

# Crossroads 1.16

Karel Kubat  
e-tunity

2005, 2006, ff.

## Abstract

Crossroads is a load balance and fail over utility for TCP based services. It is a daemon program running in user space, and features extensive configurability, polling of back ends using 'wakeup calls', detailed status reporting, 'hooks' for special actions when backend calls fail, and much more. Crossroads is service-independent: it is usable for HTTP/HTTPS, SSH, SMTP, DNS, etc. In the case of HTTP balancing, Crossroads can modify HTTP headers, e.g. to provide 'session stickiness' for back-end processes that need sessions, but aren't session-aware of other back-ends.

## Contents

<b>1</b>	<b>Introduction</b>	<b>3</b>
1.1	Obtaining Crossroads . . . . .	3
1.2	Copyright and Disclaimer . . . . .	3
1.3	Terminology . . . . .	3
1.4	Porting issues for pre-0.26 installations . . . . .	5
1.5	Porting issues for pre-1.08 installations . . . . .	5
<b>2</b>	<b>Installation for the impatient</b>	<b>5</b>
<b>3</b>	<b>Using Crossroads</b>	<b>6</b>
3.1	General Commandline Syntax . . . . .	6
3.2	Logging-related options . . . . .	6
<b>4</b>	<b>The configuration</b>	<b>7</b>
4.1	General language elements . . . . .	7
4.1.1	Empty lines and comments . . . . .	7
4.1.2	Keywords, numbers, identifiers, generic strings . . . . .	7
4.2	Service definitions . . . . .	8
4.3	Backend definitions . . . . .	10
4.3.1	General Backend Directives . . . . .	11
4.3.2	HTTP-related Backend Directives . . . . .	12
<b>5</b>	<b>Tips, Tricks and Remarks</b>	<b>15</b>
5.1	How back ends are selected in load balancing . . . . .	15
5.1.1	Bysize, byduration or byconnections? . . . . .	15
5.1.2	Averaging size and duration . . . . .	15
5.1.3	Specifying decays . . . . .	16
5.1.4	Adjusting the weights . . . . .	16

5.1.5	Throttling the number of concurrent connections . . . . .	17
5.2	HTTP Session Stickiness . . . . .	17
5.2.1	Don't use stickiness! . . . . .	18
5.2.2	But if you must.. . . . .	18
5.3	Passing the client's IP address . . . . .	19
5.3.1	Sample Crossroads configuration . . . . .	19
5.3.2	Sample Apache configuration . . . . .	19
5.4	Debugging network traffic . . . . .	20
5.5	Configuration examples . . . . .	21
5.5.1	A load balancer for three webserver back ends . . . . .	21
5.5.2	An HTTP forwarder when travelling . . . . .	24
5.5.3	SSH login with enforced idle logout . . . . .	25
<b>6</b>	<b>Benchmarking</b>	<b>25</b>
6.1	Benchmark 1: Accessing a proxy via crossroads or directly . . . . .	25
6.1.1	Results . . . . .	26
6.1.2	Discussion . . . . .	26
6.2	Benchmark 2: Crossroads versus Linux Virtual Server (LVS) . . . . .	26
6.2.1	Environment . . . . .	27
6.2.2	Tests and results . . . . .	27
<b>7</b>	<b>Compiling and Installing</b>	<b>28</b>
7.1	Prerequisites . . . . .	28
7.2	Compiling and installing . . . . .	28
7.3	Configuring crossroads . . . . .	28
7.4	A boot script . . . . .	29
7.4.1	SysV Style Startup . . . . .	29
7.4.2	BSD Style Startup . . . . .	30

## 1 Introduction

Crossroads is a daemon that basically accepts TCP connections at preconfigured ports, and given a list of 'back ends' distributes each incoming connection to one of the back ends, so that a client request is served. Additionally, crossroads maintains an internal administration of the back end connectivity: if a back end isn't usable, then the client request is handled using another back end. Crossroads will then periodically check whether a previously not usable back end has come to life yet. Also, crossroads can select back ends by estimating the load, so that balancing is achieved.

Using this approach, crossroads serves as load balancer and fail over utility. Crossroads will very likely not be as reliable as hardware based balancers, since it always will require a server to run on. This server, in turn, may become a new Single Point of Failure (SPOS). However, in situations where cost efficiency is an issue, crossroads may be a good choice. Furthermore, crossroads can be deployed in situations where a hardware based balancing already exists and augmenting service reliability is needed. Or, crossroads may be run off a diskless system, which again improves reliability of the underlying hardware.

This document describes how to use crossroads, how to configure it in order to increase the reliability of your systems, and how to compile the program from its sources.

### 1.1 Obtaining Crossroads

As quick reference, here are some important URL's for Crossroads:

- <http://crossroads.e-tunity.com> is the site that serves Crossroads. You can browse this at leisure for documentation, sources, and so on.
- <http://freshmeat.net/projects/crossr> is the Freshmeat announcement page.
- <svn://svn.e-tunity.com/crossroads> is the SVN repository; anonymous reading (fetching) is allowed. In order to commit changes, mail me<sup>1</sup> for credentials.

### 1.2 Copyright and Disclaimer

Crossroads is distributed as-is, without assumptions of fitness or usability. You are free to use crossroads to your liking. It's free, and as with everything that's free: there's also no warranty.

You are allowed to make modifications to the source code of crossroads, and you are allowed to (re)distribute crossroads, as long as you include this text, all sources, and if applicable: all your modifications, with each distribution.

While you are allowed to make any and all changes to the sources, I would appreciate hearing about them. If the changes concern new functionality or bugfixes, then I'll include them in a next release, stating full credits. If you want to seriously contribute (to which you are heartily encouraged), then mail me and I'll get you access to the Crossroads SVN repository, so that you can update and commit as you like.

### 1.3 Terminology

Throughout this document, the following terms are used: <sup>2</sup>

**A client** is a process that initiates a network connection to get contact with some service.

**A service** or **server process** or **listener** is a central application that accepts network connections from clients and services them.

---

<sup>1</sup><mailto:karel@e-tunity.com>

<sup>2</sup>Many more meanings of the terms will exist – yes, I am aware of that. I'm using the terms here in a very strict sense.

**Back ends** are locations where crossroads looks in order to service its clients. Crossroads sits 'in between' and does its tricks. Therefore, as far as the back ends are concerned, crossroads behaves like a client. As far as the true client is concerned, crossroads behaves like the service. The communication is however transparent: neither client nor back end are aware of the middle position of crossroads.

**A connection** is a network conversation between client and service, where data are transferred to and fro. As far as crossroads is concerned, success means that a connection can be established without errors on the network level. Crossroads isn't aware of service peculiarities. E.g., when a webserver answers `HTTP/1.0 500 Server Error` then crossroads will see this as a succesful connection, though the user behind a browser may think otherwise.

**Back end selection algorithms** are methods by which crossroads determines which back end it will talk to next. Crossroads has a number of built-in algorithms, which may be configured per service.

**Back end states** are the statuses of each back end that is known to crossroads. A back end may be available, (temporarily) unavailable or truly down. When a back end is temporarily unavailable, then crossroads will periodically check whether the back end has come to life yet (that is, if configured so).

**A spike** is a sudden increase in activity, leading to extra load on a given service. When crossroads is in effect and when the spike occurs in one connection, then obviously the spike will also appear at one of the back ends. However, crossroads will see the spike and will make sure that a subsequent request goes to an other back end. In contrast, when several connections arrive simultaneously and cause a spike, then crossroads will be able to distribute the connections over several back ends, thereby 'flattening out' the increase.

**Load balancing** means that incoming client requests are distributed over more than just one back end (which wouldn't be the case if you wouldn't be running crossroads). Enabling load balancing is nothing more than duplicating services over more than one back end, and having something (in this case: crossroads) distribute the requests, so that per back end the load doesn't get too high.

**An HTTP session** is a series of separate network connections that originate from one browser. E.g., to fill the display with text and images, the browser hits a website several times. An HTTP session may even span several screens. E.g., a website registration dialog may involve 3 screens that when called from the same browser, form a logical group of some sort.

**Headers** or **header lines** are specific parts of an HTTP message. Crossroads has directives to add or modify headers that are part of the request that a browser sends to server, or those that are part of the server.

**Session stickiness** means that when a browser starts an HTTP dialog, the balancer makes sure that it 'sticks' to the same back end (i.e., subsequent requests from the browser are forced to go to the same back end, instead of being balanced to other ones).

**Back end usage** is measured by crossroads in order to be able to determine back end selection. Crossroads stores information about the number of active connections, the transferred bytes and about the connection duration. These numbers can be used to estimate which back end is the least used – and therefore, presumably, the best candidate for a new request.

**Fail over** is almost always used when load balancing is in effect. The distributor of client requests (crossroads of course) can also monitor back ends, so that incase a back end is 'down', it is no longer accessed.

**Service downtime** normally occurs when a service is switched off. Downtime is obviously avoided when fail over is in effect: a back end can be taken out of service in a controlled manner, without any client noticing it.

## 1.4 Porting issues for pre-0.26 installations

As of version 0.26 the syntax of the configuration file has changed. In particular:

- The keyword `maxconnections` is now used instead of `maxclients`;
- The keyword `connectiontimeout` is now used instead of `sessiontimeout`.

Therefore when converting configuration files to the new syntax, the above keywords must be changed. (The reason for these changes is that 0.26 introduces *sticky HTTP sessions* that span multiple TCP connections, and the term *session* is used strictly in that sense – and no longer for a TCP connection.)

## 1.5 Porting issues for pre-1.08 installations

As of version 1.08, the following directives no longer are supported:

- `insertstickycookie` was replaced by the more generic directive `addclientheader`. E.g., instead of  
`insertstickycookie "XRID=100; Path=/"`;  
the syntax is now  
`addclientheader "Set-Cookie: XRID=100; Path=/"`;
- `insertrealip` was replaced by the more generic directive `setserverheader`. E.g., instead of  
`insertrealip on`;  
the syntax is now  
`setserverheader "XR-Real-IP: %r"`;  
This incidentally also makes it possible to change the header name (here: `XR-Real-IP`).

# 2 Installation for the impatient

For the impatient, here's the very-quick-but-very-superficial recipe for getting crossroads up and running:

- If you don't have SVN or don't want to use it:
  - Obtain the crossroads source archive at <http://crossroads.e-tunity.com>.
  - Change-dir to a 'sources' directory on your system and unpack the archive.
  - Change-dir into the create directory `crossroads/`.
- If you have SVN and want to go for the newest snapshot:
  - Get the latest sources and snapshots using SVN from `svn://svn.e-tunity.com/crossroads`.
  - You'll find the newest alpha version under `crossroads/trunk` and the stable versions under `crossroads/tags`, e.g. `crossroads/tags/release-1.00`.
  - Choose which you want to use: the latest stable release, or the bleeding edge alpha? In the former case, change-dir to `crossroads/tags/release-X.YY`, where `X.YY` is a release ID. In the latter case, change-dir to `crossroads/trunk`.
- Type `make install`. This installs the crossroads binary into `/usr/local/bin/`. If the compilation doesn't work on your system, check `etc/Makefile.def` for hints.
- Create a file `/etc/crossroads.conf`. In it state something like:

```
service www {  
    port 80;  
    revivinginterval 15;
```

```
    backend one {
        server 10.1.1.100:80;
    }
    backend two {
        server 10.1.1.101:80;
    }
}
```

That's off course assuming that you want to balance HTTP on port 80 to two back ends at 10.1.1.100 and 10.1.1.101.

- Type `crossroads start`.
- Surf to the machine where `crossroads` is running. You will see the pages served by the back ends 10.1.1.100 or 10.1.1.101.
- To monitor the status of `crossroads`, type `crossroads status`.

## 3 Using Crossroads

Crossroads is started from the commandline, and highly depends on `/etc/crossroads.conf` (the default configuration file). It supports a number of flags (e.g., to overrule the location of the configuration file). The actual usage information is always obtained by typing `crossroads` without any arguments. Crossroads then displays the allowed arguments.

### 3.1 General Commandline Syntax

This section shows the most basic usage. As said above, start `crossroads` without arguments to view the full listing of options.

- `crossroads start` and `crossroads stop` are typical actions that are run from system startup scripts. The meaning is self-explanatory.
- `crossroads restart` is a combination of the former two. Beware that a restart may cause discontinuity in service; it is just a shorthand for typing the 'stop' and 'start' actions after one another.
- `crossroad status` reports on each running service. Per service, the state of each back end is reported.
- `crossroads tell service backend state` is a command line way of telling crossroads that a given back end, of a given service, is in a given state. Normally crossroads maintains state information itself, but by using `crossroads tell`, a back end can be e.g. taken 'off line' for servicing.
- `crossroads configtest` tells you whether the configuration is syntactially correct.
- `crossroads services` reports on the configured services. In contrast to `crossroads status`, this option only shows what's configured – not what's up and running. Therefore, `crossroads services` doesn't report on back end states.
- `crossroads sampleconf` shows a sample configuration on screen. A good way of quickly viewing the configuration file syntax, or of getting a start for your own configuration `/etc/crossroads.conf`.

### 3.2 Logging-related options

Two 'flags' of Crossroads are specifically logging-related. This section elaborates on these flags.

First, there's flag `-a`. When present, the start and end of activity is logged using statements like

YYYY-MM-DD HH/MM/SS starting http from 61.45.32.189 to 10.1.1.1

Similarly, there are 'ending' statements. Using this flag and scanning your logs for these statements may be helpful in quickly determining your system load.

Second, there's flag `-l`. This flag selects the 'facility' of logging and defaults to `LOG_DAEMON`. You can supply a number between 0 and 7 to flag `-l` to select `LOG_LOCAL0` to `LOG_LOCAL7`. This would separate the Crossroads-related logging from other streams. Here's a very short guide; please read your Unix manpages of `syslogd` for more information.

- First edit `/etc/syslog.conf` and add a line:

```
local7.*    /var/log/crossroads.log
```

That instructs `syslogd` to send `LOG_LOCAL7` requests to the logfile `/var/log/crossroads.log`.

- Next, restart `syslogd`. On most Unices that's done by issuing `killall -1 syslogd`. (As a side-note, I tried this once on an Bull/AIX system, and the box just shut down. The `killall` command killed every process...)
- Now start `crossroads` with the flag `-l7`.
- Finally, monitor `/var/log/crossroads.log` for Crossroads' messages.

## 4 The configuration

The configuration that `crossroads` uses is normally stored in the file `/etc/crossroads.conf`. This location can be overruled using the command line flag `-c`.

This section explains the syntax of the configuration file, and what all settings do.

### 4.1 General language elements

This section describes the general elements of the `crossroads` configuration language.

#### 4.1.1 Empty lines and comments

Empty lines are of course allowed in the configuration. `Crossroads` recognizes three formats of comment:

- C-style, between `/*` and `*/`,
- C++-style, starting with `//` and ending with the end of the text line;
- Shell-style, starting with `#` and ending with the end of the text line.

Simply choose your favorite editor and use the comment that 'looks best'.<sup>3</sup>

#### 4.1.2 Keywords, numbers, identifiers, generic strings

In a configuration file, statements are identified by *keywords*, such as `service`, `verbosity`. These are reserved words.

Many keywords require an *identifier* as the argument. E.g, a `service` has a unique name, which must start with a letter or underscore, followed by zero or more letters, underscores, or digits. Therefore, in the statement `service myservice`, the keyword is `service` and the identifier is `myservice`.

Other keywords require a numeric argument. `Crossroads` knows only non-negative integer numbers, as in `port 8000`. Here, `port` is the keyword and `8000` is the number.

---

<sup>3</sup>I favor C or C++ comment. My favorite editor *emacs* can be put in `cmode` and nicely highlight what's comment and what's not. And as a bonus it will auto-indent the configuration!

Yet other keywords require 'generic strings', such as hostname specifications or system commands. Such generic strings contain any characters (including white space) up to the terminating statement character `;`. If a string must contain a semicolon, then it must be enclosed in single or double quotes:

- `This is a string;` is a string that starts at `T` and ends with `g`
- `"This is a string";` is the same, the double quotes are not necessary
- `"This is ; a string";` has double quotes to protect the inner `;`

Finally, an argument can be a 'boolean' value. Crossroads knows `true`, `false`, `yes`, `no`, `on`, `off`. The keywords `true`, `yes` and `on` all mean the same and can be used interchangeably; as can the keywords `false`, `no` and `off`.

## 4.2 Service definitions

Service definitions are blocks in the configuration file that state what is for each service. A service definition starts with `service`, followed by a unique identifier, and by statements in `{` and `}`. For example:

```
// Definition of service 'www':
service www {
    ...
    ... // statements that define the
    ... // service named 'www'
    ...
}
```

The configuration file can contain many service blocks, as long as the identifying names differ. The following list shows possible statements. Each statement must end with a semicolon, except for the `backend` statement, which has its own block (more on this later).

**The `type` statement** defines how crossroads handles the stated service. There are currently two types: `any` and `http`. The type `any` means that crossroads doesn't interpret the contents of a TCP stream, but only distributes streams over back ends. The type `http` means that crossroads has to analyze what's in the messages, does magical HTTP header tricks, and so on – all to ensure that multiple connections are treated as one session, or that the back end is notified of the client's IP address.

Unless you really need such special features, use the type `any` (the default), even for HTTP protocols.

- Syntax: `type typespecifier ;`
- Where *typespecifier* is `any` or `http`
- Default: `any`

**The `port` statement** defines to which TCP port a service 'listens'. E.g. `port 8000` says that this service will accept connections on port 8000.

- Syntax: `port number ;`
- There is no default. This is a required setting.

**The `bindto` statement** is used in situations where crossroads should only listen to the stated port at a given IP address. E.g., `bindto 127.0.0.1` causes crossroads to 'bind' the service only to the local IP address. Network connections from other hosts won't be serviced. By default, crossroads binds a service to all presently active IP addresses at the invoking host.

- Syntax: `bindto ip-address ;`



- where *ip-address* is a numeric IP address, such as 192.168.1.45, or the keyword *any*
- Default: *any*

**Verbosity statements** come in two forms: `verbosity on` or `verbosity off`. When 'on', log messages to `/var/log/messages` are generated that show what's going on.<sup>4</sup> The keyword `verbose` is an alias for `verbosity`.

- Syntax: `verbosity setting ;`
- Or: `verbose setting ;`
- where *setting* is `true`, `yes` or `on` to turn verbosity on; or `false`, `no`, `off` to turn it off.
- Default: `off`.

**The dispatch mode** controls how crossroads selects a back end from a list of active back ends. The below text shows the bare syntax. See section 5.1 for a textual explanation.

The syntax is:

- `dispatchmode roundrobin`: Simply the 'next in line' is chosen. E.g, when 3 back ends are active, then the usage series is 1, 2, 3, 1, 2, 3, and so on. Roundrobin dispatching is the default method, when no `dispatchmode` statement occurs.
- `dispatchmode random`: Random selection. Probably only for stress testing, though when used with weights (see below) it is a good distributor of new connections too.
- `dispatchmode bysize [ over connections ]`: The next back end is the one that has transferred the least number of bytes. This selection mechanism assumes that the more bytes, the heavier the load. The modifier *over connections* is optional. (The square brackets shown above are not part of the statement but indicate optionality.) When given, the load is computed as an average of the last stated number of connections. When this modifier is absent, then the load is computed over all connections since startup.
- `dispatchmode byduration [ over connections ]`: The next back end is the one that served connections for the shortest time. This mechanism assumes that the longer the connection, the heavier the load.
- `dispatchmode byconnections`: The next back end is the one with the least active connections. This mechanism assumes that each connection to a back end represents load. It is usable for e.g. database connections.
- `dispatchmode byorder`: The first back end is selected every time, unless it's unavailable. In that case the second is taken, and so on.

The selection algorithm is only used when clients are serviced that aren't part of a sticky HTTP session. This is the case during:

- all client requests of a service type `any`;
- new sessions of a service type `http`.

When type `http` is in effect and a session is underway, then the previously used back end is always selected – regardless of dispatching mode.

Your 'right' dispatch mode will depend on the type of service. Given the fact that crossroads doesn't know (and doesn't care) how to estimate load from a network traffic stream, you have to choose an appropriate dispatch mode to optimize load balancing. In most cases, `roundrobin` or `byconnections` will do the job just fine.

**A reviving interval definition** is needed when crossroads determines that a back end is temporarily unavailable. This will happen when:

---

<sup>4</sup>Actually, the messages go to `syslog(3)`, using facility `LOG_DAEMON` and priority `LOG_INFO`. In most (Linux) cases this will mean: output to `/var/log/messages`. On Mac OSX the messages go to `/var/log/system.log`.

- The back end cannot be reached (network connection fails);
- The network connection to the back end suddenly dies.

An example of the definition is `revivinginterval 10`. When this reviving interval is given, crossroads will check each 10 seconds whether unavailable back ends have woken up yet. A back end is considered awake when a network connection to that back end can successfully be established.

- Syntax: `revivinginterval number ;`
- Default: 0, meaning no revivals will occur.

**The maximum number of connections** is specified using `maxconnections`. There is one argument; the number of concurrent established connections that may be active within one service.

‘Throttling’ the number of connections is a way of preventing Denial of Service (DOS) attacks. Without a limit, numerous network connections may spawn so many server instances, that the service ultimately breaks down and becomes unavailable.

- Syntax: `maxconnections number ;`
- Default: 0, meaning that all client connections will be accepted.

**The TCP back log size** is a number that controls how many ‘waiting’ network connections may be queued, before a client simply cannot connect. The syntax is e.g. `backlog 5` to cause crossroads to have 5 waiting connections for 1 active connection. The backlog queue shouldn’t be too high, or clients will experience timeouts before they can actually connect. The queue shouldn’t be too small either, because clients would be simply rejected. Your mileage may vary.

- Syntax: `backlog number ;`
- Default: zero, which takes the operating system’s default value for socket back log size.

**Reporting based: the shared memory key.** Different crossroad invocations must ‘know’ of each others activity. E.g. `crossroad status` must be able to get to the actual state information of all running services. This is internally implemented through shared memory, which is reserved using a key.

Normally crossroads will supply a shared memory key, based on the service port and bitwise or-ed with a magic number. In situations where this conflicts with existing keys (of other programs, having their own keys), you may supply a chosen value.

The syntax is e.g. `shmkey 123456`. The actual key value doesn’t matter much, as long as it’s unique and as long as each invocation of crossroads uses it.

- Syntax: `shmkey number ;`
- Default: 0, which means that crossroads will ‘guess’ its own key, based on TCP port and a magic number.

**Connection timeouts:** Sometimes, clients simply won’t close a network connection which leads to unnecessary resource usage. To avoid this, one might state e.g. `connectiontimeout 300`. This instructs crossroads to consider a connection where nothing has happened for 300 seconds as ‘finished’. Crossroads will terminate the connection when this timeout is exceeded.

- Syntax: `connectiontimeout number ;`
- Default: 0, meaning that crossroads will not try to determine timeouts.

### 4.3 Backend definitions

Inside the service definitions as are described in the previous section, *backend definitions* must also occur. Backend definitions are started by the keyword `backend`, followed by an identifier (the back end name), and statements inside `{` and `}`:

```
service myservice {  
    ...  
    ... // statements that define the  
    ... // service named 'myservice'  
    ...  
  
    backend mybackend {  
        ...  
        ... // statements that define the  
        ... // backend named 'mybackend'  
        ...  
    }  
}
```

Each service definition must have at least one backend definition. There may be more (and probably will, if you want balancing and fail over) as long as the backend names differ. The statements in the backend definition blocks are described in the following sections.

#### 4.3.1 General Backend Directives

The following directives are used in all types of services (`any` or `http`). HTTP-specific directives are shown in section 4.3.2.

**Server:** Each back end must be identified by the network name (server name) where it is located. For example: `server 10.1.1.23`, or `server web.mydomain.org`. A TCP port specifier can follow the server name, as in `server web.mydomain.org:80`.

- Syntax: `server servername ;`
- Or: `server servername:port ;`
- There is no default. This is a required setting.

**Port:** When the `server` specifier doesn't include a TCP port, then this statement is used to define the port at which the back end expects its traffic. There is one argument, the (numeric) port number.

- Syntax: `port number ;`
- There is no default. The port must be defined either in the `server` setting or using the `port` specifier.

**Verbosity:** Similar to service specifications, a backend can have its own verbosity (`on` or `off`). When `on`, traffic to and from this back end is reported.

- Syntax: `verbosity setting ;`
- Or: `verbose setting ;`
- where *setting* is `true`, `yes` or `on` to turn verbosity on; or `false`, `no`, `off` to turn it off.
- Default: `off`.

**Maxconnections:** This setting states how many concurrent connections a back end connection may accept. Note that there is also a `maxconnections` statement for the overall service description.

The difference is that a `maxconnections` statement at the level of a service description avoids too many hits from the outside (DOS prevention). A `maxconnections` statement at the level of a back end description makes sure that this particular back end doesn't get overloaded.

- Syntax: `maxconnections number ;`
- where *number* is the maximum number of concurrent client connections.

- Default: 0, meaning that there is no limit.

**Weight:** To influence how backends are selected by size or by duration, a backend can specify its 'weight' in the process. The higher the weight, the less likely a back end will be chosen. The default is 1.

The weighing mechanism only applies to the dispatch modes `random`, `byconnections`, `bysize` and `byduration`. The weight is in fact a penalty factor. E.g., if backend A has weight 2 and backend B has weight 1, then backend B will be selected all the time, until its usage parameter is twice as large as the parameter of A. Think of it as a 'sluggishness' statement.

- Syntax: `weight number ;`
- Default: 1

**Decay:** To make sure that a 'spike' of activity doesn't influence the perceived load of a back end forever, you may specify a certain decay. E.g, the statement `decay 10` makes sure that the load that crossroads computes for this back end (be it in seconds or in bytes) is decreased by 10% each time that **an other** back end is hit. Decays are not applied to the count of concurrent connections.

This means that when a given back end is hit, then its usage data of the transferred bytes and the connection duration are updated using the actual number of bytes and actual duration. However, when a different back end is hit, then the usage data are decreased by the specified decay.

- Syntax: `decay number ;`
- where *number* is a percentage that decreases the back data when other back ends are hit
- Default: 0, meaning that no decay is applied to usage statistics.

**Event triggers:** As special 'hooks' for actions, two triggers are available: `onfailure` and `onsuccess`. The argument to the triggers is a system command that is executed when a connection with the back end either fails or succeeds.

- Syntax: `onfailure commandline ;` and `onsuccess commandline ;`
- There is no default.

**Debugging and Performance aids:** Two directives are available to log network traffic to files. They are `trafficlog` and `throughputlog`.

- Syntax: `trafficlog filename ;`
- And: `throughputlog filename ;`
- Default: none

The `trafficlog` statement causes all traffic to be logged in hexadecimal format. Each line is prefixed by B or C, depending on whether the information was received from the back end or from the client.

The `throughputlog` statement writes shorthand transmissions to its log, accompanied by timings.

### 4.3.2 HTTP-related Backend Directives

The following directives are specific for HTTP-type services; i.e., services with a specification type `http`.

It is inevitable that when Crossroads handles services of type `http`, more processing is necessary. Crossroads has to unpack the TCP payload in order to do its header magic; which leads to performance impact.

**Session Stickiness:** The directive `stickycookie` *value* causes Crossroads to unpack clients' requests, to check for *value* in the cookies. When found, the message is routed to the back end having the appropriate `stickycookie` directive.

E.g., consider the following configuration:

```
service ... {
    ...
    backend one {
        ...
        stickycookie "BalancerID=first";
    }
    backend two {
        ...
        stickycookie "BalancerID=second";
    }
}
```

When clients' messages contain cookies named `BalancerID` with the value `first`, then such messages are routed to backend one. When the value is `second` then they are routed to the backend two.

There are basically two ways to provide such cookies to a browser. First, a back end can insert such a cookie into the HTTP response. E.g., the webserver of back end one might insert a cookie named `BalancerID`, having value `first`. Second, Crossroads can insert such cookies using a carefully crafted directive `addclientheader`. See below.

**Header modification:** Crossroads understands the following header modification directives: `addclientheader`, `appendclientheader`, `setclientheader`, `addserverheader`, `appendserverheader`, `setserverheader`.

The directive names always consist of *ActionDestinationheader*, where:

- The action is `add`, `append` or `insert`.
  - Action `add` adds a header, even when headers with the same name already are present in an HTTP message. Adding headers is useful for e.g. `Set-Cookie` headers; a message may contain several of such headers.
  - Action `append` adds a header if it isn't present yet in an HTTP message. If such a header is already present, then the value is appended to the pre-existing header. This is useful for e.g. `Via` headers. Imagine an HTTP message with a header `Via: someproxy`. Then the directive `appendclientheader "Via: crossroads"` will rewrite the header to `Via: someproxy; crossroads`.
  - Action `set` overwrites headers with the same name; or adds a new header if no pre-existing is found. This is useful for e.g. `Host` headers.
- The destination is one of `client` or `server`. When the destination is `server`, then Crossroads will apply such directives to HTTP messages that originate from the browser and are being forwarded to back ends. When the destination is `client`, then Crossroads will apply such directives to backend responses that are shuttled to the browser.

The syntax of the directives is e.g. `addclientheader "X-Processed-By: Crossroads";`. The directives expect one argument; a string, consisting of a header name, a colon, and a header value. The directive ends with a semicolon.

The header value may contain one of the following formatting directives:

- `%r` is expanded to the real IP address of a client;
- `%t` is expanded to a timestamp of the local time;
- `%T` is expanded to a timestamp of Greenwich Mean Time;

- %v is expanded to the Crossroads version;
- %x (where x is any other character) is expanded to x. E.g., %% is a literal % sign.

**Common Uses** The following examples show common uses of header modifications.

**Enforcing session stickiness:** By combining `stickycookie` and `addclientheader`, HTTP session stickiness is enforced. Consider the following configuration:

```
service ... {
    ...
    backend one {
        ...
        addclientheader "Set-Cookie: BalancerID=first; path="/";
        stickycookie "BalancerID=first";
    }
    backend two {
        ...
        addclientheader "Set-Cookie: BalancerID=second; path="/";
        stickycookie "BalancerID=second";
    }
}
```

The first request of an HTTP session is balanced to either backend one or two. The server response is enriched using `addclientheader` with an appropriate cookie. A subsequent request from the same browser now has that cookie in place; and is therefore sent to the same back end where the its predecessors went.

**Hiding the server software version:** Many servers (e.g. Apache) advertize their version, as in `Server: Apache 1.27`. This potentially provides information to attackers. The following configuration hides such information:

```
service ... {
    ...
    backend one {
        ...
        setclientheader "Server: WWW-Server";
    }
}
```

**Informing the server of the clients' IP address:** Since Crossroads sits 'in the middle' between a client and a back end, the back end perceives Crossroads as its client. The following sends the true clients' IP address to the server, in a header `X-Real-IP`:

```
service ... {
    ...
    backend one {
        ...
        setserverheader "X-Real-IP: %r";
    }
}
```

**Keep-Alive Downgrading:** The directives `setclientheader` and `setserverheader` also play a key role in downgrading Keep-Alive connections to 'single-shot'. E.g., the following configuration makes sure that no Keep-Alive connections occur.

```
service ... {
    ...
    backend one {
        ...
        setserverheader "Connection: close";
        setclientheader "Connection: close";
    }
}
```

## 5 Tips, Tricks and Remarks

The following sections elaborate on the directives as described in section 4 to illustrate how crossroads works and to help you achieve the "optimal" balancing configuration.

### 5.1 How back ends are selected in load balancing

In order to tune your load balancing, you'll need to understand how crossroads computes usage, how weighing works, and so on. In this section we'll focus on the dispatching modes `bysize`, `byduration` and `byconnections` only. The other dispatching types are self-explanatory.

#### 5.1.1 Bysize, byduration or byconnections?

As stated before, crossroads doesn't know 'what a service does' and how to judge whether a given back end is very busy or not. You must therefore give the right hints:

- In general, a service which is CPU bound, will be more busy when it takes longer to process a request. The dispatch mode `byduration` is appropriate here.
- In contrast, a service which is filesystem bound, will be more busy when more data are transferred. The dispatch mode `bysize` is appropriate.
- The dispatch mode `byduration` can also be used when network latency is an issue. E.g., if your balancer has back ends that are geographically distributed, then `byduration` would be a good way to select best available back ends.
- Furthermore it is noteworthy that `dispatchmode byduration` is not usable for interactive processes such as SSH logins. Idle time of a login adds to the duration, while causing (almost) no load. Mode `byduration` should only be used for automated processes that don't wait for user interaction (e.g., SOAP calls and other HTTP requests).
- As a last remark, the dispatching mode `byconnections` can be used if you don't have other clues for load estimations.

E.g., consider a database connection. What's heavier on the back end, time-consuming connections, or connections where loads of bytes are transferred? Well, that depends. A tough `select` query that joins multiple tables can be very heavy on the back end, though the response set can be quite small - and hence the number of transferred bytes. That would suggest dispatching by duration. However, `byduration` balancing doesn't represent the true world, when interactive connections can occur where users have an idle TCP connection to the database: this consumes time, but no bytes (see the SSH login example above). In this case, the dispatch mode `byconnections` may be your best bet.

#### 5.1.2 Averaging size and duration

The configuration statement `dispatchmode bysize` or `byduration` allows an optional modifier `over number`, where the stated number represents a connection count. When this modifier is present, then crossroads will use a moving average over the last *n* connections to compute duration and size figures.

In the real world you'll always want this modifier. E.g., consider two back ends that are running for years now, and one of them is suddenly overloaded and very busy (it experiences a 'spike' in activity). When the `over` modifier is absent, then the sudden load will hardly show up in the usage figures - it will flatten out due to the large usage figures already stored in the years of service.

In contrast, when e.g. `over 3` is in effect, then a sudden load does show up - because it highly contributes to the average of three connections.

### 5.1.3 Specifying decays

Decays are also only relevant when crossroads computes the 'next best back end' by size (bytes) or duration (seconds). E.g., imagine two back ends A and B, both averaged over say 3 connections.

Now when back end A is suddenly hit by a spike, its average would go up accordingly. But the back end would never again be used, unless B also received a similar spike, because A's 'usage data' over its last three connections would forever be larger than B's data.

For that reason, you should in real situations probably always specify a decay, so that the backend selection algorithm recovers from spikes. Note that the usage data of the back end where a decay is specified, decay when **other** back ends are hit. The decay parameter is like specifying how fast your body regenerates when someone else does the work.

The below configuration illustrates this:

```
/* Definition of the service */
service soap {
    /* Local TCP port */
    port 8080;

    /* We'll select back ends by the processing
     * duration
     */
    dispatchmode byduration over 3;

    /* First back end: */
    backend A {
        /* Back end IP address and port */
        server 10.1.1.1:8080;

        /* When this back end is NOT hit because
         * the other one was less busy, then the
         * usage parameters decay 10% per connection
         */
        decay 10;
    }

    /* Second back end: */
    backend B {
        server 10.1.1.2:8080;
        decay 10;
    }
}
```

### 5.1.4 Adjusting the weights

The back end modifier `weight` is useful in situations where your back ends differ in respect to performance. E.g., your back ends may be geographically distributed, and you know that a given back end is difficult to reach and often experiences network lag.

Or you may have one primary back end, a system with a fast CPU and enough memory, and a small fall-back back end, with a slow CPU and short on memory. In that case you know in advance that the second back end should be used only rarely. Most requests should go to the big server, up to a certain load.

In such cases you will know in advance that the best performing back ends should be selected the most often. Here's where the `weight` statement comes in: you can simply increase the weight of the back ends with the least performance, so that they are selected less frequently.



E.g., consider the following configuration:

```
service soap {
  port 8080;
  dispatchmode byduration over 3;
  backend A {
    server 10.1.1.1:8080;
    decay 20;
  }
  backend B {
    server 10.1.1.2:8080;
    weight 2;
    decay 10;
  }
  backend C {
    server 10.1.1.3:8080;
    weight 4;
    decay 5;
  }
}
```

This will cause crossroads to select back ends by the processing time, averaging over the last three connections. However, backend B will kick in only when its usage is half of the usage of A (back end B is probably only half as fast as A). Backend C will kick in only when its usage is a quarter of the usage of A, which is half of the usage of B (back end C is probably very weak, and just a fall-back system incase both A and B crash). Note also that A's usage data decay much faster than B's and C's: we're assuming that this big server recovers quicker than its smaller siblings.

### 5.1.5 Throttling the number of concurrent connections

If you suspect that your service may occasionally receive 'spikes' of activity<sup>5</sup>, then it might be a good idea to protect your service by specifying a maximum number of concurrent connections. This protection can be specified on two levels:

**On the service level** a statement like `maxconnections 100;` states that the service as a whole will never service more than 100 concurrent connections. This means that all your back ends and the crossroads balancer itself will be protected from being overloaded.

**On the back end level** a statement like `maxconnections 10;` states that this particular back end will never have more than 10 concurrent connections; regardless of the overall setting on the service level. This means that this particular back end will be protected from being overloaded (regardless of what other back ends may experience).

The `maxconnections` statement, combined with a back end selection algorithm, allows very fine granularity. The `maxconnections` statement on the back end level is like a hand brake: even when you specify a back end algorithm that would protect a given back end from being used too much, a situation may occur where that back end is about to be hit. A `maxconnections` statement on the level of that back may then protect it.

## 5.2 HTTP Session Stickiness

This section focuses on HTTP session stickiness. This term refers to the ability of a balancer to route a conversation between browser and a backend farm always to the same back end. In other words: once a back end is selected by the balancer, it will remain the back end of choice, even for subsequent connections.

---

<sup>5</sup>which you should always assume

### 5.2.1 Don't use stickiness!

The rule of thumb as far as the balancer is concerned, is: **Do not use HTTP session stickiness unless you really have to**. Enabling session stickiness hampers failover, balancing and performance:

- Failover is hampered because during the session, the balancer has to assign new connections to the same back end that was selected at the start of a session. If the back end suddenly goes 'down', then the session will most likely crash. (Actually, when a back end becomes unreachable in the middle of a session, Crossroads will assign a new back end to that session. This will most likely result in a malfunction of the underlying application.)
- Balancing is hampered because at the start of the session, the balancer has selected the next-best back end. But during the session, that back end may well become overloaded. The balancer however must continue to send the requests there.
- Performance is hampered because crossroads needs to 'unpack' messages as they are passed to and fro. That's because crossroads needs to check the HTTP headers in the messages for persistence cookies.

There is a number of measures that you can take to avoid using session stickiness. E.g., session data can be 'shared' between web back ends. PHP offers functionality to store session data in a database, so that all PHP applications have access to these data. Application servers such as Websphere can be configured to replicate session data between nodes.

### 5.2.2 But if you must..

However, if you **must** use session stickiness, then proceed as follows:

- At the level of a `service` description, set the type to `http`.
- At the level of each back end description, configure the `stickycookie` and a `addclientheader` directives.

Once crossroads sees that, it will examine each HTTP message that it shuttles between client and back end:

- If there is no persistence cookie in the HTTP headers of a client's request, then the message must be the first one and a new session should be established. Crossroads selects an appropriate back end, sends the message to that back end, catches the reply, and inserts a `Set-Cookie` directive.
- If there is a persistence cookie in the HTTP headers of a client's request, then the request is part of an already established session. Crossroads analyzes the cookie and forwards the request to the appropriate back end.

Below is a short example of a configuration.

```
service www {
    port 80;
    type stickyhttp;
    revivinginterval 15;
    dispatchmode byconnections;

    backend one {
        server 10.1.1.100:80;
        stickycookie XRID=100;
        addclientheader "Set-Cookie: XRID=100; Path=/";
    }

    backend two {
```

```
server 10.1.1.101:80;
stickycookie XRID=101;
addclientheader "Set-Cookie: XRID=101; Path="/";
}
}
```

Note how the cookie names and values in the directives `stickycookie` and `addclientheader` match. That is obviously a prerequisite for stickiness.

### 5.3 Passing the client's IP address

Since Crossroads just shuttles bytes to and fro, meta-information of network connections is lost. As far as the back ends are concerned, their connections originate at the Crossroads junction. For example, standard Apache access logs will show the IP address of Crossroads.

In order to compensate for this, Crossroads can insert a special header in HTTP connections, to inform the back end of the original client's IP address. In order to enable this, the Crossroads configuration must state the following:

- The service type must be `http`, and not `any`;
- In the back end definition, the following statement must occur:  

```
addserverheader "X-Real-IP: %r";
```

You are of course free to choose the header name; the here used `X-Real-IP` is a common name for this purpose.

After this, HTTP traffic that arrives at the back ends has a new header: `X-Real-IP`, holding the client's IP address. **Note that** once the type is set to `http`, Crossroads' performance will be hampered – all passing messages will have to be unpacked and analyzed.

#### 5.3.1 Sample Crossroads configuration

The below sample configuration shows two HTTP back ends that receive the client's IP address:

```
service www {
    port 80;
    type http;
    revivinginterval 5;
    dispatchmode roundrobin;

    backend one {
        server 10.1.1.100:80;
        addserverheader "X-Real-IP: %r";
    }

    backend two {
        server 10.1.1.200:80;
        addserverheader "X-Real-IP: %r";
    }
}
```

#### 5.3.2 Sample Apache configuration

The method by which each back end analyzes the header `X-Real-IP` will obviously be different per server implementations. However, a common method with the Apache webserver is to log the client's IP address into the access log.

Often this is accomplished using the log format `custom`, defined as follows:

```
LogFormat "%h %l %u %t %D \"%r\" %>s %b" common
CustomLog logs/access_log common
```

The first line defines the format `common`, with the remote host specified by `%h`. The second line sends access information to a log file `logs/access_log`, using the previously defined format `common`.

Furtunately, Apache's `LogFormat` allows one to log contents of headers. By replacing the `%h` with `%{X-Real-IP}i`, the desired information is sent to the log. Therefore, normally you can simply redefine the `common` format to

```
LogFormat "%{X-Real-IP}i %l %u %t %D \"%r\" %>s %b" common
```

## 5.4 Debugging network traffic

Incase the traffic between client and backend must be debugged, the statement `trafficlog filename` can be issued. This causes the traffic to be dumped in hexadecimal format to the stated filename.

Traffic sent by the client is prefixed by a **C**, traffic sent by the back end is prefixed by a **B**. Below is a sample traffic dump of a browser trying to get a HTML page. The server replies that the page was not modified.

```
C 0000 47 45 54 20 68 74 74 70 3a 2f 2f 77 77 77 2e 63 GET http://www.c
C 0010 73 2e 68 65 6c 73 69 6e 6b 69 2e 66 69 2f 6c 69 s.helsinki.fi/li
C 0020 6e 75 78 2f 6c 69 6e 75 78 2d 6b 65 72 6e 65 6c nux/linux-kernel
C 0030 2f 32 30 30 31 2d 34 37 2f 30 34 31 37 2e 68 74 /2001-47/0417.ht
C 0040 6d 6c 20 48 54 54 50 2f 31 2e 31 0d 0a 43 6f 6e ml HTTP/1.1..Con
C 0050 6e 65 63 74 69 6f 6e 3a 20 63 6c 6f 73 65 0d 0a nexion: close..
.
. etcetera
.
B 0000 48 54 54 50 2f 31 2e 30 20 33 30 34 20 4e 6f 74 HTTP/1.0 304 Not
B 0010 20 4d 6f 64 69 66 69 65 64 0d 0a 44 61 74 65 3a Modified..Date:
B 0020 20 54 75 65 2c 20 31 32 20 4a 75 6c 20 32 30 30 Tue, 12 Jul 200
B 0030 35 20 30 39 3a 34 39 3a 34 37 20 47 4d 54 0d 0a 5 09:49:47 GMT..
B 0040 43 6f 6e 74 65 6e 74 2d 54 79 70 65 3a 20 74 65 Content-Type: te
B 0050 78 74 2f 68 74 6d 6c 3b 20 63 68 61 72 73 65 74 xt/html; charset
.
. etcetera
.
```

Turning on traffic dumps will *significantly* slow down crossroads.

Besides `trafficlog`, there is also a directive `throughputlog`. This directive also takes one argument, a filename. The file is appended, and the following information is logged:

- The process ID of the crossroads image that serves the TCP connection;
- The time of the request, in seconds and microseconds since start of the run;
- A **C** when the request originated at the client, or **B** when the request originated at the back end;
- The first 100 bytes of the request.

As an example, consider the following (the lines are shortened for brevity and prefixed by line numbers for clarity):

```
1 0000594 0.000001 C GET http://public.e-tunity.com/index.html...
2 0000594 0.173713 B HTTP/1.0 200 OK..Date: Fri, 18 Nov 2005 0...
3 0000594 0.278125 B width="100" bgcolor="#e0e0e0" valign="to...
```

```
4 0000595 0.000001 C GET http://public.e-tunity.com/css/style/...
5 0000594 0.944339 B /a></td>.. </tr>.</table>.</td><td class...
6 0000594 0.946356 B smallboxdown1">Download</td>.. <td class...
7 0000594 0.961102 B td><td class="smallboxodd" valign="top"><...
8 0000595 0.698215 B HTTP/1.0 304 Not Modified..Date: Fri, 18 ...
```

This tells us that:

- Line 1: PID 594 served a request that originated at the client. The corresponding time is (almost) 0 seconds, so this is really the start of the run.
- Line 2: A back end replied 0.17 seconds later, and 0.28 seconds later, it was still replying (this is the third line, again a **B**-type transmission).
- Line 4: PID 595 served a request that originated at the client. Again, the corresponding time is (almost) 0 seconds, since this is the first conversation part of this connection.
- Lines 5 to 7: This is the continuation of line 2. Line 7 is the last line of the **B** series (not visible from the example, but trust me, it is), so that we may conclude that it took the back end 0.96 seconds to serve the file `index.html` requested in line 1.
- Line 8: This is the answer to the client's request of line 4 (you can tell by the process ID number). So the back end took 0.68 seconds to confirm that the stylesheet requested in line 4 wasn't modified.

It is also worth while remembering that the start time of a **C** request is the time that crossroads sees the activity. Any latency between the true client and crossroads is obviously not included. This is illustrated by the below simple ASCII art:

```
client ---->---->---->--->*crossroads ===>====>====>
                                     \
                                     back end
                                     /
client ----<----<----<----< crossroads ===<====<====<
```

This simple picture shows a typical HTTP request that originates at a client, travels to crossroads, and is relayed via the back end. The **C** entry in a throughput log is the time when crossroads sees the request, indicated by an asterisk. The **B** entries are the times that it takes the back end to answer, indicated by `===` style lines. Therefore, the true roundtrip time will be longer than the number of seconds that are logged in the throughput log: the latency between client and crossroads isn't included in that measurement.

Summarizing, the throughput times of a client-back end connection can be analyzed using the directive `throughputlog`. In a real-world analysis, you'd probably want to write up a script to analyze the output and to compute round trip times. Such scripts are not (yet) included in Crossroads.

## 5.5 Configuration examples

As a general hint, use `crossroads sampleconf` to view the most up-to-date examples of configurations. The description below shows a few examples too.

### 5.5.1 A load balancer for three webserver back ends

The following configuration example binds crossroads to port 80 of the current server, and distributes the load over three back ends. This configuration shows most of the possible settings.

```
service www {
    /* We don't need session stickyness. */
    type any;

    /* Port on which we'll listen in this service: required. */
    port 8000;

    /* What IP address should this service listen? Default is 'any'.
     * Alternatively you can state an explicit IP address, such as
     * 127.0.0.1; that would bind the service only to 'localhost'. */
    bindto any;

    /* Verbose reporting or not. Default is off. */
    verbosity on;

    /* Dispatching mode, or: How to select a back end for an incoming
     * request. Possible values:
     * roundrobin: just the next back end in line
     * random: like roundrobin, but at random to make things more
     *          confusing. Probably only good for testing.
     * bysize: The backend that transferred the least nr of bytes
     *          is the next in line. As a modifier you can say e.g.
     *          bysize over 10, meaning that the 10 last connections will
     *          be used to compute the transfer size, instead of all
     *          transfers.
     * byduration: The backend that was active for the shortest time
     *              is the next in line. As a modifier you can say e.g.
     *              byduration of 10 to compute over the last 10 connections.
     * byconnections: The back end with the least active connections
     *                 is the next in line.
     * byorder: The first available back end is always taken.
     */
    dispatchmode byduration over 5;

    /* Interval at which we'll check whether a temporarily unavailable
     * backend has woken up.
     */
    revivinginterval 5;

    /* TCP backlog of connections. Default is 0 (no backlog, one
     * connection may be active).
     */
    backlog 5;

    /* For status reporting: a shared memory key. Default is the same
     * as the port number, OR-ed by a magic number.
     */
    shmkey 8000;

    /* This controls when crossroads should consider a connection as
     * finished even when the TCP sockets weren't closed. This is to
     * avoid hanging connections that don't do anything. NOTE THAT when
     * crossroads cuts off a connection due to timeout exceed, this is
     * not marked as a failure, but as a success. Default is 0: no timeout.
```

```
*/
connectiontimeout 300;

/* The max number of allowed client connections. When present, connections
 * won't be accepted if the max is about to be exceeded. When
 * absent, all connections will be accepted, which might be misused
 * for a DOS attack.
 */
maxconnections 300;

/* Now let's define a couple of back ends. Number 1: */
backend www_backend_1 {
    /* The server and its port, the minimum configuration. */
    server httpserver1;
    port 9010;
    /* The 'decay' of usage data of this back end. Only relevant
     * when the whole service has 'dispatchmode bysize' or
     * 'byduration'. The number is a percentage by which the usage
     * parameter is decreased upon each connection of an other back
     * end.
     */
    decay 10;

    /* To see what's happening in /var/log/messages: */
    verbosity on;
}

/* The second one: */
backend www_backend_2 {
    /* Server and port */
    server httpserver2;
    port 9011;

    /* Verbosity of reporting when this back end is active */
    verbosity on;

    /* Decay */
    decay 10;

    /* This back end is twice as weak as the first one */
    weight 2;

    /* Event triggers for system commands upon succesful activation
     * and upon failure.
     */
    onsuccess echo 'success on backend 2' | mail root;
    onfailure echo 'failure on backend 2' | mail root;
}

/* And yet another one.. this time we will dump the traffic
 * to a trace file. Furthermore we don't want more than 10 concurrent
 * connections here. Note that there's also a total maxconnections for the
 * whole service.
 */
```

```
backend www_backend_3 {
    server httpserver3;
    verbosity on;
    port 9000;
    verbosity on;
    decay 10;
    trafficlog /tmp/backend.3.log;
    maxconnections 10;
}
}
```

### 5.5.2 An HTTP forwarder when travelling

As another example, here's my `crossroads.conf` that I use on my Unix laptop. The problem that I face is that I need many HTTP proxy configurations (at home, at customers' sites and so on) but I'm too lazy to reconfigure browsers all the time.

Here's how it used to be before `crossroads`:

- At home, I would surf through a squid proxy on my local machine. The browser proxy setting is then `http://localhost:3128`.
- Sometimes I start up an SSH tunnel to our offices. The tunnel has a local port 3129, and connects to a squid proxy on our e-tunity server. Hence, the browser proxy is then `http://localhost:3129`.
- At a customer's location I need the proxy `http://10.120.34.113:8080`, because they have configured it so.
- And in yet other instances, I use a HTTP diagnostic tool Charles<sup>6</sup> that sits between browser and website and shows me what's happening. I run `charles` on my own machine and it listens to port 8888, behaving like a proxy. The browser configuration for the proxy is then `http://localhost:8888`.

Here's how it works with a `crossroads` configuration:

- I have configured my browsers to use `http://localhost:8080` as the proxy. For all situations.
- I use the following `crossroads` configuration, and let `crossroads` figure out which proxy backend works, and which doesn't. Note two particularities:
  - The statement `dispatchmode byorder`. This makes sure that once `crossroads` determines which backend works, it will stick to it. This usage of `crossroads` doesn't need to balance over more than one back end.
  - The statement `bindto 127.0.0.1` makes sure that requests from other interfaces than loopback won't get serviced.

```
service HttpProxy {
    port 8080;
    bindto 127.0.0.1;
    verbosity on;
    dispatchmode byorder;
    revivinginterval 15;

    backend Charles {
        server localhost:8888;
        verbosity on;
    }
}
```

---

<sup>6</sup><http://www.xk72.com/charles>



```
    }

    backend CustomerProxy {
        server 10.120.34.113:8080;
        verbosity on;
    }

    backend SshTunnel {
        server localhost:3129;
    }

    backend LocalSquid {
        server localhost:3128;
    }
}
```

As a final note, the commandline argument `tell` can be used to influence crossroad's own detection mechanism of back end availability detection. E.g., if in the above example the back ends `SshTunnel` and `LocalSquid` are both active, then crossroads `tell httpproxy sshtunnel down` will 'take down' the back end `SshTunnel` – and will automatically cause crossroads to switch to `LocalSquid`.

### 5.5.3 SSH login with enforced idle logout

The following example shows how crossroads 'throttles' SSH logins. Connections are accepted on port 22 (the normal SSH port) and forwarded to the actual SSH daemon which is running on port 2222.

Note the usage of the `connectiontimeout` directive. This makes sure that users are logged out after 10 minutes of inactivity. Note also the `maxconnections` setting, this makes sure that no more than 10 concurrent logins occur.

```
service Ssh {
    port 22;
    backlog 5;
    maxconnections 10;
    connectiontimeout 600;
    backend TrueSshDaemon {
        server localhost:2222;
    }
}
```

## 6 Benchmarking

This section shows how crossroads affects the transmitting of HTML data when used as an intermediate 'station' through which all data travels.

### 6.1 Benchmark 1: Accessing a proxy via crossroads or directly

The benchmark was run on a system where the following was varied:

1. A website was recursively spidered through a local squid proxy. The spidering was repeated 10 times, the total was recorded.
2. Crossroads was placed in front of the squid proxy, and the website was again recursively spidered. Again, the spidering was repeated 10 times and the total was recorded.

The crossroads configuration of the second alternative is shown below:

```
service HttpProxy {
    port 8080;
    verbosity on;
    backend LocalSquid {
        server 127.0.0.1;
        port 3128;
        verbosity on;
    }
}
```

### 6.1.1 Results

The results of this test are that crossroads causes a negligible delay, if it is statistically relevant at all. Without crossroads, the timing results are:

```
real 0m8.146s
user 0m0.130s
sys 0m0.253s
```

When using crossroads as a middle station, the results are:

```
real 0m9.481s
user 0m0.141s
sys 0m0.230s
```

### 6.1.2 Discussion

The above shown results are quite favorable to crossroads. However, one should know that situations will exist where crossroads leans towards the 'worst case' scenario, causing up to 50% delay.

E.g., imagine a test where a `wget` command retrieves a HTML document from an Apache server on `localhost`. Now we have (almost) no overhead due to network throttling, host-name lookups and so on. When this test would be run either with or without crossroads in between, then theoretically, crossroads would cause a much larger delay, because it has to read from the server, and then write the same information to `wget`. Each read/write occurs twice when crossroads sits in between.

This worst case scenario will however (fortunately) occur only very seldom in the real world:

- Normally network issues, such as the above mentioned host name lookups or throughput restrictions, will add significantly to the duration of a request. The 'twice as many' read/writes caused by crossroads are then relatively irrelevant.
- Normally a significant amount of time will be spent in a back end, due to processing (e.g., when calling a servlet on a back end). Again, this processing time will weigh much heavier than the multiple read/writes.

## 6.2 Benchmark 2: Crossroads versus Linux Virtual Server (LVS)

LVS is a kernel-based balancer that acts like a masquerading firewall: TCP packets that arrive at the balancer are sent to one of the configured back ends. LVS has the advantage over crossroads that there is no stop-and-go in the transmission; in contrast, crossroads needs to send data via an internal buffer. Crossroads has the advantage that it offers instantaneous failover because it tries to contact the back end for upon each new TCP connection; in contrast, LVS isn't aware of downtime of back ends (unless one implements an external heartbeat). Also, crossroads offers more complex balancing than LVS.

### 6.2.1 Environment

On the balancer, LVS was run on port 80, its forwarding set up for two equally weighted backends, using `ipvsadm`:

```
ipvsadm -a -t 192.168.1.250:http -r 10.1.1.100:http -m -w 1
ipvsadm -a -t 192.168.1.250:http -r 10.1.1.101:http -m -w 1
```

Crossroads was run on port 81. The configuration file is shown below:

```
service http {
    port 81;
    dispatchmode roundrobin;
    revivinginterval 5;
    backend one {
        server 10.1.1.100;
        port 80;
    }
    backend two {
        server 10.1.1.101;
        port 80;
    }
}
```

### 6.2.2 Tests and results

In the first test, ports 80 and 81 on the balancer were 'bombed' with 50 concurrent clients, each requesting a small page 50 times. The following timings where measured:

- How long it takes to establish a connection;
- How long it takes to retrieve the page.

The results of this test were:

- On average, each client took 0.12 seconds to connect to LVS, and each page was retrieved in 0.14 seconds;
- On average, each client took 0.11 seconds to connect to crossroads, and each page was retrieved in 0.13 seconds.

In this setup there seems to be no difference between the performance of LVS and crossroads!

In a second test, the size of the retrieved page was varied from 2.000 to 2.000.000 bytes. This test was taken to see whether crossroads would show performance degradation when transferring larger amounts of data.

For each page size, 30 concurrent clients were started, that retrieved the page 50 times. Again, the connect times and processing times where recorded.

The results of the total time (connect time + retrieval time) are shown in the below table:

Bytes	LVS timing	Crossroads timing
2000	0.130741688	0.12739582
20000	0.490916224	0.50376901
200000	3.799440328	4.33125273
2000000	45.25090855	45.9600728

Again, the results show that crossroads performs just as effectively as LVS, even with large data chunks!

## 7 Compiling and Installing

### 7.1 Prerequisites

The creation of crossroads requires:

- Standard Unix tools, such as `sed`, `awk`, `Perl` (5.00 or better);
- A POSIX-compliant C compiler;
- Support for SYSV IPC, networking and so on.

Basically a Linux or Apple MacOSX box will do nicely. To compile and install crossroads, follow these steps.

### 7.2 Compiling and installing

- Obtain the source distribution. It can be found on <http://crossroads.e-tunity.com>. The distribution comes as an archive `crossroads-type.tar.gz`, where *type* is *stable* or *devel*.
- Unpack the archive in a sources directory using `tar xzf crossroads-X.YY.tar.gz`. The contents spill into a subdirectory `crossroads-X.YY/`.
- Change-dir into the directory.
- Next, edit `etc/Makefile.def` and verify that all compilation settings are to your likings. The settings are explained in the file. **Note that** the default distribution of `Makefile.def` is suited for Linux or Apple MacOSX systems. On other Unices, or on non-Unix systems, you must particularly pay attention to `SET_PROC_TITLE_BY...`. When in doubt, comment out all `SET_PROC_TITLE...` settings. Crossroads will work nevertheless, but it won't show nice titles in `ps` listings. Also there's a macro `EXTRA_LIBS` to add linkage flags (an example for a Solaris build is included).
- Now crossroads is ready for compilation. Do a `make local` followed by `make install`. The latter step may have to be done by the user `root` if the `BINDIR` setting of `etc/Makefile.def` points to a root-owned directory.
- The documentation doesn't install in this process. If you want to install the documentation, then proceed as follows:
  - Optionally, `cp doc/crossroads.html htmdirectory/`; where *htmdirectory* is the destination directory for your HTML manuals;
  - Optionally, `cp doc/crossroads.pdf pdfdirectory/`; where *pdfdirectory* is the destination directory for your PDF manuals;
  - Optionally, `cp doc/crossroads.man manualdirectory/crossroads.1`, where *manualdirectory* is e.g. `/usr/man/man1`, `/usr/share/man1`, `/usr/local/man/man1`, `/usr/local/share/man1`. Any possibility is valid, as long as *manualdirectory* is one of the directories where manual pages are stored;
  - If your manual page system supports compressed manual pages, then you can save some space with `gzip manualdirectory/crossroads.1`.

### 7.3 Configuring crossroads

Now that the binary is available on your system, you need to create a suitable `/etc/crossroads.conf`. Use this manual or the output of `crossroads samplconf` to get started.

Once you have the configuration ready, start crossroads with `crossroads start`. Test the availability of your services and back ends. Monitor how crossroads is doing with:

- In one terminal, run the script:

```
while [ 1 ] ; do
    tput clear
    crossroads status
    sleep 3
done
```

**Note** that depending on your system you might need `sleep 3s`, i.e., with an `s` appended.

- In another terminal, run:

```
while [ 1 ] ; do
    tput clear
    ps ax | grep crossroads | grep -v grep
    sleep 3
done
```

**Note** that depending on your system you might need `ps -ef` instead of `ps ax`.

- In yet another terminal, run `tail -f /var/log/messages` (supply the appropriate system log file if `/var/log/messages` doesn't work for you).

Now thoroughly test the availability of your back ends through crossroads. The status display will show an updated view of which back ends are selected and how busy they are. The process list will show which crossroads daemons are running. Finally, the tailing of `/var/log/messages` shows what's going on – especially if you have `verbosity true` statements in the configuration.

## 7.4 A boot script

Finally, you may want to create a boot-time startup script. The exact procedure depends on the used Unix flavor.

### 7.4.1 SysV Style Startup

On SysV style systems, there's a startup script directory `/etc/init.d` where bootscripts for all utilities are located. You may have the `chkconfig` utility to automate the task of inserting scripts into the boot sequence, but otherwise the steps will resemble the following.

- Create a script `crossroads` in `/etc/init.d` similar to the following:

```
#!/bin/sh
/usr/local/bin/crossroads -v $@
```

The stated directory `/usr/local/bin` must correspond with the installation path. The flag `-v` causes the startup to be more 'verbose'. However, once daemonized, the verbosity is controlled by the appropriate statements in the configuration.

- Determine your 'runlevel': usually 3 when your system is running in text-mode only, or 5 when you are using a graphical interface. If your runlevel is 3, then:

```
root> cd /etc/rc.d/rc3.d
root> ln -s /etc/init.d/crossroads S99crossroads
root> ln -s /etc/init.d/crossroads K99crossroads
```

This creates startup (`S*`) and stop (`K*`) links that will be run when the system enters or leaves a given runlevel.

If your runlevel is 5, then the right `cd` command is to `/etc/rc.d/rc5.d`. Alternatively, you can create the symlinks in both runlevel directories.

### 7.4.2 BSD Style Startup

On BSD style systems, daemons are booted directly from `/etc/rc` and related scripts. In case you have a file `/etc/rc.local`, edit it, and add the statement:

```
/usr/local/bin/crossroads start
```

If your BSD system lacks `/etc/rc.local`, then you may need to start Crossroads from `/etc/rc`. Your mileage may vary.