtual server by a name that is matched by multiple options (for example, both a wildcard and a regular expression), the first matching option will be selected in the following priority order:

- exact name;
- longest name with a wildcard at the beginning, such as \*.example.com;
- longest name with a wildcard at the end, such as mail.\*;
- the first matching regular expression (in the order of appearance), including an empty name.

#### **A** Attention

To make server\_name work with TLS, you need to terminate the TLS connection. The directive matches the Host in an HTTP request, so the handshake must be completed and the connection decrypted.



# server\_name\_in\_redirect

| Syntax  | server_name_in_redirect on   off; |
|---------|-----------------------------------|
| Default | server_name_in_redirect off;      |
| Context | http, server, location            |

Enables or disables the use of the primary server name, specified by the *server\_name* directive, in *absolute redirects* issued by Angie.

| on | the primary server name, specified by the server_name directive  |
|----|--|
|    | the name from the "Host" request header field is used. If this field is not present, the IP address of the server is used. |

The use of a port in redirects is controlled by the *port in redirect* directive.

# server names hash bucket size

| Syntax  | server_names_hash_bucket_size size;          |
|---------|--|
| Default | server_names_hash_bucket_size 32   64   128; |
| Context | http   |

Sets the bucket size for the server names hash tables. The default value depends on the size of the processor's cache line. The details of setting up hash tables are provided in a separate document.

# server\_names\_hash\_max\_size

| Syntax  | server_names_hash_max_size $size;$         |
|---------|--|
| Default | <pre>server_names_hash_max_size 512;</pre> |
| Context | $\operatorname{http}$                      |

Sets the maximum size of the server names hash tables. The details of setting up hash tables are provided in a separate document.

# server\_tokens

| Syntax  | server_tokens on   off   build   string; |
|---------|--|
| Default | $server_tokens on;$                      |
| Context | http, server, location                   |

Enables or disables emitting Angie version on error pages and in the Server response header field. The build parameter enables emitting the build name, set by the respective configure parameter, along with the version.

### Added in version 1.1.0: PRO

In Angie PRO, if the directive sets a string, which may also contain variables, the error pages and the Server response header field will use the string's variable-interpolated value instead of server name, version, and build name. An empty string disables emitting the Server field.



### status\_zone

| Syntax  | $status\_zone off   zone   key zone=zone[:count];$ |
|---------|--|
| Default | _  |
| Context | server, location, if in location                   |

Allocates a shared memory zone to collect metrics for  $/status/http/location\_zones/<zone>$  and  $/status/http/server\_zones/<zone>$ .

Multiple server contexts can share the same zone for data collection; the special value off disables data collection in nested location blocks.

The single-value zone syntax aggregates all metrics for its context in the same shared memory zone:

```
server {
    listen 80;
    server_name *.example.com;
    status_zone single;
    # ...
}
```

The alternative syntax uses the following parameters:

| key              | A string with variables, whose value determines the grouping of requests in the zone. All requests producing identical values after substitution are grouped together. If substitution yields an empty value, metrics aren't updated. |
|------------------|---|
| zone             | The name of the shared memory zone.   |
| count (optional) | The maximum number of separate groups for collecting metrics. If new key values would exceed this limit, they are grouped under zone instead. The default value is 1.   |

In the following example, all requests sharing the same \$host value are grouped into the host\_zone. Metrics are tracked separately for each unique \$host until there are 10 metric groups. Once this limit is reached, any additional \$host values are included under the host\_zone:

```
server {
    listen 80;
    server_name *.example.com;

    status_zone $host zone=host_zone:10;

    location / {
        proxy_pass http://example.com;
    }
}
```

The resulting metrics are thus split between individual hosts in the API output.

# $subrequest\_output\_buffer\_size$

| Syntax  | $subrequest\_output\_buffer\_size \ size;$        |
|---------|---|
| Default | <pre>subrequest_output_buffer_size 4k   8k;</pre> |
| Context | http, server, location                            |



Sets the size of the buffer used for storing the response body of a subrequest. By default, the buffer size is equal to one memory page. This is either 4K or 8K, depending on a platform. It can be made smaller, however.

# 1 Note

The directive is applicable only for subrequests with response bodies saved into memory. For example, such subrequests are created by SSI.

# tcp nodelay

| Syntax  | tcp_nodelay on   off;  |
|---------|------------------------|
| Default | tcp_nodelay on;        |
| Context | http, server, location |

Enables or disables the use of the TCP\_NODELAY option. The option is enabled when a connection is transitioned into the keep-alive state. Additionally, it is enabled on SSL connections, for unbuffered proxying, and for WebSocket proxying.

### tcp nopush

| Syntax  | tcp_nopush on   off;   |
|---------|------------------------|
| Default | tcp_nopush off;        |
| Context | http, server, location |

Enables or disables the use of the TCP\_NOPUSH socket option on FreeBSD or the TCP\_CORK socket option on Linux. The options are enabled only when sendfile is used. Enabling the option allows

- sending the response header and the beginning of a file in one packet, on Linux and FreeBSD 4.\*;
- sending a file in full packets.

# try\_files

| Syntax  | try_files file uri;<br>try_files file = code; |
|---------|---|
| Default | _   |
| Context | server, location                              |

Checks the existence of files in the specified order and uses the first found file for request processing; the processing is performed in the current context. The path to a file is constructed from the file parameter according to the root and alias directives. It is possible to check directory's existence by specifying a slash at the end of a name, e.g. \$uri/. If none of the files were found, an internal redirect to the uri specified in the last parameter is made. For example:

```
location /images/ {
 try_files $uri /images/default.gif;
}
location = /images/default.gif {
  expires 30s;
```



The last parameter can also point to a named location, as shown in examples below. The last parameter can also be a code:

```
location / {
  try_files $uri $uri/index.html $uri.html =404;
}
```

In the following example,

```
location / {
  try_files $uri $uri/ @drupal;
}
```

the try\_files directive is equivalent to

```
location / {
  error_page 404 = @drupal;
  log_not_found off;
}
```

And here,

```
location ~ \.php$ {
  try_files $uri @drupal;

  fastcgi_pass ...;

  fastcgi_param SCRIPT_FILENAME /path/to$fastcgi_script_name;
# ...
}
```

try\_files checks the existence of the PHP file before passing the request to the FastCGI server.



```
location @drupal {
  fastcgi_pass ...;

fastcgi_param SCRIPT_FILENAME /path/to/index.php;
  fastcgi_param SCRIPT_NAME /index.php;
  fastcgi_param QUERY_STRING q=$uri&$args;

# ... other fastcgi_param
}
```

```
location / {
   try_files $uri $uri/ @wordpress;
}
location ~ \.php$ {
   try_files $uri @wordpress;
   fastcgi_pass ...;

   fastcgi_param SCRIPT_FILENAME /path/to$fastcgi_script_name;
# ... other fastcgi_param
}
location @wordpress {
   fastcgi_pass ...;

   fastcgi_param SCRIPT_FILENAME /path/to/index.php;
# ... other fastcgi_param
}
```

### types

```
Syntax types { ... }
Default types text/html html; image/gif gif; image/jpeg jpg;
Context http, server, location
```

Maps file name extensions to MIME types of responses. Extensions are case-insensitive. Several extensions can be mapped to one type, for example:

```
types {
   application/octet-stream bin exe dll;
   application/octet-stream deb;
   application/octet-stream dmg;
}
```

A sufficiently full mapping table is distributed with Angie in the conf/mime.types file.

To make a particular location emit the "application/octet-stream" MIME type for all requests, the following configuration can be used:



# types\_hash\_bucket\_size

| Syntax  | $types_hash_bucket_size \ size;$      |
|---------|---------------------------------------|
| Default | <pre>types_hash_bucket_size 64;</pre> |
| Context | http, server, location                |

Sets the bucket size for the types hash tables. The details of setting up hash tables are provided in a separate document.

# types hash max size

| Syntax  | types_hash_max_size $size;$          |
|---------|--------------------------------------|
| Default | <pre>types_hash_max_size 1024;</pre> |
| Context | http, server, location               |

Sets the maximum size of the types hash tables. The details of setting up hash tables are provided in a separate document.

# underscores in headers

| Syntax  | underscores_in_headers on   off; |
|---------|----------------------------------|
| Default | underscores_in_headers off;      |
| Context | http, server                     |

Enables or disables the use of underscores in client request header fields. When the use of underscores is disabled, request header fields whose names contain underscores are marked as invalid and become subject to the <code>ignore\_invalid\_headers</code> directive.

If the directive is specified on the server level, the value from the default server can be used.

# variables\_hash\_bucket\_size

| Syntax  | variables_hash_bucket_size $size;$        |
|---------|---|
| Default | <pre>variables_hash_bucket_size 64;</pre> |
| Context | http                                      |

Sets the bucket size for the variables hash table. The details of setting up hash tables are provided in a separate document.

# variables hash max size

| Syntax  | ${\tt variables\_hash\_max\_size}\ size;$ |
|---------|---|
| Default | <pre>variables_hash_max_size 1024;</pre>  |
| Context | http                                      |

Sets the maximum size of the variables hash table. The details of setting up hash tables are provided in a separate document.



#### **Built-in Variables**

The http\_core module supports built-in variables with names matching the Apache Server variables. First of all, these are variables representing client request header fields, such as \$http\_user\_agent, \$http\_cookie, and so on. Also, there are other variables:

# \$angie\_version

Angie version

### \$arg\_<name>

argument with the specified name in the request line

### \$args

arguments in the request line

### \$binary\_remote\_addr

client address in a binary form, value's length is always 4 bytes for IPv4 addresses or 16 bytes for IPv6 addresses

#### \$body\_bytes\_sent

number of bytes sent to the client, not counting the response header; this variable is compatible with the "%B" parameter of the mod\_log\_config Apache module

### \$bytes\_sent

number of bytes sent to a client

### \$connection

connection serial number

### \$connection\_requests

current number of requests made through a connection

#### \$connection\_time

connection time in seconds with a milliseconds resolution

### \$content\_length

"Content-Length" request header field

# \$content\_type

"Content-Type" request header field

# \$cookie\_<name>

cookie with the specified name



#### \$document\_root

root or alias directive's value for the current request

### \$document\_uri

same as \$uri

#### \$host

in this order of precedence: host name from the request line, or host name from the "Host" request header field, or the server name matching a request

#### \$hostname

host name

#### \$http\_<name>

arbitrary request header field; the name is the field name converted to lower case with dashes replaced by underscores

#### \$https

on if connection operates in SSL mode, or an empty string otherwise

#### \$is\_args

?, if a request line has arguments, or an empty string otherwise

#### \$limit\_rate

setting this variable enables response rate limiting; see *limit\_rate* 

### \$msec

current time in seconds with the milliseconds resolution

#### \$pid

PID of the worker process

# \$pipe

p if request was pipelined, . otherwise

# \$proxy\_protocol\_addr

Client address from the PROXY protocol header.

The PROXY protocol must be previously enabled by setting the proxy\_protocol parameter in the *listen* directive.

### \$proxy\_protocol\_port

Client port from the PROXY protocol header.

The PROXY protocol must be previously enabled by setting the proxy\_protocol parameter in the *listen* directive.



### \$proxy\_protocol\_server\_addr

Server address from the PROXY protocol header.

The PROXY protocol must be previously enabled by setting the proxy\_protocol parameter in the *listen* directive.

#### \$proxy\_protocol\_server\_port

Server port from the PROXY protocol header.

The PROXY protocol must be previously enabled by setting the proxy\_protocol parameter in the *listen* directive.

### \$proxy\_protocol\_tlv\_<name>

TLV from the PROXY Protocol header. The *name* can be a TLV type or its numeric value. In the latter case, the value is hexadecimal and should be prefixed with 0x:

```
$proxy_protocol_tlv_alpn
$proxy_protocol_tlv_0x01
```

SSL TLVs can also be accessed by TLV type name or its numeric value, both prefixed by ssl\_:

```
$proxy_protocol_tlv_ssl_version
$proxy_protocol_tlv_ssl_0x21
```

The following TLV type names are supported:

- alpn (0x01) upper layer protocol used over the connection
- authority (0x02) host name value passed by the client
- unique\_id (0x05) unique connection id
- netns (0x30) name of the namespace
- ssl (0x20) binary SSL TLV structure

The following SSL TLV type names are supported:

- ssl\_version (0x21) SSL version used in client connection
- ssl\_cn (0x22) SSL certificate Common Name
- ssl\_cipher (0x23) name of the used cipher
- $\bullet$  ssl\_sig\_alg (0x24) algorithm used to sign the certificate
- ssl\_key\_alg (0x25) public-key algorithm

Also, the following special SSL TLV type name is supported:

• ssl\_verify - client SSL certificate verification result, 0 if the client presented a certificate and it was successfully verified, non-zero otherwise.

The PROXY protocol must be previously enabled by setting the proxy\_protocol parameter in the *listen* directive.

\$query\_string

same as \$args



#### \$realpath\_root

an absolute pathname corresponding to the *root* or *alias* directive's value for the current request, with all symbolic links resolved to real paths

#### \$remote\_addr

client address

### \$remote\_port

client port

# \$remote\_user

user name supplied with the Basic authentication

### \$request

full original request line

### \$request\_body

Request body.

The variable's value is made available in locations processed by the proxy\_pass, fastcgi\_pass, uwsgi\_pass and scqi\_pass directives when the request body was read to a memory buffer.

# \$request\_body\_file

Name of a temporary file with the request body.

At the end of processing, the file needs to be removed. To always write the request body to a file,  $client\_body\_in\_file\_only$  needs to be enabled. When the name of a temporary file is passed in a proxied request or in a request to a FastCGI/uwsgi/SCGI server, passing the request body should be disabled by the  $proxy\_pass\_request\_body$  off,  $fastcgi\_pass\_request\_body$  off,  $fastcgi\_pass\_request\_body$  off or  $scgi\_pass\_request\_body$  off directives, respectively.

# \$request\_completion

"OK" if a request has completed, or an empty string otherwise

### \$request\_filename

file path for the current request, based on the root or alias directives, and the request URI

# \$request\_id

unique request identifier generated from 16 random bytes, in hexadecimal

#### \$request\_length

request length (including request line, header, and request body)

# \$request\_method

request method, usually GET or POST



#### \$request\_time

request processing time in seconds with a milliseconds resolution; time elapsed since the first bytes were read from the client

# \$request\_uri

full original request URI (with arguments)

#### \$scheme

request scheme, "http" or "https"

# \$sent\_http\_<name>

arbitrary response header field; the name is the field name converted to lower case with dashes replaced by underscores

#### \$sent\_trailer\_<name>

arbitrary field sent at the end of the response; the name is the field name converted to lower case with dashes replaced by underscores

### \$server\_addr

The address of the server which accepted a request. Computing the variable's value usually requires one system call. To avoid it, the *listen* directives must specify addresses and use the bind parameter.

### \$server\_name

name of the server which accepted a request

# \$server\_port

port of the server which accepted a request

# \$server\_protocol

request protocol, usually "HTTP/1.0", "HTTP/1.1", or "HTTP/2.0"

#### \$status

response status

### \$time\_iso8601

local time in the ISO 8601 standard format

### \$time\_local

local time in the Common Log Format

# \$tcpinfo\_rtt, \$tcpinfo\_rttvar, \$tcpinfo\_snd\_cwnd, \$tcpinfo\_rcv\_space

information about the client TCP connection; available on systems that support the  $TCP\_INFO$  socket option



#### \$uri

Current URI in the request, *normalized*. The value of \$uri may change during request processing, e.g. when doing internal redirects, or when using index files

# 3.3.3 Stream Module

#### Access

The module allows limiting access to certain client addresses.

# **Configuration Example**

```
server {
    ...
    deny 192.168.1.1;
    allow 192.168.1.0/24;
    allow 10.1.1.0/16;
    allow 2001:0db8::/32;
    deny all;
}
```

The rules are checked in sequence until the first match is found. In this example, access is allowed only for IPv4 networks 10.1.1.0/16 and 192.168.1.0/24 excluding the address 192.168.1.1, and for IPv6 network 2001:0db8::/32.

# **Directives**

#### allow

| Syntax  | allow $address \mid CIDR \mid$ unix:   all; |
|---------|---|
| Default | _   |
| Context | http, server, location, limit_except        |

Allows access for the specified network or address. If the special value unix: is specified (1.5.1), allows access for all UNIX domain sockets.

# deny

| Syntax  | $	ext{deny } address \mid CIDR \mid 	ext{unix:} \mid 	ext{all;}$ |
|---------|--|
| Default | _  |
| Context | http, server, location, limit_except                             |

Denies access for the specified network or address. If the special value unix: is specified (1.5.1), denies access for all UNIX domain sockets.

### Geo

The module creates variables with values depending on the client IP address.

# **Configuration Example**

```
geo $geo {
    default 0;
    127.0.0.1 2;
```



```
192.168.1.0/24 1;

10.1.0.0/16 1;

::1 2;

2001:0db8::/32 1;

}
```

#### **Directives**

#### geo

| Syntax  | geo [\$address] \$variable { } |
|---------|--------------------------------|
| Default | _                              |
| Context | stream                         |

Describes the dependency of values of the specified variable on the client IP address. By default, the address is taken from the  $\$remote\_addr$  variable, but it can also be taken from another variable, for example:

```
geo $arg_remote_addr $geo {
    ...;
}
```

# 1 Note

Since variables are evaluated only when used, the mere existence of even a large number of declared geo variables does not cause any extra costs for connection processing.

If the value of a variable does not represent a valid IP address then the "255.255.255.255" address is used.

Addresses are specified either as prefixes in CIDR notation (including individual addresses) or as ranges.

The following special parameters are also supported:

| delete  | deletes the specified network  |
|---------|--|
| default | the value set to the variable if the client address does not match any of the specified addresses. When addresses are specified in CIDR notation, "0.0.0.0/0" and ":/0" can be used instead of default. When default is not specified, the default value will be an empty string |
| include | includes a file with addresses and values. There can be several inclusions.  |
| ranges  | indicates that addresses are specified as ranges. This parameter should be the first. To speed up loading of a geo base, addresses should be put in ascending order.   |

# Example:



```
127.0.0.1/32 RU;
10.1.0.0/16 RU;
192.168.1.0/24 UK;
}
```

The conf/geo.conf file could contain the following lines:

```
10.2.0.0/16 RU;
192.168.2.0/24 RU;
```

A value of the most specific match is used. For example, for the 127.0.0.1 address the value RU will be chosen, not US.

Example with ranges:

#### **GeoIP**

Creates variables with values depending on the client IP address, using the precompiled MaxMind databases

When using the databases with IPv6 support, IPv4 addresses are looked up as IPv4-mapped IPv6 addresses.

When building from the source code, this module isn't built by default; it should be enabled with the --with-stream\_geoip\_module build option.

```
Important
This module requires the MaxMind GeoIP library.
```

### **Configuration Example**

```
stream {
    geoip_country
                          GeoIP.dat;
                          GeoLiteCity.dat;
    geoip_city
    map $geoip_city_continent_code $nearest_server {
        default
                       example.com;
        EII
                    eu.example.com;
        NA
                    na.example.com;
                    as.example.com;
        AS
    }
}
```



#### **Directives**

# geoip\_country

| Syntax  | ${	t geoip\_country} \ file;$ |
|---------|-------------------------------|
| Default | _                             |
| Context | stream                        |

Specifies a database used to determine the country depending on the client IP address. The following variables are available when using this database:

```
$geoip_country_c two-letter country code, for example, "RU", "US".
$geoip_country_c three-letter country code, for example, "RUS", "USA".
$geoip_country_n country name, for example, "Russian Federation", "United States".
```

# geoip\_city

| Syntax  | geoip_city file; |
|---------|------------------|
| Default | _                |
| Context | stream           |

Specifies a database used to determine the country, region, and city depending on the client IP address. The following variables are available when using this database:

| <pre>\$geoip_city_cont</pre> | two-letter continent code, for example, "EU", "NA".  |
|------------------------------|--|
| \$geoip_city_coun            | two-letter country code, for example, "RU", "US".  |
| \$geoip_city_coun            | three-letter country code, for example, "RUS", "USA".  |
| \$geoip_city_coun            | country name, for example, "Russian Federation", "United States".  |
| <pre>\$geoip_dma_code</pre>  | DMA region code in US (also known as "metro code"), according to the geotargeting in Google AdWords API.                                 |
| \$geoip_latitude             | latitude.  |
| <pre>\$geoip_longitude</pre> | longitude.   |
| <pre>\$geoip_region</pre>    | two-symbol country region code (region, territory, state, province, federal land and the like), for example, "48", "DC".                 |
| \$geoip_region_na            | country region name (region, territory, state, province, federal land and the like), for example, "Moscow City", "District of Columbia". |
| \$geoip_city                 | city name, for example, "Moscow", "Washington".  |
| \$geoip_postal_co            | postal code.   |

# geoip\_org

| Syntax  | ${	t geoip\_org} \; file;$ |
|---------|----------------------------|
| Default | _                          |
| Context | stream                     |

Specifies a database used to determine the organization depending on the client IP address. The following variable is available when using this database:

| <pre>\$geoip_org</pre> | organization name, for example, "The University of Melbourne". |  |
|------------------------|--|--|
|------------------------|--|--|



#### JS

The module is used to implement handlers in njs — a subset of the JavaScript language.

In our repositories, the module is built dynamically and is available as a separate package named angie-module-njs or angie-pro-module-njs.

# **Configuration Example**

```
stream {
    js_import stream.js;
    js_set $bar stream.bar;
    js_set $req_line stream.req_line;
    server {
        listen 12345;
        js_preread stream.preread;
                 $req_line;
        return
    }
    server {
        listen 12346;
        js_access stream.access;
        proxy_pass 127.0.0.1:8000;
        js_filter stream.header_inject;
}
http {
    server {
        listen 8000;
        location / {
            return 200 $http_foo\n;
    }
}
```

The stream.js file:

```
var line = '';
function bar(s) {
    var v = s.variables;
    s.log("hello from bar() handler!");
    return "bar-var" + v.remote_port + "; pid=" + v.pid;
}

function preread(s) {
    s.on('upload', function (data, flags) {
       var n = data.indexOf('\n');
       if (n != -1) {
            line = data.substr(0, n);
            s.done();
       }
       });
});
```



```
function req_line(s) {
   return line;
}
// Read HTTP request line.
// Collect bytes in 'req' until
// request line is read.
// Injects HTTP header into a client's request
var my_header = 'Foo: foo';
function header_inject(s) {
   var req = '';
   s.on('upload', function(data, flags) {
        req += data;
        var n = req.search('\n');
        if (n != -1) {
            var rest = req.substr(n + 1);
            req = req.substr(0, n + 1);
            s.send(req + my_header + '\r\n' + rest, flags);
            s.off('upload');
        }
   });
}
function access(s) {
    if (s.remoteAddress.match('^192.*')) {
        s.deny();
        return;
   }
    s.allow();
}
export default {bar, preread, req_line, header_inject, access};
```

### **Directives**

### js access

| Syntax  | <pre>js_access function   module.function;</pre> |
|---------|--|
| Default | _  |
| Context | stream, server                                   |

Sets an njs function which will be called at the access phase. Module functions can be referenced.

The function is called once at the moment when the stream session reaches the *access phase* for the first time. The function is called with the following arguments:

```
s the stream session object
```

At this phase, it is possible to perform initialization or register a callback with the s.on() method for each incoming data chunk until one of the following methods are called: s.done(), s.decline(), s.allow(). As soon as one of these methods is called, the stream session processing switches to the *next phase* and all current s.on() callbacks are dropped.



## js\_fetch\_buffer\_size

| Syntax  | <pre>js_fetch_buffer_size size;</pre> |
|---------|---------------------------------------|
| Default | <pre>js_fetch_buffer_size 16k;</pre>  |
| Context | stream, server                        |

Sets the size of the buffer used for reading and writing with Fetch API.

## js\_fetch\_ciphers

| Syntax  | <pre>js_fetch_ciphers ciphers;</pre>          |
|---------|---|
| Default | <pre>js_fetch_ciphers HIGH:!aNULL:!MD5;</pre> |
| Context | stream, server                                |

Specifies the enabled ciphers for HTTPS connections with Fetch API. The ciphers are specified in the format understood by the OpenSSL library.

The list of ciphers depends on the version of OpenSSL installed. The full list can be viewed using the openssl ciphers command.

## js\_fetch\_max\_response\_buffer\_size

| Syntax  | ${\tt js\_fetch\_max\_response\_buffer\_size}$ $size;$ |
|---------|--|
| Default | <pre>js_fetch_max_response_buffer_size 1m;</pre>       |
| Context | stream, server   |

Sets the maximum size of the response received with Fetch API.

## js fetch protocols

| Syntax  | js_fetch_protocols [TLSv1] [TLSv1.1] [TLSv1.2] [TLSv1.3]; |
|---------|---|
| Default | js_fetch_protocols TLSv1 TLSv1.1 TLSv1.2;                 |
| Context | stream, server  |

Enables the specified protocols for HTTPS connections with Fetch API.

## js fetch timeout

| Syntax  | <pre>js_fetch_timeout time;</pre> |
|---------|-----------------------------------|
| Default | <pre>js_fetch_timeout 60s;</pre>  |
| Context | stream, server                    |

Defines a timeout for reading and writing for Fetch API. The timeout is set only between two successive read/write operations, not for the whole response. If no data is transmitted within this time, the connection is closed.

## $js\_fetch\_trusted\_certificate$

| Syntax  | ${\tt js\_fetch\_trusted\_certificate}$ $file;$ |
|---------|---|
| Default | _   |
| Context | stream, server                                  |



Specifies a file with trusted CA certificates in the PEM format used to verify the HTTPS certificate with Fetch API.

## js fetch verify

| Syntax  | <pre>js_fetch_verify on   off;</pre> |
|---------|--------------------------------------|
| Default | <pre>js_fetch_verify on;</pre>       |
| Context | stream, server                       |

Enables or disables verification of the HTTPS server certificate with Fetch API.

## js\_fetch\_verify\_depth

| Syntax  | <pre>js_fetch_verify_depth number;</pre> |
|---------|--|
| Default | <pre>js_fetch_verify_depth 100;</pre>    |
| Context | stream, server                           |

Sets the verification depth in the HTTPS server certificates chain with Fetch API.

## js filter

| Syntax  | <pre>js_filter function   module.function;</pre> |
|---------|--|
| Default | _  |
| Context | stream, server                                   |

Sets a data filter. Module functions can be referenced.

The filter function is called once at the moment when the stream session reaches the *content phase*. The filter function is called with the following arguments:

```
s the stream session object
```

At this phase, it is possible to perform initialization or register a callback with the s.on() method for each incoming data chunk. The s.off() method may be used to unregister a callback and stop filtering.

## 1 Note

As the  $js\_filter$  handler returns its result immediately, it supports only synchronous operations. Thus, asynchronous operations such as ngx.fetch() or setTimeout() are not supported.

### js\_import

| Syntax  | <pre>js_import module.js   export_name from module.js;</pre> |
|---------|--|
| Default | _  |
| Context | stream, server   |

Imports a module that implements location and variable handlers in njs. The <code>export\_name</code> is used as a namespace to access module functions. If the <code>export\_name</code> is not specified, the module name will be used as a namespace.



# js\_import stream.js;

Here, the module name stream is used as

Several js import directives can be specified.

## js path

| Syntax  | <pre>js_path path;</pre> |
|---------|--------------------------|
| Default | _                        |
| Context | stream, server           |

Sets an additional path for njs modules.

## js preload object

| Syntax  | <pre>js_preload_object name.json   name from file.json;</pre> |
|---------|---|
| Default | _   |
| Context | stream, server  |

Preloads an immutable object at configure time. The *name* is used as a name of the global variable though which the object is available in njs code. If the *name* is not specified, the file name will be used instead.

```
js_preload_object map.json;
```

Here, the map is used as a name while accessing the preloaded object.

Several js\_preload\_object directives can be specified.

## js preread

| Syntax  | <pre>js_preread function   module.function;</pre> |
|---------|---|
| Default | _   |
| Context | stream, server                                    |

Sets an njs function which will be called at the preread phase. Module functions can be referenced.

The function is called once at the moment when the stream session reaches the *preread phase* for the first time. The function is called with the following arguments:

```
s the stream session object
```

At this phase, it is possible to perform initialization or register a callback with the s.on() method for each incoming data chunk until one of the following methods are called: s.done(), s.decline(), s.allow(). When one of these methods is called, the stream session switches to the *next phase* and all current s.on() callbacks are dropped.

### 1 Note

As the  $js\_preread$  handler returns its result immediately, it supports only synchronous callbacks. Thus, asynchronous callbacks such as ngx.fetch() or setTimeout() are not supported. Nevertheless, asynchronous operations are supported in s.on() callbacks in the  $preread\ phase$ .



## js\_set

| Syntax  | ${\tt js\_set}$ \$variable function   module.function; |
|---------|--|
| Default | _  |
| Context | stream, server   |

Sets an njs function for the specified variable. Module functions can be referenced.

The function is called when the variable is referenced for the first time for a given request. The exact moment depends on a *phase* at which the variable is referenced. This can be used to perform some logic not related to variable evaluation. For example, if the variable is referenced only in the *log\_format* directive, its handler will not be executed until the log phase. This handler can be used to do some cleanup right before the request is freed.

## 1 Note

As the  $js\_set$  handler returns its result immediately, it supports only synchronous callbacks. Thus, asynchronous callbacks such as ngx.fetch() or setTimeout() are not supported.

## js\_shared\_dict\_zone

| Syntax  | <pre>js_shared_dict_zone [evict];</pre> | zone=name:size | $[\mathtt{timeout} \texttt{=} time]$ | [type=string | number] |
|---------|---|----------------|--------------------------------------|--------------|---------|
| Default | _                                       |                |                                      |              |         |
| Context | stream                                  |                |                                      |              |         |

Sets the name and size of the shared memory zone that keeps the key-value dictionary shared between worker processes.

| type    | optional parameter, allows redefining the value type to number, by default the shared dictionary uses a string as a key and a value |
|---------|---|
| timeout | optional parameter, sets the time after which all shared dictionary entries are removed from the zone                               |
| evict   | optional parameter, removes the oldest key-value pair when the zone storage is exhausted  |

## Examples:

```
example.conf:
    # Creates a 1Mb dictionary with string values,
    # removes key-value pairs after 60 seconds of inactivity:
    js_shared_dict_zone zone=foo:1M timeout=60s;

# Creates a 512Kb dictionary with string values,
    # forcibly removes oldest key-value pairs when the zone is exhausted:
    js_shared_dict_zone zone=bar:512K timeout=30s evict;

# Creates a 32Kb permanent dictionary with number values:
    js_shared_dict_zone zone=num:32k type=number;
```

```
example.js:
   function get(r) {
     r.return(200, ngx.shared.foo.get(r.args.key));
```



```
function set(r) {
    r.return(200, ngx.shared.foo.set(r.args.key, r.args.value));
}

function delete(r) {
    r.return(200, ngx.shared.bar.delete(r.args.key));
}

function increment(r) {
    r.return(200, ngx.shared.num.incr(r.args.key, 2));
}
```

## js\_var

| Syntax  | js_var \$variable [value]; |
|---------|----------------------------|
| Default | _                          |
| Context | stream, server             |

Declares a writable variable. The value can contain text, variables, and their combination.

## **Session Object Properties**

Each stream njs handler receives one argument, a stream-session object.

## **Limit Conn**

The module is used to limit the number of connections per the defined key, in particular, the number of connections from a single IP address.

### **Configuration Example**

### **Directives**

## limit\_conn

| Syntax  | limit_conn zone number; |
|---------|-------------------------|
| Default | _                       |
| Context | stream, server          |



Sets the shared memory zone and the maximum allowed number of connections for a given key value. When this limit is exceeded, the server will close the connection. For example, the directives

```
limit_conn_zone $binary_remote_addr zone=addr:10m;
server {
    ...
    limit_conn addr 1;
}
```

allow only one connection per IP address at a time.

When several limit\_conn directives are specified, any configured limit will apply.

These directives are inherited from the previous configuration level if and only if there are no limit\_conn directives defined on the current level.

## limit conn dry run

| Syntax  | limit_conn_dry_run on   off;       |
|---------|------------------------------------|
| Default | <pre>limit_conn_dry_run off;</pre> |
| Context | stream, server                     |

Enables the dry run mode. In this mode, the number of connections is not limited, however, in the *shared memory zone*, the number of excessive connections is accounted as usual.

## limit\_conn\_log\_level

| Syntax  | limit_conn_log_level info   notice   warn   error; |
|---------|--|
| Default | <pre>limit_conn_log_level error;</pre>             |
| Context | stream, server                                     |

Sets the desired logging level for cases when the server limits the number of connections.

#### limit conn zone

| Syntax  | $limit\_conn\_zone \ key \ zone = name:size;$ |
|---------|---|
| Default | _   |
| Context | stream  |

Sets parameters for a shared memory zone that will keep states for various keys. In particular, the state includes the current number of connections. The key can contain text, variables, and their combinations. Connections with an empty key value are not accounted.

Usage example:

```
limit_conn_zone $binary_remote_addr zone=addr:10m;
```

Here, a client IP address is set by the \$binary\_remote\_addr variable.

The size of \$binary\_remote\_addr is 4 bytes for IPv4 addresses or 16 bytes for IPv6 addresses. The stored state always occupies 32 or 64 bytes on 32-bit platforms and 64 bytes on 64-bit platforms.

One megabyte zone can keep about 32 thousand 32-byte states or about 16 thousand 64-byte states. If the zone storage is exhausted, the server will close the connection.



#### **Built-in Variables**

\$limit\_conn\_status

keeps the result of limiting the number of connections: PASSED, REJECTED or REJECTED\_DRY\_RUN

#### Log

The module writes request logs in the specified format.

### **Configuration Example**

#### **Directives**

## access log

| Syntax  | <pre>access_log path [format [buffer=size] [gzip[=level]] [flush=time] [if=condition]]; access_log off;</pre> |
|---------|---|
| Default | access_log off;   |
| Context | stream, server  |

Sets the path, format, and configuration for a buffered log write. Several logs can be specified on the same configuration level. Logging to *syslog* can be configured by specifying the "syslog:" prefix in the first parameter. The special value off cancels all access\_log directives on the current level.

If either the buffer or gzip parameter is used, writes to log will be buffered.

#### \* Caution

The buffer size must not exceed the size of an atomic write to a disk file. For FreeBSD this size is unlimited.

When buffering is enabled, the data will be written to the file:

- if the next log line does not fit into the buffer;
- if the buffered data is older than specified by the flush parameter;
- $\bullet$  when a worker process is *re-opening log files* or is shutting down.

If the gzip parameter is used, then the buffered data will be compressed before writing to the file. The compression level can be set between 1 (fastest, less compression) and 9 (slowest, best compression). By default, the buffer size is equal to 64K bytes, and the compression level is set to 1. Since the data is compressed in atomic blocks, the log file can be decompressed or read by "zcat" at any time.

Example:

```
access_log /path/to/log.gz basic gzip flush=5m;
```

## Important

For gzip compression to work, Angie must be built with the zlib library.



The file path can contain variables, but such logs have some constraints:

- the *user* whose credentials are used by worker processes should have permissions to create files in a directory with such logs;
- buffered writes do not work;
- the file is opened and closed for each log write. However, since the descriptors of frequently used files can be stored in a cache, writing to the old file can continue during the time specified by the open log file cache directive's valid parameter.

The *if* parameter enables conditional logging. A session will not be logged if the condition evaluates to "0" or an empty string.

## log format

| Syntax  | log_format name [escape=default   json   none] string; |
|---------|--|
| Default | _  |
| Context | stream, server   |

### Specifies log format.

The escape parameter allows setting json or default characters escaping in variables, by default, default escaping is used. The none value disables escaping.

For default escaping, characters """, "\", and other characters with values less than 32 or above 126 are escaped as "\xXX". If the variable value is not found, a hyphen "-" will be logged.

For json escaping, all characters not allowed in JSON strings will be escaped: characters """ and "\" are escaped as "\"" and "\\", characters with values less than 32 are escaped as "\n", "\r", "\t", "\b", "\f", or "\u00XX".

### open log file cache

| Syntax  | $	exttt{open_log_file_cache } \max = N \; 	exttt{[inactive=} time] \; 	exttt{[min_uses=} N] \; 	exttt{[valid=} time];$ |
|---------|--|
| Default | open_log_file_cache off;   |
| Context | stream, server   |

Defines a cache that stores the file descriptors of frequently used logs whose names contain variables. The directive has the following parameters:

| max      | sets the maximum number of descriptors in a cache; if the cache becomes full the least recently used (LRU) descriptors are closed                |
|----------|--|
| inactive | sets the time after which the cached descriptor is closed if there were no access during this time; by default, 10 seconds                       |
| min_uses | sets the minimum number of file uses during the time defined by the inactive parameter to let the descriptor stay open in a cache; by default, 1 |
| valid    | sets the time after which it should be checked that the file still exists with the same name; by default, 60 seconds                             |
| off      | disables caching   |

### Usage example:

```
open_log_file_cache max=1000 inactive=20s valid=1m min_uses=2;
```



#### Map

The module creates variables whose values depend on values of other variables.

#### **Configuration Example**

```
map $remote_addr $limit {
    127.0.0.1 "";
    default $binary_remote_addr;
}
limit_conn_zone $limit zone=addr:10m;
limit_conn addr 1;
```

#### **Directives**

#### map

```
Syntax map string $variable { ... };
Default —
Context stream
```

Creates a new variable. Its value depends on the first parameter, specified as a string with variables, for example:

```
set $var1 "foo";
set $var2 "bar";

map $var1$var2 $new_variable {
    default "foobar_value";
}
```

Here, the variable **\$new\_variable** will have a value composed of the two variables **\$var1** and **\$var2**, or a default value if these variables are not defined.

## 1 Note

Since variables are evaluated only when they are used, the mere declaration even of a large number of "map" variables does not add any extra costs to request processing.

Parameters inside the *map* block specify a mapping between source and resulting values.

Source values are specified as strings or regular expressions.

Strings are matched ignoring the case.

A regular expression should either start with a "~" symbol for a case-sensitive matching, or with the "~\*" symbols for case-insensitive matching. A regular expression can contain named and positional captures that can later be used in other directives along with the resulting variable.

If a source value matches one of the names of special parameters described below, it should be prefixed with the "" symbol.

The resulting value can contain text, variable and their combination.

The following special parameters are also supported:



| default $value$ | sets the resulting value if the source value matches none of the specified variants. When <i>default</i> is not specified, the default resulting value will be an empty string. |
|-----------------|---|
| hostnames       | indicates that source values can be hostnames with a prefix or suffix mask. This parameter should be specified before the list of values.                                       |

## For example,

```
*.example.com 1;
example.* 1;
```

### The following two records

```
example.com 1;
*.example.com 1;
```

### can be combined:

```
.example.com 1;
```

| $\verb"include" \mathit{file}$ | includes a file with values. There can be several inclusions. |
|--------------------------------|---|
| volatile                       | indicates that the variable is not cacheable.                 |

If the source value matches more than one of the specified variants, e.g. both a mask and a regular expression match, the first matching variant will be chosen, in the following order of priority:

- 1. String value without a mask
- 2. Longest string value with a prefix mask, e.g. \*.example.com
- 3. Longest string value with a suffix mask, e.g. mail.\*
- 4. First matching regular expression (in order of appearance in a configuration file)
- 5. Default value (default)

## map\_hash\_bucket\_size

| Syntax  | map_hash_bucket_size size;                 |
|---------|--|
| Default | <pre>map_hash_bucket_size 32 64 128;</pre> |
| Context | stream                                     |

Sets the bucket size for the map variables hash tables. Default value depends on the processor's cache line size. The details of setting up hash tables are provided separately.

### map\_hash\_max\_size

| Syntax  | map_hash_max_size $size;$          |
|---------|------------------------------------|
| Default | <pre>map_hash_max_size 2048;</pre> |
| Context | stream                             |

Sets the maximum size of the map variables hash tables. The details of setting up hash tables are provided separately.



#### **MQTT** Preread

Enables extracting client IDs and usernames from CONNECT packets for Message Queuing Telemetry Transport (MQTT) versions 3.1.1 and 5.0.

When building from the source code, this module isn't built by default; it should be enabled with the --with-stream\_mqtt\_preread\_module build option.

In packages and images from our repos, the module is included in the build.

### **Configuration Example**

### Choosing an upstream server by the client ID:

```
stream {
    mqtt_preread on;
    upstream mqtt {
        hash $mqtt_preread_clientid;
        # ...
    }
}
```

#### **Directives**

### mqtt preread

| Syntax  | mqtt_preread on   off;       |
|---------|------------------------------|
| Default | <pre>mqtt_preread off;</pre> |
| Context | stream, server               |

Controls extracting information from CONNECT packets during the *preread phase*. If the setting is on, its surrounding context will have the following variables populated.

### **Built-in Variables**

See the details of the value semantics in the specification of MQTT versions 3.1.1 and 5.0.

\$mqtt\_preread\_clientid

Unique client ID.

\$mqtt\_preread\_username

Optional username.

#### Pass

Allows passing the accepted connection directly to any configured listening socket in HTTP, Stream, or Mail modules.

## **Configuration Example**

After the stream module handles the SSL/TLS termination, it forwards the connection to the http module:



```
http {
    server {
        listen 8000;
        location / {
            root html;
        }
    }
}
stream {
    server {
        listen 12345 ssl;
        ssl_certificate
                           domain.crt;
        ssl_certificate_key domain.key;
        pass 127.0.0.1:8000;
    }
}
```

### **Directives**

### pass

| Syntax  | pass address; |
|---------|---------------|
| Default | _             |
| Context | server        |

This directive sets the server address to which the client connection should be passed. The *address* can be given as an IP address and port:

```
pass 127.0.0.1:12345;
```

Or as a path to a UNIX domain socket:

```
pass unix:/tmp/stream.socket;
```

Also, the address can be set with variables:

```
pass $upstream;
```

### **Proxy**

Allows proxying data streams over TCP, UDP, and UNIX domain sockets.

## **Configuration Example**

```
server {
    listen 127.0.0.1:12345;
    proxy_pass 127.0.0.1:8080;
}
server {
    listen 12345;
    proxy_connect_timeout 1s;
```



```
proxy_timeout 1m;
    proxy_pass example.com:12345;
}

server {
    listen 53 udp reuseport;
    proxy_timeout 20s;
    proxy_pass dns.example.com:53;
}

server {
    listen [::1]:12345;
    proxy_pass unix:/tmp/stream.socket;
}
```

### **Directives**

### proxy bind

| Syntax  | proxy_bind address [transparent]   off; |
|---------|---|
| Default | _                                       |
| Context | stream, server                          |

Makes outgoing connections to a proxied server originate from the specified local IP address. Parameter value can contain variables. The special value off cancels the effect of the  $proxy\_bind$  directive inherited from the previous configuration level, which allows the system to auto-assign the local IP address.

The transparent parameter allows outgoing connections to a proxied server originate from a non-local IP address, for example, from a real IP address of a client:

```
proxy_bind $remote_addr transparent;
```

For this parameter to work, Angie worker processes usually need to run with *superuser* privileges. On Linux, this is not required: if the transparent parameter is specified, worker processes inherit the  $CAP\_NET\_RAW$  capability from the master process.

## Important

The kernel routing table should also be configured to intercept network traffic from the FastCGI server.

## proxy\_buffer\_size

| Syntax  | proxy_buffer_size size;           |
|---------|-----------------------------------|
| Default | <pre>proxy_buffer_size 16k;</pre> |
| Context | stream, server                    |

Sets the size of the buffer used for reading data from the proxied server. Also sets the size of the buffer used for reading data from the client.



## proxy\_connect\_timeout

| Syntax  | <pre>proxy_connect_timeout time;</pre> |
|---------|--|
| Default | <pre>proxy_connect_timeout 60s;</pre>  |
| Context | stream, server                         |

Defines a timeout for establishing a connection with a proxied server.

## proxy\_connection\_drop

| Syntax  | proxy_connection_drop time   on   off; |
|---------|--|
| Default | <pre>proxy_connection_drop off;</pre>  |
| Context | stream, server                         |

Enables termination of all sessions to the proxied server after it has been removed from the group or marked as permanently unavailable by a *reresolve* process or the *API command* DELETE.

A session is terminated when the next read or write event is processed for either the client or the proxied server.

Setting time enables a session termination timeout; with on set, sessions are dropped immediately.

### proxy\_download\_rate

| Syntax  | proxy_download_rate rate;         |
|---------|-----------------------------------|
| Default | <pre>proxy_download_rate 0;</pre> |
| Context | stream, server                    |

Limits the speed of reading the data from the proxied server. The rate is specified in bytes per second.

```
0 disables rate limiting
```

## 1 Note

The limit is set per a connection, so if Angie simultaneously opens two connections to the proxied server, the overall rate will be twice as much as the specified limit.

Parameter value can contain variables. It may be useful in cases where rate should be limited depending on a certain condition:

```
map $slow $rate {
    1    4k;
    2   8k;
}
proxy_download_rate $rate;
```

## proxy\_half\_close

| Syntax  | proxy_half_close on   off;       |
|---------|----------------------------------|
| Default | <pre>proxy_half_close off;</pre> |
| Context | stream, server                   |



Enables or disables closing each direction of a TCP connection independently ("TCP half-close"). If enabled, proxying over TCP will be kept until both sides close the connection.

### proxy next upstream

| Syntax  | <pre>proxy_next_upstream on   off;</pre> |
|---------|--|
| Default | <pre>proxy_next_upstream on;</pre>       |
| Context | stream, server                           |

When a connection to the proxied server cannot be established, determines whether a client connection will be passed to the next server in the *upstream pool*.

Passing a connection to the next server can be limited by the *number of tries* and by time.

### proxy next upstream timeout

| Syntax  | $\verb"proxy_next_upstream_time" out $time$;$ |
|---------|---|
| Default | <pre>proxy_next_upstream_timeout 0;</pre>     |
| Context | stream, server                                |

Limits the time allowed to pass a connection to the *next* server.

| 0 turns off this limitation |
|-----------------------------|
|-----------------------------|

### proxy\_next\_upstream\_tries

| Syntax  | proxy_next_upstream_tries number;       |
|---------|---|
| Default | <pre>proxy_next_upstream_tries 0;</pre> |
| Context | stream, server                          |

Limits the number of possible tries for passing a connection to the *next* server.

| 0 turns off this limitation |
|-----------------------------|
|-----------------------------|

#### proxy\_pass

| Syntax  | proxy_pass address; |
|---------|---------------------|
| Default | _                   |
| Context | server              |

Sets the address of a proxied server. The address can be specified as a domain name or IP address, and a port:

```
proxy_pass localhost:12345;
```

or as a UNIX domain socket path:

```
proxy_pass unix:/tmp/stream.socket;
```



If a domain name resolves to several addresses, all of them will be used in a round-robin fashion. In addition, an address can be specified as a *server group*.

The address can also be specified using variables:

```
proxy_pass $upstream;
```

In this case, the server name is searched among the described *server groups* and, if not found, is determined using a *resolver*.

### proxy protocol

| Syntax  | proxy_protocol on   off;       |
|---------|--------------------------------|
| Default | <pre>proxy_protocol off;</pre> |
| Context | stream, server                 |

Enables the PROXY protocol for connections to a proxied server.

### proxy\_requests

| Syntax  | proxy_requests number;       |
|---------|------------------------------|
| Default | <pre>proxy_requests 0;</pre> |
| Context | stream, server               |

Sets the number of client datagrams at which binding between a client and existing UDP stream session is dropped. After receiving the specified number of datagrams, next datagram from the same client starts a new session. The session terminates when all client datagrams are transmitted to a proxied server and the expected *number of responses* is received, or when it reaches a *timeout*.

## proxy responses

| Syntax  | proxy_responses number; |
|---------|-------------------------|
| Default | _                       |
| Context | stream, server          |

Sets the number of datagrams expected from the proxied server in response to a client datagram if the UDP protocol is used. The number serves as a hint for session termination. By default, the number of datagrams is not limited.

If zero value is specified, no response is expected. However, if a response is received and the session is still not finished, the response will be handled.

## proxy socket keepalive

| Syntax  | proxy_socket_keepalive on   off;       |
|---------|--|
| Default | <pre>proxy_socket_keepalive off;</pre> |
| Context | stream, server                         |

Configures the "TCP keepalive" behavior for outgoing connections to a proxied server.

| 11 11 | By default, the operating system's settings are in effect for the socket. |
|-------|---|
| on    | The SO_KEEPALIVE socket option is turned on for the socket.               |



### proxy\_ssl

| Syntax  | <pre>proxy_ssl on   off;</pre> |
|---------|--------------------------------|
| Default | <pre>proxy_ssl off;</pre>      |
| Context | stream, server                 |

Enables the SSL/TLS protocol for connections to a proxied server.

## proxy ssl certificate

| Syntax  | proxy_ssl_certificate file [file]; |
|---------|------------------------------------|
| Default | _                                  |
| Context | stream, server                     |

Specifies a file with the certificate in the PEM format used for authentication to a proxied server. Variables can be used in the file name.

Added in version 1.2.0.

When proxy ssl ntls enabled, directive accepts two arguments instead of one:

```
server {
    proxy_ssl_ntls on;

proxy_ssl_certificate sign.crt enc.crt;
    proxy_ssl_certificate_key sign.key enc.key;

proxy_ssl_ciphers "ECC-SM2-WITH-SM4-SM3:ECDHE-SM2-WITH-SM4-SM3:RSA";

proxy_pass backend:12345;
}
```

## proxy ssl certificate key

```
Syntax proxy_ssl_certificate_key file [file];
Default —
Context stream, server
```

Specifies a file with the secret key in the PEM format used for authentication to a proxied server. Variables can be used in the file name.

Added in version 1.2.0.

When proxy ssl ntls enabled, directive accepts two arguments instead of one:

```
server {
    proxy_ssl_ntls on;

proxy_ssl_certificate sign.crt enc.crt;
    proxy_ssl_certificate_key sign.key enc.key;

proxy_ssl_ciphers "ECC-SM2-WITH-SM4-SM3:ECDHE-SM2-WITH-SM4-SM3:RSA";

proxy_pass backend:12345;
}
```



## proxy\_ssl\_ciphers

| Syntax  | <pre>proxy_ssl_ciphers ciphers;</pre> |
|---------|---------------------------------------|
| Default | <pre>proxy_ssl_ciphers DEFAULT;</pre> |
| Context | stream, server                        |

Specifies the enabled ciphers for requests to a proxied server. The ciphers are specified in the format understood by the OpenSSL library.

The list of ciphers depends on the version of OpenSSL installed. The full list can be viewed using the openssl ciphers command.

### **A** Attention

The proxy\_ssl\_ciphers directive does not configure ciphers for TLS 1.3 when using OpenSSL. To tune TLS 1.3 ciphers with OpenSSL, use the proxy\_ssl\_conf\_command directive, which was added to support advanced SSL configuration.

- In LibreSSL, TLS 1.3 ciphers can be configured using proxy\_ssl\_ciphers.
- In BoringSSL, TLS 1.3 ciphers cannot be configured at all.

## proxy\_ssl\_conf\_command

| Syntax  | proxy_ssl_conf_command name value; |
|---------|------------------------------------|
| Default | _                                  |
| Context | stream, server                     |

Sets arbitrary OpenSSL configuration commands when establishing a connection with the proxied server.

## Important

The directive is supported when using OpenSSL 1.0.2 or higher. To configure TLS 1.3 ciphers with OpenSSL, use the ciphersuites command.

Several proxy\_ssl\_conf\_command directives can be specified on the same level. These directives are inherited from the previous configuration level if and only if there are no proxy\_ssl\_conf\_command directives defined on the current level.

## \* Caution

Note that configuring OpenSSL directly might result in unexpected behavior.

## proxy\_ssl\_crl

| Syntax  | proxy_ssl_crl file; |
|---------|---------------------|
| Default | _                   |
| Context | stream, server      |

Specifies a file with revoked certificates (CRL) in the PEM format used to *verify* the certificate of the proxied server.



## proxy\_ssl\_name

| Syntax  | $proxy_sl_name name;$                           |
|---------|---|
| Default | <pre>proxy_ssl_name host from proxy_pass;</pre> |
| Context | stream, server                                  |

Allows overriding the server name used to *verify* the certificate of the proxied server and to be *passed through SNI* when establishing a connection with the proxied server. The server name can also be specified using variables.

By default, the host part of the address specified in the proxy pass directive is used.

## proxy ssl ntls

Added in version 1.2.0.

```
Syntax proxy_ssl_ntls on | off;
Default proxy_ssl_ntls off;
Context stream, server
```

Enables client-side support for NTLS using the TongSuo TLS library.

```
server {
   proxy_ssl_ntls on;

proxy_ssl_certificate sign.crt enc.crt;
  proxy_ssl_certificate_key sign.key enc.key;

proxy_ssl_ciphers "ECC-SM2-WITH-SM4-SM3:ECDHE-SM2-WITH-SM4-SM3:RSA";

proxy_pass backend:12345;
}
```

### proxy\_ssl\_password\_file

```
Syntax proxy_ssl_password_file file;
Default —
Context stream, server
```

Specifies a file with passphrases for *secret keys* where each passphrase is specified on a separate line. Passphrases are tried in turn when loading the key.



## proxy\_ssl\_protocols

| Syntax  | proxy_ssl_protocols [SSLv2] [SSLv3] [TLSv1] [TLSv1.1] [TLSv1.2] [TLSv1.3]; |
|---------|--|
| Default | <pre>proxy_ssl_protocols TLSv1.2 TLSv1.3;</pre>                            |
| Context | stream, server   |

Changed in version 1.2.0: TLSv1.3 parameter added to default set.

Enables the specified protocols for requests to a proxied server.

## proxy\_ssl\_server\_name

| Syntax  | proxy_ssl_server_name on   off;       |
|---------|---------------------------------------|
| Default | <pre>proxy_ssl_server_name off;</pre> |
| Context | stream, server                        |

Enables or disables passing the server name set by the proxy\_ssl\_name directive via the Server Name Indication TLS extension (SNI, RFC 6066) while establishing a connection with the proxied server.

## proxy\_ssl\_session\_reuse

| Syntax  | proxy_ssl_session_reuse on   off;      |
|---------|--|
| Default | <pre>proxy_ssl_session_reuse on;</pre> |
| Context | stream, server                         |

Determines whether SSL sessions can be reused when working with the proxied server. If the errors " $SSL3\_GET\_FINISHED:digest\ check\ failed$ " appear in the logs, try disabling session reuse.

## proxy ssl trusted certificate

| Syntax  | $	ext{proxy\_ssl\_trusted\_certificate} \ file;$ |
|---------|--|
| Default | _  |
| Context | stream, server                                   |

Specifies a file with trusted CA certificates in the PEM format used to *verify* the certificate of the proxied server.

## proxy\_ssl\_verify

| Syntax  | proxy_ssl_verify on   off;       |
|---------|----------------------------------|
| Default | <pre>proxy_ssl_verify off;</pre> |
| Context | stream, server                   |

Enables or disables verification of the proxied server certificate.

## proxy\_ssl\_verify\_depth

| Syntax  | proxy_ssl_verify_depth number;       |
|---------|--------------------------------------|
| Default | <pre>proxy_ssl_verify_depth 1;</pre> |
| Context | stream, server                       |



Sets the verification depth in the proxied server certificates chain.

## proxy timeout

| Syntax  | proxy_timeout time;           |
|---------|-------------------------------|
| Default | <pre>proxy_timeout 10m;</pre> |
| Context | stream, server                |

Sets the timeout between two successive read or write operations on client or proxied server connections. If no data is transmitted within this time, the connection is closed.

## upstream\_probe\_timeout (PRO)

Added in version 1.4.0: PRO

| Syntax  | $\verb"upstream_probe_time" out $time$;$ |
|---------|--|
| Default | upstream_probe_timeout 50s;              |
| Context | server                                   |

Sets the maximum inactivity time of an established server connection for probes configured using the upstream probe (PRO) directive; if this limit is exceeded, the connection will be closed.

## proxy\_upload\_rate

| Syntax  | proxy_upload_rate rate;         |
|---------|---------------------------------|
| Default | <pre>proxy_upload_rate 0;</pre> |
| Context | stream, server                  |

Limits the speed of reading the data from the client. The rate is specified in bytes per second.

```
0 disables rate limiting
```

## 1 Note

The limit is set per a connection, so if the client simultaneously opens two connections, the overall rate will be twice as much as the specified limit.

Parameter value can contain variables. It may be useful in cases where rate should be limited depending on a certain condition:

```
map $slow $rate {
    1    4k;
    2   8k;
}
proxy_upload_rate $rate;
```

#### **RDP Preread**

When using the RDP protocol, this module allows extracting cookies, which are used for session identification and management, before making a load balancing decision.



When building from the source code, this module isn't built by default; it should be enabled with the --with-stream\_rdp\_preread\_module build option.

In packages and images from our repos, the module is included in the build.

#### **Configuration Example**

## Binding to the Cookie-Issuing Server

This uses the learn mode of the sticky directive:

```
stream {
    rdp_preread on;
    upstream rdp {
        server 127.0.0.1:3390 sid=a;
        server 127.0.0.1:3391 sid=b;

        sticky learn lookup=$rdp_cookie create=$rdp_cookie zone=sessions:1m;
    }
}
```

#### **Directives**

## rdp preread

| Syntax  | rdp_preread on   off; |
|---------|-----------------------|
| Default | rdp_preread off;      |
| Context | stream, server        |

Controls extracting information from RDP protocol cookies during the *preread stage*. If the setting is on, its surrounding context will have the following variables populated.

#### **Built-in Variables**

The semantics of cookie values depend on the RDP protocol version.

#### \$rdp\_cookie

The entire cookie value.

## \$rdp\_cookie\_<name>

The value of the cookie field with the specified name.

### RealIP

The module is used to change the client address and port to the ones sent in the PROXY protocol header. The PROXY protocol must be previously enabled by setting the *proxy\_protocol* parameter in the *listen* directive.

When building from the source code, this module isn't built by default; it should be enabled with the --with-stream\_realip\_module build option.

In packages and images from our repos, the module is included in the build.



### **Configuration Example**

```
listen 12345 proxy_protocol;

set_real_ip_from 192.168.1.0/24;
set_real_ip_from 192.168.2.1;
set_real_ip_from 2001:0db8::/32;
```

#### **Directives**

```
set real ip from
```

| Syntax  | set_real_ip_from address   CIDR   unix:; |
|---------|--|
| Default | _  |
| Context | stream, server                           |

Defines trusted addresses that are known to send correct replacement addresses. If the special value unix: is specified, all UNIX domain sockets will be trusted.

#### **Built-in Variables**

```
$realip_remote_addr
```

keeps the original client address

```
$realip_remote_port
```

keeps the original client port

### Return

The module allows sending a specified value to the client and then closing the connection.

### **Configuration Example**

```
server {
    listen 12345;
    return $time_iso8601;
}
```

#### **Directives**

#### return

| Syntax  | return value; |
|---------|---------------|
| Default | _             |
| Context | server        |

Specifies a value to send to the client. The value can contain text, variables, and their combination.

### Set

The module allows setting a value for a variable.



### **Configuration Example**

```
server {
    listen 12345;
    set $true 1;
}
```

#### **Directives**

set

| Syntax  | set \$variable value; |
|---------|-----------------------|
| Default | _                     |
| Context | server                |

Sets a value for the specified variable. The value can contain text, variables, and their combination.

### **Split Clients**

The module emits variables for A/B testing, canary releases, or other scenarios that require routing a percentage of clients to one server or configuration while directing the rest elsewhere.

### **Configuration Example**

### **Directives**

## split\_clients

```
Syntax split_clients string $variable { ... }

Default —
Context stream
```

Creates a *\$variable* by hashing the *string*; the variables in *string* are substituted, the result is hashed, then the hash is mapped to the *\$variable*'s string value.

The hash function uses MurmurHash2 (32-bit), and its entire value range (0 to 4294967295) is mapped to buckets in order of appearance; the percentages determine the size of the buckets. A wildcard (\*) may occur last; hashes that fall outside other buckets are mapped to its assigned value.

An example:



Here, the hashed values of the interpolated \$remote\_addrAAA string are distributed as follows:

- $\bullet$  values 0 to 21474835 (a sample of 0.5%) yield .one
- values 21474836 to 107374180 (a sample of 2%) yield .two
- values 107374181 to 4294967295 (all other values) yield "" (an empty string)

#### **SSL**

Provides the necessary support for a stream proxy server to work with the SSL/TLS protocol.

When building from the source code, this module isn't built by default; it should be enabled with the --with-stream\_ssl\_module build option.

In packages and images from our repos, the module is included in the build.

```
• Important
This module requires the OpenSSL library.
```

### **Configuration Example**

To reduce the processor load it is recommended to

- set the number of worker processes equal to the number of processors,
- enable the *shared* session cache,
- disable the *built-in* session cache,
- and possibly increase the session *lifetime* (by default, 5 minutes):



#### **Directives**

## ssl\_alpn

```
Syntax ssl_alpn protocol ...;

Default —
Context stream, server
```

Specifies the list of supported ALPN protocols. One of the protocols must be *negotiated* if the client uses ALPN:

#### ssl certificate

```
Syntax ssl_certificate file;
Default —
Context stream, server
```

Specifies a file with the certificate in the PEM format for the given server. If intermediate certificates should be specified in addition to a primary certificate, they should be specified in the same file in the following order: the primary certificate comes first, then the intermediate certificates. A secret key in the PEM format may be placed in the same file.

This directive can be specified multiple times to load certificates of different types, for example, RSA and ECDSA:

Only OpenSSL 1.0.2 or higher supports separate certificate chains for different certificates. With older versions, only one certificate chain can be used.

## Important

Variables can be used in the file name when using OpenSSL 1.0.2 or higher:



```
ssl_certificate $ssl_server_name.crt;
ssl_certificate_key $ssl_server_name.key;
```

Note that using variables implies that a certificate will be loaded for each SSL handshake, and this may have a negative impact on performance.

The value "data: \$variable" can be specified instead of the file, which loads a certificate from a variable without using intermediate files.

Note that inappropriate use of this syntax may have its security implications, such as writing secret key data to *error log*.

## ssl\_certificate\_key

| Syntax  | ${	t ssl\_certificate\_key} \; file;$ |
|---------|---------------------------------------|
| Default | _                                     |
| Context | stream, server                        |

Specifies a file with the secret key in the PEM format for the given server.

## Important

Variables can be used in the file name when using OpenSSL 1.0.2 or higher.

The value "engine: name: id" can be specified instead of the file, which loads a secret key with a specified id from the OpenSSL engine name.

The value "data: \$variable" can be specified instead of the file, which loads a secret key from a variable without using intermediate files. Note that inappropriate use of this syntax may have its security implications, such as writing secret key data to error log.

### ssl ciphers

| Syntax  | ssl_ciphers ciphers;          |
|---------|-------------------------------|
| Default | ssl_ciphers HIGH:!aNULL:!MD5; |
| Context | stream, server                |

Specifies the enabled ciphers. The ciphers are specified in the format understood by the OpenSSL library, for example:

```
ssl_ciphers ALL:!aNULL:!EXPORT56:RC4+RSA:+HIGH:+MEDIUM:+LOW:+SSLv2:+EXP;
```

The list of ciphers depends on the version of OpenSSL installed. The full list can be viewed using the openssl ciphers command.

#### **A** Attention

The ssl\_ciphers directive does *not* configure ciphers for TLS 1.3 when using OpenSSL. To tune TLS 1.3 ciphers with OpenSSL, use the *ssl\_conf\_command* directive, which was added to support advanced SSL configuration.

- In LibreSSL, TLS 1.3 ciphers can be configured using ssl\_ciphers.
- In BoringSSL, TLS 1.3 ciphers cannot be configured at all.



## ssl\_client\_certificate

| Syntax  | ${\tt ssl\_client\_certificate}\ file;$ |
|---------|---|
| Default | _                                       |
| Context | stream, server                          |

Specifies a file with trusted CA certificates in the PEM format used to verify client certificates and OCSP responses if  $ssl\_stapling$  is enabled.

The list of certificates will be sent to clients. If this is not desired, the *ssl\_trusted\_certificate* directive can be used.

## ssl conf command

| Syntax  | ssl_conf_command name value; |
|---------|------------------------------|
| Default | _                            |
| Context | stream, server               |

Sets arbitrary OpenSSL configuration commands.

## Important

The directive is supported when using OpenSSL 1.0.2 or higher. To configure TLS 1.3 ciphers with OpenSSL, use the ciphersuites command.

Several  $ssl\_conf\_command$  directives can be specified on the same level:

```
ssl_conf_command Options PrioritizeChaCha;
ssl_conf_command Ciphersuites TLS_CHACHA20_POLY1305_SHA256;
```

These directives are inherited from the previous configuration level if and only if there are no  $ssl\ conf\ command\ directives\ defined\ on\ the\ current\ level.$ 

## \* Caution

Note that configuring OpenSSL directly might result in unexpected behavior.

## ssl\_crl

| Syntax  | $	exttt{ssl\_crl}\ \mathit{file};$ |
|---------|------------------------------------|
| Default | _                                  |
| Context | stream, server                     |

Specifies a file with revoked certificates (CRL) in the PEM format used to verify client certificates.

## ssl dhparam

| Syntax  | ${\tt ssl\_dhparam}\;file;$ |
|---------|-----------------------------|
| Default | _                           |
| Context | stream, server              |



Specifies a file with DH parameters for DHE ciphers.

## \* Caution

By default no parameters are set, and therefore DHE ciphers will not be used.

## ssl\_early\_data

| Syntax  | ssl_early_data on   off; |
|---------|--------------------------|
| Default | ssl_early_data off;      |
| Context | stream, server           |

Enables or disables TLS 1.3 early data.

## Important

The directive is supported when using OpenSSL 1.1.1 or higher or BoringSSL.

## ssl\_ecdh\_curve

| Syntax  | ssl_ecdh_curve curve; |
|---------|-----------------------|
| Default | ssl_ecdh_curve auto;  |
| Context | stream, server        |

Specifies a curve for ECDHE ciphers.

## Important

When using OpenSSL 1.0.2 or higher, it is possible to specify multiple curves, for example:

```
ssl_ecdh_curve prime256v1:secp384r1;
```

The special value auto instructs Angie to use a list built into the OpenSSL library when using OpenSSL 1.0.2 or higher, or prime256v1 with older versions.

## Important

When using OpenSSL 1.0.2 or higher, this directive sets the list of curves supported by the server. Thus, in order for ECDSA certificates to work, it is important to include the curves used in the certificates.

#### ssl handshake timeout

| Syntax  | ${\tt ssl\_handshake\_timeout}\ time;$ |
|---------|--|
| Default | ssl_handshake_timeout 60s;             |
| Context | stream, server                         |

Specifies a timeout for the SSL handshake to complete.



### ssl\_ocsp

| Syntax  | ssl_ocsp on   off   leaf; |
|---------|---------------------------|
| Default | ssl_ocsp off;             |
| Context | http, server              |

Enables OCSP validation of the client certificate chain. The leaf parameter enables validation of the client certificate only.

For the OCSP validation to work, the ssl\_verify\_client directive should be set to on or optional.

To resolve the OCSP responder hostname, the resolver directive should also be specified.

#### Example:

## ssl\_ocsp\_cache

| Syntax  | ssl_ocsp_cache off   [shared:name:size]; |
|---------|--|
| Default | ssl_ocsp_cache off;                      |
| Context | http, server                             |

Sets name and size of the cache that stores client certificates status for OCSP validation. The cache is shared between all worker processes. A cache with the same name can be used in several virtual servers.

The off parameter prohibits the use of the cache.

## ssl\_ocsp\_responder

| Syntax  | ssl_ocsp_responder $uri;$ |
|---------|---------------------------|
| Default | _                         |
| Context | http, server              |

Overrides the URI of the OCSP responder specified in the "Authority Information Access" certificate extension for validation of client certificates.

Only http:// OCSP responders are supported:

```
ssl_ocsp_responder http://ocsp.example.com/;
```

### ssl ntls

Added in version 1.2.0.

```
Syntax ssl_ntls on | off;
Default ssl_ntls off;
Context stream, server
```

Enables server-side support for NTLS using TongSuo library.

```
listen ... ssl;
ssl_ntls on;
```



### Important

Build Angie using the --with-ntls build option and link with NTLS-enabled SSL library

```
./configure --with-openssl=../Tongsuo-8.3.0 \
    --with-openssl-opt=enable-ntls \
    --with-ntls
```

### ssl password file

```
Syntax ssl_password_file file;
Default —
Context stream, server
```

Specifies a file with passphrases for *secret keys* where each passphrase is specified on a separate line. Passphrases are tried in turn when loading the key.

### Example:

```
stream {
    ssl_password_file /etc/keys/global.pass;
    ...

server {
        listen 127.0.0.1:12345;
        ssl_certificate_key /etc/keys/first.key;
}

server {
        listen 127.0.0.1:12346;

        # named pipe can also be used instead of a file
        ssl_password_file /etc/keys/fifo;
        ssl_certificate_key /etc/keys/second.key;
}
```

## ssl prefer server ciphers

```
Syntax ssl_prefer_server_ciphers on | off;
Default ssl_prefer_server_ciphers off;
Context stream, server
```

Specifies that server ciphers should be preferred over client ciphers when the SSLv3 and TLS protocols are used.

### ssl\_protocols

| Syntax  | ssl_protocols [SSLv2] [SSLv3] [TLSv1] [TLSv1.1] [TLSv1.2] [TLSv1.3]; |
|---------|--|
| Default | ssl_protocols TLSv1.2 TLSv1.3;                                       |
| Context | stream, server   |

Changed in version 1.2.0: TLSv1.3 parameter added to default set.



Enables the specified protocols.

## Important

The TLSv1.1 and TLSv1.2 parameters work only when OpenSSL 1.0.1 or higher is used.

The TLSv1.3 parameter works only when OpenSSL 1.1.1 or higher is used.

## ssl\_session\_cache

| Syntax  | ssl_session_cache off   none   [builtin[:size]] [shared:name:size]; |
|---------|---|
| Default | ssl_session_cache none;   |
| Context | stream, server  |

Sets the types and sizes of caches that store session parameters. A cache can be of any of the following types:

| off     | the use of a session cache is strictly prohibited: Angie explicitly tells a client that sessions may not be reused.  |
|---------|--|
| none    | the use of a session cache is gently disallowed: Angie tells a client that sessions may be reused, but does not actually store session parameters in the cache.  |
| builtin | a cache built in OpenSSL; used by one worker process only. The cache size is specified in sessions. If size is not given, it is equal to 20480 sessions. Use of the built-in cache can cause memory fragmentation.   |
| shared  | a cache shared between all worker processes. The cache size is specified in bytes; one megabyte can store about 4000 sessions. Each shared cache should have an arbitrary name. A cache with the same name can be used in several servers. It is also used to automatically generate, store, and periodically rotate TLS session ticket keys unless configured explicitly using the $ssl\_session\_ticket\_key$ directive. |

Both cache types can be used simultaneously, for example:

```
ssl_session_cache builtin:1000 shared:SSL:10m;
```

but using only shared cache without the built-in cache should be more efficient.

### ssl session ticket key

| Syntax  | ssl_session_ticket_key file; |
|---------|------------------------------|
| Default | _                            |
| Context | stream, server               |

Sets a file with the secret key used to encrypt and decrypt TLS session tickets. The directive is necessary if the same key has to be shared between multiple servers. By default, a randomly generated key is used.

If several keys are specified, only the first key is used to encrypt TLS session tickets. This allows configuring key rotation, for example:

```
ssl_session_ticket_key current.key;
ssl_session_ticket_key previous.key;
```

The file must contain 80 or 48 bytes of random data and can be created using the following command:



```
openssl rand 80 > ticket.key
```

Depending on the file size either AES256 (for 80-byte keys) or AES128 (for 48-byte keys) is used for encryption.

## ssl\_session\_tickets

| Syntax  | ssl_session_tickets on   off; |
|---------|-------------------------------|
| Default | ssl_session_tickets on;       |
| Context | stream, server                |

Enables or disables session resumption through TLS session tickets.

## $ssl\_session\_timeout$

| Syntax  | ${\tt ssl\_session\_timeout}\ time;$ |
|---------|--------------------------------------|
| Default | ssl_session_timeout 5m;              |
| Context | stream, server                       |

Specifies a time during which a client may reuse the session parameters.

### ssl stapling

| Syntax  | ssl_stapling on   off; |
|---------|------------------------|
| Default | ssl_stapling off;      |
| Context | http, server           |

Enables or disables stapling of OCSP responses by the server. Example:

```
ssl_stapling on; resolver 127.0.0.53;
```

For the OCSP stapling to work, the certificate of the server certificate issuer should be known. If the  $ssl\_certificate$  file does not contain intermediate certificates, the certificate of the server certificate issuer should be present in the  $ssl\_trusted\_certificate$  file.

#### **A** Attention

For the resolution of the OCSP responder hostname, the resolver directive should also be specified.

### ssl\_stapling\_file

| Syntax  | ${\tt ssl\_stapling\_file}$ $file;$ |
|---------|-------------------------------------|
| Default | _                                   |
| Context | http, server                        |

When set, the stapled OCSP response will be taken from the specified file instead of querying the OCSP responder specified in the server certificate.

The file should be in the DER format as produced by the openssl ocsp command.



## ssl\_stapling\_responder

| Syntax  | ssl_stapling_responder uri; |
|---------|-----------------------------|
| Default | _                           |
| Context | http, server                |

Overrides the URI of the OCSP responder specified in the "Authority Information Access" certificate extension.

Only http:// OCSP responders are supported:

```
ssl_stapling_responder http://ocsp.example.com/;
```

## ssl stapling verify

| Syntax  | ssl_stapling_verify on   off; |
|---------|-------------------------------|
| Default | ssl_stapling_verify off;      |
| Context | http, server                  |

Enables or disables verification of OCSP responses by the server.

For verification to work, the certificate of the server certificate issuer, the root certificate, and all intermediate certificates should be configured as trusted using the *ssl trusted certificate* directive.

## ssl\_trusted\_certificate

| Syntax  | ${	t ssl\_trusted\_certificate} \ file;$ |
|---------|--|
| Default | _  |
| Context | stream, server                           |

Specifies a file with trusted CA certificates in the PEM format used to verify client certificates.

In contrast to the certificate set by  $ssl\_client\_certificate$ , the list of these certificates will not be sent to clients.

## ssl\_verify\_client

| Syntax  | ssl_verify_client on   off   optional   optional_no_ca; |
|---------|---|
| Default | ssl_verify_client off;                                  |
| Context | stream, server  |

Enables verification of client certificates. The verification result is stored in the \$ssl\_client\_verify variable. If an error has occurred during the client certificate verification or a client has not presented the required certificate, the connection is closed.

| optional       | requests the client certificate and verifies it if the certificate is present.       |
|----------------|--|
| optional_no_ca | requests the client certificate but does not require it to be signed by a trusted CA |
|                | certificate. This is intended for the use in cases when a service that is external   |
|                | to Angie performs the actual certificate verification.                               |



## ssl\_verify\_depth

| Syntax  | ${\tt ssl\_verify\_depth}\ number;$ |
|---------|-------------------------------------|
| Default | ssl_verify_depth 1;                 |
| Context | stream, server                      |

Sets the verification depth in the client certificates chain.

### **Built-in Variables**

The stream\_ssl module supports the following variables:

#### \$ssl\_alpn\_protocol

returns the protocol selected by ALPN during the SSL handshake, or an empty string otherwise.

#### \$ssl\_cipher

returns the name of the cipher used for an established SSL connection.

#### \$ssl\_ciphers

returns the list of ciphers supported by the client. Known ciphers are listed by names, unknown are shown in hexadecimal, for example:

AES128-SHA:AES256-SHA:0x00ff

## Important

The variable is fully supported only when using OpenSSL version 1.0.2 or higher. With older versions, the variable is available only for new sessions and lists only known ciphers.

### \$ssl\_client\_cert

returns the client certificate in the PEM format for an established SSL connection, with each line except the first prepended with the tab character.

#### \$ssl\_client\_fingerprint

returns the SHA1 fingerprint of the client certificate for an established SSL connection.

## \$ssl\_client\_i\_dn

returns the "issuer DN" string of the client certificate for an established SSL connection according to RFC 2253.

### \$ssl\_client\_raw\_cert

returns the client certificate in the PEM format for an established SSL connection.

### \$ssl\_client\_s\_dn

returns the "subject DN" string of the client certificate for an established SSL connection according to RFC 2253.



#### \$ssl\_client\_serial

returns the serial number of the client certificate for an established SSL connection.

### \$ssl\_client\_v\_end

returns the end date of the client certificate.

#### \$ssl\_client\_v\_remain

returns the number of days until the client certificate expires.

#### \$ssl\_client\_v\_start

returns the start date of the client certificate.

#### \$ssl\_client\_verify

returns the result of client certificate verification: SUCCESS, FAILED:reason and NONE if a certificate was not present.

#### \$ssl\_curve

returns the negotiated curve used for SSL handshake key exchange process. Known curves are listed by names, unknown are shown in hexadecimal, for example:

prime256v1

## Important

The variable is supported only when using OpenSSL version 3.0 or higher. With older versions, the variable value will be an empty string.

#### \$ssl\_curves

returns the list of curves supported by the client. Known curves are listed by names, unknown are shown in hexadecimal, for example:

0x001d:prime 256v1:secp 521r1:secp 384r1

## Important

The variable is supported only when using OpenSSL version 1.0.2 or higher. With older versions, the variable value will be an empty string.

The variable is available only for new sessions.

### \$ssl\_early\_data

returns "1" if TLS 1.3 early data is used and the handshake is not complete, otherwise "".

## \$ssl\_protocol

returns the protocol of an established SSL connection.



# \$ssl\_server\_cert\_type

takes the values RSA, DSA, ECDSA, ED448, ED25519, SM2, RSA-PSS, or unknown depending on the type of server certificate and key.

#### \$ssl\_server\_name

returns the server name requested through SNI.

#### \$ssl\_session\_id

returns the session identifier of an established SSL connection.

#### \$ssl\_session\_reused

returns r if an SSL session was reused, or "." otherwise

### **SSL** Preread

Enables extracting information from the ClientHello message without terminating SSL/TLS, such as the server name requested via SNI or protocols advertised in ALPN.

When building from the source code, this module isn't built by default; it should be enabled with the --with-stream\_ssl\_preread\_module build option.

In packages and images from our repos, the module is included in the build.

### **Configuration Example**

# Selecting an upstream by server name

```
map $ssl_preread_server_name $name {
   backend.example.com
                            backend;
    default
                             backend2;
}
upstream backend {
   server 192.168.0.1:12345;
    server 192.168.0.2:12345;
}
upstream backend2 {
   server 192.168.0.3:12345;
   server 192.168.0.4:12345;
}
server {
   listen
               12346;
   proxy_pass $name;
   ssl_preread on;
```

# Selecting a server by protocol



```
server {
   listen 9000;
   proxy_pass $proxy;
   ssl_preread on;
}
```

#### Selecting a server by SSL version

#### **Directives**

# ssl\_preread

```
Syntax ssl_preread on | off;
Default ssl_preread off;
Context stream, server
```

Enables extracting information from the ClientHello message at the preread phase.

#### **Built-in Variables**

```
$ssl_preread_protocol
```

Highest SSL version supported by the client.

```
$ssl_preread_server_name
```

Server name requested via SNI.

```
$ssl_preread_alpn_protocols
```

List of protocols advertised by the client through ALPN. The values are comma separated.

# **Upstream**

The module is used to define groups of servers that can be referenced by the proxy pass directive.

#### **Configuration Example**

```
upstream backend {
   hash $remote_addr consistent;
   zone backend 1m;
```



```
server backend1.example.com:1935 weight=5;
server unix:/tmp/backend3;
server backend3.example.com service=_example._tcp resolve;

server backup1.example.com:1935 backup;
server backup2.example.com:1935 backup;
}

resolver 127.0.0.53 status_zone=resolver;

server {
    listen 1936;
    proxy_pass backend;
}
```

#### **Directives**

#### upstream

```
Syntax upstream name { ... }
Default —
Context stream
```

Defines a group of servers. Servers can listen on different ports. In addition, servers listening on TCP and UNIX domain sockets can be mixed.

#### Example:

By default, requests are distributed between the servers using a weighted round-robin balancing method. In the above example, each 7 requests will be distributed as follows: 5 requests go to backend1.example.com and one request to each of the second and third servers.

If an error occurs during communication with a server, the request will be passed to the next server, and so on until all of the functioning servers will be tried. If a successful response could not be obtained from any of the servers, the client will receive the result of the communication with the last server.

#### server

| Syntax  | server address [parameters]; |
|---------|------------------------------|
| Default | _                            |
| Context | upstream                     |

Defines the address and other parameters of a server. The address can be specified as a domain name or IP address with an obligatory port, or as a UNIX domain socket path specified after the unix: prefix. A domain name that resolves to several IP addresses defines multiple servers at once.



The following parameters can be defined:

| ${\tt weight} {=} number$   | sets the weight of the server; by default, 1.                                       |
|-----------------------------|---|
| ${\tt max\_conns} = number$ | limits the maximum number of simultaneous active connections to the proxied         |
|                             | server. Default value is 0, meaning there is no limit. If the server group does not |
|                             | reside in the <i>shared memory</i> , the limitation works per each worker process.  |

max\_fails=number — sets the number of unsuccessful attempts to communicate with the server that should happen in the duration set by fail\_timeout to consider the server unavailable; it is then retried after the same duration.

Here, an unsuccessful attempt is an error or timeout while establishing a connection with the server.

### 1 Note

If a server in an upstream resolves into multiple peers, its max\_fails setting applies to each peer individually.

If an upstream contains only one peer after all its server directives are resolved, the max\_fails setting has no effect and will be ignored.

| max_fails=1 | the default number of unsuccessful attempts |
|-------------|---|
| max_fails=0 | disables the accounting of attempts         |

 $\mathtt{fail\_timeout}$ =time — sets the period of time during which a number of unsuccessful attempts to communicate with the server  $(max\_fails)$  should happen to consider the server unavailable. The server then becomes unavailable for the same amount of time before it is retried.

By default, this is set to 10 seconds.

#### Note

If a server in an upstream resolves into multiple peers, its fail\_timeout setting applies to each peer individually.

If an upstream contains only one peer after all its server directives are resolved, the fail\_timeout setting has no effect and will be ignored.

| backup      | marks the server as a backup server. It will be passed requests when the primary servers are unavailable.  |  |
|-------------|--|--|
| down        | marks the server as permanently unavailable.   |  |
| drain (PRO) | sets the server to draining; this means it receives only requests from the sessions that were bound earlier with <i>sticky</i> . Otherwise it behaves similarly to down. |  |

# \* Caution

The backup parameter cannot be used along with the hash and random load balancing methods.

The down and drain options are mutually exclusive.

Added in version 1.3.0.



| resolve      | Enables monitoring changes to the list of IP addresses that corresponds to a domain name, updating it without a configuration reload. The group should be stored in a <i>shared memory zone</i> ; also, you need to define a <i>resolver</i> .   |
|--------------|--|
| service=name | Enables resolving DNS SRV records and sets the service name. For this parameter to work, specify the resolve server parameter, providing a hostname without a port number.  If there are no dots in the service name, the name is formed according to the RFC standard: the service name is prefixed with _, then _tcp is added after a dot. Thus, the service name http will result in _httptcp.  Angie resolves the SRV records by combining the normalized service name and the hostname and obtaining the list of servers for the combination via DNS, along with their priorities and weights.  • Top-priority SRV records (ones that share the minimum priority value) resolve into primary servers, and other records become backup servers. If backup is set with server, top-priority SRV records resolve into backup servers, and other records are ignored.  • Weight influences the selection of servers by the assigned capacity: higher weights receive more requests. If set by both the server directive and the SRV record, the weight set by server is used. |

This example will look up the \_http.\_tcp.backend.example.com record:

server backend.example.com service=http resolve;

Added in version 1.4.0.

| slow_start=time sets the time to recover the weight for a server that goes back onlin |   |
|---|---|
|   | balancing uses the round-robin or least_conn method.                                  |
|   | If the value is set and the server is again considered available and healthy as       |
|   | defined by max_fails and upstream_probe (PRO), the server will steadily recover       |
|   | its designated weight within the allocated timeframe.                                 |
|   | If the value isn't set, the server in a similar situation will recover its designated |
|   | weight immediately.   |



If there's only one server in an upstream, slow\_start has no effect and will be ignored.

# state (PRO)

Added in version 1.4.0: PRO

| Syntax  | $\mathtt{state}\ \mathit{file};$ |
|---------|----------------------------------|
| Default | _                                |
| Context | upstream                         |

Specifies the *file* where the upstream's server list is persisted. When installing from our packages, a designated /var/lib/angie/state/ (/var/db/angie/state/ on FreeBSD) directory with appropriate permissions is created to store these files, so you will only need to add the file's basename in the configuration:



```
upstream backend {
   zone backend 1m;
   state /var/lib/angie/state/<FILE NAME>;
}
```

The format of this server list is similar to **server**. The contents of the file change whenever there is any modification to servers in the /config/stream/upstreams/ section via the configuration API. The file is read at Angie start or configuration reload.

### Caution

For the state directive to be used in an upstream block, the block should have no server directives; instead, it must have a shared memory zone (zone).

#### zone

| Syntax  | zone name [size]; |
|---------|-------------------|
| Default | _                 |
| Context | upstream          |

Defines the name and size of the shared memory zone that keeps the group's configuration and run-time state that are shared between worker processes. Several groups may share the same zone. In this case, it is enough to specify the size only once.

### feedback (PRO)

Added in version 1.7.0: PRO

| Syntax  | <pre>feedback variable [invers [last_byte];</pre> | ] [factor=number] [account=condition_variable] |
|---------|---|--|
| Default | _   |  |
| Context | upstream  |  |

Enables a feedback-based load balancing mechanism for the upstream. It adjusts the load balancing decisions dynamically, multiplying each peer's weight by its average feedback value that is affected by the value of a *variable* over time and is subject to an optional condition.

The following parameters are accepted:



| variable | The variable from which the feedback value is taken. It should represent a performance or health metric, and is intended to be supplied by the peer.  The value is assessed at each response from the peer and factored into the rolling average according to inverse and factor settings.   |  |
|----------|--|--|
| inverse  | If set, the feedback value is interpreted inversely, meaning lower values indicate better performance.   |  |
| factor   | The factor by which the feedback value is weighted when calculating the average. Valid values are integers between 0 and 99. By default — 90. The average feedback is calculated using the exponential moving average formula. The larger is the factor, the less is the average affected by new values; if the factor is set to 90, the result has 90% of the previous value and only 10% of the new value. |  |
| account  | Specifies a condition variable that controls how connections are included in the calculation. The average is updated with the feedback value only if the condition variable isn't "" or "0".  1 Note  By default, traffic from probes isn't included in the calculation; combining the \$upstream_probe variable with account allows to include them or even exclude everything else.                        |  |

### Example:

```
upstream backend {
    zone backend 1m;
    feedback $feedback_value factor=80 account=$condition_value;
    server backend1.example.com:1935 weight=1;
    server backend2.example.com:1935 weight=2;
}
map $protocol $feedback_value {
    "TCP"
                               100;
    "UDP"
                               75;
                               10;
    default
}
map $upstream_probe $condition_value {
    "high_priority" "1";
    "low_priority" "0";
    default
                    "1";
}
```

This categorizes servers into different feedback levels based on specific protocols used for different sessions, and also adds a condition mapped from  $\$upstream\_probe$  to account only for the high\_priority probe or regular client sessions.

### hash

| Syntax  | $	ext{hash } key 	ext{ [consistent]};$ |
|---------|--|
| Default | _                                      |
| Context | upstream                               |



Specifies a load balancing method for a server group where the client-server mapping is based on the hashed key value. The key can contain text, variables, and their combinations (1.11.2). Usage example:

# hash \$remote\_addr;

Note that adding or removing a server from the group may result in remapping most of the keys to different servers. The method is compatible with the Cache::Memcached Perl library.

If the consistent parameter is specified, the ketama consistent hashing method will be used instead. The method ensures that only a few keys will be remapped to different servers when a server is added to or removed from the group. This helps to achieve a higher cache hit ratio for caching servers. The method is compatible with the Cache::Memcached::Fast Perl library with the ketama\_points parameter set to 160.

# least conn

| Syntax  | least_conn; |
|---------|-------------|
| Default | _           |
| Context | upstream    |

Specifies that a group should use a load balancing method where a connection is passed to the server with the least number of active connections, taking into account weights of servers. If there are several such servers, they are tried in turn using a weighted round-robin balancing method.

# least time (PRO)

| Syntax  | <pre>least_time connect   [account=condition_variable];</pre> | first_byte | last_byte | [factor=number] |
|---------|---|------------|-----------|-----------------|
| Default | _   |            |           |                 |
| Context | upstream  |            |           |                 |

Sets the load balancing method for a group where the probability of forwarding a connection to an active server is inversely proportional to the average time it takes to respond; the smaller the response time, the more conections the server will receive.

| connect    | The directive accounts for the average time to establish the connection.       |
|------------|--|
| first_byte | The directive uses the average time to receive the first byte of the response. |
| last_byte  | The directive uses the average time to receive the entire response.            |

#### Added in version 1.7.0: PRO

| factor  | Serves the same purpose as response_time_factor (PRO) and overrides it if set.  |
|---------|---|
| account | Specifies a condition variable that controls which connections should be included in the calculation. The average is updated only if the condition variable for the connection isn't "" or "0". |
|         | 1 Note  |
|         | By default, <i>probes</i> aren't included in the calculation; combining the <i>\$up-stream_probe</i> variable with account allows to include them or even exclude everything else.              |

The respective moving averages, adjusted for factor and account, are also presented as connect\_time, first\_byte\_time, and last\_byte\_time in the health object of the server among the stream upstream



metrics in the API.

#### random

| Syntax  | random [two]; |
|---------|---------------|
| Default | _             |
| Context | upstream      |

Specifies that a group should use a load balancing method where a request is passed to a randomly selected server, taking into account weights of servers.

The optional two parameter instructs Angie to randomly select two servers and then choose a server using the specified method. The default method is least\_conn which passes a request to a server with the least number of active connections.

# response time factor (PRO)

| Syntax  | response_time_factor number; |
|---------|------------------------------|
| Default | response_time_factor 90;     |
| Context | upstream                     |

Sets the smoothing factor for the least time (PRO) load balancing method, using the **previous** value when calculating the average response time according to the formula of the exponential weighted moving average.

The larger the specified number, the less new values influence the average; if 90 is specified, 90% of the previous value will be taken, and only 10% of the new value. Acceptable values range from 0 to 99 inclusive.

The respective moving averages are presented as connect\_time (time to establish the connection), first\_byte\_time (time to receive the first byte of the response), and last\_byte\_time (time to receive the complete response) in the health object of the server among the stream upstream metrics in the API.



#### Note

Only successful responses are considered in the calculation; what constitutes an unsuccessful response is determined by the proxy next upstream directives.

#### sticky

Added in version 1.6.0: Angie

Added in version 1.6.0: Angie PRO

| Syntax  | <pre>sticky route \$variable; sticky learn zone=zone create=\$create_var1 lookup=\$lookup_var1 [connect] [timeout=time];</pre> |
|---------|--|
| Default | _  |
| Context | upstream   |

Configures the binding of client sessions to proxied servers in the mode specified by the first parameter; to drain requests from servers that have sticky defined, use the drain option in the server block.



#### Attention

The sticky directive must be used after all directives that set the load balancing method; otherwise, it won't work.

#### route mode

This mode uses predefined route identifiers that can be embedded in any connection properties Angie can access. It is less flexible because it relies on predefined values but can suit better if such identifiers are already in place.

Here, when a connection is established with the proxied server, it can assign a route to the client and return its identifier in a manner that they both are aware of. The value of the *sid* parameter of the *server* directive must be used as the route identifier. Note that the parameter is additionally hashed if the *sticky secret* directive is set.

Subsequent connections from clients that wish to use this route must contain the identifier issued by the server in a way that ensures it ends up in Angie variables.

The directive lists specific variables used for routing. To select the server where the incoming connection is routed, the first non-empty variable is used; it is then compared with the *sid* parameter of the *server* directive. If selecting a server fails or the chosen server can't accept the connection, another server is selected according to the configured balancing method.

Here, Angie looks for the identifier in a custom **\$route** variable, which is mapped from *\$ssl preread server name* (note that *ssl preread* must be enabled):

```
stream {
    map $ssl_preread_server_name $route {
        a.example.com
                                  a;
        b.example.com
        default
    upstream backend {
        server 127.0.0.1:8081 sid=a;
        server 127.0.0.1:8082 sid=b;
        sticky route $route;
    }
    server {
        listen 127.0.0.1:8080;
        ssl_preread on;
        proxy_pass backend;
    }
}
```

#### learn mode (PRO)

This mode uses a dynamically generated key to associate a client with a particular proxied server; it's more flexible because it assigns servers on the go, stores sessions in a shared memory zone, and supports different ways of passing session identifiers.

Here, a session is created based on the connection properties from the proxied server. The create and



lookup parameters list variables indicating how new sessions are created and existing sessions are looked up. Both parameters can occur multiple times.

The session identifier is the value of the first non-empty variable specified with create; for example, this could be the *name of the proxied server*.

Sessions are stored in a shared memory zone; its name and size are set by the zone parameter. If a session has been inactive for the time set by timeout, it is deleted. The default is 1 hour.

Subsequent connections from clients that wish to use the session must contain its identifier, ensuring that it ends up in a non-empty variable specified with lookup; its value will then be matched against sessions in shared memory. If selecting a server fails or the chosen server can't accept the connection, another server is selected according to the configured balancing method.

The connect parameter allows creating a session immediately after the connection to the proxied server was established. Without it, a session is created only after processing the connection.

In the example, Angie creates and looks up sessions, using the \$rdp cookie variable:

```
stream {
    upstream backend {
        server 127.0.0.1:3390 sid=a;
        server 127.0.0.1:3391 sid=b;

        sticky learn lookup=$rdp_cookie create=$rdp_cookie zone=sessions:1m;
}

server {
        listen 127.0.0.1:3389;
        ssl_preread on;
        proxy_pass backend;
}
```

### sticky strict

Added in version 1.6.0: Angie

Added in version 1.6.0: Angie PRO

```
Syntax sticky_strict on | off;
Default sticky_strict off;
Context upstream
```

When enabled, causes Angie to return a connection error to the client if the desired server is unavailable, rather than using any other available server as it would when no servers in the upstream are available.

# sticky\_secret

Added in version 1.6.0: Angie

Added in version 1.6.0: Angie PRO



| Syntax  | sticky_secret $string$ ; |
|---------|--------------------------|
| Default | _                        |
| Context | upstream                 |

Adds the *string* as the salt value to the MD5 hashing function for the *sticky* directive in the route mode. The *string* may contain variables, for example, *\$remote addr*:

```
upstream backend {
    server 127.0.0.1:8081 sid=a;
    server 127.0.0.1:8082 sid=b;

    sticky route $route;
    sticky_secret my_secret.$remote_addr;
}
```

Salt is appended to the value being hashed; to verify the hashing mechanism independently:

```
$ echo -n "<VALUE><SALT>" | md5sum
```

#### **Built-in Variables**

The stream\_upstream module supports the following built-in variables:

#### \$upstream\_addr

stores the IP address and port, or the path to the UNIX domain socket of the upstream server. If several servers were contacted during request processing, their addresses are separated by commas, e.g.:

```
192.168.1.1:1935, 192.168.1.2:1935, unix:/tmp/sock
```

If a server cannot be selected, the variable keeps the *name* of the *server group*.

#### \$upstream\_bytes\_received

number of bytes received from an upstream server. Values from several connections are separated by commas like addresses in the  $\$upstream\_addr$  variable.

### \$upstream\_bytes\_sent

number of bytes sent to an upstream server. Values from several connections are separated by commas like addresses in the  $\$upstream \ addr$  variable.

```
$upstream_connect_time
```

time to connect to the upstream server; the time is kept in seconds with millisecond resolution. Times of several connections are separated by commas like addresses in the *\$upstream addr* variable.

```
$upstream_first_byte_time
```

time to receive the first byte of data; the time is kept in seconds with millisecond resolution. Times of several connections are separated by commas like addresses in the *\$upstream addr* variable.

```
$upstream_session_time_<name>
```

session duration in seconds with millisecond resolution. Times of several connections are separated by commas like addresses in the  $\$upstream\_addr$  variable.



#### \$upstream\_sticky\_status

Status of sticky connections.

| 11 11 | Connection routed to upstream without sticky enabled.  |
|-------|--|
| NEW   | Connection without sticky information.   |
| HIT   | Connection with sticky information routed to the desired backend.                                  |
| MISS  | Connection with sticky information routed to the backend selected by the load balancing algorithm. |

Values from multiple connections are separated by commas and colons, similar to addresses in the  $\$upstream\_addr$  variable.

# **Upstream Probe**

The module implements active health probes for  $stream\_upstream$ .

### **Configuration Example**

```
server {
    listen ...;
    # ...
    proxy_pass backend;
    upstream_probe_timeout 1s;
    upstream_probe backend_probe
        port=12345
        interval=5s
        test=$good
        essential
        fails=3
        passes=3
        max_response=512k
        mode=onfail
        "send=data:GET / HTTP/1.0\r\n\r\n";
}
```

### 1 Note

According to RFC 2616 (HTTP/1.1) and RFC 9110 (HTTP Semantics), HTTP headers must be separated by a CRLF sequence (rn) rather than just n.

#### **Directives**

# upstream\_probe (PRO)

Added in version 1.4.0: PRO

| Syntax  | $\verb"upstream_probe" name [port=number] [interval=time] [test=condition]$   |
|---------|---|
|         | $[\texttt{essential} \ [\texttt{persistent}]] \ [\texttt{fails} = number] \ [\texttt{passes} = number] \ [\texttt{max\_response} = size]$ |
|         | $[\verb mode=always   idle   onfail  [udp] [send=string];$  |
| Default | _   |
| Context | server  |



Defines an active health probe for peers within the *upstream* groups that are specified for *proxy\_pass* in the same location context with the *upstream\_probe* directive. Subsequently, Angie regularly probes each peer of the upstream group according to the parameters configured here.

A peer's probe is passed if the request to the peer succeeds, considering all parameter settings of the upstream\_probe directive and the settings that control how upstreams are used by the directive's location context, including the proxy\_next\_upstream directive.

To make use of the probes, the upstream must have a shared memory zone (zone). One upstream may be configured with several probes.

The following parameters are accepted:

| name         | Mandatory name of the probe.   |
|--------------|--|
| port         | Alternative port number for the probe request.   |
| interval     | Interval between probes. By default $-5s$ .  |
| test         | The condition for the probe, defined as a string of variables. If the variables' substitution yields "" or "0", the probe is not passed.   |
| essential    | If set, the initial state of the peer is being checked, so the peer doesn't receive client requests until the probe is passed.   |
| persistent   | Setting this parameter requires enabling essential first; persistent peers that were deemed healthy prior to a configuration reload start receiving requests without being required to pass this probe first.  |
| fails        | Number of subsequent failed probes that renders the peer unhealthy. By default $-1$ .  |
| passes       | Number of subsequent passed probes that renders the peer unhealthy. By default $-1$ .  |
| max_response | Maximum memory size for the response. If a zero value is specified, response waiting is disabled. By default $-256k$ .   |
| mode         | <ul> <li>Probe mode, depending on the peers' health:</li> <li>always — peers are probed regardless of their state;</li> <li>idle — probes affect unhealthy peers and peers where interval has elapsed since the last client request.</li> <li>onfail — only unhealthy peers are probed.</li> <li>By default — always.</li> </ul> |
| udp          | If specified, the UDP protocol is used for probing. By default, TCP is used for probing.   |
| send         | Data sent for the check; this can be a string with the prefix data: or a file name with data (specified absolutely or relative to the /usr/local/angie/ directory).  |

#### Example:

```
upstream backend {
    zone backend 1m;

    server a.example.com;
    server b.example.com;
}

map $upstream_probe_response $good {
    ~200    "1";
    default   "";
}

server {
    listen ...;

# ...
    proxy_pass backend;
```



```
upstream_probe_timeout 1s;

upstream_probe backend_probe
    port=12345
    interval=5s
    test=$good
    essential
    persistent
    fails=3
    passes=3
    max_response=512k
    mode=onfail
    "send=data:GET / HTTP/1.0\r\n\r\n";
}
```

#### Details of probe operation:

- Initially, the peer won't receive client requests until it passes all essential probes configured for it, skipping persistent ones if the configuration was reloaded and the peer was deemed healthy prior to that. If there are no such probes, the peer is considered healthy.
- The peer is considered unhealthy and won't receive client requests, if *any* of the probes configured for it hits fails or the peer reaches *max\_fails*.
- For an unhealthy peer to be considered healthy again, all probes configured for it must reach their respective passes; after that, max fails is also considered.

#### **Built-in Variables**

The stream\_upstream module supports the following built-in variables:

```
$upstream_probe (PRO)
```

Name of the currently active upstream probe.

```
$upstream_probe_response (PRO)
```

Contents of the response received during an active probe configured by upstream probe.

The core stream module implements basic functionality for handling TCP and UDP connections: this includes defining server blocks, traffic routing, configuring proxying, SSL/TLS support, and managing connections for streaming services, such as databases, DNS, and other protocols that operate over TCP and UDP.

The other modules in this section extend this functionality, allowing you to flexibly configure and optimize the stream server for various scenarios and requirements.

When building from the source code, this module isn't built by default; it should be enabled with the --with-stream build option. In packages and images from our repos, the module is included in the build.

### **Configuration Example**

```
worker_processes auto;
error_log /var/log/angie/error.log info;
events {
   worker_connections 1024;
}
```



```
stream {
    upstream backend {
        hash $remote_addr consistent;
        server backend1.example.com:12345 weight=5;
        server 127.0.0.1:12345
                                           max_fails=3 fail_timeout=30s;
        server unix:/tmp/backend3;
    upstream dns {
       server 192.168.0.1:53535;
       server dns.example.com:53;
    }
    server {
        listen 12345;
        proxy_connect_timeout 1s;
        proxy_timeout 3s;
        proxy_pass backend;
    server {
        listen 127.0.0.1:53 udp reuseport;
        proxy_timeout 20s;
        proxy_pass dns;
    server {
        listen [::1]:12345;
        proxy_pass unix:/tmp/stream.socket;
}
```

### **Directives**

#### listen

| Syntax  | listen $address[:port]$ [ssl] [udp] [proxy_protocol] [setfib= $numb$   | ber] |
|---------|--|------|
|         | $[fastopen=number] \qquad [backlog=number] \qquad [rcvbuf=size] \qquad [sndbuf=size]$  | ize] |
|         | [accept_filter=filter] [deferred] [bind] [ipv6only=on   off] [reusepo   so_keepalive=on off [keepidle]:[keepintvl]:[keepcnt]]; | rt]  |
| Default | _  |      |
| Context | server   |      |

Sets the address and port for the socket on which the server will accept connections. It is possible to specify just the port, so Angie listens on all available IPv4 (and IPv6, if enabled) interfaces. The address can also be a hostname, for example:

```
listen 127.0.0.1:12345;
listen *:12345;
listen 12345;  # same as *:12345
listen localhost:12345;
```

IPv6 addresses are specified in square brackets:



```
listen [::1]:12345;
listen [::]:12345;
```

UNIX domain sockets are specified with the unix: prefix:

```
listen unix:/var/run/angie.sock;
```

Port ranges are specified with the first and last port separated by a hyphen:

```
listen 127.0.0.1:12345-12399;
listen 12345-12399;
```

# Important

Different servers must listen on different address:port pairs.

| ssl            | allows specifying that all connections accepted on this port should work in SSL mode.   |
|----------------|---|
| udp            | configures a listening socket for working with datagrams. In order to handle packets from the same address and port in the same session, the <i>reuseport</i> parameter should also be specified. |
| proxy_protocol | allows specifying that all connections accepted on this port should use the PROXY protocol.   |

The listen directive can have several additional parameters specific to socket-related system calls.

| $\mathtt{setfib} = number$ | sets the associated routing table, FIB (the SO_SETFIB option) for the listening |  |  |  |
|----------------------------|---|--|--|--|
|                            | socket. This currently works only on FreeBSD.                                   |  |  |  |
| ${\tt fastopen} = number$  | enables "TCP Fast Open" for the listening socket and limits the maximum length  |  |  |  |
|                            | for the queue of connections that have not yet completed the three-way hand-    |  |  |  |
|                            | shake.  |  |  |  |

# \* Caution

Do not enable this feature unless the server can handle receiving the same SYN packet with data more than once.



| backlog=number              | sets the backlog parameter in the listen() call that limits the maximum length for the queue of pending connections. By default, backlog is set to -1 on FreeBSD, DragonFly BSD, and macOS, and to 511 on other platforms.  |
|-----------------------------|---|
| ${	t rcvbuf} = size$        | sets the receive buffer size (the SO_RCVBUF option) for the listening socket.   |
| $\mathtt{sndbuf} = size$    | sets the send buffer size (the SO_SNDBUF option) for the listening socket.  |
| ${\tt accept\_filter=}filt$ | Sets the name of accept filter (the SO_ACCEPTFILTER option) for the listening socket that filters incoming connections before passing them to accept(). This works only on FreeBSD and NetBSD 5.0+. Acceptable values are dataready and httpready.  |
| deferred                    | instructs to use a deferred accept() (the TCP_DEFER_ACCEPT socket option) on Linux.   |
| bind                        | this parameter instructs to make a separate bind() call for a given address:port pair. The fact is that if there are several listen directives with the same port but different addresses, and one of the listen directives listens on all addresses for the given port (*:port), Angie will bind() only to *:port. It should be noted that the getsockname() system call will be made in this case to determine the address that accepted the connection. If the setfib, fastopen, backlog, rcvbuf, sndbuf, accept_filter, deferred, ipv6only, reuseport, or so_keepalive parameters are used then for a given address:port pair a separate bind() call will always be made. |
| ipv6only=on off             | this parameter determines (via the IPV6_V60NLY socket option) whether an IPv6 socket listening on a wildcard address [::] will accept only IPv6 connections or both IPv6 and IPv4 connections. This parameter is turned on by default. It can only be set once on start.  |
| reuseport                   | this parameter instructs to create an individual listening socket for each worker process (using the SO_REUSEPORT socket option on Linux 3.9+ and DragonFly BSD, or SO_REUSEPORT_LB on FreeBSD 12+), allowing a kernel to distribute incoming connections between worker processes. This currently works only on Linux 3.9+, DragonFly BSD, and FreeBSD 12+.  |

# \* Caution

Inappropriate use of this option may have its security implications.

 $\verb|so_keepalive=on| off | [`keepidle]: [keepintvl]: [keepintvl] ` configures the "TCP keepalive" behavior for the listening socket.$ 

| 1.1 | if this parameter is omitted then the operating system's settings will be in effect for the socket |
|-----|--|
| on  | the $SO\_KEEPALIVE$ option is turned on for the socket   |
| off | the $SO\_KEEPALIVE$ option is turned off for the socket  |

Some operating systems support setting of TCP keepalive parameters on a per-socket basis using the TCP\_KEEPIDLE, TCP\_KEEPINTVL, and TCP\_KEEPCNT socket options. On such systems (currently, Linux 2.4+, NetBSD 5+, and FreeBSD 9.0-STABLE), they can be configured using the keepidle, keepintvl, and keepcnt parameters. One or two parameters may be omitted, in which case the system default setting for the corresponding socket option will be in effect.

For example,

# so\_keepalive=30m::10

will set the idle timeout  $(TCP\_KEEPIDLE)$  to 30 minutes, leave the probe interval  $(TCP\_KEEPINTVL)$  at its system default, and set the probes count  $(TCP\_KEEPINTVL)$  to 10 probes.



# preread\_buffer\_size

| Syntax  | <pre>preread_buffer_size size;</pre> |
|---------|--------------------------------------|
| Default | <pre>preread_buffer_size 16k;</pre>  |
| Context | stream, server                       |

Specifies a size of the *preread* buffer.

#### preread timeout

| Syntax  | preread_timeout timeout;        |
|---------|---------------------------------|
| Default | <pre>preread_timeout 30s;</pre> |
| Context | stream, server                  |

Specifies a timeout of the *preread* phase.

# proxy\_protocol\_timeout

| Syntax  | <pre>proxy_protocol_timeout timeout;</pre> |
|---------|--|
| Default | <pre>proxy_protocol_timeout 30s;</pre>     |
| Context | stream, server                             |

Specifies a *timeout* for reading the PROXY protocol header to complete. If no entire header is transmitted within this time, the connection is closed.

#### resolver

| Syntax  | resolver address [status_zone=zone]; | [valid=time] | [ipv4=on | c | off] | [ipv6=on | off] |
|---------|--------------------------------------|--------------|----------|---|------|----------|------|
| Default | _                                    |              |          |   |      |          |      |
| Context | stream, server, upstream             |              |          |   |      |          |      |

Configures name servers used to resolve names of upstream servers into addresses, for example:

```
resolver 127.0.0.53 [::1]:5353;
```

The address can be specified as a domain name or IP address, with an optional port. If port is not specified, the port 53 is used. Name servers are queried in a round-robin fashion.

By default, Angie caches answers using the TTL value of a response. The optional valid parameter allows overriding it:

| nal valid parameter allows overriding cached entry validity |
|---|
|---|

```
resolver 127.0.0.53 [::1]:5353 valid=30s;
```

By default, Angie will look up both IPv4 and IPv6 addresses while resolving.

| ipv4=off | disables looking up of IPv4 addresses |
|----------|---------------------------------------|
| ipv6=off | disables looking up of IPv6 addresses |



| status_zone | optional parameter; enables the collection of DNS server request and response |
|-------------|---|
|             | metrics $(/status/resolvers/)$ in the specified zone.                         |

# 🗘 Tip

To prevent DNS spoofing, it is recommended configuring DNS servers in a properly secured trusted local network.

# 🗘 Tip

When running in Docker, use its internal DNS server address such as 127.0.0.11.

# resolver\_timeout

| Syntax  | resolver_timeout $time;$ |
|---------|--------------------------|
| Default | resolver_timeout 30s;    |
| Context | stream, server, upstream |

Sets a timeout for name resolution, for example:

```
resolver_timeout 5s;
```

### server

| Syntax  | server { } |
|---------|------------|
| Default | _          |
| Context | stream     |

Sets the configuration for a server.

# server\_name

| Syntax  | server_name name; |
|---------|-------------------|
| Default | server_name "";   |
| Context | server            |

Sets names of a virtual server.

### **A** Attention

In the stream module, server\_name relies on Server Name Indication (SNI) and only works with TLS connections. To use it, you must configure TLS termination or *enable TLS preread* in the corresponding server block.

Example configuration:



```
server {
    listen 443 ssl;
    server_name example.com www.example.com;
    ssl_certificate /etc/angie/cert.pem;
    ssl_certificate_key /etc/angie/key.pem;
}
```

The first name becomes the primary server name.

Server names can include an asterisk (\*) to replace the first or last part of a name:

```
server {
    server_name example.com *.example.com www.example.*;
}
```

These names are called wildcard names.

You can also use regular expressions in server names by preceding the name with a tilde (~):

```
server {
    server_name www.example.com ~^www\d+\.example\.com$;
}
```

Regular expressions may include captures that can be used in other directives:

```
server {
    server_name ~^(www\.)?(.+)$;

    proxy_pass www.$2:12345;
}
```

Named captures in regular expressions create variables that can be used in other directives:

```
server {
    server_name ~^(www\.)?(?<domain>.+)$;

    proxy_pass www.$domain:12345;
}
```

If the directive's parameter is set to \$hostname, the machine's hostname is inserted.

When searching for a virtual server by name, if the name matches more than one of the specified variants (e.g., both a wildcard name and a regular expression match), the first matching variant will be chosen in the following order of priority:

- The exact name
- ullet The longest wildcard name starting with an asterisk, e.g., \*.example.com
- The longest wildcard name ending with an asterisk, e.g., mail.\*
- The first matching regular expression (in order of appearance in the configuration file)

server\_names\_hash\_bucket\_size

```
Syntax server_names_hash_bucket_size size;
Default server_names_hash_bucket_size 32|64|128;
Context stream
```

Sets the bucket size for the server names hash tables. The default value depends on the size of the processor's cache line.



### server\_names\_hash\_max\_size

| Syntax  | server_names_hash_max_size $size;$ |
|---------|------------------------------------|
| Default | server_names_hash_max_size 512;    |
| Context | stream                             |

Sets the maximum size of the server names hash tables.

### status zone

| Syntax  | status_zone zone   key zone=zone[:count]; |
|---------|---|
| Default | _   |
| Context | server                                    |

Allocates a shared memory zone to collect metrics for /status/stream/server\_zones/<zone>.

Multiple server contexts can share the same zone for data collection.

The single-value zone syntax aggregates all metrics for its context in the same shared memory zone:

```
server {
    listen 80;
    server_name *.example.com;

    status_zone single;
    # ...
}
```

The alternative syntax uses the following parameters:

| key              | A string with variables, whose value determines the grouping of connections in the zone. All connections producing identical values after substitution are grouped together. If substitution yields an empty value, metrics aren't updated. |
|------------------|---|
| zone             | The name of the shared memory zone.   |
| count (optional) | The maximum number of separate groups for collecting metrics. If new <i>key</i> values would exceed this limit, they are grouped under <i>zone</i> instead. The default value is 1.   |

In the following example, all connections sharing the same \$server\_addr value are grouped into the host\_zone. Metrics are tracked separately for each unique \$server\_addr until there are 10 metric groups. Once this limit is reached, any additional \$server\_addr values are included under the server\_zone:

```
stream {
    upstream backend {
        server 192.168.0.1:3306;
        server 192.168.0.2:3306;
        # ...
}
    server {
        listen 3306;
        proxy_pass backend;
```



```
status_zone $server_addr zone=server_zone:10;
}
```

The resulting metrics are thus split between individual servers in the API output.

#### stream

| Syntax  | stream { } |
|---------|------------|
| Default | _          |
| Context | main       |

Provides the configuration file context in which the stream server directives are specified.

# tcp\_nodelay

| Syntax  | tcp_nodelay on   off; |
|---------|-----------------------|
| Default | tcp_nodelay on;       |
| Context | stream, server        |

Enables or disables the use of the TCP\_NODELAY option. The option is enabled for both client and proxied server connections.

### variables hash bucket size

| Syntax  | variables_hash_bucket_size $size;$ |
|---------|------------------------------------|
| Default | variables_hash_bucket_size 64;     |
| Context | stream                             |

Sets the bucket size for the variables hash table. The details of setting up hash tables are provided in a separate document.

# variables hash max size

| Syntax  | variables_hash_max_size $size;$          |
|---------|--|
| Default | <pre>variables_hash_max_size 1024;</pre> |
| Context | stream                                   |

Sets the maximum size of the variables hash table. The details of setting up hash tables are provided in a separate document.

# **Built-in Variables**

The *stream core* module supports following variables:

### \$angie\_version

Angie version



#### \$binary\_remote\_addr

client address in a binary form, value's length is always 4 bytes for IPv4 addresses or 16 bytes for IPv6 addresses

#### \$bytes\_received

number of bytes received from a client

#### \$bytes\_sent

number of bytes sent to a client

#### \$connection

connection serial number

### \$hostname

host name

#### \$msec

current time in seconds with the milliseconds resolution

#### \$pid

PID of the worker process

#### \$protocol

protocol used to communicate with the client: TCP or UDP

#### \$proxy\_protocol\_addr

client address from the PROXY protocol header The PROXY protocol must be previously enabled by setting the *proxy protocol* parameter in the *listen* directive.

#### \$proxy\_protocol\_port

client port from the PROXY protocol header The PROXY protocol must be previously enabled by setting the proxy protocol parameter in the listen directive.

#### \$proxy\_protocol\_server\_addr

server address from the PROXY protocol header The PROXY protocol must be previously enabled by setting the *proxy\_protocol* parameter in the *listen* directive.

# \$proxy\_protocol\_server\_port

server port from the PROXY protocol header The PROXY protocol must be previously enabled by setting the *proxy protocol* parameter in the *listen* directive.

### \$proxy\_protocol\_tlv\_<name>

TLV from the PROXY Protocol header. The *name* can be a TLV type or its numeric value. In the latter case, the value is hexadecimal and should be prefixed with  $\theta x$ :

\$proxy protocol tlv alpn \$proxy protocol tlv 0x01



SSL TLVs can also be accessed by TLV type name or its numeric value, both prefixed by ssl:

 $proxy protocol tlv ssl_version \\proxy_protocol_tlv_ssl_0x21$ 

The following TLV type names are supported:

- alpn (0x01) upper layer protocol used over the connection
- authority (0x02) host name value passed by the client
- $unique\_id$  (0x05) unique connection id
- netns (0x30) name of the namespace
- ssl (0x20) binary SSL TLV structure

The following SSL TLV type names are supported:

- $ssl\ version\ (0x21)$  SSL version used in client connection
- ssl cn (0x22) SSL certificate Common Name
- ssl cipher (0x23) name of the used cipher
- $ssl\_sig\_alg~(0x24)$  algorithm used to sign the certificate
- $ssl\_key\_alg~(0x25)$  public-key algorithm

Also, the following special SSL TLV type name is supported:

• ssl\_verify - client SSL certificate verification result, 0 if the client presented a certificate and it was successfully verified, non-zero otherwise.

The PROXY protocol must be previously enabled by setting the proxy\_protocol parameter in the listen directive.

#### \$remote\_addr

client address

### \$remote\_port

client port

#### \$server\_addr

an address of the server which accepted a connection Computing a value of this variable usually requires one system call. To avoid a system call, the *listen* directives must specify addresses and use the bind parameter.

#### \$server\_port

port of the server which accepted a connection

# \$session\_time

session duration in seconds with a milliseconds resolution

#### \$status

session status, can be one of the following:



| 200 | session completed successfully   |
|-----|--|
| 400 | client data could not be parsed, for example, the PROXY protocol header            |
| 403 | access forbidden, for example, when access is limited for certain client addresses |
| 500 | internal server error  |
| 502 | bad gateway, for example, if an upstream server could not be selected or reached   |
| 503 | service unavailable, for example, when access is limited by the number of connec-  |
|     | tions  |

#### \$time\_iso8601

local time in the ISO 8601 standard format

#### \$time\_local

local time in the Common Log Format

### 3.3.4 Mail Module

#### **Auth HTTP**

The module enables subrequest-based authentication by sending an additional HTTP request before processing the main request. If the subrequest returns a 2xx status, the main request proceeds; if it returns 401 or 403, the appropriate error is sent to the user, while any other response triggers a 500 error. This approach is typically used to delegate authentication to external services, unify authentication across applications, or integrate with third-party systems like OAuth or LDAP.

#### **Directives**

### auth\_http

| Syntax  | auth_http uri; |
|---------|----------------|
| Default | _              |
| Context | mail, server   |

Sets the URL of the HTTP authentication server. The protocol is described below.

# auth\_http\_header

| Syntax  | auth_http_header header value; |
|---------|--------------------------------|
| Default | _                              |
| Context | mail, server                   |

Appends the specified header to requests sent to the authentication server. This header can be used as the shared secret to verify that the request comes from Angie. For example:

```
auth_http_header X-Auth-Key "secret_string";
```

# auth\_http\_pass\_client\_cert

| Syntax  | auth_http_pass_client_cert on   off;       |
|---------|--|
| Default | <pre>auth_http_pass_client_cert off;</pre> |
| Context | mail, server                               |



Appends the "Auth-SSL-Cert" header with the *client* certificate in the PEM format (urlencoded) to requests sent to the authentication server.

### auth http timeout

| Syntax  | auth_http_timeout $time$ ;        |
|---------|-----------------------------------|
| Default | <pre>auth_http_timeout 60s;</pre> |
| Context | mail, server                      |

Sets the timeout for communication with the authentication server.

#### **Protocol**

The HTTP protocol is used to communicate with the authentication server. The data in the response body is ignored, the information is passed only in the headers.

### Examples of requests and responses:

### Request:

```
GET /auth HTTP/1.0
Host: localhost
Auth-Method: plain # plain/apop/cram-md5/external
Auth-User: user
Auth-Pass: password
Auth-Protocol: imap # imap/pop3/smtp
Auth-Login-Attempt: 1
Client-IP: 192.0.2.42
Client-Host: client.example.org
```

### Good response:

```
HTTP/1.0 200 OK
Auth-Status: OK
Auth-Server: 198.51.100.1
Auth-Port: 143
```

#### Bad response:

```
HTTP/1.0 200 OK
Auth-Status: Invalid login or password
Auth-Wait: 3
```

If there is no "Auth-Wait" header, an error will be returned and the connection will be closed. The current implementation allocates memory for each authentication attempt. The memory is freed only at the end of a session. Therefore, the number of invalid authentication attempts in a single session must be limited — the server must respond without the "Auth-Wait" header after 10-20 attempts (the attempt number is passed in the "Auth-Login-Attempt" header).

When the APOP or CRAM-MD5 are used, request-response will look as follows:

```
GET /auth HTTP/1.0
Host: localhost
Auth-Method: apop
Auth-User: user
Auth-Salt: <238188073.1163692009@mail.example.com>
Auth-Pass: auth_response
Auth-Protocol: imap
```



Auth-Login-Attempt: 1
Client-IP: 192.0.2.42
Client-Host: client.example.org

#### Good response:

HTTP/1.0 200 OK
Auth-Status: OK
Auth-Server: 198.51.100.1
Auth-Port: 143
Auth-Pass: plain-text-pass

If the "Auth-User" header exists in the response, it overrides the username used to authenticate with the backend

For the SMTP, the response additionally takes into account the "Auth-Error-Code" header — if exists, it is used as a response code in case of an error. Otherwise, the 535 5.7.0 code will be added to the "Auth-Status" header.

For example, if the following response is received from the authentication server:

```
HTTP/1.0 200 OK
Auth-Status: Temporary server problem, try again later
Auth-Error-Code: 451 4.3.0
Auth-Wait: 3
```

then the SMTP client will receive an error

```
451 4.3.0 Temporary server problem, try again later
```

If proxying SMTP does not require authentication, the request will look as follows:

For the SSL/TLS client connection, the "Auth-SSL" header is added, and "Auth-SSL-Verify" will contain the result of client certificate verification, if enabled: SUCCESS, FAILED:reason, and NONE if a certificate was not present.

When the client certificate was present, its details are passed in the following request headers: "Auth-SSL-Subject", "Auth-SSL-Issuer", "Auth-SSL-Serial", and "Auth-SSL-Fingerprint". If  $auth\_http\_pass\_client\_cert$  is enabled, the certificate itself is passed in the "Auth-SSL-Cert" header. The protocol and cipher of the established connection are passed in the "Auth-SSL-Protocol" and "Auth-SSL-Cipher" headers. The request will look as follows:

```
GET /auth HTTP/1.0
Host: localhost
Auth-Method: plain
Auth-User: user
Auth-Pass: password
```



Auth-Protocol: imap Auth-Login-Attempt: 1 Client-IP: 192.0.2.42

Auth-SSL: on

Auth-SSL-Protocol: TLSv1.3

Auth-SSL-Cipher: TLS\_AES\_256\_GCM\_SHA384

Auth-SSL-Verify: SUCCESS

Auth-SSL-Subject: /CN=example.com Auth-SSL-Issuer: /CN=example.com Auth-SSL-Serial: CO7AD56B846B5BFF

Auth-SSL-Fingerprint: 29d6a80a123d13355ed16b4b04605e29cb55a5ad

When the *PROXY protocol* is used, its details are passed in the following request headers: "Proxy-Protocol-Addr", "Proxy-Protocol-Server-Addr", and "Proxy-Protocol-Server-Port".

#### **IMAP**

The module enables IMAP mail protocol support, allowing the server to interact with mail storage systems. It establishes connections to IMAP servers, processes common commands like listing mailboxes and retrieving messages, and provides secure authentication while managing message statuses.

#### **Directives**

### imap auth

| Syntax  | imap_auth method;           |
|---------|-----------------------------|
| Default | <pre>imap_auth plain;</pre> |
| Context | mail, server                |

Sets permitted methods of authentication for IMAP clients. Supported methods are:

| plain    | LOGIN, AUTH=PLAIN   |
|----------|---|
| login    | AUTH=LOGIN  |
| cram-md5 | AUTH=CRAM-MD5. In order for this method to work, the password must be stored unencrypted. |
| external | $\mathrm{AUTH} = \mathrm{EXTERNAL}$   |

Plain text authentication methods (the LOGIN command, AUTH=PLAIN, and AUTH=LOGIN) are always enabled, though if the plain and login methods are not specified, AUTH=PLAIN and AUTH=LOGIN will not be automatically included in *imap capabilities*.

### imap\_capabilities

| Syntax  | <pre>imap_capabilities extension;</pre>               |
|---------|---|
| Default | <pre>imap_capabilities IMAP4 IMAP4rev1 UIDPLUS;</pre> |
| Context | mail, server  |

Sets the IMAP protocol extensions list that is passed to the client in response to the CAPABILITY command. The authentication methods specified in the *imap\_auth* directive and STARTTLS are automatically added to this list depending on the *starttls* directive value.

It makes sense to specify the extensions supported by the IMAP backends to which the clients are proxied (if these extensions are related to commands used after the authentication, when Angie transparently proxies a client connection to the backend).



# imap\_client\_buffer

| Syntax  | <pre>imap_client_buffer size;</pre>  |
|---------|--------------------------------------|
| Default | <pre>imap_client_buffer 4k 8k;</pre> |
| Context | mail, server                         |

Sets the size of the buffer used for reading IMAP commands. By default, the buffer size is equal to one memory page. This is either 4K or 8K, depending on a platform.

#### POP3

The module enables POP3 mail protocol support, allowing the server to download messages from mail servers. It connects to POP3 servers, retrieves message headers and content, provides secure authentication, and manages message statuses like downloaded or deleted.

#### **Directives**

### pop3 auth

| Syntax  | pop3_auth method;           |
|---------|-----------------------------|
| Default | <pre>pop3_auth plain;</pre> |
| Context | mail, server                |

Sets permitted methods of authentication for POP3 clients. Supported methods are:

| plain    | USER/PASS, AUTH PLAIN, AUTH LOGIN   |
|----------|---|
| apop     | APOP. In order for this method to work, the password must be stored unencrypted.          |
| cram-md5 | AUTH=CRAM-MD5. In order for this method to work, the password must be stored unencrypted. |
| external | $\mathrm{AUTH} = \mathrm{EXTERNAL}$   |

Plain text authentication methods (USER/PASS, AUTH PLAIN and AUTH LOGIN) are always enabled, though if the plain method is not specified, AUTH PLAIN and AUTH LOGIN will not be automatically included in pop3 capabilities.

#### pop3 capabilities

| Syntax  | pop3_capabilities extension;     |
|---------|----------------------------------|
| Default | pop3_capabilities TOP USER UIDL; |
| Context | mail, server                     |

Sets the POP3 protocol extensions list that is passed to the client in response to the CAPA command. The authentication methods specified in the *pop3\_auth* directive (SASL extension) and STLS are automatically added to this list depending on the *starttls* directive value.

It makes sense to specify the extensions supported by the POP3 backends to which the clients are proxied (if these extensions are related to commands used after the authentication, when Angie transparently proxies the client connection to the backend).



#### **Proxy**

The module enables support for mail protocols (POP3, IMAP, SMTP), allowing the server to act as a proxy between clients and backend mail servers. It establishes connections to backend servers, handles secure authentication with plain text, SSL/TLS, or STARTTLS, routes client traffic appropriately, and supports flexible authentication and server selection.

#### **Directives**

# proxy\_buffer

| Syntax  | proxy_buffer $size$ ;          |
|---------|--------------------------------|
| Default | <pre>proxy_buffer 4k 8k;</pre> |
| Context | mail, server                   |

Sets the size of the buffer used for proxying. By default, the buffer size is equal to one memory page. Depending on a platform, it is either 4K or 8K.

### proxy\_pass\_error\_message

| Syntax  | proxy_pass_error_message on   off;       |
|---------|--|
| Default | <pre>proxy_pass_error_message off;</pre> |
| Context | mail, server                             |

Indicates whether to pass the error message obtained during the authentication on the backend to the client.

Usually, if the authentication in Angie is a success, the backend cannot return an error. If it nevertheless returns an error, it means some internal error has occurred. In such case the backend message can contain information that should not be shown to the client. However, responding with an error for the correct password is a normal behavior for some POP3 servers. The directive should be enabled in this case.

### proxy protocol

| Syntax  | proxy_protocol on   off;       |
|---------|--------------------------------|
| Default | <pre>proxy_protocol off;</pre> |
| Context | mail, server                   |

Enables the PROXY protocol for connections to a backend.

#### proxy smtp auth

| Syntax  | proxy_smtp_auth on   off;       |
|---------|---------------------------------|
| Default | <pre>proxy_smtp_auth off;</pre> |
| Context | mail, server                    |

Enables or disables user authentication on the SMTP backend using the AUTH command.

If XCLIENT is also enabled, then the XCLIENT command will not send the LOGIN parameter.



### proxy\_timeout

| Syntax  | <pre>proxy_timeout time;</pre> |
|---------|--------------------------------|
| Default | <pre>proxy_timeout 24h;</pre>  |
| Context | mail, server                   |

Sets the timeout between two successive read or write operations on client or proxied server connections. If no data is transmitted within this time, the connection is closed.

#### xclient

| Syntax  | xclient on   off; |
|---------|-------------------|
| Default | xclient on;       |
| Context | mail, server      |

Enables or disables the passing of the XCLIENT command with client parameters when connecting to the SMTP backend.

With XCLIENT, the MTA is able to write client information to the log and apply various limitations based on this data.

If XCLIENT is enabled then Angie passes the following commands when connecting to the backend:

- EHLO with the server name
- XCLIENT
- EHLO or HELO, as passed by the client

If the name *found* by the client IP address points to the same address, it is passed in the NAME parameter of the XCLIENT command. If the name could not be found, points to a different address, or *resolver* is not specified, then [UNAVAILABLE] is passed in the NAME parameter. If an error has occurred in the process of resolving, the [TEMPUNAVAIL] value is used.

If XCLIENT is disabled, Angie passes the EHLO command with the *server name* when connecting to the backend if the client has passed EHLO, or HELO with the server name, otherwise.

#### RealIP

The module is used to change the client address and port to the ones sent in the PROXY protocol header. The PROXY protocol must be previously enabled by setting the proxy\_protocol parameter in the *listen* directive.

### **Configuration Example**

```
listen 110 proxy_protocol;

set_real_ip_from 192.168.1.0/24;
set_real_ip_from 192.168.2.1;
set_real_ip_from 2001:0db8::/32;
```

#### **Directives**

```
set_real_ip_from
```

| Syntax  | set_real_ip_from address   CIDR   unix:; |
|---------|--|
| Default | _  |
| Context | mail, server                             |



Defines trusted addresses that are known to send correct replacement addresses. If the special value unix: is specified, all UNIX domain sockets will be trusted.

#### **SMTP**

The module enables support for the SMTP mail protocol, allowing the server to proxy outgoing email traffic between clients and backend mail servers. It establishes connections to SMTP servers, supports secure authentication with LOGIN or PLAIN methods, provides STARTTLS and SSL/TLS encryption, and routes client requests based on authentication results.

#### **Directives**

#### smtp auth

| Syntax  | smtp_auth method;                 |
|---------|-----------------------------------|
| Default | <pre>smtp_auth plain login;</pre> |
| Context | mail, server                      |

Sets permitted methods of SASL authentication for SMTP clients. Supported methods are:

| plain    | AUTH PLAIN  |
|----------|---|
| login    | AUTH LOGIN  |
| cram-md5 | AUTH CRAM-MD5. In order for this method to work, the password must be stored unencrypted. |
| external | AUTH EXTERNAL   |
| none     | Authentication is not required  |

Plain text authentication methods (AUTH PLAIN and AUTH LOGIN) are always enabled, though if the plain and login methods are not specified, AUTH PLAIN and AUTH LOGIN will not be automatically included in smtp capabilities.

### smtp capabilities

| Syntax  | smtp_capabilities extension; |
|---------|------------------------------|
| Default | _                            |
| Context | mail, server                 |

Sets the SMTP protocol extensions list that is passed to the client in response to the EHLO command. The authentication methods specified in the  $smtp\_auth$  directive and STARTTLS are automatically added to this list depending on the starttls directive value.

It makes sense to specify the extensions supported by the MTA to which the clients are proxied (if these extensions are related to commands used after the authentication, when Angie transparently proxies the client connection to the backend).

# smtp\_client\_buffer

| Syntax  | <pre>smtp_client_buffer size;</pre>  |
|---------|--------------------------------------|
| Default | <pre>smtp_client_buffer 4k 8k;</pre> |
| Context | mail, server                         |

Sets the size of the buffer used for reading SMTP commands. By default, the buffer size is equal to one memory page. This is either 4K or 8K, depending on a platform.



### smtp\_greeting\_delay

| Syntax  | <pre>smtp_greeting_delay time;</pre> |
|---------|--------------------------------------|
| Default | <pre>smtp_greeting_delay 0;</pre>    |
| Context | mail, server                         |

Allows setting a delay before sending an SMTP greeting in order to reject clients who fail to wait for the greeting before sending SMTP commands.

#### **SSL**

The module enables SSL/TLS encryption support for mail proxy protocols (POP3, IMAP, SMTP), allowing secure communication between clients and the server. It provides SSL/TLS encryption for incoming connections, supports STARTTLS upgrades, manages certificates and keys, and controls SSL settings such as ciphers and protocol versions.

When building from the source code, this module isn't built by default; it should be enabled with the --with-mail\_ssl\_module build option.

In packages and images from our repos, the module is included in the build.

```
Important
This module requires the OpenSSL library.
```

### **Configuration Example**

To reduce the processor load it is recommended to

- set the number of worker processes equal to the number of processors,
- enable the *shared* session cache,
- disable the *built-in* session cache,
- and possibly increase the session *lifetime* (by default, 5 minutes):



#### **Directives**

### ssl\_certificate

| Syntax  | ${	t ssl\_certificate} \ file;$ |
|---------|---------------------------------|
| Default | _                               |
| Context | mail, server                    |

Specifies a file with the certificate in the PEM format for the given server. If intermediate certificates should be specified in addition to a primary certificate, they should be specified in the same file in the following order: the primary certificate comes first, then the intermediate certificates. A secret key in the PEM format may be placed in the same file.

This directive can be specified multiple times to load certificates of different types, for example, RSA and ECDSA:

Only OpenSSL 1.0.2 or higher supports separate certificate chains for different certificates. With older versions, only one certificate chain can be used.

The value "data: certificate" can be specified instead of the file, which loads a certificate without using intermediate files.

Note that inappropriate use of this syntax may have its security implications, such as writing secret key data to *error log*.

### ssl certificate key

| Syntax  | ${	t ssl\_certificate\_key} \; file;$ |
|---------|---------------------------------------|
| Default | _                                     |
| Context | mail, server                          |

Specifies a file with the secret key in the PEM format for the given server.

The value "engine: name: id" can be specified instead of the file, which loads a secret key with a specified id from the OpenSSL engine name.

The value "data: key" can be specified instead of the file, which loads a secret key without using intermediate files. Note that inappropriate use of this syntax may have its security implications, such as writing secret key data to error log.

### ssl\_ciphers

| Syntax  | ssl_ciphers ciphers;          |
|---------|-------------------------------|
| Default | ssl_ciphers HIGH:!aNULL:!MD5; |
| Context | mail, server                  |



Specifies the enabled ciphers. The ciphers are specified in the format understood by the OpenSSL library, for example:

```
ssl_ciphers ALL:!aNULL:!EXPORT56:RC4+RSA:+HIGH:+MEDIUM:+LOW:+SSLv2:+EXP;
```

The list of ciphers depends on the version of OpenSSL installed. The full list can be viewed using the openssl ciphers command.

# **A**ttention

The  $ssl\_ciphers$  directive does not configure ciphers for TLS 1.3 when using OpenSSL. To tune TLS 1.3 ciphers with OpenSSL, use the  $ssl\_conf\_command$  directive, which was added to support advanced SSL configuration.

- In LibreSSL, TLS 1.3 ciphers can be configured using ssl\_ciphers.
- In BoringSSL, TLS 1.3 ciphers cannot be configured at all.

### ssl client certificate

| Syntax  | ${\tt ssl\_client\_certificate}$ $file;$ |
|---------|--|
| Default | _  |
| Context | mail, server                             |

Specifies a file with trusted CA certificates in the PEM format used to verify client certificates.

The list of certificates will be sent to clients. If this is not desired, the *ssl\_trusted\_certificate* directive can be used.

# ssl\_conf\_command

| Syntax  | ssl_conf_command name value; |
|---------|------------------------------|
| Default | _                            |
| Context | mail, server                 |

Sets arbitrary OpenSSL configuration commands.

# Important

The directive is supported when using OpenSSL 1.0.2 or higher. To configure TLS 1.3 ciphers with OpenSSL, use the ciphersuites command.

Several  $ssl\_conf\_command$  directives can be specified on the same level:

```
ssl_conf_command Options PrioritizeChaCha;
ssl_conf_command Ciphersuites TLS_CHACHA20_POLY1305_SHA256;
```

These directives are inherited from the previous configuration level if and only if there are no ssl conf command directives defined on the current level.

### \* Caution

Note that configuring OpenSSL directly might result in unexpected behavior.



### ssl\_crl

| Syntax  | ssl_crl file; |
|---------|---------------|
| Default | _             |
| Context | mail, server  |

Specifies a file with revoked certificates (CRL) in the PEM format used to verify client certificates.

#### ssl dhparam

| Syntax  | ${	t ssl\_dhparam} \; file;$ |
|---------|------------------------------|
| Default | _                            |
| Context | mail, server                 |

Specifies a file with DH parameters for DHE ciphers.

### \* Caution

By default no parameters are set, and therefore DHE ciphers will not be used.

### ssl ecdh curve

| Syntax  | ssl_ecdh_curve curve; |
|---------|-----------------------|
| Default | ssl_ecdh_curve auto;  |
| Context | mail, server          |

Specifies a curve for ECDHE ciphers.

# Important

When using OpenSSL 1.0.2 or higher, it is possible to specify multiple curves, for example:

```
ssl_ecdh_curve prime256v1:secp384r1;
```

The special value auto instructs Angie to use a list built into the OpenSSL library when using OpenSSL 1.0.2 or higher, or prime256v1 with older versions.

### Important

When using OpenSSL 1.0.2 or higher, this directive sets the list of curves supported by the server. Thus, in order for ECDSA certificates to work, it is important to include the curves used in the certificates.

## ssl\_password\_file

| Syntax  | ${	t ssl_password_file} \ file;$ |
|---------|----------------------------------|
| Default | _                                |
| Context | mail, server                     |



Specifies a file with passphrases for *secret keys* where each passphrase is specified on a separate line. Passphrases are tried in turn when loading the key.

#### Example:

```
mail {
    ssl_password_file /etc/keys/global.pass;
    ...

    server {
        server_name mail1.example.com;
        ssl_certificate_key /etc/keys/first.key;
    }

    server {
        server_name mail2.example.com;

        # named pipe can also be used instead of a file
        ssl_password_file /etc/keys/fifo;
        ssl_certificate_key /etc/keys/second.key;
    }
}
```

### ssl\_prefer\_server\_ciphers

| Syntax  | ssl_prefer_server_ciphers on   off; |
|---------|-------------------------------------|
| Default | ssl_prefer_server_ciphers off;      |
| Context | mail, server                        |

Specifies that server ciphers should be preferred over client ciphers when the SSLv3 and TLS protocols are used.

### ssl protocols

| Syntax  | ssl_protocols [SSLv2] [SSLv3] [TLSv1] [TLSv1.1] [TLSv1.2] [TLSv1.3]; |
|---------|--|
| Default | ssl_protocols TLSv1.2 TLSv1.3;                                       |
| Context | mail, server   |

Changed in version 1.2.0: TLSv1.3 parameter added to default set.

Enables the specified protocols.

# Important

The TLSv1.1 and TLSv1.2 parameters work only when OpenSSL 1.0.1 or higher is used.

The TLSv1.3 parameters works only when OpenSSL 1.1.1 or higher is used.

### ssl session cache

| Syntax  | ssl_session_cache off   none   [builtin[:size]] [shared:name:size]; |
|---------|---|
| Default | ssl_session_cache none;   |
| Context | mail, server  |



Sets the types and sizes of caches that store session parameters. A cache can be of any of the following types:

| off     | the use of a session cache is strictly prohibited: Angie explicitly tells a client that sessions may not be reused.  |
|---------|--|
| none    | the use of a session cache is gently disallowed: Angie tells a client that sessions may be reused, but does not actually store session parameters in the cache.  |
| builtin | a cache built in OpenSSL; used by one worker process only. The cache size is specified in sessions. If size is not given, it is equal to 20480 sessions. Use of the built-in cache can cause memory fragmentation.   |
| shared  | a cache shared between all worker processes. The cache size is specified in bytes; one megabyte can store about 4000 sessions. Each shared cache should have an arbitrary name. A cache with the same name can be used in several servers. It is also used to automatically generate, store, and periodically rotate TLS session ticket keys unless configured explicitly using the $ssl\_session\_ticket\_key$ directive. |

Both cache types can be used simultaneously, for example:

```
ssl_session_cache builtin:1000 shared:SSL:10m;
```

but using only shared cache without the built-in cache should be more efficient.

## ssl session ticket key

| Syntax  | ssl_session_ticket_key file; |
|---------|------------------------------|
| Default | _                            |
| Context | mail, server                 |

Sets a file with the secret key used to encrypt and decrypt TLS session tickets. The directive is necessary if the same key has to be shared between multiple servers. By default, a randomly generated key is used.

If several keys are specified, only the first key is used to encrypt TLS session tickets. This allows configuring key rotation, for example:

```
ssl_session_ticket_key current.key;
ssl_session_ticket_key previous.key;
```

The file must contain 80 or 48 bytes of random data and can be created using the following command:

```
openssl rand 80 > ticket.key
```

Depending on the file size either AES256 (for 80-byte keys) or AES128 (for 48-byte keys) is used for encryption.

### ssl\_session\_tickets

| Syntax  | ssl_session_tickets on   off; |
|---------|-------------------------------|
| Default | ssl_session_tickets on;       |
| Context | mail, server                  |

Enables or disables session resumption through TLS session tickets.



### ssl\_session\_timeout

| Syntax  | ${\tt ssl\_session\_timeout}\ time;$ |
|---------|--------------------------------------|
| Default | ssl_session_timeout 5m;              |
| Context | mail, server                         |

Specifies a time during which a client may reuse the session parameters.

### ssl trusted certificate

| Syntax  | ${\tt ssl\_trusted\_certificate}$ $file;$ |
|---------|---|
| Default | _   |
| Context | mail, server                              |

Specifies a file with trusted CA certificates in the PEM format used to verify client certificates.

In contrast to the certificate set by  $ssl\_client\_certificate$ , the list of these certificates will not be sent to clients.

# $ssl\_verify\_client$

| Syntax  | ssl_verify_client on   off   optional   optional_no_ca; |
|---------|---|
| Default | ssl_verify_client off;                                  |
| Context | mail, server  |

Enables verification of client certificates. The verification result is passed in the "Auth-SSL-Verify" header of the *authentication* request. If an error occurs during client certificate verification or a client does not provide the required certificate, the connection is closed.

| optional       | requests the client certificate and verifies it if the certificate is present  |
|----------------|--|
| optional_no_ca | requests the client certificate but does not require it to be signed by a trusted CA certificate. This is intended for the use in cases when a service that is external to |
|                | Angie performs the actual certificate verification. The contents of the certificate is accessible through requests <i>sent</i> to the authentication server.               |

# ssl\_verify\_depth

| Syntax  | ssl_verify_depth number; |
|---------|--------------------------|
| Default | ssl_verify_depth 1;      |
| Context | mail, server             |

Sets the verification depth in the client certificates chain.

## starttls

| Syntax  | starttls on   off   only; |
|---------|---------------------------|
| Default | starttls off;             |
| Context | mail, server              |



| on   | allow usage of the STLS command for the POP3 and the STARTTLS command for the IMAP and SMTP; |
|------|--|
| off  | deny usage of the STLS and STARTTLS commands;  |
| only | require preliminary TLS transition.  |

The core mail module implements basic functionality for a mail proxy server: this includes support for SMTP, IMAP, and POP3 protocols, configuring server blocks, mail request routing, user authentication, and SSL/TLS support for securing mail connections.

The other modules in this section extend this functionality, allowing you to flexibly configure and optimize the mail server for various scenarios and requirements.

When building from the source code, this module isn't built by default; it should be enabled with the --with-mail build option. In packages and images from our repos, the module is included in the build.

### **Configuration Example**

```
worker_processes auto;
error_log /var/log/angie/error.log info;
events {
    worker_connections 1024;
mail {
   server_name
                      mail.example.com;
                      localhost:9000/cgi-bin/auth.cgi;
   auth_http
    imap_capabilities IMAP4rev1 UIDPLUS IDLE LITERAL+ QUOTA;
                      plain apop cram-md5;
   pop3_auth
   pop3_capabilities LAST TOP USER PIPELINING UIDL;
                      login plain cram-md5;
   smtp_capabilities "SIZE 10485760" ENHANCEDSTATUSCODES 8BITMIME DSN;
   xclient
                      off;
    server {
        listen
                 25;
        protocol smtp;
   }
    server {
        listen
                110;
        protocol pop3;
        proxy_pass_error_message on;
    server {
        listen
                143;
        protocol imap;
   server {
        listen
                 587;
        protocol smtp;
   }
}
```



#### listen

| Syntax  | listen $address[:port]$ [ssl] [proxy_protocol] [backlog= [rcvbuf= $size$ ] [sndbuf= $size$ ] [bind] [ipv6only=on   off] [rev [so_keepalive=on off [keepidle]:[keepintvl]:[keepcnt]]; | number] |
|---------|--|---------|
| Default | _  |         |
| Context | server   |         |

Sets the address and port for the socket on which the server will accept requests. It is possible to specify just the port, so Angie listens on all available IPv4 (and IPv6, if enabled) interfaces. The address can also be a hostname, for example:

```
listen 127.0.0.1:110;
listen *:110;
listen 110;  # same as *:110
listen localhost:110;
```

IPv6 addresses are specified in square brackets:

```
listen [::1]:110;
listen [::]:110;
```

UNIX domain sockets are specified with the unix: prefix:

```
listen unix:/var/run/angie.sock;
```

# Important

Different servers must listen on different address:port pairs.

| ssl            | allows specifying that all connections accepted on this port should work in SSL mode.  |
|----------------|--|
| proxy_protocol | allows specifying that all connections accepted on this port should use the PROXY protocol. Obtained information is passed to the <i>authentication server</i> and can be used to <i>change the client address</i> . |

The listen directive can have several additional parameters specific to socket-related system calls.



| backlog=number                    | sets the <i>backlog</i> parameter in the listen() call that limits the maximum length for the queue of pending connections. By default, backlog is set to -1 on FreeBSD, DragonFly BSD, and macOS, and to 511 on other platforms.  |
|-----------------------------------|--|
| ${\tt rcvbuf} = size$             | sets the receive buffer size (the SO_RCVBUF option) for the listening socket.  |
| $\mathtt{sndbuf} \texttt{=} size$ | sets the send buffer size (the SO_SNDBUF option) for the listening socket.   |
| bind                              | this parameter instructs to make a separate bind() call for a given address:port pair. The fact is that if there are several listen directives with the same port but different addresses, and one of the listen directives listens on all addresses for the given port (*:port), Angie will bind() only to *:port. It should be noted that the getsockname() system call will be made in this case to determine the address that accepted the connection. If the backlog, rcvbuf, sndbuf, ipv6only, reuseport, or so_keepalive parameters are used then for a given address:port pair a separate bind() call will always be made. |
| ipv6only=on off                   | this parameter determines (via the $IPV6\_V6ONLY$ socket option) whether an IPv6 socket listening on a wildcard address $[::]$ will accept only IPv6 connections or both IPv6 and IPv4 connections. This parameter is turned on by default. It can only be set once on start.  |

 $so_{keepalive=on \mid off \mid [keepidle] : [keepintvl] : [keepcnt] configures the "TCP keepalive" behavior for the listening socket.$ 

| 11  | if this parameter is omitted then the operating system's settings will be in effect for the socket |
|-----|--|
| on  | the SO_KEEPALIVE option is turned on for the socket  |
| off | the SO_KEEPALIVE option is turned off for the socket   |

Some operating systems support setting of TCP keepalive parameters on a per-socket basis using the TCP\_KEEPIDLE, TCP\_KEEPINTVL, and TCP\_KEEPCNT socket options. On such systems (currently, Linux 2.4+, NetBSD 5+, and FreeBSD 9.0-STABLE), they can be configured using the keepidle, keepintvl, and keepcnt parameters. One or two parameters may be omitted, in which case the system default setting for the corresponding socket option will be in effect.

For example,

```
so_keepalive=30m::10
```

will set the idle timeout  $(TCP\_KEEPIDLE)$  to 30 minutes, leave the probe interval  $(TCP\_KEEPINTVL)$  at its system default, and set the probes count  $(TCP\_KEEPINTVL)$  to 10 probes.

#### mail

| Syntax  | mail { } |
|---------|----------|
| Default | _        |
| Context | main     |

Provides the configuration file context in which the mail server directives are specified.

## max commands

Added in version 1.7.0.

| Syntax  | max_commands number;          |
|---------|-------------------------------|
| Default | <pre>max_commands 1000;</pre> |
| Context | mail, server                  |



Sets the maximum number of commands issued during authentication to enhance protection against DoS attacks.

#### max errors

| Syntax  | max_errors number;       |
|---------|--------------------------|
| Default | <pre>max_errors 5;</pre> |
| Context | mail, server             |

Sets the number of protocol errors after which the connection is closed.

#### protocol

| Syntax  | protocol imap   pop3   smtp; |
|---------|------------------------------|
| Default | _                            |
| Context | server                       |

Sets the protocol for a proxied server. Supported protocols are IMAP, POP3, and SMTP.

If the directive is not set, the protocol can be detected automatically based on the well-known port specified in the listen directive:

```
imap: 143, 993
pop3: 110, 995
smtp: 25, 587, 465
```

When building from source, unnecessary protocols can be disabled using the --without-mail\_imap\_module, --without-mail\_pop3\_module, and --without-mail\_smtp\_module build options.

#### resolver

| Syntax  | resolver $address$ [status_zone= $zone$ ]; | $[{\tt valid=} time]$ | [ipv4=on | off] | [ipv6=on | off] |
|---------|--|-----------------------|----------|------|----------|------|
| Default | resolver off;                              |                       |          |      |          |      |
| Context | mail, server                               |                       |          |      |          |      |

Configures name servers used to find the client's hostname to pass it to the *authentication server*, and in the *XCLIENT* command when proxying SMTP. For example:

```
resolver 127.0.0.53 [::1]:5353;
```

The address can be specified as a domain name or IP address, with an optional port. If port is not specified, the port 53 is used. Name servers are queried in a round-robin fashion.

By default, Angie caches answers using the TTL value of a response. The optional valid parameter allows overriding it:

valid optional valid parameter allows overriding cached entry validity

```
resolver 127.0.0.53 [::1]:5353 valid=30s;
```



By default, Angie will look up both IPv4 and IPv6 addresses while resolving.

| ipv4=off    | disables looking up of IPv4 addresses                                |
|-------------|--|
| ipv6=off    | disables looking up of IPv6 addresses                                |
| status_zone | optional parameter, enables statistics collection for specified zone |

# 🗘 Tip

To prevent DNS spoofing, it is recommended configuring DNS servers in a properly secured trusted local network.

# 🗘 Tip

When running in Docker, use its internal DNS server address such as 127.0.0.11.

# resolver\_timeout

| Syntax  | resolver_timeout $time;$ |
|---------|--------------------------|
| Default | resolver_timeout 30s;    |
| Context | mail, server             |

Sets a timeout for DNS operations, for example:

```
resolver_timeout 5s;
```

#### server

| Syntax  | server { } |
|---------|------------|
| Default | _          |
| Context | mail       |

Sets the configuration for a server.

### server\_name

| Syntax  | server_name name;     |
|---------|-----------------------|
| Default | server_name hostname; |
| Context | mail, server          |

Sets the server name that is used:

- in the initial POP3/SMTP server greeting;
- in the salt during the SASL CRAM-MD5 authentication;
- in the EHLO command when connecting to the SMTP backend, if the passing of the XCLIENT command is enabled.

If the directive is not specified, the machine's hostname is used.



#### timeout

| Syntax  | timeout time; |
|---------|---------------|
| Default | timeout 60s;  |
| Context | mail, server  |

Sets the timeout that is used before proxying to the backend starts.

### 3.3.5 Google PerfTools

The module enables profiling of Angie worker processes using Google Performance Tools. The module is intended for Angie developers and allows them to analyze and optimize the performance of the server by providing detailed insights into memory usage, CPU usage, and other performance-related metrics.

When building from the source code, this module isn't built by default; it should be enabled with the --with-google\_perftools\_module build option.

```
Important
```

This module requires the gperftools library.

### **Configuration Example**

```
google_perftools_profiles /var/log/angie/perftools;
```

Profiles will be stored as /var/log/angie/perftools.<worker PID> files.

#### **Directives**

### google perftools profiles

| Syntax  | google_perftools_profiles filename prefix; |
|---------|--|
| Default | _  |
| Context | main                                       |

Sets the prefix of the filename where the profiling information for the worker process Angie will be stored. The worker process ID is appended at the end of the name after a dot, for example: /var/log/angie/perftools.1234.

### 3.3.6 WASM Module

### **WAMR**

The module enables integration with WebAssembly Micro Runtime for executing WASM code, adding a number of runtime-specific directives to the  $wasm\_modules$  context.

In our repositories, the module is built dynamically and is available as a separate package named angie-module-wamr.

#### **Configuration Example**

```
wasm_modules {
   wamr_heap_size 16k;
   wamr_stack_size 16k;
```



```
load fft_transform.wasm id=fft;
}
```

wamr\_heap\_size

| Syntax  | $\verb wamr_heap_size   size;$ |
|---------|--------------------------------|
| Default | <pre>wamr_heap_size 8k;</pre>  |
| Context | wasm_modules                   |

Sets the heap size for an individual module instance.

```
wamr global heap size
```

| Syntax  | wamr_global_heap_size $size;$        |
|---------|--------------------------------------|
| Default | <pre>wamr_global_heap_size 1m;</pre> |
| Context | wasm_modules                         |

Sets the heap size for the entire WAMR runtime.

```
wamr_stack_size
```

| Syntax  | wamr_stack_size $size;$ |
|---------|-------------------------|
| Default | wamr_stack_size 8k;     |
| Context | wasm_modules            |

Sets the stack size for an individual module instance.

#### Wasmtime

The module enables integration with the Wasmtime runtime for executing WASM code, adding a number of runtime-specific directives to the  $wasm\_modules$  context.

In our repositories, the module is built dynamically and is available as a separate package named angie-module-wasmtime.

### **Configuration Example**

```
wasm_modules {
   wasmtime_stack_size 8k;
   wasmtime_enable_wasi on;
   load fft_transform.wasm id=fft;
}
```



### wasmtime enable wasi

| Syntax  | wasmtime_enable_wasi on   off; |
|---------|--------------------------------|
| Default | wasmtime_enable_wasi on;       |
| Context | wasm_modules                   |

Enables or disables the use of WebAssembly System Interface APIs that provide basic POSIX-like functionality to WASM modules running on Angie.



Angie-specific APIs can be whitelisted using the *load* directive.

## wasmtime\_stack\_size

| Syntax  | wasmtime_stack_size $size;$        |
|---------|------------------------------------|
| Default | <pre>wasmtime_stack_size 8k;</pre> |
| Context | wasm_modules                       |

Sets the max\_wasm\_stack value to a specific size, thus limiting the maximum amount of stack space available for executing WASM code.

The core module that implements basic WASM functionality in Angie: this includes support for loading alternative runtimes and WASM modules, as well as configuring their features and limits.

The other modules in this section extend this functionality, allowing you to flexibly configure and optimize the WASM features for various scenarios and requirements.

In our repositories, the module is built dynamically and is available as a separate package named angie-module-wasm.

### **Configuration Example**

```
# These directives load the core functionality
load_module modules/ngx_wasm_module.so;
load_module modules/ngx_wasm_core_module.so;

load_module modules/ngx_wasmtime_module.so;

# Available here: https://git.angie.software/web-server/angie-wasm
load_module modules/ngx_http_wasm_host_module.so;

events {
}

wasm_modules {
    #use wasmtime;
    load ngx_http_handler.wasm id=handler;
    load ngx_http_vars.wasm id=vars type=reactor;
```



```
http {
    wasm_var vars "ngx:wasi/var-utils#sum-entry" $rvar $arg_a $arg_b $arg_c $arg_d;
    server {
        listen *:8080;
        location / {
            return 200 "sum('$arg_a','$arg_b','$arg_c','$arg_d')=$rvar\n";
        }
        location /wasm {
            client_max_body_size 20M;
            wasm_content handler "ngx:wasi/http-handler-entry#handle-request";
        }
    }
}
```

#### load

| Syntax  | $\texttt{load} \ file \ \texttt{id} = id \ [\texttt{fs} = host\_path: guest\_path] \ [\texttt{api} = api] \ [\texttt{type} = \texttt{command} \   \ \texttt{reactor}]$ |
|---------|--|
| Default | _  |
| Context | wasm_modules   |

Loads a module from a disk file and assigns it an id (required). During loading, the module is verified to ensure it can be instantiated.

The directive has the following parameters:

| fs   | Allows the guest to access a directory on the host. Can be specified multiple times for different directories.  |
|------|---|
| api  | Explicitly restricts the list of APIs available to the module by listing them. If the module attempts to use restricted APIs (i.e. not listed here), an "API not found" error will be returned.  By default, the module has access to all APIs.   |
| type | <ul> <li>Controls the lifecycle of the loaded module.</li> <li>In command mode, the machine runs once and its state is destroyed after execution.</li> <li>In reactor mode, the machine effectively runs indefinitely, allowing multiple code executions. This requires careful memory management: if resources aren't cleaned up, memory leaks can occur.</li> </ul> |



# $wasm\_modules$

| Syntax  | ${\tt wasm\_modules} \ \{ \ \dots \ \};$ |
|---------|--|
| Default | _  |
| Context | main                                     |

A top-level directive that provides the configuration file context in which the WASM directives should be specified. It can contain directives for loading WASM modules and configuring runtime-specific settings.

# 3.3.7 Core Module

CoreResponsible for managing service files, processes, and other Angie modules.

# 3.3.8 HTTP Modules

| HTTP Core functionality for processing HTTP requests and responses, man              | aging the HTTP serve     |
|--|--------------------------|
| Access control based on IPs and CIDR ranges.   |                          |
| ACME Automatic retrieval of SSL certificates using the ACME protocol.                |                          |
| Addition Inserts a predefined snippet before or after the response body.             |                          |
| API RESTful HTTP interface to obtain basic information about the web                 | server and its statistic |
| Auth Basic Basic HTTP authentication for access control based on the username        | e and password.          |
| Auth Request Authorization using a subrequest to an external HTTP service.           |                          |
| AutoIndex Automatic directory listing without an index file.                         |                          |
| Browser (deprecated) Browser identification based on the User-Agent header.          |                          |
| Charset Configuration and conversion of response encoding.                           |                          |
| DAV File management on the server using the WebDAV protocol.                         |                          |
| Empty GIF Serves a one-pixel transparent GIF.  |                          |
| FastCGI Proxying requests to a FastCGI server.                                       |                          |
| FLV Pseudo-streaming of Flash Video (FLV) files.                                     |                          |
| Geo Conversion of IP addresses into predefined variable values.                      |                          |
| GeoIP Retrieves IP address data using geolocation by MaxMind GeoIP data              | abases.                  |
| gRPC Proxying requests to a gRPC server.   |                          |
| GunZIP Decompression of compressed GZip responses for modification or in c           | cases where the client   |
| GZip Compression of responses using the GZip method to save traffic.                 |                          |
| GZip Static Serves static files pre-compressed using the GZip method.                |                          |
| Headers Modification of response header fields.                                      |                          |
| Handles requests using the HTTP/2 protocol.  |                          |
| Handles requests using the HTTP/3 protocol.  |                          |
| Image Filter Image transformation.   |                          |
| Index Configuration of index files that serve requests with a trailing slash (       | /).                      |
| JS Handlers to extend functionality by implementing additional logic in              | njs, a subset of the Ja  |
| Limit Conn Limiting the number of concurrent requests (active connections) to p      | orotect against overload |
| Limit Req Limiting the request rate to protect against overload and password g       | uessing.                 |
| Log Configuration of request logs to track resource access for monitoring            | and analysis.            |
| Map Converts variables based on predefined key-value pairs.                          |                          |
| Memcached Retrieval of responses from a Memcached server.                            |                          |
| Mirroring requests to other servers.   |                          |
| MP4 Pseudo-streaming of MP4 files.   |                          |
| Perl Handlers to extend functionality by implementing additional logic in            | the Perl language.       |
| Prometheus Server metrics in a Prometheus-compatible format for monitoring and       | d statistics collection. |
| Proxy Reverse proxying requests to other HTTP servers.                               |                          |
| Random Index Random selection of an index file for requests with a trailing slash (/ | <b>'</b> ).              |
| RealIP Client address and port identification when running behind another p          | proxy server.            |
| Referer header value validation.   |                          |



# Table 1 – continued from previous page

| Rewrite                  | Conditional URI modification, redirects, variable setting, and configuration selection.      |
|--------------------------|--|
| SCGI                     | Proxying requests to a SCGI server.  |
| Secure Link              | Secure link creation with the ability to restrict access time.                               |
| Slice                    | Splitting requests into multiple sub-requests for better caching of large responses.         |
| Split Clients            | Creating variables for A/B testing, canary releases, sharding, and other scenarios that requ |
| SSI                      | Processing SSI (Server Side Includes) commands in responses.                                 |
| SSL                      | SSL/TLS configuration to handle HTTPS requests.  |
| Stub Status (deprecated) | Global connection and request counters in a text-based format.                               |
| Sub                      | Search and replace snippets in server responses.   |
| Upstream                 | Configures groups of proxied servers for load balancing.                                     |
| $Upstream\ Probe$        | Configures health probes for groups of proxied servers.                                      |
| UserID                   | Issuing and processing unique client identifier cookies for session tracking and analytics.  |
| uWSGI                    | Proxying requests to a uWSGI server.   |
| XSLT                     | XML transformation using XSLT stylesheets.   |

# 3.3.9 Streaming Modules

| Stream          | Basic stream server functionality for balancing TCP and UDP protocols at the L4 level.    |
|-----------------|---|
| Access          | Access control based on IPs and CIDR ranges.  |
| Geo             | Conversion of IP addresses into predefined variable values.                               |
| GeoIP           | Retrieves IP address data using geolocation by MaxMind GeoIP databases.                   |
| JS              | Handlers to extend functionality by implementing additional logic in njs, a subset of the |
|                 | JavaScript language.  |
| Limit           | Limiting the number of concurrent requests (active connections) to protect against over-  |
| Conn            | load.   |
| Log             | Configuration of request logs to track resource access for monitoring and analysis.       |
| Map             | Converts variables based on predefined key-value pairs.                                   |
| MQTT            | Reads the client identifier and username from an MQTT connection before making a          |
| Preread         | load balancing decision.  |
| Pass            | Configures passing accepted connections directly to a configured listening socket.        |
| Proxy           | Configures proxying to other servers.   |
| RDP Pre-        | Reads cookies from an RDP connection before making a load balancing decision.             |
| read            |   |
| RealIP          | Client address and port identification when running behind another proxy server.          |
| Return          | Sends a specified value to the client upon connection without further proxying.           |
| Set             | Sets predefined variable values.  |
| Split           | Creating variables for A/B testing, canary releases, sharding, and other scenarios that   |
| Clients         | require a proportional group split.   |
| SSL             | Terminates SSL/TLS and DTLS protocols.  |
| SSL $Pre-$      | Extracts information from the ClientHello message without terminating SSL/TLS and         |
| read            | before making a load balancing decision.  |
| Upstream        | Configures groups of proxied servers for load balancing.                                  |
| Up-             | Configures health probes for groups of proxied servers.                                   |
| $stream\ Probe$ |   |



# 3.3.10 Mail Modules

| Mail         | Basic mail proxy server functionality.   |
|--------------|--|
| $Auth\ HTTI$ | User authentication and server selection for subsequent proxying using HTTP requests |
|              | to an external server.   |
| IMAP         | Support for the IMAP protocol.   |
| POP3         | Support for the POP3 protocol.   |
| Proxy        | Configures proxying to other servers.  |
| RealIP       | Client address and port identification when running behind another proxy server.     |
| SMTP         | Support for the SMTP protocol.   |
| SSL          | Support for the SSL/TLS and StartTLS protocols.                                      |

# 3.3.11 Google PerfTools Module

Google PerfToo Integration with the Google Performance Tools library for application profiling and performance analysis.

# 3.3.12 WASM Modules

| WASM     | Basic WASM functionality to enable running WASM code in Angie. |
|----------|--|
| WAMR     | Integration with WebAssembly Micro Runtime.                    |
| Wasmtime | Integration with the Wasmtime runtime.                         |



| CHAPTER | 4 |
|---------|---|
|---------|---|

Intellectual Property Rights

The documentation for the software product Angie PRO is the intellectual property of Web Server, LLC. The documentation was created as a result of modifications (revisions) to the documentation for the software product Nginx.



| ۸   | 1 1/14 \ 100   |
|---|--|
| absolute_redirect (http), 279 accept_mutex (core), 16 accept_mutex_delay (core), 16 access_log (http), 120 access_log (stream), 329 acme (http), 26 acme_client (http), 26 acme_client_path (http), 28 acme_dns_port (http), 28 acme_hook (http), 29 add_after_body (http), 31 add_before_body (http), 31 add_header (http), 104 add_trailer (http), 104 add_trailer (http), 31 aio (http), 280 aio_write (http), 281 allow (http), 25 allow (stream), 317 ancient_browser (http), 66 ancient_browser_value (http), 66 api (http), 31 api_config_files (http), 32 auth_basic_user_file (http), 62 auth_delay (http), 281 auth_http, 384 auth_http_header, 384 auth_http_pass_client_cert, 384 auth_http_timeout, 385 auth_request_set (http), 63 auto_redirect (http), 63 auto_redirect (http), 64 autoindex_exact_size (http), 64 autoindex_localtime (http), 64 | C  charset (http), 67  charset_map (http), 67  charset_types (http), 68  chunked_transfer_encoding (http), 282  client_body_buffer_size (http), 282  client_body_in_file_only (http), 283  client_body_temp_path (http), 283  client_body_timeout (http), 283  client_body_timeout (http), 283  client_header_buffer_size (http), 283  client_header_timeout (http), 284  client_max_body_size (http), 284  connection_pool_size (http), 284  connection_pool_size (http), 69  daw_access (http), 69  dav_methods (http), 69  debug_connection (core), 16  debug_points (core), 17  default_type (http), 284  deny (http), 25  deny (stream), 317  directio (http), 284  directio_alignment (http), 285  disable_symlinks (http), 285  E  empty_gif (http), 70  env (core), 17  error_log (core), 18  error_page (http), 286  etag (http), 287  events (core), 18  expires (http), 105 |
| B backup_switch (http), 231 bind_conn (http), 232   | F fastcgi_bind (http), 71 fastcgi_buffer_size (http), 71   |
|   |  |



| $fastcgi\_buffering (http), 71$                   | geoip_country $(stream)$ , 320  |
|---|---|
| ${	t fastcgi\_buffers} \ (http), \ 72$            | $geoip\_org(http), 91$  |
| $	extsf{fastcgi_busy_buffers_size}\ (http),72$    | $geoip\_org (stream), 320$  |
| $fastcgi\_cache\ (http),\ 72$                     | $geoip\_proxy (http), 91$   |
| $fastcgi\_cache\_background\_update\ (http),\ 72$ | $geoip\_proxy\_recursive\ (http),\ 91$                                |
| fastcgi_cache_bypass (http), 73                   | google_perftools_profiles, 404  |
| $fastcgi\_cache\_key~(http), 73$                  | grpc_bind (http), 92  |
| fastcgi_cache_lock $(http)$ , 73                  | grpc_buffer_size (http), 92   |
| fastcgi_cache_lock_age $(http)$ , 73              | grpc_connect_timeout (http), 92                                       |
| fastcgi_cache_lock_timeout (http), 74             | grpc_connection_drop (http), 93                                       |
| fastcgi_cache_max_range_offset (http), 74         | grpc_hide_header (http), 93   |
| fastcgi_cache_methods $(http)$ , 74               | grpc_ignore_headers (http), 93  |
| fastcgi_cache_min_uses $(http)$ , 74              | grpc_intercept_errors (http), 93                                      |
| fastcgi_cache_path (http), 74                     | grpc_next_upstream (http), 94   |
| fastcgi_cache_revalidate (http), 75               | grpc_next_upstream_timeout (http), 94                                 |
| fastcgi_cache_use_stale $(http)$ , 76             | grpc_next_upstream_tries (http), 95                                   |
| fastcgi_cache_valid $(http)$ , 76                 | grpc_pass (http), 95  |
| fastcgi_catch_stderr $(http)$ , 77                | grpc_pass_header (http), 95   |
| fastcgi_connect_timeout (http), 77                | grpc_read_timeout (http), 96  |
| fastcgi_connection_drop $(http)$ , 77             | grpc_send_timeout (http), 96  |
| fastcgi_force_ranges (http), 78                   | grpc_set_header (http), 96  |
| fastegi_hide_header $(http)$ , 78                 | grpc_set_header (http), 96 grpc_socket_keepalive (http), 96           |
|   | grpc_ssl_certificate (http), 97                                       |
| fastcgi_ignore_client_abort (http), 78            | grpc_ssl_certificate (http), 97 grpc_ssl_certificate_cache (http), 97 |
| fastcgi_ignore_headers (http), 78                 | <b>-</b> /-   |
| fastcgi_index (http), 79                          | grpc_ssl_certificate_key (http), 97                                   |
| fastcgi_intercept_errors (http), 79               | grpc_ssl_ciphers (http), 97   |
| fastcgi_keep_conn (http), 79                      | grpc_ssl_conf_command (http), 98                                      |
| fastcgi_limit_rate (http), 79                     | grpc_ssl_crl (http), 98   |
| fastcgi_max_temp_file_size (http), 80             | grpc_ssl_name (http), 98  |
| fastcgi_next_upstream (http), 80                  | grpc_ssl_password_file (http), 99                                     |
| fastcgi_next_upstream_timeout (http), 81          | grpc_ssl_protocols (http), 99   |
| fastcgi_next_upstream_tries (http), 81            | grpc_ssl_server_name (http), 99                                       |
| fastcgi_no_cache (http), 82                       | grpc_ssl_session_reuse (http), 99                                     |
| fastcgi_param (http), 82                          | grpc_ssl_trusted_certificate (http), 99                               |
| $fastcgi_pass(http), 82$                          | grpc_ssl_verify (http), 100   |
| $fastcgi_pass_header(http), 83$                   | grpc_ssl_verify_depth (http), 100                                     |
| $fastcgi_pass_request_body (http), 83$            | gunzip $(http)$ , $100$   |
| $	exttt{fastcgi_pass_request_headers}\ (http),83$ | gunzip_buffers $(http)$ , $100$                                       |
| $	extsf{fastcgi_read_timeout} \ (http),  83$      | gzip(http), 101   |
| fastcgi_request_buffering $(http),84$             | $\texttt{gzip\_buffers}\ (http), 101$                                 |
| ${	t fastcgi\_send\_lowat} \ (http),  84$         | $\texttt{gzip\_comp\_level}\ (http),\ 101$                            |
| ${	t fastcgi\_send\_timeout} \ (http),  84$       | $\texttt{gzip\_disable}\ (http), 102$                                 |
| fastcgi_socket_keepalive $(http),84$              | $gzip_http_version (http), 102$                                       |
| ${	t fastcgi\_split\_path\_info}\ (http),85$      | $\texttt{gzip\_min\_length}\ (http),\ 102$                            |
| $	extsf{fastcgi\_store} \ (http),  85$            | $\texttt{gzip\_proxied}\ (http), 102$                                 |
| $fastcgi\_store\_access\ (http),\ 86$             | $gzip\_static\ (http),\ 104$  |
| $fastcgi\_temp\_file\_write\_size\ (http),\ 86$   | $gzip\_types(http), 103$  |
| $fastcgi\_temp\_path\ (http),\ 86$                | $gzip\_vary(http), 103$   |
| feedback $(http)$ , 233                           |   |
| feedback (stream), 364                            | Н   |
| flv $(http)$ , $87$                               | $\mathtt{hash}\ (http), 234$  |
|   | hash $(stream)$ , 365   |
| G   | http (http), 287  |
| geo (http), 88                                    | http2 (http), 272   |
| geo (stream), 318                                 | http2_body_preread_size (http), 272                                   |
| = ,   |   |
| geoip_city (http), 90                             | http2_chunk_size (http), 272  |
| geoip_city (stream), 320                          | http2_max_concurrent_pushes (http), 272                               |
| $geoip\_country\ (http),\ 90$                     | $http2_max_concurrent_streams\ (http),\ 273$                          |



```
http2_push (http), 273
                                                 js\_set (http), 113
http2_push_preload (http), 273
                                                 js_set (stream), 325
http2_recv_buffer_size (http), 273
                                                 js\_shared\_dict\_zone (http), 114
http3 (http), 275
                                                 js_shared_dict_zone (stream), 326
http3_hq (http), 275
                                                 js_var (http), 115
http3_max_concurrent_streams (http), 275
                                                 js_var (stream), 327
http3_max_table_capacity (http), 275
                                                 Κ
http3_stream_buffer_size (http), 276
                                                 keepalive (http), 235
                                                 keepalive_disable (http), 288
\mathtt{if}\ (http),\ 188
                                                 keepalive_requests (http), 236, 288
if_modified_since (http), 287
                                                 keepalive_time (http), 237, 289
ignore_invalid_headers (http), 287
                                                 keepalive_timeout (http), 237, 289
image_filter(http), 106
image_filter_buffer (http), 107
image_filter_interlace (http), 107
                                                 large_client_header_buffers (http), 289
image_filter_jpeg_quality (http), 107
                                                 least_conn (http), 237
image_filter_sharpen (http), 107
                                                 least_conn (stream), 366
image_filter_transparency (http), 107
                                                 least_time (http), 237
image_filter_webp_quality (http), 107
                                                 least_time (stream), 366
imap_auth, 387
                                                 limit\_conn (http), 115
imap\_capabilities, 387
                                                 limit_conn (stream), 327
imap_client_buffer, 387
                                                 limit_conn_dry_run (http), 116
include (core), 18
                                                 limit_conn_dry_run (stream), 328
index (http), 108
                                                 limit_conn_log_level (http), 116
internal (http), 287
                                                 limit_conn_log_level (stream), 328
ip_hash (http), 234
                                                 limit_conn_status (http), 117
                                                 limit\_conn\_zone (http), 117
J
                                                 limit_conn_zone (stream), 328
js\_access (stream), 322
                                                 limit_except (http), 289
js\_body\_filter(http), 110
                                                 limit_rate (http), 290
                                                 limit_rate_after (http), 290
js\_content (http), 111
js_fetch_buffer_size (http), 111
                                                 limit_req (http), 118
js_fetch_buffer_size (stream), 322
                                                 limit_req_dry_run(http), 118
js_fetch_ciphers (http), 111
                                                 limit_req_log_level (http), 119
js_fetch_ciphers (stream), 323
                                                 limit_req_status\ (http),\ 119
js_fetch_max_response_buffer_size (http), 111
                                                 limit_req_zone (http), 119
js_fetch_max_response_buffer_size (stream),
                                                 lingering_close (http), 291
        323
                                                 lingering_time (http), 291
js\_fetch\_protocols (http), 111
                                                 lingering_timeout (http), 291
js_fetch_protocols (stream), 323
                                                 listen, 400
js\_fetch\_timeout (http), 112
                                                 listen (http), 291
js_fetch_timeout (stream), 323
                                                 listen (stream), 374
js_fetch_trusted_certificate (http), 112
                                                 load, 407
js_fetch_trusted_certificate (stream), 323
                                                 load_module (core), 19
                                                 location (http), 295
js\_fetch\_verify(http), 112
js_fetch_verify (stream), 324
                                                 lock_file (core), 19
js_fetch_verify_depth (http), 112
                                                 log_format (http), 121
                                                 log_format (stream), 330
js_fetch_verify_depth (stream), 324
js_filter (stream), 324
                                                 log_not_found (http), 297
js_header_filter(http), 112
                                                 log\_subrequest (http), 297
js\_import (http), 113
                                                 M
js_import (stream), 324
js_path(http), 113
                                                 mail. 401
js_path (stream), 325
                                                 map (http), 123
js_preload_object (http), 113
                                                 map (stream), 331
js_preload_object (stream), 325
                                                 map_hash_bucket_size (http), 124
js_preread (stream), 325
                                                 map_hash_bucket_size (stream), 332
```



| $\mathtt{map\_hash\_max\_size}\ (http),\ 124$            | preread_timeout (stream), 377                   |
|--|---|
| map_hash_max_size (stream), 332                          | prometheus (http), 154                          |
| master_process (core), 19                                | prometheus_template (http), 154                 |
| max_commands, 401  | protocol, 402                                   |
| max_errors, 402  | $proxy\_bind (http), 156$                       |
| $	exttt{max\_headers} (http), 297$                       | proxy_bind (stream), 335                        |
| $\max\_{ranges\ (http),\ 297}$                           | proxy_buffer, 389                               |
| $memcached\_bind\ (http),\ 125$                          | proxy_buffer_size (http), 157                   |
| $memcached\_buffer\_size\ (http),\ 125$                  | proxy_buffer_size (stream), 335                 |
| ${\tt memcached\_connect\_timeout}\ (http),\ 125$        | $proxy\_buffering (http), 157$                  |
| ${\tt memcached\_gzip\_flag} \ (http), \ 125$            | proxy_buffers (http), 157                       |
| ${\tt memcached\_next\_upstream}\ (http),\ 125$          | $proxy_busy_buffers_size\ (http),\ 158$         |
| ${\tt memcached\_next\_upstream\_timeout}\ (http),\ 126$ | $proxy\_cache\ (http),\ 158$                    |
| ${\tt memcached\_next\_upstream\_tries}~(http),~126$     | $proxy_cache_background_update\ (http),\ 159$   |
| $memcached_pass\ (http),\ 127$                           | $proxy\_cache\_bypass\ (http),\ 159$            |
| $memcached\_read\_timeout (http), 127$                   | $proxy_cache_convert_head (http), 159$          |
| $memcached\_send\_timeout\ (http),\ 127$                 | $proxy_cache_key (http), 159$                   |
| ${\tt memcached\_socket\_keepalive}\ (http),\ 127$       | $proxy\_cache\_lock (http), 160$                |
| $merge\_slashes~(http),~297$                             | $proxy_cache_lock_age (http), 160$              |
| $min_delete_depth\ (http),\ 70$                          | $proxy_cache_lock_timeout (http), 160$          |
| mirror $(http)$ , $128$                                  | proxy_cache_max_range_offset (http), 160        |
| $mirror\_request\_body (http), 128$                      | $proxy_cache_methods(http), 160$                |
| $modern\_browser\ (http),\ 66$                           | $proxy_cache_min_uses (http), 161$              |
| modern_browser_value (http), 66                          | $proxy_cache_path (http), 161$                  |
| mp4 (http), 130  | $proxy_cache_revalidate (http), 163$            |
| $\mathtt{mp4\_buffer\_size}\ (http),\ 130$               | $proxy_cache_use_stale\ (http),\ 163$           |
| $mp4\_limit\_rate\ (http),\ 130$                         | $proxy\_cache\_valid\ (http),\ 163$             |
| $\mathtt{mp4\_limit\_rate\_after}\ (http), 130$          | $proxy\_connect\_timeout (http), 164$           |
| $\mathtt{mp4\_max\_buffer\_size}\ (http), 130$           | $proxy\_connect\_timeout (stream), 335$         |
| $\mathtt{mp4\_start\_key\_frame}\ (http), 131$           | $proxy\_connection\_drop\ (http),\ 164$         |
| mqtt_preread (stream), 333                               | proxy_connection_drop (stream), 336             |
| ${\tt msie\_padding}\ (http),\ 298$                      | $proxy\_cookie\_domain (http), 165$             |
| ${\tt msie\_refresh}\ (http),\ 298$                      | $proxy\_cookie\_flags\ (http),\ 165$            |
| $\mathtt{multi\_accept}\ (core),\ 19$                    | $proxy\_cookie\_path\ (http),\ 166$             |
| ^  | $proxy_download_rate (stream), 336$             |
| O  | $proxy\_force\_ranges\ (http),\ 166$            |
| open_file_cache $(http)$ , 298                           | $proxy_half_close (stream), 336$                |
| open_file_cache_errors (http), 299                       | $proxy_headers_hash_bucket_size\ (http),\ 167$  |
| open_file_cache_min_uses (http), 299                     | $proxy_headers_hash_max_size\ (http),\ 167$     |
| open_file_cache_valid $(http)$ , $299$                   | $proxy_hide_header (http), 167$                 |
| open_log_file_cache (http), 122                          | $proxy_http3_hq (http), 167$                    |
| open_log_file_cache (stream), 330                        | $proxy_http3_max_concurrent_streams$ (http),    |
| $\mathtt{output\_buffers}\ (http),\ 299$                 | 168   |
| override_charset $(http)$ , $68$                         | $proxy_http3_max_table_capacity\ (http),\ 168$  |
| _  | proxy_http3_stream_buffer_size $(http)$ , $168$ |
| P  | $proxy_http_version\ (http),\ 167$              |
| pass $(stream)$ , 334                                    | $proxy\_ignore\_client\_abort\ (http),\ 168$    |
| pcre_jit (core), 20                                      | $proxy\_ignore\_headers\ (http),\ 169$          |
| perl(http), 132  | $proxy_intercept_errors (http), 169$            |
| perl_modules (http), 133                                 | $proxy_limit_rate\ (http),\ 169$                |
| perl_require $(http)$ , 133                              | $proxy_max_temp_file_size\ (http),\ 169$        |
| perl_set (http), 133                                     | $proxy\_method\ (http),\ 170$                   |
| pid (core), 20   | $proxy_next_upstream\ (http),\ 170$             |
| pop3_auth, 388   | proxy_next_upstream (stream), 337               |
| pop3_capabilities, 388                                   | $proxy_next_upstream_timeout(http), 171$        |
| port_in_redirect (http), 299                             | proxy_next_upstream_timeout (stream), 337       |
| postpone_output $(http)$ , 300                           | $proxy_next_upstream_tries\ (http),\ 171$       |
| preread_buffer_size (stream), 376                        | $proxy_next_upstream_tries\ (stream),\ 337$     |
|  |   |



```
{\tt proxy\_no\_cache}~(http),~171
                                                proxy_store_access (http), 183
proxy_pass(http), 172
                                                proxy_temp_file_write_size (http), 183
                                                proxy_temp_path (http), 184
proxy_pass (stream), 337
proxy_pass_error_message, 389
                                                proxy_timeout, 389
proxy_pass_header(http), 173
                                                proxy_timeout (stream), 343
\verb"proxy_pass_request_body" (http), 173
                                                proxy_upload_rate (stream), 343
proxy_pass_request_headers (http), 173
                                                Q
proxy_pass_trailers (http), 174
proxy_protocol, 389
                                                queue (http), 238
proxy_protocol (stream), 338
                                                quic_active_connection_id_limit (http), 276
proxy_protocol_timeout (stream), 377
                                                quic_bpf (http), 276
proxy_quic_active_connection_id_limit
                                                quic_gso (http), 276
        (http), 174
                                                quic_host_key(http), 276
proxy_quic_gso(http), 174
                                                quic_retry (http), 277
proxy_quic_host_key(http), 174
proxy_read_timeout (http), 174
proxy_redirect (http), 175
                                                random (http), 238
proxy_request_buffering (http), 176
                                                random (stream), 367
proxy_requests (stream), 338
                                                random_index (http), 185
proxy_responses (stream), 338
                                                rdp_preread (stream), 344
proxy_send_lowat (http), 176
                                                read_ahead (http), 300
proxy_send_timeout (http), 176
                                                real_ip_header (http), 185
proxy_set_body(http), 177
                                                real_ip_recursive (http), 186
proxy_set_header (http), 177
                                                recursive_error_pages (http), 300
proxy_smtp_auth, 389
                                                referer_hash_bucket_size (http), 186
proxy_socket_keepalive(http), 177
                                                referer_hash_max_size (http), 186
proxy_socket_keepalive (stream), 338
                                                request_pool_size (http), 300
proxy_ssl (stream), 338
                                                reset_timedout_connection (http), 300
proxy_ssl_certificate(http), 178
                                                resolver, 402
proxy_ssl_certificate (stream), 339
                                                resolver (http), 301
proxy_ssl_certificate_cache (http), 178
                                                resolver (stream), 377
proxy_ssl_certificate_key(http), 179
                                                resolver\_timeout, 403
proxy_ssl_certificate_key (stream), 339
                                                resolver_timeout (http), 302
proxy_ssl\_ciphers (http), 179
                                                resolver_timeout (stream), 378
proxy_ssl_ciphers (stream), 339
                                                response_time_factor (http), 239
proxy_ssl_conf_command(http), 179
                                                response_time_factor (stream), 367
proxy_ssl_conf_command (stream), 340
                                                return (http), 189
proxy_ssl_crl (http), 180
                                                return (stream), 345
proxy_ssl_crl\ (stream),\ 340
                                                rewrite (http), 189
proxy_ssl_name\ (http),\ 180
                                                rewrite_log (http), 190
proxy_ssl_name\ (stream),\ 340
                                                root (http), 302
proxy_ssl_ntls(http), 180
proxy_ssl_ntls (stream), 341
                                                S
proxy_ssl_password_file (http), 181
                                                satisfy (http), 302
proxy_ssl_password_file (stream), 341
                                                scgi_bind(http), 192
proxy_ssl_protocols (http), 181
                                                scgi_buffer_size(http), 192
proxy_ssl_protocols (stream), 341
                                                scgi_buffering (http), 192
proxy_ssl_server_name(http), 181
                                                scgi_buffers (http), 193
proxy_ssl_server_name (stream), 342
                                                scgi_busy_buffers_size (http), 193
proxy_ssl_session_reuse (http), 181
                                                scgi\_cache(http), 193
proxy_ssl_session_reuse (stream), 342
                                                scgi_cache_background_update (http), 194
proxy_ssl_trusted_certificate (http), 182
                                                scgi_cache_bypass(http), 194
proxy_ssl_trusted_certificate (stream), 342
                                                scgi_cache_key (http), 194
proxy_ssl_verify(http), 182
                                                scgi_cache_lock (http), 194
proxy_ssl_verify (stream), 342
                                                scgi_cache_lock_age (http), 194
proxy_ssl_verify_depth(http), 182
                                                scgi\_cache\_lock\_timeout (http), 195
proxy_ssl_verify_depth (stream), 342
                                                scgi_cache_max_range_offset (http), 195
proxy_store (http), 182
                                                scgi_cache_methods (http), 195
```



```
scgi_cache_min_uses(http), 195
                                                 smtp_client_buffer, 391
{\tt scgi\_cache\_path}\ (http),\ 195
                                                 smtp_greeting_delay, 391
scgi_cache_revalidate (http), 197
                                                 source_charset (http), 68
scgi_cache_use_stale(http), 197
                                                 split_clients (http), 210
scgi_cache_valid(http), 197
                                                 split_clients (stream), 346
\verb|scgi_connect_timeout| (http), \, 198
                                                 ssi(http), 210
                                                 ssi_last_modified (http), 211
scgi_connection_drop (http), 198
scgi_force_ranges (http), 199
                                                 ssi_min_file_chunk (http), 211
scgi_hide_header (http), 199
                                                 ssi_silent_errors (http), 211
scgi_ignore_client_abort (http), 199
                                                 ssi\_types(http), 211
                                                 ssi_value_length (http), 211
scgi_ignore_headers (http), 199
scgi_intercept_errors (http), 200
                                                 ssl_alpn (stream), 348
scgi_limit_rate (http), 200
                                                 ssl_buffer_size (http), 215
scgi_max_temp_file_size(http), 200
                                                 ssl\_certificate, 393
                                                 ssl\_certificate (http), 216
scgi_next_upstream\ (http),\ 201
scgi_next_upstream_timeout (http), 201
                                                 ssl_certificate (stream), 348
scgi_next_upstream_tries (http), 202
                                                 ssl_certificate_cache (http), 217
scgi_no_cache (http), 202
                                                 ssl_certificate_key, 393
scgi_param (http), 202
                                                 ssl_certificate_key(http), 217
                                                 ssl_certificate_key (stream), 349
scgi_pass(http), 203
\verb|scgi_pass_header|(http),\,203|
                                                 ssl_ciphers, 393
scgi_pass_request_body (http), 203
                                                 ssl\_ciphers (http), 218
scgi_pass_request_headers (http), 203
                                                 ssl_ciphers (stream), 349
scgi_read_timeout (http), 203
                                                 ssl_client_certificate, 394
scgi\_request\_buffering (http), 204
                                                 {\tt ssl\_client\_certificate}\ (http),\,218
                                                 ssl_client_certificate (stream), 349
scgi_send_timeout (http), 204
scgi_socket_keepalive (http), 204
                                                 ssl_conf_command, 394
scgi_store (http), 204
                                                 ssl\_conf\_command (http), 218
                                                 ssl_conf_command (stream), 350
scgi\_store\_access (http), 205
scgi_temp_file_write_size (http), 206
                                                 ssl_crl, 394
                                                 ssl_crl(http), 219
scgi_temp_path(http), 206
secure_link (http), 206
                                                 ssl_crl (stream), 350
secure\_link\_md5 (http), 207
                                                 ssl_dhparam, 395
secure_link_secret (http), 208
                                                 ssl_dhparam (http), 219
send_lowat(http), 302
                                                 ssl_dhparam (stream), 350
send\_timeout (http), 303
                                                 ssl_early_data (http), 219
sendfile (http), 303
                                                 ssl_early_data (stream), 351
sendfile_max_chunk (http), 303
                                                 ssl_ecdh_curve, 395
                                                 ssl\_ecdh\_curve~(http),~220
server, 403
server (http), 239, 303
                                                 ssl_ecdh_curve (stream), 351
server (stream), 361, 378
                                                 ssl_engine (core), 20
server_name, 403
                                                 ssl_handshake_timeout (stream), 351
server_name (http), 304
                                                 ssl_ntls(http), 220
server_name (stream), 378
                                                 ssl_ntls (stream), 352
                                                 ssl_object_cache_inheritable (core), 21
server_name_in_redirect (http), 305
server_names_hash_bucket_size (http), 306
                                                 ssl_ocsp(http), 221
server_names_hash_bucket_size (stream), 379
                                                 ssl_ocsp (stream), 351
                                                 ssl_ocsp_cache\ (http),\ 221
server_names_hash_max_size (http), 306
server_names_hash_max_size (stream), 379
                                                 ssl_ocsp_cache (stream), 352
server_tokens (http), 306
                                                 ssl_ocsp_responder (http), 221
set (http), 190
                                                 ssl_ocsp_responder (stream), 352
set (stream), 346
                                                 ssl_password_file, 395
set_real_ip_from, 390
                                                 ssl_password_file (http), 221
set_real_ip_from (http), 185
                                                 ssl_password_file (stream), 353
set_real_ip_from (stream), 345
                                                 ssl_prefer_server_ciphers, 396
                                                 ssl\_prefer\_server\_ciphers (http), 222
slice (http), 209
smtp_auth, 391
                                                 ssl_prefer_server_ciphers (stream), 353
smtp_capabilities, 391
                                                 ssl_preread (stream), 360
```



| ssl_protocols, 396   | timer_resolution $(core)$ , 21                |
|--|---|
| $ssl\_protocols\ (http),\ 222$   | $try_files (http), 308$                       |
| ssl_protocols (stream), 353  | types $(http)$ , 310                          |
| $ssl\_reject\_handshake (http), 222$   | types_hash_bucket_size $(http)$ , 310         |
| ssl_session_cache, 396   | types_hash_max_size $(http)$ , 311            |
| $ssl_session_cache\ (http),\ 223$  |   |
| ssl_session_cache (stream), 354  | U   |
| ssl_session_ticket_key, 397  | underscores_in_headers (http), 311            |
| ssl_session_ticket_key (http), 223   | uninitialized_variable_warn (http), 190       |
| ssl_session_ticket_key (stream), 354   | upstream (http), 246                          |
| ssl_session_tickets, 397   | upstream (stream), 361                        |
| ssl_session_tickets (http), 224  | upstream_probe $(http)$ , 249                 |
| ssl_session_tickets (stream), 355  | upstream_probe (stream), 371                  |
| ssl_session_timeout, 397   | upstream_probe_timeout (stream), 343          |
| ssl_session_timeout (http), 224  | use (core), 22                                |
| ssl_session_timeout (stream), 355  | user (core), 22                               |
| ssl_stapling $(http)$ , 224  | userid $(http)$ , 251                         |
| ssl_stapling (stream), 355   |   |
| ssl_stapling_file (http), 224  | userid_domain (http), 251                     |
| ssl_stapling_file (stream), 355  | userid_expires (http), 252                    |
| ssl_stapling_responder (http), 225   | userid_flags (http), 252                      |
| ssl_stapling_responder (mtp), 355 ssl_stapling_responder (stream), 355       | userid_mark (http), 252                       |
| ssl_stapling_responder ( $stream$ ), 355 ssl_stapling_verify ( $http$ ), 225 | userid_name (http), 252                       |
| ssl_stapling_verify (mtp), 226 ssl_stapling_verify (stream), 356             | userid_p3p (http), 252                        |
|  | userid_path (http), 253                       |
| ssl_trusted_certificate, 398   | userid_service (http), 253                    |
| ssl_trusted_certificate (http), 225  | uwsgi_bind (http), 254                        |
| ssl_trusted_certificate (stream), 356  | uwsgi_buffer_size (http), 254                 |
| ssl_verify_client, 398   | uwsgi_buffering $(http)$ , 254                |
| ssl_verify_client (http), 225  | uwsgi_buffers (http), 255                     |
| ssl_verify_client (stream), 356  | uwsgi_busy_buffers_size $(http)$ , $255$      |
| ssl_verify_depth, 398  | uwsgi_cache $(http)$ , $255$                  |
| ssl_verify_depth (http), 226   | $uwsgi_cache_background_update\ (http),\ 255$ |
| ssl_verify_depth (stream), 356   | uwsgi_cache_bypass $(http), 255$              |
| starttls, 398  | uwsgi_cache_key $(http)$ , $256$              |
| state $(http)$ , 241   | uwsgi_cache_lock $(http)$ , 256               |
| state (stream), 363  | uwsgi_cache_lock_age $(http)$ , $256$         |
| $status\_zone (http), 306$   | uwsgi_cache_lock_timeout $(http)$ , 256       |
| status_zone (stream), 380  | uwsgi_cache_max_range_offset $(http)$ , 257   |
| sticky (http), 242   | uwsgi_cache_methods $(http)$ , 257            |
| sticky (stream), 367   | uwsgi_cache_min_uses $(http)$ , 257           |
| $sticky\_secret(http), 245$  | uwsgi_cache_path $(http)$ , 257               |
| $sticky\_secret (stream), 369$   | uwsgi_cache_revalidate $(http)$ , 258         |
| $sticky\_strict\ (http),\ 245$   | uwsgi_cache_use_stale (http), 258             |
| sticky_strict (stream), 369  | uwsgi_cache_valid (http), 259                 |
| stream (stream), 381   | uwsgi_connect_timeout (http), 260             |
| $stub\_status\ (http),\ 229$   | uwsgi_connection_drop (http), 260             |
| $sub_filter(http), 230$  | uwsgi_force_ranges (http), 260                |
| $sub_filter_last_modified (http), 230$                                       | uwsgi_hide_header (http), 260                 |
| $sub_filter_once(http), 230$   | uwsgi_ignore_client_abort (http), 261         |
| sub_filter_types (http), 231   | uwsgi_ignore_headers (http), 261              |
| subrequest_output_buffer_size (http), 307                                    | uwsgi_intercept_errors (http), 261            |
|  | uwsgi_limit_rate (http), 261                  |
| T  | uwsgi_max_temp_file_size (http), 262          |
| tcp_nodelay (http), 308  | uwsgi_modifier1 (http), 262                   |
| tcp_nodelay (stream), 381  | uwsgi_modifier2 (http), 262                   |
| tcp_noquelay (stream), 381<br>tcp_nopush (http), 308                         | uwsgi_next_upstream (http), 263               |
| thread_pool $(core)$ , 21  | = ' ' - ' '                                   |
| = ' ' ' '  | uwsgi_next_upstream_timeout (http), 263       |
| timeout, 403   | uwsgi_next_upstream_tries $(http)$ , 264      |



```
uwsgi_no_cache (http), 264
uwsgi_param (http), 264
uwsgi_pass (http), 264
uwsgi_pass_header (http), 265
uwsgi_pass_request_body (http), 265
uwsgi_pass_request_headers (http), 265
uwsgi_read_timeout (http), 265
uwsgi_request_buffering (http), 266
uwsgi_send_timeout (http), 266
uwsgi_socket_keepalive (http), 266
uwsgi_ssl_certificate (http), 266
uwsgi_ssl_certificate_cache (http), 267
uwsgi_ssl_certificate_key (http), 267
uwsgi_ssl_ciphers (http), 267
uwsgi_ssl_conf_command (http), 268
uwsgi_ssl_crl (http), 268
uwsgi_ssl_name (http), 268
uwsgi_ssl_password_file (http), 269
uwsgi_ssl_protocols (http), 269
uwsgi_ssl_server_name (http), 269
{\tt uwsgi\_ssl\_session\_reuse}\ (http),\ 269
uwsgi_ssl_trusted_certificate (http), 269
uwsgi_ssl_verify (http), 269
uwsgi_ssl_verify_depth(http), 270
uwsgi_store(http), 270
uwsgi_store_access (http), 271
uwsgi_temp_file_write_size (http), 271
uwsgi_temp_path (http), 271
valid_referers (http), 187
variables_hash_bucket_size (http), 311
variables_hash_bucket_size (stream), 381
variables_hash_max_size (http), 311
variables_hash_max_size (stream), 381
W
wamr_global_heap_size, 405
wamr_heap_size, 405
wamr_stack_size, 405
wasm_modules, 407
wasmtime_enable_wasi, 406
wasmtime_stack_size, 406
worker_aio_requests (core), 22
worker_connections (core), 22
worker\_cpu\_affinity(core), 22
worker_priority (core), 23
worker_processes (core), 23
worker_rlimit_core (core), 24
worker_rlimit_nofile (core), 24
worker_shutdown_timeout (core), 24
working_directory (core), 24
X
xclient, 390
xml_entities (http), 278
xslt_last_modified (http), 278
xslt_param(http), 278
```

```
{\tt xslt\_string\_param}\;(http),\,278
xslt_stylesheet (http), 279
xslt\_types(http), 279
Ζ
zone (http), 246
zone (stream), 364
```