

































































































































































































































































































































































































## memcached\_gzip\_flag

SYNTAX: **memcached\_gzip\_flag** *flag*;

DEFAULT —

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.3.6.

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: **memcached\_next\_upstream** *error* | *timeout* | *invalid\_response* | *not\_found* | *off* ...;

DEFAULT *error timeout*

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\_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.

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. The cases of *error*, *timeout* and *invalid\_response* are always considered unsuccessful attempts, even if they are not specified in the directive. The case of *not\_found* is never considered an unsuccessful attempt.

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 0

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.7.5.

Limits the time during which a request can be passed to the [next server](#). The 0 value turns off this limitation.

### **memcached\_next\_upstream\_tries**

SYNTAX: **memcached\_next\_upstream\_tries** *number*;

DEFAULT 0

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.7.5.

Limits the number of possible tries for passing a request to the [next server](#). The 0 value turns off this limitation.

### **memcached\_pass**

SYNTAX: **memcached\_pass** *address*;

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](#).

### **memcached\_read\_timeout**

SYNTAX: **memcached\_read\_timeout** *time*;

DEFAULT 60s

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: **memcached\_send\_timeout** *time*;

DEFAULT 60s

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 off

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.15.6.

Configures the “TCP keepalive” behavior for outgoing connections to a memcached server. By default, the operating system’s settings are in effect for the socket. If the directive is set to the value “on”, the `SO_KEEPALIVE` socket option is turned on for the socket.

## **2.32.4 Embedded Variables**

*\$memcached\_key*

Defines a key for obtaining response from a memcached server.

## 2.33 Module ngx\_http\_mirror\_module

2.33.1 Summary	195
2.33.2 Example Configuration	195
2.33.3 Directives	195
mirror	195
mirror_request_body	195

### 2.33.1 Summary

The `ngx_http_mirror_module` module (1.13.4) implements mirroring of an original request by creating background mirror subrequests. Responses to mirror subrequests are ignored.

### 2.33.2 Example Configuration

```
location / {
    mirror /mirror;
    proxy_pass http://backend;
}

location = /mirror {
    internal;
    proxy_pass http://test_backend$request_uri;
}
```

### 2.33.3 Directives

#### mirror

SYNTAX: **mirror** *uri* | off;  
 DEFAULT off  
 CONTEXT: http, server, location

Sets the URI to which an original request will be mirrored. Several mirrors can be specified on the same level.

#### mirror\_request\_body

SYNTAX: **mirror\_request\_body** on | off;  
 DEFAULT 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 [proxy\\_request\\_buffering](#), [fastcgi\\_request\\_buffering](#), [scgi\\_request\\_buffering](#), and [uwsgi\\_request\\_buffering](#) 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;
}
```

## 2.34 Module ngx\_http\_mp4\_module

2.34.1 Summary	197
2.34.2 Example Configuration	198
2.34.3 Directives	198
mp4	198
mp4_buffer_size	198
mp4_max_buffer_size	198
mp4_limit_rate	199
mp4_limit_rate_after	199

### 2.34.1 Summary

The `ngx_http_mp4_module` module provides pseudo-streaming server-side support for MP4 files. Such files typically have the `.mp4`, `.m4v`, or `.m4a` filename extensions.

Pseudo-streaming works in alliance with a compatible Flash 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 nginx to simply start sending back the file contents. If the metadata are located at the end of the file, nginx 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 nginx do this on every such request.

The module also supports the `end` argument of an HTTP request (1.5.13) 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, nginx 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 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.

This module is not built by default, it should be enabled with the `--with-http_mp4_module` configuration parameter.

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 [ngx-http\\_flv\\_module](#) module.

## 2.34.2 Example Configuration

```
location /video/ {
    mp4;
    mp4_buffer_size      1m;
    mp4_max_buffer_size  5m;
    mp4_limit_rate       on;
    mp4_limit_rate_after 30s;
}
```

## 2.34.3 Directives

### mp4

SYNTAX: **mp4**;

DEFAULT —

CONTEXT: location

Turns on module processing in a surrounding location.

### mp4\_buffer\_size

SYNTAX: **mp4\_buffer\_size** *size*;

DEFAULT 512K

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 10M

CONTEXT: http, server, location

During metadata processing, a larger buffer may become necessary. Its size cannot exceed the specified *size*, or else nginx 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 off

CONTEXT: http, server, location

Limits the rate of response transmission to a client. The rate is limited based on the average bitrate of the MP4 file served. To calculate the rate, the bitrate is multiplied by the specified *factor*. The special value “on” corresponds to the factor of 1.1. The special value “off” 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.

This directive is available as part of our [commercial subscription](#).

### **mp4\_limit\_rate\_after**

SYNTAX: **mp4\_limit\_rate\_after** *time*;

DEFAULT 60s

CONTEXT: http, server, location

Sets the initial amount of media data (measured in playback time) after which the further transmission of the response to a client will be rate limited.

This directive is available as part of our [commercial subscription](#).

## 2.35 Module ngx\_http\_perl\_module

2.35.1	Summary	200
2.35.2	Known Issues	200
2.35.3	Example Configuration	201
2.35.4	Directives	201
	perl	201
	perl_modules	202
	perl_require	202
	perl_set	202
2.35.5	Calling Perl from SSI	202
2.35.6	The \$r Request Object Methods	202

### 2.35.1 Summary

The `ngx_http_perl_module` module is used to implement location and variable handlers in Perl and insert Perl calls into SSI.

This module is not built by default, it should be enabled with the `--with-http_perl_module` configuration parameter.

This module requires [Perl](#) version 5.6.1 or higher. The C compiler should be compatible with the one used to build Perl.

### 2.35.2 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 `-Dusymalloc=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:usymalloc
usemultiplicity='define';
usymalloc='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.

### 2.35.3 Example Configuration

```
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 "";
        }

    ';

    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__
```

### 2.35.4 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.

## 2.35.5 Calling Perl from SSI

An SSI command calling Perl has the following format:

```
<!--# perl sub="module::function" arg="parameter1" arg="parameter2" ...
-->
```

## 2.35.6 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, nginx 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 nginx 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.

Redirections to named locations are currently not supported.

`$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, nginx continues to process other requests. After the specified time has elapsed, nginx 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.

## 2.36 Module ngx\_http\_proxy\_module

2.36.1	Summary	207
2.36.2	Example Configuration	207
2.36.3	Directives	208
	proxy_bind	208
	proxy_buffer_size	208
	proxy_buffering	208
	proxy_buffers	209
	proxy_busy_buffers_size	209
	proxy_cache	209
	proxy_cache_background_update	209
	proxy_cache_bypass	210
	proxy_cache_convert_head	210
	proxy_cache_key	210
	proxy_cache_lock	210
	proxy_cache_lock_age	211
	proxy_cache_lock_timeout	211
	proxy_cache_max_range_offset	211
	proxy_cache_methods	211
	proxy_cache_min_uses	212
	proxy_cache_path	212
	proxy_cache_purge	214
	proxy_cache_revalidate	214
	proxy_cache_use_stale	214
	proxy_cache_valid	215
	proxy_connect_timeout	216
	proxy_cookie_domain	216
	proxy_cookie_path	217
	proxy_force_ranges	218
	proxy_headers_hash_bucket_size	218
	proxy_headers_hash_max_size	218
	proxy_hide_header	218
	proxy_http_version	219
	proxy_ignore_client_abort	219
	proxy_ignore_headers	219
	proxy_intercept_errors	219
	proxy_limit_rate	220
	proxy_max_temp_file_size	220
	proxy_method	220
	proxy_next_upstream	221
	proxy_next_upstream_timeout	222
	proxy_next_upstream_tries	222
	proxy_no_cache	222
	proxy_pass	222
	proxy_pass_header	224

proxy_pass_request_body . . . . .	224
proxy_pass_request_headers . . . . .	224
proxy_read_timeout . . . . .	225
proxy_redirect . . . . .	225
proxy_request_buffering . . . . .	226
proxy_send_lowat . . . . .	227
proxy_send_timeout . . . . .	227
proxy_set_body . . . . .	227
proxy_set_header . . . . .	227
proxy_socket_keepalive . . . . .	228
proxy_ssl_certificate . . . . .	228
proxy_ssl_certificate_key . . . . .	229
proxy_ssl_ciphers . . . . .	229
proxy_ssl_crl . . . . .	229
proxy_ssl_name . . . . .	229
proxy_ssl_password_file . . . . .	230
proxy_ssl_protocols . . . . .	230
proxy_ssl_server_name . . . . .	230
proxy_ssl_session_reuse . . . . .	230
proxy_ssl_trusted_certificate . . . . .	230
proxy_ssl_verify . . . . .	231
proxy_ssl_verify_depth . . . . .	231
proxy_store . . . . .	231
proxy_store_access . . . . .	232
proxy_temp_file_write_size . . . . .	232
proxy_temp_path . . . . .	233
2.36.4 Embedded Variables . . . . .	233

### 2.36.1 Summary

The `ngx_http_proxy_module` module allows passing requests to another server.

### 2.36.2 Example Configuration

```
location / {
    proxy_pass      http://localhost:8000;
    proxy_set_header Host      $host;
    proxy_set_header X-Real-IP $remote_addr;
}
```

## 2.36.3 Directives

### proxy\_bind

SYNTAX: **proxy\_bind** *address* [transparent] | off;

DEFAULT —

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 0.8.22.

Makes outgoing connections to a proxied server originate from the specified local IP address with an optional port (1.11.2). Parameter value can contain variables (1.3.12). The special value `off` (1.3.12) 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 (1.11.0) 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 nginx worker processes with the [superuser](#) privileges. On Linux it is not required (1.13.8) as if the `transparent` parameter is specified, worker processes inherit the `CAP_NET_RAW` capability from the master process. It is also necessary to configure kernel routing table to intercept network traffic from the proxied server.

### proxy\_buffer\_size

SYNTAX: **proxy\_buffer\_size** *size*;

DEFAULT 4k | 8k

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 on

CONTEXT: http, server, location

Enables or disables buffering of responses from the proxied server.

When buffering is enabled, nginx 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.

When buffering is disabled, the response is passed to a client synchronously, immediately as it is received. nginx will not try to read the whole response from the proxied server. The maximum size of the data that nginx can receive from the server at a time is set by the [proxy\\_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 [proxy\\_ignore\\_headers](#) directive.

### proxy\_buffers

SYNTAX: **proxy\_buffers** *number size*;  
DEFAULT 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 8k|16k  
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;  
DEFAULT 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 (1.7.9). The `off` parameter disables caching inherited from the previous configuration level.

### proxy\_cache\_background\_update

SYNTAX: **proxy\_cache\_background\_update** on | off;  
DEFAULT off  
CONTEXT: http, server, location  
THIS DIRECTIVE APPEARED IN VERSION 1.11.10.

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.

### proxy\_cache\_bypass

SYNTAX: **proxy\_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:

```
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 on

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.9.7.

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 *\$scheme\$proxy\_host\$request\_uri*

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 off

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.1.12.

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 5s

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.7.8.

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: **proxy\_cache\_lock\_timeout** *time*;

DEFAULT 5s

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.1.12.

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.

Before 1.7.8, the response could be cached.

### proxy\_cache\_max\_range\_offset

SYNTAX: **proxy\_cache\_max\_range\_offset** *number*;

DEFAULT —

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.11.6.

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 GET HEAD

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 0.7.59.

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 1

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* [inactive=*time*]  
 [max\_size=*size*] [manager\_files=*number*] [manager\_sleep=*time*]  
 [manager\_threshold=*time*] [loader\_files=*number*]  
 [loader\_sleep=*time*] [loader\_threshold=*time*]  
 [purger=on|off] [purger\_files=*number*] [purger\_sleep=*time*]  
 [purger\_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

```
proxy_cache_path /data/nginx/cache levels=1:2 keys_zone=one:10m;
```

file names in a cache will look like this:

```
/data/nginx/cache/c/29/b7f54b2df7773722d382f4809d65029c
```

A cached response is first written to a temporary file, and then the file is renamed. Starting from version 0.8.9, 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 (1.7.10). 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. If the value is set to *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.

As part of [commercial subscription](#), the shared memory zone also stores extended cache [information](#), thus, it is required to specify a larger zone size for the same number of keys. For example, one megabyte zone can store about 4 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.

The special “cache manager” process monitors the maximum cache size set by the `max_size` parameter. When this size is exceeded, it removes the least recently used data. The data is removed in iterations configured by `manager_files`, `manager_threshold`, and `manager_sleep` parameters (1.11.5). During one iteration no more than `manager_files` items are deleted (by default, 100). The duration of one iteration is limited by the `manager_threshold` parameter (by default, 200 milliseconds). Between iterations, a pause configured by the `manager_sleep` parameter (by default, 50 milliseconds) is made.

A minute after the start 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. During one iteration no more than `loader_files` items are loaded (by default, 100). Besides, the duration of one iteration is limited by the `loader_threshold` parameter (by default, 200 milliseconds). Between iterations, a pause configured by the `loader_sleep` parameter (by default, 50 milliseconds) is made.

Additionally, the following parameters are available as part of our [commercial subscription](#):

`purger=on|off`

Instructs whether cache entries that match a [wildcard key](#) will be removed from the disk by the cache purger (1.7.12). Setting the parameter to `on` (default is `off`) will activate the “cache purger” process that permanently iterates through all cache entries and deletes the entries that match the wildcard key.

`purger_files=number`

Sets the number of items that will be scanned during one iteration (1.7.12). By default, `purger_files` is set to 10.

`purger_threshold=number`

Sets the duration of one iteration (1.7.12). By default, `purger_threshold` is set to 50 milliseconds.

`purger_sleep=number`

Sets a pause between iterations (1.7.12). By default, `purger_sleep` is set to 50 milliseconds.

In versions 1.7.3, 1.7.7, and 1.11.10 cache header format has been changed. Previously cached responses will be considered invalid after upgrading to a newer nginx version.

## proxy\_cache\_purge

SYNTAX: **proxy\_cache\_purge** string ...;

DEFAULT —

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.5.7.

Defines conditions under which the request will be considered a cache purge request. If at least one value of the string parameters is not empty and is not equal to “0” then the cache entry with a corresponding [cache key](#) is removed. The result of successful operation is indicated by returning the 204 No Content response.

If the [cache key](#) of a purge request ends with an asterisk (“\*”), all cache entries matching the wildcard key will be removed from the cache. However, these entries will remain on the disk until they are deleted for either [inactivity](#), or processed by the [cache purger](#) (1.7.12), or a client attempts to access them.

Example configuration:

```
proxy_cache_path /data/nginx/cache keys_zone=cache_zone:10m;

map $request_method $purge_method {
    PURGE    1;
    default  0;
}

server {
    ...
    location / {
        proxy_pass http://backend;
        proxy_cache cache_zone;
        proxy_cache_key $uri;
        proxy_cache_purge $purge_method;
    }
}
```

This functionality is available as part of our [commercial subscription](#).

## proxy\_cache\_revalidate

SYNTAX: **proxy\_cache\_revalidate** on | off;

DEFAULT off

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.5.7.

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 off

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.

The `error` parameter also permits using a stale cached response if a proxied server to process a request cannot be selected.

Additionally, the `updating` 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 (1.11.10). This has lower priority than using the directive parameters.

- 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 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:

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

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 (1.7.7). If the header includes the `Vary` field with another value, such a response will be cached taking into account the corresponding request header fields (1.7.7).

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 60s

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\_cookie\_domain

SYNTAX: **proxy\_cookie\_domain** *off*;

SYNTAX: **proxy\_cookie\_domain** *domain replacement*;

DEFAULT *off*

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.1.15.

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 from the “~” 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;
```

There could be several `proxy_cookie_domain` directives:

```
proxy_cookie_domain localhost example.org;  
proxy_cookie_domain ~\.[a-z]+\.[a-z]+)$ $1;
```

The `off` parameter cancels the effect of all `proxy_cookie_domain` directives on the current level:

```
proxy_cookie_domain off;  
proxy_cookie_domain localhost example.org;  
proxy_cookie_domain www.example.org example.org;
```

## proxy\_cookie\_path

SYNTAX: **proxy\_cookie\_path** *off*;

SYNTAX: **proxy\_cookie\_path** *path replacement*;

DEFAULT `off`

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.1.15.

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 from the “~” symbol for a case-sensitive matching, or from 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;
```

There could be several `proxy_cookie_path` directives:

```
proxy_cookie_path /one/ /;  
proxy_cookie_path / /two/;
```

The `off` parameter cancels the effect of all `proxy_cookie_path` directives on the current level:

```
proxy_cookie_path off;
proxy_cookie_path /two/ /;
proxy_cookie_path ~*^/user/([^/]+) /u/$1;
```

### **proxy\_force\_ranges**

SYNTAX: **proxy\_force\_ranges** on | off;

DEFAULT off

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.7.7.

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: **proxy\_headers\_hash\_bucket\_size** *size*;

DEFAULT 64

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 in a separate [document](#).

### **proxy\_headers\_hash\_max\_size**

SYNTAX: **proxy\_headers\_hash\_max\_size** *size*;

DEFAULT 512

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 in a separate [document](#).

### **proxy\_hide\_header**

SYNTAX: **proxy\_hide\_header** *field*;

DEFAULT —

CONTEXT: http, server, location

By default, nginx 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;

DEFAULT 1.0

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.1.4.

Sets the HTTP protocol version for proxying. By default, version 1.0 is used. Version 1.1 is recommended for use with [keepalive](#) connections and [NTLM authentication](#).

## proxy\_ignore\_client\_abort

SYNTAX: **proxy\_ignore\_client\_abort** on | off;

DEFAULT off

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 (1.1.6), X-Accel-Buffering (1.1.6), X-Accel-Charset (1.1.6), Expires, Cache-Control, Set-Cookie (0.8.44), and Vary (1.7.7).

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 off

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 nginx for processing with the [error\\_page](#) directive.

### proxy\_limit\_rate

SYNTAX: **proxy\_limit\_rate** *rate*;  
DEFAULT 0  
CONTEXT: http, server, location  
THIS DIRECTIVE APPEARED IN VERSION 1.7.7.

Limits the speed of reading the response from the proxied server. The *rate* is specified in bytes per second. The zero value disables rate limiting. The limit is set per a request, and so if nginx 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: **proxy\_max\_temp\_file\_size** *size*;  
DEFAULT 1024m  
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.

The zero value disables buffering of responses to temporary files.

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 (1.11.6).

## 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 error timeout

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\_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 (1.11.13);

### 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 (1.9.13); 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. The cases of `error`, `timeout` and `invalid_header` are always considered unsuccessful attempts, even if they are not specified in the directive. The cases of `http_500`, `http_502`, `http_503`, `http_504`, and `http_429` are considered unsuccessful attempts only if they

are specified in the directive. The cases of `http_403` and `http_404` are 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 0  
CONTEXT: http, server, location  
THIS DIRECTIVE APPEARED IN VERSION 1.7.5.

Limits the time during which a request can be passed to the [next server](#). The 0 value turns off this limitation.

### **proxy\_next\_upstream\_tries**

SYNTAX: **proxy\_next\_upstream\_tries** *number*;  
DEFAULT 0  
CONTEXT: http, server, location  
THIS DIRECTIVE APPEARED IN VERSION 1.7.5.

Limits the number of possible tries for passing a request to the [next server](#). The 0 value 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: **proxy\_pass** *URL*;  
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;  
}
```

Before version 1.1.12, if `proxy_pass` is specified without a URI, the original request URI might be passed instead of the changed URI in some cases.

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 and is supported since version 1.3.13.

### `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 on

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 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\_read\_timeout

SYNTAX: **proxy\_read\_timeout** *time*;

DEFAULT 60s

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: **proxy\_redirect** default;

SYNTAX: **proxy\_redirect** off;

SYNTAX: **proxy\_redirect** *redirect replacement*;

DEFAULT default

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/;
```

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 (1.1.11) variables:

```
proxy_redirect http://$proxy_host:8000/ /;
```

The directive can be specified (1.1.11) using regular expressions. In this case, *redirect* should either start with the “~” 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_redirect ~^(http://[^\:]+\):\d+(/.+)$ $1$2;
proxy_redirect ~*/user/([^\:]+)/(.+)$ http://$1.example.com/$2;
```

There could be several `proxy_redirect` directives:

```
proxy_redirect default;
proxy_redirect http://localhost:8000/ /;
proxy_redirect http://www.example.com/ /;
```

The `off` parameter cancels the effect of all `proxy_redirect` directives on the current level:

```
proxy_redirect off;
proxy_redirect default;
proxy_redirect http://localhost:8000/ /;
proxy_redirect http://www.example.com/ /;
```

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: **proxy\_request\_buffering** on | off;

DEFAULT on

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.7.11.

Enables or disables buffering of a client request body.

When buffering is enabled, the entire request body is [read](#) from the client before sending the request to a proxied server.

When buffering is disabled, the request body is sent to the proxied server immediately as it is received. In this case, the request cannot be passed to the [next server](#) if nginx 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: **proxy\_send\_lowat** *size*;

DEFAULT 0

CONTEXT: http, server, location

If the directive is set to a non-zero value, nginx 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*.

This directive is ignored on Linux, Solaris, and Windows.

### proxy\_send\_timeout

SYNTAX: **proxy\_send\_timeout** *time*;

DEFAULT 60s

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 Host \$proxy\_host

DEFAULT Connection close

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 level if and only if there are no `proxy_set_header` directives defined on the current level. By default, only two fields are redefined:

```
proxy_set_header Host      $proxy_host;
proxy_set_header Connection close;
```

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: **proxy\_socket\_keepalive** on | off;

DEFAULT off

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.15.6.

Configures the “TCP keepalive” behavior for outgoing connections to a proxied server. By default, the operating system’s settings are in effect for the socket. If the directive is set to the value “on”, the `SO_KEEPALIVE` socket option is turned on for the socket.

### proxy\_ssl\_certificate

SYNTAX: **proxy\_ssl\_certificate** *file*;

DEFAULT —

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.7.8.

Specifies a *file* with the certificate in the PEM format used for authentication to a proxied HTTPS server.

### proxy\_ssl\_certificate\_key

SYNTAX: **proxy\_ssl\_certificate\_key** *file*;

DEFAULT —

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.7.8.

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* (1.7.9), which loads a secret key with a specified *id* from the OpenSSL engine *name*.

### proxy\_ssl\_ciphers

SYNTAX: **proxy\_ssl\_ciphers** *ciphers*;

DEFAULT DEFAULT

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.5.6.

Specifies the enabled ciphers for requests to a proxied HTTPS server. The ciphers are specified in the format understood by the OpenSSL library.

The full list can be viewed using the “`openssl ciphers`” command.

### proxy\_ssl\_crl

SYNTAX: **proxy\_ssl\_crl** *file*;

DEFAULT —

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.7.0.

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 `$proxy_host`

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.7.0.

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\_password\_file

SYNTAX: **proxy\_ssl\_password\_file** *file*;

DEFAULT —

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.7.8.

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 TLSv1 TLSv1.1 TLSv1.2

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.5.6.

Enables the specified protocols for requests to a proxied HTTPS server.

### proxy\_ssl\_server\_name

SYNTAX: **proxy\_ssl\_server\_name** on | off;

DEFAULT off

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.7.0.

Enables or disables passing of the server name through [TLS Server Name Indication extension](#) (SNI, RFC 6066) when establishing a connection with the proxied HTTPS server.

### proxy\_ssl\_session\_reuse

SYNTAX: **proxy\_ssl\_session\_reuse** on | off;

DEFAULT on

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: **proxy\_ssl\_trusted\_certificate** *file*;

DEFAULT —

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.7.0.

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 off  
CONTEXT: http, server, location  
THIS DIRECTIVE APPEARED IN VERSION 1.7.0.

Enables or disables verification of the proxied HTTPS server certificate.

## proxy\_ssl\_verify\_depth

SYNTAX: **proxy\_ssl\_verify\_depth** *number*;  
DEFAULT 1  
CONTEXT: http, server, location  
THIS DIRECTIVE APPEARED IN VERSION 1.7.0.

Sets the verification depth in the proxied HTTPS server certificates chain.

## proxy\_store

SYNTAX: **proxy\_store** on | off | *string*;  
DEFAULT off  
CONTEXT: http, server, location

Enables saving of files to a disk. The `on` parameter saves files with paths corresponding to the directives [alias](#) or [root](#). The `off` parameter disables saving of files. In addition, 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. 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 [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;
    error_page          404 = /fetch$uri;
}

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;
    error_page    404 = @fetch;
}

location @fetch {
    internal;

    proxy_pass     http://backend;
    proxy_store    on;
    proxy_store_access user:rw group:rw all:r;
    proxy_temp_path /data/temp;

    root          /data/www;
}
```

### proxy\_store\_access

SYNTAX: **proxy\_store\_access** *users:permissions ...*;

DEFAULT `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 `8k|16k`

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: **proxy\_temp\_path** *path* [*level1* [*level2* [*level3*]]];

DEFAULT `proxy_temp`

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/nginx/proxy_temp 1 2;
```

a temporary file might look like this:

```
/spool/nginx/proxy_temp/7/45/00000123457
```

See also the `use_temp_path` parameter of the [proxy\\_cache\\_path](#) directive.

### 2.36.4 Embedded Variables

The `ngx_http_proxy_module` module supports embedded 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\_forwarded\_for* variable is equal to the *\$remote\_addr* variable.

## 2.37 Module ngx\_http\_random\_index\_module

2.37.1 Summary	234
2.37.2 Example Configuration	234
2.37.3 Directives	234
random_index	234

### 2.37.1 Summary

The `ngx_http_random_index_module` 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 [ngx\\_http\\_index\\_module](#) module.

This module is not built by default, it should be enabled with the `--with-http_random_index_module` configuration parameter.

### 2.37.2 Example Configuration

```
location / {
    random_index on;
}
```

### 2.37.3 Directives

#### random\_index

SYNTAX: **random\_index** on | off;

DEFAULT off

CONTEXT: location

Enables or disables module processing in a surrounding location.

## 2.38 Module ngx\_http\_realip\_module

2.38.1 Summary	235
2.38.2 Example Configuration	235
2.38.3 Directives	235
set_real_ip_from	235
real_ip_header	235
real_ip_recursive	236
2.38.4 Embedded Variables	236

### 2.38.1 Summary

The `ngx_http_realip_module` module is used to change the client address and optional port to those sent in the specified header field.

This module is not built by default, it should be enabled with the `--with-http_realip_module` configuration parameter.

### 2.38.2 Example Configuration

```
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;
```

### 2.38.3 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 (1.13.1).

IPv6 addresses are supported starting from versions 1.3.0 and 1.2.1.

#### real\_ip\_header

SYNTAX: **real\_ip\_header** *field* | `X-Real-IP` | `X-Forwarded-For` | `proxy_protocol`;

DEFAULT `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 (1.11.0). The address and port should be specified according to [RFC 3986](#).

The `proxy_protocol` parameter (1.5.12) 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.

### **realip\_recursive**

SYNTAX: **realip\_recursive** on | off;

DEFAULT off

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSIONS 1.3.0 AND 1.2.1.

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 [realip\\_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.

## **2.38.4 Embedded Variables**

*\$realip\_remote\_addr*

keeps the original client address (1.9.7)

*\$realip\_remote\_port*

keeps the original client port (1.11.0)

## 2.39 Module ngx\_http\_referer\_module

2.39.1 Summary	237
2.39.2 Example Configuration	237
2.39.3 Directives	237
<a href="#">referer_hash_bucket_size</a>	237
<a href="#">referer_hash_max_size</a>	237
<a href="#">valid_referers</a>	238
2.39.4 Embedded Variables	238

### 2.39.1 Summary

The `ngx_http_referer_module` 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.

### 2.39.2 Example Configuration

```
valid_referers none blocked server_names
               *.example.com example.* www.example.org/galleries/
               ~\.google\.;

if ($invalid_referer) {
    return 403;
}
```

### 2.39.3 Directives

#### `referer_hash_bucket_size`

SYNTAX: **referer\_hash\_bucket\_size** *size*;

DEFAULT 64

CONTEXT: server, location

THIS DIRECTIVE APPEARED IN VERSION 1.0.5.

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 2048

CONTEXT: server, location

THIS DIRECTIVE APPEARED IN VERSION 1.0.5.

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 embedded `$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 string

defines a server name and an optional URI prefix. A server name can have an “`*`” at the beginning or end. During the checking, the server’s port in the `Referer` field is ignored;

regular 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:

```
valid_referers none blocked server_names
                *.example.com example.* www.example.org/galleries/
                ~\.google\.;
```

## 2.39.4 Embedded Variables

`$invalid_referer`

Empty string, if the `Referer` request header field value is considered [valid](#), otherwise “1”.

## 2.40 Module ngx\_http\_rewrite\_module

2.40.1 Summary	239
2.40.2 Directives	239
break	239
if	240
return	241
rewrite	241
rewrite_log	243
set	243
uninitialized_variable_warn	243
2.40.3 Internal Implementation	243

### 2.40.1 Summary

The `ngx_http_rewrite_module` 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.

### 2.40.2 Directives

#### break

SYNTAX: **break**;

DEFAULT —

CONTEXT: server, location, if

Stops processing the current set of `ngx_http_rewrite_module` 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”;

Before version 1.0.1, any string starting with “0” was considered a false value.

- 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;
}

```

A value of the *\$invalid\_referer* embedded variable is set by the [valid\\_referers](#) directive.

## return

SYNTAX: **return** *code* [*text*];

SYNTAX: **return** *code* *URL*;

SYNTAX: **return** *URL*;

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.

Starting from version 0.8.42, 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.

Only the following codes could be returned before version 0.7.51: 204, 400, 402 — 406, 408, 410, 411, 413, 416, and 500 — 504.

The code 307 was not treated as a redirect until versions 1.1.16 and 1.0.13.

The code 308 was not treated as a redirect until version 1.13.0.

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 `ngx_http_rewrite_module` directives and starts a search for a new location matching the changed URI;

**break**

stops processing the current set of `ngx_http_rewrite_module` directives as with the `break` 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 nginx 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 off

CONTEXT: http, server, location, if

Enables or disables logging of `ngx_http_rewrite_module` 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 on

CONTEXT: http, server, location, if

Controls whether warnings about uninitialized variables are logged.

### 2.40.3 Internal Implementation

The `ngx_http_rewrite_module` 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 "/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 `ngx_http_rewrite_module` 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
```

## 2.41 Module ngx\_http\_scgi\_module

2.41.1	Summary	246
2.41.2	Example Configuration	246
2.41.3	Directives	246
	scgi_bind	246
	scgi_buffer_size	246
	scgi_buffering	247
	scgi_buffers	247
	scgi_busy_buffers_size	247
	scgi_cache	248
	scgi_cache_background_update	248
	scgi_cache_bypass	248
	scgi_cache_key	248
	scgi_cache_lock	248
	scgi_cache_lock_age	249
	scgi_cache_lock_timeout	249
	scgi_cache_max_range_offset	249
	scgi_cache_methods	249
	scgi_cache_min_uses	250
	scgi_cache_path	250
	scgi_cache_purge	252
	scgi_cache_revalidate	252
	scgi_cache_use_stale	252
	scgi_cache_valid	253
	scgi_connect_timeout	254
	scgi_force_ranges	254
	scgi_hide_header	254
	scgi_ignore_client_abort	255
	scgi_ignore_headers	255
	scgi_intercept_errors	255
	scgi_limit_rate	255
	scgi_max_temp_file_size	256
	scgi_next_upstream	256
	scgi_next_upstream_timeout	257
	scgi_next_upstream_tries	257
	scgi_no_cache	258
	scgi_param	258
	scgi_pass	258
	scgi_pass_header	259
	scgi_pass_request_body	259
	scgi_pass_request_headers	259
	scgi_read_timeout	259
	scgi_request_buffering	260
	scgi_send_timeout	260
	scgi_socket_keepalive	260

<a href="#">scgi_store</a>	260
<a href="#">scgi_store_access</a>	261
<a href="#">scgi_temp_file_write_size</a>	261
<a href="#">scgi_temp_path</a>	262

### 2.41.1 Summary

The `ngx_http_scgi_module` module allows passing requests to an SCGI server.

### 2.41.2 Example Configuration

```
location / {
    include    scgi_params;
    scgi_pass  localhost:9000;
}
```

### 2.41.3 Directives

#### `scgi_bind`

SYNTAX: **scgi\_bind** *address* [transparent] | off;

DEFAULT —

CONTEXT: http, server, location

Makes outgoing connections to an SCGI server originate from the specified local IP address with an optional port (1.11.2). Parameter value can contain variables (1.3.12). The special value `off` (1.3.12) 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 (1.11.0) allows outgoing connections to an 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 nginx worker processes with the [superuser](#) privileges. On Linux it is not required (1.13.8) as if the `transparent` parameter is specified, worker processes inherit the `CAP_NET_RAW` capability from the master process. It is also necessary to configure kernel routing table to intercept network traffic from the SCGI server.

#### `scgi_buffer_size`

SYNTAX: **scgi\_buffer\_size** *size*;

DEFAULT 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** on | off;

DEFAULT on

CONTEXT: http, server, location

Enables or disables buffering of responses from the SCGI server.

When buffering is enabled, nginx 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.

When buffering is disabled, the response is passed to a client synchronously, immediately as it is received. nginx will not try to read the whole response from the SCGI server. The maximum size of the data that nginx 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 [scgi\\_ignore\\_headers](#) directive.

### scgi\_buffers

SYNTAX: **scgi\_buffers** *number size*;

DEFAULT 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 8k | 16k

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 *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 (1.7.9). The *off* parameter disables caching inherited from the previous configuration level.

## scgi\_cache\_background\_update

SYNTAX: **scgi\_cache\_background\_update** *on* | *off*;

DEFAULT *off*

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.11.10.

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.

## scgi\_cache\_bypass

SYNTAX: **scgi\_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:

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

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 *off*

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.1.12.

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 an 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: **scgi\_cache\_lock\_age** *time*;

DEFAULT 5s

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.7.8.

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 5s

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.1.12.

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.

Before 1.7.8, the response could be cached.

### **scgi\_cache\_max\_range\_offset**

SYNTAX: **scgi\_cache\_max\_range\_offset** *number*;

DEFAULT —

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.11.6.

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 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 1

CONTEXT: http, server, location

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

### scgi\_cache\_path

SYNTAX: **scgi\_cache\_path** *path* [levels=*levels*] [use\_temp\_path=on|off]  
keys\_zone=*name:size* [inactive=*time*] [max\_size=*size*]  
[manager\_files=*number*] [manager\_sleep=*time*]  
[manager\_threshold=*time*] [loader\_files=*number*]  
[loader\_sleep=*time*] [loader\_threshold=*time*]  
[purger=on|off] [purger\_files=*number*] [purger\_sleep=*time*]  
[purger\_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/nginx/cache levels=1:2 keys_zone=one:10m;
```

file names in a cache will look like this:

```
/data/nginx/cache/c/29/b7f54b2df7773722d382f4809d65029c
```

A cached response is first written to a temporary file, and then the file is renamed. Starting from version 0.8.9, 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 (1.7.10). 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. If the value is set to *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.

As part of [commercial subscription](#), the shared memory zone also stores extended cache [information](#), thus, it is required to specify a larger zone size

for the same number of keys. For example, one megabyte zone can store about 4 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.

The special “cache manager” process monitors the maximum cache size set by the `max_size` parameter. When this size is exceeded, it removes the least recently used data. The data is removed in iterations configured by `manager_files`, `manager_threshold`, and `manager_sleep` parameters (1.11.5). During one iteration no more than `manager_files` items are deleted (by default, 100). The duration of one iteration is limited by the `manager_threshold` parameter (by default, 200 milliseconds). Between iterations, a pause configured by the `manager_sleep` parameter (by default, 50 milliseconds) is made.

A minute after the start 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. During one iteration no more than `loader_files` items are loaded (by default, 100). Besides, the duration of one iteration is limited by the `loader_threshold` parameter (by default, 200 milliseconds). Between iterations, a pause configured by the `loader_sleep` parameter (by default, 50 milliseconds) is made.

Additionally, the following parameters are available as part of our [commercial subscription](#):

`purger=on|off`

Instructs whether cache entries that match a [wildcard key](#) will be removed from the disk by the cache purger (1.7.12). Setting the parameter to `on` (default is `off`) will activate the “cache purger” process that permanently iterates through all cache entries and deletes the entries that match the wildcard key.

`purger_files=number`

Sets the number of items that will be scanned during one iteration (1.7.12). By default, `purger_files` is set to 10.

`purger_threshold=number`

Sets the duration of one iteration (1.7.12). By default, `purger_threshold` is set to 50 milliseconds.

`purger_sleep=number`

Sets a pause between iterations (1.7.12). By default, `purger_sleep` is set to 50 milliseconds.

In versions 1.7.3, 1.7.7, and 1.11.10 cache header format has been changed. Previously cached responses will be considered invalid after upgrading to a newer nginx version.

## scgi\_cache\_purge

SYNTAX: **scgi\_cache\_purge**string ...;

DEFAULT —

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.5.7.

Defines conditions under which the request will be considered a cache purge request. If at least one value of the string parameters is not empty and is not equal to “0” then the cache entry with a corresponding [cache key](#) is removed. The result of successful operation is indicated by returning the 204 No Content response.

If the [cache key](#) of a purge request ends with an asterisk (“\*”), all cache entries matching the wildcard key will be removed from the cache. However, these entries will remain on the disk until they are deleted for either [inactivity](#), or processed by the [cache purger](#) (1.7.12), or a client attempts to access them.

Example configuration:

```

scgi_cache_path /data/nginx/cache keys_zone=cache_zone:10m;

map $request_method $purge_method {
    PURGE    1;
    default  0;
}

server {
    ...
    location / {
        scgi_pass          backend;
        scgi_cache          cache_zone;
        scgi_cache_key      $uri;
        scgi_cache_purge    $purge_method;
    }
}

```

This functionality is available as part of our [commercial subscription](#).

## scgi\_cache\_revalidate

SYNTAX: **scgi\_cache\_revalidate** on | off;

DEFAULT off

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.5.7.

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: **scgi\_cache\_use\_stale** error | timeout | invalid\_header |  
updating | http\_500 | http\_503 | http\_403 | http\_404 |  
http\_429 | off ...;

DEFAULT off

CONTEXT: http, server, location

Determines in which cases a stale cached response can be used when an error occurs during communication with the SCGI server. The directive's parameters match the parameters of the [scgi\\_next\\_upstream](#) directive.

The `error` parameter also permits using a stale cached response if an SCGI server to process a request cannot be selected.

Additionally, the `updating` 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 (1.11.10). This has lower priority than using the directive parameters.

- 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

```
scgi_cache_valid 200 302 10m;  
scgi_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

```
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:

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

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 (1.7.7). If the header includes the `Vary` field with another value, such a response will be cached taking into account the corresponding request header fields (1.7.7).

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 60s  
CONTEXT: http, server, location

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

### **scgi\_force\_ranges**

SYNTAX: **scgi\_force\_ranges** on | off;  
DEFAULT off  
CONTEXT: http, server, location  
THIS DIRECTIVE APPEARED IN VERSION 1.7.7.

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: **scgi\_hide\_header** *field*;  
DEFAULT —  
CONTEXT: http, server, location

By default, nginx does not pass the header fields `Status` and `X-Accel-...` from the response of an 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: **scgi\_ignore\_client\_abort** on | off;

DEFAULT off

CONTEXT: http, server, location

Determines whether the connection with an 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 (1.1.6), X-Accel-Buffering (1.1.6), X-Accel-Charset (1.1.6), Expires, Cache-Control, Set-Cookie (0.8.44), and Vary (1.7.7).

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.

### scgi\_intercept\_errors

SYNTAX: **scgi\_intercept\_errors** on | off;

DEFAULT off

CONTEXT: http, server, location

Determines whether an SCGI server responses with codes greater than or equal to 300 should be passed to a client or be intercepted and redirected to nginx for processing with the [error\\_page](#) directive.

### scgi\_limit\_rate

SYNTAX: **scgi\_limit\_rate** *rate*;

DEFAULT 0

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.7.7.

Limits the speed of reading the response from the SCGI server. The *rate* is specified in bytes per second. The zero value disables rate limiting. The limit is set per a request, and so if nginx 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 1024m

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.

The zero value 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 error timeout

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 (1.11.13);

`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 (1.9.13); 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. The cases of `error`, `timeout` and `invalid_header` are always considered unsuccessful attempts, even if they are not specified in the directive. The cases of `http_500`, `http_503`, and `http_429` are considered unsuccessful attempts only if they are specified in the directive. The cases of `http_403` and `http_404` are 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: **scgi\_next\_upstream\_timeout** *time*;  
DEFAULT 0  
CONTEXT: http, server, location  
THIS DIRECTIVE APPEARED IN VERSION 1.7.5.

Limits the time during which a request can be passed to the [next server](#). The 0 value turns off this limitation.

### **scgi\_next\_upstream\_tries**

SYNTAX: **scgi\_next\_upstream\_tries** *number*;  
DEFAULT 0  
CONTEXT: http, server, location  
THIS DIRECTIVE APPEARED IN VERSION 1.7.5.

Limits the number of possible tries for passing a request to the [next server](#). The 0 value 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:

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

Can be used along with the [scgi\\_cache\\_bypass](#) directive.

## scgi\_param

SYNTAX: **scgi\_param** *parameter value* [if\_not\_empty];

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 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` (1.1.11) 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: **scgi\_pass** *address*;

DEFAULT —

CONTEXT: location, if in location

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

```
scgi_pass localhost:9000;
```

or as a UNIX-domain socket path:

```
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](#).

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 an SCGI server to a client.

### scgi\_pass\_request\_body

SYNTAX: **scgi\_pass\_request\_body** on | off;

DEFAULT on

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 on

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: **scgi\_read\_timeout** *time*;

DEFAULT 60s

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 on

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.7.11.

Enables or disables buffering of a client request body.

When buffering is enabled, the entire request body is [read](#) from the client before sending the request to an SCGI server.

When buffering is disabled, the request body is sent to the SCGI server immediately as it is received. In this case, the request cannot be passed to the [next server](#) if nginx 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: **scgi\_send\_timeout** *time*;

DEFAULT 60s

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 off

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.15.6.

Configures the “TCP keepalive” behavior for outgoing connections to an SCGI server. By default, the operating system’s settings are in effect for the socket. If the directive is set to the value “on”, the `SO_KEEPALIVE` socket option is turned on for the socket.

### scgi\_store

SYNTAX: **scgi\_store** on | off | *string*;

DEFAULT off

CONTEXT: http, server, location

Enables saving of files to a disk. The `on` parameter saves files with paths corresponding to the directives [alias](#) or [root](#). The `off` parameter disables saving of files. In addition, 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. 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 [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/ {
    root          /data/www;
    error_page    404 = /fetch$uri;
}

location /fetch/ {
    internal;

    scgi_pass     backend:9000;
    ...

    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 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: **scgi\_temp\_file\_write\_size** *size*;

DEFAULT 8k|16k

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 `scgi_temp`

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/nginx/scgi_temp 1 2;
```

a temporary file might look like this:

```
/spool/nginx/scgi_temp/7/45/00000123457
```

See also the `use_temp_path` parameter of the [scgi\\_cache\\_path](#) directive.

## 2.42 Module ngx\_http\_secure\_link\_module

2.42.1 Summary	263
2.42.2 Directives	263
<a href="#">secure_link</a>	263
<a href="#">secure_link_md5</a>	264
<a href="#">secure_link_secret</a>	264
2.42.3 Embedded Variables	265

### 2.42.1 Summary

The `ngx_http_secure_link_module` (0.7.18) 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 (0.8.50) is enabled by the [secure\\_link](#) and [secure\\_link\\_md5](#) directives and is also used to limit lifetime of links.

This module is not built by default, it should be enabled with the `--with-http_secure_link_module` configuration parameter.

### 2.42.2 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
```

### 2.42.3 Embedded 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.

## 2.43 Module ngx\_http\_session\_log\_module

2.43.1 Summary	266
2.43.2 Example Configuration	266
2.43.3 Directives	266
session_log	266
session_log_format	266
session_log_zone	267
2.43.4 Embedded Variables	267

### 2.43.1 Summary

The `ngx_http_session_log_module` module enables logging sessions (that is, aggregates of multiple HTTP requests) instead of individual HTTP requests.

This module is available as part of our [commercial subscription](#).

### 2.43.2 Example Configuration

The following configuration sets up a session log and maps requests to sessions according to the request client address and User-Agent request header field:

```
session_log_zone /path/to/log format=combined
                 zone=one:1m timeout=30s
                 md5=$binary_remote_addr$http_user_agent;

location /media/ {
    session_log one;
}
```

### 2.43.3 Directives

#### session\_log

SYNTAX: **session\_log** *name* | `off`;

DEFAULT `off`

CONTEXT: `http`, `server`, `location`

Enables the use of the specified session log. The special value `off` cancels all `session_log` directives inherited from the previous configuration level.

#### session\_log\_format

SYNTAX: **session\_log\_format** *name string* ...;

DEFAULT `combined "..."`

CONTEXT: `http`

Specifies the output format of a log. The value of the *\$body\_bytes\_sent* variable is aggregated across all requests in a session. The values of all other variables available for logging correspond to the first request in a session.

### session\_log\_zone

SYNTAX: **session\_log\_zone** *path* zone=*name:size* [*format=format*]  
[*timeout=time*] [*id=id*] [*md5=md5*];

DEFAULT —

CONTEXT: http

Sets the path to a log file and configures the shared memory zone that is used to store currently active sessions.

A session is considered active for as long as the time elapsed since the last request in the session does not exceed the specified *timeout* (by default, 30 seconds). Once a session is no longer active, it is written to the log.

The *id* parameter identifies the session to which a request is mapped. The *id* parameter is set to the hexadecimal representation of an MD5 hash (for example, obtained from a cookie using variables). If this parameter is not specified or does not represent the valid MD5 hash, nginx computes the MD5 hash from the value of the *md5* parameter and creates a new session using this hash. Both the *id* and *md5* parameters can contain variables.

The *format* parameter sets the custom session log format configured by the [session\\_log\\_format](#) directive. If *format* is not specified, the predefined “combined” format is used.

## 2.43.4 Embedded Variables

The `ngx_http_session_log_module` module supports two embedded variables:

*\$session\_log\_id*

current session ID;

*\$session\_log\_binary\_id*

current session ID in binary form (16 bytes).

## 2.44 Module ngx\_http\_slice\_module

2.44.1 Summary	268
2.44.2 Example Configuration	268
2.44.3 Directives	268
<a href="#">slice</a>	268
2.44.4 Embedded Variables	268

### 2.44.1 Summary

The `ngx_http_slice_module` module (1.9.8) 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.

This module is not built by default, it should be enabled with the `--with-http_slice_module` configuration parameter.

### 2.44.2 Example Configuration

```
location / {
    slice          1m;
    proxy_cache    cache;
    proxy_cache_key $uri$is_args$args$slice_range;
    proxy_set_header Range $slice_range;
    proxy_cache_valid 200 206 1h;
    proxy_pass      http://localhost:8000;
}
```

In this example, the response is split into 1-megabyte cacheable slices.

### 2.44.3 Directives

#### `slice`

SYNTAX: **slice** *size*;

DEFAULT 0

CONTEXT: http, server, location

Sets the *size* of the slice. The zero value disables splitting responses into slices. 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](#).

### 2.44.4 Embedded Variables

The `ngx_http_slice_module` module supports the following embedded variables:

*\$slice\_range*

the current slice range in [HTTP byte range](#) format, for example,  
bytes=0-1048575.

## 2.45 Module ngx\_http\_split\_clients\_module

2.45.1 Summary	270
2.45.2 Example Configuration	270
2.45.3 Directives	270
split_clients	270

### 2.45.1 Summary

The `ngx_http_split_clients_module` module creates variables suitable for A/B testing, also known as split testing.

### 2.45.2 Example Configuration

```
http {
    split_clients "${remote_addr}AAA" $variant {
        0.5% .one;
        2.0% .two;
        *    "";
    }

    server {
        location / {
            index index${variant}.html;
        }
    }
}
```

### 2.45.3 Directives

#### split\_clients

SYNTAX: **split\_clients** *string* *\$variable* { ... }

DEFAULT —

CONTEXT: http

Creates a variable for A/B testing, for example:

```
split_clients "${remote_addr}AAA" $variant {
    0.5% .one;
    2.0% .two;
    *    "";
}
```

The value of the original string is hashed using MurmurHash2. In the example given, hash values from 0 to 21474835 (0.5%) correspond to the value `".one"` of the `$variant` variable, hash values from 21474836 to 107374180 (2%) correspond to the value `".two"`, and hash values from 107374181 to 4294967295 correspond to the value `""` (an empty string).

## 2.46 Module ngx\_http\_ssi\_module

2.46.1	Summary	271
2.46.2	Example Configuration	271
2.46.3	Directives	271
	ssi	271
	ssi_last_modified	271
	ssi_min_file_chunk	272
	ssi_silent_errors	272
	ssi_types	272
	ssi_value_length	272
2.46.4	SSI Commands	272
2.46.5	Embedded Variables	276

### 2.46.1 Summary

The `ngx_http_ssi_module` module is a filter that processes SSI (Server Side Includes) commands in responses passing through it. Currently, the list of supported SSI commands is incomplete.

### 2.46.2 Example Configuration

```
location / {
    ssi on;
    ...
}
```

### 2.46.3 Directives

#### ssi

SYNTAX: **ssi** on | off;

DEFAULT 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 off

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.5.1.

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 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 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 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 (0.8.29).

### **ssi\_value\_length**

SYNTAX: **ssi\_value\_length** *length*;

DEFAULT 256

CONTEXT: http, server, location

Sets the maximum length of parameter values in SSI commands.

## **2.46.4 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 -->
```

```
<!--# 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:

```
<!--# if expr="$name = /(.)@(?P<domain>.+)/" -->
  <!--# echo var="1" -->
  <!--# echo var="domain" -->
<!--# endif -->
```

`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 (1.13.10):

```
location /remote/ {
    subrequest_output_buffer_size 64k;
    ...
}
```

Prior to version 1.13.10, only the results of responses obtained using the [ngx\\_http\\_proxy\\_module](#), [ngx\\_http\\_memcached\\_module](#), [ngx\\_http\\_fastcgi\\_module](#) (1.5.6), [ngx\\_http\\_uwsgi\\_module](#) (1.5.6), and [ngx\\_http\\_scgi\\_module](#) (1.5.6) modules could be written into variables. The maximum size of the response was set with the [proxy\\_buffer\\_size](#), [memcached\\_buffer\\_size](#), [fastcgi\\_buffer\\_size](#), [uwsgi\\_buffer\\_size](#), and [scgi\\_buffer\\_size](#) directives.

**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.

### 2.46.5 Embedded Variables

The `ngx_http_ssi_module` module supports two embedded 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.

## 2.47 Module ngx\_http\_ssl\_module

2.47.1	Summary	277
2.47.2	Example Configuration	277
2.47.3	Directives	278
	ssl	278
	ssl_buffer_size	278
	ssl_certificate	279
	ssl_certificate_key	279
	ssl_ciphers	280
	ssl_client_certificate	280
	ssl_crl	280
	ssl_dhparam	281
	ssl_early_data	281
	ssl_ecdh_curve	281
	ssl_password_file	282
	ssl_prefer_server_ciphers	282
	ssl_protocols	282
	ssl_session_cache	283
	ssl_session_ticket_key	283
	ssl_session_tickets	284
	ssl_session_timeout	284
	ssl_stapling	284
	ssl_stapling_file	285
	ssl_stapling_responder	285
	ssl_stapling_verify	285
	ssl_trusted_certificate	285
	ssl_verify_client	286
	ssl_verify_depth	286
2.47.4	Error Processing	286
2.47.5	Embedded Variables	286

### 2.47.1 Summary

The `ngx_http_ssl_module` module provides the necessary support for HTTPS.

This module is not built by default, it should be enabled with the `--with-http_ssl_module` configuration parameter.

This module requires the [OpenSSL](#) library.

### 2.47.2 Example Configuration

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,
- 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 TLSv1.1 TLSv1.2;
        ssl_ciphers      AES128-SHA:AES256-SHA:RC4-SHA:DES-CBC3-SHA:RC4-MD5;
        ssl_certificate   /usr/local/nginx/conf/cert.pem;
        ssl_certificate_key /usr/local/nginx/conf/cert.key;
        ssl_session_cache shared:SSL:10m;
        ssl_session_timeout 10m;

        ...
    }
}
```

### 2.47.3 Directives

#### ssl

SYNTAX: **ssl** on | off;  
DEFAULT off  
CONTEXT: http, server

This directive was made obsolete in version 1.15.0. The `ssl` parameter of the [listen](#) directive should be used instead.

#### ssl\_buffer\_size

SYNTAX: **ssl\_buffer\_size** *size*;  
DEFAULT 16k  
CONTEXT: http, server  
THIS DIRECTIVE APPEARED IN VERSION 1.5.9.

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.

Since version 1.11.0, this directive can be specified multiple times to load certificates of different types, for example, RSA and ECDSA:

```
server {
    listen      443 ssl;
    server_name example.com;

    ssl_certificate      example.com.rsa.crt;
    ssl_certificate_key  example.com.rsa.key;

    ssl_certificate      example.com.ecdsa.crt;
    ssl_certificate_key  example.com.ecdsa.key;

    ...
}
```

Only OpenSSL 1.0.2 or higher supports separate [certificate chains](#) for different certificates. With older versions, only one certificate chain can be used.

Since version 1.15.9, 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* (1.15.10), 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](#).

It should be kept in mind that due to the HTTPS protocol limitations for maximum interoperability virtual servers should listen on [different IP addresses](#).

## 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.

The value `engine:name:id` can be specified instead of the *file* (1.7.9), 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* (1.15.10), 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](#).

Since version 1.15.9, variables can be used in the *file* name when using OpenSSL 1.0.2 or higher.

## ssl\_ciphers

SYNTAX: **ssl\_ciphers** *ciphers*;

DEFAULT 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 full list can be viewed using the “`openssl ciphers`” command.

The previous versions of nginx used [different](#) ciphers by default.

## 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\_crl

SYNTAX: **ssl\_crl** *file*;

DEFAULT —

CONTEXT: http, server

THIS DIRECTIVE APPEARED IN VERSION 0.8.7.

Specifies a *file* with revoked certificates (CRL) in the PEM format used to [verify](#) client certificates.

## ssl\_dhparam

SYNTAX: **ssl\_dhparam** *file*;  
DEFAULT —  
CONTEXT: http, server  
THIS DIRECTIVE APPEARED IN VERSION 0.7.2.

Specifies a *file* with DH parameters for DHE ciphers.

By default no parameters are set, and therefore DHE ciphers will not be used.

Prior to version 1.11.0, builtin parameters were used by default.

## ssl\_early\_data

SYNTAX: **ssl\_early\_data** on | off;  
DEFAULT off  
CONTEXT: http, server  
THIS DIRECTIVE APPEARED IN VERSION 1.15.3.

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;
```

The directive is supported when using OpenSSL 1.1.1 or higher (1.15.4) and [BoringSSL](#).

## ssl\_ecdh\_curve

SYNTAX: **ssl\_ecdh\_curve** *curve*;  
DEFAULT auto  
CONTEXT: http, server  
THIS DIRECTIVE APPEARED IN VERSIONS 1.1.0 AND 1.0.6.

Specifies a *curve* for ECDHE ciphers.

When using OpenSSL 1.0.2 or higher, it is possible to specify multiple curves (1.11.0), for example:

```
ssl_ecdh_curve prime256v1:secp384r1;
```

The special value `auto` (1.11.0) instructs nginx to use a list built into the OpenSSL library when using OpenSSL 1.0.2 or higher, or `prime256v1` with older versions.

Prior to version 1.11.0, the prime256v1 curve was used by default.

### ssl\_password\_file

SYNTAX: **ssl\_password\_file** *file*;

DEFAULT —

CONTEXT: http, server

THIS DIRECTIVE APPEARED IN VERSION 1.7.3.

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 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] [TLSv1.1] [TLSv1.2]  
[TLSv1.3];

DEFAULT TLSv1 TLSv1.1 TLSv1.2

CONTEXT: http, server

Enables the specified protocols.

The TLSv1.1 and TLSv1.2 parameters (1.1.13, 1.0.12) work only when OpenSSL 1.0.1 or higher is used.

The `TLSv1.3` parameter (1.13.0) works only when OpenSSL 1.1.1 built with TLSv1.3 support is used.

### `ssl_session_cache`

SYNTAX: **`ssl_session_cache`** `off` | `none` | [`builtin[:size]`]  
           [`shared:name:size`];  
 DEFAULT `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: nginx explicitly tells a client that sessions may not be reused.

`none`

the use of a session cache is gently disallowed: nginx 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.

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`  
 THIS DIRECTIVE APPEARED IN VERSION 1.5.7.

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, 1.11.8) or AES128 (for 48-byte keys) is used for encryption.

### ssl\_session\_tickets

SYNTAX: **ssl\_session\_tickets** on | off;

DEFAULT on

CONTEXT: http, server

THIS DIRECTIVE APPEARED IN VERSION 1.5.9.

Enables or disables session resumption through [TLS session tickets](#).

### ssl\_session\_timeout

SYNTAX: **ssl\_session\_timeout** *time*;

DEFAULT 5m

CONTEXT: http, server

Specifies a time during which a client may reuse the session parameters.

### ssl\_stapling

SYNTAX: **ssl\_stapling** on | off;

DEFAULT off

CONTEXT: http, server

THIS DIRECTIVE APPEARED IN VERSION 1.3.7.

Enables or disables [stapling of OCSP responses](#) by the server. Example:

```
ssl_stapling on;  
resolver 192.0.2.1;
```

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.

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

THIS DIRECTIVE APPEARED IN VERSION 1.3.7.

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** *url*;

DEFAULT —

CONTEXT: http, server

THIS DIRECTIVE APPEARED IN VERSION 1.3.7.

Overrides the URL 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 off

CONTEXT: http, server

THIS DIRECTIVE APPEARED IN VERSION 1.3.7.

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: **ssl\_trusted\_certificate** *file*;

DEFAULT —

CONTEXT: http, server

THIS DIRECTIVE APPEARED IN VERSION 1.3.7.

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 off

CONTEXT: http, server

Enables verification of client certificates. The verification result is stored in the [\\$ssl\\_client\\_verify](#) variable.

The `optional` parameter (0.8.7+) requests the client certificate and verifies it if the certificate is present.

The `optional_no_ca` parameter (1.3.8, 1.2.5) 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 nginx performs the actual certificate verification. The contents of the certificate is accessible through the [\\$ssl\\_client\\_cert](#) variable.

## ssl\_verify\_depth

SYNTAX: **ssl\_verify\_depth** *number*;

DEFAULT 1

CONTEXT: http, server

Sets the verification depth in the client certificates chain.

### 2.47.4 Error Processing

The `ngx_http_ssl_module` 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.

### 2.47.5 Embedded Variables

The `ngx_http_ssl_module` module supports several embedded variables:

*\$ssl\_cipher*

returns the string of ciphers used for an established SSL connection;

*\$ssl\_ciphers*

returns the list of ciphers supported by the client (1.11.7). Known ciphers are listed by names, unknown are shown in hexadecimal, for example:

`AES128-SHA:AES256-SHA:0x00ff`

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 (1.13.5);

*\$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; this is intended for the use in the [proxy\\_set\\_header](#) directive;

The variable is deprecated, the *\$ssl\_client\_escaped\_cert* variable should be used instead.

*\$ssl\_client\_fingerprint*

returns the SHA1 fingerprint of the client certificate for an established SSL connection (1.7.1);

*\$ssl\_client\_i\_dn*

returns the “issuer DN” string of the client certificate for an established SSL connection according to [RFC 2253](#) (1.11.6);

*\$ssl\_client\_i\_dn\_legacy*

returns the “issuer DN” string of the client certificate for an established SSL connection;

Prior to version 1.11.6, the variable name was *\$ssl\_client\_i\_dn*.

*\$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](#) (1.11.6);

*\$ssl\_client\_s\_dn\_legacy*

returns the “subject DN” string of the client certificate for an established SSL connection;

Prior to version 1.11.6, the variable name was *\$ssl\_client\_s\_dn*.

*\$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 (1.11.7);

*\$ssl\_client\_v\_remain*

returns the number of days until the client certificate expires (1.11.7);

*\$ssl\_client\_v\_start*

returns the start date of the client certificate (1.11.7);

*\$ssl\_client\_verify*

returns the result of client certificate verification: “SUCCESS”, “FAILED: *reason*”, and “NONE” if a certificate was not present;

Prior to version 1.11.7, the “FAILED” result did not contain the *reason* string.

*\$ssl\_curves*

returns the list of curves supported by the client (1.11.7). Known curves are listed by names, unknown are shown in hexadecimal, for example:

```
0x001d:prime256v1:secp521r1:secp384r1
```

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 “” (1.15.3).

*\$ssl\_protocol*

returns the protocol of an established SSL connection;

*\$ssl\_server\_name*

returns the server name requested through [SNI](#) (1.7.0);

*\$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 (1.5.11).

## 2.48 Module ngx\_http\_status\_module

2.48.1 Summary	289
2.48.2 Example Configuration	289
2.48.3 Directives	290
status	290
status_format	290
status_zone	291
2.48.4 Data	291
2.48.5 Compatibility	298

### 2.48.1 Summary

The ngx\_http\_status\_module module provides access to various status information.

This module was available as part of our [commercial subscription](#) until 1.13.10. It was superseded by the [ngx\\_http\\_api\\_module](#) module in 1.13.3.

### 2.48.2 Example Configuration

```
http {
    upstream backend {
        zone http_backend 64k;

        server backend1.example.com weight=5;
        server backend2.example.com;
    }

    proxy_cache_path /data/nginx/cache_backend keys_zone=cache_backend:10m;

    server {
        server_name backend.example.com;

        location / {
            proxy_pass http://backend;
            proxy_cache cache_backend;

            health_check;
        }

        status_zone server_backend;
    }

    server {
        listen 127.0.0.1;

        location /upstream_conf {
            upstream_conf;
        }

        location /status {
            status;
        }

        location = /status.html {
        }
    }
}
```

```

    }
}

stream {
    upstream backend {
        zone stream_backend 64k;

        server backend1.example.com:12345 weight=5;
        server backend2.example.com:12345;
    }

    server {
        listen      127.0.0.1:12345;
        proxy_pass  backend;
        status_zone server_backend;
        health_check;
    }
}

```

Examples of status requests with this configuration:

```

http://127.0.0.1/status
http://127.0.0.1/status/nginx_version
http://127.0.0.1/status/caches/cache_backend
http://127.0.0.1/status/upstreams
http://127.0.0.1/status/upstreams/backend
http://127.0.0.1/status/upstreams/backend/peers/1
http://127.0.0.1/status/upstreams/backend/peers/1/weight
http://127.0.0.1/status/stream
http://127.0.0.1/status/stream/upstreams
http://127.0.0.1/status/stream/upstreams/backend
http://127.0.0.1/status/stream/upstreams/backend/peers/1
http://127.0.0.1/status/stream/upstreams/backend/peers/1/weight

```

The simple monitoring page is shipped with this distribution, accessible as “/status.html” in the default configuration. It requires the locations “/status” and “/status.html” to be configured as shown above.

### 2.48.3 Directives

#### status

SYNTAX: **status**;

DEFAULT —

CONTEXT: location

The status information will be accessible from the surrounding location. Access to this location should be [limited](#).

#### status\_format

SYNTAX: **status\_format** json;

SYNTAX: **status\_format** jsonp [*callback*];

DEFAULT json

CONTEXT: http, server, location

By default, status information is output in the JSON format.

Alternatively, data may be output as JSONP. The *callback* parameter specifies the name of a callback function. The value can contain variables. If

parameter is omitted, or the computed value is an empty string, then “ngx\_status\_jsonp\_callback” is used.

### status\_zone

SYNTAX: **status\_zone** *zone*;

DEFAULT —

CONTEXT: server

Enables collection of virtual [http](#) or [stream](#) (1.7.11) server status information in the specified *zone*. Several servers may share the same zone.

## 2.48.4 Data

The following status information is provided:

version

Version of the provided data set. The current version is 8.

nginx\_version

Version of nginx.

nginx\_build

Name of nginx build.

address

The address of the server that accepted status request.

generation

The total number of configuration [reloads](#).

load\_timestamp

Time of the last reload of configuration, in milliseconds since Epoch.

timestamp

Current time in milliseconds since Epoch.

pid

The ID of the worker process that handled status request.

ppid

The ID of the master process that started the [worker process](#).

processes

respawned

The total number of abnormally terminated and respawned child processes.

connections

accepted

The total number of accepted client connections.

dropped

The total number of dropped client connections.

active

The current number of active client connections.

idle  
The current number of idle client connections.

ssl

handshakes  
The total number of successful SSL handshakes.

handshakes\_failed  
The total number of failed SSL handshakes.

session\_reuses  
The total number of session reuses during SSL handshake.

requests

total  
The total number of client requests.

current  
The current number of client requests.

server\_zones

For each [status\\_zone](#):

processing  
The number of client requests that are currently being processed.

requests  
The total number of client requests received from clients.

responses

total  
The total number of responses sent to clients.

1xx, 2xx, 3xx, 4xx, 5xx  
The number of responses with status codes 1xx, 2xx, 3xx, 4xx, and 5xx.

discarded  
The total number of requests completed without sending a response.

received  
The total number of bytes received from clients.

sent  
The total number of bytes sent to clients.

slabs

For each shared memory zone that uses slab allocator:

pages

used  
The current number of used memory pages.

free  
The current number of free memory pages.

slots  
For each memory slot size (8, 16, 32, 64, 128, etc.) the following data are provided:

used

The current number of used memory slots.

free

The current number of free memory slots.

reqs

The total number of attempts to allocate memory of specified size.

fails

The number of unsuccessful attempts to allocate memory of specified size.

upstreams

For each [dynamically configurable group](#), the following data are provided:

peers

For each [server](#), the following data are provided:

id

The ID of the server.

server

An [address](#) of the server.

name

The name of the server specified in the [server](#) directive.

service

The [service](#) parameter value of the [server](#) directive.

backup

A boolean value indicating whether the server is a [backup](#) server.

weight

[Weight](#) of the server.

state

Current state, which may be one of “up”, “draining”, “down”, “unavail”, “checking”, or “unhealthy”.

active

The current number of active connections.

max\_conns

The [max\\_conns](#) limit for the server.

requests

The total number of client requests forwarded to this server.

responses

total

The total number of responses obtained from this server.

1xx, 2xx, 3xx, 4xx, 5xx

The number of responses with status codes 1xx, 2xx, 3xx, 4xx, and 5xx.

sent

The total number of bytes sent to this server.

received

The total number of bytes received from this server.

fails

The total number of unsuccessful attempts to communicate with the server.

unavail

How many times the server became unavailable for client requests (state “unavail”) due to the number of unsuccessful attempts reaching the [max\\_fails](#) threshold.

health\_checks

checks

The total number of [health check](#) requests made.

fails

The number of failed health checks.

unhealthy

How many times the server became unhealthy (state “unhealthy”).

last\_passed

Boolean indicating if the last health check request was successful and passed [tests](#).

downtime

Total time the server was in the “unavail”, “checking”, and “unhealthy” states.

downstart

The time (in milliseconds since Epoch) when the server became “unavail”, “checking”, or “unhealthy”.

selected

The time (in milliseconds since Epoch) when the server was last selected to process a request (1.7.5).

header\_time

The average time to get the [response header](#) from the server (1.7.10). Prior to version 1.11.6, the field was available only when using the [least\\_time](#) load balancing method.

response\_time

The average time to get the [full response](#) from the server (1.7.10). Prior to version 1.11.6, the field was available only when using the [least\\_time](#) load balancing method.

keepalive

The current number of idle [keepalive](#) connections.

zombies

The current number of servers removed from the group but still processing active client requests.

zone

The name of the shared memory [zone](#) that keeps the group’s configuration and run-time state.

### queue

For the requests [queue](#), the following data are provided:

#### size

The current number of requests in the queue.

#### max\_size

The maximum number of requests that can be in the queue at the same time.

#### overflows

The total number of requests rejected due to the queue overflow.

### caches

For each cache (configured by [proxy\\_cache\\_path](#) and the likes):

#### size

The current size of the cache.

#### max\_size

The limit on the maximum size of the cache specified in the configuration.

#### cold

A boolean value indicating whether the “cache loader” process is still loading data from disk into the cache.

#### hit, stale, updating, revalidated

##### responses

The total number of responses read from the cache (hits, or stale responses due to [proxy\\_cache\\_use\\_stale](#) and the likes).

##### bytes

The total number of bytes read from the cache.

#### miss, expired, bypass

##### responses

The total number of responses not taken from the cache (misses, expires, or bypasses due to [proxy\\_cache\\_bypass](#) and the likes).

##### bytes

The total number of bytes read from the proxied server.

##### responses\_written

The total number of responses written to the cache.

##### bytes\_written

The total number of bytes written to the cache.

### stream

#### server\_zones

For each [status\\_zone](#):

##### processing

The number of client connections that are currently being processed.

##### connections

The total number of connections accepted from clients.

`sessions``total`

The total number of completed client sessions.

`2xx, 4xx, 5xx`

The number of sessions completed with [status codes](#) 2xx, 4xx, or 5xx.

`discarded`

The total number of connections completed without creating a session.

`received`

The total number of bytes received from clients.

`sent`

The total number of bytes sent to clients.

`upstreams`

For each [dynamically configurable group](#), the following data are provided:

`peers`

For each [server](#) the following data are provided:

`id`

The ID of the server.

`server`

An [address](#) of the server.

`name`

The name of the server specified in the [server](#) directive.

`service`

The [service](#) parameter value of the [server](#) directive.

`backup`

A boolean value indicating whether the server is a [backup](#) server.

`weight`

[Weight](#) of the server.

`state`

Current state, which may be one of “up”, “down”, “unavail”, “checking”, or “unhealthy”.

`active`

The current number of connections.

`max_conns`

The [max\\_conns](#) limit for the server.

`connections`

The total number of client connections forwarded to this server.

`connect_time`

The average time to connect to the upstream server. Prior to version 1.11.6, the field was available only when using the [least\\_time](#) load balancing method.

`first_byte_time`

The average time to receive the first byte of data. Prior to version 1.11.6, the field was available only when using the [least\\_time](#) load balancing method.

`response_time`

The average time to receive the last byte of data. Prior to version 1.11.6, the field was available only when using the [least\\_time](#) load balancing method.

`sent`

The total number of bytes sent to this server.

`received`

The total number of bytes received from this server.

`fails`

The total number of unsuccessful attempts to communicate with the server.

`unavail`

How many times the server became unavailable for client connections (state “unavail”) due to the number of unsuccessful attempts reaching the [max\\_fails](#) threshold.

`health_checks``checks`

The total number of [health check](#) requests made.

`fails`

The number of failed health checks.

`unhealthy`

How many times the server became unhealthy (state “unhealthy”).

`last_passed`

Boolean indicating if the last health check request was successful and passed [tests](#).

`downtime`

Total time the server was in the “unavail”, “checking”, and “unhealthy” states.

`downstart`

The time (in milliseconds since Epoch) when the server became “unavail”, “checking”, or “unhealthy”.

`selected`

The time (in milliseconds since Epoch) when the server was last selected to process a connection.

`zombies`

The current number of servers removed from the group but still processing active client connections.

`zone`

The name of the shared memory [zone](#) that keeps the group’s configuration and run-time state.

## 2.48.5 Compatibility

- The `zone` field in `http` and `stream` upstreams was added in [version 8](#).
- The `slabs` status data were added in [version 8](#).
- The `checking` state was added in [version 8](#).
- The `name` and `service` fields in `http` and `stream` upstreams were added in [version 8](#).
- The `nginx_build` and `ppid` fields were added in [version 8](#).
- The `sessions` status data and the `discarded` field in stream `server_zones` were added in [version 7](#).
- The `zombies` field was moved from nginx `debug` version in [version 6](#).
- The `ssl` status data were added in [version 6](#).
- The `discarded` field in `server_zones` was added in [version 6](#).
- The `queue` status data were added in [version 6](#).
- The `pid` field was added in [version 6](#).
- The list of servers in `upstreams` was moved into `peers` in [version 6](#).
- The `keepalive` field of an upstream server was removed in [version 5](#).
- The `stream` status data were added in [version 5](#).
- The `generation` field was added in [version 5](#).
- The `respawned` field in `processes` was added in [version 5](#).
- The `header_time` and `response_time` fields in `upstreams` were added in [version 5](#).
- The `selected` field in `upstreams` was added in [version 4](#).
- The `draining` state in `upstreams` was added in [version 4](#).
- The `id` and `max_conns` fields in `upstreams` were added in [version 3](#).
- The `revalidated` field in `caches` was added in [version 3](#).
- The `server_zones`, `caches`, and `load_timestamp` status data were added in [version 2](#).

## 2.49 Module ngx\_http\_stub\_status\_module

2.49.1 Summary	299
2.49.2 Example Configuration	299
2.49.3 Directives	299
stub_status	299
2.49.4 Data	300
2.49.5 Embedded Variables	300

### 2.49.1 Summary

The ngx\_http\_stub\_status\_module module provides access to basic status information.

This module is not built by default, it should be enabled with the `--with-http_stub_status_module` configuration parameter.

### 2.49.2 Example Configuration

```
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
```

### 2.49.3 Directives

#### stub\_status

SYNTAX: **stub\_status;**

DEFAULT —

CONTEXT: server, location

The basic status information will be accessible from the surrounding location.

In versions prior to 1.7.5, the directive syntax required an arbitrary argument, for example, “`stub_status on`”.

### 2.49.4 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 nginx is reading the request header.

`Writing`

The current number of connections where nginx is writing the response back to the client.

`Waiting`

The current number of idle client connections waiting for a request.

### 2.49.5 Embedded Variables

The `ngx_http_stub_status_module` module supports the following embedded variables (1.3.14):

*`$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.

## 2.50 Module ngx\_http\_sub\_module

2.50.1 Summary	301
2.50.2 Example Configuration	301
2.50.3 Directives	301
sub_filter	301
sub_filter_last_modified	301
sub_filter_once	302
sub_filter_types	302

### 2.50.1 Summary

The ngx\_http\_sub\_module module is a filter that modifies a response by replacing one specified string by another.

This module is not built by default, it should be enabled with the `--with-http_sub_module` configuration parameter.

### 2.50.2 Example Configuration

```
location / {
    sub_filter ' <a href="http://127.0.0.1:8080/' ' <a href="https://$host/';
    sub_filter ' \w+) $ $route;
}

map $request_uri $route_uri {
    ~jsessionid=.+\. (?P<route>\w+) $ $route;
}

upstream backend {
    server backend1.example.com route=a;
    server backend2.example.com route=b;

    sticky route $route_cookie $route_uri;
}
```

Here, the route is taken from the “JSESSIONID” cookie if present in a request. Otherwise, the route from the URI is used.

`learn`

When the `learn` method (1.7.1) is used, nginx analyzes upstream server responses and learns server-initiated sessions usually passed in an HTTP cookie.

```
upstream backend {
    server backend1.example.com:8080;
    server backend2.example.com:8081;

    sticky learn
        create=$upstream_cookie_examplecookie
        lookup=$cookie_examplecookie
        zone=client_sessions:1m;
}
```

In the example, the upstream server creates a session by setting the cookie “EXAMPLECOOKIE” in the response. Further requests with this cookie will be passed to the same server. If the server cannot process the request, the new server is selected as if the client has not been bound yet.

The parameters `create` and `lookup` specify variables that indicate how new sessions are created and existing sessions are searched, respectively. Both parameters may be specified more than once, in which case the first non-empty variable is used.

Sessions are stored in a shared memory zone, whose *name* and *size* are configured by the `zone` parameter. One megabyte zone can store about 4000 sessions on the 64-bit platform. The sessions that are not accessed during the time specified by the `timeout` parameter get removed from the zone. By default, `timeout` is set to 10 minutes.

The `header` parameter (1.13.1) allows creating a session right after receiving response headers from the upstream server.

The `sync` parameter (1.13.8) enables [synchronization](#) of the shared

memory zone.

This directive is available as part of our [commercial subscription](#).

### sticky\_cookie\_insert

SYNTAX: **sticky\_cookie\_insert** *name* [expires=*time*] [domain=*domain*]  
[path=*path*];

DEFAULT —

CONTEXT: upstream

This directive is obsolete since version 1.5.7. An equivalent [sticky](#) directive with a new syntax should be used instead:

```
sticky cookie    name    [expires=time]    [domain=domain]  
[path=path];
```

## 2.51.4 Embedded Variables

The `ngx_http_upstream_module` module supports the following embedded variables:

#### *\$upstream\_addr*

keeps 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 (1.11.4). 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 (1.15.8). 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 (0.8.3). The status can be either “MISS”, “BYPASS”, “EXPIRED”, “STALE”, “UPDATING”, “REVALIDATED”, or “HIT”.

*\$upstream\_connect\_time*

keeps time spent on establishing a connection with the upstream server (1.9.1); 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 (1.7.1). Only the cookies from the response of the last server are saved.

*\$upstream\_header\_time*

keeps time spent on receiving the response header from the upstream server (1.7.10); 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*

keep 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*

keeps time the request spent in the upstream [queue](#) (1.13.9); 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\_response\_length*

keeps the length of the response obtained from the upstream server (0.7.27); 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\_trailer\_name*

keeps fields from the end of the response obtained from the upstream server (1.13.10).

## 2.52 Module ngx\_http\_upstream\_conf\_module

2.52.1 Summary	318
2.52.2 Example Configuration	318
2.52.3 Directives	318
upstream_conf	318

### 2.52.1 Summary

The `ngx_http_upstream_conf_module` module allows configuring upstream server groups on-the-fly via a simple HTTP interface without the need of restarting nginx. The [http](#) or [stream](#) server group must reside in the shared memory.

This module was available as part of our [commercial subscription](#) until 1.13.10. It was superseded by the [ngx\\_http\\_api\\_module](#) module in 1.13.3.

### 2.52.2 Example Configuration

```
upstream backend {
    zone upstream_backend 64k;

    ...
}

server {
    location /upstream_conf {
        upstream_conf;
        allow 127.0.0.1;
        deny all;
    }
}
```

### 2.52.3 Directives

#### upstream\_conf

SYNTAX: **upstream\_conf**;

DEFAULT —

CONTEXT: location

Turns on the HTTP interface of upstream configuration in the surrounding location. Access to this location should be [limited](#).

Configuration commands can be used to:

- view the group configuration;
- view, modify, or remove a server;
- add a new server.

Since addresses in a group are not required to be unique, specific servers in a group are referenced by their IDs. IDs are assigned automatically and shown when adding a new server or viewing the group configuration.

A configuration command consists of parameters passed as request arguments, for example:

```
http://127.0.0.1/upstream_conf?upstream=backend
```

The following parameters are supported:

`stream=`

Selects a [stream](#) upstream server group. Without this parameter, selects an [http](#) upstream server group.

`upstream=name`

Selects a group to work with. This parameter is mandatory.

`id=number`

Selects a server for viewing, modifying, or removing.

`remove=`

Removes a server from the group.

`add=`

Adds a new server to the group.

`backup=`

Required to add a backup server.

Before version 1.7.2, `backup=` was also required to view, modify, or remove existing backup servers.

`server=address`

Same as the “address” parameter of the [http](#) or [stream](#) upstream server. When adding a server, it is possible to specify it as a domain name. In this case, changes of the IP addresses that correspond to a domain name will be monitored and automatically applied to the upstream configuration without the need of restarting nginx (1.7.2). This requires the “resolver” directive in the [http](#) or [stream](#) block. See also the “resolve” parameter of the [http](#) or [stream](#) upstream server.

`service=name`

Same as the “service” parameter of the [http](#) or [stream](#) upstream server (1.9.13).

`weight=number`

Same as the “weight” parameter of the [http](#) or [stream](#) upstream server.

`max_conns=number`

Same as the “max\_conns” parameter of the [http](#) or [stream](#) upstream server.

`max_fails=number`

Same as the “max\_fails” parameter of the [http](#) or [stream](#) upstream server.

`fail_timeout=`*time*

Same as the “fail\_timeout” parameter of the [http](#) or [stream](#) upstream server.

`slow_start=`*time*

Same as the “slow\_start” parameter of the [http](#) or [stream](#) upstream server.

`down=`

Same as the “down” parameter of the [http](#) or [stream](#) upstream server.

`drain=`

Puts the [http](#) upstream server into the “draining” mode (1.7.5). In this mode, only requests [bound](#) to the server will be proxied to it.

`up=`

The opposite of the “down” parameter of the [http](#) or [stream](#) upstream server.

`route=`*string*

Same as the “route” parameter of the [http](#) upstream server.

The first three parameters select an object. This can be either the whole [http](#) or [stream](#) upstream server group, or a specific server. Without other parameters, the configuration of the selected group or server is shown.

For example, to view the configuration of the whole group, send:

```
http://127.0.0.1/upstream_conf?upstream=backend
```

To view the configuration of a specific server, also specify its ID:

```
http://127.0.0.1/upstream_conf?upstream=backend&id=42
```

To add a new server, specify its address in the “server=” parameter. Without other parameters specified, a server will be added with other parameters set to their default values (see the [http](#) or [stream](#) “server” directive).

For example, to add a new primary server, send:

```
http://127.0.0.1/upstream_conf?add=&upstream=backend&server=127.0.0.1:8080
```

To add a new backup server, send:

```
http://127.0.0.1/upstream_conf?add=&upstream=backend&backup=&server=127.0.0.1:8080
```

To add a new primary server, set its parameters to non-default values and mark it as “down”, send:

```
http://127.0.0.1/upstream_conf?add=&upstream=backend&server=127.0.0.1:8080&weight=2&down=
```

To remove a server, specify its ID:

```
http://127.0.0.1/upstream_conf?remove=&upstream=backend&id=42
```

To mark an existing server as “down”, send:

```
http://127.0.0.1/upstream_conf?upstream=backend&id=42&down=
```

To modify the address of an existing server, send:

```
http://127.0.0.1/upstream_conf?upstream=backend&id=42&server=192.0.2.3:8123
```

To modify other parameters of an existing server, send:

```
http://127.0.0.1/upstream_conf?upstream=backend&id=42&max_fails=3&weight=4
```

The above examples are for an [http](#) upstream server group. Similar examples for a [stream](#) upstream server group require the “stream=” parameter.

## 2.53 Module ngx\_http\_upstream\_hc\_module

2.53.1 Summary	322
2.53.2 Example Configuration	322
2.53.3 Directives	323
health_check	323
match	324

### 2.53.1 Summary

The `ngx_http_upstream_hc_module` module allows enabling periodic health checks of the servers in a [group](#) referenced in the surrounding location. The server group must reside in the [shared memory](#).

If a health check fails, the server will be considered unhealthy. If several health checks are defined for the same group of servers, a single failure of any check will make the corresponding server be considered unhealthy. Client requests are not passed to unhealthy servers and servers in the “checking” state.

Please note that most of the variables will have empty values when used with health checks.

This module is available as part of our [commercial subscription](#).

### 2.53.2 Example Configuration

```
upstream dynamic {
    zone upstream_dynamic 64k;

    server backend1.example.com      weight=5;
    server backend2.example.com:8080 fail_timeout=5s slow_start=30s;
    server 192.0.2.1                  max_fails=3;

    server backup1.example.com:8080  backup;
    server backup2.example.com:8080  backup;
}

server {
    location / {
        proxy_pass http://dynamic;
        health_check;
    }
}
```

With this configuration, nginx will send “/” requests to each server in the backend group every five seconds. If any communication error or timeout occurs, or a proxied server responds with the status code other than 2xx or 3xx, the health check will fail, and the server will be considered unhealthy.

Health checks can be configured to test the status code of a response, presence of certain header fields and their values, and the body contents. Tests are configured separately using the [match](#) directive and referenced in the `match` parameter of the [health\\_check](#) directive:

```
http {
    server {
        ...
        location / {
            proxy_pass http://backend;
            health_check match=welcome;
        }

        match welcome {
            status 200;
            header Content-Type = text/html;
            body ~ "Welcome to nginx!";
        }
    }
}
```

This configuration shows that in order for a health check to pass, the response to a health check request should succeed, have status 200, and contain “Welcome to nginx!” in the body.

### 2.53.3 Directives

#### health\_check

SYNTAX: **health\_check** [*parameters*];

DEFAULT —

CONTEXT: location

Enables periodic health checks of the servers in a [group](#) referenced in the surrounding location.

The following optional parameters are supported:

*interval=time*

sets the interval between two consecutive health checks, by default, 5 seconds.

*jitter=time*

sets the time within which each health check will be randomly delayed, by default, there is no delay.

*fails=number*

sets the number of consecutive failed health checks of a particular server after which this server will be considered unhealthy, by default, 1.

*passes=number*

sets the number of consecutive passed health checks of a particular server after which the server will be considered healthy, by default, 1.

*uri=uri*

defines the URI used in health check requests, by default, “/”.

*mandatory*

sets the initial “checking” state for a server until the first health check is completed (1.11.7). Client requests are not passed to servers in the “checking” state. If the parameter is not specified, the server will be initially considered healthy.

`match=name`

specifies the `match` block configuring the tests that a response should pass in order for a health check to pass. By default, the response should have status code 2xx or 3xx.

`port=number`

defines the port used when connecting to a server to perform a health check (1.9.7). By default, equals the [server](#) port.

## **match**

SYNTAX: **match** *name* { ... }

DEFAULT —

CONTEXT: http

Defines the named test set used to verify responses to health check requests. The following items can be tested in a response:

```
status 200;
    status is 200
status ! 500;
    status is not 500
status 200 204;
    status is 200 or 204
status ! 301 302;
    status is neither 301 nor 302
status 200-399;
    status is in the range from 200 to 399
status ! 400-599;
    status is not in the range from 400 to 599
status 301-303 307;
    status is either 301, 302, 303, or 307

header Content-Type = text/html;
    header contains Content-Type with value text/html
header Content-Type != text/html;
    header contains Content-Type with value other than text/html
header Connection ~ close;
    header contains Connection with value matching regular expression
    close
header Connection !~ close;
    header contains Connection with value not matching regular
    expression close
header Host;
    header contains Host
header ! X-Accel-Redirect;
    header lacks X-Accel-Redirect

body ~ "Welcome to nginx!";
```

body matches regular expression “Welcome to nginx!”  
body !~ "Welcome to nginx!";  
body does not match regular expression “Welcome to nginx!”

require *\$variable* ...;  
all specified variables are not empty and not equal to “0” (1.15.9).

If several tests are specified, the response matches only if it matches all tests.

Only the first 256k of the response body are examined.

Examples:

```
# status is 200, content type is "text/html",  
# and body contains "Welcome to nginx!"  
match welcome {  
    status 200;  
    header Content-Type = text/html;  
    body ~ "Welcome to nginx!";  
}
```

```
# status is not one of 301, 302, 303, or 307, and header does not have "Refresh  
:"  
match not_redirect {  
    status ! 301-303 307;  
    header ! Refresh;  
}
```

```
# status ok and not in maintenance mode  
match server_ok {  
    status 200-399;  
    body !~ "maintenance mode";  
}
```

```
# status is 200 or 204  
map $upstream_status $good_status {  
    200 1;  
    204 1;  
}  
  
match server_ok {  
    require $good_status;  
}
```

## 2.54 Module ngx\_http\_userid\_module

2.54.1 Summary	326
2.54.2 Example Configuration	326
2.54.3 Directives	326
userid	326
userid_domain	327
userid_expires	327
userid_mark	327
userid_name	327
userid_p3p	328
userid_path	328
userid_service	328
2.54.4 Embedded Variables	328

### 2.54.1 Summary

The `ngx_http_userid_module` module sets cookies suitable for client identification. Received and set cookies can be logged using the embedded variables `$uid_got` and `$uid_set`. This module is compatible with the `mod_uid` module for Apache.

### 2.54.2 Example Configuration

```
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"';
```

### 2.54.3 Directives

#### userid

SYNTAX: **userid** on | v1 | log | off;

DEFAULT 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: **userid\_domain** *name* | none;  
DEFAULT 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: **userid\_expires** *time* | max | off;  
DEFAULT 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\_mark**

SYNTAX: **userid\_mark** *letter* | *digit* | = | off;  
DEFAULT off  
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 uid  
CONTEXT: http, server, location

Sets the cookie name.

### userid\_p3p

SYNTAX: **userid\_p3p** *string* | none;  
DEFAULT 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: **userid\_path** *path*;  
DEFAULT /  
CONTEXT: http, server, location

Defines a path for which the cookie is set.

### userid\_service

SYNTAX: **userid\_service** *number*;  
DEFAULT 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.

## 2.54.4 Embedded Variables

The ngx\_http\_userid\_module module supports the following embedded 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.

## 2.55 Module ngx\_http\_uwsgi\_module

2.55.1	Summary	330
2.55.2	Example Configuration	330
2.55.3	Directives	330
	uwsgi_bind	330
	uwsgi_buffer_size	331
	uwsgi_buffering	331
	uwsgi_buffers	331
	uwsgi_busy_buffers_size	332
	uwsgi_cache	332
	uwsgi_cache_background_update	332
	uwsgi_cache_bypass	332
	uwsgi_cache_key	333
	uwsgi_cache_lock	333
	uwsgi_cache_lock_age	333
	uwsgi_cache_lock_timeout	333
	uwsgi_cache_max_range_offset	334
	uwsgi_cache_methods	334
	uwsgi_cache_min_uses	334
	uwsgi_cache_path	334
	uwsgi_cache_purge	336
	uwsgi_cache_revalidate	337
	uwsgi_cache_use_stale	337
	uwsgi_cache_valid	337
	uwsgi_connect_timeout	338
	uwsgi_force_ranges	339
	uwsgi_hide_header	339
	uwsgi_ignore_client_abort	339
	uwsgi_ignore_headers	339
	uwsgi_intercept_errors	340
	uwsgi_limit_rate	340
	uwsgi_max_temp_file_size	340
	uwsgi_modifier1	340
	uwsgi_modifier2	341
	uwsgi_next_upstream	341
	uwsgi_next_upstream_timeout	342
	uwsgi_next_upstream_tries	342
	uwsgi_no_cache	342
	uwsgi_param	342
	uwsgi_pass	343
	uwsgi_pass_header	343
	uwsgi_pass_request_body	344
	uwsgi_pass_request_headers	344
	uwsgi_read_timeout	344
	uwsgi_request_buffering	344

<a href="#">uwsgi_send_timeout</a>	345
<a href="#">uwsgi_socket_keepalive</a>	345
<a href="#">uwsgi_ssl_certificate</a>	345
<a href="#">uwsgi_ssl_certificate_key</a>	345
<a href="#">uwsgi_ssl_ciphers</a>	345
<a href="#">uwsgi_ssl_crl</a>	346
<a href="#">uwsgi_ssl_name</a>	346
<a href="#">uwsgi_ssl_password_file</a>	346
<a href="#">uwsgi_ssl_protocols</a>	346
<a href="#">uwsgi_ssl_server_name</a>	347
<a href="#">uwsgi_ssl_session_reuse</a>	347
<a href="#">uwsgi_ssl_trusted_certificate</a>	347
<a href="#">uwsgi_ssl_verify</a>	347
<a href="#">uwsgi_ssl_verify_depth</a>	347
<a href="#">uwsgi_store</a>	348
<a href="#">uwsgi_store_access</a>	348
<a href="#">uwsgi_temp_file_write_size</a>	349
<a href="#">uwsgi_temp_path</a>	349

### 2.55.1 Summary

The `ngx_http_uwsgi_module` module allows passing requests to a uwsgi server.

### 2.55.2 Example Configuration

```
location / {
    include    uwsgi_params;
    uwsgi_pass localhost:9000;
}
```

### 2.55.3 Directives

#### `uwsgi_bind`

SYNTAX: **uwsgi\_bind** *address* [transparent] | off;

DEFAULT —

CONTEXT: http, server, location

Makes outgoing connections to a uwsgi server originate from the specified local IP address with an optional port (1.11.2). Parameter value can contain variables (1.3.12). The special value `off` (1.3.12) cancels the effect of the `uwsgi_bind` directive inherited from the previous configuration level, which allows the system to auto-assign the local IP address and port.

The `transparent` parameter (1.11.0) 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 nginx worker processes with the [superuser](#) privileges. On Linux it is not required (1.13.8) as if the `transparent` parameter is specified, worker processes inherit the `CAP_NET_RAW` capability from the master process. It is also necessary to configure kernel routing table to intercept network traffic from the uwsgi server.

### uwsgi\_buffer\_size

SYNTAX: **uwsgi\_buffer\_size** *size*;  
DEFAULT 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** on | off;  
DEFAULT on  
CONTEXT: http, server, location

Enables or disables buffering of responses from the uwsgi server.

When buffering is enabled, nginx 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.

When buffering is disabled, the response is passed to a client synchronously, immediately as it is received. nginx will not try to read the whole response from the uwsgi server. The maximum size of the data that nginx can receive from the server at a time is set by the [uwsgi\\_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 [uwsgi\\_ignore\\_headers](#) directive.

### uwsgi\_buffers

SYNTAX: **uwsgi\_buffers** *number size*;  
DEFAULT 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 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 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 (1.7.9). The `off` parameter disables caching inherited from the previous configuration level.

### uwsgi\_cache\_background\_update

SYNTAX: **uwsgi\_cache\_background\_update** on | off;  
 DEFAULT off  
 CONTEXT: http, server, location  
 THIS DIRECTIVE APPEARED IN VERSION 1.11.10.

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.

### uwsgi\_cache\_bypass

SYNTAX: **uwsgi\_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:

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

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 off

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.1.12.

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: **uwsgi\_cache\_lock\_age** *time*;

DEFAULT 5s

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.7.8.

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: **uwsgi\_cache\_lock\_timeout** *time*;

DEFAULT 5s

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.1.12.

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.

Before 1.7.8, the response could be cached.

## uwsgi\_cache\_max\_range\_offset

SYNTAX: **uwsgi\_cache\_max\_range\_offset** *number*;

DEFAULT —

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.11.6.

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 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 1

CONTEXT: http, server, location

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

## uwsgi\_cache\_path

SYNTAX: **uwsgi\_cache\_path** *path* [levels=*levels*]  
[use\_temp\_path=on|off] keys\_zone=*name:size* [inactive=*time*]  
[max\_size=*size*] [manager\_files=*number*] [manager\_sleep=*time*]  
[manager\_threshold=*time*] [loader\_files=*number*]  
[loader\_sleep=*time*] [loader\_threshold=*time*]  
[purger=on|off] [purger\_files=*number*] [purger\_sleep=*time*]  
[purger\_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/nginx/cache levels=1:2 keys_zone=one:10m;
```

file names in a cache will look like this:

```
/data/nginx/cache/c/29/b7f54b2df7773722d382f4809d65029c
```

A cached response is first written to a temporary file, and then the file is renamed. Starting from version 0.8.9, 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 (1.7.10). If this parameter is omitted or set to the value `on`, the directory set by the [uwsgi\\_temp\\_path](#) directive for the given location will be used. If the value is set to `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.

As part of [commercial subscription](#), the shared memory zone also stores extended cache [information](#), thus, it is required to specify a larger zone size for the same number of keys. For example, one megabyte zone can store about 4 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.

The special “cache manager” process monitors the maximum cache size set by the `max_size` parameter. When this size is exceeded, it removes the least recently used data. The data is removed in iterations configured by `manager_files`, `manager_threshold`, and `manager_sleep` parameters (1.11.5). During one iteration no more than `manager_files` items are deleted (by default, 100). The duration of one iteration is limited by the `manager_threshold` parameter (by default, 200 milliseconds). Between iterations, a pause configured by the `manager_sleep` parameter (by default, 50 milliseconds) is made.

A minute after the start 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. During one iteration no more than `loader_files` items are loaded (by default, 100). Besides, the duration of one iteration is limited by the `loader_threshold` parameter (by default, 200 milliseconds). Between iterations, a pause configured by the `loader_sleep` parameter (by default, 50 milliseconds) is made.

Additionally, the following parameters are available as part of our [commercial subscription](#):

```
purger=on|off
```

Instructs whether cache entries that match a [wildcard key](#) will be removed from the disk by the cache purger (1.7.12). Setting the parameter to `on` (default is `off`) will activate the “cache purger” process that permanently iterates through all cache entries and deletes the entries that match the wildcard key.

`purger_files=number`

Sets the number of items that will be scanned during one iteration (1.7.12). By default, `purger_files` is set to 10.

`purger_threshold=number`

Sets the duration of one iteration (1.7.12). By default, `purger_threshold` is set to 50 milliseconds.

`purger_sleep=number`

Sets a pause between iterations (1.7.12). By default, `purger_sleep` is set to 50 milliseconds.

In versions 1.7.3, 1.7.7, and 1.11.10 cache header format has been changed. Previously cached responses will be considered invalid after upgrading to a newer nginx version.

## uwsgi\_cache\_purge

SYNTAX: **uwsgi\_cache\_purge**string ...;

DEFAULT —

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.5.7.

Defines conditions under which the request will be considered a cache purge request. If at least one value of the string parameters is not empty and is not equal to “0” then the cache entry with a corresponding [cache key](#) is removed. The result of successful operation is indicated by returning the 204 No Content response.

If the [cache key](#) of a purge request ends with an asterisk (“\*”), all cache entries matching the wildcard key will be removed from the cache. However, these entries will remain on the disk until they are deleted for either [inactivity](#), or processed by the [cache purger](#) (1.7.12), or a client attempts to access them.

Example configuration:

```
uwsgi_cache_path /data/nginx/cache keys_zone=cache_zone:10m;

map $request_method $purge_method {
    PURGE    1;
    default 0;
}

server {
    ...
    location / {
        uwsgi_pass          backend;
        uwsgi_cache          cache_zone;
        uwsgi_cache_key      $uri;
        uwsgi_cache_purge    $purge_method;
    }
}
```

```
}

```

This functionality is available as part of our [commercial subscription](#).

### **uwsgi\_cache\_revalidate**

SYNTAX: **uwsgi\_cache\_revalidate** on | off;

DEFAULT off

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.5.7.

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 off

CONTEXT: http, server, location

Determines in which cases a stale cached response can be used when an error occurs during communication with the uwsgi server. The directive's parameters match the parameters of the [uwsgi\\_next\\_upstream](#) directive.

The error parameter also permits using a stale cached response if a uwsgi server to process a request cannot be selected.

Additionally, the updating 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 (1.11.10). This has lower priority than using the directive parameters.

- 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

```
uwsgi_cache_valid 200 302 10m;  
uwsgi_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

```
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;
```

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 (1.7.7). If the header includes the `Vary` field with another value, such a response will be cached taking into account the corresponding request header fields (1.7.7).

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 60s

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\_force\_ranges

SYNTAX: **uwsgi\_force\_ranges** on | off;

DEFAULT off

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.7.7.

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: **uwsgi\_hide\_header** *field*;

DEFAULT —

CONTEXT: http, server, location

By default, nginx 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: **uwsgi\_ignore\_client\_abort** on | off;

DEFAULT off

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` (1.1.6), `X-Accel-Buffering` (1.1.6), `X-Accel-Charset` (1.1.6), `Expires`, `Cache-Control`, `Set-Cookie` (0.8.44), and `Vary` (1.7.7).

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: **`uwsgi_intercept_errors`** `on` | `off`;  
DEFAULT `off`  
CONTEXT: `http`, `server`, `location`

Determines whether a uwsgi server responses with codes greater than or equal to 300 should be passed to a client or be intercepted and redirected to nginx for processing with the [error\\_page](#) directive.

### `uwsgi_limit_rate`

SYNTAX: **`uwsgi_limit_rate`** *rate*;  
DEFAULT `0`  
CONTEXT: `http`, `server`, `location`  
THIS DIRECTIVE APPEARED IN VERSION 1.7.7.

Limits the speed of reading the response from the uwsgi server. The *rate* is specified in bytes per second. The zero value disables rate limiting. The limit is set per a request, and so if nginx 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 `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.

The zero value disables buffering of responses to temporary files.

This restriction does not apply to responses that will be [cached](#) or [stored](#) on disk.

### `uwsgi_modifier1`

SYNTAX: **`uwsgi_modifier1`** *number*;  
DEFAULT `0`  
CONTEXT: `http`, `server`, `location`

Sets the value of the `modifier1` field in the [uwsgi packet header](#).

### **uwsgi\_modifier2**

SYNTAX: **uwsgi\_modifier2** *number*;

DEFAULT 0

CONTEXT: http, server, location

Sets the value of the `modifier2` 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 error timeout

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 (1.11.13);

#### **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 (1.9.13); 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. The cases of `error`, `timeout` and `invalid_header` are always considered unsuccessful attempts, even if they are not specified in the directive. The cases of `http_500`, `http_503`, and `http_429` are considered unsuccessful attempts only if they are specified in the directive. The cases of `http_403` and `http_404` are 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 0  
CONTEXT: http, server, location  
THIS DIRECTIVE APPEARED IN VERSION 1.7.5.

Limits the time during which a request can be passed to the [next server](#). The 0 value turns off this limitation.

### **uwsgi\_next\_upstream\_tries**

SYNTAX: **uwsgi\_next\_upstream\_tries** *number*;  
DEFAULT 0  
CONTEXT: http, server, location  
THIS DIRECTIVE APPEARED IN VERSION 1.7.5.

Limits the number of possible tries for passing a request to the [next server](#). The 0 value 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:

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

Can be used along with the [uwsgi\\_cache\\_bypass](#) directive.

### **uwsgi\_param**

SYNTAX: **uwsgi\_param** *parameter value* [*if\_not\_empty*];  
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 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` (1.1.11) 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: **uwsgi\_pass** [*protocol://*]*address*;

DEFAULT —

CONTEXT: location, if in location

Sets the protocol and address of a uwsgi server. 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](#).

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](#).

Secured uwsgi protocol is supported since version 1.5.8.

## 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: **uwsgi\_pass\_request\_body** on | off;

DEFAULT on

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 on

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](#) directive.

### **uwsgi\_read\_timeout**

SYNTAX: **uwsgi\_read\_timeout** *time*;

DEFAULT 60s

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 on

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.7.11.

Enables or disables buffering of a client request body.

When buffering is enabled, the entire request body is [read](#) from the client before sending the request to a uwsgi server.

When buffering is disabled, the request body is sent to the uwsgi server immediately as it is received. In this case, the request cannot be passed to the [next server](#) if nginx 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 60s

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 off

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.15.6.

Configures the “TCP keepalive” behavior for outgoing connections to a uwsgi server. By default, the operating system’s settings are in effect for the socket. If the directive is set to the value “on”, the `SO_KEEPALIVE` socket option is turned on for the socket.

## uwsgi\_ssl\_certificate

SYNTAX: **uwsgi\_ssl\_certificate** *file*;

DEFAULT —

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.7.8.

Specifies a *file* with the certificate in the PEM format used for authentication to a secured uwsgi server.

## uwsgi\_ssl\_certificate\_key

SYNTAX: **uwsgi\_ssl\_certificate\_key** *file*;

DEFAULT —

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.7.8.

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* (1.7.9), which loads a secret key with a specified *id* from the OpenSSL engine *name*.

## uwsgi\_ssl\_ciphers

SYNTAX: **uwsgi\_ssl\_ciphers** *ciphers*;

DEFAULT DEFAULT

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.5.8.

Specifies the enabled ciphers for requests to a secured uwsgi server. The ciphers are specified in the format understood by the OpenSSL library.

The full list can be viewed using the “`openssl ciphers`” command.

### **uwsgi\_ssl\_crl**

SYNTAX: **uwsgi\_ssl\_crl** *file*;  
DEFAULT —  
CONTEXT: http, server, location  
THIS DIRECTIVE APPEARED IN VERSION 1.7.0.

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 host from uwsgi\_pass  
CONTEXT: http, server, location  
THIS DIRECTIVE APPEARED IN VERSION 1.7.0.

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 from [uwsgi\\_pass](#) is used.

### **uwsgi\_ssl\_password\_file**

SYNTAX: **uwsgi\_ssl\_password\_file** *file*;  
DEFAULT —  
CONTEXT: http, server, location  
THIS DIRECTIVE APPEARED IN VERSION 1.7.8.

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 TLSv1 TLSv1.1 TLSv1.2  
CONTEXT: http, server, location  
THIS DIRECTIVE APPEARED IN VERSION 1.5.8.

Enables the specified protocols for requests to a secured uwsgi server.

## uwsgi\_ssl\_server\_name

SYNTAX: **uwsgi\_ssl\_server\_name** on | off;

DEFAULT off

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.7.0.

Enables or disables passing of the server name through [TLS Server Name Indication extension](#) (SNI, RFC 6066) when establishing a connection with the secured uwsgi server.

## uwsgi\_ssl\_session\_reuse

SYNTAX: **uwsgi\_ssl\_session\_reuse** on | off;

DEFAULT on

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.5.8.

Determines whether SSL sessions can be reused when working with a secured uwsgi server. If the errors “SSL3\_GET\_FINISHED:digest check failed” appear in the logs, try disabling session reuse.

## uwsgi\_ssl\_trusted\_certificate

SYNTAX: **uwsgi\_ssl\_trusted\_certificate** *file*;

DEFAULT —

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.7.0.

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: **uwsgi\_ssl\_verify** on | off;

DEFAULT off

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.7.0.

Enables or disables verification of the secured uwsgi server certificate.

## uwsgi\_ssl\_verify\_depth

SYNTAX: **uwsgi\_ssl\_verify\_depth** *number*;

DEFAULT 1

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.7.0.

Sets the verification depth in the secured uwsgi server certificates chain.

## uwsgi\_store

SYNTAX: **uwsgi\_store** on | off | *string*;

DEFAULT off

CONTEXT: http, server, location

Enables saving of files to a disk. The `on` parameter saves files with paths corresponding to the directives `alias` or `root`. The `off` parameter disables saving of files. In addition, 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. 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 `uwsgi_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;
    error_page    404 = /fetch$uri;
}

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 user:rw

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 8k|16k

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 uwsgi\_temp

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/nginx/uwsgi_temp 1 2;
```

a temporary file might look like this:

```
/spool/nginx/uwsgi_temp/7/45/00000123457
```

See also the `use_temp_path` parameter of the [uwsgi\\_cache\\_path](#) directive.

## 2.56 Module ngx\_http\_v2\_module

2.56.1	Summary	350
2.56.2	Known Issues	350
2.56.3	Example Configuration	350
2.56.4	Directives	351
	<a href="#">http2_body_preread_size</a>	351
	<a href="#">http2_chunk_size</a>	351
	<a href="#">http2_idle_timeout</a>	351
	<a href="#">http2_max_concurrent_pushes</a>	351
	<a href="#">http2_max_concurrent_streams</a>	351
	<a href="#">http2_max_field_size</a>	352
	<a href="#">http2_max_header_size</a>	352
	<a href="#">http2_max_requests</a>	352
	<a href="#">http2_push</a>	352
	<a href="#">http2_push_preload</a>	353
	<a href="#">http2_recv_buffer_size</a>	353
	<a href="#">http2_recv_timeout</a>	353
2.56.5	Embedded Variables	353

### 2.56.1 Summary

The `ngx_http_v2_module` module (1.9.5) provides support for [HTTP/2](#) and supersedes the `ngx_http_spdy_module` module.

This module is not built by default, it should be enabled with the `--with-http_v2_module` configuration parameter.

### 2.56.2 Known Issues

Before version 1.9.14, buffering of a client request body could not be disabled regardless of [proxy\\_request\\_buffering](#), [fastcgi\\_request\\_buffering](#), [uwsgi\\_request\\_buffering](#), and [scgi\\_request\\_buffering](#) directive values.

### 2.56.3 Example Configuration

```
server {
    listen 443 ssl http2;

    ssl_certificate server.crt;
    ssl_certificate_key server.key;
}
```

Note that accepting HTTP/2 connections over TLS requires the “Application-Layer Protocol Negotiation” (ALPN) TLS extension support, which is available only since [OpenSSL](#) version 1.0.2. Using the “Next Protocol Negotiation” (NPN) TLS extension for this purpose (available since OpenSSL version 1.0.1) is not guaranteed to work.

Also note that if the [ssl\\_prefer\\_server\\_ciphers](#) directive is set to the value “on”, the [ciphers](#) should be configured to comply with [RFC 7540, Appendix A](#) black list and supported by clients.

## 2.56.4 Directives

### http2\_body\_preread\_size

SYNTAX: **http2\_body\_preread\_size** *size*;  
DEFAULT 64k  
CONTEXT: http, server  
THIS DIRECTIVE APPEARED IN VERSION 1.11.0.

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 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\_idle\_timeout

SYNTAX: **http2\_idle\_timeout** *time*;  
DEFAULT 3m  
CONTEXT: http, server

Sets the timeout of inactivity after which the connection is closed.

### http2\_max\_concurrent\_pushes

SYNTAX: **http2\_max\_concurrent\_pushes** *number*;  
DEFAULT 10  
CONTEXT: http, server  
THIS DIRECTIVE APPEARED IN VERSION 1.13.9.

Limits the maximum number of concurrent [push](#) requests in a connection.

### http2\_max\_concurrent\_streams

SYNTAX: **http2\_max\_concurrent\_streams** *number*;  
DEFAULT 128  
CONTEXT: http, server

Sets the maximum number of concurrent HTTP/2 streams in a connection.

## http2\_max\_field\_size

SYNTAX: **http2\_max\_field\_size** *size*;  
DEFAULT 4k  
CONTEXT: http, server

Limits the maximum size of an [HPACK](#)-compressed request header field. The limit applies equally to both name and value. Note that if Huffman encoding is applied, the actual size of decompressed name and value strings may be larger. For most requests, the default limit should be enough.

## http2\_max\_header\_size

SYNTAX: **http2\_max\_header\_size** *size*;  
DEFAULT 16k  
CONTEXT: http, server

Limits the maximum size of the entire request header list after [HPACK](#) decompression. For most requests, the default limit should be enough.

## http2\_max\_requests

SYNTAX: **http2\_max\_requests** *number*;  
DEFAULT 1000  
CONTEXT: http, server  
THIS DIRECTIVE APPEARED IN VERSION 1.11.6.

Sets the maximum number of requests (including [push](#) requests) that can be served through one HTTP/2 connection, after which the next client request will lead to connection closing and the need of establishing a new connection.

## http2\_push

SYNTAX: **http2\_push** *uri* | *off*;  
DEFAULT *off*  
CONTEXT: http, server, location  
THIS DIRECTIVE APPEARED IN VERSION 1.13.9.

Pre-emptively 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

SYNTAX: **http2\_push\_preload** on | off;

DEFAULT off

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.13.9.

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 256k

CONTEXT: http

Sets the size of the per [worker](#) input buffer.

## http2\_recv\_timeout

SYNTAX: **http2\_recv\_timeout** *time*;

DEFAULT 30s

CONTEXT: http, server

Sets the timeout for expecting more data from the client, after which the connection is closed.

## 2.56.5 Embedded Variables

The `ngx_http_v2_module` module supports the following embedded variables:

*\$http2*

negotiated protocol identifier: “h2” for HTTP/2 over TLS, “h2c” for HTTP/2 over cleartext TCP, or an empty string otherwise.

## 2.57 Module ngx\_http\_xslt\_module

2.57.1 Summary	354
2.57.2 Example Configuration	354
2.57.3 Directives	354
xml_entities	354
xslt_last_modified	355
xslt_param	355
xslt_string_param	355
xslt_stylesheet	355
xslt_types	356

### 2.57.1 Summary

The `ngx_http_xslt_module` (0.7.8+) is a filter that transforms XML responses using one or more XSLT stylesheets.

This module is not built by default, it should be enabled with the `--with-http_xslt_module` configuration parameter.

This module requires the [libxml2](#) and [libxslt](#) libraries.

### 2.57.2 Example Configuration

```
location / {
    xml_entities    /site/dtd/entities.dtd;
    xslt_stylesheet /site/xslt/one.xslt param=value;
    xslt_stylesheet /site/xslt/two.xslt;
}
```

### 2.57.3 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 off

CONTEXT: http, server, location

THIS DIRECTIVE APPEARED IN VERSION 1.5.1.

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

THIS DIRECTIVE APPEARED IN VERSION 1.1.18.

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 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

THIS DIRECTIVE APPEARED IN VERSION 1.1.18.

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 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:

```
location / {
    xslt_stylesheet /site/xslt/one.xslt
                   $arg_xslt_params
                   param1=' $value1':param2=value2
                   param3=value3;
}
```

It is possible to specify several stylesheets. They will be applied sequentially in the specified order.

### **xslt\_types**

SYNTAX: **xslt\_types** *mime-type* ...;

DEFAULT `text/xml`

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 (0.8.29). If the transformation result is an HTML response, its MIME type is changed to “`text/html`”.

# Chapter 3

## Stream server modules

### 3.1 Module ngx\_stream\_core\_module

3.1.1	Summary	357
3.1.2	Example Configuration	357
3.1.3	Directives	358
	listen	358
	preread_buffer_size	360
	preread_timeout	360
	proxy_protocol_timeout	361
	resolver	361
	resolver_timeout	361
	server	362
	stream	362
	tcp_nodelay	362
	variables_hash_bucket_size	362
	variables_hash_max_size	362
3.1.4	Embedded Variables	363

#### 3.1.1 Summary

The `ngx_stream_core_module` module is available since version 1.9.0. This module is not built by default, it should be enabled with the `--with-stream` configuration parameter.

#### 3.1.2 Example Configuration

```
worker_processes auto;

error_log /var/log/nginx/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;
    }
}

```

### 3.1.3 Directives

#### listen

SYNTAX: **listen** *address:port* [ssl] [udp] [proxy\_protocol]  
 [backlog=*number*] [rcvbuf=*size*] [sndbuf=*size*] [bind]  
 [ipv6only=on|off] [reuseport]  
 [so\_keepalive=on|off|[*keepidle*]:[*keepintvl*]:[*keepcnt*]];

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. 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/nginx.sock;

```

Port ranges (1.15.10) are specified with the first and last port separated by a hyphen:

```
listen 127.0.0.1:12345-12399;  
listen 12345-12399;
```

The `ssl` parameter allows specifying that all connections accepted on this port should work in SSL mode.

The `udp` parameter configures a listening socket for working with datagrams (1.9.13).

The `proxy_protocol` parameter (1.11.4) allows specifying that all connections accepted on this port should use the [PROXY protocol](#).

The PROXY protocol version 2 is supported since version 1.13.11.

The `listen` directive can have several additional parameters specific to socket-related system calls.

`backlog=number`

sets the `backlog` parameter in the `listen` call that limits the maximum length for the queue of pending connections (1.9.2). 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 (1.11.13).

`sndbuf=size`

sets the send buffer size (the `SO_SNDBUF` option) for the listening socket (1.11.13).

`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`), nginx 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 `ipv6only` 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.

`reuseport`

this parameter (1.9.1) 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+ (1.15.1).

Inappropriate use of this option may have its security [implications](#).

`so_keepalive=on|off[[keepidle]:[keepintvl]:[keepcnt]`

this parameter configures the “TCP keepalive” behavior for the listening socket. If this parameter is omitted then the operating system’s settings will be in effect for the socket. If it is set to the value “on”, the `SO_KEEPALIVE` option is turned on for the socket. If it is set to the value “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_KEEPIIDLE`, `TCP_KEEPIINTVL`, 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_KEEPIIDLE`) to 30 minutes, leave the probe interval (`TCP_KEEPIINTVL`) at its system default, and set the probes count (`TCP_KEEPCNT`) to 10 probes.

Different servers must listen on different *address:port* pairs.

### **preread\_buffer\_size**

SYNTAX: **preread\_buffer\_size** *size*;

DEFAULT 16k

CONTEXT: stream, server

THIS DIRECTIVE APPEARED IN VERSION 1.11.5.

Specifies a *size* of the preread buffer.

### **preread\_timeout**

SYNTAX: **preread\_timeout** *timeout*;

DEFAULT 30s

CONTEXT: stream, server

THIS DIRECTIVE APPEARED IN VERSION 1.11.5.

Specifies a *timeout* of the preread phase.

## proxy\_protocol\_timeout

SYNTAX: **proxy\_protocol\_timeout** *timeout*;

DEFAULT 30s

CONTEXT: stream, server

THIS DIRECTIVE APPEARED IN VERSION 1.11.4.

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* ... [valid=*time*] [ipv6=on|off];

DEFAULT —

CONTEXT: stream, server

THIS DIRECTIVE APPEARED IN VERSION 1.11.3.

Configures name servers used to resolve names of upstream servers into addresses, for example:

```
resolver 127.0.0.1 [::1]:5353;
```

An address can be specified as a domain name or IP address, and an optional port. If port is not specified, the port 53 is used. Name servers are queried in a round-robin fashion.

By default, nginx will look up both IPv4 and IPv6 addresses while resolving. If looking up of IPv6 addresses is not desired, the `ipv6=off` parameter can be specified.

By default, nginx caches answers using the TTL value of a response. The optional `valid` parameter allows overriding it:

```
resolver 127.0.0.1 [::1]:5353 valid=30s;
```

Before version 1.11.3, this directive was available as part of our [commercial subscription](#).

## resolver\_timeout

SYNTAX: **resolver\_timeout** *time*;

DEFAULT 30s

CONTEXT: stream, server

THIS DIRECTIVE APPEARED IN VERSION 1.11.3.

Sets a timeout for name resolution, for example:

```
resolver_timeout 5s;
```

Before version 1.11.3, this directive was available as part of our [commercial subscription](#).

### server

SYNTAX: **server** { ... }

DEFAULT —

CONTEXT: stream

Sets the configuration for a server.

### 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 on

CONTEXT: stream, server

THIS DIRECTIVE APPEARED IN VERSION 1.9.4.

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 64

CONTEXT: stream

THIS DIRECTIVE APPEARED IN VERSION 1.11.2.

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 1024

CONTEXT: stream

THIS DIRECTIVE APPEARED IN VERSION 1.11.2.

Sets the maximum *size* of the variables hash table. The details of setting up hash tables are provided in a separate [document](#).

### 3.1.4 Embedded Variables

The `ngx_stream_core_module` module supports variables since 1.11.2.

*\$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 (1.11.4)

*\$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

*\$nginx\_version*

nginx version

*\$pid*

PID of the worker process

*\$protocol*

protocol used to communicate with the client: TCP or UDP (1.11.4)

*\$proxy\_protocol\_addr*

client address from the PROXY protocol header, or an empty string otherwise (1.11.4)

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, or an empty string otherwise (1.11.4)

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 (1.11.4);

*\$status*

session status (1.11.4), 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 connections](#)

*\$time\_iso8601*

local time in the ISO 8601 standard format

*\$time\_local*

local time in the Common Log Format

## 3.2 Module ngx\_stream\_access\_module

3.2.1	Summary	365
3.2.2	Example Configuration	365
3.2.3	Directives	365
	allow	365
	deny	365

### 3.2.1 Summary

The `ngx_stream_access_module` module (1.9.2) allows limiting access to certain client addresses.

### 3.2.2 Example Configuration

```
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`.

### 3.2.3 Directives

#### allow

SYNTAX: **allow** *address* | *CIDR* | `unix:` | `all`;

DEFAULT —

CONTEXT: stream, server

Allows access for the specified network or address. If the special value `unix:` is specified, allows access for all UNIX-domain sockets.

#### deny

SYNTAX: **deny** *address* | *CIDR* | `unix:` | `all`;

DEFAULT —

CONTEXT: stream, server

Denies access for the specified network or address. If the special value `unix:` is specified, denies access for all UNIX-domain sockets.

## 3.3 Module ngx\_stream\_geo\_module

3.3.1	Summary	366
3.3.2	Example Configuration	366
3.3.3	Directives	366
	geo	366

### 3.3.1 Summary

The `ngx_stream_geo_module` module (1.11.3) creates variables with values depending on the client IP address.

### 3.3.2 Example Configuration

```
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;
}
```

### 3.3.3 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 {
    ...;
}
```

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`

a 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:

```
geo $country {
    default      ZZ;
    include      conf/geo.conf;
    delete      127.0.0.0/16;

    127.0.0.0/24  US;
    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:

```
geo $country {
    ranges;
    default      ZZ;
    127.0.0.0-127.0.0.0    US;
    127.0.0.1-127.0.0.1    RU;
    127.0.0.1-127.0.0.255  US;
    10.1.0.0-10.1.255.255  RU;
    192.168.1.0-192.168.1.255 UK;
}
```

## 3.4 Module ngx\_stream\_geoip\_module

3.4.1	Summary	368
3.4.2	Example Configuration	368
3.4.3	Directives	368
	<a href="#">geoip_country</a>	368
	<a href="#">geoip_city</a>	369
	<a href="#">geoip_org</a>	370

### 3.4.1 Summary

The `ngx_stream_geoip_module` module (1.11.3) 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.

This module is not built by default, it should be enabled with the `--with-stream_geoip_module` configuration parameter.

This module requires the [MaxMind GeoIP](#) library.

### 3.4.2 Example Configuration

```
stream {
    geoip_country      GeoIP.dat;
    geoip_city         GeoLiteCity.dat;

    map $geoip_city_continent_code $nearest_server {
        default        example.com;
        EU             eu.example.com;
        NA             na.example.com;
        AS             as.example.com;
    }
    ...
}
```

### 3.4.3 Directives

#### `geoip_country`

SYNTAX: **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\_code*

two-letter country code, for example, “RU”, “US”.

*\$geoip\_country\_code3*

three-letter country code, for example, “RUS”, “USA”.

*\$geoip\_country\_name*

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:

*\$geoip\_area\_code*

telephone area code (US only).

This variable may contain outdated information since the corresponding database field is deprecated.

*\$geoip\_city\_continent\_code*

two-letter continent code, for example, “EU”, “NA”.

*\$geoip\_city\_country\_code*

two-letter country code, for example, “RU”, “US”.

*\$geoip\_city\_country\_code3*

three-letter country code, for example, “RUS”, “USA”.

*\$geoip\_city\_country\_name*

country name, for example, “Russian Federation”, “United States”.

*\$geoip\_dma\_code*

DMA region code in US (also known as “metro code”), according to the [geotargeting](#) in Google AdWords API.

*\$geoip\_latitude*

latitude.

*\$geoip\_longitude*

longitude.

*\$geoip\_region*

two-symbol country region code (region, territory, state, province, federal land and the like), for example, “48”, “DC”.

*\$geoip\_region\_name*

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\_code*

postal code.

**geoip\_org**

SYNTAX: **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:

*\$geoip\_org*

organization name, for example, “The University of Melbourne”.

## 3.5 Module ngx\_stream\_js\_module

3.5.1	Summary	371
3.5.2	Example Configuration	371
3.5.3	Directives	372
	js_access	372
	js_filter	373
	js_include	373
	js_path	373
	js_preread	373
	js_set	373
3.5.4	Session Object Properties	373

### 3.5.1 Summary

The ngx\_stream\_js\_module module is used to implement handlers in njs — a subset of the JavaScript language.

This module is not built by default. Download and install instructions are available [here](#).

### 3.5.2 Example Configuration

```
load_module modules/ngx_stream_js_module.so;
...

stream {
    js_include stream.js;

    js_set $bar bar;
    js_set $req_line req_line;

    server {
        listen 12345;

        js_preread preread;
        return $req_line;
    }

    server {
        listen 12346;

        js_access access;
        proxy_pass 127.0.0.1:8000;
        js_filter 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.abort();
        return;
    }

    s.allow();
}

```

### 3.5.3 Directives

#### js\_access

SYNTAX: **js\_access** *function*;

DEFAULT —

CONTEXT: stream, server

Sets an njs function which will be called at the access phase.

### **js\_filter**

SYNTAX: **js\_filter** *function*;

DEFAULT —

CONTEXT: stream, server

Sets a data filter.

### **js\_include**

SYNTAX: **js\_include** *file*;

DEFAULT —

CONTEXT: stream

Specifies a file that implements server and variable handlers in njs.

### **js\_path**

SYNTAX: **js\_path** *path*;

DEFAULT —

CONTEXT: http

THIS DIRECTIVE APPEARED IN VERSION 0.3.0.

Sets an additional path for njs modules.

### **js\_preread**

SYNTAX: **js\_preread** *function*;

DEFAULT —

CONTEXT: stream, server

Sets an njs function which will be called at the preread phase.

### **js\_set**

SYNTAX: **js\_set** *\$variable function*;

DEFAULT —

CONTEXT: stream

Sets an njs function for the specified variable.

## **3.5.4 Session Object Properties**

Each stream njs handler receives one argument, a stream session object.

## 3.6 Module ngx\_stream\_keyval\_module

3.6.1	Summary	374
3.6.2	Example Configuration	374
3.6.3	Directives	374
	keyval	374
	keyval_zone	375

### 3.6.1 Summary

The ngx\_stream\_keyval\_module module (1.13.7) creates variables with values taken from key-value pairs managed by the [API](#).

This module is available as part of our [commercial subscription](#).

### 3.6.2 Example Configuration

```
http {
    server {
        ...
        location /api {
            api write=on;
        }
    }
}

stream {

    keyval_zone zone=one:32k state=one.keyval;
    keyval      $ssl_server_name $name zone=one;

    server {
        listen          12345 ssl;
        proxy_pass       $name;
        ssl_certificate  /usr/local/nginx/conf/cert.pem;
        ssl_certificate_key /usr/local/nginx/conf/cert.key;
    }
}
```

### 3.6.3 Directives

#### keyval

SYNTAX: **keyval** *key* *\$variable* zone=*name*;

DEFAULT —

CONTEXT: stream

Creates a new *\$variable* whose value is looked up by the *key* in the key-value database. Strings are matched ignoring the case. The database is stored in a shared memory zone specified by the zone parameter.

## keyval\_zone

SYNTAX: **keyval\_zone** zone=*name:size* [state=*file*] [timeout=*time*] [sync];

DEFAULT —

CONTEXT: stream

Sets the *name* and *size* of the shared memory zone that keeps the key-value database. Key-value pairs are managed by the [API](#).

The optional *state* parameter specifies a *file* that keeps the current state of the key-value database in the JSON format and makes it persistent across nginx restarts.

The optional *timeout* parameter (1.15.0) sets the time after which key-value pairs are removed from the zone.

The optional *sync* parameter (1.15.0) enables [synchronization](#) of the shared memory zone. The synchronization requires the *timeout* parameter to be set.

If the synchronization is enabled, removal of key-value pairs (no matter [one](#) or [all](#)) will be performed only on a target cluster node. The same key-value pairs on other cluster nodes will be removed upon *timeout*.

## 3.7 Module ngx\_stream\_limit\_conn\_module

3.7.1	Summary	376
3.7.2	Example Configuration	376
3.7.3	Directives	376
	<a href="#">limit_conn</a>	376
	<a href="#">limit_conn_log_level</a>	377
	<a href="#">limit_conn_zone</a>	377

### 3.7.1 Summary

The `ngx_stream_limit_conn_module` module (1.9.3) is used to limit the number of connections per the defined key, in particular, the number of connections from a single IP address.

### 3.7.2 Example Configuration

```
stream {
    limit_conn_zone $binary_remote_addr zone=addr:10m;

    ...

    server {
        ...

        limit_conn          addr 1;
        limit_conn_log_level error;
    }
}
```

### 3.7.3 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 an IP address at a time.

When several `limit_conn` directives are specified, any configured limit will apply.

The directives are inherited from the previous level if and only if there are no `limit_conn` directives on the current level.

### **limit\_conn\_log\_level**

SYNTAX: **limit\_conn\_log\_level** info | notice | warn | error;

DEFAULT error

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 (1.11.2). Connections with an empty key value are not accounted. Usage example:

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

Here, the key is a client IP address 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.

## 3.8 Module ngx\_stream\_log\_module

3.8.1	Summary	378
3.8.2	Example Configuration	378
3.8.3	Directives	378
	access_log	378
	log_format	379
	open_log_file_cache	379

### 3.8.1 Summary

The ngx\_stream\_log\_module module (1.11.4) writes session logs in the specified format.

### 3.8.2 Example Configuration

```
log_format basic '$remote_addr [$time_local] '
                '$protocol $status $bytes_sent $bytes_received '
                '$session_time';

access_log /spool/logs/nginx-access.log basic buffer=32k;
```

### 3.8.3 Directives

#### access\_log

SYNTAX: **access\_log** *path format* [buffer=*size*] [gzip[=*level*]]  
[flush=*time*] [if=*condition*];

SYNTAX: **access\_log** off;

DEFAULT off

CONTEXT: stream, server

Sets the path, [format](#), and configuration for a buffered log write. Several logs can be specified on the same 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.

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 basic gzip flush=5m;
```

For gzip compression to work, nginx 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

Specifies the log format, for example:

```
log_format proxy '$remote_addr [$time_local] '
                 '$protocol $status $bytes_sent $bytes_received '
                 '$session_time "$upstream_addr" '
                 '"$upstream_bytes_sent" "$upstream_bytes_received" '
                 '$upstream_connect_time";
```

The `escape` parameter (1.11.8) allows setting `json` or `default` characters escaping in variables, by default, `default` escaping is used. The `none` parameter (1.13.10) disables escaping.

## open\_log\_file\_cache

SYNTAX: **open\_log\_file\_cache** *max=N* [`inactive=time`] [`min_uses=N`] [`valid=time`];

SYNTAX: **open\_log\_file\_cache** *off*;

DEFAULT *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;
```

## 3.9 Module ngx\_stream\_map\_module

3.9.1	Summary	381
3.9.2	Example Configuration	381
3.9.3	Directives	381
	map	381
	map_hash_bucket_size	382
	map_hash_max_size	383

### 3.9.1 Summary

The ngx\_stream\_map\_module module (1.11.2) creates variables whose values depend on values of other variables.

### 3.9.2 Example Configuration

```
map $remote_addr $limit {
    127.0.0.1      "";
    default        $binary_remote_addr;
}

limit_conn_zone $limit zone=addr:10m;
limit_conn addr 1;
```

### 3.9.3 Directives

#### map

SYNTAX: **map** *string* *\$variable* { ... }

DEFAULT —

CONTEXT: stream

Creates a new variable whose value depends on values of one or more of the source variables specified in the first parameter.

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 connection 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 from the “~” symbol for a case-sensitive matching, or from 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 `default` 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:

```
*.example.com 1;
example.*     1;
```

The following two records

```
example.com 1;
*.example.com 1;
```

can be combined:

```
.example.com 1;
```

This parameter should be specified before the list of values.

`include` *file*

includes a file with values. There can be several inclusions.

`volatile`

indicates that the variable is not cacheable (1.11.7).

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

## **map\_hash\_bucket\_size**

SYNTAX: **map\_hash\_bucket\_size** *size*;

DEFAULT 32 | 64 | 128

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 in a separate [document](#).

### **map\_hash\_max\_size**

SYNTAX: **map\_hash\_max\_size** *size*;

DEFAULT 2048

CONTEXT: stream

Sets the maximum *size* of the [map](#) variables hash tables. The details of setting up hash tables are provided in a separate [document](#).

## 3.10 Module ngx\_stream\_proxy\_module

3.10.1	Summary	384
3.10.2	Example Configuration	384
3.10.3	Directives	385
	proxy_bind	385
	proxy_buffer_size	385
	proxy_connect_timeout	386
	proxy_download_rate	386
	proxy_next_upstream	386
	proxy_next_upstream_timeout	386
	proxy_next_upstream_tries	386
	proxy_pass	387
	proxy_protocol	387
	proxy_requests	387
	proxy_responses	388
	proxy_session_drop	388
	proxy_socket_keepalive	388
	proxy_ssl	388
	proxy_ssl_certificate	389
	proxy_ssl_certificate_key	389
	proxy_ssl_ciphers	389
	proxy_ssl_crl	389
	proxy_ssl_name	389
	proxy_ssl_password_file	390
	proxy_ssl_protocols	390
	proxy_ssl_server_name	390
	proxy_ssl_session_reuse	390
	proxy_ssl_trusted_certificate	390
	proxy_ssl_verify	391
	proxy_ssl_verify_depth	391
	proxy_timeout	391
	proxy_upload_rate	391

### 3.10.1 Summary

The ngx\_stream\_proxy\_module module (1.9.0) allows proxying data streams over TCP, UDP (1.9.13), and UNIX-domain sockets.

### 3.10.2 Example Configuration

```
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;
}
```

### 3.10.3 Directives

#### proxy\_bind

SYNTAX: **proxy\_bind** *address* [transparent] | off;

DEFAULT —

CONTEXT: stream, server

THIS DIRECTIVE APPEARED IN VERSION 1.9.2.

Makes outgoing connections to a proxied server originate from the specified local IP *address*. Parameter value can contain variables (1.11.2). 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 (1.11.0) 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 nginx worker processes with the [superuser](#) privileges. On Linux it is not required (1.13.8) as if the `transparent` parameter is specified, worker processes inherit the `CAP_NET_RAW` capability from the master process. It is also necessary to configure kernel routing table to intercept network traffic from the proxied server.

#### proxy\_buffer\_size

SYNTAX: **proxy\_buffer\_size** *size*;

DEFAULT 16k

CONTEXT: stream, server

THIS DIRECTIVE APPEARED IN VERSION 1.9.4.

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: **proxy\_connect\_timeout** *time*;

DEFAULT 60s

CONTEXT: stream, server

Defines a timeout for establishing a connection with a proxied server.

### proxy\_download\_rate

SYNTAX: **proxy\_download\_rate** *rate*;

DEFAULT 0

CONTEXT: stream, server

THIS DIRECTIVE APPEARED IN VERSION 1.9.3.

Limits the speed of reading the data from the proxied server. The *rate* is specified in bytes per second. The zero value disables rate limiting. The limit is set per a connection, so if nginx simultaneously opens two connections to the proxied server, the overall rate will be twice as much as the specified limit.

### proxy\_next\_upstream

SYNTAX: **proxy\_next\_upstream** on | off;

DEFAULT on

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.

Passing a connection 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 0

CONTEXT: stream, server

Limits the time allowed to pass a connection to the [next server](#). The 0 value turns off this limitation.

### proxy\_next\_upstream\_tries

SYNTAX: **proxy\_next\_upstream\_tries** *number*;

DEFAULT 0

CONTEXT: stream, server

Limits the number of possible tries for passing a connection to the [next server](#). The 0 value 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 (1.11.3):

```
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 off

CONTEXT: stream, server

THIS DIRECTIVE APPEARED IN VERSION 1.9.2.

Enables the [PROXY protocol](#) for connections to a proxied server.

## proxy\_requests

SYNTAX: **proxy\_requests** *number*;

DEFAULT 0

CONTEXT: stream, server

THIS DIRECTIVE APPEARED IN VERSION 1.15.7.

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

THIS DIRECTIVE APPEARED IN VERSION 1.9.13.

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\_session\_drop

SYNTAX: **proxy\_session\_drop** on | off;

DEFAULT off

CONTEXT: stream, server

THIS DIRECTIVE APPEARED IN VERSION 1.15.8.

Enables terminating all sessions to a proxied server after it was removed from the group or marked as permanently unavailable. This can occur because of [re-resolve](#) or with the API [DELETE](#) command. A server can be marked as permanently unavailable if it is considered [unhealthy](#) or with the API [PATCH](#) command. Each session is terminated when the next read or write event is processed for the client or proxied server.

This directive is available as part of our [commercial subscription](#).

## proxy\_socket\_keepalive

SYNTAX: **proxy\_socket\_keepalive** on | off;

DEFAULT off

CONTEXT: stream, server

THIS DIRECTIVE APPEARED IN VERSION 1.15.6.

Configures the “TCP keepalive” behavior for outgoing connections to a proxied server. By default, the operating system’s settings are in effect for the socket. If the directive is set to the value “on”, the `SO_KEEPALIVE` socket option is turned on for the socket.

## proxy\_ssl

SYNTAX: **proxy\_ssl** on | off;

DEFAULT off

CONTEXT: stream, server

Enables the SSL/TLS protocol for connections to a proxied server.

### proxy\_ssl\_certificate

SYNTAX: **proxy\_ssl\_certificate** *file*;

DEFAULT —

CONTEXT: stream, server

Specifies a *file* with the certificate in the PEM format used for authentication to a proxied server.

### proxy\_ssl\_certificate\_key

SYNTAX: **proxy\_ssl\_certificate\_key** *file*;

DEFAULT —

CONTEXT: stream, server

Specifies a *file* with the secret key in the PEM format used for authentication to a proxied server.

### proxy\_ssl\_ciphers

SYNTAX: **proxy\_ssl\_ciphers** *ciphers*;

DEFAULT DEFAULT

CONTEXT: stream, server

Specifies the enabled ciphers for connections to a proxied server. The ciphers are specified in the format understood by the OpenSSL library.

The full list can be viewed using the “`openssl ciphers`” command.

### 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\_ssl\_name** *name*;

DEFAULT host from proxy\_pass

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 (1.11.3).

By default, the host part of the [proxy\\_pass](#) address is used.

### 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 TLSv1 TLSv1.1 TLSv1.2

CONTEXT: stream, server

Enables the specified protocols for connections to a proxied server.

### proxy\_ssl\_server\_name

SYNTAX: **proxy\_ssl\_server\_name** on | off;

DEFAULT off

CONTEXT: stream, server

Enables or disables passing of the server name through [TLS Server Name Indication extension](#) (SNI, RFC 6066) when establishing a connection with the proxied server.

### proxy\_ssl\_session\_reuse

SYNTAX: **proxy\_ssl\_session\_reuse** on | off;

DEFAULT on

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: **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 off

CONTEXT: stream, server

Enables or disables verification of the proxied server certificate.

### **proxy\_ssl\_verify\_depth**

SYNTAX: **proxy\_ssl\_verify\_depth** *number*;

DEFAULT 1

CONTEXT: stream, server

Sets the verification depth in the proxied server certificates chain.

### **proxy\_timeout**

SYNTAX: **proxy\_timeout** *timeout*;

DEFAULT 10m

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.

### **proxy\_upload\_rate**

SYNTAX: **proxy\_upload\_rate** *rate*;

DEFAULT 0

CONTEXT: stream, server

THIS DIRECTIVE APPEARED IN VERSION 1.9.3.

Limits the speed of reading the data from the client. The *rate* is specified in bytes per second. The zero value disables rate limiting. 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.

## 3.11 Module ngx\_stream\_realip\_module

3.11.1 Summary	392
3.11.2 Example Configuration	392
3.11.3 Directives	392
set_real_ip_from	392
3.11.4 Embedded Variables	392

### 3.11.1 Summary

The `ngx_stream_realip_module` module is used to change the client address and port to the ones sent in the PROXY protocol header (1.11.4). The PROXY protocol must be previously enabled by setting the `proxy_protocol` parameter in the `listen` directive.

This module is not built by default, it should be enabled with the `--with-stream_realip_module` configuration parameter.

### 3.11.2 Example Configuration

```
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;
```

### 3.11.3 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.

### 3.11.4 Embedded Variables

*\$realip\_remote\_addr*

keeps the original client address

*\$realip\_remote\_port*

keeps the original client port

## 3.12 Module ngx\_stream\_return\_module

3.12.1 Summary	393
3.12.2 Example Configuration	393
3.12.3 Directives	393
return	393

### 3.12.1 Summary

The ngx\_stream\_return\_module module (1.11.2) allows sending a specified value to the client and then closing the connection.

### 3.12.2 Example Configuration

```
server {
    listen 12345;
    return $time_iso8601;
}
```

### 3.12.3 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.

## 3.13 Module ngx\_stream\_split\_clients\_module

3.13.1 Summary	394
3.13.2 Example Configuration	394
3.13.3 Directives	394
split_clients	394

### 3.13.1 Summary

The `ngx_stream_split_clients_module` module (1.11.3) creates variables suitable for A/B testing, also known as split testing.

### 3.13.2 Example Configuration

```
stream {
    ...
    split_clients "${remote_addr}AAA" $upstream {
        0.5%      feature_test1;
        2.0%      feature_test2;
        *         production;
    }

    server {
        ...
        proxy_pass $upstream;
    }
}
```

### 3.13.3 Directives

#### split\_clients

SYNTAX: **split\_clients** *string \$variable* { ... }

DEFAULT —

CONTEXT: stream

Creates a variable for A/B testing, for example:

```
split_clients "${remote_addr}AAA" $variant {
    0.5%      .one;
    2.0%      .two;
    *         "";
}
```

The value of the original string is hashed using MurmurHash2. In the example given, hash values from 0 to 21474835 (0.5%) correspond to the value `".one"` of the `$variant` variable, hash values from 21474836 to 107374180 (2%) correspond to the value `".two"`, and hash values from 107374181 to 4294967295 correspond to the value `""` (an empty string).

## 3.14 Module ngx\_stream\_ssl\_module

3.14.1	Summary	395
3.14.2	Example Configuration	395
3.14.3	Directives	396
	ssl_certificate	396
	ssl_certificate_key	397
	ssl_ciphers	397
	ssl_client_certificate	397
	ssl_crl	398
	ssl_dhparam	398
	ssl_ecdh_curve	398
	ssl_handshake_timeout	398
	ssl_password_file	399
	ssl_prefer_server_ciphers	399
	ssl_protocols	399
	ssl_session_cache	400
	ssl_session_ticket_key	400
	ssl_session_tickets	401
	ssl_session_timeout	401
	ssl_trusted_certificate	401
	ssl_verify_client	401
	ssl_verify_depth	402
3.14.4	Embedded Variables	402

### 3.14.1 Summary

The `ngx_stream_ssl_module` module (1.9.0) provides the necessary support for a stream proxy server to work with the SSL/TLS protocol. This module is not built by default, it should be enabled with the `--with-stream_ssl_module` configuration parameter.

### 3.14.2 Example Configuration

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):

```
worker_processes auto;

stream {
    ...
}
```

```
server {
    listen          12345 ssl;

    ssl_protocols   TLSv1 TLSv1.1 TLSv1.2;
    ssl_ciphers     AES128-SHA:AES256-SHA:RC4-SHA:DES-CBC3-SHA:RC4-MD5;
    ssl_certificate  /usr/local/nginx/conf/cert.pem;
    ssl_certificate_key /usr/local/nginx/conf/cert.key;
    ssl_session_cache    shared:SSL:10m;
    ssl_session_timeout 10m;

    ...
}
```

### 3.14.3 Directives

#### 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.

Since version 1.11.0, this directive can be specified multiple times to load certificates of different types, for example, RSA and ECDSA:

```
server {
    listen          12345 ssl;

    ssl_certificate    example.com.rsa.crt;
    ssl_certificate_key example.com.rsa.key;

    ssl_certificate    example.com.ecdsa.crt;
    ssl_certificate_key example.com.ecdsa.key;

    ...
}
```

Only OpenSSL 1.0.2 or higher supports separate certificate chains for different certificates. With older versions, only one certificate chain can be used.

Since version 1.15.9, 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* (1.15.10), 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: **ssl\_certificate\_key** *file*;

DEFAULT —

CONTEXT: stream, 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:$variable` can be specified instead of the *file* (1.15.10), 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](#).

Since version 1.15.9, variables can be used in the *file* name when using OpenSSL 1.0.2 or higher.

### ssl\_ciphers

SYNTAX: **ssl\_ciphers** *ciphers*;

DEFAULT 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 full list can be viewed using the “`openssl ciphers`” command.

### ssl\_client\_certificate

SYNTAX: **ssl\_client\_certificate** *file*;

DEFAULT —

CONTEXT: stream, server

THIS DIRECTIVE APPEARED IN VERSION 1.11.8.

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\_crl

SYNTAX: **ssl\_crl** *file*;

DEFAULT —

CONTEXT: stream, server

THIS DIRECTIVE APPEARED IN VERSION 1.11.8.

Specifies a *file* with revoked certificates (CRL) in the PEM format used to [verify](#) client certificates.

## ssl\_dhparam

SYNTAX: **ssl\_dhparam** *file*;

DEFAULT —

CONTEXT: stream, server

Specifies a *file* with DH parameters for DHE ciphers.

By default no parameters are set, and therefore DHE ciphers will not be used.

Prior to version 1.11.0, builtin parameters were used by default.

## ssl\_ecdh\_curve

SYNTAX: **ssl\_ecdh\_curve** *curve*;

DEFAULT `auto`

CONTEXT: stream, server

Specifies a *curve* for ECDHE ciphers.

When using OpenSSL 1.0.2 or higher, it is possible to specify multiple curves (1.11.0), for example:

```
ssl_ecdh_curve prime256v1:secp384r1;
```

The special value `auto` (1.11.0) instructs nginx to use a list built into the OpenSSL library when using OpenSSL 1.0.2 or higher, or `prime256v1` with older versions.

Prior to version 1.11.0, the `prime256v1` curve was used by default.

## ssl\_handshake\_timeout

SYNTAX: **ssl\_handshake\_timeout** *time*;

DEFAULT `60s`

CONTEXT: stream, server

Specifies a timeout for the SSL handshake to complete.

## 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 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 TLSv1 TLSv1.1 TLSv1.2

CONTEXT: stream, server

Enables the specified protocols.

The TLSv1.1 and TLSv1.2 parameters work only when OpenSSL 1.0.1 or higher is used.

The TLSv1.3 parameter (1.13.0) works only when OpenSSL 1.1.1 built with TLSv1.3 support is used.

## ssl\_session\_cache

SYNTAX: **ssl\_session\_cache** *off* | *none* | [*builtin[:size]*]  
           [*shared:name:size*];

DEFAULT *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: nginx explicitly tells a client that sessions may not be reused.

*none*

the use of a session cache is gently disallowed: nginx 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.

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, 1.11.8) or AES128 (for 48-byte keys) is used for encryption.

### ssl\_session\_tickets

SYNTAX: **ssl\_session\_tickets** on | off;

DEFAULT on

CONTEXT: stream, server

Enables or disables session resumption through [TLS session tickets](#).

### ssl\_session\_timeout

SYNTAX: **ssl\_session\_timeout** *time*;

DEFAULT 5m

CONTEXT: stream, server

Specifies a time during which a client may reuse the session parameters.

### ssl\_trusted\_certificate

SYNTAX: **ssl\_trusted\_certificate** *file*;

DEFAULT —

CONTEXT: stream, server

THIS DIRECTIVE APPEARED IN VERSION 1.11.8.

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 off

CONTEXT: stream, server

THIS DIRECTIVE APPEARED IN VERSION 1.11.8.

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.

The `optional` parameter requests the client certificate and verifies it if the certificate is present.

The `optional_no_ca` parameter 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 nginx performs the actual

certificate verification. The contents of the certificate is accessible through the `$ssl_client_cert` variable.

### `ssl_verify_depth`

SYNTAX: **`ssl_verify_depth`** *number*;

DEFAULT 1

CONTEXT: stream, server

THIS DIRECTIVE APPEARED IN VERSION 1.11.8.

Sets the verification depth in the client certificates chain.

## 3.14.4 Embedded Variables

The `ngx_stream_ssl_module` module supports variables since 1.11.2.

`$ssl_cipher`

returns the string of ciphers used for an established SSL connection;

`$ssl_ciphers`

returns the list of ciphers supported by the client (1.11.7). Known ciphers are listed by names, unknown are shown in hexadecimal, for example:

```
AES128-SHA:AES256-SHA:0x00ff
```

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 (1.11.8);

`$ssl_client_fingerprint`

returns the SHA1 fingerprint of the client certificate for an established SSL connection (1.11.8);

`$ssl_client_i_dn`

returns the “issuer DN” string of the client certificate for an established SSL connection according to [RFC 2253](#) (1.11.8);

`$ssl_client_raw_cert`

returns the client certificate in the PEM format for an established SSL connection (1.11.8);

`$ssl_client_s_dn`

returns the “subject DN” string of the client certificate for an established SSL connection according to [RFC 2253](#) (1.11.8);

`$ssl_client_serial`

returns the serial number of the client certificate for an established SSL connection (1.11.8);

*\$ssl\_client\_v\_end*

returns the end date of the client certificate (1.11.8);

*\$ssl\_client\_v\_remain*

returns the number of days until the client certificate expires (1.11.8);

*\$ssl\_client\_v\_start*

returns the start date of the client certificate (1.11.8);

*\$ssl\_client\_verify*

returns the result of client certificate verification (1.11.8): “SUCCESS”, “FAILED:reason”, and “NONE” if a certificate was not present;

*\$ssl\_curves*

returns the list of curves supported by the client (1.11.7). Known curves are listed by names, unknown are shown in hexadecimal, for example:

0x001d:prime256v1:secp521r1:secp384r1

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\_protocol*

returns the protocol of an established SSL connection;

*\$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.

## 3.15 Module ngx\_stream\_ssl\_preread\_module

3.15.1 Summary	404
3.15.2 Example Configuration	404
3.15.3 Directives	405
ssl_preread	405
3.15.4 Embedded Variables	405

### 3.15.1 Summary

The `ngx_stream_ssl_preread_module` module (1.11.5) allows extracting information from the [ClientHello](#) message without terminating SSL/TLS, for example, the sever name requested through [SNI](#) or protocols advertised in [ALPN](#). This module is not built by default, it should be enabled with the `--with-stream-ssl_preread_module` configuration parameter.

### 3.15.2 Example Configuration

Selecting an upstream based on 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 an upstream based on protocol:

```
map $ssl_preread_alpn_protocols $proxy {
    ~\bh2\b          127.0.0.1:8001;
    ~\bhttp/1.1\b    127.0.0.1:8002;
    ~\bxmlpp-client\b 127.0.0.1:8003;
}

server {
    listen      9000;
    proxy_pass  $proxy;
    ssl_preread on;
}
```

Selecting an upstream based on SSL protocol version:

```
map $ssl_preread_protocol $upstream {
    ""      ssh.example.com:22;
    "TLSv1.2" new.example.com:443;
    default  tls.example.com:443;
}

# ssh and https on the same port
server {
    listen      192.168.0.1:443;
    proxy_pass  $upstream;
    ssl_preread on;
}
```

### 3.15.3 Directives

#### ssl\_preread

SYNTAX: **ssl\_preread** on | off;

DEFAULT off

CONTEXT: stream, server

Enables extracting information from the ClientHello message at the preread phase.

### 3.15.4 Embedded Variables

*\$ssl\_preread\_protocol*

the highest SSL protocol version supported by the client (1.15.2)

*\$ssl\_preread\_server\_name*

server name requested through SNI

*\$ssl\_preread\_alpn\_protocols*

list of protocols advertised by the client through ALPN (1.13.10). The values are separated by commas.

## 3.16 Module ngx\_stream\_upstream\_module

3.16.1 Summary	406
3.16.2 Example Configuration	406
3.16.3 Directives	407
upstream	407
server	407
zone	409
state	410
hash	410
least_conn	411
least_time	411
random	411
3.16.4 Embedded Variables	412

### 3.16.1 Summary

The ngx\_stream\_upstream\_module module (1.9.0) is used to define groups of servers that can be referenced by the [proxy\\_pass](#) directive.

### 3.16.2 Example Configuration

```
upstream backend {
    hash $remote_addr consistent;

    server backend1.example.com:12345 weight=5;
    server backend2.example.com:12345;
    server unix:/tmp/backend3;

    server backup1.example.com:12345 backup;
    server backup2.example.com:12345 backup;
}

server {
    listen 12346;
    proxy_pass backend;
}
```

Dynamically configurable group with periodic [health checks](#) is available as part of our [commercial subscription](#):

```
resolver 10.0.0.1;

upstream dynamic {
    zone upstream_dynamic 64k;

    server backend1.example.com:12345 weight=5;
    server backend2.example.com:12345 fail_timeout=5s slow_start=30s;
    server 192.0.2.1:12345 max_fails=3;
    server backend3.example.com:12345 resolve;
    server backend4.example.com service=http resolve;

    server backup1.example.com:12345 backup;
    server backup2.example.com:12345 backup;
}
```

```
server {  
    listen 12346;  
    proxy_pass dynamic;  
    health_check;  
}
```

### 3.16.3 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:

```
upstream backend {  
    server backend1.example.com:12345 weight=5;  
    server 127.0.0.1:12345 max_fails=3 fail_timeout=30s;  
    server unix:/tmp/backend2;  
    server backend3.example.com:12345 resolve;  
  
    server backup1.example.com:12345 backup;  
}
```

By default, connections are distributed between the servers using a weighted round-robin balancing method. In the above example, each 7 connections will be distributed as follows: 5 connections go to `backend1.example.com:12345` and one connection to each of the second and third servers. If an error occurs during communication with a server, the connection will be passed to the next server, and so on until all of the functioning servers will be tried. If communication with all servers fails, the connection will be closed.

#### 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:

`weight=number`

sets the weight of the server, by default, 1.

`max_conns=number`

limits the maximum *number* of simultaneous connections to the proxied server (1.11.5). Default value is zero, meaning there is no limit. If the server group does not reside in the [shared memory](#), the limitation works per each worker process.

Prior to version 1.11.5, this parameter was available as part of our [commercial subscription](#).

`max_fails=number`

sets the number of unsuccessful attempts to communicate with the server that should happen in the duration set by the `fail_timeout` parameter to consider the server unavailable for a duration also set by the `fail_timeout` parameter. By default, the number of unsuccessful attempts is set to 1. The zero value disables the accounting of attempts. Here, an unsuccessful attempt is an error or timeout while establishing a connection with the server.

`fail_timeout=time`

sets

- the time during which the specified number of unsuccessful attempts to communicate with the server should happen to consider the server unavailable;
- and the period of time the server will be considered unavailable.

By default, the parameter is set to 10 seconds.

`backup`

marks the server as a backup server. Connections to the backup server will be passed when the primary servers are unavailable.

`down`

marks the server as permanently unavailable.

Additionally, the following parameters are available as part of our [commercial subscription](#):

`resolve`

monitors changes of the IP addresses that correspond to a domain name of the server, and automatically modifies the upstream configuration without the need of restarting nginx. The server group must reside in the [shared memory](#).

In order for this parameter to work, the [resolver](#) directive must be specified in the [stream](#) block. Example:

```
stream {
    resolver 10.0.0.1;

    upstream u {
        zone ...;
        ...
        server example.com:12345 resolve;
    }
}
```

`service=name`

enables resolving of DNS [SRV](#) records and sets the service *name* (1.9.13). In order for this parameter to work, it is necessary to specify the [resolve](#) parameter for the server and specify a hostname without a port number. If the service name does not contain a dot (“.”), then the [RFC](#)-compliant name is constructed and the TCP protocol is added to the service prefix. For example, to look up the `_http._tcp.backend.example.com` SRV record, it is necessary to specify the directive:

```
server backend.example.com service=http resolve;
```

If the service name contains one or more dots, then the name is constructed by joining the service prefix and the server name. For example, to look up the `_http._tcp.backend.example.com` and `server1.backend.example.com` SRV records, it is necessary to specify the directives:

```
server backend.example.com service=_http._tcp resolve;  
server example.com service=server1.backend resolve;
```

Highest-priority SRV records (records with the same lowest-number priority value) are resolved as primary servers, the rest of SRV records are resolved as backup servers. If the [backup](#) parameter is specified for the server, high-priority SRV records are resolved as backup servers, the rest of SRV records are ignored.

`slow_start=time`

sets the *time* during which the server will recover its weight from zero to a nominal value, when unhealthy server becomes [healthy](#), or when the server becomes available after a period of time it was considered [unavailable](#). Default value is zero, i.e. slow start is disabled.

The parameter cannot be used along with the [hash](#) load balancing method.

If there is only a single server in a group, `max_fails`, `fail_timeout` and `slow_start` parameters are ignored, and such a server will never be considered unavailable.

## 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.

Additionally, as part of our [commercial subscription](#), such groups allow changing the group membership or modifying the settings of a particular server without the need of restarting nginx. The configuration is accessible via the [API](#) module (1.13.3).

Prior to version 1.13.3, the configuration was accessible only via a special location handled by [upstream\\_conf](#).

## state

SYNTAX: **state** *file*;

DEFAULT —

CONTEXT: upstream

THIS DIRECTIVE APPEARED IN VERSION 1.9.7.

Specifies a *file* that keeps the state of the dynamically configurable group. Examples:

```
state /var/lib/nginx/state/servers.conf; # path for Linux
state /var/db/nginx/state/servers.conf; # path for FreeBSD
```

The state is currently limited to the list of servers with their parameters. The file is read when parsing the configuration and is updated each time the upstream configuration is [changed](#). Changing the file content directly should be avoided. The directive cannot be used along with the [server](#) directive.

Changes made during [configuration reload](#) or [binary upgrade](#) can be lost.

This directive is available as part of our [commercial subscription](#).

## hash

SYNTAX: **hash** *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 (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

SYNTAX: **least\_time** connect | first\_byte | last\_byte [inflight];

DEFAULT —

CONTEXT: upstream

Specifies that a group should use a load balancing method where a connection is passed to the server with the least average time and 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.

If the *connect* parameter is specified, time to [connect](#) to the upstream server is used. If the *first\_byte* parameter is specified, time to receive the [first byte](#) of data is used. If the *last\_byte* is specified, time to receive the [last byte](#) of data is used. If the *inflight* parameter is specified (1.11.6), incomplete connections are also taken into account.

Prior to version 1.11.6, incomplete connections were taken into account by default.

This directive is available as part of our [commercial subscription](#).

### random

SYNTAX: **random** [two [*method*]];

DEFAULT —

CONTEXT: upstream

THIS DIRECTIVE APPEARED IN VERSION 1.15.1.

Specifies that a group should use a load balancing method where a connection is passed to a randomly selected server, taking into account weights of servers.

The optional *two* parameter instructs nginx to randomly select [two](#) servers and then choose a server using the specified *method*. The default method is

`least_conn` which passes a connection to a server with the least number of active connections.

The `least_time` method passes a connection to a server with the least average time and least number of active connections. If `least_time=connect` parameter is specified, time to [connect](#) to the upstream server is used. If `least_time=first_byte` parameter is specified, time to receive the [first byte](#) of data is used. If `least_time=last_byte` is specified, time to receive the [last byte](#) of data is used.

The `least_time` method is available as a part of our [commercial subscription](#).

### 3.16.4 Embedded Variables

The `ngx_stream_upstream_module` module supports the following embedded variables:

*`$upstream_addr`*

keeps the IP address and port, or the path to the UNIX-domain socket of the upstream server (1.11.4). If several servers were contacted during proxying, their addresses are separated by commas, e.g. “192.168.1.1:12345, 192.168.1.2:12345, 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 (1.11.4). 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 (1.11.4). 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 (1.11.4); 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 (1.11.4); 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`*

session duration in seconds with millisecond resolution (1.11.4). Times of several connections are separated by commas like addresses in the [\\$upstream\\_addr](#) variable.

## 3.17 Module ngx\_stream\_upstream\_hc\_module

3.17.1 Summary	413
3.17.2 Example Configuration	413
3.17.3 Directives	414
health_check	414
health_check_timeout	415
match	415

### 3.17.1 Summary

The `ngx_stream_upstream_hc_module` module (1.9.0) allows enabling periodic health checks of the servers in a [group](#). The server group must reside in the [shared memory](#).

If a health check fails, the server will be considered unhealthy. If several health checks are defined for the same group of servers, a single failure of any check will make the corresponding server be considered unhealthy. Client connections are not passed to unhealthy servers and servers in the “checking” state.

This module is available as part of our [commercial subscription](#).

### 3.17.2 Example Configuration

```
upstream tcp {
    zone upstream_tcp 64k;

    server backend1.example.com:12345 weight=5;
    server backend2.example.com:12345 fail_timeout=5s slow_start=30s;
    server 192.0.2.1:12345 max_fails=3;

    server backup1.example.com:12345 backup;
    server backup2.example.com:12345 backup;
}

server {
    listen 12346;
    proxy_pass tcp;
    health_check;
}
```

With this configuration, nginx will check the ability to establish a TCP connection to each server in the `tcp` group every five seconds. When a connection to the server cannot be established, the health check will fail, and the server will be considered unhealthy.

Health checks can be configured for the UDP protocol:

```
upstream dns_upstream {
```

```

zone    dns_zone 64k;

server dns1.example.com:53;
server dns2.example.com:53;
server dns3.example.com:53;
}

server {
    listen      53 udp;
    proxy_pass  dns_upstream;
    health_check udp;
}

```

In this case, the absence of ICMP “Destination Unreachable” message is expected in reply to the sent string “nginx health check”.

Health checks can also be configured to test data obtained from the server. Tests are configured separately using the [match](#) directive and referenced in the match parameter of the [health\\_check](#) directive.

### 3.17.3 Directives

#### health\_check

SYNTAX: **health\_check** [*parameters*];

DEFAULT —

CONTEXT: server

Enables periodic health checks of the servers in a [group](#).

The following optional parameters are supported:

*interval=time*

sets the interval between two consecutive health checks, by default, 5 seconds.

*jitter=time*

sets the time within which each health check will be randomly delayed, by default, there is no delay.

*fails=number*

sets the number of consecutive failed health checks of a particular server after which this server will be considered unhealthy, by default, 1.

*passes=number*

sets the number of consecutive passed health checks of a particular server after which the server will be considered healthy, by default, 1.

*mandatory*

sets the initial “checking” state for a server until the first health check is completed (1.11.7). Client connections are not passed to servers in the “checking” state. If the parameter is not specified, the server will be initially considered healthy.

*match=name*

specifies the [match](#) block configuring the tests that a successful connection should pass in order for a health check to pass. By default, for TCP, only the ability to establish a TCP connection with the server is checked. For [UDP](#), the absence of ICMP “Destination

Unreachable” message is expected in reply to the sent string “nginx health check”.

Prior to version 1.11.7, by default, UDP health check required a match block with the [send](#) and [expect](#) parameters.

`port=number`

defines the port used when connecting to a server to perform a health check (1.9.7). By default, equals the [server](#) port.

`udp`

specifies that the UDP protocol should be used for health checks instead of the default TCP protocol (1.9.13).

### health\_check\_timeout

SYNTAX: **health\_check\_timeout** *timeout*;

DEFAULT 5s

CONTEXT: stream, server

Overrides the [proxy\\_timeout](#) value for health checks.

### match

SYNTAX: **match** *name* { ... }

DEFAULT —

CONTEXT: stream

Defines the named test set used to verify server responses to health checks. The following parameters can be configured:

`send string;`

sends a *string* to the server;

`expect string | ~ regex;`

a literal string (1.9.12) or a regular expression that the data obtained from the server should match. The regular expression is specified with the preceding “~\*” modifier (for case-insensitive matching), or the “~” modifier (for case-sensitive matching).

Both `send` and `expect` parameters can contain hexadecimal literals with the prefix “\x” followed by two hex digits, for example, “\x80” (1.9.12).

Health check is passed if:

- the TCP connection was successfully established;
- the *string* from the `send` parameter, if specified, was sent;
- the data obtained from the server matched the string or regular expression from the `expect` parameter, if specified;
- the time elapsed does not exceed the value specified in the [health\\_check\\_timeout](#) directive.

Example:

```
upstream backend {
    zone    upstream_backend 10m;
    server  127.0.0.1:12345;
}

match http {
    send    "GET / HTTP/1.0\r\nHost: localhost\r\n\r\n";
    expect ~ "200 OK";
}

server {
    listen      12346;
    proxy_pass  backend;
    health_check match=http;
}
```

Only the first [proxy\\_buffer\\_size](#) bytes of data obtained from the server are examined.

## 3.18 Module ngx\_stream\_zone\_sync\_module

3.18.1	Summary	417
3.18.2	Example Configuration	417
3.18.3	Directives	418
	zone_sync	418
	zone_sync_buffers	419
	zone_sync_connect_retry_interval	419
	zone_sync_connect_timeout	419
	zone_sync_interval	419
	zone_sync_recv_buffer_size	419
	zone_sync_server	420
	zone_sync_ssl	420
	zone_sync_ssl_certificate	420
	zone_sync_ssl_certificate_key	421
	zone_sync_ssl_ciphers	421
	zone_sync_ssl_crl	421
	zone_sync_ssl_name	421
	zone_sync_ssl_password_file	421
	zone_sync_ssl_protocols	422
	zone_sync_ssl_server_name	422
	zone_sync_ssl_trusted_certificate	422
	zone_sync_ssl_verify	422
	zone_sync_ssl_verify_depth	422
	zone_sync_timeout	422
3.18.4	API endpoints	423
3.18.5	Starting, stopping, removing a cluster node	423

### 3.18.1 Summary

The `ngx_stream_zone_sync_module` module (1.13.8) provides the necessary support for synchronizing contents of [shared memory zones](#) between nodes of a cluster. To enable synchronization for a particular zone, a corresponding module must support this feature. Currently, it is possible to synchronize HTTP [sticky](#) sessions, information about [excessive HTTP requests](#), and key-value pairs both in [http](#) and [stream](#).

This module is available as part of our [commercial subscription](#).

### 3.18.2 Example Configuration

Minimal configuration:

```
http {
    ...

    upstream backend {
```

```

        server backend1.example.com:8080;
        server backend2.example.com:8081;

        sticky learn
            create=$upstream_cookie_examplecookie
            lookup=$cookie_examplecookie
            zone=client_sessions:1m sync;
    }
    ...
}

stream {
    ...

    server {
        zone_sync;

        listen 127.0.0.1:8090;

        # cluster of 2 nodes
        zone_sync_server a.example.com;
        zone_sync_server b.example.com;
    }
}

```

A more complex configuration with SSL enabled and with cluster members defined by DNS:

```

...
stream {
    ...

    resolver 127.0.0.1 valid=10s;

    server {
        zone_sync;

        # the name resolves to multiple addresses that correspond to cluster
        # nodes
        zone_sync_server cluster.example.com resolve;

        listen 127.0.0.1:4433 ssl;

        ssl_certificate      localhost.crt;
        ssl_certificate_key  localhost.key;

        zone_sync_ssl on;

        zone_sync_ssl_certificate      localhost.crt;
        zone_sync_ssl_certificate_key  localhost.key;
    }
}

```

### 3.18.3 Directives

#### zone\_sync

SYNTAX: **zone\_sync;**

DEFAULT —

CONTEXT: server

Enables the synchronization of shared memory zones between cluster nodes. Cluster nodes are defined using [zone\\_sync\\_server](#) directives.

### **zone\_sync\_buffers**

SYNTAX: **zone\_sync\_buffers** *number size*;

DEFAULT 256 4k|8k

CONTEXT: stream, server

Sets the *number* and *size* of the per-zone buffers used for pushing zone contents. By default, the buffer size is equal to one memory page. This is either 4K or 8K, depending on a platform.

### **zone\_sync\_connect\_retry\_interval**

SYNTAX: **zone\_sync\_connect\_retry\_interval** *time*;

DEFAULT 1s

CONTEXT: stream, server

Defines an interval between connection attempts to another cluster node.

### **zone\_sync\_connect\_timeout**

SYNTAX: **zone\_sync\_connect\_timeout** *time*;

DEFAULT 5s

CONTEXT: stream, server

Defines a timeout for establishing a connection with another cluster node.

### **zone\_sync\_interval**

SYNTAX: **zone\_sync\_interval** *time*;

DEFAULT 1s

CONTEXT: stream, server

Defines an interval for polling updates in a shared memory zone.

### **zone\_sync\_recv\_buffer\_size**

SYNTAX: **zone\_sync\_recv\_buffer\_size** *size*;

DEFAULT 4k|8k

CONTEXT: stream, server

Sets *size* of a per-connection receive buffer used to parse incoming stream of synchronization messages. By default, the buffer size is equal to one memory page. This is either 4K or 8K, depending on a platform.

### zone\_sync\_server

SYNTAX: **zone\_sync\_server** *address* [resolve];

DEFAULT —

CONTEXT: server

Defines the *address* of a cluster node. The address can be specified as a domain name or IP address with a mandatory port, or as a UNIX-domain socket path specified after the “unix:” prefix. A domain name that resolves to several IP addresses defines multiple nodes at once.

The *resolve* parameter instructs nginx to monitor changes of the IP addresses that correspond to a domain name of the node and automatically modify the configuration without the need of restarting nginx.

Cluster nodes are specified either dynamically as a single `zone_sync_server` directive with the *resolve* parameter, or statically as a series of several directives without the parameter.

Each cluster node should be specified only once.

All cluster nodes should use the same configuration.

In order for the *resolve* parameter to work, the [resolver](#) directive must be specified in the [stream](#) block. Example:

```
stream {
    resolver 10.0.0.1;

    server {
        zone_sync;
        zone_sync_server cluster.example.com resolve;
        ...
    }
}
```

### zone\_sync\_ssl

SYNTAX: **zone\_sync\_ssl** on | off;

DEFAULT off

CONTEXT: stream, server

Enables the SSL/TLS protocol for connections to another cluster server.

### zone\_sync\_ssl\_certificate

SYNTAX: **zone\_sync\_ssl\_certificate** *file*;

DEFAULT —

CONTEXT: stream, server

Specifies a *file* with the certificate in the PEM format used for authentication to another cluster server.

### zone\_sync\_ssl\_certificate\_key

SYNTAX: **zone\_sync\_ssl\_certificate\_key** *file*;

DEFAULT —

CONTEXT: stream, server

Specifies a *file* with the secret key in the PEM format used for authentication to another cluster server.

### zone\_sync\_ssl\_ciphers

SYNTAX: **zone\_sync\_ssl\_ciphers** *ciphers*;

DEFAULT DEFAULT

CONTEXT: stream, server

Specifies the enabled ciphers for connections to another cluster server. The ciphers are specified in the format understood by the OpenSSL library.

The full list can be viewed using the “`openssl ciphers`” command.

### zone\_sync\_ssl\_crl

SYNTAX: **zone\_sync\_ssl\_crl** *file*;

DEFAULT —

CONTEXT: stream, server

Specifies a *file* with revoked certificates (CRL) in the PEM format used to [verify](#) the certificate of another cluster server.

### zone\_sync\_ssl\_name

SYNTAX: **zone\_sync\_ssl\_name** *name*;

DEFAULT host from zone\_sync\_server

CONTEXT: stream, server

THIS DIRECTIVE APPEARED IN VERSION 1.15.7.

Allows overriding the server name used to [verify](#) the certificate of a cluster server and to be [passed through SNI](#) when establishing a connection with the cluster server.

By default, the host part of the [zone\\_sync\\_server](#) address is used, or resolved IP address if the resolve parameter is specified.

### zone\_sync\_ssl\_password\_file

SYNTAX: **zone\_sync\_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.

### zone\_sync\_ssl\_protocols

SYNTAX: **zone\_sync\_ssl\_protocols** [SSLv2] [SSLv3] [TLSv1] [TLSv1.1]  
[TLSv1.2] [TLSv1.3];

DEFAULT TLSv1 TLSv1.1 TLSv1.2

CONTEXT: stream, server

Enables the specified protocols for connections to another cluster server.

### zone\_sync\_ssl\_server\_name

SYNTAX: **zone\_sync\_ssl\_server\_name** on | off;

DEFAULT off

CONTEXT: stream, server

THIS DIRECTIVE APPEARED IN VERSION 1.15.7.

Enables or disables passing of the server name through [TLS Server Name Indication extension](#) (SNI, RFC 6066) when establishing a connection with another cluster server.

### zone\_sync\_ssl\_trusted\_certificate

SYNTAX: **zone\_sync\_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 another cluster server.

### zone\_sync\_ssl\_verify

SYNTAX: **zone\_sync\_ssl\_verify** on | off;

DEFAULT off

CONTEXT: stream, server

Enables or disables verification of another cluster server certificate.

### zone\_sync\_ssl\_verify\_depth

SYNTAX: **zone\_sync\_ssl\_verify\_depth** *number*;

DEFAULT 1

CONTEXT: stream, server

Sets the verification depth in another cluster server certificates chain.

### zone\_sync\_timeout

SYNTAX: **zone\_sync\_timeout** *timeout*;

DEFAULT 5s

CONTEXT: stream, server

Sets the *timeout* between two successive read or write operations on connection to another cluster node. If no data is transmitted within this time, the connection is closed.

### 3.18.4 API endpoints

The synchronization status of a node is available via the [/stream/zone-sync/](#) endpoint of the API which returns the [following](#) metrics.

### 3.18.5 Starting, stopping, removing a cluster node

To start a new node, update a DNS record of a cluster hostname with the IP address of the new node and start an instance. The new node will discover other nodes from DNS or static configuration and will start sending updates to them. Other nodes will eventually discover the new node using DNS and start pushing updates to it. In case of static configuration, other nodes need to be reloaded in order to send updates to the new node.

To stop a node, send the `QUIT` signal to the instance. The node will finish zone synchronization and gracefully close open connections.

To remove a node, update a DNS record of a cluster hostname and remove the IP address of the node. All other nodes will eventually discover that the node is removed, close connections to the node, and will no longer try to connect to it. After the node is removed, it can be stopped as described above. In case of static configuration, other nodes need to be reloaded in order to stop sending updates to the removed node.

# Chapter 4

## Mail server modules

### 4.1 Module ngx\_mail\_core\_module

4.1.1	Summary	424
4.1.2	Example Configuration	424
4.1.3	Directives	425
	listen	425
	mail	427
	protocol	427
	resolver	427
	resolver_timeout	428
	server	428
	server_name	428
	timeout	428

#### 4.1.1 Summary

This module is not built by default, it should be enabled with the `--with-mail` configuration parameter.

#### 4.1.2 Example Configuration

```
worker_processes 1;

error_log /var/log/nginx/error.log info;

events {
    worker_connections 1024;
}

mail {
    server_name      mail.example.com;
    auth_http        localhost:9000/cgi-bin/nginxauth.cgi;

    imap_capabilities IMAP4rev1 UIDPLUS IDLE LITERAL+ QUOTA;

    pop3_auth        plain apop cram-md5;
    pop3_capabilities LAST TOP USER PIPELINING UIDL;
```

```

smtp_auth      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;
}
}

```

### 4.1.3 Directives

#### listen

SYNTAX: **listen** *address:port* [*ssl*] [*backlog=number*] [*rcvbuf=size*] [*sndbuf=size*] [*bind*] [*ipv6only=on|off*] [*so\_keepalive=on|off*][*keepidle*]:[*keepintvl*]:[*keepcnt*];

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. 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 (0.7.58) are specified in square brackets:

```

listen [::1]:110;
listen [::]:110;

```

UNIX-domain sockets (1.3.5) are specified with the “unix:” prefix:

```

listen unix:/var/run/nginx.sock;

```

Different servers must listen on different *address:port* pairs.

The *ssl* parameter allows specifying that all connections accepted on this port should work in SSL mode.

The *listen* directive can have several additional parameters specific to socket-related system calls.

`backlog=number`

sets the `backlog` parameter in the `listen` call that limits the maximum length for the queue of pending connections (1.9.2). 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 (1.11.13).

`sndbuf=size`

sets the send buffer size (the `SO_SNDBUF` option) for the listening socket (1.11.13).

`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`), nginx 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 `ipv6only` 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|off[[keepidle]:[keepintvl]:[keepcnt]`

this parameter configures the “TCP keepalive” behavior for the listening socket. If this parameter is omitted then the operating system’s settings will be in effect for the socket. If it is set to the value “on”, the `SO_KEEPALIVE` option is turned on for the socket. If it is set to the value “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.

## mail

SYNTAX: **mail** { ... }

DEFAULT —

CONTEXT: main

Provides the configuration file context in which the mail server directives are specified.

## 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

Unnecessary protocols can be disabled using the [configuration](#) parameters `--without-mail_imap_module`, `--without-mail_pop3_module`, and `--without-mail_smtp_module`.

## resolver

SYNTAX: **resolver** address ... [valid=*time*];

SYNTAX: **resolver** off;

DEFAULT 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 127.0.0.1 [::1]:5353;
```

An address can be specified as a domain name or IP address, and an optional port (1.3.1, 1.2.2). If port is not specified, the port 53 is used. Name servers are queried in a round-robin fashion.

Before version 1.1.7, only a single name server could be configured. Specifying name servers using IPv6 addresses is supported starting from versions 1.3.1 and 1.2.2.

By default, nginx caches answers using the TTL value of a response. An optional `valid` parameter allows overriding it:

```
resolver 127.0.0.1 [::1]:5353 valid=30s;
```

Before version 1.1.9, tuning of caching time was not possible, and nginx always cached answers for the duration of 5 minutes.

The special value `off` disables resolving.

### resolver\_timeout

SYNTAX: **resolver\_timeout** *time*;

DEFAULT 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 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 60s

CONTEXT: mail, server

Sets the timeout that is used before proxying to the backend starts.

## 4.2 Module ngx\_mail\_auth\_http\_module

4.2.1 Directives . . . . .	429
auth_http . . . . .	429
auth_http_header . . . . .	429
auth_http_pass_client_cert . . . . .	429
auth_http_timeout . . . . .	429
4.2.2 Protocol . . . . .	430

### 4.2.1 Directives

#### auth\_http

SYNTAX: **auth\_http** *URL*;

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 nginx. For example:

```
auth_http_header X-Auth-Key "secret_string";
```

#### auth\_http\_pass\_client\_cert

SYNTAX: **auth\_http\_pass\_client\_cert** on | off;

DEFAULT off

CONTEXT: mail, server

THIS DIRECTIVE APPEARED IN VERSION 1.7.11.

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 60s

CONTEXT: mail, server

Sets the timeout for communication with the authentication server.

### 4.2.2 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:

```
GET /auth HTTP/1.0
Host: localhost
Auth-Method: none
Auth-User:
Auth-Pass:
Auth-Protocol: smtp
Auth-Login-Attempt: 1
Client-IP: 192.0.2.42
Client-Host: client.example.org
Auth-SMTP-Helo: client.example.org
Auth-SMTP-From: MAIL FROM: <>
Auth-SMTP-To: RCPT TO: <postmaster@mail.example.com>
```

For the SSL/TLS client connection (1.7.11), 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.

Prior to version 1.11.7, the “FAILED” result did not contain the *reason* string.

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 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-Verify: SUCCESS
Auth-SSL-Subject: /CN=example.com
Auth-SSL-Issuer: /CN=example.com
Auth-SSL-Serial: C07AD56B846B5BFF
Auth-SSL-Fingerprint: 29d6a80a123d13355ed16b4b04605e29cb55a5ad
```

## 4.3 Module ngx\_mail\_proxy\_module

4.3.1 Directives . . . . .	433
<a href="#">proxy_buffer</a> . . . . .	433
<a href="#">proxy_pass_error_message</a> . . . . .	433
<a href="#">proxy_timeout</a> . . . . .	433
<a href="#">xclient</a> . . . . .	434

### 4.3.1 Directives

#### proxy\_buffer

SYNTAX: **proxy\_buffer** *size*;  
 DEFAULT 4k|8k  
 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 off  
 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 nginx 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. For example, CommuniGatePro informs a user about [mailbox overflow](#) or other events by periodically outputting the [authentication error](#). The directive should be enabled in this case.

#### proxy\_timeout

SYNTAX: **proxy\_timeout** *timeout*;  
 DEFAULT 24h  
 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 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 nginx 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, the [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 then nginx 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.

## 4.4 Module ngx\_mail\_ssl\_module

4.4.1	Summary	435
4.4.2	Example Configuration	435
4.4.3	Directives	436
	ssl	436
	ssl_certificate	436
	ssl_certificate_key	437
	ssl_ciphers	437
	ssl_client_certificate	437
	ssl_crl	438
	ssl_dhparam	438
	ssl_ecdh_curve	438
	ssl_password_file	438
	ssl_prefer_server_ciphers	439
	ssl_protocols	439
	ssl_session_cache	439
	ssl_session_ticket_key	440
	ssl_session_tickets	441
	ssl_session_timeout	441
	ssl_trusted_certificate	441
	ssl_verify_client	441
	ssl_verify_depth	441
	starttls	442

### 4.4.1 Summary

The ngx\_mail\_ssl\_module module provides the necessary support for a mail proxy server to work with the SSL/TLS protocol.

This module is not built by default, it should be enabled with the `--with-mail_ssl_module` configuration parameter.

This module requires the [OpenSSL](#) library.

### 4.4.2 Example Configuration

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):

```
worker_processes auto;

mail {
    ...

    server {
        listen          993 ssl;

        ssl_protocols   TLSv1 TLSv1.1 TLSv1.2;
        ssl_ciphers      AES128-SHA:AES256-SHA:RC4-SHA:DES-CBC3-SHA:RC4-MD5;
        ssl_certificate   /usr/local/nginx/conf/cert.pem;
        ssl_certificate_key /usr/local/nginx/conf/cert.key;
        ssl_session_cache shared:SSL:10m;
        ssl_session_timeout 10m;

        ...
    }
}
```

### 4.4.3 Directives

#### ssl

SYNTAX: **ssl** on | off;

DEFAULT off

CONTEXT: mail, server

This directive was made obsolete in version 1.15.0. The `ssl` parameter of the [listen](#) directive should be used instead.

#### ssl\_certificate

SYNTAX: **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.

Since version 1.11.0, this directive can be specified multiple times to load certificates of different types, for example, RSA and ECDSA:

```
server {
    listen          993 ssl;

    ssl_certificate   example.com.rsa.crt;
    ssl_certificate_key example.com.rsa.key;

    ssl_certificate   example.com.ecdsa.crt;
    ssl_certificate_key example.com.ecdsa.key;

    ...
}
```

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* (1.15.10), 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: **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* (1.7.9), 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* (1.15.10), 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 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 full list can be viewed using the “`openssl ciphers`” command.

The previous versions of nginx used [different](#) ciphers by default.

### ssl\_client\_certificate

SYNTAX: **ssl\_client\_certificate** *file*;

DEFAULT —

CONTEXT: mail, server

THIS DIRECTIVE APPEARED IN VERSION 1.7.11.

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\_crl

SYNTAX: **ssl\_crl** *file*;

DEFAULT —

CONTEXT: mail, server

THIS DIRECTIVE APPEARED IN VERSION 1.7.11.

Specifies a *file* with revoked certificates (CRL) in the PEM format used to [verify](#) client certificates.

## ssl\_dhparam

SYNTAX: **ssl\_dhparam** *file*;

DEFAULT —

CONTEXT: mail, server

THIS DIRECTIVE APPEARED IN VERSION 0.7.2.

Specifies a *file* with DH parameters for DHE ciphers.

By default no parameters are set, and therefore DHE ciphers will not be used.

Prior to version 1.11.0, builtin parameters were used by default.

## ssl\_ecdh\_curve

SYNTAX: **ssl\_ecdh\_curve** *curve*;

DEFAULT `auto`

CONTEXT: mail, server

THIS DIRECTIVE APPEARED IN VERSIONS 1.1.0 AND 1.0.6.

Specifies a *curve* for ECDHE ciphers.

When using OpenSSL 1.0.2 or higher, it is possible to specify multiple curves (1.11.0), for example:

```
ssl_ecdh_curve prime256v1:secp384r1;
```

The special value `auto` (1.11.0) instructs nginx to use a list built into the OpenSSL library when using OpenSSL 1.0.2 or higher, or `prime256v1` with older versions.

Prior to version 1.11.0, the `prime256v1` curve was used by default.

## ssl\_password\_file

SYNTAX: **ssl\_password\_file** *file*;

DEFAULT —

CONTEXT: mail, server

THIS DIRECTIVE APPEARED IN VERSION 1.7.3.

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 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 TLSv1 TLSv1.1 TLSv1.2

CONTEXT: mail, server

Enables the specified protocols.

The TLSv1.1 and TLSv1.2 parameters (1.1.13, 1.0.12) work only when OpenSSL 1.0.1 or higher is used.

The TLSv1.3 parameter (1.13.0) works only when OpenSSL 1.1.1 built with TLSv1.3 support is used.

### ssl\_session\_cache

SYNTAX: **ssl\_session\_cache** off | none | [builtin[:size]]  
[shared:name:size];

DEFAULT 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: nginx explicitly tells a client that sessions may not be reused.

`none`

the use of a session cache is gently disallowed: nginx 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.

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

THIS DIRECTIVE APPEARED IN VERSION 1.5.7.

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, 1.11.8) or AES128 (for 48-byte keys) is used for encryption.

### ssl\_session\_tickets

SYNTAX: **ssl\_session\_tickets** on | off;

DEFAULT on

CONTEXT: mail, server

THIS DIRECTIVE APPEARED IN VERSION 1.5.9.

Enables or disables session resumption through [TLS session tickets](#).

### ssl\_session\_timeout

SYNTAX: **ssl\_session\_timeout** *time*;

DEFAULT 5m

CONTEXT: mail, server

Specifies a time during which a client may reuse the session parameters.

### ssl\_trusted\_certificate

SYNTAX: **ssl\_trusted\_certificate** *file*;

DEFAULT —

CONTEXT: mail, server

THIS DIRECTIVE APPEARED IN VERSION 1.7.11.

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 off

CONTEXT: mail, server

THIS DIRECTIVE APPEARED IN VERSION 1.7.11.

Enables verification of client certificates. The verification result is passed in the Auth-SSL-Verify header of the [authentication](#) request.

The *optional* parameter requests the client certificate and verifies it if the certificate is present.

The *optional\_no\_ca* parameter 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 nginx performs the actual certificate verification. The contents of the certificate is accessible through requests [sent](#) to the authentication server.

### ssl\_verify\_depth

SYNTAX: **ssl\_verify\_depth** *number*;

DEFAULT 1

CONTEXT: mail, server

THIS DIRECTIVE APPEARED IN VERSION 1.7.11.

Sets the verification depth in the client certificates chain.

### **starttls**

SYNTAX: **starttls** on | off | only;

DEFAULT 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.

## 4.5 Module ngx\_mail\_imap\_module

4.5.1 Directives . . . . .	443
imap_auth . . . . .	443
imap_capabilities . . . . .	443
imap_client_buffer . . . . .	443

### 4.5.1 Directives

#### imap\_auth

SYNTAX: **imap\_auth** *method* ...;

DEFAULT plain

CONTEXT: mail, server

Sets permitted methods of authentication for IMAP clients. Supported methods are:

login

[AUTH=LOGIN](#)

plain

[AUTH=PLAIN](#)

cram-md5

[AUTH=CRAM-MD5](#). In order for this method to work, the password must be stored unencrypted.

external

[AUTH=EXTERNAL](#) (1.11.6).

#### imap\_capabilities

SYNTAX: **imap\_capabilities** *extension* ...;

DEFAULT IMAP4 IMAP4rev1 UIDPLUS

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 nginx transparently proxies a client connection to the backend).

The current list of standardized extensions is published at [www.iana.org](http://www.iana.org).

#### imap\_client\_buffer

SYNTAX: **imap\_client\_buffer** *size*;

DEFAULT 4k | 8k

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.

## 4.6 Module ngx\_mail\_pop3\_module

4.6.1 Directives . . . . .	445
<a href="#">pop3_auth</a> . . . . .	445
<a href="#">pop3_capabilities</a> . . . . .	445

### 4.6.1 Directives

#### pop3\_auth

SYNTAX: **pop3\_auth** *method* ...;

DEFAULT plain

CONTEXT: mail, server

Sets permitted methods of authentication for POP3 clients. Supported methods are:

plain

[USER/PASS](#), [AUTH PLAIN](#), [AUTH LOGIN](#). It is not possible to disable these methods.

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

[AUTH EXTERNAL](#) (1.11.6).

#### pop3\_capabilities

SYNTAX: **pop3\_capabilities** *extension* ...;

DEFAULT 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 nginx transparently proxies the client connection to the backend).

The current list of standardized extensions is published at [www.iana.org](http://www.iana.org).

## 4.7 Module ngx\_mail\_smtp\_module

4.7.1 Directives . . . . .	446
smtp_auth . . . . .	446
smtp_capabilities . . . . .	446
smtp_client_buffer . . . . .	447
smtp_greeting_delay . . . . .	447

### 4.7.1 Directives

#### smtp\_auth

SYNTAX: **smtp\_auth** *method* ...;

DEFAULT login plain

CONTEXT: mail, server

Sets permitted methods of [SASL authentication](#) for SMTP clients. Supported methods are:

login

[AUTH LOGIN](#)

plain

[AUTH PLAIN](#)

cram-md5

[AUTH CRAM-MD5](#). In order for this method to work, the password must be stored unencrypted.

external

[AUTH EXTERNAL](#) (1.11.6).

none

Authentication is not required.

#### 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 nginx transparently proxies the client connection to the backend).

The current list of standardized extensions is published at [www.iana.org](http://www.iana.org).

**smtp\_client\_buffer**

SYNTAX: **smtp\_client\_buffer** *size*;

DEFAULT 4k | 8k

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: **smtp\_greeting\_delay** *time*;

DEFAULT 0

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.

# Chapter 5

## Miscellaneous

### 5.1 High Availability support for NGINX Plus

5.1.1	High Availability support	448
5.1.2	Configuring HA setup	449
5.1.3	Check scripts	450
5.1.4	Checking the status of HA setup	451
5.1.5	Forcing state change	451
5.1.6	Adding more virtual IP addresses	451
5.1.7	Troubleshooting keepalived and VRRP	452
5.1.8	Miscellaneous	453

#### 5.1.1 High Availability support

NGINX-HA-Keepalived is a solution for fast and easy configuration of NGINX Plus in an active-passive high-availability (HA) setup. It is based on keepalived.

The keepalived project provides a keepalive facility for Linux servers, an implementation of the VRRP protocol to manage virtual routers (virtual IP addresses), and a health check facility to determine if a service (web server, PHP back end, database server, etc.) is up and operational. If a service on a node fails a configurable number of health checks, keepalived reassigns the virtual IP address of the node to a secondary node.

The VRRP protocol ensures that one of participating nodes is master. The backup node listens for VRRP advertisement packets from the master node. If it does not receive an advertisement packet for a period longer than three times the configured advertisement interval, the backup node takes over as master and assigns the configured virtual IP addresses to itself.

### 5.1.2 Configuring HA setup

Run the `nginx-ha-setup` script (available in the `nginx-ha-keepalived` package, must be installed separately) on both nodes as the root user.

The script configures a high-availability NGINX Plus environment with an active-passive pair of nodes acting as master and backup. It prompts for the following data:

- IP address of the local and remote nodes (one of which will be configured as a master, the other one as a backup).
- One free IP address to be used as the cluster endpoint's (floating) virtual IP address.

The configuration of the keepalived daemon is recorded in a text file, `/etc-keepalived/keepalived.conf`. The configuration blocks in the file control notification settings, the virtual IP addresses to manage, and the health checks to use to test the services that rely on virtual IP addresses. Following is the configuration created by the `nginx-ha-setup` script on a CentOS 7 machine:

```
vrrp_script chk_nginx_service {
    script "/usr/libexec/keepalived/nginx-ha-check"
    interval 3
    weight 50
}

vrrp_instance VI_1 {
    interface eth0
    state BACKUP
    priority 101
    virtual_router_id 51
    advert_int 1
    unicast_src_ip 192.168.100.100
    unicast_peer {
        192.168.100.101
    }
    authentication {
        auth_type PASS
        auth_pass f8f0e5114cbe031a3e1e622daf18f82a
    }
    virtual_ipaddress {
        192.168.100.150
    }
    track_script {
        chk_nginx_service
    }
    notify "/usr/libexec/keepalived/nginx-ha-notify"
}
```

The configuration shown above is self-explanatory, but a few items are worth noting:

- Each node in the HA setup needs its own copy of the configuration file, with values for the `priority`, `unicast_src_ip`, and `unicast_peer` directives that are appropriate to the node's status (master or backup).

- The `priority` directive controls which host becomes the master, as explained in the next section.
- The `notify` directive names the notification script included in the distribution, which can be used to generate syslog messages (or other notifications) when a state transition or fault occurs.
- The value 51 for the `virtual_router_id` directive in the `vrrp_instance VI_1` block is a sample value.
- If you have multiple pairs of keepalived instances (or other VRRP instances) running in your local network, create a `vrrp_instance` block for each one, with a unique name (like `VI_1` in the sample) and `virtual_router_id` number.

### 5.1.3 Check scripts

There is no fencing mechanism in keepalived. If the two nodes in a pair are not aware of each other, each assumes it is the master and assigns the virtual IP address to itself. To prevent this situation, the `chk_nginx_service` script is executed regularly to check its exit code and adjust the node's priority as necessary. Code 0 indicates correct operation, and code 1 (or any nonzero code) indicates an error.

In the default configuration of the `chk_nginx_service` script, the `weight` directive is set to 50, which means that when the check script succeeds:

- The priority of the first node (which has a base priority of 101) is set to 151.
- The priority of the second node (which has a base priority of 100) is set to 150.

The first node has higher priority (151 in this case) and becomes master.

Use the `interval` directive to specify how often the check script executes, in seconds (it is set to 3 in the default configuration). Note that the check also fails when the timeout is reached (by default, the timeout is the same as the check interval).

Use the `rise` and `fall` directives to specify how many times the script must succeed or fail before action is taken (they are not set in the default configuration).

The default script provided with the `nginx-ha-keepalived` package checks if nginx is up. We recommend creating additional scripts as appropriate for your local setup.

### 5.1.4 Checking the status of HA setup

To see which node is currently the master for a given virtual IP address, run the `ip addr show` command for the interface on which the `vrp_instance` is defined (in the following commands, interface `eth0`):

```
centos7-1 # ip addr show eth0
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP
    qlen 1000
    link/ether 52:54:00:33:a5:a5 brd ff:ff:ff:ff:ff:ff
    inet 192.168.100.100/24 brd 192.168.122.255 scope global dynamic eth0
        valid_lft 3071sec preferred_lft 3071sec
    inet 192.168.100.150/32 scope global eth0
        valid_lft forever preferred_lft forever

centos7-2 # ip addr show eth0
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP
    qlen 1000
    link/ether 52:54:00:33:a5:87 brd ff:ff:ff:ff:ff:ff
    inet 192.168.100.101/24 brd 192.168.122.255 scope global eth0
        valid_lft forever preferred_lft forever
```

In this output, the defined virtual IP address (192.168.100.150) is currently assigned to the host with real IP address of 192.168.100.100.

When a host's HA state changes, `nginx-ha-keepalived` writes it to the `/var/run/nginx-ha-keepalived.state` file:

```
centos7-1 # cat /var/run/nginx-ha-keepalived.state
STATE=MASTER

centos7-2 # cat /var/run/nginx-ha-keepalived.state
STATE=BACKUP
```

### 5.1.5 Forcing state change

To force the master node to switch to backup state, run the following command on it:

```
# service keepalived stop
```

As it shuts down, `keepalived` sends a VRRP packet with priority 0 to the backup node, which causes the backup node to take over the virtual IP address.

### 5.1.6 Adding more virtual IP addresses

The configuration created by `nginx-ha-setup` is very basic, and makes a single IP address highly available. To make more than one IP address highly available, add each new IP address to the `virtual_ipaddress` block in the `/etc/keepalived/keepalived.conf` configuration file. Then run the `service keepalived reload` command on both nodes to reload the `keepalived` service:

```
virtual_ipaddress {
    192.168.100.150
```

```

    192.168.100.200
    1234:5678:9abc:def::1/64
}

```

As indicated in this example, keepalived can be utilized in dual-stack IPv4/IPv6 environments to fail over both IPv4 and IPv6 addresses.

The syntax in the `virtualipaddress` block replicates the syntax of the `ip` utility.

### 5.1.7 Troubleshooting keepalived and VRRP

The keepalived daemon logs to syslog. On CentOS, RHEL, and SLES-based systems, the output is typically written to `/var/log/messages`, whereas on Ubuntu and Debian-based systems it is written to `/var/log/syslog`. Log entries record events such as startup of the keepalived daemon and state transitions. Here are a few sample entries that show the keepalived daemon starting up, and the node transitioning a VRRP instance to the master state:

```

Feb 27 14:42:04 centos7-1 systemd: Starting LVS and VRRP High Availability
Monitor...
Feb 27 14:42:04 centos7-1 Keepalived[19242]: Starting Keepalived v1.2.15
(02/26,2015)
Feb 27 14:42:04 centos7-1 Keepalived[19243]: Starting VRRP child process, pid
=19244
Feb 27 14:42:04 centos7-1 Keepalived_vrrp[19244]: Registering Kernel netlink
reflector
Feb 27 14:42:04 centos7-1 Keepalived_vrrp[19244]: Registering Kernel netlink
command channel
Feb 27 14:42:04 centos7-1 Keepalived_vrrp[19244]: Registering gratuitous ARP
shared channel
Feb 27 14:42:05 centos7-1 systemd: Started LVS and VRRP High Availability
Monitor.
Feb 27 14:42:05 centos7-1 Keepalived_vrrp[19244]: Opening file '/etc/keepalived
/keepalived.conf'.
Feb 27 14:42:05 centos7-1 Keepalived_vrrp[19244]: Truncating auth_pass to 8
characters
Feb 27 14:42:05 centos7-1 Keepalived_vrrp[19244]: Configuration is using :
64631 Bytes
Feb 27 14:42:05 centos7-1 Keepalived_vrrp[19244]: Using LinkWatch kernel
netlink reflector...
Feb 27 14:42:05 centos7-1 Keepalived_vrrp[19244]: VRRP_Instance(VI_1) Entering
BACKUP STATE
Feb 27 14:42:05 centos7-1 Keepalived_vrrp[19244]: VRRP sockpool: [ifindex(2),
proto(112), unicast(1), fd(14,15)]
Feb 27 14:42:05 centos7-1 nginx-ha-keepalived: Transition to state 'BACKUP' on
VRRP instance 'VI_1'.
Feb 27 14:42:05 centos7-1 Keepalived_vrrp[19244]: VRRP_Script(chk_nginx_service
) succeeded
Feb 27 14:42:06 centos7-1 Keepalived_vrrp[19244]: VRRP_Instance(VI_1) forcing a
new MASTER election
Feb 27 14:42:06 centos7-1 Keepalived_vrrp[19244]: VRRP_Instance(VI_1) forcing a
new MASTER election
Feb 27 14:42:07 centos7-1 Keepalived_vrrp[19244]: VRRP_Instance(VI_1)
Transition to MASTER STATE
Feb 27 14:42:08 centos7-1 Keepalived_vrrp[19244]: VRRP_Instance(VI_1) Entering
MASTER STATE
Feb 27 14:42:08 centos7-1 Keepalived_vrrp[19244]: VRRP_Instance(VI_1) setting
protocol VIPs.
Feb 27 14:42:08 centos7-1 Keepalived_vrrp[19244]: VRRP_Instance(VI_1) Sending
gratuitous ARPs on eth0 for 192.168.100.150
Feb 27 14:42:08 centos7-1 nginx-ha-keepalived: Transition to state 'MASTER' on
VRRP instance 'VI_1'.

```

```
Feb 27 14:42:13 centos7-1 Keepalived_vrrp[19244]: VRRP_Instance(VI_1) Sending gratuitous ARPs on eth0 for 192.168.100.150
```

If the system log does not explain the source of a problem, run the `tcpdump` command with the following parameters to display the VRRP advertisements that are sent on the local network:

```
# tcpdump -vvv -ni eth0 proto vrrp
```

If you have multiple VRRP instances on the local network and want to filter the traffic for select hosts, include the `host` parameter to specify the IP address that is defined in the `unicast_peer` block, as in the following example:

```
centos7-1 # tcpdump -vvv -ni eth0 proto vrrp and host 192.168.100.101
tcpdump: listening on eth0, link-type EN10MB (Ethernet), capture size 65535 bytes
14:48:27.188100 IP (tos 0xc0, ttl 255, id 382, offset 0, flags [none], proto VRRP (112), length 40)
  192.168.100.100 > 192.168.100.101: vrrp 192.168.100.100 > 192.168.100.101: VRRPv2, Advertisement, vrid 51, prio 151, authtype simple, intvl 1s, length 20, addrs: 192.168.100.150 auth "f8f0e511"
```

Several fields in the output are useful for debugging:

- `authtype` - the type of authentication in use (`authentication directive`)
- `vrid` - the virtual router ID (`virtual_router_id directive`)
- `prio` - the node's priority (`priority directive`)
- `intvl` - the frequency at which advertisements are sent (`advert_int directive`)
- `auth` - the authentication token sent (`auth_pass directive`)

### 5.1.8 Miscellaneous

Note that NGINX configuration files on both nodes must define the services that are being made highly available. Keeping the configuration files in sync is outside the scope of the provided clustering software.

The `nginx-ha-keepalived` package comes with numerous configuration examples, in the `/usr/share/doc/nginx-ha-keepalived/` directory. They show how to configure numerous aspects of an HA setup.

## 5.2 Command-line parameters

### 5.2.1 Overview . . . . . 454

#### 5.2.1 Overview

nginx supports the following command-line parameters:

- `-? | -h` — print help for command-line parameters.
- `-c file` — use an alternative configuration *file* instead of a default file.
- `-g directives` — set [global configuration directives](#), for example,

```
nginx -g "pid /var/run/nginx.pid; worker_processes `sysctl -n hw.ncpu`;"
```

- `-p prefix` — set nginx path prefix, i.e. a directory that will keep server files (default value is `/usr/local/nginx`).
- `-q` — suppress non-error messages during configuration testing.
- `-s signal` — send a *signal* to the master process. The argument *signal* can be one of:
  - `stop` — shut down quickly
  - `quit` — shut down gracefully
  - `reload` — reload configuration, start the new worker process with a new configuration, gracefully shut down old worker processes.
  - `reopen` — reopen log files
- `-t` — test the configuration file: nginx checks the configuration for correct syntax, and then tries to open files referred in the configuration.
- `-T` — same as `-t`, but additionally dump configuration files to standard output (1.9.2).
- `-v` — print nginx version.
- `-V` — print nginx version, compiler version, and configure parameters.

# Appendix A

## Changelog for NGINX Plus

This appendix contains the most important changes that may apply to both NGINX Plus and nginx/OSS. Full changelog for nginx/OSS is available in the packages and by the following link: <http://nginx.org/en/CHANGES>

- NGINX Plus R18 (1.15.10), released Apr 9, 2019
  - Added support for dynamic SSL certificate loading, either from [file](#) or from [key-value](#) storage (variable should be prefixed with `data:` for the latter case).
  - Active health checks extended with additional logic of verifying arbitrary variables (the `require` parameter of the [match](#) directive).
  - Clustering enhancement: a single [zone sync](#) configuration can now be used for all instances in a cluster with the help of wildcard support in the [listen](#) directive.
  - Port ranges can now be used in the [listen](#) directive.
  - Stream proxy: added an ability to explicitly close existing connections to the particular upstream server after it was removed from the group due to health check failure, API call, or re-resolve action (the [proxy\\_session\\_drop](#) directive).
  - New variable, [\\$upstream\\_bytes\\_sent](#), contains number of bytes sent to an upstream server.
- NGINX Plus R17 (1.15.7), released Dec 11, 2018
  - Added [support](#) for TLS v1.3 [early data](#). Check out the [\\$ssl\\_early\\_data](#) variable which can be used to protect against [replay attacks](#) at the application layer.
  - Introduced two-stage rate limiting (the `delay` parameter of the [limit\\_req](#) directive).
  - Added [support](#) for retrieving JSON Web Key (JWK) set from a subrequest.
  - Added support for the Ed25519 and Ed448 cryptographic [algorithms](#) to the [JSON Web Token authorization module](#).
  - Added an option to enable TCP keepalives for outgoing connections to proxied servers (the [proxy\\_socket\\_keepalive](#) and friends).
  - Fine-grained control over persistent connections to upstreams with the [keepalive\\_timeout](#) and [keepalive\\_requests](#) directives.
  - Added ability to restrict UDP session to a particular number of packets with the [proxy\\_requests](#) directive.
- NGINX Plus R16 (1.15.2), released Sep 5, 2018
  - IMPORTANT: [status](#) and [upstream\\_conf](#) modules deprecated since R13 were finally removed in favor to the [API](#) module.

- Shared zones synchronization extended to support [keyval](#) and [limit\\_req](#) modules.
  - Key-value pairs can now be expired after configured timeout (the `timeout` parameter of the [keyval\\_zone](#) directive).
  - Introduced new load balancing algorithm: [random](#) with two choices (optional).
  - UDP load balancing extended to support multiple incoming datagrams from a client within a single session, thus allowing to proxy/load balance of more complex applications.
  - Added support for PROXY protocol version 2 in [HTTP](#) and [stream](#) modules.
  - New variable, [\\$ssl\\_preread\\_protocol](#), of stream module contains the highest SSL protocol version supported by the client.
- NGINX Plus R15 (1.13.10), released Apr 10, 2018
    - Added [gRPC proxy](#) support.
    - Added [HTTP/2 push](#) support.
    - Introduced [shared zone synchronization](#) between NGINX Plus nodes. Currently, it is possible to synchronize [sticky](#) sessions when using [learn](#) method.
    - Added an ability to completely disable any escaping in access log (the `escape=none` parameter of the [log\\_format](#) directive).
    - It is no longer required to run nginx under superuser when using the [proxy\\_bind](#) directive with `transparent` parameter on Linux, as worker processes now inherit the `CAP_NET_RAW` capability from the master process.
    - Added the [auth\\_jwt\\_leeway](#) directive used to configure a leeway to account for clock skew when verifying exp and nbf claims, as per [RFC 7519](#).
    - Stream [SSL pre-read](#) module now can extract list of protocols advertised by the client through ALPN (the [\\$ssl\\_preread\\_alpn\\_protocols](#) variable).
    - New variable, [\\$upstream\\_queue\\_time](#), contains time the request spent in the upstream [queue](#).
  - NGINX Plus R14 (1.13.7), released Dec 12, 2017
    - Added refactored status dashboard v2 (`dashboard.html` in packages) that uses recently introduced [API](#) subsystem instead of older [status](#) and [upstream\\_conf](#) interfaces. Previous dashboard v1 is still available (`status.html` in packages). Please note that older status and upstream\_conf interfaces, as well as dashboard v1, are going to be removed in NGINX Plus R16.
    - Dynamic key-value pairs support added to the [stream](#) proxy. [API](#) version has changed to 2 (see the ["Compatibility"](#) section for details).
    - The [auth\\_jwt\\_header\\_set](#) and [auth\\_jwt\\_claim\\_set](#) directives can now handle multiple values, providing complex JWT claims support.
    - Additional cryptographic [algorithms](#) were introduced in the [JSON Web Token authorization module](#).
    - Upstream servers configured with the [resolve](#) parameter are now being pre-resolved on configuration reload.
    - The [drain](#) parameter of the [server](#) directive can now be specified in configuration.
    - New variable: [\\$ssl\\_client\\_escaped\\_cert](#).
    - Swagger UI was adjusted to support custom ports in URLs.

- NGINX Plus R13 (1.13.4), released Aug 29, 2017
  - Introduced new [API](#) for accessing various status information, configuring upstream server groups on-the-fly, and managing key-value pairs. Swagger specification and UI is bundled in the `nginx-plus` packages for easy try-out.
  - Introduced new module that creates variables based on [key-value pairs](#), dynamically managed by the new [API](#).
  - Introduced new module for [mirroring](#) requests by creating background mirror subrequests and ignoring their responses.
  - Added support for HTTP [trailers](#).
  - New [worker\\_shutdown\\_timeout](#) directive allows configuring a timeout for graceful shutdown of worker processes.
  - [Sticky](#) cookie with `expires` parameter now also includes the `max-age` attribute defined in [RFC 6265](#).
  - [Sticky learn](#) session affinity mechanism extended with the `header` parameter used to indicate that a session should be created immediately after receiving response headers from upstream server.
  - The [proxy\\_next\\_upstream](#) directive extended with the new `http_429` parameter ("Too Many Requests").
  - Added the ability to specify network buffer sizes in [stream](#) and [mail](#) modules.
  - SSL renegotiation is now allowed on backend connections.
  - Added initial support for TLS 1.3.
- NGINX Plus R12 (1.11.10), released Mar 14, 2017
  - Status module [dataset](#) updated with `nginx_build` name (`nginx_build`), shared zones usage statistics (the `slabs/ subtree`), and additional upstream fields (`name`, `service`).
  - Status dashboard now shows NGINX Plus version, response time metrics, shared zones memory usage, and server names in upstreams.
  - Added [support](#) for the `stale-while-revalidate` and `stale-if-error` Cache-Control extensions, as defined by [RFC 5861](#).
  - Introduced flexible control of caching byte-range responses (the [proxy\\_cache\\_max\\_range\\_offset](#) directive).
  - Cache header `Vary` and `ETag` lengths increased to 128 bytes. Note that the on-disk cache format has changed, so cached content will be invalidated after the upgrade.
  - Introduced the mandatory parameter in the `health_check` directive, requiring all new servers to pass the associated health check before accepting real traffic.
  - UDP health checks now may be configured without specifying `match` block.
  - Stream module now supports client SSL certificates [verification](#).
  - Added a number of [SSL variables](#) representing various details about client certificates and capabilities (`$ssl_client_v_end`, `$ssl_client_v_start`, `$ssl_client_v_remain`, `$ssl_curves`, `$ssl_ciphers`). The `$ssl_client_verify` variable was extended to include a reason of failure.
  - The `$ssl_client_i_dn` and `$ssl_client_s_dn` variables are now compliant with [RFC 2253](#); legacy variants are available as `$ssl_client_i_dn_legacy` and `$ssl_client_s_dn_legacy`, accordingly.
  - Support for accessing arbitrary JWT fields as [variables](#).
  - Added support for JSON escaping in access logs (the `escape` parameter of the [log\\_format](#) directive).
  - WebP support added to the [image filter](#) module.
  - Duplicate configuration parts excluded from `nginx -T` output.
  - Various improvements in memory usage and performance, including upstream [queue](#) optimization.

- NGINX Plus R11 (1.11.5), released Oct 25, 2016
  - Introduced dynamic modules binary compatibility between NGINX Plus and corresponding version of nginx/OSS.
  - Stream module enhancements (custom [logging](#) with a number of additional [variables](#), [PROXY protocol](#) support for incoming connections, support for [obtaining](#) real IP address and port from PROXY protocol header, ability to [extract server name](#) from SNI to a variable for various purposes, e.g. custom routing).
  - Status module [dataset](#) updated with additional stream metrics ([sessions](#), [discarded](#)).
  - Cache manager improved to support iterative operations mode when deleting old cache files, reducing the disk load (see the [manager\\_files](#), [manager\\_threshold](#), and [manager\\_sleep](#) parameters of the [proxy\\_cache\\_path](#) directive).
  - Added support for using variables in the domain parameter of the [sticky](#) directive.
  - New variable: [\\$upstream\\_bytes\\_received](#).
- NGINX Plus R10 (1.11.3), released Aug 23, 2016
  - New dynamic module: [ModSecurity](#) (package name is `nginx-plus-module-modsecurity`). This is the early release candidate of ModSecurity 3.0.
  - New dynamic module: [nginScript](#) (package name is `nginx-plus-module-njs`).
  - Support for client authorization using the [JSON Web Token](#) (JWT).
  - Stream module enhancements (embedded [variables](#), [resolver](#) support, [map](#) module, [geo](#) and [geoip](#) modules, [A/B testing](#) support).
  - Support for multiple [SSL certificate](#) types per SSL server or SNI name (e.g., RSA and ECDSA).
  - Transparent proxy mode support (the [transparent](#) parameter of the [proxy\\_bind](#) directive).
  - Support for the `IP_BIND_ADDRESS_NO_PORT` socket option where available, allowing for many more upstream connections.
  - HTTP/2 improvements: unbuffered upload support, general bugfixes.
  - New variables: [\\$request\\_id](#), [\\$proxy\\_protocol\\_port](#), [\\$realip\\_remote\\_port](#).
  - Lua module updated to version 0.10.6 (`nginx-plus-extras`, `nginx-plus-module-lua`).
  - Passenger module updated to version 5.0.30 (`nginx-plus-extras`, `nginx-plus-module-passenger`).
  - `headers-more` module updated to version 0.31 (`nginx-plus-extras`, `nginx-plus-module-headers-more`).
  - `set-misc` module updated to version 0.31 (`nginx-plus-extras`, `nginx-plus-module-headers-more`).

NGINX Plus R10 will be the last release to provide the NGINX Plus Extras package. Users should migrate to the NGINX Plus package and use the equivalent dynamic modules.

- NGINX Plus R9 (1.9.13), released Apr 12, 2016
  - Introduced a number of standalone packages with dynamic modules for NGINX Plus (both official and third-party). Packages with official modules:

- \* `nginx-plus-module-geoip` ([doc](#))
- \* `nginx-plus-module-image-filter` ([doc](#))
- \* `nginx-plus-module-perl` ([doc](#))
- \* `nginx-plus-module-xslt` ([doc](#))

Packages with third-party modules:

- \* `nginx-plus-module-headers-more` ([site](#))
- \* `nginx-plus-module-lua` ([site](#))
- \* `nginx-plus-module-passenger` ([site](#))
- \* `nginx-plus-module-rtmp` ([site](#))
- \* `nginx-plus-module-set-misc` ([site](#))
- UDP proxy support added to the [stream](#) module.
- Added support for retrieving upstream servers configuration via DNS SRV records (the [service](#) parameter of the [server](#) directive).
- Resolver: added support for TCP fallback on retrieving large DNS responses.
- Change: requests with [non-idempotent](#) method (POST, LOCK, PATCH) are not passed to the next server in upstream group if a request has already been sent to an upstream server. Enabling the `non_idempotent` option in the [proxy\\_next\\_upstream](#) directive explicitly allows retrying such requests.
- Cache: improved meta-data accounting.
- Automatic binding of worker processes to available CPUs (the `auto` parameter of the [worker\\_cpu\\_affinity](#) directive).
- Some write operations can now be [offloaded](#) to [thread pools](#).
- Added support for [customizing](#) the Server response header field, as well as the signature in standard error messages.
- Lua module updated to version 0.10.2 (`nginx-plus-extras`, `nginx-plus-module-lua`).
- Passenger module updated to version 5.0.26 (`nginx-plus-extras`, `nginx-plus-module-passenger`).
- `headers-more` module updated to version 0.29 (`nginx-plus-extras`, `nginx-plus-module-headers-more`).
- Updated status dashboard.

- NGINX Plus R8 (1.9.9), released Dec 29, 2015

- [HTTP/2](#) support is now included into the `nginx-plus` and `nginx-plus-extras` packages. The `nginx-plus-http2` and `nginx-plus-lua` packages are deprecated.
- Caching improvements, including support of caching [HEAD](#) requests and more effective caching of big responses with the [slice](#) module.
- Dynamically configured upstream groups now can be configured to [keep states](#) between reloads.
- Support for arbitrary port in health check requests (the `port` parameter of the `health_check` directive).
- Enhancement in the [real IP](#) module: the `$realip_remote_addr` variable.
- Enhancement in [syslog](#) logging: the `nohostname` parameter.
- Lua module updated to version 0.9.20 (`nginx-plus-extras`).
- The `lua-resty-redis` Lua module updated to version 0.21 (`nginx-plus-extras`).
- Passenger module updated to version 5.0.22 (`nginx-plus-extras`).
- `headers-more` module updated to version 0.28 (`nginx-plus-extras`).
- Updated status dashboard.

- NGINX Plus R7 (1.9.4), released Sep 15, 2015
  - Introduced separate family of `nginx-plus-http2` packages with HTTP/2 support included in favor of SPDY. General `nginx-plus` packages still have SPDY support. Please refer to the [listen](#) directive documentation for the instructions on how to enable HTTP/2.
  - [TCP proxy](#) enhancements ([access](#) control; connection [limiting](#); [upload](#) and [download](#) bandwidth control; client-side [PROXY protocol](#) support; ability to [choose](#) local IP address for outgoing connections; the `backlog` parameter of the [listen](#) directive; the [tcp\\_nodelay](#) directive).
  - More efficient connections distribution between worker processes (the `reuseport` parameter of the [listen](#) directive).
  - Introduced [thread pools](#) used for multi-threaded reading and sending files without blocking worker processes.
  - Enhanced support for [modifying HTTP responses](#) (multiple substitutions support, variables support in search strings).
  - A number of additional metrics in the new version (6) of the [status dataset](#) (SSL handshakes and upstream queue overflows in particular).
  - Updated status dashboard.
  - Additional arguments to playlists in the [HLS module](#) (`start`, `end` and `offset`).
  - Support for proxying requests with [NTLM authentication](#).
  - New command-line switch to dump configuration to standard output: `-T`.
  - Added `lua-resty-redis` Lua module (`nginx-plus-extras`).
  - Lua module updated to version 0.9.16 (`nginx-plus-lua`, `nginx-plus-extras`).
  - Passenger module updated to version 5.0.15 (`nginx-plus-extras`).
  - `headers-more` module updated to version 0.26 (`nginx-plus-extras`).
  - `set-misc` module updated to version 0.29 (`nginx-plus-extras`).
- NGINX Plus R6 (1.7.11), released Apr 14, 2015
  - [TCP proxy](#) enhancements (health checks, dynamic reconfiguration, SSL support, logging, status counters).
  - New [least\\_time](#) load balancing method.
  - Unbuffered upload support ([proxy\\_request\\_buffering](#) and friends).
  - Proxy SSL authentication support for [http](#) and [uwsgi](#).
  - Proxy cache enhancements (variables support in [proxy\\_cache](#), `use_temp_path` parameter in [proxy\\_cache\\_path](#)).
  - [Client SSL certificates](#) support in mail proxy.
  - Autoindex module enhancement (the [autoindex\\_format](#) directive).
  - New status dashboard.
  - Lua module updated to version 0.9.16rc1 (`nginx-plus-lua`, `nginx-plus-extras`).
  - Passenger module updated to version 4.0.59 (`nginx-plus-extras`).
  - `set-misc` module updated to version 0.28 (`nginx-plus-extras`).
- NGINX Plus R5 (1.7.7), released Dec 1, 2014
  - New TCP proxying and load balancing mode (the [stream](#) module).

- [Sticky](#) session timeout now applies from the most recent request in the session.
  - Upstream “draining” can be used to remove an upstream server without interrupting any user sessions (the `drain` command of the [upstream\\_conf](#) dynamic configuration interface).
  - Improved control over request retries in the event of failure, based on [number of tries](#) and [time](#). Also available for fastcgi, uwsgi, scgi and memcached modules.
  - Caching: the `Vary` response header is correctly handled (multiple variants of the same resource can be cached). Note that the on-disk cache format has changed, so cached content will be invalidated after the upgrade.
  - Caching: improved support for [byte-range](#) requests.
  - Ability to control upstream bandwidth with the [proxy\\_limit\\_rate](#) directive.
  - Lua module updated to version 0.9.13 (nginx-plus-lua, nginx-plus-extras).
  - Passenger module updated to version 4.0.53 (nginx-plus-extras).
- NGINX Plus R4 (1.7.3), released Jul 22, 2014
    - [MP4](#) module now supports the `end` query argument which sets the end point of playback.
    - Added the ability to [verify](#) backend SSL certificates.
    - Added support for [SNI](#) while working with SSL backends.
    - Added conditional logging for requests (the `if` parameter of the [access\\_log](#) directive).
    - New load balancing method based on [user-defined keys](#) with optional consistency.
    - Cache revalidation now uses `If-None-Match` header if possible.
    - Passphrases for SSL private keys can now be stored in an [external file](#).
    - Introduced a new session affinity mechanism ([sticky learn](#)) based on server-initiated sessions.
    - Added the ability to retrieve a subset of the [extended status](#) data.
    - Lua module updated to version 0.9.10 (nginx-plus-lua, nginx-plus-extras).
    - Passenger module updated to version 4.0.45 (nginx-plus-extras).
- NGINX Plus R3 (1.5.12), released Apr 2, 2014
    - SPDY protocol updated to version 3.1. SPDY/2 is no longer supported.
    - Added [PROXY protocol](#) support (the `proxy_protocol` parameter of the [listen](#) directive).
    - IPv6 support added to [resolver](#).
    - DNS names in upstream groups are periodically re-resolved (the `resolve` parameter of the [server](#) directive).
    - Introduced limiting connections to [upstream servers](#) (the `max_conns` parameter) with optional support for [connections queue](#).
- NGINX Plus R2 (1.5.7), released Dec 12, 2013
    - Enhanced [sticky routing](#) support.
    - Additional [status metrics](#) for virtual hosts and cache zones.
    - [Cache purge](#) support (also available for [FastCGI](#)).
    - Added support for [cache revalidation](#).

- New module: [ngx\\_http\\_auth\\_request\\_module](#) (authorization based on the result of a subrequest).
- NGINX Plus R1 (1.5.3), released Aug 12, 2013
  - Enhanced [status](#) monitoring.
  - Load balancing: [slow start](#) feature.
  - Added syslog support for both [error\\_log](#) and [access\\_log](#).
  - Support for [Apple HTTP Live Streaming](#).
- NGINX Plus 1.5.0-2, released May 27, 2013
  - Added support for active healthchecks.
- NGINX Plus 1.5.0, released May 7, 2013
  - Security: fixed CVE-2013-2028.
- NGINX Plus 1.3.16, released Apr 19, 2013
  - Added SPDY support.
- NGINX Plus 1.3.13, released Feb 22, 2013
  - Added [sticky](#) sessions support.
  - Added support for proxying WebSocket connections.
- NGINX Plus 1.3.11, released Jan 18, 2013
  - Added base module [ngx\\_http\\_gunzip\\_module](#).
  - New extra module: [ngx\\_http\\_f4f\\_module](#) (Adobe HDS Dynamic Streaming).
  - New extra module: [ngx\\_http\\_session\\_log\\_module](#) (aggregated session logging).
- NGINX Plus 1.3.9-2, released Dec 20, 2012
  - License information updated.
  - End-User License Agreement added to the package.
- NGINX Plus 1.3.9, released Nov 27, 2012
  - Added [dynamic upstream management](#) feature.
  - PDF documentation bundled into package.
- NGINX Plus 1.3.7, released Oct 18, 2012
  - Initial release of NGINX Plus package.

# Appendix B

## Legal Notices

Open source components included in the NGINX Plus (package name is `nginx-plus`) are:

- `nginx/OSS` (1.15.10), distributed under 2-clause BSD license.  
<http://nginx.org/>

Copyright © 2002-2019 Igor Sysoev

Copyright © 2011-2019 Nginx, Inc.

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR AND CONTRIBUTORS “AS IS” AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

- Internal MD5 implementation (used only if no system MD5 support was found), based on Alexander Peslyak’s public domain implementation:  
This is an OpenSSL-compatible implementation of the RSA Data Security, Inc. MD5 Message-Digest Algorithm (RFC 1321).

Homepage:

<http://openwall.info/wiki/people/solar/software/public-domain-source-code/md5>

Author: Alexander Peslyak, better known as Solar Designer <solar at openwall.com>

This software was written by Alexander Peslyak in 2001. No copyright is claimed, and the software is hereby placed in the public domain. In case this attempt to disclaim copyright and place the software in the public domain is deemed null and void, then the software is Copyright © 2001 Alexander Peslyak and it is hereby released to the general public under the following terms:

1. Redistribution and use in source and binary forms, with or without modification, are permitted.
2. There's ABSOLUTELY NO WARRANTY, express or implied.

(This is a heavily cut-down "BSD license".)

- MurmurHash algorithm (version 2), distributed under MIT license.  
<https://sites.google.com/site/murmurhash/>

Copyright © Austin Appleby

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

The following components (babel-core, babel-polyfill, babel-plugin-react-css-modules, history, preact, webpack, npm-font-open-sans, whatwg-fetch, preact-portal) are used in status monitoring dashboard v2 only (dashboard.html in nginx-plus package):

- babel-core, Babel compiler core (6.26.0), distributed under MIT license.  
<https://github.com/babel/babel/tree/master/packages/babel-core>

Copyright © 2014-2017 Sebastian McKenzie <sebmck@gmail.com>

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

- The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

- babel-polyfill, provides polyfills necessary for a full ES2015+ environment (6.26.0), distributed under MIT license.  
<https://github.com/babel/babel/tree/master/packages/babel-polyfill>

Copyright © 2014-2017 Sebastian McKenzie <sebmck@gmail.com>

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

- The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

- babel-plugin-react-css-modules, transforms styleName to className using compile time CSS module resolution (3.3.2), distributed under 3-clause BSD license.

<https://github.com/gajus/babel-plugin-react-css-modules>

Copyright © 2016, Gajus Kuizinas (<http://gajus.com/>)

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of the Gajus Kuizinas (<http://gajus.com/>) nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL ANYONE BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

- history, manage session history with JavaScript (4.7.2), distributed under MIT license.

<https://github.com/ReactTraining/history>

Copyright © 2015-2016 Michael Jackson

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

- The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR

OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

- preact, fast 3kb React alternative with the same ES6 API (8.2.6), distributed under MIT license.

<https://github.com/developit/preact>

Copyright © 2017 Jason Miller

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

- The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

- webpack, a bundler for javascript and friends (3.9.1), distributed under MIT license.

<https://github.com/webpack/webpack>

Copyright © JS Foundation and other contributors

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the 'Software'), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

- The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED 'AS IS', WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

- npm-font-open-sans, Open Sans font family - incl. usage of CSS, SCSS, LESS (1.1.0), distributed under Apache 2.0 license.

<https://github.com/dasrick/npm-font-open-sans>

Copyright © Steve Matteson

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at: <http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

- whatwg-fetch, a window.fetch JavaScript polyfill (2.0.3), distributed under MIT license.  
<https://github.com/github/fetch>

Copyright © 2014-2016 GitHub, Inc.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

- The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

- preact-portal, render Preact components in a SPACE (1.1.2), distributed under MIT license.  
<https://github.com/developit/preact-portal>

Copyright © 2015 Jason Miller

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

- The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

The following component (Swagger UI) is distributed as a set of standalone files in the `/usr/share/nginx/html/swagger-ui` directory:

- Swagger UI (3.5.0), distributed under Apache 2.0 license.  
<https://github.com/swagger-api/swagger-ui>

Copyright 2017 SmartBear Software

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at <http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

Optional add-on and third-party modules provided with NGINX Plus may include additional open-source components. The licenses for these components are included in the installation package for each module.

# Index

absolute\_redirect, [21](#)  
accept\_mutex, [7](#)  
accept\_mutex\_delay, [7](#)  
access\_log, [183](#), [378](#)  
add\_after\_body, [59](#)  
add\_before\_body, [59](#)  
add\_header, [160](#)  
add\_trailer, [160](#)  
addition\_types, [60](#)  
aio, [21](#)  
aio\_write, [22](#)  
alias, [22](#)  
allow, [57](#), [365](#)  
ancient\_browser, [108](#)  
ancient\_browser\_value, [108](#)  
api, [62](#)  
auth\_basic, [97](#)  
auth\_basic\_user\_file, [97](#)  
auth\_http, [429](#)  
auth\_http\_header, [429](#)  
auth\_http\_pass\_client\_cert, [429](#)  
auth\_http\_timeout, [429](#)  
auth\_jwt, [100](#)  
auth\_jwt\_claim\_set, [100](#)  
auth\_jwt\_header\_set, [100](#)  
auth\_jwt\_key\_file, [101](#)  
auth\_jwt\_key\_request, [101](#)  
auth\_jwt\_leeway, [101](#)  
auth\_request, [103](#)  
auth\_request\_set, [104](#)  
autoindex, [105](#)  
autoindex\_exact\_size, [105](#)  
autoindex\_format, [105](#)  
autoindex\_localtime, [106](#)  
  
break, [239](#)  
  
charset, [109](#)  
charset\_map, [110](#)  
charset\_types, [111](#)  
  
chunked\_transfer\_encoding, [23](#)  
client\_body\_buffer\_size, [23](#)  
client\_body\_in\_file\_only, [23](#)  
client\_body\_in\_single\_buffer, [24](#)  
client\_body\_temp\_path, [24](#)  
client\_body\_timeout, [24](#)  
client\_header\_buffer\_size, [25](#)  
client\_header\_timeout, [25](#)  
client\_max\_body\_size, [25](#)  
connection\_pool\_size, [25](#)  
create\_full\_put\_path, [112](#)  
  
daemon, [7](#)  
dav\_access, [113](#)  
dav\_methods, [113](#)  
debug\_connection, [8](#)  
debug\_points, [8](#)  
default\_type, [26](#)  
deny, [57](#), [365](#)  
directio, [26](#)  
directio\_alignment, [26](#)  
disable\_symlinks, [26](#)  
  
empty\_gif, [115](#)  
env, [8](#)  
error\_log, [9](#)  
error\_page, [27](#)  
etag, [28](#)  
events, [10](#)  
expires, [161](#)  
  
f4f, [116](#)  
f4f\_buffer\_size, [116](#)  
fastcgi\_bind, [118](#)  
fastcgi\_buffer\_size, [119](#)  
fastcgi\_buffering, [119](#)  
fastcgi\_buffers, [119](#)  
fastcgi\_busy\_buffers\_size, [120](#)  
fastcgi\_cache, [120](#)  
fastcgi\_cache\_background\_update, [120](#)

fastcgi\_cache\_bypass, 120  
fastcgi\_cache\_key, 121  
fastcgi\_cache\_lock, 121  
fastcgi\_cache\_lock\_age, 121  
fastcgi\_cache\_lock\_timeout, 121  
fastcgi\_cache\_max\_range\_offset, 122  
fastcgi\_cache\_methods, 122  
fastcgi\_cache\_min\_uses, 122  
fastcgi\_cache\_path, 122  
fastcgi\_cache\_purge, 124  
fastcgi\_cache\_revalidate, 125  
fastcgi\_cache\_use\_stale, 125  
fastcgi\_cache\_valid, 125  
fastcgi\_catch\_stderr, 126  
fastcgi\_connect\_timeout, 127  
fastcgi\_force\_ranges, 127  
fastcgi\_hide\_header, 127  
fastcgi\_ignore\_client\_abort, 127  
fastcgi\_ignore\_headers, 128  
fastcgi\_index, 128  
fastcgi\_intercept\_errors, 128  
fastcgi\_keep\_conn, 129  
fastcgi\_limit\_rate, 129  
fastcgi\_max\_temp\_file\_size, 129  
fastcgi\_next\_upstream, 129  
fastcgi\_next\_upstream\_timeout, 130  
fastcgi\_next\_upstream\_tries, 131  
fastcgi\_no\_cache, 131  
fastcgi\_param, 131  
fastcgi\_pass, 132  
fastcgi\_pass\_header, 132  
fastcgi\_pass\_request\_body, 132  
fastcgi\_pass\_request\_headers, 133  
fastcgi\_read\_timeout, 133  
fastcgi\_request\_buffering, 133  
fastcgi\_send\_lowat, 133  
fastcgi\_send\_timeout, 134  
fastcgi\_socket\_keepalive, 134  
fastcgi\_split\_path\_info, 134  
fastcgi\_store, 134  
fastcgi\_store\_access, 135  
fastcgi\_temp\_file\_write\_size, 136  
fastcgi\_temp\_path, 136  
flv, 138  
geo, 139, 366  
geoip\_city, 143, 369  
geoip\_country, 142, 368  
geoip\_org, 144, 370  
geoip\_proxy, 144  
geoip\_proxy\_recursive, 144  
grpc\_bind, 146  
grpc\_buffer\_size, 146  
grpc\_connect\_timeout, 146  
grpc\_hide\_header, 146  
grpc\_ignore\_headers, 147  
grpc\_intercept\_errors, 147  
grpc\_next\_upstream, 147  
grpc\_next\_upstream\_timeout, 148  
grpc\_next\_upstream\_tries, 148  
grpc\_pass, 149  
grpc\_pass\_header, 149  
grpc\_read\_timeout, 149  
grpc\_send\_timeout, 150  
grpc\_set\_header, 150  
grpc\_socket\_keepalive, 150  
grpc\_ssl\_certificate, 150  
grpc\_ssl\_certificate\_key, 151  
grpc\_ssl\_ciphers, 151  
grpc\_ssl\_crl, 151  
grpc\_ssl\_name, 151  
grpc\_ssl\_password\_file, 151  
grpc\_ssl\_protocols, 152  
grpc\_ssl\_server\_name, 152  
grpc\_ssl\_session\_reuse, 152  
grpc\_ssl\_trusted\_certificate, 152  
grpc\_ssl\_verify, 152  
grpc\_ssl\_verify\_depth, 152  
gunzip, 154  
gunzip\_buffers, 154  
gzip, 155  
gzip\_buffers, 156  
gzip\_comp\_level, 156  
gzip\_disable, 156  
gzip\_http\_version, 156  
gzip\_min\_length, 156  
gzip\_proxied, 157  
gzip\_static, 159  
gzip\_types, 157  
gzip\_vary, 158  
hash, 308, 410  
health\_check, 323, 414  
health\_check\_timeout, 415

- hls, 164
- hls\_buffers, 164
- hls\_forward\_args, 164
- hls\_fragment, 165
- hls\_mp4\_buffer\_size, 165
- hls\_mp4\_max\_buffer\_size, 166
- http, 29
- http2\_body\_preread\_size, 351
- http2\_chunk\_size, 351
- http2\_idle\_timeout, 351
- http2\_max\_concurrent\_pushes, 351
- http2\_max\_concurrent\_streams, 351
- http2\_max\_field\_size, 352
- http2\_max\_header\_size, 352
- http2\_max\_requests, 352
- http2\_push, 352
- http2\_push\_preload, 353
- http2\_recv\_buffer\_size, 353
- http2\_recv\_timeout, 353
- if, 240
- if\_modified\_since, 29
- ignore\_invalid\_headers, 29
- image\_filter, 168
- image\_filter\_buffer, 169
- image\_filter\_interlace, 169
- image\_filter\_jpeg\_quality, 169
- image\_filter\_sharpen, 169
- image\_filter\_transparency, 169
- image\_filter\_webp\_quality, 170
- imap\_auth, 443
- imap\_capabilities, 443
- imap\_client\_buffer, 443
- include, 10
- index, 171
- internal, 29
- ip\_hash, 308
- js\_access, 372
- js\_content, 173
- js\_filter, 373
- js\_include, 173, 373
- js\_path, 174, 373
- js\_preread, 373
- js\_set, 174, 373
- keepalive, 309
- keepalive\_disable, 30
- keepalive\_requests, 30, 310
- keepalive\_timeout, 31, 311
- keyval, 175, 374
- keyval\_zone, 175, 375
- large\_client\_header\_buffers, 31
- least\_conn, 311, 411
- least\_time, 312, 411
- limit\_conn, 177, 376
- limit\_conn\_log\_level, 178, 377
- limit\_conn\_status, 178
- limit\_conn\_zone, 178, 377
- limit\_except, 31
- limit\_rate, 32
- limit\_rate\_after, 32
- limit\_req, 180
- limit\_req\_log\_level, 181
- limit\_req\_status, 181
- limit\_req\_zone, 182
- limit\_zone, 179
- lingering\_close, 33
- lingering\_time, 33
- lingering\_timeout, 33
- listen, 34, 358, 425
- load\_module, 10
- location, 37
- lock\_file, 10
- log\_format, 185, 379
- log\_not\_found, 38
- log\_subrequest, 39
- mail, 427
- map, 187, 381
- map\_hash\_bucket\_size, 189, 382
- map\_hash\_max\_size, 189, 383
- master\_process, 11
- match, 324, 415
- max\_ranges, 39
- memcached\_bind, 190
- memcached\_buffer\_size, 191
- memcached\_connect\_timeout, 191
- memcached\_force\_ranges, 191
- memcached\_gzip\_flag, 192
- memcached\_next\_upstream, 192
- memcached\_next\_upstream\_timeout, 192

memcached\_next\_upstream\_tries, 193  
memcached\_pass, 193  
memcached\_read\_timeout, 193  
memcached\_send\_timeout, 193  
memcached\_socket\_keepalive, 194  
merge\_slashes, 39  
min\_delete\_depth, 113  
mirror, 195  
mirror\_request\_body, 195  
modern\_browser, 108  
modern\_browser\_value, 108  
mp4, 198  
mp4\_buffer\_size, 198  
mp4\_limit\_rate, 199  
mp4\_limit\_rate\_after, 199  
mp4\_max\_buffer\_size, 198  
msie\_padding, 40  
msie\_refresh, 40  
multi\_accept, 11  
  
ntlm, 311  
  
open\_file\_cache, 40  
open\_file\_cache\_errors, 41  
open\_file\_cache\_min\_uses, 41  
open\_file\_cache\_valid, 41  
open\_log\_file\_cache, 186, 379  
output\_buffers, 41  
override\_charset, 111  
  
pcre\_jit, 11  
perl, 201  
perl\_modules, 202  
perl\_require, 202  
perl\_set, 202  
pid, 11  
pop3\_auth, 445  
pop3\_capabilities, 445  
port\_in\_redirect, 41  
postpone\_output, 42  
pread\_buffer\_size, 360  
pread\_timeout, 360  
protocol, 427  
proxy\_bind, 208, 385  
proxy\_buffer, 433  
proxy\_buffer\_size, 208, 385  
proxy\_buffering, 208  
proxy\_buffers, 209  
proxy\_busy\_buffers\_size, 209  
proxy\_cache, 209  
proxy\_cache\_background\_update, 209  
proxy\_cache\_bypass, 210  
proxy\_cache\_convert\_head, 210  
proxy\_cache\_key, 210  
proxy\_cache\_lock, 210  
proxy\_cache\_lock\_age, 211  
proxy\_cache\_lock\_timeout, 211  
proxy\_cache\_max\_range\_offset, 211  
proxy\_cache\_methods, 211  
proxy\_cache\_min\_uses, 212  
proxy\_cache\_path, 212  
proxy\_cache\_purge, 214  
proxy\_cache\_revalidate, 214  
proxy\_cache\_use\_stale, 214  
proxy\_cache\_valid, 215  
proxy\_connect\_timeout, 216, 386  
proxy\_cookie\_domain, 216  
proxy\_cookie\_path, 217  
proxy\_download\_rate, 386  
proxy\_force\_ranges, 218  
proxy\_headers\_hash\_bucket\_size, 218  
proxy\_headers\_hash\_max\_size, 218  
proxy\_hide\_header, 218  
proxy\_http\_version, 219  
proxy\_ignore\_client\_abort, 219  
proxy\_ignore\_headers, 219  
proxy\_intercept\_errors, 219  
proxy\_limit\_rate, 220  
proxy\_max\_temp\_file\_size, 220  
proxy\_method, 220  
proxy\_next\_upstream, 221, 386  
proxy\_next\_upstream\_timeout, 222, 386  
proxy\_next\_upstream\_tries, 222, 386  
proxy\_no\_cache, 222  
proxy\_pass, 222, 387  
proxy\_pass\_error\_message, 433  
proxy\_pass\_header, 224  
proxy\_pass\_request\_body, 224  
proxy\_pass\_request\_headers, 224  
proxy\_protocol, 387  
proxy\_protocol\_timeout, 361  
proxy\_read\_timeout, 225

proxy\_redirect, 225  
proxy\_request\_buffering, 226  
proxy\_requests, 387  
proxy\_responses, 388  
proxy\_send\_lowat, 227  
proxy\_send\_timeout, 227  
proxy\_session\_drop, 388  
proxy\_set\_body, 227  
proxy\_set\_header, 227  
proxy\_socket\_keepalive, 228, 388  
proxy\_ssl, 388  
proxy\_ssl\_certificate, 228, 389  
proxy\_ssl\_certificate\_key, 229, 389  
proxy\_ssl\_ciphers, 229, 389  
proxy\_ssl\_crl, 229, 389  
proxy\_ssl\_name, 229, 389  
proxy\_ssl\_password\_file, 230, 390  
proxy\_ssl\_protocols, 230, 390  
proxy\_ssl\_server\_name, 230, 390  
proxy\_ssl\_session\_reuse, 230, 390  
proxy\_ssl\_trusted\_certificate, 230, 390  
proxy\_ssl\_verify, 231, 391  
proxy\_ssl\_verify\_depth, 231, 391  
proxy\_store, 231  
proxy\_store\_access, 232  
proxy\_temp\_file\_write\_size, 232  
proxy\_temp\_path, 233  
proxy\_timeout, 391, 433  
proxy\_upload\_rate, 391  
  
queue, 312  
  
random, 313, 411  
random\_index, 234  
read\_ahead, 42  
real\_ip\_header, 235  
real\_ip\_recursive, 236  
recursive\_error\_pages, 42  
referer\_hash\_bucket\_size, 237  
referer\_hash\_max\_size, 237  
request\_pool\_size, 42  
reset\_timedout\_connection, 42  
resolver, 43, 361, 427  
resolver\_timeout, 43, 361, 428  
return, 241, 393  
rewrite, 241  
rewrite\_log, 243  
  
root, 44  
  
satisfy, 44  
scgi\_bind, 246  
scgi\_buffer\_size, 246  
scgi\_buffering, 247  
scgi\_buffers, 247  
scgi\_busy\_buffers\_size, 247  
scgi\_cache, 248  
scgi\_cache\_background\_update, 248  
scgi\_cache\_bypass, 248  
scgi\_cache\_key, 248  
scgi\_cache\_lock, 248  
scgi\_cache\_lock\_age, 249  
scgi\_cache\_lock\_timeout, 249  
scgi\_cache\_max\_range\_offset, 249  
scgi\_cache\_methods, 249  
scgi\_cache\_min\_uses, 250  
scgi\_cache\_path, 250  
scgi\_cache\_purge, 252  
scgi\_cache\_revalidate, 252  
scgi\_cache\_use\_stale, 252  
scgi\_cache\_valid, 253  
scgi\_connect\_timeout, 254  
scgi\_force\_ranges, 254  
scgi\_hide\_header, 254  
scgi\_ignore\_client\_abort, 255  
scgi\_ignore\_headers, 255  
scgi\_intercept\_errors, 255  
scgi\_limit\_rate, 255  
scgi\_max\_temp\_file\_size, 256  
scgi\_next\_upstream, 256  
scgi\_next\_upstream\_timeout, 257  
scgi\_next\_upstream\_tries, 257  
scgi\_no\_cache, 258  
scgi\_param, 258  
scgi\_pass, 258  
scgi\_pass\_header, 259  
scgi\_pass\_request\_body, 259  
scgi\_pass\_request\_headers, 259  
scgi\_read\_timeout, 259  
scgi\_request\_buffering, 260  
scgi\_send\_timeout, 260  
scgi\_socket\_keepalive, 260  
scgi\_store, 260  
scgi\_store\_access, 261  
scgi\_temp\_file\_write\_size, 261

scgi\_temp\_path, 262  
secure\_link, 263  
secure\_link\_md5, 264  
secure\_link\_secret, 264  
send\_lowat, 44  
send\_timeout, 45  
sendfile, 45  
sendfile\_max\_chunk, 45  
server, 46, 304, 362, 407, 428  
server\_name, 46, 428  
server\_name\_in\_redirect, 48  
server\_names\_hash\_bucket\_size, 48  
server\_names\_hash\_max\_size, 48  
server\_tokens, 48  
session\_log, 266  
session\_log\_format, 266  
session\_log\_zone, 267  
set, 243  
set\_real\_ip\_from, 235, 392  
slice, 268  
smtp\_auth, 446  
smtp\_capabilities, 446  
smtp\_client\_buffer, 447  
smtp\_greeting\_delay, 447  
source\_charset, 111  
split\_clients, 270, 394  
ssi, 271  
ssi\_last\_modified, 271  
ssi\_min\_file\_chunk, 272  
ssi\_silent\_errors, 272  
ssi\_types, 272  
ssi\_value\_length, 272  
ssl, 278, 436  
ssl\_buffer\_size, 278  
ssl\_certificate, 279, 396, 436  
ssl\_certificate\_key, 279, 397, 437  
ssl\_ciphers, 280, 397, 437  
ssl\_client\_certificate, 280, 397, 437  
ssl\_crl, 280, 398, 438  
ssl\_dhparam, 281, 398, 438  
ssl\_early\_data, 281  
sslecdh\_curve, 281, 398, 438  
ssl\_engine, 11  
ssl\_handshake\_timeout, 398  
ssl\_password\_file, 282, 399, 438  
ssl\_prefer\_server\_ciphers, 282, 399, 439  
ssl\_preread, 405  
ssl\_protocols, 282, 399, 439  
ssl\_session\_cache, 283, 400, 439  
ssl\_session\_ticket\_key, 283, 400, 440  
ssl\_session\_tickets, 284, 401, 441  
ssl\_session\_timeout, 284, 401, 441  
ssl\_stapling, 284  
ssl\_stapling\_file, 285  
ssl\_stapling\_responder, 285  
ssl\_stapling\_verify, 285  
ssl\_trusted\_certificate, 285, 401, 441  
ssl\_verify\_client, 286, 401, 441  
ssl\_verify\_depth, 286, 402, 441  
starttls, 442  
state, 307, 410  
status, 290  
status\_format, 290  
status\_zone, 63, 291  
sticky, 313  
sticky\_cookie\_insert, 316  
stream, 362  
stub\_status, 299  
sub\_filter, 301  
sub\_filter\_last\_modified, 301  
sub\_filter\_once, 302  
sub\_filter\_types, 302  
subrequest\_output\_buffer\_size, 49  
  
tcp\_nodelay, 49, 362  
tcp\_nopush, 49  
thread\_pool, 12  
timeout, 428  
timer\_resolution, 12  
try\_files, 49  
types, 51  
types\_hash\_bucket\_size, 52  
types\_hash\_max\_size, 52  
  
underscores\_in\_headers, 52  
uninitialized\_variable\_warn, 243  
upstream, 304, 407  
upstream\_conf, 318  
use, 12  
user, 13  
userid, 326  
userid\_domain, 327  
userid\_expires, 327

- userid\_mark, [327](#)
- userid\_name, [327](#)
- userid\_p3p, [328](#)
- userid\_path, [328](#)
- userid\_service, [328](#)
- uwsgi\_bind, [330](#)
- uwsgi\_buffer\_size, [331](#)
- uwsgi\_buffering, [331](#)
- uwsgi\_buffers, [331](#)
- uwsgi\_busy\_buffers\_size, [332](#)
- uwsgi\_cache, [332](#)
- uwsgi\_cache\_background\_update, [332](#)
- uwsgi\_cache\_bypass, [332](#)
- uwsgi\_cache\_key, [333](#)
- uwsgi\_cache\_lock, [333](#)
- uwsgi\_cache\_lock\_age, [333](#)
- uwsgi\_cache\_lock\_timeout, [333](#)
- uwsgi\_cache\_max\_range\_offset, [334](#)
- uwsgi\_cache\_methods, [334](#)
- uwsgi\_cache\_min\_uses, [334](#)
- uwsgi\_cache\_path, [334](#)
- uwsgi\_cache\_purge, [336](#)
- uwsgi\_cache\_revalidate, [337](#)
- uwsgi\_cache\_use\_stale, [337](#)
- uwsgi\_cache\_valid, [337](#)
- uwsgi\_connect\_timeout, [338](#)
- uwsgi\_force\_ranges, [339](#)
- uwsgi\_hide\_header, [339](#)
- uwsgi\_ignore\_client\_abort, [339](#)
- uwsgi\_ignore\_headers, [339](#)
- uwsgi\_intercept\_errors, [340](#)
- uwsgi\_limit\_rate, [340](#)
- uwsgi\_max\_temp\_file\_size, [340](#)
- uwsgi\_modifier1, [340](#)
- uwsgi\_modifier2, [341](#)
- uwsgi\_next\_upstream, [341](#)
- uwsgi\_next\_upstream\_timeout, [342](#)
- uwsgi\_next\_upstream\_tries, [342](#)
- uwsgi\_no\_cache, [342](#)
- uwsgi\_param, [342](#)
- uwsgi\_pass, [343](#)
- uwsgi\_pass\_header, [343](#)
- uwsgi\_pass\_request\_body, [344](#)
- uwsgi\_pass\_request\_headers, [344](#)
- uwsgi\_read\_timeout, [344](#)
- uwsgi\_request\_buffering, [344](#)
- uwsgi\_send\_timeout, [345](#)
- uwsgi\_socket\_keepalive, [345](#)
- uwsgi\_ssl\_certificate, [345](#)
- uwsgi\_ssl\_certificate\_key, [345](#)
- uwsgi\_ssl\_ciphers, [345](#)
- uwsgi\_ssl\_crl, [346](#)
- uwsgi\_ssl\_name, [346](#)
- uwsgi\_ssl\_password\_file, [346](#)
- uwsgi\_ssl\_protocols, [346](#)
- uwsgi\_ssl\_server\_name, [347](#)
- uwsgi\_ssl\_session\_reuse, [347](#)
- uwsgi\_ssl\_trusted\_certificate, [347](#)
- uwsgi\_ssl\_verify, [347](#)
- uwsgi\_ssl\_verify\_depth, [347](#)
- uwsgi\_store, [348](#)
- uwsgi\_store\_access, [348](#)
- uwsgi\_temp\_file\_write\_size, [349](#)
- uwsgi\_temp\_path, [349](#)
- valid\_referers, [238](#)
- variables\_hash\_bucket\_size, [53](#), [362](#)
- variables\_hash\_max\_size, [53](#), [362](#)
- worker\_aio\_requests, [13](#)
- worker\_connections, [13](#)
- worker\_cpu\_affinity, [13](#)
- worker\_priority, [14](#)
- worker\_processes, [14](#)
- worker\_rlimit\_core, [15](#)
- worker\_rlimit\_nofile, [15](#)
- worker\_shutdown\_timeout, [15](#)
- working\_directory, [15](#)
- xclient, [434](#)
- xml\_entities, [354](#)
- xslt\_last\_modified, [355](#)
- xslt\_param, [355](#)
- xslt\_string\_param, [355](#)
- xslt\_stylesheet, [355](#)
- xslt\_types, [356](#)
- zone, [307](#), [409](#)
- zone\_sync, [418](#)
- zone\_sync\_buffers, [419](#)
- zone\_sync\_connect\_retry\_interval, [419](#)
- zone\_sync\_connect\_timeout, [419](#)
- zone\_sync\_interval, [419](#)
- zone\_sync\_recv\_buffer\_size, [419](#)

---

`zone_sync_server`, [420](#)  
`zone_sync_ssl`, [420](#)  
`zone_sync_ssl_certificate`, [420](#)  
`zone_sync_ssl_certificate_key`, [421](#)  
`zone_sync_ssl_ciphers`, [421](#)  
`zone_sync_ssl_crl`, [421](#)  
`zone_sync_ssl_name`, [421](#)  
`zone_sync_ssl_password_file`, [421](#)  
`zone_sync_ssl_protocols`, [422](#)  
`zone_sync_ssl_server_name`, [422](#)  
`zone_sync_ssl_trusted_certificate`, [422](#)  
`zone_sync_ssl_verify`, [422](#)  
`zone_sync_ssl_verify_depth`, [422](#)  
`zone_sync_timeout`, [422](#)