

# BIRD Programmer's Documentation

---

Ondrej Filip <*feela@network.cz*>, Pavel Machek <*pavel@ucw.cz*>, Martin Mares <*mj@ucw.cz*>, Ondrej Zajicek <*santiago@crfreenet.org*>

This document contains programmer's documentation for the BIRD Internet Routing Daemon project.

# Contents

<b>1</b>	<b>BIRD Design</b>	<b>3</b>
1.1	Introduction . . . . .	3
1.2	Design goals . . . . .	3
1.3	Architecture . . . . .	4
1.4	Implementation . . . . .	4
<b>2</b>	<b>Core</b>	<b>6</b>
2.1	Forwarding Information Base . . . . .	6
2.2	Routing tables . . . . .	8
2.3	Route attribute cache . . . . .	14
2.4	Routing protocols . . . . .	19
2.5	Graceful restart recovery . . . . .	23
2.6	Protocol hooks . . . . .	26
2.7	Interfaces . . . . .	32
2.8	Neighbor cache . . . . .	35
2.9	Command line interface . . . . .	37
2.10	Object locks . . . . .	38
<b>3</b>	<b>Configuration</b>	<b>40</b>
3.1	Configuration manager . . . . .	40
3.2	Lexical analyzer . . . . .	43
3.3	Parser . . . . .	45
<b>4</b>	<b>Filters</b>	<b>46</b>
4.1	Filters . . . . .	46
4.2	Trie for prefix sets . . . . .	48
<b>5</b>	<b>Protocols</b>	<b>51</b>
5.1	The Babel protocol . . . . .	51
5.2	Bidirectional Forwarding Detection . . . . .	54
5.3	Border Gateway Protocol . . . . .	55
5.4	Open Shortest Path First (OSPF) . . . . .	62
5.5	Pipe . . . . .	70
5.6	Router Advertisements . . . . .	70
5.7	Routing Information Protocol (RIP) . . . . .	70
5.8	RPKI To Router (RPKI-RTR) . . . . .	73
5.9	Static . . . . .	79
5.10	Direct . . . . .	79
<b>6</b>	<b>System dependent parts</b>	<b>80</b>
6.1	Introduction . . . . .	80
6.2	Logging . . . . .	80
6.3	Kernel synchronization . . . . .	81
<b>7</b>	<b>Library functions</b>	<b>83</b>
7.1	IP addresses . . . . .	83
7.2	Linked lists . . . . .	87

---

7.3	Miscellaneous functions. . . . .	89
7.4	Message authentication codes . . . . .	93
7.5	Flow specification (flowspec) . . . . .	95
<b>8</b>	<b>Resources</b>	<b>102</b>
8.1	Introduction . . . . .	102
8.2	Resource pools . . . . .	102
8.3	Memory blocks . . . . .	104
8.4	Linear memory pools . . . . .	105
8.5	Slabs . . . . .	107
8.6	Events . . . . .	108
8.7	Sockets . . . . .	109

# Chapter 1: BIRD Design

## 1.1 Introduction

This document describes the internal workings of BIRD, its architecture, design decisions and rationale behind them. It also contains documentation on all the essential components of the system and their interfaces.

Routing daemons are complicated things which need to act in real time to complex sequences of external events, respond correctly even to the most erroneous behavior of their environment and still handle enormous amount of data with reasonable speed. Due to all of this, their design is very tricky as one needs to carefully balance between efficiency, stability and (last, but not least) simplicity of the program and it would be possible to write literally hundreds of pages about all of these issues. In accordance to the famous quote of Anton Chekhov "Shortness is a sister of talent", we've tried to write a much shorter document highlighting the most important stuff and leaving the boring technical details better explained by the program source itself together with comments contained therein.

## 1.2 Design goals

When planning the architecture of BIRD, we've taken a close look at the other existing routing daemons and also at some of the operating systems used on dedicated routers, gathered all important features and added lots of new ones to overcome their shortcomings and to better match the requirements of routing in today's Internet: IPv6, policy routing, route filtering and so on. From this planning, the following set of design goals has arisen:

- *Support all the standard routing protocols and make it easy to add new ones.* This leads to modularity and clean separation between the core and the protocols.
- *Support both IPv4 and IPv6 in the same source tree, re-using most of the code.* This leads to abstraction of IP addresses and operations on them.
- *Minimize OS dependent code to make porting as easy as possible.* Unfortunately, such code cannot be avoided at all as the details of communication with the IP stack differ from OS to OS and they often vary even between different versions of the same OS. But we can isolate such code in special modules and do the porting by changing or replacing just these modules. Also, don't rely on specific features of various operating systems, but be able to make use of them if they are available.
- *Allow multiple routing tables.* Easily solvable by abstracting out routing tables and the corresponding operations.
- *Offer powerful route filtering.* There already were several attempts to incorporate route filters to a dynamic router, but most of them have used simple sequences of filtering rules which were very inflexible and hard to use for non-trivial filters. We've decided to employ a simple loop-free programming language having access to all the route attributes and being able to modify the most of them.
- *Support easy configuration and re-configuration.* Most routers use a simple configuration language designed ad hoc with no structure at all and allow online changes of configuration by using their command-line interface, thus any complex re-configurations are hard to achieve without replacing the configuration file and restarting the whole router. We've decided to use a more general approach: to have a configuration defined in a context-free language with blocks and nesting, to perform all configuration changes by editing the configuration file, but to be able to read the new configuration and smoothly adapt to it without disturbing parts of the routing process which are not affected by the change.
- *Be able to be controlled online.* In addition to the online reconfiguration, a routing daemon should be able to communicate with the user and with many other programs (primarily scripts used for network maintenance) in order to make it possible to inspect contents of routing tables, status of all routing protocols and also to control their behavior (disable, enable or reset a protocol without restarting all

the others). To achieve this, we implement a simple command-line protocol based on those used by FTP and SMTP (that is textual commands and textual replies accompanied by a numeric code which makes them both readable to a human and easy to recognize in software).

- *Respond to all events in real time.* A typical solution to this problem is to use lots of threads to separate the workings of all the routing protocols and also of the user interface parts and to hope that the scheduler will assign time to them in a fair enough manner. This is surely a good solution, but we have resisted the temptation and preferred to avoid the overhead of threading and the large number of locks involved and preferred a event driven architecture with our own scheduling of events. An unpleasant consequence of such an approach is that long lasting tasks must be split to more parts linked by special events or timers to make the CPU available for other tasks as well.

## 1.3 Architecture

The requirements set above have lead to a simple modular architecture containing the following types of modules:

### Core modules

implement the core functions of BIRD: taking care of routing tables, keeping protocol status, interacting with the user using the Command-Line Interface (to be called CLI in the rest of this document) etc.

### Library modules

form a large set of various library functions implementing several data abstractions, utility functions and also functions which are a part of the standard libraries on some systems, but missing on other ones.

### Resource management modules

take care of resources, their allocation and automatic freeing when the module having requested shuts itself down.

### Configuration modules

are fragments of lexical analyzer, grammar rules and the corresponding snippets of C code. For each group of code modules (core, each protocol, filters) there exist a configuration module taking care of all the related configuration stuff.

### The filter

implements the route filtering language.

### Protocol modules

implement the individual routing protocols.

### System-dependent modules

implement the interface between BIRD and specific operating systems.

### The client

is a simple program providing an easy, though friendly interface to the CLI.

## 1.4 Implementation

BIRD has been written in GNU C. We've considered using C++, but we've preferred the simplicity and straightforward nature of C which gives us fine control over all implementation details and on the other hand enough instruments to build the abstractions we need.

The modules are statically linked to produce a single executable file (except for the client which stands on its own).

The building process is controlled by a set of Makefiles for GNU Make, intermixed with several Perl and shell scripts.

The initial configuration of the daemon, detection of system features and selection of the right modules to include for the particular OS and the set of protocols the user has chosen is performed by a configure script

generated by GNU Autoconf.

The parser of the configuration is generated by the GNU Bison.

The documentation is generated using **SGMLtools** with our own DTD and mapping rules which produce both an online version in HTML and a neatly formatted one for printing (first converted from SGML to  $\text{\LaTeX}$  and then processed by  $\text{\TeX}$  and **dvips** to get a PostScript file).

The comments from C sources which form a part of the programmer's documentation are extracted using a modified version of the **kernel-doc** tool.

If you want to work on BIRD, it's highly recommended to configure it with a **--enable-debug** switch which enables some internal consistency checks and it also links BIRD with a memory allocation checking library if you have one (either **efence** or **dmalloc**).

# Chapter 2: Core

## 2.1 Forwarding Information Base

FIB is a data structure designed for storage of routes indexed by their network prefixes. It supports insertion, deletion, searching by prefix, ‘routing’ (in CIDR sense, that is searching for a longest prefix matching a given IP address) and (which makes the structure very tricky to implement) asynchronous reading, that is enumerating the contents of a FIB while other modules add, modify or remove entries.

Internally, each FIB is represented as a collection of nodes of type `fib_node` indexed using a sophisticated hashing mechanism. We use two-stage hashing where we calculate a 16-bit primary hash key independent on hash table size and then we just divide the primary keys modulo table size to get a real hash key used for determining the bucket containing the node. The lists of nodes in each bucket are sorted according to the primary hash key, hence if we keep the total number of buckets to be a power of two, re-hashing of the structure keeps the relative order of the nodes.

To get the asynchronous reading consistent over node deletions, we need to keep a list of readers for each node. When a node gets deleted, its readers are automatically moved to the next node in the table.

Basic FIB operations are performed by functions defined by this module, enumerating of FIB contents is accomplished by using the `FIB_WALK()` macro or `FIB_ITERATE_START()` if you want to do it asynchronously.

For simple iteration just place the body of the loop between `FIB_WALK()` and `FIB_WALK_END()`. You can’t modify the FIB during the iteration (you can modify data in the node, but not add or remove nodes).

If you need more freedom, you can use the `FIB_ITERATE_*`() group of macros. First, you initialize an iterator with `FIB_ITERATE_INIT()`. Then you can put the loop body in between `FIB_ITERATE_START()` and `FIB_ITERATE_END()`. In addition, the iteration can be suspended by calling `FIB_ITERATE_PUT()`. This’ll link the iterator inside the FIB. While suspended, you may modify the FIB, exit the current function, etc. To resume the iteration, enter the loop again. You can use `FIB_ITERATE_UNLINK()` to unlink the iterator (while iteration is suspended) in cases like premature end of FIB iteration.

Note that the iterator must not be destroyed when the iteration is suspended, the FIB would then contain a pointer to invalid memory. Therefore, after each `FIB_ITERATE_INIT()` or `FIB_ITERATE_PUT()` there must be either `FIB_ITERATE_START()` or `FIB_ITERATE_UNLINK()` before the iterator is destroyed.

---

### Function

void *fib\_init* (struct fib \* *f*, pool \* *p*, uint *addr\_type*, uint *node\_size*, uint *node\_offset*, uint *hash\_order*, fib\_init\_fn *init*) – initialize a new FIB

### Arguments

struct fib \* *f*  
the FIB to be initialized (the structure itself being allocated by the caller)

pool \* *p*  
pool to allocate the nodes in

uint *addr\_type*  
– undescribed –

uint *node\_size*  
node size to be used (each node consists of a standard header `fib_node` followed by user data)

uint *node\_offset*  
– undescribed –

uint *hash\_order*  
initial hash order (a binary logarithm of hash table size), 0 to use default order (recommended)

fib\_init\_fn *init*  
pointer a function to be called to initialize a newly created node

**Description**

This function initializes a newly allocated FIB and prepares it for use.

---

**Function**

`void * fib_find (struct fib * f, const net_addr * a)` – search for FIB node by prefix

**Arguments**

struct fib \* *f*  
    FIB to search in

const net\_addr \* *a*  
    – undescribed –

**Description**

Search for a FIB node corresponding to the given prefix, return a pointer to it or NULL if no such node exists.

---

**Function**

`void * fib_get (struct fib * f, const net_addr * a)` – find or create a FIB node

**Arguments**

struct fib \* *f*  
    FIB to work with

const net\_addr \* *a*  
    – undescribed –

**Description**

Search for a FIB node corresponding to the given prefix and return a pointer to it. If no such node exists, create it.

---

**Function**

`void * fib_route (struct fib * f, const net_addr * n)` – CIDR routing lookup

**Arguments**

struct fib \* *f*  
    FIB to search in

const net\_addr \* *n*  
    network address

**Description**

Search for a FIB node with longest prefix matching the given network, that is a node which a CIDR router would use for routing that network.

---

**Function**

`void fib_delete (struct fib * f, void * E)` – delete a FIB node

**Arguments**

struct fib \* *f*  
    FIB to delete from

void \* *E*  
    entry to delete

**Description**

This function removes the given entry from the FIB, taking care of all the asynchronous readers by shifting them to the next node in the canonical reading order.



---

**Function**

void *fib\_free* (struct fib \* *f*) – delete a FIB

**Arguments**

struct fib \* *f*  
FIB to be deleted

**Description**

This function deletes a FIB – it frees all memory associated with it and all its entries.

---

**Function**

void *fib\_check* (struct fib \* *f*) – audit a FIB

**Arguments**

struct fib \* *f*  
FIB to be checked

**Description**

This debugging function audits a FIB by checking its internal consistency. Use when you suspect somebody of corrupting innocent data structures.

## 2.2 Routing tables

Routing tables are probably the most important structures BIRD uses. They hold all the information about known networks, the associated routes and their attributes.

There are multiple routing tables (a primary one together with any number of secondary ones if requested by the configuration). Each table is basically a FIB containing entries describing the individual destination networks. For each network (represented by structure **net**), there is a one-way linked list of route entries (**rte**), the first entry on the list being the best one (i.e., the one we currently use for routing), the order of the other ones is undetermined.

The **rte** contains information specific to the route (preference, protocol metrics, time of last modification etc.) and a pointer to a **rta** structure (see the route attribute module for a precise explanation) holding the remaining route attributes which are expected to be shared by multiple routes in order to conserve memory.

---

**Function**

int *net\_roa\_check* (rtable \* *tab*, const net\_addr \* *n*, u32 *asn*) – check validity of route origination in a ROA table

**Arguments**

rtable \* *tab*  
ROA table  
  
const net\_addr \* *n*  
network prefix to check  
  
u32 *asn*  
AS number of network prefix

**Description**

Implements RFC 6483 route validation for the given network prefix. The procedure is to find all candidate ROAs - ROAs whose prefixes cover the given network prefix. If there is no candidate ROA, return ROA\_UNKNOWN. If there is a candidate ROA with matching ASN and maxlen field greater than or equal to the given prefix length, return ROA\_VALID. Otherwise, return ROA\_INVALID. If caller cannot determine origin AS, 0 could be used (in that case ROA\_VALID cannot happen). Table *tab* must have type NET\_ROA4 or NET\_ROA6, network *n* must have type NET\_IP4 or NET\_IP6, respectively.

---

**Function**

`rte * rte_find (net * net, struct rte_src * src)` – find a route

**Arguments**

`net * net`  
network node

`struct rte_src * src`  
route source

**Description**

The `rte_find()` function returns a route for destination `net` which is from route source `src`.

---

**Function**

`rte * rte_get_temp (rta * a)` – get a temporary `rte`

**Arguments**

`rta * a`  
attributes to assign to the new route (a `rta`; in case it's un-cached, `rte_update()` will create a cached copy automatically)

**Description**

Create a temporary `rte` and bind it with the attributes `a`. Also set route preference to the default preference set for the protocol.

---

**Function**

`rte * rte_cow_rta (rte * r, linpool * lp)` – get a private writable copy of `rte` with writable `rta`

**Arguments**

`rte * r`  
a route entry to be copied

`linpool * lp`  
a linpool from which to allocate `rta`

**Description**

`rte_cow_rta()` takes a `rte` and prepares it and associated `rta` for modification. There are three possibilities: First, both `rte` and `rta` are private copies, in that case they are returned unchanged. Second, `rte` is private copy, but `rta` is cached, in that case `rta` is duplicated using `rta_do_cow()`. Third, both `rte` is shared and `rta` is cached, in that case both structures are duplicated by `rte_do_cow()` and `rta_do_cow()`.

Note that in the second case, cached `rta` loses one reference, while private copy created by `rta_do_cow()` is a shallow copy sharing indirect data (eattrs, nexthops, ...) with it. To work properly, original shared `rta` should have another reference during the life of created private copy.

**Result**

a pointer to the new writable `rte` with writable `rta`.

---

**Function**

`void rte_announce (rtable * tab, unsigned type, net * net, rte * new, rte * old, rte * new_best, rte * old_best, rte * before_old)` – announce a routing table change

**Arguments**

`rtable * tab`  
     table the route has been added to

`unsigned type`  
     type of route announcement (RA\_OPTIMAL or RA\_ANY)

`net * net`  
     network in question

`rte * new`  
     the new route to be announced

`rte * old`  
     the previous route for the same network

`rte * new_best`  
     the new best route for the same network

`rte * old_best`  
     the previous best route for the same network

`rte * before_old`  
     The previous route before *old* for the same network. If *before\_old* is NULL *old* was the first.

**Description**

This function gets a routing table update and announces it to all protocols that accepts given type of route announcement and are connected to the same table by their announcement hooks.

Route announcement of type RA\_OPTIMAL si generated when optimal route (in routing table *tab*) changes. In that case *old* stores the old optimal route.

Route announcement of type RA\_ANY si generated when any route (in routing table *tab*) changes In that case *old* stores the old route from the same protocol.

For each appropriate protocol, we first call its *import\_control()* hook which performs basic checks on the route (each protocol has a right to veto or force accept of the route before any filter is asked) and adds default values of attributes specific to the new protocol (metrics, tags etc.). Then it consults the protocol's export filter and if it accepts the route, the *rt\_notify()* hook of the protocol gets called.

**Function**

`void rte_free (rte * e)` – delete a **rte**

**Arguments**

`rte * e`  
     **rte** to be deleted

**Description**

*rte\_free()* deletes the given **rte** from the routing table it's linked to.

**Function**

`void rte_update2 (struct channel * c, const net_addr * n, rte * new, struct rte_src * src)` – enter a new update to a routing table

**Arguments**

`struct channel * c`  
     channel doing the update

`const net_addr * n`  
     – undescribed –

```

rte * new
    a rte representing the new route or NULL for route removal.

struct rte_src * src
    protocol originating the update

```

### Description

This function is called by the routing protocols whenever they discover a new route or wish to update/remove an existing route. The right announcement sequence is to build route attributes first (either un-cached with *aflags* set to zero or a cached one using *rta\_lookup()*; in this case please note that you need to increase the use count of the attributes yourself by calling *rta\_clone()*), call *rte\_get\_temp()* to obtain a temporary **rte**, fill in all the appropriate data and finally submit the new **rte** by calling *rte\_update()*.

*src* specifies the protocol that originally created the route and the meaning of protocol-dependent data of *new*. If *new* is not NULL, *src* have to be the same value as *new->attrs->proto*. *p* specifies the protocol that called *rte\_update()*. In most cases it is the same protocol as *src*. *rte\_update()* stores *p* in *new->sender*;

When *rte\_update()* gets any route, it automatically validates it (checks, whether the network and next hop address are valid IP addresses and also whether a normal routing protocol doesn't try to smuggle a host or link scope route to the table), converts all protocol dependent attributes stored in the **rte** to temporary extended attributes, consults import filters of the protocol to see if the route should be accepted and/or its attributes modified, stores the temporary attributes back to the **rte**.

Now, having a "public" version of the route, we automatically find any old route defined by the protocol *src* for network *n*, replace it by the new one (or removing it if *new* is NULL), recalculate the optimal route for this destination and finally broadcast the change (if any) to all routing protocols by calling *rte\_announce()*. All memory used for attribute lists and other temporary allocations is taken from a special linear pool *rte\_update\_pool* and freed when *rte\_update()* finishes.

### Function

void *rt\_refresh\_begin* (rtable \* *t*, struct channel \* *c*) – start a refresh cycle

### Arguments

```

rtable * t
    related routing table c related channel

struct channel * c
    – undescribed –

```

### Description

This function starts a refresh cycle for given routing table and announce hook. The refresh cycle is a sequence where the protocol sends all its valid routes to the routing table (by *rte\_update()*). After that, all protocol routes (more precisely routes with *c* as *sender*) not sent during the refresh cycle but still in the table from the past are pruned. This is implemented by marking all related routes as stale by REF\_STALE flag in *rt\_refresh\_begin()*, then marking all related stale routes with REF\_DISCARD flag in *rt\_refresh\_end()* and then removing such routes in the prune loop.

### Function

void *rt\_refresh\_end* (rtable \* *t*, struct channel \* *c*) – end a refresh cycle

### Arguments

```

rtable * t
    related routing table

struct channel * c
    related channel

```

### Description

This function ends a refresh cycle for given routing table and announce hook. See *rt\_refresh\_begin()* for description of refresh cycles.

---

**Function**

void *rt\_dump* (rte \* *e*) – dump a route

**Arguments**

rte \* *e*  
rte to be dumped

**Description**

This functions dumps contents of a **rte** to debug output.

---

**Function**

void *rt\_dump* (rtable \* *t*) – dump a routing table

**Arguments**

rtable \* *t*  
routing table to be dumped

**Description**

This function dumps contents of a given routing table to debug output.

---

**Function**

void *rt\_dump\_all* (*void*) – dump all routing tables

**Description**

This function dumps contents of all routing tables to debug output.

---

**Function**

void *rt\_init* (*void*) – initialize routing tables

**Description**

This function is called during BIRD startup. It initializes the routing table module.

---

**Function**

void *rt\_prune\_table* (rtable \* *tab*) – prune a routing table

**Arguments**

rtable \* *tab*  
– undescribed –

**Description**

The prune loop scans routing tables and removes routes belonging to flushing protocols, discarded routes and also stale network entries. It is called from *rt\_event()*. The event is rescheduled if the current iteration do not finish the table. The pruning is directed by the prune state (*prune\_state*), specifying whether the prune cycle is scheduled or running, and there is also a persistent pruning iterator (*prune\_fit*).

The prune loop is used also for channel flushing. For this purpose, the channels to flush are marked before the iteration and notified after the iteration.

---

**Function**

void *rt\_lock\_table* (rtable \* *r*) – lock a routing table

**Arguments**

rtable \* *r*  
routing table to be locked

**Description**

Lock a routing table, because it's in use by a protocol, preventing it from being freed when it gets undefined in a new configuration.

---

**Function**

void *rt\_unlock\_table* (rtable \* *r*) – unlock a routing table

**Arguments**

rtable \* *r*  
routing table to be unlocked

**Description**

Unlock a routing table formerly locked by *rt\_lock\_table()*, that is decrease its use count and delete it if it's scheduled for deletion by configuration changes.

---

**Function**

void *rt\_commit* (struct config \* *new*, struct config \* *old*) – commit new routing table configuration

**Arguments**

struct config \* *new*  
new configuration

struct config \* *old*  
original configuration or NULL if it's boot time config

**Description**

Scan differences between *old* and *new* configuration and modify the routing tables according to these changes. If *new* defines a previously unknown table, create it, if it omits a table existing in *old*, schedule it for deletion (it gets deleted when all protocols disconnect from it by calling *rt\_unlock\_table()*), if it exists in both configurations, leave it unchanged.

---

**Function**

int *rt\_feed\_channel* (struct channel \* *c*) – advertise all routes to a channel

**Arguments**

struct channel \* *c*  
channel to be fed

**Description**

This function performs one pass of advertisement of routes to a channel that is in the ES\_FEEDING state. It is called by the protocol code as long as it has something to do. (We avoid transferring all the routes in single pass in order not to monopolize CPU time.)

---

**Function**

void *rt\_feed\_channel\_abort* (struct channel \* *c*) – abort protocol feeding

**Arguments**

struct channel \* *c*  
channel

**Description**

This function is called by the protocol code when the protocol stops or ceases to exist during the feeding.

**Function**

`net * net_find (rtable * tab, net_addr * addr)` – find a network entry

**Arguments**

`rtable * tab`  
a routing table

`net_addr * addr`  
address of the network

**Description**

`net_find()` looks up the given network in routing table `tab` and returns a pointer to its `net` entry or `NULL` if no such network exists.

**Function**

`net * net_get (rtable * tab, net_addr * addr)` – obtain a network entry

**Arguments**

`rtable * tab`  
a routing table

`net_addr * addr`  
address of the network

**Description**

`net_get()` looks up the given network in routing table `tab` and returns a pointer to its `net` entry. If no such entry exists, it's created.

**Function**

`rte * rte_cow (rte * r)` – copy a route for writing

**Arguments**

`rte * r`  
a route entry to be copied

**Description**

`rte_cow()` takes a `rte` and prepares it for modification. The exact action taken depends on the flags of the `rte` – if it's a temporary entry, it's just returned unchanged, else a new temporary entry with the same contents is created.

The primary use of this function is inside the filter machinery – when a filter wants to modify `rte` contents (to change the preference or to attach another set of attributes), it must ensure that the `rte` is not shared with anyone else (and especially that it isn't stored in any routing table).

**Result**

a pointer to the new writable `rte`.

## 2.3 Route attribute cache

Each route entry carries a set of route attributes. Several of them vary from route to route, but most attributes are usually common for a large number of routes. To conserve memory, we've decided to store only the varying ones directly in the `rte` and hold the rest in a special structure called `rta` which is shared among all the `rte`'s with these attributes.

Each `rta` contains all the static attributes of the route (i.e., those which are always present) as structure members and a list of dynamic attributes represented by a linked list of `ea_list` structures, each of them

consisting of an array of **eaattr**'s containing the individual attributes. An attribute can be specified more than once in the **ea\_list** chain and in such case the first occurrence overrides the others. This semantics is used especially when someone (for example a filter) wishes to alter values of several dynamic attributes, but it wants to preserve the original attribute lists maintained by another module.

Each **eaattr** contains an attribute identifier (split to protocol ID and per-protocol attribute ID), protocol dependent flags, a type code (consisting of several bit fields describing attribute characteristics) and either an embedded 32-bit value or a pointer to a **adata** structure holding attribute contents.

There exist two variants of **rta**'s – cached and un-cached ones. Un-cached **rta**'s can have arbitrarily complex structure of **ea\_list**'s and they can be modified by any module in the route processing chain. Cached **rta**'s have their attribute lists normalized (that means at most one **ea\_list** is present and its values are sorted in order to speed up searching), they are stored in a hash table to make fast lookup possible and they are provided with a use count to allow sharing.

Routing tables always contain only cached **rta**'s.

### Function

**struct nexthop \* *nexthop\_merge*** (**struct nexthop \* *x***, **struct nexthop \* *y***, **int *rx***, **int *ry***, **int *max***, **linpool \* *lp***) – merge nexthop lists

### Arguments

**struct nexthop \* *x***  
list 1

**struct nexthop \* *y***  
list 2

**int *rx***  
reusability of list *x*

**int *ry***  
reusability of list *y*

**int *max***  
max number of nexthops

**linpool \* *lp***  
linpool for allocating nexthops

### Description

The *nexthop\_merge()* function takes two nexthop lists *x* and *y* and merges them, eliminating possible duplicates. The input lists must be sorted and the result is sorted too. The number of nexthops in result is limited by *max*. New nodes are allocated from linpool *lp*.

The arguments *rx* and *ry* specify whether corresponding input lists may be consumed by the function (i.e. their nodes reused in the resulting list), in that case the caller should not access these lists after that. To eliminate issues with deallocation of these lists, the caller should use some form of bulk deallocation (e.g. stack or linpool) to free these nodes when the resulting list is no longer needed. When reusability is not set, the corresponding lists are not modified nor linked from the resulting list.

### Function

**eaattr \* *ea\_find*** (**ea\_list \* *e***, **unsigned *id***) – find an extended attribute

### Arguments

**ea\_list \* *e***  
attribute list to search in

**unsigned *id***  
attribute ID to search for

### Description

Given an extended attribute list, *ea\_find()* searches for a first occurrence of an attribute with specified ID, returning either a pointer to its **eaattr** structure or NULL if no such attribute exists.



---

**Function**

`eaattr * ea_walk (struct ea_walk_state * s, uint id, uint max)` – walk through extended attributes

**Arguments**

`struct ea_walk_state * s`  
walk state structure

`uint id`  
start of attribute ID interval

`uint max`  
length of attribute ID interval

**Description**

Given an extended attribute list, `ea_walk()` walks through the list looking for first occurrences of attributes with ID in specified interval from `id` to `(id + max - 1)`, returning pointers to found `eaattr` structures, storing its walk state in `s` for subsequent calls.

The function `ea_walk()` is supposed to be called in a loop, with initially zeroed walk state structure `s` with filled the initial extended attribute list, returning one found attribute in each call or `NULL` when no other attribute exists. The extended attribute list or the arguments should not be modified between calls. The maximum value of `max` is 128.

---

**Function**

`int ea_get_int (ea_list * e, unsigned id, int def)` – fetch an integer attribute

**Arguments**

`ea_list * e`  
attribute list

`unsigned id`  
attribute ID

`int def`  
default value

**Description**

This function is a shortcut for retrieving a value of an integer attribute by calling `ea_find()` to find the attribute, extracting its value or returning a provided default if no such attribute is present.

---

**Function**

`void ea_sort (ea_list * e)` – sort an attribute list

**Arguments**

`ea_list * e`  
list to be sorted

**Description**

This function takes a `ea_list` chain and sorts the attributes within each of its entries.

If an attribute occurs multiple times in a single `ea_list`, `ea_sort()` leaves only the first (the only significant) occurrence.

---

**Function**

`unsigned ea_scan (ea_list * e)` – estimate attribute list size

**Arguments**

`ea_list * e`  
attribute list

**Description**

This function calculates an upper bound of the size of a given `ea_list` after merging with `ea_merge()`.

---

**Function**

void *ea\_merge* (ea\_list \* *e*, ea\_list \* *t*) – merge segments of an attribute list

**Arguments**

ea\_list \* *e*  
attribute list

ea\_list \* *t*  
buffer to store the result to

**Description**

This function takes a possibly multi-segment attribute list and merges all of its segments to one.

The primary use of this function is for **ea\_list** normalization: first call *ea\_scan()* to determine how much memory will the result take, then allocate a buffer (usually using *alloca()*), merge the segments with *ea\_merge()* and finally sort and prune the result by calling *ea\_sort()*.

---

**Function**

int *ea\_same* (ea\_list \* *x*, ea\_list \* *y*) – compare two **ea\_list**'s

**Arguments**

ea\_list \* *x*  
attribute list

ea\_list \* *y*  
attribute list

**Description**

*ea\_same()* compares two normalized attribute lists *x* and *y* and returns 1 if they contain the same attributes, 0 otherwise.

---

**Function**

void *ea\_show* (struct cli \* *c*, eattr \* *e*) – print an **eattr** to CLI

**Arguments**

struct cli \* *c*  
destination CLI

eattr \* *e*  
attribute to be printed

**Description**

This function takes an extended attribute represented by its **eattr** structure and prints it to the CLI according to the type information.

If the protocol defining the attribute provides its own *get\_attr()* hook, it's consulted first.

---

**Function**

void *ea\_dump* (ea\_list \* *e*) – dump an extended attribute

**Arguments**

ea\_list \* *e*  
attribute to be dumped

**Description**

*ea\_dump()* dumps contents of the extended attribute given to the debug output.

---

**Function**

uint *ea\_hash* (ea\_list \* *e*) – calculate an **ea\_list** hash key

**Arguments**

ea\_list \* *e*  
attribute list

**Description**

*ea\_hash()* takes an extended attribute list and calculated a hopefully uniformly distributed hash value from its contents.

---

**Function**

ea\_list \* *ea\_append* (ea\_list \* *to*, ea\_list \* *what*) – concatenate **ea\_list**'s

**Arguments**

ea\_list \* *to*  
destination list (can be NULL)

ea\_list \* *what*  
list to be appended (can be NULL)

**Description**

This function appends the **ea\_list** *what* at the end of **ea\_list** *to* and returns a pointer to the resulting list.

---

**Function**

rta \* *rta\_lookup* (rta \* *o*) – look up a **rta** in attribute cache

**Arguments**

rta \* *o*  
a un-cached **rta**

**Description**

*rta\_lookup()* gets an un-cached **rta** structure and returns its cached counterpart. It starts with examining the attribute cache to see whether there exists a matching entry. If such an entry exists, it's returned and its use count is incremented, else a new entry is created with use count set to 1.

The extended attribute lists attached to the **rta** are automatically converted to the normalized form.

---

**Function**

void *rta\_dump* (rta \* *a*) – dump route attributes

**Arguments**

rta \* *a*  
attribute structure to dump

**Description**

This function takes a **rta** and dumps its contents to the debug output.

---

**Function**

void *rta\_dump\_all* (void) – dump attribute cache

**Description**

This function dumps the whole contents of route attribute cache to the debug output.

**Function**

void *rta\_init* (*void*) – initialize route attribute cache

**Description**

This function is called during initialization of the routing table module to set up the internals of the attribute cache.

**Function**

*rta \* rta\_clone* (*rta \* r*) – clone route attributes

**Arguments**

*rta \* r*  
a **rta** to be cloned

**Description**

*rta\_clone()* takes a cached **rta** and returns its identical cached copy. Currently it works by just returning the original **rta** with its use count incremented.

**Function**

void *rta\_free* (*rta \* r*) – free route attributes

**Arguments**

*rta \* r*  
a **rta** to be freed

**Description**

If you stop using a **rta** (for example when deleting a route which uses it), you need to call *rta\_free()* to notify the attribute cache the attribute is no longer in use and can be freed if you were the last user (which *rta\_free()* tests by inspecting the use count).

## 2.4 Routing protocols

### 2.4.1 Introduction

The routing protocols are the bird's heart and a fine amount of code is dedicated to their management and for providing support functions to them. (-: Actually, this is the reason why the directory with sources of the core code is called **nest** :-).

When talking about protocols, one need to distinguish between *protocols* and protocol *instances*. A protocol exists exactly once, not depending on whether it's configured or not and it can have an arbitrary number of instances corresponding to its "incarnations" requested by the configuration file. Each instance is completely autonomous, has its own configuration, its own status, its own set of routes and its own set of interfaces it works on.

A protocol is represented by a **protocol** structure containing all the basic information (protocol name, default settings and pointers to most of the protocol hooks). All these structures are linked in the *protocol\_list* list. Each instance has its own **proto** structure describing all its properties: protocol type, configuration, a resource pool where all resources belonging to the instance live, various protocol attributes (take a look at the declaration of **proto** in **protocol.h**), protocol states (see below for what do they mean), connections to routing tables, filters attached to the protocol and finally a set of pointers to the rest of protocol hooks (they are the same for all instances of the protocol, but in order to avoid extra indirections when calling the hooks from the fast path, they are stored directly in **proto**). The instance is always linked in both the global instance list (*proto\_list*) and a per-status list (either *active\_proto\_list* for running protocols, *initial\_proto\_list* for protocols being initialized or *flush\_proto\_list* when the protocol is being shut down).

The protocol hooks are described in the next chapter, for more information about configuration of protocols, please refer to the configuration chapter and also to the description of the *proto\_commit* function.

### 2.4.2 Protocol states

As startup and shutdown of each protocol are complex processes which can be affected by lots of external events (user's actions, reconfigurations, behavior of neighboring routers etc.), we have decided to supervise them by a pair of simple state machines – the protocol state machine and a core state machine.

The *protocol state machine* corresponds to internal state of the protocol and the protocol can alter its state whenever it wants to. There are the following states:

#### PS\_DOWN

The protocol is down and waits for being woken up by calling its `start()` hook.

#### PS\_START

The protocol is waiting for connection with the rest of the network. It's active, it has resources allocated, but it still doesn't want any routes since it doesn't know what to do with them.

#### PS\_UP

The protocol is up and running. It communicates with the core, delivers routes to tables and wants to hear announcement about route changes.

#### PS\_STOP

The protocol has been shut down (either by being asked by the core code to do so or due to having encountered a protocol error).

Unless the protocol is in the PS\_DOWN state, it can decide to change its state by calling the *proto\_notify\_state* function.

At any time, the core code can ask the protocol to shut itself down by calling its `stop()` hook.

### 2.4.3 Functions of the protocol module

The protocol module provides the following functions:

---

#### Function

`struct channel * proto_find_channel_by_table (struct proto * p, struct rtable * t)` – find channel connected to a routing table

#### Arguments

`struct proto * p`  
protocol instance

`struct rtable * t`  
routing table

#### Description

Returns pointer to channel or NULL

---

#### Function

`struct channel * proto_find_channel_by_name (struct proto * p, const char * n)` – find channel by its name

#### Arguments

`struct proto * p`  
protocol instance

`const char * n`  
channel name

#### Description

Returns pointer to channel or NULL

---

**Function**

struct channel \* *proto\_add\_channel* (struct proto \* *p*, struct channel\_config \* *cf*) – connect protocol to a routing table

**Arguments**

struct proto \* *p*  
     protocol instance

struct channel\_config \* *cf*  
     channel configuration

**Description**

This function creates a channel between the protocol instance *p* and the routing table specified in the configuration *cf*, making the protocol hear all changes in the table and allowing the protocol to update routes in the table.

The channel is linked in the protocol channel list and when active also in the table channel list. Channels are allocated from the global resource pool (*proto\_pool*) and they are automatically freed when the protocol is removed.

---

**Function**

void *channel\_request\_feeding* (struct channel \* *c*) – request feeding routes to the channel

**Arguments**

struct channel \* *c*  
     given channel

**Description**

Sometimes it is needed to send again all routes to the channel. This is called feeding and can be requested by this function. This would cause channel export state transition to ES\_FEEDING (during feeding) and when completed, it will switch back to ES\_READY. This function can be called even when feeding is already running, in that case it is restarted.

---

**Function**

void \* *proto\_new* (struct proto\_config \* *cf*) – create a new protocol instance

**Arguments**

struct proto\_config \* *cf*  
     – undescribed –

**Description**

When a new configuration has been read in, the core code starts initializing all the protocol instances configured by calling their *init()* hooks with the corresponding instance configuration. The initialization code of the protocol is expected to create a new instance according to the configuration by calling this function and then modifying the default settings to values wanted by the protocol.

---

**Function**

void \* *proto\_config\_new* (struct protocol \* *pr*, int *class*) – create a new protocol configuration

**Arguments**

struct protocol \* *pr*  
     protocol the configuration will belong to

int *class*  
     SYM\_PROTO or SYM\_TEMPLATE

**Description**

Whenever the configuration file says that a new instance of a routing protocol should be created, the parser calls *proto\_config\_new()* to create a configuration entry for this instance (a structure starting with the **proto\_config** header containing all the generic items followed by protocol-specific ones). Also, the configuration entry gets added to the list of protocol instances kept in the configuration.

The function is also used to create protocol templates (when class **SYM\_TEMPLATE** is specified), the only difference is that templates are not added to the list of protocol instances and therefore not initialized during *protos\_commit()*.

**Function**

void *proto\_copy\_config* (struct proto\_config \* *dest*, struct proto\_config \* *src*) – copy a protocol configuration

**Arguments**

struct proto\_config \* *dest*  
destination protocol configuration

struct proto\_config \* *src*  
source protocol configuration

**Description**

Whenever a new instance of a routing protocol is created from the template, *proto\_copy\_config()* is called to copy a content of the source protocol configuration to the new protocol configuration. Name, class and a node in *protos* list of *dest* are kept intact. *copy\_config()* protocol hook is used to copy protocol-specific data.

**Function**

void *protos\_preconfig* (struct config \* *c*) – pre-configuration processing

**Arguments**

struct config \* *c*  
new configuration

**Description**

This function calls the *preconfig()* hooks of all routing protocols available to prepare them for reading of the new configuration.

**Function**

void *protos\_commit* (struct config \* *new*, struct config \* *old*, int *force\_reconfig*, int *type*) – commit new protocol configuration

**Arguments**

struct config \* *new*  
new configuration

struct config \* *old*  
old configuration or NULL if it's boot time config

int *force\_reconfig*  
force restart of all protocols (used for example when the router ID changes)

int *type*  
type of reconfiguration (**RECONFIG\_SOFT** or **RECONFIG\_HARD**)

**Description**

Scan differences between *old* and *new* configuration and adjust all protocol instances to conform to the new configuration.

When a protocol exists in the new configuration, but it doesn't in the original one, it's immediately started. When a collision with the other running protocol would arise, the new protocol will be temporarily stopped by the locking mechanism.

When a protocol exists in the old configuration, but it doesn't in the new one, it's shut down and deleted after the shutdown completes.

When a protocol exists in both configurations, the core decides whether it's possible to reconfigure it dynamically - it checks all the core properties of the protocol (changes in filters are ignored if type is RECONFIG\_SOFT) and if they match, it asks the *reconfigure()* hook of the protocol to see if the protocol is able to switch to the new configuration. If it isn't possible, the protocol is shut down and a new instance is started with the new configuration after the shutdown is completed.

## 2.5 Graceful restart recovery

Graceful restart of a router is a process when the routing plane (e.g. BIRD) restarts but both the forwarding plane (e.g kernel routing table) and routing neighbors keep proper routes, and therefore uninterrupted packet forwarding is maintained.

BIRD implements graceful restart recovery by deferring export of routes to protocols until routing tables are refilled with the expected content. After start, protocols generate routes as usual, but routes are not propagated to them, until protocols report that they generated all routes. After that, graceful restart recovery is finished and the export (and the initial feed) to protocols is enabled.

When graceful restart recovery need is detected during initialization, then enabled protocols are marked with *gr\_recovery* flag before start. Such protocols then decide how to proceed with graceful restart, participation is voluntary. Protocols could lock the recovery for each channel by function *channel\_graceful\_restart\_lock()* (state stored in *gr\_lock* flag), which means that they want to postpone the end of the recovery until they converge and then unlock it. They also could set *gr\_wait* before advancing to PS\_UP, which means that the core should defer route export to that channel until the end of the recovery. This should be done by protocols that expect their neighbors to keep the proper routes (kernel table, BGP sessions with BGP graceful restart capability).

The graceful restart recovery is finished when either all graceful restart locks are unlocked or when graceful restart wait timer fires.

**Function**

void *graceful\_restart\_recovery* (void) – request initial graceful restart recovery

**Graceful restart recovery**

Called by the platform initialization code if the need for recovery after graceful restart is detected during boot. Have to be called before *protos\_commit()*.

**Function**

void *graceful\_restart\_init* (void) – initialize graceful restart

**Description**

When graceful restart recovery was requested, the function starts an active phase of the recovery and initializes graceful restart wait timer. The function have to be called after *protos\_commit()*.

**Function**

void *graceful\_restart\_done* (timer \*t *UNUSED*) – finalize graceful restart

**Arguments**

timer \*t *UNUSED*  
– undescribed –



**Description**

When there are no locks on graceful restart, the functions finalizes the graceful restart recovery. Protocols postponing route export until the end of the recovery are awakened and the export to them is enabled. All other related state is cleared. The function is also called when the graceful restart wait timer fires (but there are still some locks).

**Function**

`void channel_graceful_restart_lock (struct channel * c)` – lock graceful restart by channel

**Arguments**

struct channel \* c  
– undescribed –

**Description**

This function allows a protocol to postpone the end of graceful restart recovery until it converges. The lock is removed when the protocol calls `channel_graceful_restart_unlock()` or when the channel is closed. The function have to be called during the initial phase of graceful restart recovery and only for protocols that are part of graceful restart (i.e. their `gr_recovery` is set), which means it should be called from protocol start hooks.

**Function**

`void channel_graceful_restart_unlock (struct channel * c)` – unlock graceful restart by channel

**Arguments**

struct channel \* c  
– undescribed –

**Description**

This function unlocks a lock from `channel_graceful_restart_lock()`. It is also automatically called when the lock holding protocol went down.

**Function**

`void protos_dump_all (void)` – dump status of all protocols

**Description**

This function dumps status of all existing protocol instances to the debug output. It involves printing of general status information such as protocol states, its position on the protocol lists and also calling of a `dump()` hook of the protocol to print the internals.

**Function**

`void proto_build (struct protocol * p)` – make a single protocol available

**Arguments**

struct protocol \* p  
the protocol

**Description**

After the platform specific initialization code uses `protos_build()` to add all the standard protocols, it should call `proto_build()` for all platform specific protocols to inform the core that they exist.

**Function**

`void protos_build (void)` – build a protocol list

**Description**

This function is called during BIRD startup to insert all standard protocols to the global protocol list. Insertion of platform specific protocols (such as the kernel syncer) is in the domain of competence of the platform dependent startup code.

---

**Function**

void *proto\_set\_message* (struct proto \* *p*, char \* *msg*, int *len*) – set administrative message to protocol

**Arguments**

struct proto \* *p*  
     protocol

char \* *msg*  
     message

int *len*  
     message length (-1 for NULL-terminated string)

**Description**

The function sets administrative message (string) related to protocol state change. It is called by the nest code for manual enable/disable/restart commands all routes to the protocol, and by protocol-specific code when the protocol state change is initiated by the protocol. Using NULL message clears the last message. The message string may be either NULL-terminated or with an explicit length.

---

**Function**

void *channel\_notify\_limit* (struct channel \* *c*, struct channel\_limit \* *l*, int *dir*, u32 *rt\_count*)

**Arguments**

struct channel \* *c*  
     channel

struct channel\_limit \* *l*  
     limit being hit

int *dir*  
     limit direction (PLD\_\*)

u32 *rt\_count*  
     the number of routes

**Description**

The function is called by the route processing core when limit *l* is breached. It activates the limit and takes appropriate action according to *l*->action.

---

**Function**

void *proto\_notify\_state* (struct proto \* *p*, uint *state*) – notify core about protocol state change

**Arguments**

struct proto \* *p*  
     protocol the state of which has changed

uint *state*  
     – undescribed –

**Description**

Whenever a state of a protocol changes due to some event internal to the protocol (i.e., not inside a *start()* or *shutdown()* hook), it should immediately notify the core about the change by calling *proto\_notify\_state()* which will write the new state to the **proto** structure and take all the actions necessary to adapt to the new state. State change to PS\_DOWN immediately frees resources of protocol and might execute start callback of protocol; therefore, it should be used at tail positions of protocol callbacks.

## 2.6 Protocol hooks

Each protocol can provide a rich set of hook functions referred to by pointers in either the `proto` or `protocol` structure. They are called by the core whenever it wants the protocol to perform some action or to notify the protocol about any change of its environment. All of the hooks can be set to `NULL` which means to ignore the change or to take a default action.

### Function

void *preconfig* (struct protocol \* *p*, struct config \* *c*) – protocol preconfiguration

### Arguments

struct protocol \* *p*  
a routing protocol

struct config \* *c*  
new configuration

### Description

The *preconfig()* hook is called before parsing of a new configuration.

### Function

void *postconfig* (struct proto\_config \* *c*) – instance post-configuration

### Arguments

struct proto\_config \* *c*  
instance configuration

### Description

The *postconfig()* hook is called for each configured instance after parsing of the new configuration is finished.

### Function

struct proto \* *init* (struct proto\_config \* *c*) – initialize an instance

### Arguments

struct proto\_config \* *c*  
instance configuration

### Description

The *init()* hook is called by the core to create a protocol instance according to supplied protocol configuration.

### Result

a pointer to the instance created

### Function

int *reconfigure* (struct proto \* *p*, struct proto\_config \* *c*) – request instance reconfiguration

### Arguments

struct proto \* *p*  
an instance

struct proto\_config \* *c*  
new configuration

### Description

The core calls the *reconfigure()* hook whenever it wants to ask the protocol for switching to a new configuration. If the reconfiguration is possible, the hook returns 1. Otherwise, it returns 0 and the core will shut down the instance and start a new one with the new configuration.

After the protocol confirms reconfiguration, it must no longer keep any references to the old configuration since the memory it's stored in can be re-used at any time.

---

**Function**

void *dump* (struct proto \* *p*) – dump protocol state

**Arguments**

struct proto \* *p*  
an instance

**Description**

This hook dumps the complete state of the instance to the debug output.

---

**Function**

void *dump\_attrs* (rte \* *e*) – dump protocol-dependent attributes

**Arguments**

rte \* *e*  
a route entry

**Description**

This hook dumps all attributes in the *rte* which belong to this protocol to the debug output.

---

**Function**

int *start* (struct proto \* *p*) – request instance startup

**Arguments**

struct proto \* *p*  
protocol instance

**Description**

The *start()* hook is called by the core when it wishes to start the instance. Multitable protocols should lock their tables here.

**Result**

new protocol state

---

**Function**

int *shutdown* (struct proto \* *p*) – request instance shutdown

**Arguments**

struct proto \* *p*  
protocol instance

**Description**

The *stop()* hook is called by the core when it wishes to shut the instance down for some reason.

**Returns**

new protocol state

---

**Function**

void *cleanup* (struct proto \* *p*) – request instance cleanup

**Arguments**

struct proto \* *p*  
protocol instance

**Description**

The *cleanup()* hook is called by the core when the protocol became hungry/down, i.e. all protocol ahooks and routes are flushed. Multitable protocols should unlock their tables here.

---

**Function**

void *get\_status* (struct proto \* *p*, byte \* *buf*) – get instance status

**Arguments**

struct proto \* *p*  
     protocol instance

byte \* *buf*  
     buffer to be filled with the status string

**Description**

This hook is called by the core if it wishes to obtain an brief one-line user friendly representation of the status of the instance to be printed by the <cf/show protocols/ command.

---

**Function**

void *get\_route\_info* (rte \* *e*, byte \* *buf*, ea\_list \* *attrs*) – get route information

**Arguments**

rte \* *e*  
     a route entry

byte \* *buf*  
     buffer to be filled with the resulting string

ea\_list \* *attrs*  
     extended attributes of the route

**Description**

This hook is called to fill the buffer *buf* with a brief user friendly representation of metrics of a route belonging to this protocol.

---

**Function**

int *get\_attr* (eattr \* *a*, byte \* *buf*, int *buflen*) – get attribute information

**Arguments**

eattr \* *a*  
     an extended attribute

byte \* *buf*  
     buffer to be filled with attribute information

int *buflen*  
     a length of the *buf* parameter

**Description**

The *get\_attr()* hook is called by the core to obtain a user friendly representation of an extended route attribute. It can either leave the whole conversion to the core (by returning **GA\_UNKNOWN**), fill in only attribute name (and let the core format the attribute value automatically according to the type field; by returning **GA\_NAME**) or doing the whole conversion (used in case the value requires extra care; return **GA\_FULL**).

---

**Function**

void *if\_notify* (struct proto \* *p*, unsigned *flags*, struct iface \* *i*) – notify instance about interface changes

**Arguments**

struct proto \* *p*  
     protocol instance

unsigned *flags*  
     interface change flags

struct iface \* *i*  
     the interface in question

**Description**

This hook is called whenever any network interface changes its status. The change is described by a combination of status bits (`IF_CHANGE_XXX`) in the *flags* parameter.

**Function**

void *ifa\_notify* (struct proto \* *p*, unsigned *flags*, struct ifa \* *a*) – notify instance about interface address changes

**Arguments**

struct proto \* *p*  
     protocol instance

unsigned *flags*  
     address change flags

struct ifa \* *a*  
     the interface address

**Description**

This hook is called to notify the protocol instance about an interface acquiring or losing one of its addresses. The change is described by a combination of status bits (`IF_CHANGE_XXX`) in the *flags* parameter.

**Function**

void *rt\_notify* (struct proto \* *p*, net \* *net*, rte \* *new*, rte \* *old*, ea\_list \* *attrs*) – notify instance about routing table change

**Arguments**

struct proto \* *p*  
     protocol instance

net \* *net*  
     a network entry

rte \* *new*  
     new route for the network

rte \* *old*  
     old route for the network

ea\_list \* *attrs*  
     extended attributes associated with the *new* entry

**Description**

The *rt\_notify()* hook is called to inform the protocol instance about changes in the connected routing table *table*, that is a route *old* belonging to network *net* being replaced by a new route *new* with extended attributes *attrs*. Either *new* or *old* or both can be NULL if the corresponding route doesn't exist.

If the type of route announcement is `RA_OPTIMAL`, it is an announcement of optimal route change, *new* stores the new optimal route and *old* stores the old optimal route.

If the type of route announcement is `RA_ANY`, it is an announcement of any route change, *new* stores the new route and *old* stores the old route from the same protocol.

*p->accept\_ra\_types* specifies which kind of route announcements protocol wants to receive.

---

**Function**

void *neigh\_notify* (neighbor \* *neigh*) – notify instance about neighbor status change

**Arguments**

neighbor \* *neigh*  
a neighbor cache entry

**Description**

The *neigh\_notify()* hook is called by the neighbor cache whenever a neighbor changes its state, that is it gets disconnected or a sticky neighbor gets connected.

---

**Function**

ea\_list \* *make\_tmp\_attrs* (rte \* *e*, struct linpool \* *pool*) – convert embedded attributes to temporary ones

**Arguments**

rte \* *e*  
route entry  
  
struct linpool \* *pool*  
linear pool to allocate attribute memory in

**Description**

This hook is called by the routing table functions if they need to convert the protocol attributes embedded directly in the **rte** to temporary extended attributes in order to distribute them to other protocols or to filters. *make\_tmp\_attrs()* creates an **ea\_list** in the linear pool *pool*, fills it with values of the temporary attributes and returns a pointer to it.

---

**Function**

void *store\_tmp\_attrs* (rte \* *e*, ea\_list \* *attrs*) – convert temporary attributes to embedded ones

**Arguments**

rte \* *e*  
route entry  
  
ea\_list \* *attrs*  
temporary attributes to be converted

**Description**

This hook is an exact opposite of *make\_tmp\_attrs()* – it takes a list of extended attributes and converts them to attributes embedded in the **rte** corresponding to this protocol.

You must be prepared for any of the attributes being missing from the list and use default values instead.

---

**Function**

int *import\_control* (struct proto \* *p*, rte \*\* *e*, ea\_list \*\* *attrs*, struct linpool \* *pool*) – pre-filtering decisions on route import

**Arguments**

struct proto \* *p*  
protocol instance the route is going to be imported to  
  
rte \*\* *e*  
the route in question  
  
ea\_list \*\* *attrs*  
extended attributes of the route

```
struct linpool * pool
    linear pool for allocation of all temporary data
```

### Description

The *import\_control()* hook is called as the first step of exporting a route from a routing table to the protocol instance. It can modify route attributes and force acceptance or rejection of the route regardless of user-specified filters. See *rte\_announce()* for a complete description of the route distribution process.

The standard use of this hook is to reject routes having originated from the same instance and to set default values of the protocol's metrics.

### Result

1 if the route has to be accepted, -1 if rejected and 0 if it should be passed to the filters.

### Function

int *rte\_recalculate* (struct rtable \* *table*, struct network \* *net*, struct rte \* *new*, struct rte \* *old*, struct rte \* *old\_best*) – prepare routes for comparison

### Arguments

```
struct rtable * table
    a routing table

struct network * net
    a network entry

struct rte * new
    new route for the network

struct rte * old
    old route for the network

struct rte * old_best
    old best route for the network (may be NULL)
```

### Description

This hook is called when a route change (from *old* to *new* for a *net* entry) is propagated to a *table*. It may be used to prepare routes for comparison by *rte\_better()* in the best route selection. *new* may or may not be in *net->routes* list, *old* is not there.

### Result

1 if the ordering implied by *rte\_better()* changes enough that full best route calculation have to be done, 0 otherwise.

### Function

int *rte\_better* (rte \* *new*, rte \* *old*) – compare metrics of two routes

### Arguments

```
rte * new
    the new route

rte * old
    the original route
```

### Description

This hook gets called when the routing table contains two routes for the same network which have originated from different instances of a single protocol and it wants to select which one is preferred over the other one. Protocols usually decide according to route metrics.

### Result

1 if *new* is better (more preferred) than *old*, 0 otherwise.



**Function**

int *rte\_same* (rte \* *e1*, rte \* *e2*) – compare two routes

**Arguments**

rte \* *e1*  
route

rte \* *e2*  
route

**Description**

The *rte\_same()* hook tests whether the routes *e1* and *e2* belonging to the same protocol instance have identical contents. Contents of **rta**, all the extended attributes and **rte** preference are checked by the core code, no need to take care of them here.

**Result**

1 if *e1* is identical to *e2*, 0 otherwise.

**Function**

void *rte\_insert* (net \* *n*, rte \* *e*) – notify instance about route insertion

**Arguments**

net \* *n*  
network

rte \* *e*  
route

**Description**

This hook is called whenever a **rte** belonging to the instance is accepted for insertion to a routing table. Please avoid using this function in new protocols.

**Function**

void *rte\_remove* (net \* *n*, rte \* *e*) – notify instance about route removal

**Arguments**

net \* *n*  
network

rte \* *e*  
route

**Description**

This hook is called whenever a **rte** belonging to the instance is removed from a routing table. Please avoid using this function in new protocols.

## 2.7 Interfaces

The interface module keeps track of all network interfaces in the system and their addresses.

Each interface is represented by an **iface** structure which carries interface capability flags (**IF\_MULTIACCESS**, **IF\_BROADCAST** etc.), MTU, interface name and index and finally a linked list of network prefixes assigned to the interface, each one represented by struct **ifa**.

The interface module keeps a ‘soft-up’ state for each **iface** which is a conjunction of link being up, the interface being of a ‘sane’ type and at least one IP address assigned to it.

---

**Function**

void *ifa\_dump* (struct ifa \* *a*) – dump interface address

**Arguments**

struct ifa \* *a*  
interface address descriptor

**Description**

This function dumps contents of an *ifa* to the debug output.

---

**Function**

void *if\_dump* (struct iface \* *i*) – dump interface

**Arguments**

struct iface \* *i*  
interface to dump

**Description**

This function dumps all information associated with a given network interface to the debug output.

---

**Function**

void *if\_dump\_all* (*void*) – dump all interfaces

**Description**

This function dumps information about all known network interfaces to the debug output.

---

**Function**

void *if\_delete* (struct iface \* *old*) – remove interface

**Arguments**

struct iface \* *old*  
interface

**Description**

This function is called by the low-level platform dependent code whenever it notices an interface disappears. It is just a shorthand for *if\_update()*.

---

**Function**

struct iface \* *if\_update* (struct iface \* *new*) – update interface status

**Arguments**

struct iface \* *new*  
new interface status

**Description**

*if\_update()* is called by the low-level platform dependent code whenever it notices an interface change. There exist two types of interface updates – synchronous and asynchronous ones. In the synchronous case, the low-level code calls *if\_start\_update()*, scans all interfaces reported by the OS, uses *if\_update()* and *ifa\_update()* to pass them to the core and then it finishes the update sequence by calling *if\_end\_update()*. When working asynchronously, the sysdep code calls *if\_update()* and *ifa\_update()* whenever it notices a change. *if\_update()* will automatically notify all other modules about the change.

---

**Function**

void *if\_feed\_baby* (struct proto \* *p*) – advertise interfaces to a new protocol

**Arguments**

struct proto \* *p*  
protocol to feed

**Description**

When a new protocol starts, this function sends it a series of notifications about all existing interfaces.

---

**Function**

struct iface \* *if\_find\_by\_index* (unsigned *idx*) – find interface by ifindex

**Arguments**

unsigned *idx*  
ifindex

**Description**

This function finds an **iface** structure corresponding to an interface of the given index *idx*. Returns a pointer to the structure or NULL if no such structure exists.

---

**Function**

struct iface \* *if\_find\_by\_name* (char \* *name*) – find interface by name

**Arguments**

char \* *name*  
interface name

**Description**

This function finds an **iface** structure corresponding to an interface of the given name *name*. Returns a pointer to the structure or NULL if no such structure exists.

---

**Function**

struct ifa \* *ifa\_update* (struct ifa \* *a*) – update interface address

**Arguments**

struct ifa \* *a*  
new interface address

**Description**

This function adds address information to a network interface. It's called by the platform dependent code during the interface update process described under *if\_update()*.

---

**Function**

void *ifa\_delete* (struct ifa \* *a*) – remove interface address

**Arguments**

struct ifa \* *a*  
interface address

**Description**

This function removes address information from a network interface. It's called by the platform dependent code during the interface update process described under *if\_update()*.

---

**Function**

void *if\_init* (*void*) – initialize interface module

**Description**

This function is called during BIRD startup to initialize all data structures of the interface module.

## 2.8 Neighbor cache

Most routing protocols need to associate their internal state data with neighboring routers, check whether an address given as the next hop attribute of a route is really an address of a directly connected host and which interface is it connected through. Also, they often need to be notified when a neighbor ceases to exist or when their long awaited neighbor becomes connected. The neighbor cache is there to solve all these problems.

The neighbor cache maintains a collection of neighbor entries. Each entry represents one IP address corresponding to either our directly connected neighbor or our own end of the link (when the scope of the address is set to `SCOPE_HOST`) together with per-neighbor data belonging to a single protocol.

Active entries represent known neighbors and are stored in a hash table (to allow fast retrieval based on the IP address of the node) and two linked lists: one global and one per-interface (allowing quick processing of interface change events). Inactive entries exist only when the protocol has explicitly requested it via the `NEF_STICKY` flag because it wishes to be notified when the node will again become a neighbor. Such entries are enqueued in a special list which is walked whenever an interface changes its state to up. Neighbor entry VRF association is implied by respective protocol.

When a neighbor event occurs (a neighbor gets disconnected or a sticky inactive neighbor becomes connected), the protocol hook *neigh\_notify()* is called to advertise the change.

### Function

neighbor \* *neigh\_find* (struct proto \* *p*, ip\_addr \* *a*, unsigned *flags*) – find or create a neighbor entry.

### Arguments

struct proto \* *p*  
     protocol which asks for the entry.

ip\_addr \* *a*  
     pointer to IP address of the node to be searched for.

unsigned *flags*  
     0 or `NEF_STICKY` if you want to create a sticky entry.

### Description

Search the neighbor cache for a node with given IP address. If it's found, a pointer to the neighbor entry is returned. If no such entry exists and the node is directly connected on one of our active interfaces, a new entry is created and returned to the caller with protocol-dependent fields initialized to zero. If the node is not connected directly or *\*a* is not a valid unicast IP address, *neigh\_find()* returns `NULL`.

### Function

void *neigh\_dump* (neighbor \* *n*) – dump specified neighbor entry.

### Arguments

neighbor \* *n*  
     the entry to dump

### Description

This functions dumps the contents of a given neighbor entry to debug output.

### Function

void *neigh\_dump\_all* (*void*) – dump all neighbor entries.

### Description

This function dumps the contents of the neighbor cache to debug output.

---

**Function**

void *neigh\_if\_up* (struct iface \* *i*)

**Arguments**

struct iface \* *i*  
interface in question

**Description**

Tell the neighbor cache that a new interface became up.

The neighbor cache wakes up all inactive sticky neighbors with addresses belonging to prefixes of the interface *i*.

---

**Function**

void *neigh\_if\_down* (struct iface \* *i*) – notify neighbor cache about interface down event

**Arguments**

struct iface \* *i*  
the interface in question

**Description**

Notify the neighbor cache that an interface has ceased to exist.

It causes all entries belonging to neighbors connected to this interface to be flushed.

---

**Function**

void *neigh\_if\_link* (struct iface \* *i*) – notify neighbor cache about interface link change

**Arguments**

struct iface \* *i*  
the interface in question

**Description**

Notify the neighbor cache that an interface changed link state. All owners of neighbor entries connected to this interface are notified.

---

**Function**

void *neigh\_ifa\_update* (struct ifa \* *a*)

**Arguments**

struct ifa \* *a*  
interface address in question

**Description**

Tell the neighbor cache that an address was added or removed.

The neighbor cache wakes up all inactive sticky neighbors with addresses belonging to prefixes of the interface belonging to *ifa* and causes all unreachable neighbors to be flushed.

---

**Function**

void *neigh\_prune* (*void*) – prune neighbor cache

**Description**

*neigh\_prune()* examines all neighbor entries cached and removes those corresponding to inactive protocols. It's called whenever a protocol is shut down to get rid of all its heritage.

**Function**

void *neigh\_init* (pool \* *if\_pool*) – initialize the neighbor cache.

**Arguments**

pool \* *if\_pool*  
resource pool to be used for neighbor entries.

**Description**

This function is called during BIRD startup to initialize the neighbor cache module.

## 2.9 Command line interface

This module takes care of the BIRD's command-line interface (CLI). The CLI exists to provide a way to control BIRD remotely and to inspect its status. It uses a very simple textual protocol over a stream connection provided by the platform dependent code (on UNIX systems, it's a UNIX domain socket).

Each session of the CLI consists of a sequence of request and replies, slightly resembling the FTP and SMTP protocols. Requests are commands encoded as a single line of text, replies are sequences of lines starting with a four-digit code followed by either a space (if it's the last line of the reply) or a minus sign (when the reply is going to continue with the next line), the rest of the line contains a textual message semantics of which depends on the numeric code. If a reply line has the same code as the previous one and it's a continuation line, the whole prefix can be replaced by a single white space character.

Reply codes starting with 0 stand for 'action successfully completed' messages, 1 means 'table entry', 8 'runtime error' and 9 'syntax error'.

Each CLI session is internally represented by a `cli` structure and a resource pool containing all resources associated with the connection, so that it can be easily freed whenever the connection gets closed, not depending on the current state of command processing.

The CLI commands are declared as a part of the configuration grammar by using the `CF_CLI` macro. When a command is received, it is processed by the same lexical analyzer and parser as used for the configuration, but it's switched to a special mode by prepending a fake token to the text, so that it uses only the CLI command rules. Then the parser invokes an execution routine corresponding to the command, which either constructs the whole reply and returns it back or (in case it expects the reply will be long) it prints a partial reply and asks the CLI module (using the *cont* hook) to call it again when the output is transferred to the user.

The *this\_cli* variable points to a `cli` structure of the session being currently parsed, but it's of course available only in command handlers not entered using the *cont* hook.

TX buffer management works as follows: At `cli.tx.buf` there is a list of TX buffers (struct `cli_out`), `cli.tx.write` is the buffer currently used by the producer (*cli\_printf()*, *cli\_alloc\_out()*) and `cli.tx.pos` is the buffer currently used by the consumer (*cli\_write()*, in system dependent code). The producer uses `cli_out.wpos` ptr as the current write position and the consumer uses `cli_out.outpos` ptr as the current read position. When the producer produces something, it calls *cli\_write\_trigger()*. If there is not enough space in the current buffer, the producer allocates the new one. When the consumer processes everything in the buffer queue, it calls *cli\_written()*, tha frees all buffers (except the first one) and schedules `cli.event` .

**Function**

void *cli\_printf* (cli \* *c*, int *code*, char \* *msg*, ... ...) – send reply to a CLI connection

**Arguments**

cli \* *c*  
CLI connection

int *code*  
numeric code of the reply, negative for continuation lines

```
char * msg
    a printf()-like formatting string.

... ...
    variable arguments
```

**Description**

This function send a single line of reply to a given CLI connection. It works in all aspects like *bsprintf()* except that it automatically prepends the reply line prefix.

Please note that if the connection can be already busy sending some data in which case *cli\_printf()* stores the output to a temporary buffer, so please avoid sending a large batch of replies without waiting for the buffers to be flushed.

If you want to write to the current CLI output, you can use the *cli\_msg()* macro instead.

**Function**

void *cli\_init* (*void*) – initialize the CLI module

**Description**

This function is called during BIRD startup to initialize the internal data structures of the CLI module.

## 2.10 Object locks

The lock module provides a simple mechanism for avoiding conflicts between various protocols which would like to use a single physical resource (for example a network port). It would be easy to say that such collisions can occur only when the user specifies an invalid configuration and therefore he deserves to get what he has asked for, but unfortunately they can also arise legitimately when the daemon is reconfigured and there exists (although for a short time period only) an old protocol instance being shut down and a new one willing to start up on the same interface.

The solution is very simple: when any protocol wishes to use a network port or some other non-shareable resource, it asks the core to lock it and it doesn't use the resource until it's notified that it has acquired the lock.

Object locks are represented by `object_lock` structures which are in turn a kind of resource. Lockable resources are uniquely determined by resource type (`OBJLOCK_UDP` for a UDP port etc.), IP address (usually a broadcast or multicast address the port is bound to), port number, interface and optional instance ID.

**Function**

struct `object_lock` \* *olock\_new* (*pool* \* *p*) – create an object lock

**Arguments**

```
pool * p
    resource pool to create the lock in.
```

**Description**

The *olock\_new()* function creates a new resource of type `object_lock` and returns a pointer to it. After filling in the structure, the caller should call *olock\_acquire()* to do the real locking.

**Function**

void *olock\_acquire* (struct `object_lock` \* *l*) – acquire a lock

**Arguments**

```
struct object_lock * l
    the lock to acquire
```

**Description**

This function attempts to acquire exclusive access to the non-shareable resource described by the lock *l*. It returns immediately, but as soon as the resource becomes available, it calls the *hook()* function set up by the caller.

When you want to release the resource, just *rfree()* the lock.

---

**Function**

void *olock\_init* (*void*) – initialize the object lock mechanism

**Description**

This function is called during BIRD startup. It initializes all the internal data structures of the lock module.



# Chapter 3: Configuration

## 3.1 Configuration manager

Configuration of BIRD is complex, yet straightforward. There are three modules taking care of the configuration: config manager (which takes care of storage of the config information and controls switching between configs), lexical analyzer and parser.

The configuration manager stores each config as a `config` structure accompanied by a linear pool from which all information associated with the config and pointed to by the `config` structure is allocated.

There can exist up to four different configurations at one time: an active one (pointed to by `config`), configuration we are just switching from (`old_config`), one queued for the next reconfiguration (`future_config`; if there is one and the user wants to reconfigure once again, we just free the previous queued config and replace it with the new one) and finally a config being parsed (`new_config`). The stored `old_config` is also used for undo reconfiguration, which works in a similar way. Reconfiguration could also have timeout (using `config_timer`) and undo is automatically called if the new configuration is not confirmed later. The new config (`new_config`) and associated linear pool (`cfg_mem`) is non-NULL only during parsing.

Loading of new configuration is very simple: just call `config_alloc()` to get a new `config` structure, then use `config_parse()` to parse a configuration file and fill all fields of the structure and finally ask the config manager to switch to the new config by calling `config_commit()`.

CLI commands are parsed in a very similar way – there is also a stripped-down `config` structure associated with them and they are lex-ed and parsed by the same functions, only a special fake token is prepended before the command text to make the parser recognize only the rules corresponding to CLI commands.

---

### Function

struct config \* *config\_alloc* (const char \* *name*) – allocate a new configuration

### Arguments

const char \* *name*  
name of the config

### Description

This function creates new `config` structure, attaches a resource pool and a linear memory pool to it and makes it available for further use. Returns a pointer to the structure.

---

### Function

int *config\_parse* (struct config \* *c*) – parse a configuration

### Arguments

struct config \* *c*  
configuration

### Description

*config\_parse()* reads input by calling a hook function pointed to by *cf\_read\_hook* and parses it according to the configuration grammar. It also calls all the preconfig and postconfig hooks before, resp. after parsing.

### Result

1 if the config has been parsed successfully, 0 if any error has occurred (such as anybody calling *cf\_error()*) and the *err\_msg* field has been set to the error message.

---

### Function

int *cli\_parse* (struct config \* *c*) – parse a CLI command

**Arguments**

struct config \* *c*  
temporary config structure

**Description**

*cli\_parse()* is similar to *config\_parse()*, but instead of a configuration, it parses a CLI command. See the CLI module for more information.

---

**Function**

void *config\_free* (struct config \* *c*) – free a configuration

**Arguments**

struct config \* *c*  
configuration to be freed

**Description**

This function takes a `config` structure and frees all resources associated with it.

---

**Function**

int *config\_commit* (struct config \* *c*, int *type*, uint *timeout*) – commit a configuration

**Arguments**

struct config \* *c*  
new configuration

int *type*  
type of reconfiguration (RECONFIG\_SOFT or RECONFIG\_HARD)

uint *timeout*  
timeout for undo (in seconds; or 0 for no timeout)

**Description**

When a configuration is parsed and prepared for use, the *config\_commit()* function starts the process of reconfiguration. It checks whether there is already a reconfiguration in progress in which case it just queues the new config for later processing. Else it notifies all modules about the new configuration by calling their *commit()* functions which can either accept it immediately or call *config\_add\_obstacle()* to report that they need some time to complete the reconfiguration. After all such obstacles are removed using *config\_del\_obstacle()*, the old configuration is freed and everything runs according to the new one.

When *timeout* is nonzero, the undo timer is activated with given timeout. The timer is deactivated when *config\_commit()*, *config\_confirm()* or *config\_undo()* is called.

**Result**

CONF\_DONE if the configuration has been accepted immediately, CONF\_PROGRESS if it will take some time to switch to it, CONF\_QUEUED if it's been queued due to another reconfiguration being in progress now or CONF\_SHUTDOWN if BIRD is in shutdown mode and no new configurations are accepted.

---

**Function**

int *config\_confirm* (void) – confirm a committed configuration

**Description**

When the undo timer is activated by *config\_commit()* with nonzero timeout, this function can be used to deactivate it and therefore confirm the current configuration.

**Result**

CONF\_CONFIRM when the current configuration is confirmed, CONF\_NONE when there is nothing to confirm (i.e. undo timer is not active).

**Function**

int *config\_undo* (*void*) – undo a configuration

**Description**

Function *config\_undo()* can be used to change the current configuration back to stored *old\_config*. If no reconfiguration is running, this stored configuration is committed in the same way as a new configuration in *config\_commit()*. If there is already a reconfiguration in progress and no next reconfiguration is scheduled, then the undo is scheduled for later processing as usual, but if another reconfiguration is already scheduled, then such reconfiguration is removed instead (i.e. undo is applied on the last commit that scheduled it).

**Result**

*CONF\_DONE* if the configuration has been accepted immediately, *CONF\_PROGRESS* if it will take some time to switch to it, *CONF\_QUEUED* if it's been queued due to another reconfiguration being in progress now, *CONF\_UNQUEUED* if a scheduled reconfiguration is removed, *CONF\_NOTHING* if there is no relevant configuration to undo (the previous config request was *config\_undo()* too) or *CONF\_SHUTDOWN* if BIRD is in shutdown mode and no new configuration changes are accepted.

**Function**

void *order\_shutdown* (*void*) – order BIRD shutdown

**Description**

This function initiates shutdown of BIRD. It's accomplished by asking for switching to an empty configuration.

**Function**

void *cf\_error* (const char \* *msg*, ... ...) – report a configuration error

**Arguments**

const char \* *msg*  
     printf-like format string

... ...  
     variable arguments

**Description**

*cf\_error()* can be called during execution of *config\_parse()*, that is from the parser, a preconfig hook or a postconfig hook, to report an error in the configuration.

**Function**

char \* *cfg\_strdup* (const char \* *c*) – copy a string to config memory

**Arguments**

const char \* *c*  
     string to copy

**Description**

*cfg\_strdup()* creates a new copy of the string in the memory pool associated with the configuration being currently parsed. It's often used when a string literal occurs in the configuration and we want to preserve it for further use.

## 3.2 Lexical analyzer

The lexical analyzer used for configuration files and CLI commands is generated using the **flex** tool accompanied by a couple of functions maintaining the hash tables containing information about symbols and keywords.

Each symbol is represented by a **symbol** structure containing name of the symbol, its lexical scope, symbol class (**SYM\_PROTO** for a name of a protocol, **SYM\_CONSTANT** for a constant etc.) and class dependent data. When an unknown symbol is encountered, it's automatically added to the symbol table with class **SYM\_VOID**. The keyword tables are generated from the grammar templates using the **gen\_keywords.m4** script.

### Function

void *cf\_lex\_unwind* (*void*) – unwind lexer state during error

### Lexical analyzer

*cf\_lex\_unwind()* frees the internal state on IFS stack when the lexical analyzer is terminated by *cf\_error()*.

### Function

struct symbol \* *cf\_find\_symbol* (struct config \* *cfg*, byte \* *c*) – find a symbol by name

### Arguments

struct config \* *cfg*  
     specified config

byte \* *c*  
     symbol name

### Description

This functions searches the symbol table in the config *cfg* for a symbol of given name. First it examines the current scope, then the second recent one and so on until it either finds the symbol and returns a pointer to its **symbol** structure or reaches the end of the scope chain and returns **NULL** to signify no match.

### Function

struct symbol \* *cf\_get\_symbol* (byte \* *c*) – get a symbol by name

### Arguments

byte \* *c*  
     symbol name

### Description

This functions searches the symbol table of the currently parsed config (*new\_config*) for a symbol of given name. It returns either the already existing symbol or a newly allocated undefined (**SYM\_VOID**) symbol if no existing symbol is found.

### Function

struct symbol \* *cf\_define\_symbol* (struct symbol \* *sym*, int *type*, void \* *def*) – define meaning of a symbol

### Arguments

struct symbol \* *sym*  
     symbol to be defined

int *type*  
     symbol class to assign

void \* *def*  
     class dependent data

**Description**

Defines new meaning of a symbol. If the symbol is an undefined one (`SYM_VOID`), it's just re-defined to the new type. If it's defined in different scope, a new symbol in current scope is created and the meaning is assigned to it. If it's already defined in the current scope, an error is reported via `cf_error()`.

**Result**

Pointer to the newly defined symbol. If we are in the top-level scope, it's the same *sym* as passed to the function.

---

**Function**

`void cf_lex_init (int is_cli, struct config * c)` – initialize the lexer

**Arguments**

`int is_cli`  
true if we're going to parse CLI command, false for configuration

`struct config * c`  
configuration structure

**Description**

`cf_lex_init()` initializes the lexical analyzer and prepares it for parsing of a new input.

---

**Function**

`void cf_push_scope (struct symbol * sym)` – enter new scope

**Arguments**

`struct symbol * sym`  
symbol representing scope name

**Description**

If we want to enter a new scope to process declarations inside a nested block, we can just call `cf_push_scope()` to push a new scope onto the scope stack which will cause all new symbols to be defined in this scope and all existing symbols to be sought for in all scopes stored on the stack.

---

**Function**

`void cf_pop_scope (void)` – leave a scope

**Description**

`cf_pop_scope()` pops the topmost scope from the scope stack, leaving all its symbols in the symbol table, but making them invisible to the rest of the config.

---

**Function**

`char * cf_symbol_class_name (struct symbol * sym)` – get name of a symbol class

**Arguments**

`struct symbol * sym`  
symbol

**Description**

This function returns a string representing the class of the given symbol.

## 3.3 Parser

Both the configuration and CLI commands are analyzed using a syntax driven parser generated by the `bison` tool from a grammar which is constructed from information gathered from grammar snippets by the `gen_parser.m4` script.

Grammar snippets are files (usually with extension `.Y`) contributed by various BIRD modules in order to provide information about syntax of their configuration and their CLI commands. Each snippet consists of several sections, each of them starting with a special keyword: `CF_HDR` for a list of `#include` directives needed by the C code, `CF_DEFINES` for a list of C declarations, `CF_DECLS` for `bison` declarations including keyword definitions specified as `CF_KEYWORDS`, `CF_GRAMMAR` for the grammar rules, `CF_CODE` for auxiliary C code and finally `CF_END` at the end of the snippet.

To create references between the snippets, it's possible to define multi-part rules by utilizing the `CF_ADDTO` macro which adds a new alternative to a multi-part rule.

CLI commands are defined using a `CF_CLI` macro. Its parameters are: the list of keywords determining the command, the list of parameters, help text for the parameters and help text for the command.

Values of `enum` filter types can be defined using `CF_ENUM` with the following parameters: name of filter type, prefix common for all literals of this type and names of all the possible values.

# Chapter 4: Filters

## 4.1 Filters

You can find sources of the filter language in `filter/` directory. File `filter/config.Y` contains filter grammar and basically translates the source from user into a tree of `f_inst` structures. These trees are later interpreted using code in `filter/filter.c`.

A filter is represented by a tree of `f_inst` structures, one structure per "instruction". Each `f_inst` contains `code`, `aux` value which is usually the data type this instruction operates on and two generic arguments (`a1`, `a2`). Some instructions contain pointer(s) to other instructions in their (`a1`, `a2`) fields.

Filters use a `f_val` structure for their data. Each `f_val` contains type and value (types are constants prefixed with `T_`). Few of the types are special; `T_RETURN` can be or-ed with a type to indicate that return from a function or from the whole filter should be forced. Important thing about `f_val`'s is that they may be copied with a simple `=`. That's fine for all currently defined types: strings are read-only (and therefore okay), paths are copied for each operation (okay too).

---

### Function

int *val\_compare* (struct `f_val` *v1*, struct `f_val` *v2*) – compare two values

### Arguments

struct `f_val` *v1*  
first value

struct `f_val` *v2*  
second value

### Description

Compares two values and returns -1, 0, 1 on <, =, > or `CMP_ERROR` on error. Tree module relies on this giving consistent results so that it can be used for building balanced trees.

---

### Function

int *val\_same* (struct `f_val` *v1*, struct `f_val` *v2*) – compare two values

### Arguments

struct `f_val` *v1*  
first value

struct `f_val` *v2*  
second value

### Description

Compares two values and returns 1 if they are same and 0 if not. Comparison of values of different types is valid and returns 0.

---

### Function

int *val\_in\_range* (struct `f_val` *v1*, struct `f_val` *v2*) – implement `~` operator

### Arguments

struct `f_val` *v1*  
element

struct `f_val` *v2*  
set

### Description

Checks if *v1* is element (`~` operator) of *v2*.

---

**Function**

struct f\_val *interpret* (struct f\_inst \* *what*)

**Arguments**

struct f\_inst \* *what*  
filter to interpret

**Description**

Interpret given tree of filter instructions. This is core function of filter system and does all the hard work.

**Each instruction has 4 fields**

code (which is instruction code), aux (which is extension to instruction code, typically type), arg1 and arg2 - arguments. Depending on instruction, arguments are either integers, or pointers to instruction trees. Common instructions like +, that have two expressions as arguments use TWOARGS macro to get both of them evaluated.

f\_val structures are copied around, so there are no problems with memory management.

---

**Function**

int *f\_run* (struct filter \* *filter*, struct rte \*\* *rte*, struct ea\_list \*\* *tmp\_attrs*, struct linpool \* *tmp\_pool*, int *flags*) – run a filter for a route

**Arguments**

struct filter \* *filter*  
filter to run

struct rte \*\* *rte*  
route being filtered, may be modified

struct ea\_list \*\* *tmp\_attrs*  
temporary attributes, prepared by caller or generated by *f\_run()*

struct linpool \* *tmp\_pool*  
all filter allocations go from this pool

int *flags*  
flags

**Description**

If filter needs to modify the route, there are several possibilities. *rte* might be read-only (with REF\_COW flag), in that case rw copy is obtained by *rte\_cow()* and *rte* is replaced. If *rte* is originally rw, it may be directly modified (and it is never copied).

The returned *rte* may reuse the (possibly cached, cloned) *rta*, or (if *rta* was modified) contains a modified uncached *rta*, which uses parts allocated from *tmp\_pool* and parts shared from original *rta*. There is one exception - if *rte* is rw but contains a cached *rta* and that is modified, *rta* in returned *rte* is also cached.

Ownership of cached *rtas* is consistent with *rte*, i.e. if a new *rte* is returned, it has its own clone of cached *rta* (and cached *rta* of read-only source *rte* is intact), if *rte* is modified in place, old cached *rta* is possibly freed.

---

**Function**

int *filter\_same* (struct filter \* *new*, struct filter \* *old*) – compare two filters

**Arguments**

struct filter \* *new*  
first filter to be compared

struct filter \* *old*  
second filter to be compared, notice that this filter is damaged while comparing.

**Description**

Returns 1 in case filters are same, otherwise 0. If there are underlying bugs, it will rather say 0 on same filters than say 1 on different.



**Function**

```
struct f_tree * find_tree (struct f_tree * t, struct f_val val)
```

**Arguments**

```
struct f_tree * t
    tree to search in

struct f_val val
    value to find
```

**Description**

Search for given value in the tree. It relies on the fact that sorted tree is populated by **f\_val** structures (that can be compared by *val\_compare()*). In each node of tree, either single value (then *t->from==t->to*) or range is present.

Both set matching and *switch()* { } construction is implemented using this function, thus both are as fast as they can be.

**Function**

```
struct f_tree * build_tree (struct f_tree * from)
```

**Arguments**

```
struct f_tree * from
    degenerated tree (linked by tree->left) to be transformed into form suitable for find_tree()
```

**Description**

Transforms degenerated tree into balanced tree.

**Function**

```
int same_tree (struct f_tree * t1, struct f_tree * t2)
```

**Arguments**

```
struct f_tree * t1
    first tree to be compared

struct f_tree * t2
    second one
```

**Description**

Compares two trees and returns 1 if they are same

## 4.2 Trie for prefix sets

We use a (compressed) trie to represent prefix sets. Every node in the trie represents one prefix (**addr/plen**) and **plen** also indicates the index of the bit in the address that is used to branch at the node. If we need to represent just a set of prefixes, it would be simple, but we have to represent a set of prefix patterns. Each prefix pattern consists of **ppaddr/pplen** and two integers: **low** and **high**, and a prefix **ppaddr/plen** matches that pattern if the first  $\text{MIN}(\text{plen}, \text{pplen})$  bits of **ppaddr** and **ppaddr** are the same and  $\text{low} \leq \text{plen} \leq \text{high}$ .

We use a bitmask (**accept**) to represent accepted prefix lengths at a node. As there are 33 prefix lengths (0..32 for IPv4), but there is just one prefix of zero length in the whole trie so we have **zero** flag in **f\_trie** (indicating whether the trie accepts prefix 0.0.0.0/0) as a special case, and **accept** bitmask represents accepted prefix lengths from 1 to 32.

There are two cases in prefix matching - a match when the length of the prefix is smaller than the length of the prefix pattern, ( $\text{plen} < \text{pplen}$ ) and otherwise. The second case is simple - we just walk through the

trie and look at every visited node whether that prefix accepts our prefix length (`pplen`). The first case is tricky - we don't want to examine every descendant of a final node, so (when we create the trie) we have to propagate that information from nodes to their ascendants.

Suppose that we have two masks (M1 and M2) for a node. Mask M1 represents accepted prefix lengths by just the node and mask M2 represents accepted prefix lengths by the node or any of its descendants. Therefore M2 is a bitwise or of M1 and children's M2 and this is a maintained invariant during trie building. Basically, when we want to match a prefix, we walk through the trie, check mask M1 for our prefix length and when we came to final node, we check mask M2.

There are two differences in the real implementation. First, we use a compressed trie so there is a case that we skip our final node (if it is not in the trie) and we came to node that is either extension of our prefix, or completely out of path. In the first case, we also have to check M2.

Second, we really need not to maintain two separate bitmasks. Checks for mask M1 are always larger than `applen` and we need just the first `pplen` bits of mask M2 (if trie compression hadn't been used it would suffice to know just `$applen`-th bit), so we have to store them together in `accept` mask - the first `pplen` bits of mask M2 and then mask M1.

There are four cases when we walk through a trie:

- we are in NULL - we are out of path (prefixes are inconsistent)
- we are in the wanted (final) node (node length == `pplen`)
- we are beyond the end of path (node length > `pplen`)
- we are still on path and keep walking (node length < `pplen`)

The walking code in `trie_match_prefix()` is structured according to these cases.

### Function

`struct f_trie * f_new_trie (linpool * lp, uint node_size)` – allocates and returns a new empty trie

### Arguments

`linpool * lp`  
linear pool to allocate items from

`uint node_size`  
node size to be used (`f_trie_node` and user data)

### Function

`void * trie_add_prefix (struct f_trie * t, const net_addr * net, uint l, uint h)`

### Arguments

`struct f_trie * t`  
trie to add to

`const net_addr * net`  
IP network prefix

`uint l`  
prefix lower bound

`uint h`  
prefix upper bound

### Description

Adds prefix (prefix pattern) `n` to trie `t`. `l` and `h` are lower and upper bounds on accepted prefix lengths, both inclusive.  $0 \leq l, h \leq 32$  (128 for IPv6).

Returns a pointer to the allocated node. The function can return a pointer to an existing node if `px` and `pplen` are the same. If `px/pplen == 0/0` (or `::/0`), a pointer to the root node is returned.

---

**Function**

int *trie\_match\_net* (struct f\_trie \* *t*, const net\_addr \* *n*)

**Arguments**

struct f\_trie \* *t*  
    trie  
  
const net\_addr \* *n*  
    net address

**Description**

Tries to find a matching net in the trie such that prefix *n* matches that prefix pattern. Returns 1 if there is such prefix pattern in the trie.

---

**Function**

int *trie\_same* (struct f\_trie \* *t1*, struct f\_trie \* *t2*)

**Arguments**

struct f\_trie \* *t1*  
    first trie to be compared  
  
struct f\_trie \* *t2*  
    second one

**Description**

Compares two tries and returns 1 if they are same

---

**Function**

void *trie\_format* (struct f\_trie \* *t*, buffer \* *buf*)

**Arguments**

struct f\_trie \* *t*  
    trie to be formatted  
  
buffer \* *buf*  
    destination buffer

**Description**

Prints the trie to the supplied buffer.

# Chapter 5: Protocols

## 5.1 The Babel protocol

Babel (RFC6126) is a loop-avoiding distance-vector routing protocol that is robust and efficient both in ordinary wired networks and in wireless mesh networks.

The Babel protocol keeps state for each neighbour in a `babel_neighbor` struct, tracking received Hello and I Heard You (IHU) messages. A `babel_interface` struct keeps hello and update times for each interface, and a separate hello seqno is maintained for each interface.

For each prefix, Babel keeps track of both the possible routes (with next hop and router IDs), as well as the feasibility distance for each prefix and router id. The prefix itself is tracked in a `babel_entry` struct, while the possible routes for the prefix are tracked as `babel_route` entries and the feasibility distance is maintained through `babel_source` structures.

The main route selection is done in `babel_select_route()`. This is called when an entry is updated by receiving updates from the network or when modified by internal timers. The function selects from feasible and reachable routes the one with the lowest metric to be announced to the core.

---

### Function

void *babel\_announce\_rte* (struct babel\_proto \* *p*, struct babel\_entry \* *e*) – announce selected route to the core

### Arguments

struct babel\_proto \* *p*  
    Babel protocol instance

struct babel\_entry \* *e*  
    Babel route entry to announce

### Description

This function announces a Babel entry to the core if it has a selected incoming path, and retracts it otherwise. If there is no selected route but the entry is valid and ours, the unreachable route is announced instead.

---

### Function

void *babel\_select\_route* (struct babel\_proto \* *p*, struct babel\_entry \* *e*, struct babel\_route \* *mod*) – select best route for given route entry

### Arguments

struct babel\_proto \* *p*  
    Babel protocol instance

struct babel\_entry \* *e*  
    Babel entry to select the best route for

struct babel\_route \* *mod*  
    Babel route that was modified or NULL if unspecified

### Description

Select the best reachable and feasible route for a given prefix among the routes received from peers, and propagate it to the nest. This just selects the reachable and feasible route with the lowest metric, but keeps selected the old one in case of tie.

If no feasible route is available for a prefix that previously had a route selected, a seqno request is sent to try to get a valid route. If the entry is valid and not owned by us, the unreachable route is announced to the nest (to blackhole packets going to it, as per section 2.8). It is later removed by *babel\_expire\_routes()*. Otherwise, the route is just removed from the nest.

Argument *mod* is used to optimize best route calculation. When specified, the function can assume that only the *mod* route was modified to avoid full best route selection and announcement when non-best route

was modified in minor way. The caller is advised to not call *babel\_select\_route()* when no change is done (e.g. periodic route updates) to avoid unnecessary announcements of the same best route. The caller is not required to call the function in case of a retraction of a non-best route.

Note that the function does not active triggered updates. That is done by *babel\_rt\_notify()* when the change is propagated back to Babel.

### Function

void *babel\_send\_update\_* (struct babel\_iface \* *ifa*, btime *changed*, struct fib \* *rtable*) – send route table updates

### Arguments

struct babel\_iface \* *ifa*  
 Interface to transmit on

btime *changed*  
 Only send entries changed since this time

struct fib \* *rtable*  
 – undescribed –

### Description

This function produces update TLVs for all entries changed since the time indicated by the **changed** parameter and queues them for transmission on the selected interface. During the process, the feasibility distance for each transmitted entry is updated.

### Function

void *babel\_handle\_update* (union babel\_msg \* *m*, struct babel\_iface \* *ifa*) – handle incoming route updates

### Arguments

union babel\_msg \* *m*  
 Incoming update TLV

struct babel\_iface \* *ifa*  
 Interface the update was received on

### Description

This function is called as a handler for update TLVs and handles the updating and maintenance of route entries in Babel's internal routing cache. The handling follows the actions described in the Babel RFC, and at the end of each update handling, *babel\_select\_route()* is called on the affected entry to optionally update the selected routes and propagate them to the core.

### Function

void *babel\_iface\_timer* (timer \* *t*) – Babel interface timer handler

### Arguments

timer \* *t*  
 Timer

### Description

This function is called by the per-interface timer and triggers sending of periodic Hello's and both triggered and periodic updates. Periodic Hello's and updates are simply handled by setting the next\_{hello,regular} variables on the interface, and triggering an update (and resetting the variable) whenever 'now' exceeds that value.

For triggered updates, *babel\_trigger\_iface\_update()* will set the want\_triggered field on the interface to a timestamp value. If this is set (and the next\_triggered time has passed; this is a rate limiting mechanism), *babel\_send\_update()* will be called with this timestamp as the second parameter. This causes updates to be send consisting of only the routes that have changed since the time saved in want\_triggered.

Mostly when an update is triggered, the route being modified will be set to the value of 'now' at the time of the trigger; the >= comparison for selecting which routes to send in the update will make sure this is included.

---

**Function**

void *babel\_timer* (timer \* *t*) – global timer hook

**Arguments**

timer \* *t*  
Timer

**Description**

This function is called by the global protocol instance timer and handles expiration of routes and neighbours as well as pruning of the seqno request cache.

---

**Function**

uint *babel\_write\_queue* (struct babel\_iface \* *ifa*, list \* *queue*) – Write a TLV queue to a transmission buffer

**Arguments**

struct babel\_iface \* *ifa*  
Interface holding the transmission buffer

list \* *queue*  
TLV queue to write (containing internal-format TLVs)

**Description**

This function writes a packet to the interface transmission buffer with as many TLVs from the *queue* as will fit in the buffer. It returns the number of bytes written (NOT counting the packet header). The function is called by *babel\_send\_queue()* and *babel\_send\_unicast()* to construct packets for transmission, and uses per-TLV helper functions to convert the internal-format TLVs to their wire representations. The TLVs in the queue are freed after they are written to the buffer.

---

**Function**

void *babel\_send\_unicast* (union babel\_msg \* *msg*, struct babel\_iface \* *ifa*, ip\_addr *dest*) – send a single TLV via unicast to a destination

**Arguments**

union babel\_msg \* *msg*  
TLV to send

struct babel\_iface \* *ifa*  
Interface to send via

ip\_addr *dest*  
Destination of the TLV

**Description**

This function is used to send a single TLV via unicast to a designated receiver. This is used for replying to certain incoming requests, and for sending unicast requests to refresh routes before they expire.

---

**Function**

void *babel\_enqueue* (union babel\_msg \* *msg*, struct babel\_iface \* *ifa*) – enqueue a TLV for transmission on an interface

**Arguments**

```
union babel_msg * msg
    TLV to enqueue (in internal TLV format)

struct babel_iface * ifa
    Interface to enqueue to
```

**Description**

This function is called to enqueue a TLV for subsequent transmission on an interface. The transmission event is triggered whenever a TLV is enqueued; this ensures that TLVs will be transmitted in a timely manner, but that TLVs which are enqueued in rapid succession can be transmitted together in one packet.

**Function**

void *babel\_process\_packet* (struct babel\_pkt\_header \* *pkt*, int *len*, ip\_addr *saddr*, struct babel\_iface \* *ifa*) – process incoming data packet

**Arguments**

```
struct babel_pkt_header * pkt
    Pointer to the packet data

int len
    Length of received packet

ip_addr saddr
    Address of packet sender

struct babel_iface * ifa
    Interface packet was received on.
```

**Description**

This function is the main processing hook of incoming Babel packets. It checks that the packet header is well-formed, then processes the TLVs contained in the packet. This is done in two passes: First all TLVs are parsed into the internal TLV format. If a TLV parser fails, processing of the rest of the packet is aborted. After the parsing step, the TLV handlers are called for each parsed TLV in order.

## 5.2 Bidirectional Forwarding Detection

The BFD protocol is implemented in three files: `bfd.c` containing the protocol logic and the protocol glue with BIRD core, `packets.c` handling BFD packet processing, RX, TX and protocol sockets. `io.c` then contains generic code for the event loop, threads and event sources (sockets, microsecond timers). This generic code will be merged to the main BIRD I/O code in the future.

The BFD implementation uses a separate thread with an internal event loop for handling the protocol logic, which requires high-res and low-latency timing, so it is not affected by the rest of BIRD, which has several low-granularity hooks in the main loop, uses second-based timers and cannot offer good latency. The core of BFD protocol (the code related to BFD sessions, interfaces and packets) runs in the BFD thread, while the rest (the code related to BFD requests, BFD neighbors and the protocol glue) runs in the main thread. BFD sessions are represented by structure `bfd_session` that contains a state related to the session and two timers (TX timer for periodic packets and hold timer for session timeout). These sessions are allocated from `session_slab` and are accessible by two hash tables, `session_hash_id` (by session ID) and `session_hash_ip` (by IP addresses of neighbors). Slab and both hashes are in the main protocol structure `bfd_proto`. The protocol logic related to BFD sessions is implemented in internal functions `bfd_session_*`(), which are expected to be called from the context of BFD thread, and external functions `bfd_add_session()`, `bfd_remove_session()` and `bfd_reconfigure_session()`, which form an interface to the BFD core for the rest and are expected to be called from the context of main thread.

Each BFD session has an associated BFD interface, represented by structure `bfd_iface`. A BFD interface contains a socket used for TX (the one for RX is shared in `bfd_proto`), an interface configuration and

reference counter. Compared to interface structures of other protocols, these structures are not created and removed based on interface notification events, but according to the needs of BFD sessions. When a new session is created, it requests a proper BFD interface by function *bfd\_get\_iface()*, which either finds an existing one in *iface\_list* (from *bfd\_proto*) or allocates a new one. When a session is removed, an associated *iface* is discharged by *bfd\_free\_iface()*.

BFD requests are the external API for the other protocols. When a protocol wants a BFD session, it calls *bfd\_request\_session()*, which creates a structure *bfd\_request* containing appropriate information and an notify hook. This structure is a resource associated with the caller's resource pool. When a BFD protocol is available, a BFD request is submitted to the protocol, an appropriate BFD session is found or created and the request is attached to the session. When a session changes state, all attached requests (and related protocols) are notified. Note that BFD requests do not depend on BFD protocol running. When the BFD protocol is stopped or removed (or not available from beginning), related BFD requests are stored in *bfd\_wait\_list*, where waits for a new protocol.

BFD neighbors are just a way to statically configure BFD sessions without requests from other protocol. Structures *bfd\_neighbor* are part of BFD configuration (like static routes in the static protocol). BFD neighbors are handled by BFD protocol like it is a BFD client – when a BFD neighbor is ready, the protocol just creates a BFD request like any other protocol.

The protocol uses a new generic event loop (structure *birdloop*) from *io.c*, which supports sockets, timers and events like the main loop. A *birdloop* is associated with a thread (field *thread*) in which event hooks are executed. Most functions for setting event sources (like *sk\_start()* or *tm\_start()*) must be called from the context of that thread. *Birdloop* allows to temporarily acquire the context of that thread for the main thread by calling *birdloop\_enter()* and then *birdloop\_leave()*, which also ensures mutual exclusion with all event hooks. Note that resources associated with a *birdloop* (like timers) should be attached to the independent resource pool, detached from the main resource tree.

There are two kinds of interaction between the BFD core (running in the BFD thread) and the rest of BFD (running in the main thread). The first kind are configuration calls from main thread to the BFD thread (like *bfd\_add\_session()*). These calls are synchronous and use *birdloop\_enter()* mechanism for mutual exclusion. The second kind is a notification about session changes from the BFD thread to the main thread. This is done in an asynchronous way, sessions with pending notifications are linked (in the BFD thread) to *notify\_list* in *bfd\_proto*, and then *bfd\_notify\_hook()* in the main thread is activated using *bfd\_notify\_kick()* and a pipe. The hook then processes scheduled sessions and calls hooks from associated BFD requests. This *notify\_list* (and state fields in structure *bfd\_session*) is protected by a spinlock in *bfd\_proto* and functions *bfd\_lock\_sessions()* / *bfd\_unlock\_sessions()*.

There are few data races (accessing *p->p.debug* from *TRACE()* from the BFD thread and accessing some private fields of *bfd\_session* from *bfd\_show\_sessions()* from the main thread, but these are harmless (i hope).

TODO: document functions and access restrictions for fields in BFD structures.

Supported standards: - RFC 5880 - main BFD standard - RFC 5881 - BFD for IP links - RFC 5882 - generic application of BFD - RFC 5883 - BFD for multihop paths

## 5.3 Border Gateway Protocol

The BGP protocol is implemented in three parts: *bgp.c* which takes care of the connection and most of the interface with BIRD core, *packets.c* handling both incoming and outgoing BGP packets and *attrs.c* containing functions for manipulation with BGP attribute lists.

As opposed to the other existing routing daemons, BIRD has a sophisticated core architecture which is able to keep all the information needed by BGP in the primary routing table, therefore no complex data structures like a central BGP table are needed. This increases memory footprint of a BGP router with many connections, but not too much and, which is more important, it makes BGP much easier to implement.

Each instance of BGP (corresponding to a single BGP peer) is described by a *bgp\_proto* structure to which are attached individual connections represented by *bgp\_connection* (usually, there exists only one connection, but during BGP session setup, there can be more of them). The connections are handled according to the BGP state machine defined in the RFC with all the timers and all the parameters configurable.

In incoming direction, we listen on the connection's socket and each time we receive some input, we pass it to *bgp\_rx()*. It decodes packet headers and the markers and passes complete packets to *bgp\_rx\_packet()*



which distributes the packet according to its type.

In outgoing direction, we gather all the routing updates and sort them to buckets (`bgp_bucket`) according to their attributes (we keep a hash table for fast comparison of `rta`'s and a `fib` which helps us to find if we already have another route for the same destination queued for sending, so that we can replace it with the new one immediately instead of sending both updates). There also exists a special bucket holding all the route withdrawals which cannot be queued anywhere else as they don't have any attributes. If we have any packet to send (due to either new routes or the connection tracking code wanting to send a Open, Keepalive or Notification message), we call `bgp_schedule_packet()` which sets the corresponding bit in a `packet_to_send` bit field in `bgp_conn` and as soon as the transmit socket buffer becomes empty, we call `bgp_fire_tx()`. It inspects state of all the packet type bits and calls the corresponding `bgp_create_xx()` functions, eventually rescheduling the same packet type if we have more data of the same type to send.

The processing of attributes consists of two functions: `bgp_decode_attrs()` for checking of the attribute blocks and translating them to the language of BIRD's extended attributes and `bgp_encode_attrs()` which does the converse. Both functions are built around a `bgp_attr_table` array describing all important characteristics of all known attributes. Unknown transitive attributes are attached to the route as `EAF_TYPE_OPAQUE` byte streams.

BGP protocol implements graceful restart in both restarting (local restart) and receiving (neighbor restart) roles. The first is handled mostly by the graceful restart code in the nest, BGP protocol just handles capabilities, sets `gr_wait` and locks graceful restart until end-of-RIB mark is received. The second is implemented by internal restart of the BGP state to `BS_IDLE` and protocol state to `PS_START`, but keeping the protocol up from the core point of view and therefore maintaining received routes. Routing table refresh cycle (`rt_refresh_begin()`, `rt_refresh_end()`) is used for removing stale routes after reestablishment of BGP session during graceful restart.

Supported standards: <itemize> <item> <rfc id="4271"> - Border Gateway Protocol 4 (BGP) <item> <rfc id="1997"> - BGP Communities Attribute <item> <rfc id="2385"> - Protection of BGP Sessions via TCP MD5 Signature <item> <rfc id="2545"> - Use of BGP Multiprotocol Extensions for IPv6 <item> <rfc id="2918"> - Route Refresh Capability <item> <rfc id="3107"> - Carrying Label Information in BGP <item> <rfc id="4360"> - BGP Extended Communities Attribute <item> <rfc id="4364"> - BGP/MPLS IPv4 Virtual Private Networks <item> <rfc id="4456"> - BGP Route Reflection <item> <rfc id="4486"> - Subcodes for BGP Cease Notification Message <item> <rfc id="4659"> - BGP/MPLS IPv6 Virtual Private Networks <item> <rfc id="4724"> - Graceful Restart Mechanism for BGP <item> <rfc id="4760"> - Multiprotocol extensions for BGP <item> <rfc id="4798"> - Connecting IPv6 Islands over IPv4 MPLS <item> <rfc id="5065"> - AS confederations for BGP <item> <rfc id="5082"> - Generalized TTL Security Mechanism <item> <rfc id="5492"> - Capabilities Advertisement with BGP <item> <rfc id="5549"> - Advertising IPv4 NLRI with an IPv6 Next Hop <item> <rfc id="5575"> - Dissemination of Flow Specification Rules <item> <rfc id="5668"> - 4-Octet AS Specific BGP Extended Community <item> <rfc id="6286"> - AS-Wide Unique BGP Identifier <item> <rfc id="6608"> - Subcodes for BGP Finite State Machine Error <item> <rfc id="6793"> - BGP Support for 4-Octet AS Numbers <item> <rfc id="7313"> - Enhanced Route Refresh Capability for BGP <item> <rfc id="7606"> - Revised Error Handling for BGP UPDATE Messages <item> <rfc id="7911"> - Advertisement of Multiple Paths in BGP <item> <rfc id="7947"> - Internet Exchange BGP Route Server <item> <rfc id="8092"> - BGP Large Communities Attribute <item> <rfc id="8203"> - BGP Administrative Shutdown Communication <item> <rfc id="8212"> - Default EBGp Route Propagation Behavior without Policies </itemize>

---

## Function

`int bgp_open (struct bgp_proto * p)` – open a BGP instance

## Arguments

`struct bgp_proto * p`  
BGP instance

## Description

This function allocates and configures shared BGP resources, mainly listening sockets. Should be called as the last step during initialization (when lock is acquired and neighbor is ready). When error, caller should change state to `PS_DOWN` and return immediately.

---

**Function**

void *bgp\_close* (struct *bgp\_proto* \* *p*) – close a BGP instance

**Arguments**

struct *bgp\_proto* \* *p*  
BGP instance

**Description**

This function frees and deconfigures shared BGP resources.

---

**Function**

void *bgp\_start\_timer* (timer \* *t*, uint *value*) – start a BGP timer

**Arguments**

timer \* *t*  
timer  
  
uint *value*  
time (in seconds) to fire (0 to disable the timer)

**Description**

This functions calls *tm\_start()* on *t* with time *value* and the amount of randomization suggested by the BGP standard. Please use it for all BGP timers.

---

**Function**

void *bgp\_close\_conn* (struct *bgp\_conn* \* *conn*) – close a BGP connection

**Arguments**

struct *bgp\_conn* \* *conn*  
connection to close

**Description**

This function takes a connection described by the *bgp\_conn* structure, closes its socket and frees all resources associated with it.

---

**Function**

void *bgp\_update\_startup\_delay* (struct *bgp\_proto* \* *p*) – update a startup delay

**Arguments**

struct *bgp\_proto* \* *p*  
BGP instance

**Description**

This function updates a startup delay that is used to postpone next BGP connect. It also handles *disable\_after\_error* and might stop BGP instance when error happened and *disable\_after\_error* is on. It should be called when BGP protocol error happened.

---

**Function**

void *bgp\_handle\_graceful\_restart* (struct *bgp\_proto* \* *p*) – handle detected BGP graceful restart

**Arguments**

struct *bgp\_proto* \* *p*  
BGP instance

**Description**

This function is called when a BGP graceful restart of the neighbor is detected (when the TCP connection fails or when a new TCP connection appears). The function activates processing of the restart - starts routing table refresh cycle and activates BGP restart timer. The protocol state goes back to *PS\_START*, but changing BGP state back to *BS\_IDLE* is left for the caller.

---

**Function**

void *bgp\_graceful\_restart\_done* (struct *bgp\_channel* \* *c*) – finish active BGP graceful restart

**Arguments**

struct *bgp\_channel* \* *c*  
BGP channel

**Description**

This function is called when the active BGP graceful restart of the neighbor should be finished for channel *c* - either successfully (the neighbor sends all paths and reports end-of-RIB for given AFI/SAFI on the new session) or unsuccessfully (the neighbor does not support BGP graceful restart on the new session). The function ends the routing table refresh cycle.

---

**Function**

void *bgp\_graceful\_restart\_timeout* (timer \* *t*) – timeout of graceful restart 'restart timer'

**Arguments**

timer \* *t*  
timer

**Description**

This function is a timeout hook for *gr\_timer*, implementing BGP restart time limit for reestablishment of the BGP session after the graceful restart. When fired, we just proceed with the usual protocol restart.

---

**Function**

void *bgp\_refresh\_begin* (struct *bgp\_channel* \* *c*) – start incoming enhanced route refresh sequence

**Arguments**

struct *bgp\_channel* \* *c*  
BGP channel

**Description**

This function is called when an incoming enhanced route refresh sequence is started by the neighbor, demarcated by the BoRR packet. The function updates the load state and starts the routing table refresh cycle. Note that graceful restart also uses routing table refresh cycle, but RFC 7313 and load states ensure that these two sequences do not overlap.

---

**Function**

void *bgp\_refresh\_end* (struct *bgp\_channel* \* *c*) – finish incoming enhanced route refresh sequence

**Arguments**

struct *bgp\_channel* \* *c*  
BGP channel

**Description**

This function is called when an incoming enhanced route refresh sequence is finished by the neighbor, demarcated by the EoRR packet. The function updates the load state and ends the routing table refresh cycle. Routes not received during the sequence are removed by the nest.

---

**Function**

void *bgp\_connect* (struct *bgp\_proto* \* *p*) – initiate an outgoing connection

**Arguments**

struct *bgp\_proto* \* *p*  
BGP instance

**Description**

The *bgp\_connect()* function creates a new **bgp\_conn** and initiates a TCP connection to the peer. The rest of connection setup is governed by the BGP state machine as described in the standard.

---

**Function**

struct *bgp\_proto* \* *bgp\_find\_proto* (sock \* *sk*) – find existing proto for incoming connection

**Arguments**

sock \* *sk*  
TCP socket

---

**Function**

int *bgp\_incoming\_connection* (sock \* *sk*, uint dummy *UNUSED*) – handle an incoming connection

**Arguments**

sock \* *sk*  
TCP socket

uint dummy *UNUSED*  
– undescribed –

---

**Description**

This function serves as a socket hook for accepting of new BGP connections. It searches a BGP instance corresponding to the peer which has connected and if such an instance exists, it creates a **bgp\_conn** structure, attaches it to the instance and either sends an Open message or (if there already is an active connection) it closes the new connection by sending a Notification message.

---

**Function**

void *bgp\_error* (struct *bgp\_conn* \* *c*, uint *code*, uint *subcode*, byte \* *data*, int *len*) – report a protocol error

**Arguments**

struct *bgp\_conn* \* *c*  
connection

uint *code*  
error code (according to the RFC)

uint *subcode*  
error sub-code

byte \* *data*  
data to be passed in the Notification message

int *len*  
length of the data

**Description**

*bgp\_error()* sends a notification packet to tell the other side that a protocol error has occurred (including the data considered erroneous if possible) and closes the connection.

**Function**

void *bgp\_store\_error* (struct bgp\_proto \* *p*, struct bgp\_conn \* *c*, u8 *class*, u32 *code*) – store last error for status report

**Arguments**

struct bgp\_proto \* *p*  
     BGP instance

struct bgp\_conn \* *c*  
     connection

u8 *class*  
     error class (BE\_xxx constants)

u32 *code*  
     error code (class specific)

**Description**

*bgp\_store\_error()* decides whether given error is interesting enough and store that error to last\_error variables of *p*

**Function**

int *bgp\_fire\_tx* (struct bgp\_conn \* *conn*) – transmit packets

**Arguments**

struct bgp\_conn \* *conn*  
     connection

**Description**

Whenever the transmit buffers of the underlying TCP connection are free and we have any packets queued for sending, the socket functions call *bgp\_fire\_tx()* which takes care of selecting the highest priority packet queued (Notification > Keepalive > Open > Update), assembling its header and body and sending it to the connection.

**Function**

void *bgp\_schedule\_packet* (struct bgp\_conn \* *conn*, struct bgp\_channel \* *c*, int *type*) – schedule a packet for transmission

**Arguments**

struct bgp\_conn \* *conn*  
     connection

struct bgp\_channel \* *c*  
     channel

int *type*  
     packet type

**Description**

Schedule a packet of type *type* to be sent as soon as possible.

---

**Function**

const char \* *bgp\_error\_dsc* (uint *code*, uint *subcode*) – return BGP error description

**Arguments**

uint *code*  
BGP error code

uint *subcode*  
BGP error subcode

**Description**

*bgp\_error\_dsc()* returns error description for BGP errors which might be static string or given temporary buffer.

---

**Function**

void *bgp\_rx\_packet* (struct bgp\_conn \* *conn*, byte \* *pkt*, uint *len*) – handle a received packet

**Arguments**

struct bgp\_conn \* *conn*  
BGP connection

byte \* *pkt*  
start of the packet

uint *len*  
packet size

**Description**

*bgp\_rx\_packet()* takes a newly received packet and calls the corresponding packet handler according to the packet type.

---

**Function**

int *bgp\_rx* (sock \* *sk*, uint *size*) – handle received data

**Arguments**

sock \* *sk*  
socket

uint *size*  
amount of data received

**Description**

*bgp\_rx()* is called by the socket layer whenever new data arrive from the underlying TCP connection. It assembles the data fragments to packets, checks their headers and framing and passes complete packets to *bgp\_rx\_packet()*.

---

**Function**

ea\_list \* *bgp\_export\_attrs* (struct bgp\_export\_state \* *s*, ea\_list \* *attrs*) – export BGP attributes

**Arguments**

struct bgp\_export\_state \* *s*  
BGP export state

ea\_list \* *attrs*  
a list of extended attributes

**Description**

The *bgp\_export\_attrs()* function takes a list of attributes and merges it to one newly allocated and sorted segment. Attributes are validated and normalized by type-specific export hooks and attribute flags are updated. Some attributes may be eliminated (e.g. unknown non-transitive attributes, or empty community sets).

**Result**

one sorted attribute list segment, or NULL if attributes are unsuitable.

**Function**

int *bgp\_encode\_attrs* (struct *bgp\_write\_state* \* *s*, ea\_list \* *attrs*, byte \* *buf*, byte \* *end*) – encode BGP attributes

**Arguments**

struct *bgp\_write\_state* \* *s*  
     BGP write state

ea\_list \* *attrs*  
     a list of extended attributes

byte \* *buf*  
     buffer

byte \* *end*  
     buffer end

**Description**

The *bgp\_encode\_attrs()* function takes a list of extended attributes and converts it to its BGP representation (a part of an Update message).

**Result**

Length of the attribute block generated or -1 if not enough space.

**Function**

ea\_list \* *bgp\_decode\_attrs* (struct *bgp\_parse\_state* \* *s*, byte \* *data*, uint *len*) – check and decode BGP attributes

**Arguments**

struct *bgp\_parse\_state* \* *s*  
     BGP parse state

byte \* *data*  
     start of attribute block

uint *len*  
     length of attribute block

**Description**

This function takes a BGP attribute block (a part of an Update message), checks its consistency and converts it to a list of BIRD route attributes represented by an (uncached) *rta*.

## 5.4 Open Shortest Path First (OSPF)

The OSPF protocol is quite complicated and its complex implementation is split to many files. In *ospf.c*, you will find mainly the interface for communication with the core (e.g., reconfiguration hooks, shutdown and initialisation and so on). File *iface.c* contains the interface state machine and functions for allocation and deallocation of OSPF's interface data structures. Source *neighbor.c* includes the neighbor state machine and functions for election of Designated Router and Backup Designated router. In *packet.c*, you

will find various functions for sending and receiving generic OSPF packets. There are also routines for authentication and checksumming. In `hello.c`, there are routines for sending and receiving of hello packets as well as functions for maintaining wait times and the inactivity timer. Files `lsreq.c`, `lsack.c`, `dbdes.c` contain functions for sending and receiving of link-state requests, link-state acknowledgements and database descriptions respectively. In `lsupd.c`, there are functions for sending and receiving of link-state updates and also the flooding algorithm. Source `topology.c` is a place where routines for searching LSAs in the link-state database, adding and deleting them reside, there also are functions for originating of various types of LSAs (router LSA, net LSA, external LSA). File `rt.c` contains routines for calculating the routing table. `lsalib.c` is a set of various functions for working with the LSAs (endianity conversions, calculation of checksum etc.). One instance of the protocol is able to hold LSA databases for multiple OSPF areas, to exchange routing information between multiple neighbors and to calculate the routing tables. The core structure is `ospf_proto` to which multiple `ospf_area` and `ospf_iface` structures are connected. `ospf_proto` is also connected to `top_hash_graph` which is a dynamic hashing structure that describes the link-state database. It allows fast search, addition and deletion. Each LSA is kept in two pieces: header and body. Both of them are kept in the endianity of the CPU.

In OSPFv2 specification, it is implied that there is one IP prefix for each physical network/interface (unless it is an ptp link). But in modern systems, there might be more independent IP prefixes associated with an interface. To handle this situation, we have one `ospf_iface` for each active IP prefix (instead for each active iface); This behaves like virtual interface for the purpose of OSPF. If we receive packet, we associate it with a proper virtual interface mainly according to its source address.

OSPF keeps one socket per `ospf_iface`. This allows us (compared to one socket approach) to evade problems with a limit of multicast groups per socket and with sending multicast packets to appropriate interface in a portable way. The socket is associated with underlying physical iface and should not receive packets received on other ifaces (unfortunately, this is not true on BSD). Generally, one packet can be received by more sockets (for example, if there are more `ospf_iface` on one physical iface), therefore we explicitly filter received packets according to src/dst IP address and received iface.

Vlinks are implemented using particularly degenerate form of `ospf_iface`, which has several exceptions: it does not have its iface or socket (it copies these from 'parent' `ospf_iface`) and it is present in iface list even when down (it is not freed in `ospf_iface_down()`).

The heart beat of ospf is `ospf_disp()`. It is called at regular intervals (`ospf_proto->tick`). It is responsible for aging and flushing of LSAs in the database, updating topology information in LSAs and for routing table calculation.

To every `ospf_iface`, we connect one or more `ospf_neighbor`'s – a structure containing many timers and queues for building adjacency and for exchange of routing messages.

BIRD's OSPF implementation respects RFC2328 in every detail, but some of internal algorithms do differ. The RFC recommends making a snapshot of the link-state database when a new adjacency is forming and sending the database description packets based on the information in this snapshot. The database can be quite large in some networks, so rather we walk through a `slist` structure which allows us to continue even if the actual LSA we were working with is deleted. New LSAs are added at the tail of this `slist`.

We also do not keep a separate OSPF routing table, because the core helps us by being able to recognize when a route is updated to an identical one and it suppresses the update automatically. Due to this, we can flush all the routes we have recalculated and also those we have deleted to the core's routing table and the core will take care of the rest. This simplifies the process and conserves memory.

Supported standards: - RFC 2328 - main OSPFv2 standard - RFC 5340 - main OSPFv3 standard - RFC 3101 - OSPFv2 NSSA areas - RFC 5709 - OSPFv2 HMAC-SHA Cryptographic Authentication - RFC 5838 - OSPFv3 Support of Address Families - RFC 6549 - OSPFv2 Multi-Instance Extensions - RFC 6987 - OSPF Stub Router Advertisement

## Function

`void ospf_disp (timer * timer)` – invokes routing table calculation, aging and also `area_disp()`

## Arguments

`timer * timer`

timer usually called every `ospf_proto->tick` second, `timer->data` point to `ospf_proto`



**Function**

int *ospf\_import\_control* (struct proto \* *P*, rte \*\* *new*, ea\_list \*\* *attrs*, struct linpool \* *pool*) – accept or reject new route from nest's routing table

**Arguments**

struct proto \* *P*  
     OSPF protocol instance

rte \*\* *new*  
     the new route

ea\_list \*\* *attrs*  
     list of attributes

struct linpool \* *pool*  
     pool for allocation of attributes

**Description**

Its quite simple. It does not accept our own routes and leaves the decision on import to the filters.

**Function**

int *ospf\_shutdown* (struct proto \* *P*) – Finish of OSPF instance

**Arguments**

struct proto \* *P*  
     OSPF protocol instance

**Description**

RFC does not define any action that should be taken before router shutdown. To make my neighbors react as fast as possible, I send them hello packet with empty neighbor list. They should start their neighbor state machine with event NEIGHBOR\_1WAY.

**Function**

int *ospf\_reconfigure* (struct proto \* *P*, struct proto\_config \* *CF*) – reconfiguration hook

**Arguments**

struct proto \* *P*  
     current instance of protocol (with old configuration)

struct proto\_config \* *CF*  
     – undescribed –

**Description**

This hook tries to be a little bit intelligent. Instance of OSPF will survive change of many constants like hello interval, password change, addition or deletion of some neighbor on nonbroadcast network, cost of interface, etc.

**Function**

struct top\_hash\_entry \* *ospf\_install\_lsa* (struct ospf\_proto \* *p*, struct ospf\_lsa\_header \* *lsa*, u32 *type*, u32 *domain*, void \* *body*) – install new LSA into database

**Arguments**

struct ospf\_proto \* *p*  
OSPF protocol instance

struct ospf\_lsa\_header \* *lsa*  
LSA header

u32 *type*  
type of LSA

u32 *domain*  
domain of LSA

void \* *body*  
pointer to LSA body

**Description**

This function ensures installing new LSA received in LS update into LSA database. Old instance is replaced. Several actions are taken to detect if new routing table calculation is necessary. This is described in 13.2 of RFC 2328. This function is for received LSA only, locally originated LSAs are installed by *ospf\_originate\_lsa()*.

The LSA body in *body* is expected to be mb.allocated by the caller and its ownership is transferred to the LSA entry structure.

**Function**

void *ospf\_advance\_lsa* (struct ospf\_proto \* *p*, struct top\_hash\_entry \* *en*, struct ospf\_lsa\_header \* *lsa*, u32 *type*, u32 *domain*, void \* *body*) – handle received unexpected self-originated LSA

**Arguments**

struct ospf\_proto \* *p*  
OSPF protocol instance

struct top\_hash\_entry \* *en*  
current LSA entry or NULL

struct ospf\_lsa\_header \* *lsa*  
new LSA header

u32 *type*  
type of LSA

u32 *domain*  
domain of LSA

void \* *body*  
pointer to LSA body

**Description**

This function handles received unexpected self-originated LSA (*lsa*, *body*) by either advancing sequence number of the local LSA instance (*en*) and propagating it, or installing the received LSA and immediately flushing it (if there is no local LSA; i.e., *en* is NULL or MaxAge).

The LSA body in *body* is expected to be mb.allocated by the caller and its ownership is transferred to the LSA entry structure or it is freed.

**Function**

struct top\_hash\_entry \* *ospf\_originate\_lsa* (struct ospf\_proto \* *p*, struct ospf\_new\_lsa \* *lsa*) – originate new LSA

**Arguments**

```
struct ospf_proto * p
    OSPF protocol instance

struct ospf_new_lsa * lsa
    New LSA specification
```

**Description**

This function prepares a new LSA, installs it into the LSA database and floods it. If the new LSA cannot be originated now (because the old instance was originated within `MinLSInterval`, or because the LSA seqnum is currently wrapping), the origination is instead scheduled for later. If the new LSA is equivalent to the current LSA, the origination is skipped. In all cases, the corresponding LSA entry is returned. The new LSA is based on the LSA specification (*lsa*) and the LSA body from `lsab` buffer of *p*, which is emptied after the call. The opposite of this function is *ospf\_flush\_lsa()*.

**Function**

void *ospf\_flush\_lsa* (struct ospf\_proto \* *p*, struct top\_hash\_entry \* *en*) – flush LSA from OSPF domain

**Arguments**

```
struct ospf_proto * p
    OSPF protocol instance

struct top_hash_entry * en
    LSA entry to flush
```

**Description**

This function flushes *en* from the OSPF domain by setting its age to `LSA_MAXAGE` and flooding it. That also triggers subsequent events in LSA lifecycle leading to removal of the LSA from the LSA database (e.g. the LSA content is freed when flushing is acknowledged by neighbors). The function does nothing if the LSA is already being flushed. LSA entries are not immediately removed when being flushed, the caller may assume that *en* still exists after the call. The function is the opposite of *ospf\_originate\_lsa()* and is supposed to do the right thing even in cases of postponed origination.

**Function**

void *ospf\_update\_lsadb* (struct ospf\_proto \* *p*) – update LSA database

**Arguments**

```
struct ospf_proto * p
    OSPF protocol instance
```

**Description**

This function is periodically invoked from *ospf\_disp()*. It does some periodic or postponed processing related to LSA entries. It originates postponed LSAs scheduled by *ospf\_originate\_lsa()*. It continues in flushing processes started by *ospf\_flush\_lsa()*. It also periodically refreshes locally originated LSAs – when the current instance is older `LSREFRESHTIME`, a new instance is originated. Finally, it also ages stored LSAs and flushes ones that reached `LSA_MAXAGE`.

The RFC 2328 says that a router should periodically check checksums of all stored LSAs to detect hardware problems. This is not implemented.

**Function**

void *ospf\_originate\_ext\_lsa* (struct ospf\_proto \* *p*, struct ospf\_area \* *oa*, ort \* *nf*, u8 *mode*, u32 *metric*, u32 *ebit*, ip\_addr *fwaddr*, u32 *tag*, int *pbit*) – new route received from nest and filters

**Arguments**

struct ospf\_proto \* *p*  
 OSPF protocol instance

struct ospf\_area \* *oa*  
 ospf\_area for which LSA is originated

ort \* *nf*  
 network prefix and mask

u8 *mode*  
 the mode of the LSA (LSA\_M\_EXPORT or LSA\_M\_RTCLALC)

u32 *metric*  
 the metric of a route

u32 *ebit*  
 E-bit for route metric (bool)

ip\_addr *fwaddr*  
 the forwarding address

u32 *tag*  
 the route tag

int *pbit*  
 P-bit for NSSA LSAs (bool), ignored for external LSAs

**Description**

If I receive a message that new route is installed, I try to originate an external LSA. If *oa* is an NSSA area, NSSA-LSA is originated instead. *oa* should not be a stub area. *src* does not specify whether the LSA is external or NSSA, but it specifies the source of origination - the export from *ospf\_rt\_notify()*, or the NSSA-EXT translation.

**Function**

struct top\_graph \* *ospf\_top\_new* (struct ospf\_proto \* *p*, pool \* *pool*) – allocated new topology database

**Arguments**

struct ospf\_proto \* *p*  
 OSPF protocol instance

pool \* *pool*  
 pool for allocation

**Description**

This dynamically hashed structure is used for keeping LSAs. Mainly it is used for the LSA database of the OSPF protocol, but also for LSA retransmission and request lists of OSPF neighbors.

**Function**

void *ospf\_neigh\_chstate* (struct ospf\_neighbor \* *n*, u8 *state*) – handles changes related to new or lod state of neighbor

**Arguments**

struct ospf\_neighbor \* *n*  
 OSPF neighbor

u8 *state*  
 new state

**Description**

Many actions have to be taken according to a change of state of a neighbor. It starts rxmt timers, call interface state machine etc.

---

**Function**

void *ospf\_neigh\_sm* (struct *ospf\_neighbor* \* *n*, int *event*) – ospf neighbor state machine

**Arguments**

struct *ospf\_neighbor* \* *n*  
neighbor

int *event*  
actual event

**Description**

This part implements the neighbor state machine as described in 10.3 of RFC 2328. The only difference is that state `NEIGHBOR_ATTEMPT` is not used. We discover neighbors on nonbroadcast networks in the same way as on broadcast networks. The only difference is in sending hello packets. These are sent to IPs listed in *ospf\_iface->nbma.list* .

---

**Function**

void *ospf\_dr\_election* (struct *ospf\_iface* \* *ifa*) – (Backup) Designated Router election

**Arguments**

struct *ospf\_iface* \* *ifa*  
actual interface

**Description**

When the wait timer fires, it is time to elect (Backup) Designated Router. Structure describing me is added to this list so every electing router has the same list. Backup Designated Router is elected before Designated Router. This process is described in 9.4 of RFC 2328. The function is supposed to be called only from *ospf\_iface\_sm()* as a part of the interface state machine.

---

**Function**

void *ospf\_iface\_chstate* (struct *ospf\_iface* \* *ifa*, u8 *state*) – handle changes of interface state

**Arguments**

struct *ospf\_iface* \* *ifa*  
OSPF interface

u8 *state*  
new state

**Description**

Many actions must be taken according to interface state changes. New network LSAs must be originated, flushed, new multicast sockets to listen for messages for ALLDRouters have to be opened, etc.

---

**Function**

void *ospf\_iface\_sm* (struct *ospf\_iface* \* *ifa*, int *event*) – OSPF interface state machine

**Arguments**

struct *ospf\_iface* \* *ifa*  
OSPF interface

int *event*  
event coming to state machine

**Description**

This fully respects 9.3 of RFC 2328 except we have slightly different handling of DOWN and LOOP state. We remove interfaces that are DOWN. DOWN state is used when an interface is waiting for a link. LOOP state is used when an interface does not have a link.

---

**Function**

int *ospf\_rx\_hook* (sock \* *sk*, uint *len*)

**Arguments**

sock \* *sk*  
     socket we received the packet.

uint *len*  
     length of the packet

**Description**

This is the entry point for messages from neighbors. Many checks (like authentication, checksums, size) are done before the packet is passed to non generic functions.

---

**Function**

int *lsa\_validate* (struct ospf\_lsa\_header \* *lsa*, u32 *lsa\_type*, int *ospf2*, void \* *body*) – check whether given LSA is valid

**Arguments**

struct ospf\_lsa\_header \* *lsa*  
     LSA header

u32 *lsa\_type*  
     internal LSA type (LSA\_T\_XXX)

int *ospf2*  
     true for OSPFv2, false for OSPFv3

void \* *body*  
     pointer to LSA body

**Description**

Checks internal structure of given LSA body (minimal length, consistency). Returns true if valid.

---

**Function**

void *ospf\_send\_dbdes* (struct ospf\_proto \* *p*, struct ospf\_neighbor \* *n*) – transmit database description packet

**Arguments**

struct ospf\_proto \* *p*  
     OSPF protocol instance

struct ospf\_neighbor \* *n*  
     neighbor

**Description**

Sending of a database description packet is described in 10.8 of RFC 2328. Reception of each packet is acknowledged in the sequence number of another. When I send a packet to a neighbor I keep a copy in a buffer. If the neighbor does not reply, I don't create a new packet but just send the content of the buffer.

---

**Function**

void *ospf\_rt\_spf* (struct ospf\_proto \* *p*) – calculate internal routes

**Arguments**

struct ospf\_proto \* *p*  
     OSPF protocol instance

**Description**

Calculation of internal paths in an area is described in 16.1 of RFC 2328. It's based on Dijkstra's shortest path tree algorithms. This function is invoked from *ospf\_disp()*.

## 5.5 Pipe

The Pipe protocol is very simple. It just connects to two routing tables using *proto\_add\_announce\_hook()* and whenever it receives a *rt\_notify()* about a change in one of the tables, it converts it to a *rte\_update()* in the other one.

To avoid pipe loops, Pipe keeps a 'being updated' flag in each routing table.

A pipe has two announce hooks, the first connected to the main table, the second connected to the peer table. When a new route is announced on the main table, it gets checked by an export filter in ahook 1, and, after that, it is announced to the peer table via *rte\_update()*, an import filter in ahook 2 is called. When a new route is announced in the peer table, an export filter in ahook2 and an import filter in ahook 1 are used. Obviously, there is no need in filtering the same route twice, so both import filters are set to accept, while user configured 'import' and 'export' filters are used as export filters in ahooks 2 and 1. Route limits are handled similarly, but on the import side of ahooks.

## 5.6 Router Advertisements

The RAdv protocol is implemented in two files: **radv.c** containing the interface with BIRD core and the protocol logic and **packets.c** handling low level protocol stuff (RX, TX and packet formats). The protocol does not export any routes.

The RAdv is structured in the usual way - for each handled interface there is a structure **radv\_iface** that contains a state related to that interface together with its resources (a socket, a timer). There is also a prepared RA stored in a TX buffer of the socket associated with an iface. These iface structures are created and removed according to iface events from BIRD core handled by *radv\_if\_notify()* callback.

The main logic of RAdv consists of two functions: *radv\_iface\_notify()*, which processes asynchronous events (specified by RA\_EV\_\* codes), and *radv\_timer()*, which triggers sending RAs and computes the next timeout.

The RAdv protocol could receive routes (through *radv\_import\_control()* and *radv\_rt\_notify()*), but only the configured trigger route is tracked (in **active** var). When a radv protocol is reconfigured, the connected routing table is examined (in *radv\_check\_active()*) to have proper **active** value in case of the specified trigger prefix was changed.

Supported standards: - RFC 4861 - main RA standard - RFC 4191 - Default Router Preferences and More-Specific Routes - RFC 6106 - DNS extensions (RDDNS, DNSSL)

## 5.7 Routing Information Protocol (RIP)

The RIP protocol is implemented in two files: **rip.c** containing the protocol logic, route management and the protocol glue with BIRD core, and **packets.c** handling RIP packet processing, RX, TX and protocol sockets.

Each instance of RIP is described by a structure **rip\_proto**, which contains an internal RIP routing table, a list of protocol interfaces and the main timer responsible for RIP routing table cleanup.

RIP internal routing table contains incoming and outgoing routes. For each network (represented by structure **rip\_entry**) there is one outgoing route stored directly in **rip\_entry** and an one-way linked list of incoming routes (structures **rip\_rte**). The list contains incoming routes from different RIP neighbors, but only routes with the lowest metric are stored (i.e., all stored incoming routes have the same metric).

Note that RIP itself does not select outgoing route, that is done by the core routing table. When a new incoming route is received, it is propagated to the RIP table by *rip\_update\_rte()* and possibly stored in the list of incoming routes. Then the change may be propagated to the core by *rip\_announce\_rte()*. The core selects the best route and propagate it to RIP by *rip\_rt\_notify()*, which updates outgoing route part of **rip\_entry** and possibly triggers route propagation by *rip\_trigger\_update()*.

RIP interfaces are represented by structures **rip\_iface**. A RIP interface contains a per-interface socket, a list of associated neighbors, interface configuration, and state information related to scheduled interface events and running update sessions. RIP interfaces are added and removed based on core interface notifications.

There are two RIP interface events - regular updates and triggered updates. Both are managed from the RIP interface timer (*rip\_iface\_timer()*). Regular updates are called at fixed interval and propagate the whole routing table, while triggered updates are scheduled by *rip\_trigger\_update()* due to some routing table change

and propagate only the routes modified since the time they were scheduled. There are also unicast-destined requested updates, but these are sent directly as a reaction to received RIP request message. The update session is started by *rip\_send\_table()*. There may be at most one active update session per interface, as the associated state (including the fib iterator) is stored directly in *rip\_iface* structure.

RIP neighbors are represented by structures *rip\_neighbor*. Compared to neighbor handling in other routing protocols, RIP does not have explicit neighbor discovery and adjacency maintenance, which makes the *rip\_neighbor* related code a bit peculiar. RIP neighbors are interlinked with core neighbor structures (*neighbor*) and use core neighbor notifications to ensure that RIP neighbors are timely removed. RIP neighbors are added based on received route notifications and removed based on core neighbor and RIP interface events.

RIP neighbors are linked by RIP routes and use counter to track the number of associated routes, but when these RIP routes timeout, associated RIP neighbor is still alive (with zero counter). When RIP neighbor is removed but still has some associated routes, it is not freed, just changed to detached state (core neighbors and RIP ifaces are unlinked), then during the main timer cleanup phase the associated routes are removed and the *rip\_neighbor* structure is finally freed.

Supported standards: - RFC 1058 - RIPv1 - RFC 2453 - RIPv2 - RFC 2080 - RIPvng - RFC 4822 - RIP cryptographic authentication

### Function

void *rip\_announce\_rte* (struct *rip\_proto* \* *p*, struct *rip\_entry* \* *en*) – announce route from RIP routing table to the core

### Arguments

struct *rip\_proto* \* *p*  
RIP instance

struct *rip\_entry* \* *en*  
related network

### Description

The function takes a list of incoming routes from *en*, prepare appropriate *rte* for the core and propagate it by *rte\_update()*.

### Function

void *rip\_update\_rte* (struct *rip\_proto* \* *p*, net\_addr \* *n*, struct *rip\_rte* \* *new*) – enter a route update to RIP routing table

### Arguments

struct *rip\_proto* \* *p*  
RIP instance

net\_addr \* *n*  
– undscribed –

struct *rip\_rte* \* *new*  
a *rip\_rte* representing the new route

### Description

The function is called by the RIP packet processing code whenever it receives a reachable route. The appropriate routing table entry is found and the list of incoming routes is updated. Eventually, the change is also propagated to the core by *rip\_announce\_rte()*. Note that for unreachable routes, *rip\_withdraw\_rte()* should be called instead of *rip\_update\_rte()*.



---

**Function**

void *rip\_withdraw\_rte* (struct rip\_proto \* *p*, net\_addr \* *n*, struct rip\_neighbor \* *from*) – enter a route withdraw to RIP routing table

**Arguments**

struct rip\_proto \* *p*  
     RIP instance

net\_addr \* *n*  
     – undescribed –

struct rip\_neighbor \* *from*  
     a **rip\_neighbor** propagating the withdraw

**Description**

The function is called by the RIP packet processing code whenever it receives an unreachable route. The incoming route for given network from nbr *from* is removed. Eventually, the change is also propagated by *rip\_announce\_rte()*.

---

**Function**

void *rip\_timer* (timer \* *t*) – RIP main timer hook

**Arguments**

timer \* *t*  
     timer

**Description**

The RIP main timer is responsible for routing table maintenance. Invalid or expired routes (**rip\_rte**) are removed and garbage collection of stale routing table entries (**rip\_entry**) is done. Changes are propagated to core tables, route reload is also done here. Note that garbage collection uses a maximal GC time, while interfaces maintain an illusion of per-interface GC times in *rip\_send\_response()*.

Keeping incoming routes and the selected outgoing route are two independent functions, therefore after garbage collection some entries now considered invalid (RIP\_ENTRY\_DUMMY) still may have non-empty list of incoming routes, while some valid entries (representing an outgoing route) may have that list empty. The main timer is not scheduled periodically but it uses the time of the current next event and the minimal interval of any possible event to compute the time of the next run.

---

**Function**

void *rip\_iface\_timer* (timer \* *t*) – RIP interface timer hook

**Arguments**

timer \* *t*  
     timer

**Description**

RIP interface timers are responsible for scheduling both regular and triggered updates. Fixed, delay-independent period is used for regular updates, while minimal separating interval is enforced for triggered updates. The function also ensures that a new update is not started when the old one is still running.

---

**Function**

void *rip\_send\_table* (struct rip\_proto \* *p*, struct rip\_iface \* *ifa*, ip\_addr *addr*, btime *changed*) – RIP interface timer hook

**Arguments**

```

struct rip_proto * p
    RIP instance

struct rip_iface * ifa
    RIP interface

ip_addr addr
    destination IP address

btime changed
    time limit for triggered updates

```

**Description**

The function activates an update session and starts sending routing update packets (using *rip\_send\_response()*). The session may be finished during the call or may continue in *rip\_tx\_hook()* until all appropriate routes are transmitted. Note that there may be at most one active update session per interface, the function will terminate the old active session before activating the new one.

## 5.8 RPKI To Router (RPKI-RTR)

The RPKI-RTR protocol is implemented in several files: `rpki.c` containing the routes handling, protocol logic, timer events, cache connection, reconfiguration, configuration and protocol glue with BIRD core, `packets.c` containing the RPKI packets handling and finally all transports files: `transport.c`, `tcp_transport.c` and `ssh_transport.c`.

The `transport.c` is a middle layer and interface for each specific transport. Transport is a way how to wrap a communication with a cache server. There is supported an unprotected TCP transport and an encrypted SSHv2 transport. The SSH transport requires LibSSH library. LibSSH is loading dynamically using *dlopen()* function. SSH support is integrated in `sysdep/unix/io.c`. Each transport must implement an initialization function, an open function and a socket identification function. That's all.

This implementation is based on the RTRlib (<http://rpki.realmv6.org/>). The BIRD takes over files `packets.c`, `rtr.c` (inside `rpki.c`), `transport.c`, `tcp_transport.c` and `ssh_transport.c` from RTRlib.

A RPKI-RTR connection is described by a structure `rpki_cache`. The main logic is located in *rpki\_cache\_change\_state()* function. There is a state machine. The standard starting state flow looks like **Down** > **Connecting** > **Sync-Start** > **Sync-Running** > **Established** and then the last three states are periodically repeated.

**Connecting** state establishes the transport connection. The state from a call *rpki\_cache\_change\_state(CONNECTING)* to a call *rpki\_connected\_hook()*

**Sync-Start** state starts with sending **Reset Query** or **Serial Query** and then waits for **Cache Response**. The state from *rpki\_connected\_hook()* to *rpki\_handle\_cache\_response\_pdu()*

During **Sync-Running** BIRD receives data with IPv4/IPv6 Prefixes from cache server. The state starts from *rpki\_handle\_cache\_response\_pdu()* and ends in *rpki\_handle\_end\_of\_data\_pdu()*.

**Established** state means that BIRD has synced all data with cache server. Schedules a refresh timer event that invokes **Sync-Start**. Schedules Expire timer event and stops a Retry timer event.

**Transport Error** state means that we have some troubles with a network connection. We cannot connect to a cache server or we wait too long for some expected PDU for received - **Cache Response** or **End of Data**. It closes current connection and schedules a Retry timer event.

**Fatal Protocol Error** is occurred e.g. by received a bad Session ID. We restart a protocol, so all ROAs are flushed immediately.

The RPKI-RTR protocol (RFC 6810 bis) defines configurable refresh, retry and expire intervals. For maintaining a connection are used timer events that are scheduled by *rpki\_schedule\_next\_refresh()*, *rpki\_schedule\_next\_retry()* and *rpki\_schedule\_next\_expire()* functions.

A Refresh timer event performs a sync of **Established** connection. So it shifts state to **Sync-Start**. If at the beginning of second call of a refresh event is connection in **Sync-Start** state then we didn't receive a

Cache Response from a cache server and we invoke **Transport Error** state.

A Retry timer event attempts to connect cache server. It is activated after **Transport Error** state and terminated by reaching **Established** state. If cache connection is still connecting to the cache server at the beginning of an event call then the Retry timer event invokes **Transport Error** state.

An Expire timer event checks expiration of ROAs. If a last successful sync was more ago than the expire interval then the Expire timer event invokes a protocol restart thereby removes all ROAs learned from that cache server and continue trying to connect to cache server. The Expire event is activated by initial successful loading of ROAs, receiving End of Data PDU.

A reconfiguration of cache connection works well without restarting when we change only intervals values.

Supported standards: - RFC 6810 - main RPKI-RTR standard - RFC 6810 bis - an explicit timing parameters and protocol version number negotiation

### Function

const char \* *rpki\_cache\_state\_to\_str* (enum rpki\_cache\_state *state*) – give a text representation of cache state

### Arguments

enum rpki\_cache\_state *state*  
A cache state

### Description

The function converts logic cache state into string.

### Function

void *rpki\_start\_cache* (struct rpki\_cache \* *cache*) – connect to a cache server

### Arguments

struct rpki\_cache \* *cache*  
RPKI connection instance

### Description

This function is a high level method to kick up a connection to a cache server.

### Function

void *rpki\_force\_restart\_proto* (struct rpki\_proto \* *p*) – force shutdown and start protocol again

### Arguments

struct rpki\_proto \* *p*  
RPKI protocol instance

### Description

This function calls shutdown and frees all protocol resources as well. After calling this function should be no operations with protocol data, they could be freed already.

### Function

void *rpki\_cache\_change\_state* (struct rpki\_cache \* *cache*, const enum rpki\_cache\_state *new\_state*) – check and change cache state

### Arguments

struct rpki\_cache \* *cache*  
RPKI cache instance  
  
const enum rpki\_cache\_state *new\_state*  
suggested new state

### Description

This function makes transitions between internal states. It represents the core of logic management of RPKI protocol. Cannot transit into the same state as cache is in already.

---

**Function**

void *rpki\_refresh\_hook* (timer \* *tm*) – control a scheduling of downloading data from cache server

**Arguments**

timer \* *tm*  
refresh timer with cache connection instance in data

**Description**

This function is periodically called during **ESTABLISHED** or **SYNC\*** state cache connection. The first refresh schedule is invoked after receiving a **End of Data** PDU and has run by some **ERROR** is occurred.

---

**Function**

void *rpki\_retry\_hook* (timer \* *tm*) – control a scheduling of retrying connection to cache server

**Arguments**

timer \* *tm*  
retry timer with cache connection instance in data

**Description**

This function is periodically called during **ERROR\*** state cache connection. The first retry schedule is invoked after any **ERROR\*** state occurred and ends by reaching of **ESTABLISHED** state again.

---

**Function**

void *rpki\_expire\_hook* (timer \* *tm*) – control a expiration of ROA entries

**Arguments**

timer \* *tm*  
expire timer with cache connection instance in data

**Description**

This function is scheduled after received a **End of Data** PDU. A waiting interval is calculated dynamically by last update. If we reach an expiration time then we invoke a restarting of the protocol.

---

**Function**

const char \* *rpki\_check\_refresh\_interval* (uint *seconds*) – check validity of refresh interval value

**Arguments**

uint *seconds*  
suggested value

**Description**

This function validates value and should return **NULL**. If the check doesn't pass then returns error message.

---

**Function**

const char \* *rpki\_check\_retry\_interval* (uint *seconds*) – check validity of retry interval value

**Arguments**

uint *seconds*  
suggested value

**Description**

This function validates value and should return **NULL**. If the check doesn't pass then returns error message.

---

**Function**

const char \* *rpki\_check\_expire\_interval* (uint *seconds*) – check validity of expire interval value

**Arguments**

uint *seconds*  
suggested value

**Description**

This function validates value and should return NULL. If the check doesn't pass then returns error message.

---

**Function**

const char \* *rpki\_get\_cache\_ident* (struct rpki\_cache \* *cache*) – give a text representation of cache server name

**Arguments**

struct rpki\_cache \* *cache*  
RPKI connection instance

**Description**

The function converts cache connection into string.

---

**Function**

int *rpki\_reconfigure\_cache* (struct rpki\_proto \*p *UNUSED*, struct rpki\_cache \* *cache*, struct rpki\_config \* *new*, struct rpki\_config \* *old*) – a cache reconfiguration

**Arguments**

struct rpki\_proto \*p *UNUSED*  
– undescribed –  
  
struct rpki\_cache \* *cache*  
a cache connection  
  
struct rpki\_config \* *new*  
new RPKI configuration  
  
struct rpki\_config \* *old*  
old RPKI configuration

**Description**

This function reconfigures existing single cache server connection with new existing configuration. Generally, a change of time intervals could be reconfigured without restarting and all others changes requires a restart of protocol. Returns `NEED_TO_RESTART` or `SUCCESSFUL_RECONF`.

---

**Function**

int *rpki\_reconfigure* (struct proto \* *P*, struct proto\_config \* *CF*) – a protocol reconfiguration hook

**Arguments**

struct proto \* *P*  
a protocol instance  
  
struct proto\_config \* *CF*  
a new protocol configuration

**Description**

This function reconfigures whole protocol. It sets new protocol configuration into a protocol structure. Returns `NEED_TO_RESTART` or `SUCCESSFUL_RECONF`.

---

**Function**

void *rpki\_check\_config* (struct rpki\_config \* *cf*) – check and complete configuration of RPKI protocol

**Arguments**

struct rpki\_config \* *cf*  
RPKI configuration

**Description**

This function is called at the end of parsing RPKI protocol configuration.

---

**Function**

struct pdu\_header \* *rpki\_pdu\_back\_to\_network\_byte\_order* (struct pdu\_header \* *out*, const struct pdu\_header \* *in*) – convert host-byte order PDU back to network-byte order

**Arguments**

struct pdu\_header \* *out*  
allocated memory for writing a converted PDU of size *in->len*

const struct pdu\_header \* *in*  
host-byte order PDU

**Assumed**

A == ntohs(ntoh(A))

---

**Function**

int *rpki\_check\_receive\_packet* (struct rpki\_cache \* *cache*, const struct pdu\_header \* *pdu*) – make a basic validation of received RPKI PDU header

**Arguments**

struct rpki\_cache \* *cache*  
cache connection instance

const struct pdu\_header \* *pdu*  
RPKI PDU in network byte order

**Description**

This function checks protocol version, PDU type, version and size. If all is all right then function returns RPKI\_SUCCESS otherwise sends Error PDU and returns RPKI\_ERROR.

---

**Function**

net\_addr\_union \* *rpki\_prefix\_pdu\_2\_net\_addr* (const struct pdu\_header \* *pdu*, net\_addr\_union \* *n*) – convert IPv4/IPv6 Prefix PDU into net\_addr\_union

**Arguments**

const struct pdu\_header \* *pdu*  
host byte order IPv4/IPv6 Prefix PDU

net\_addr\_union \* *n*  
allocated net\_addr\_union for save ROA

**Description**

This function reads ROA data from IPv4/IPv6 Prefix PDU and write them into net\_addr\_roa4 or net\_addr\_roa6 data structure.

**Function**

void *rpki\_rx\_packet* (struct rpki\_cache \* *cache*, struct pdu\_header \* *pdu*) – process a received RPKI PDU

**Arguments**

struct rpki\_cache \* *cache*  
     RPKI connection instance

struct pdu\_header \* *pdu*  
     a RPKI PDU in network byte order

**Function**

int *rpki\_send\_error\_pdu* (struct rpki\_cache \* *cache*, const enum pdu\_error\_type *error\_code*, const u32 *err\_pdu\_len*, const struct pdu\_header \* *erroneous\_pdu*, const char \* *fmt*, ... ..) – send RPKI Error PDU

**Arguments**

struct rpki\_cache \* *cache*  
     RPKI connection instance

const enum pdu\_error\_type *error\_code*  
     PDU Error type

const u32 *err\_pdu\_len*  
     length of *erroneous\_pdu*

const struct pdu\_header \* *erroneous\_pdu*  
     optional network byte-order PDU that invokes Error by us or NULL

const char \* *fmt*  
     optional description text of error or NULL

... ..  
     variable arguments

**Description**

This function prepares Error PDU and sends it to a cache server.

**Function**

ip\_addr *rpki\_hostname\_autoresolv* (const char \* *host*) – auto-resolve an IP address from a hostname

**Arguments**

const char \* *host*  
     domain name of host, e.g. "rpki-validator.realmv6.org"

**Description**

This function resolves an IP address from a hostname. Returns *ip\_addr* structure with IP address or *IPA\_NONE*.

**Function**

int *rpki\_tr\_open* (struct rpki\_tr\_sock \* *tr*) – prepare and open a socket connection

**Arguments**

struct rpki\_tr\_sock \* *tr*  
     initialized transport socket

**Description**

Prepare and open a socket connection specified by *tr* that must be initialized before. This function ends with a calling the *sk\_open()* function. Returns *RPKI\_TR\_SUCCESS* or *RPKI\_TR\_ERROR*.

**Function**

void *rpki\_tr\_close* (struct rpki\_tr\_sock \* *tr*) – close socket and prepare it for possible next open

**Arguments**

struct rpki\_tr\_sock \* *tr*  
successfully opened transport socket

**Description**

Close socket and free resources.

**Function**

const char \* *rpki\_tr\_ident* (struct rpki\_tr\_sock \* *tr*) – Returns a string identifier for the rpki transport socket

**Arguments**

struct rpki\_tr\_sock \* *tr*  
successfully opened transport socket

**Description**

Returns a \0 terminated string identifier for the socket endpoint, e.g. "<host>:<port>". Memory is allocated inside *tr* structure.

**Function**

void *rpki\_tr\_tcp\_init* (struct rpki\_tr\_sock \* *tr*) – initializes the RPKI transport structure for a TCP connection

**Arguments**

struct rpki\_tr\_sock \* *tr*  
allocated RPKI transport structure

**Function**

void *rpki\_tr\_ssh\_init* (struct rpki\_tr\_sock \* *tr*) – initializes the RPKI transport structure for a SSH connection

**Arguments**

struct rpki\_tr\_sock \* *tr*  
allocated RPKI transport structure

## 5.9 Static

The Static protocol is implemented in a straightforward way. It keeps a list of static routes. Routes of dest RTD\_UNICAST have associated sticky node in the neighbor cache to be notified about gaining or losing the neighbor and about interface-related events (e.g. link down). They may also have a BFD request if associated with a BFD session. When a route is notified, *static\_decide()* is used to see whether the route activeness is changed. In such case, the route is marked as dirty and scheduled to be announced or withdrawn, which is done asynchronously from event hook. Routes of other types (e.g. black holes) are announced all the time. Multipath routes are a bit tricky. To represent additional next hops, dummy static\_route nodes are used, which are chained using *mp\_next* field and link to the master node by *mp\_head* field. Each next hop has a separate neighbor entry and an activeness state, but the master node is used for most purposes. Note that most functions DO NOT accept dummy nodes as arguments.

The only other thing worth mentioning is that when asked for reconfiguration, Static not only compares the two configurations, but it also calculates difference between the lists of static routes and it just inserts the newly added routes, removes the obsolete ones and reannounces changed ones.

## 5.10 Direct

The Direct protocol works by converting all *ifa\_notify()* events it receives to *rte\_update()* calls for the corresponding network.



# Chapter 6: System dependent parts

## 6.1 Introduction

We've tried to make BIRD as portable as possible, but unfortunately communication with the network stack differs from one OS to another, so we need at least some OS specific code. The good news is that this code is isolated in a small set of modules:

**config.h**

is a header file with configuration information, definition of the standard set of types and so on.

**Startup module**

controls BIRD startup. Common for a family of OS's (e.g., for all Unices).

**Logging module**

manages the system logs. [per OS family]

**IO module**

gives an implementation of sockets, timers and the global event queue. [per OS family]

**KRT module**

implements the Kernel and Device protocols. This is the most arcane part of the system dependent stuff and some functions differ even between various releases of a single OS.

## 6.2 Logging

The Logging module offers a simple set of functions for writing messages to system logs and to the debug output. Message classes used by this module are described in `birdlib.h` and also in the user's manual.

---

### Function

void *log\_commit* (int *class*, buffer \* *buf*) – commit a log message

### Arguments

int *class*

message class information (L\_DEBUG to L\_BUG, see `lib/birdlib.h`)

buffer \* *buf*

message to write

### Description

This function writes a message prepared in the log buffer to the log file (as specified in the configuration). The log buffer is reset after that. The log message is a full line, *log\_commit()* terminates it.

The message class is an integer, not a first char of a string like in *log()*, so it should be written like \*L\_INFO.

---

### Function

void *log\_msg* (const char \* *msg*, ... ...) – log a message

### Arguments

const char \* *msg*

printf-like formatting string with message class information prepended (L\_DEBUG to L\_BUG, see `lib/birdlib.h`)

... ..

variable arguments

### Description

This function formats a message according to the format string *msg* and writes it to the corresponding log file (as specified in the configuration). Please note that the message is automatically formatted as a full line, no need to include `\n` inside. It is essentially a sequence of *log\_reset()*, *logn()* and *log\_commit()*.

**Function**

void *bug* (const char \* *msg*, ... ...) – report an internal error

**Arguments**

const char \* *msg*  
     a printf-like error message

... ...  
     variable arguments

**Description**

This function logs an internal error and aborts execution of the program.

**Function**

void *die* (const char \* *msg*, ... ...) – report a fatal error

**Arguments**

const char \* *msg*  
     a printf-like error message

... ...  
     variable arguments

**Description**

This function logs a fatal error and aborts execution of the program.

**Function**

void *debug* (const char \* *msg*, ... ...) – write to debug output

**Arguments**

const char \* *msg*  
     a printf-like message

... ...  
     variable arguments

**Description**

This function formats the message *msg* and prints it out to the debugging output. No newline character is appended.

## 6.3 Kernel synchronization

This system dependent module implements the Kernel and Device protocol, that is synchronization of interface lists and routing tables with the OS kernel.

The whole kernel synchronization is a bit messy and touches some internals of the routing table engine, because routing table maintenance is a typical example of the proverbial compatibility between different Unices and we want to keep the overhead of our KRT business as low as possible and avoid maintaining a local routing table copy.

The kernel syncer can work in three different modes (according to system config header): Either with a single routing table and single KRT protocol [traditional UNIX] or with many routing tables and separate KRT protocols for all of them or with many routing tables, but every scan including all tables, so we start

separate KRT protocols which cooperate with each other [Linux]. In this case, we keep only a single scan timer.

We use FIB node flags in the routing table to keep track of route synchronization status. We also attach temporary `rte`'s to the routing table, but it cannot do any harm to the rest of BIRD since table synchronization is an atomic process.

When starting up, we cheat by looking if there is another KRT instance to be initialized later and performing table scan only once for all the instances.

The code uses OS-dependent parts for kernel updates and scans. These parts are in more specific sysdep directories (e.g. `sysdep/linux`) in functions `krt_sys_*` and `kif_sys_*` (and some others like `krt_replace_rte()`) and `krt-sys.h` header file. This is also used for platform specific protocol options and route attributes.

There was also an old code that used traditional UNIX `ioctl`s for these tasks. It was unmaintained and later removed. For reference, see `sysdep/krt-*` files in commit [396dfa9042305f62da1f56589c4b98fac57fc2f6](#)

# Chapter 7: Library functions

## 7.1 IP addresses

BIRD uses its own abstraction of IP address in order to share the same code for both IPv4 and IPv6. IP addresses are represented as entities of type `ip_addr` which are never to be treated as numbers and instead they must be manipulated using the following functions and macros.

---

### Function

`char * ip_scope_text (uint scope)` – get textual representation of address scope

### Arguments

uint *scope*  
scope (SCOPE\_xxx)

### Description

Returns a pointer to a textual name of the scope given.

---

### Function

`int ipa_equal (ip_addr x, ip_addr y)` – compare two IP addresses for equality

### Arguments

ip\_addr *x*  
IP address

ip\_addr *y*  
IP address

### Description

`ipa_equal()` returns 1 if *x* and *y* represent the same IP address, else 0.

---

### Function

`int ipa_nonzero (ip_addr x)` – test if an IP address is defined

### Arguments

ip\_addr *x*  
IP address

### Description

`ipa_nonzero` returns 1 if *x* is a defined IP address (not all bits are zero), else 0. The undefined all-zero address is reachable as a `IPA_NONE` macro.

---

### Function

`ip_addr ipa_and (ip_addr x, ip_addr y)` – compute bitwise and of two IP addresses

### Arguments

ip\_addr *x*  
IP address

ip\_addr *y*  
IP address

### Description

This function returns a bitwise and of *x* and *y*. It's primarily used for network masking.

---

**Function**

`ip_addr ipa_or (ip_addr x, ip_addr y)` – compute bitwise or of two IP addresses

**Arguments**

`ip_addr x`  
IP address

`ip_addr y`  
IP address

**Description**

This function returns a bitwise or of  $x$  and  $y$ .

---

**Function**

`ip_addr ipa_xor (ip_addr x, ip_addr y)` – compute bitwise xor of two IP addresses

**Arguments**

`ip_addr x`  
IP address

`ip_addr y`  
IP address

**Description**

This function returns a bitwise xor of  $x$  and  $y$ .

---

**Function**

`ip_addr ipa_not (ip_addr x)` – compute bitwise negation of two IP addresses

**Arguments**

`ip_addr x`  
IP address

**Description**

This function returns a bitwise negation of  $x$ .

---

**Function**

`ip_addr ipa_mkmask (int x)` – create a netmask

**Arguments**

`int x`  
prefix length

**Description**

This function returns an `ip_addr` corresponding of a netmask of an address prefix of size  $x$ .

---

**Function**

`int ipa_masklen (ip_addr x)` – calculate netmask length

**Arguments**

`ip_addr x`  
IP address

**Description**

This function checks whether  $x$  represents a valid netmask and returns the size of the associate network prefix or -1 for invalid mask.

---

**Function**

int *ipa\_hash* (ip\_addr *x*) – hash IP addresses

**Arguments**

ip\_addr *x*  
IP address

**Description**

*ipa\_hash()* returns a 16-bit hash value of the IP address *x*.

---

**Function**

void *ipa\_hton* (ip\_addr *x*) – convert IP address to network order

**Arguments**

ip\_addr *x*  
IP address

**Description**

Converts the IP address *x* to the network byte order.  
Beware, this is a macro and it alters the argument!

---

**Function**

void *ipa\_ntoh* (ip\_addr *x*) – convert IP address to host order

**Arguments**

ip\_addr *x*  
IP address

**Description**

Converts the IP address *x* from the network byte order.  
Beware, this is a macro and it alters the argument!

---

**Function**

int *ipa\_classify* (ip\_addr *x*) – classify an IP address

**Arguments**

ip\_addr *x*  
IP address

**Description**

*ipa\_classify()* returns an address class of *x*, that is a bitwise or of address type (IADDR\_INVALID, IADDR\_HOST, IADDR\_BROADCAST, IADDR\_MULTICAST) with address scope (SCOPE\_HOST to SCOPE\_UNIVERSE) or -1 (IADDR\_INVALID) for an invalid address.

---

**Function**

ip4\_addr *ip4\_class\_mask* (ip4\_addr *x*) – guess netmask according to address class

**Arguments**

ip4\_addr *x*  
IPv4 address

**Description**

This function (available in IPv4 version only) returns a network mask according to the address class of *x*. Although classful addressing is nowadays obsolete, there still live routing protocols transferring no prefix lengths nor netmasks and this function could be useful to them.

---

**Function**

u32 *ipa\_from\_u32* (ip\_addr *x*) – convert IPv4 address to an integer

**Arguments**

ip\_addr *x*  
IP address

**Description**

This function takes an IPv4 address and returns its numeric representation.

---

**Function**

ip\_addr *ipa\_to\_u32* (u32 *x*) – convert integer to IPv4 address

**Arguments**

u32 *x*  
a 32-bit integer

**Description**

*ipa\_to\_u32()* takes a numeric representation of an IPv4 address and converts it to the corresponding ip\_addr.

---

**Function**

int *ipa\_compare* (ip\_addr *x*, ip\_addr *y*) – compare two IP addresses for order

**Arguments**

ip\_addr *x*  
IP address

ip\_addr *y*  
IP address

**Description**

The *ipa\_compare()* function takes two IP addresses and returns -1 if *x* is less than *y* in canonical ordering (lexicographical order of the bit strings), 1 if *x* is greater than *y* and 0 if they are the same.

---

**Function**

ip\_addr *ipa\_build6* (u32 *a1*, u32 *a2*, u32 *a3*, u32 *a4*) – build an IPv6 address from parts

**Arguments**

u32 *a1*  
part #1

u32 *a2*  
part #2

u32 *a3*  
part #3

u32 *a4*  
part #4

**Description**

*ipa\_build()* takes *a1* to *a4* and assembles them to a single IPv6 address. It's used for example when a protocol wants to bind its socket to a hard-wired multicast address.

**Function**

`char * ip_ntop (ip_addr a, char * buf)` – convert IP address to textual representation

**Arguments**

`ip_addr a`  
IP address

`char * buf`  
buffer of size at least `STD_ADDRESS_P_LENGTH`

**Description**

This function takes an IP address and creates its textual representation for presenting to the user.

**Function**

`char * ip_ntox (ip_addr a, char * buf)` – convert IP address to hexadecimal representation

**Arguments**

`ip_addr a`  
IP address

`char * buf`  
buffer of size at least `STD_ADDRESS_P_LENGTH`

**Description**

This function takes an IP address and creates its hexadecimal textual representation. Primary use: debugging dumps.

**Function**

`int ip_pton (char * a, ip_addr * o)` – parse textual representation of IP address

**Arguments**

`char * a`  
textual representation

`ip_addr * o`  
where to put the resulting address

**Description**

This function parses a textual IP address representation and stores the decoded address to a variable pointed to by `o`. Returns 0 if a parse error has occurred, else 1.

## 7.2 Linked lists

The BIRD library provides a set of functions for operating on linked lists. The lists are internally represented as standard doubly linked lists with synthetic head and tail which makes all the basic operations run in constant time and contain no extra end-of-list checks. Each list is described by a `list` structure, nodes can have any format as long as they start with a `node` structure. If you want your nodes to belong to multiple lists at once, you can embed multiple `node` structures in them and use the `SKIP_BACK()` macro to calculate a pointer to the start of the structure from a `node` pointer, but beware of obscurity.

There also exist safe linked lists (`slist`, `snode` and all functions being prefixed with `s_`) which support asynchronous walking very similar to that used in the `fib` structure.



---

**Function**

LIST\_INLINE void *add\_tail* (list \* *l*, node \* *n*) – append a node to a list

**Arguments**

list \* *l*  
linked list

node \* *n*  
list node

**Description**

*add\_tail()* takes a node *n* and appends it at the end of the list *l*.

---

**Function**

LIST\_INLINE void *add\_head* (list \* *l*, node \* *n*) – prepend a node to a list

**Arguments**

list \* *l*  
linked list

node \* *n*  
list node

**Description**

*add\_head()* takes a node *n* and prepends it at the start of the list *l*.

---

**Function**

LIST\_INLINE void *insert\_node* (node \* *n*, node \* *after*) – insert a node to a list

**Arguments**

node \* *n*  
a new list node

node \* *after*  
a node of a list

**Description**

Inserts a node *n* to a linked list after an already inserted node *after*.

---

**Function**

LIST\_INLINE void *rem\_node* (node \* *n*) – remove a node from a list

**Arguments**

node \* *n*  
node to be removed

**Description**

Removes a node *n* from the list it's linked in. Afterwards, node *n* is cleared.

**Function**

LIST\_INLINE void *replace\_node* (node \* *old*, node \* *new*) – replace a node in a list with another one

**Arguments**

node \* *old*  
node to be removed

node \* *new*  
node to be inserted

**Description**

Replaces node *old* in the list it's linked in with node *new*. Node *old* may be a copy of the original node, which is not accessed through the list. The function could be called with *old* == *new*, which just fixes neighbors' pointers in the case that the node was reallocated.

**Function**

LIST\_INLINE void *init\_list* (list \* *l*) – create an empty list

**Arguments**

list \* *l*  
list

**Description**

*init\_list()* takes a `list` structure and initializes its fields, so that it represents an empty list.

**Function**

LIST\_INLINE void *add\_tail\_list* (list \* *to*, list \* *l*) – concatenate two lists

**Arguments**

list \* *to*  
destination list

list \* *l*  
source list

**Description**

This function appends all elements of the list *l* to the list *to* in constant time.

## 7.3 Miscellaneous functions.

**Function**

int *ipsum\_verify* (void \* *frag*, uint *len*, ... ...) – verify an IP checksum

**Arguments**

void \* *frag*  
first packet fragment

uint *len*  
length in bytes

... ..  
variable arguments

**Description**

This function verifies whether a given fragmented packet has correct one's complement checksum as used by the IP protocol.

It uses all the clever tricks described in RFC 1071 to speed up checksum calculation as much as possible.

**Result**

1 if the checksum is correct, 0 else.

---

**Function**

u16 *ipsum\_calculate* (void \* *frag*, uint *len*, ... ...) – compute an IP checksum

**Arguments**

void \* *frag*  
first packet fragment

uint *len*  
length in bytes

... ...  
variable arguments

**Description**

This function calculates a one's complement checksum of a given fragmented packet. It uses all the clever tricks described in RFC 1071 to speed up checksum calculation as much as possible.

---

**Function**

u32 *u32\_mkmask* (uint *n*) – create a bit mask

**Arguments**

uint *n*  
number of bits

**Description**

*u32\_mkmask()* returns an unsigned 32-bit integer which binary representation consists of *n* ones followed by zeroes.

---

**Function**

uint *u32\_masklen* (u32 *x*) – calculate length of a bit mask

**Arguments**

u32 *x*  
bit mask

**Description**

This function checks whether the given integer *x* represents a valid bit mask (binary representation contains first ones, then zeroes) and returns the number of ones or 255 if the mask is invalid.

---

**Function**

u32 *u32\_log2* (u32 *v*) – compute a binary logarithm.

**Arguments**

u32 *v*  
number

**Description**

This function computes a integral part of binary logarithm of given integer *v* and returns it. The computed value is also an index of the most significant non-zero bit position.

---

**Function**

int *patmatch* (byte \* *p*, byte \* *s*) – match shell-like patterns

**Arguments**

byte \* *p*  
pattern

byte \* *s*  
string

**Description**

*patmatch()* returns whether given string *s* matches the given shell-like pattern *p*. The patterns consist of characters (which are matched literally), question marks which match any single character, asterisks which match any (possibly empty) string of characters and backslashes which are used to escape any special characters and force them to be treated literally.

The matching process is not optimized with respect to time, so please avoid using this function for complex patterns.

---

**Function**

int *bvsnprintf* (char \* *buf*, int *size*, const char \* *fmt*, va\_list *args*) – BIRD's *vsnprintf()*

**Arguments**

char \* *buf*  
destination buffer

int *size*  
size of the buffer

const char \* *fmt*  
format string

va\_list *args*  
a list of arguments to be formatted

**Description**

This functions acts like ordinary *sprintf()* except that it checks available

**space to avoid buffer overflows and it allows some more format specifiers**

**I** for formatting of IP addresses (width of 1 is automatically replaced by standard IP address width which depends on whether we use IPv4 or IPv6; **I4** or **I6** can be used for explicit ip4\_addr / ip6\_addr arguments, **N** for generic network addresses (net\_addr \*), **R** for Router / Network ID (u32 value printed as IPv4 address), **1R** for 64bit Router / Network ID (u64

**value printed as eight**

-separated octets), **t** for time values (btime) with specified subsecond precision, and **m** resp. **M** for error messages (uses *strerror()* to translate *errno* code to message text). On the other hand, it doesn't support floating point numbers.

**Result**

number of characters of the output string or -1 if the buffer space was insufficient.

---

**Function**

int *bvsprintf* (char \* *buf*, const char \* *fmt*, va\_list *args*) – BIRD's *vsprintf()*

**Arguments**

char \* *buf*  
buffer

const char \* *fmt*  
format string

va\_list *args*  
a list of arguments to be formatted

**Description**

This function is equivalent to *bvsnprintf()* with an infinite buffer size. Please use carefully only when you are absolutely sure the buffer won't overflow.

---

**Function**

int *bsprintf* (char \* *buf*, const char \* *fmt*, ... ...) – BIRD's *sprintf()*

**Arguments**

char \* *buf*  
buffer

const char \* *fmt*  
format string

... ...  
variable arguments

**Description**

This function is equivalent to *bvsnprintf()* with an infinite buffer size and variable arguments instead of a *va\_list*. Please use carefully only when you are absolutely sure the buffer won't overflow.

---

**Function**

int *bsnprintf* (char \* *buf*, int *size*, const char \* *fmt*, ... ...) – BIRD's *snprintf()*

**Arguments**

char \* *buf*  
buffer

int *size*  
buffer size

const char \* *fmt*  
format string

... ...  
variable arguments

**Description**

This function is equivalent to *bsnprintf()* with variable arguments instead of a *va\_list*.

---

**Function**

void \* *xmalloc* (uint *size*) – malloc with checking

**Arguments**

uint *size*  
block size

**Description**

This function is equivalent to *malloc()* except that in case of failure it calls *die()* to quit the program instead of returning a NULL pointer.

Wherever possible, please use the memory resources instead.

---

**Function**

void \* *xrealloc* (void \* *ptr*, uint *size*) – realloc with checking

**Arguments**

void \* *ptr*  
original memory block

uint *size*  
block size

**Description**

This function is equivalent to *realloc()* except that in case of failure it calls *die()* to quit the program instead of returning a NULL pointer.

Wherever possible, please use the memory resources instead.

## 7.4 Message authentication codes

MAC algorithms are simple cryptographic tools for message authentication. They use shared a secret key and message text to generate authentication code, which is then passed with the message to the other side, where the code is verified. There are multiple families of MAC algorithms based on different cryptographic primitives, BIRD implements two MAC families which use hash functions.

The first family is simply a cryptographic hash camouflaged as MAC algorithm. Originally supposed to be (m|k)-hash (message is concatenated with key, and that is hashed), but later it turned out that a raw hash is more practical. This is used for cryptographic authentication in OSPFv2, RIP and BFD.

The second family is the standard HMAC (RFC 2104), using inner and outer hash to process key and message. HMAC (with SHA) is used in advanced OSPF and RIP authentication (RFC 5709, RFC 4822).

---

**Function**

void *mac\_init* (struct mac\_context \* *ctx*, uint *id*, const byte \* *key*, uint *keylen*) – initialize MAC algorithm

**Arguments**

struct mac\_context \* *ctx*  
context to initialize

uint *id*  
MAC algorithm ID

const byte \* *key*  
MAC key

uint *keylen*  
MAC key length

**Description**

Initialize MAC context *ctx* for algorithm *id* (e.g., `ALG_HMAC_SHA1`), with key *key* of length *keylen*. After that, message data could be added using *mac\_update()* function.

---

**Function**

void *mac\_update* (struct mac\_context \* *ctx*, const byte \* *data*, uint *datalen*) – add more data to MAC algorithm

**Arguments**

struct mac\_context \* *ctx*  
MAC context

const byte \* *data*  
data to add

uint *datalen*  
length of data

**Description**

Push another *datalen* bytes of data pointed to by *data* into the MAC algorithm currently in *ctx*. Can be called multiple times for the same MAC context. It has the same effect as concatenating all the data together and passing them at once.

---

**Function**

byte \* *mac\_final* (struct mac\_context \* *ctx*) – finalize MAC algorithm

**Arguments**

struct mac\_context \* *ctx*  
MAC context

**Description**

Finish MAC computation and return a pointer to the result. No more *mac\_update*() calls could be done, but the context may be reinitialized later.

Note that the returned pointer points into data in the *ctx* context. If it ceases to exist, the pointer becomes invalid.

---

**Function**

void *mac\_cleanup* (struct mac\_context \* *ctx*) – cleanup MAC context

**Arguments**

struct mac\_context \* *ctx*  
MAC context

**Description**

Cleanup MAC context after computation (by filling with zeros). Not strictly necessary, just to erase sensitive data from stack. This also invalidates the pointer returned by *mac\_final*().

---

**Function**

void *mac\_fill* (uint *id*, const byte \* *key*, uint *keylen*, const byte \* *data*, uint *datalen*, byte \* *mac*) – compute and fill MAC

**Arguments**

uint *id*  
MAC algorithm ID

const byte \* *key*  
secret key

uint *keylen*  
key length

const byte \* *data*  
message data

```
uint datalen
    message length

byte * mac
    place to fill MAC
```

**Description**

Compute MAC for specified key *key* and message *data* using algorithm *id* and copy it to buffer *mac*. *mac\_fill()* is a shortcut function doing all usual steps for transmitted messages.

**Function**

int *mac\_verify* (uint *id*, const byte \* *key*, uint *keylen*, const byte \* *data*, uint *datalen*, const byte \* *mac*) – compute and verify MAC

**Arguments**

```
uint id
    MAC algorithm ID

const byte * key
    secret key

uint keylen
    key length

const byte * data
    message data

uint datalen
    message length

const byte * mac
    received MAC
```

**Description**

Compute MAC for specified key *key* and message *data* using algorithm *id* and compare it with received *mac*, return whether they are the same. *mac\_verify()* is a shortcut function doing all usual steps for received messages.

## 7.5 Flow specification (flowspec)

Flowspec are rules (RFC 5575) for firewalls disseminated using BGP protocol. The `flowspec.c` is a library for handling flowspec binary streams and flowspec data structures. You will find there functions for validation incoming flowspec binary streams, iterators for jumping over components, functions for handling a length and functions for formatting flowspec data structure into user-friendly text representation.

In this library, you will find also flowspec builder. In `confbase.Y`, there are grammar's rules for parsing and building new flowspec data structure from BIRD's configuration files and from BIRD's command line interface. `Finalize` function will assemble final `net_addr_flow4` or `net_addr_flow6` data structure.

The data structures `net_addr_flow4` and `net_addr_flow6` are defined in `net.h` file. The attribute `length` is size of whole data structure plus binary stream representation of flowspec including a compressed encoded length of flowspec.

Sometimes in code, it is used expression flowspec type, it should mean flowspec component type.

**Function**

const char \* *flow\_type\_str* (enum flow\_type *type*, int *ipv6*) – get stringified flowspec name of component



**Arguments**

enum flow\_type *type*  
     flowspec component type

int *ipv6*  
     IPv4/IPv6 decide flag, use zero for IPv4 and one for IPv6

**Description**

This function returns flowspec name of component *type* in string.

---

**Function**

uint *flow\_write\_length* (byte \* *data*, u16 *len*) – write compressed length value

**Arguments**

byte \* *data*  
     destination buffer to write

u16 *len*  
     the value of the length (0 to 0xffff) for writing

**Description**

This function writes appropriate as (1- or 2-bytes) the value of *len* into buffer *data*. The function returns number of written bytes, thus 1 or 2 bytes.

---

**Function**

const byte \* *flow4\_first\_part* (const net\_addr\_flow4 \* *f*) – get position of the first flowspec component

**Arguments**

const net\_addr\_flow4 \* *f*  
     flowspec data structure `net_addr_flow4`

**Description**

This function return a position to the beginning of the first flowspec component in IPv4 flowspec *f*.

---

**Function**

const byte \* *flow6\_first\_part* (const net\_addr\_flow6 \* *f*) – get position of the first flowspec component

**Arguments**

const net\_addr\_flow6 \* *f*  
     flowspec data structure `net_addr_flow6`

**Description**

This function return a position to the beginning of the first flowspec component in IPv6 flowspec *f*.

---

**Function**

const byte \* *flow4\_next\_part* (const byte \* *pos*, const byte \* *end*) – an iterator over flowspec components in flowspec binary stream

**Arguments**

const byte \* *pos*  
     the beginning of a previous or the first component in flowspec binary stream

const byte \* *end*  
     the last valid byte in scanned flowspec binary stream

**Description**

This function returns a position to the beginning of the next component (to a component type byte) in flowspec binary stream or NULL for the end.

---

**Function**

const byte \* *flow6\_next\_part* (const byte \* *pos*, const byte \* *end*) – an iterator over flowspec components in flowspec binary stream

**Arguments**

const byte \* *pos*  
the beginning of a previous or the first component in flowspec binary stream

const byte \* *end*  
the last valid byte in scanned flowspec binary stream

**Description**

This function returns a position to the beginning of the next component (to a component type byte) in flowspec binary stream or NULL for the end.

---

**Function**

const char \* *flow\_validated\_state\_str* (enum flow\_validated\_state *code*) – return a textual description of validation process

**Arguments**

enum flow\_validated\_state *code*  
validation result

**Description**

This function return well described validation state in string.

---

**Function**

void *flow\_check\_cf\_bmk\_values* (struct flow\_builder \* *fb*, u8 *neg*, u32 *val*, u32 *mask*) – check value/bitmask part of flowspec component

**Arguments**

struct flow\_builder \* *fb*  
flow builder instance

u8 *neg*  
negation operand

u32 *val*  
value from value/mask pair

u32 *mask*  
bitmap mask from value/mask pair

**Description**

This function checks value/bitmask pair. If some problem will appear, the function calls *cf\_error()* function with a textual description of reason to failing of validation.

---

**Function**

void *flow\_check\_cf\_value\_length* (struct flow\_builder \* *fb*, u32 *val*) – check value by flowspec component type

**Arguments**

struct flow\_builder \* *fb*  
flow builder instance

u32 *val*  
value

**Description**

This function checks if the value is in range of component's type support. If some problem will appear, the function calls *cf\_error()* function with a textual description of reason to failing of validation.

---

**Function**

enum flow\_validated\_state *flow4\_validate* (const byte \* *nlri*, uint *len*) – check untrustworthy IPv4 flowspec data stream

**Arguments**

const byte \* *nlri*  
flowspec data stream without compressed encoded length value

uint *len*  
length of *nlri*

**Description**

This function checks meaningfulness of binary flowspec. It should return FLOW\_ST\_VALID or FLOW\_ST\_UNKNOWN\_COMPONENT. If some problem appears, it returns some other FLOW\_ST\_XXX state.

---

**Function**

enum flow\_validated\_state *flow6\_validate* (const byte \* *nlri*, uint *len*) – check untrustworthy IPv6 flowspec data stream

**Arguments**

const byte \* *nlri*  
flowspec binary stream without encoded length value

uint *len*  
length of *nlri*

**Description**

This function checks meaningfulness of binary flowspec. It should return FLOW\_ST\_VALID or FLOW\_ST\_UNKNOWN\_COMPONENT. If some problem appears, it returns some other FLOW\_ST\_XXX state.

---

**Function**

void *flow4\_validate\_cf* (net\_addr\_flow4 \* *f*) – validate flowspec data structure **net\_addr\_flow4** in parsing time

**Arguments**

net\_addr\_flow4 \* *f*  
flowspec data structure **net\_addr\_flow4**

**Description**

Check if *f* is valid flowspec data structure. Can call *cf\_error()* function with a textual description of reason to failing of validation.

---

**Function**

void *flow6\_validate\_cf* (net\_addr\_flow6 \* *f*) – validate flowspec data structure **net\_addr\_flow6** in parsing time

**Arguments**

net\_addr\_flow6 \* *f*  
flowspec data structure **net\_addr\_flow6**

**Description**

Check if *f* is valid flowspec data structure. Can call *cf\_error()* function with a textual description of reason to failing of validation.

**Function**

struct flow\_builder \* *flow\_builder\_init* (pool \* *pool*) – constructor for flowspec builder instance

**Arguments**

pool \* *pool*  
memory pool

**Description**

This function prepares flowspec builder instance using memory pool *pool*.

**Function**

int *flow\_builder4\_add\_pfx* (struct flow\_builder \* *fb*, const net\_addr\_ip4 \* *n4*) – add IPv4 prefix

**Arguments**

struct flow\_builder \* *fb*  
flowspec builder instance  
  
const net\_addr\_ip4 \* *n4*  
net address of type IPv4

**Description**

This function add IPv4 prefix into flowspec builder instance.

**Function**

int *flow\_builder6\_add\_pfx* (struct flow\_builder \* *fb*, const net\_addr\_ip6 \* *n6*, u32 *poffset*) – add IPv6 prefix

**Arguments**

struct flow\_builder \* *fb*  
flowspec builder instance  
  
const net\_addr\_ip6 \* *n6*  
net address of type IPv4  
  
u32 *poffset*  
prefix offset for *n6*

**Description**

This function add IPv4 prefix into flowspec builder instance. This function should return 1 for successful adding, otherwise returns 0.

**Function**

int *flow\_builder\_add\_op\_val* (struct flow\_builder \* *fb*, byte *op*, u32 *value*) – add operator/value pair

**Arguments**

struct flow\_builder \* *fb*  
flowspec builder instance  
  
byte *op*  
operator  
  
u32 *value*  
value

**Description**

This function add operator/value pair as a part of a flowspec component. It is required to set appropriate flowspec component type using function *flow\_builder\_set\_type()*. This function should return 1 for successful adding, otherwise returns 0.

---

**Function**

int *flow\_builder\_add\_val\_mask* (struct flow\_builder \* *fb*, byte *op*, u32 *value*, u32 *mask*) – add value/bitmask pair

**Arguments**

struct flow\_builder \* *fb*  
flowspec builder instance

byte *op*  
operator

u32 *value*  
value

u32 *mask*  
bitmask

**Description**

It is required to set appropriate flowspec component type using function *flow\_builder\_set\_type()*. This function should return 1 for successful adding, otherwise returns 0.

---

**Function**

void *flow\_builder\_set\_type* (struct flow\_builder \* *fb*, enum flow\_type *type*) – set type of next flowspec component

**Arguments**

struct flow\_builder \* *fb*  
flowspec builder instance

enum flow\_type *type*  
flowspec component type

**Description**

This function sets type of next flowspec component. It is necessary to call this function before each changing of adding flowspec component.

---

**Function**

net\_addr\_flow4 \* *flow\_builder4\_finalize* (struct flow\_builder \* *fb*, linpool \* *lpool*) – assemble final flowspec data structure **net\_addr\_flow4**

**Arguments**

struct flow\_builder \* *fb*  
flowspec builder instance

linpool \* *lpool*  
linear memory pool

**Description**

This function returns final flowspec data structure **net\_addr\_flow4** allocated onto *lpool* linear memory pool.

---

**Function**

net\_addr\_flow6 \* *flow\_builder6\_finalize* (struct flow\_builder \* *fb*, linpool \* *lpool*) – assemble final flowspec data structure **net\_addr\_flow6**

**Arguments**

struct flow\_builder \* *fb*  
     flowspec builder instance

linpool \* *lpool*  
     linear memory pool for allocation of

**Description**

This function returns final flowspec data structure `net_addr_flow6` allocated onto *lpool* linear memory pool.

---

**Function**

void *flow\_builder\_clear* (struct flow\_builder \* *fb*) – flush flowspec builder instance for another flowspec creation

**Arguments**

struct flow\_builder \* *fb*  
     flowspec builder instance

**Description**

This function flushes all data from builder but it maintains pre-allocated buffer space.

---

**Function**

uint *flow4\_net\_format* (char \* *buf*, uint *blen*, const net\_addr\_flow4 \* *f*) – stringify flowspec data structure `net_addr_flow4`

**Arguments**

char \* *buf*  
     pre-allocated buffer for writing a stringify net address flowspec

uint *blen*  
     free allocated space in *buf*

const net\_addr\_flow4 \* *f*  
     flowspec data structure `net_addr_flow4` for stringify

**Description**

This function writes stringified *f* into *buf*. The function returns number of written chars. If final string is too large, the string will ends the with ' ...}' sequence and zero-terminator.

---

**Function**

uint *flow6\_net\_format* (char \* *buf*, uint *blen*, const net\_addr\_flow6 \* *f*) – stringify flowspec data structure `net_addr_flow6`

**Arguments**

char \* *buf*  
     pre-allocated buffer for writing a stringify net address flowspec

uint *blen*  
     free allocated space in *buf*

const net\_addr\_flow6 \* *f*  
     flowspec data structure `net_addr_flow4` for stringify

**Description**

This function writes stringified *f* into *buf*. The function returns number of written chars. If final string is too large, the string will ends the with ' ...}' sequence and zero-terminator.

# Chapter 8: Resources

## 8.1 Introduction

Most large software projects implemented in classical procedural programming languages usually end up with lots of code taking care of resource allocation and deallocation. Bugs in such code are often very difficult to find, because they cause only ‘resource leakage’, that is keeping a lot of memory and other resources which nobody references to.

We’ve tried to solve this problem by employing a resource tracking system which keeps track of all the resources allocated by all the modules of BIRD, deallocates everything automatically when a module shuts down and it is able to print out the list of resources and the corresponding modules they are allocated by.

Each allocated resource (from now we’ll speak about allocated resources only) is represented by a structure starting with a standard header (struct **resource**) consisting of a list node (resources are often linked to various lists) and a pointer to **resclass** – a resource class structure pointing to functions implementing generic resource operations (such as freeing of the resource) for the particular resource type.

There exist the following types of resources:

- *Resource pools* (**pool**)
- *Memory blocks*
- *Linear memory pools* (**linpool**)
- *Slabs* (**slab**)
- *Events* (**event**)
- *Timers* (**timer**)
- *Sockets* (**socket**)

## 8.2 Resource pools

Resource pools (**pool**) are just containers holding a list of other resources. Freeing a pool causes all the listed resources to be freed as well. Each existing **resource** is linked to some pool except for a root pool which isn’t linked anywhere, so all the resources form a tree structure with internal nodes corresponding to pools and leaves being the other resources.

Example: Almost all modules of BIRD have their private pool which is freed upon shutdown of the module.

---

### Function

`pool * rp_new (pool * p, const char * name)` – create a resource pool

### Arguments

`pool * p`  
parent pool

`const char * name`  
pool name (to be included in debugging dumps)

### Description

`rp_new()` creates a new resource pool inside the specified parent pool.

---

### Function

`void rmove (void * res, pool * p)` – move a resource

**Arguments**

void \* *res*  
     resource

pool \* *p*  
     pool to move the resource to

**Description**

*rmove()* moves a resource from one pool to another.

---

**Function**

void *rfree* (void \* *res*) – free a resource

**Arguments**

void \* *res*  
     resource

**Description**

*rfree()* frees the given resource and all information associated with it. In case it's a resource pool, it also frees all the objects living inside the pool.

It works by calling a class-specific freeing function.

---

**Function**

void *rdump* (void \* *res*) – dump a resource

**Arguments**

void \* *res*  
     resource

**Description**

This function prints out all available information about the given resource to the debugging output.

It works by calling a class-specific dump function.

---

**Function**

void \* *ralloc* (pool \* *p*, struct resclass \* *c*) – create a resource

**Arguments**

pool \* *p*  
     pool to create the resource in

struct resclass \* *c*  
     class of the new resource

**Description**

This function is called by the resource classes to create a new resource of the specified class and link it to the given pool. Allocated memory is zeroed. Size of the resource structure is taken from the *size* field of the **resclass**.

---

**Function**

void *rlookup* (unsigned long *a*) – look up a memory location

**Arguments**

unsigned long *a*  
     memory address

**Description**

This function examines all existing resources to see whether the address *a* is inside any resource. It's used for debugging purposes only.

It works by calling a class-specific lookup function for each resource.



**Function**

void *resource\_init* (*void*) – initialize the resource manager

**Description**

This function is called during BIRD startup. It initializes all data structures of the resource manager and creates the root pool.

## 8.3 Memory blocks

Memory blocks are pieces of contiguous allocated memory. They are a bit non-standard since they are represented not by a pointer to **resource**, but by a void pointer to the start of data of the memory block. All memory block functions know how to locate the header given the data pointer.

Example: All "unique" data structures such as hash tables are allocated as memory blocks.

**Function**

void \* *mb\_alloc* (pool \* *p*, unsigned *size*) – allocate a memory block

**Arguments**

pool \* *p*  
pool

unsigned *size*  
size of the block

**Description**

*mb\_alloc()* allocates memory of a given size and creates a memory block resource representing this memory chunk in the pool *p*.

Please note that *mb\_alloc()* returns a pointer to the memory chunk, not to the resource, hence you have to free it using *mb\_free()*, not *rfree()*.

**Function**

void \* *mb\_allocc* (pool \* *p*, unsigned *size*) – allocate and clear a memory block

**Arguments**

pool \* *p*  
pool

unsigned *size*  
size of the block

**Description**

*mb\_allocc()* allocates memory of a given size, initializes it to zeroes and creates a memory block resource representing this memory chunk in the pool *p*.

Please note that *mb\_allocc()* returns a pointer to the memory chunk, not to the resource, hence you have to free it using *mb\_free()*, not *rfree()*.

**Function**

void \* *mb\_realloc* (void \* *m*, unsigned *size*) – reallocate a memory block

**Arguments**

void \* *m*  
memory block

unsigned *size*  
 new size of the block

**Description**

*mb\_realloc()* changes the size of the memory block *m* to a given size. The contents will be unchanged to the minimum of the old and new sizes; newly allocated memory will be uninitialized. Contrary to *realloc()* behavior, *m* must be non-NULL, because the resource pool is inherited from it.

Like *mb\_alloc()*, *mb\_realloc()* also returns a pointer to the memory chunk, not to the resource, hence you have to free it using *mb\_free()*, not *rfree()*.

**Function**

void *mb\_free* (void \* *m*) – free a memory block

**Arguments**

void \* *m*  
 memory block

**Description**

*mb\_free()* frees all memory associated with the block *m*.

## 8.4 Linear memory pools

Linear memory pools are collections of memory blocks which support very fast allocation of new blocks, but are able to free only the whole collection at once (or in stack order).

Example: Each configuration is described by a complex system of structures, linked lists and function trees which are all allocated from a single linear pool, thus they can be freed at once when the configuration is no longer used.

**Function**

linpool \* *lp\_new* (pool \* *p*, uint *blk*) – create a new linear memory pool

**Arguments**

pool \* *p*  
 pool  
 uint *blk*  
 block size

**Description**

*lp\_new()* creates a new linear memory pool resource inside the pool *p*. The linear pool consists of a list of memory chunks of size at least *blk*.

**Function**

void \* *lp\_alloc* (linpool \* *m*, uint *size*) – allocate memory from a linpool

**Arguments**

linpool \* *m*  
 linear memory pool  
 uint *size*  
 amount of memory

**Description**

*lp\_alloc()* allocates *size* bytes of memory from a linpool *m* and it returns a pointer to the allocated memory. It works by trying to find free space in the last memory chunk associated with the linpool and creating a new chunk of the standard size (as specified during *lp\_new()*) if the free space is too small to satisfy the allocation. If *size* is too large to fit in a standard size chunk, an "overflow" chunk is created for it instead.

---

**Function**

`void * lp_allocu (linpool * m, uint size)` – allocate unaligned memory from a `linpool`

**Arguments**

`linpool * m`  
linear memory pool

`uint size`  
amount of memory

**Description**

`lp_allocu()` allocates *size* bytes of memory from a `linpool` *m* and it returns a pointer to the allocated memory. It doesn't attempt to align the memory block, giving a very efficient way how to allocate strings without any space overhead.

---

**Function**

`void * lp_allocz (linpool * m, uint size)` – allocate cleared memory from a `linpool`

**Arguments**

`linpool * m`  
linear memory pool

`uint size`  
amount of memory

**Description**

This function is identical to `lp_alloc()` except that it clears the allocated memory block.

---

**Function**

`void lp_flush (linpool * m)` – flush a linear memory pool

**Arguments**

`linpool * m`  
linear memory pool

**Description**

This function frees the whole contents of the given `linpool` *m*, but leaves the pool itself.

---

**Function**

`void lp_save (linpool * m, lp_state * p)` – save the state of a linear memory pool

**Arguments**

`linpool * m`  
linear memory pool

`lp_state * p`  
state buffer

**Description**

This function saves the state of a linear memory pool. Saved state can be used later to restore the pool (to free memory allocated since).

**Function**

`void lp_restore (linpool * m, lp_state * p)` – restore the state of a linear memory pool

**Arguments**

`linpool * m`  
linear memory pool

`lp_state * p`  
saved state

**Description**

This function restores the state of a linear memory pool, freeing all memory allocated since the state was saved. Note that the function cannot un-free the memory, therefore the function also invalidates other states that were saved between (on the same pool).

## 8.5 Slabs

Slabs are collections of memory blocks of a fixed size. They support very fast allocation and freeing of such blocks, prevent memory fragmentation and optimize L2 cache usage. Slabs have been invented by Jeff Bonwick and published in USENIX proceedings as ‘The Slab Allocator: An Object-Caching Kernel Memory Allocator’. Our implementation follows this article except that we don’t use constructors and destructors.

When the `DEBUGGING` switch is turned on, we automatically fill all newly allocated and freed blocks with a special pattern to make detection of use of uninitialized or already freed memory easier.

Example: Nodes of a FIB are allocated from a per-FIB Slab.

**Function**

`slab * sl_new (pool * p, uint size)` – create a new Slab

**Arguments**

`pool * p`  
resource pool

`uint size`  
block size

**Description**

This function creates a new Slab resource from which objects of size *size* can be allocated.

**Function**

`void * sl_alloc (slab * s)` – allocate an object from Slab

**Arguments**

`slab * s`  
slab

**Description**

`sl_alloc()` allocates space for a single object from the Slab and returns a pointer to the object.

**Function**

`void sl_free (slab * s, void * oo)` – return a free object back to a Slab

**Arguments**

slab \* *s*  
     slab

void \* *oo*  
     object returned by *sl\_alloc()*

**Description**

This function frees memory associated with the object *oo* and returns it back to the Slab *s*.

## 8.6 Events

Events are there to keep track of deferred execution. Since BIRD is single-threaded, it requires long lasting tasks to be split to smaller parts, so that no module can monopolize the CPU. To split such a task, just create an **event** resource, point it to the function you want to have called and call *ev\_schedule()* to ask the core to run the event when nothing more important requires attention.

You can also define your own event lists (the **event\_list** structure), enqueue your events in them and explicitly ask to run them.

---

**Function**

event \* *ev\_new* (pool \* *p*) – create a new event

**Arguments**

pool \* *p*  
     resource pool

**Description**

This function creates a new event resource. To use it, you need to fill the structure fields and call *ev\_schedule()*.

---

**Function**

void *ev\_run* (event \* *e*) – run an event

**Arguments**

event \* *e*  
     an event

**Description**

This function explicitly runs the event *e* (calls its hook function) and removes it from an event list if it's linked to any.

From the hook function, you can call *ev\_enqueue()* or *ev\_schedule()* to re-add the event.

---

**Function**

void *ev\_enqueue* (event\_list \* *l*, event \* *e*) – enqueue an event

**Arguments**

event\_list \* *l*  
     an event list

event \* *e*  
     an event

**Description**

*ev\_enqueue()* stores the event *e* to the specified event list *l* which can be run by calling *ev\_run\_list()*.

**Function**

void *ev\_schedule* (event \* *e*) – schedule an event

**Arguments**

event \* *e*  
an event

**Description**

This function schedules an event by enqueueing it to a system-wide event list which is run by the platform dependent code whenever appropriate.

**Function**

int *ev\_run\_list* (event\_list \* *l*) – run an event list

**Arguments**

event\_list \* *l*  
an event list

**Description**

This function calls *ev\_run()* for all events enqueued in the list *l*.

## 8.7 Sockets

Socket resources represent network connections. Their data structure (**socket**) contains a lot of fields defining the exact type of the socket, the local and remote addresses and ports, pointers to socket buffers and finally pointers to hook functions to be called when new data have arrived to the receive buffer (*rx\_hook*), when the contents of the transmit buffer have been transmitted (*tx\_hook*) and when an error or connection close occurs (*err\_hook*).

Freeing of sockets from inside socket hooks is perfectly safe.

**Function**

int *sk\_setup\_multicast* (sock \* *s*) – enable multicast for given socket

**Arguments**

sock \* *s*  
socket

**Description**

Prepare transmission of multicast packets for given datagram socket. The socket must have defined *iface*.

**Result**

0 for success, -1 for an error.

**Function**

int *sk\_join\_group* (sock \* *s*, ip\_addr *maddr*) – join multicast group for given socket

**Arguments**

sock \* *s*  
socket  
  
ip\_addr *maddr*  
multicast address

**Description**

Join multicast group for given datagram socket and associated interface. The socket must have defined *iface*.

**Result**

0 for success, -1 for an error.

---

**Function**

int *sk\_leave\_group* (sock \* *s*, ip\_addr *maddr*) – leave multicast group for given socket

**Arguments**

sock \* *s*  
    socket

ip\_addr *maddr*  
    multicast address

**Description**

Leave multicast group for given datagram socket and associated interface. The socket must have defined *iface*.

**Result**

0 for success, -1 for an error.

---

**Function**

int *sk\_setup\_broadcast* (sock \* *s*) – enable broadcast for given socket

**Arguments**

sock \* *s*  
    socket

**Description**

Allow reception and transmission of broadcast packets for given datagram socket. The socket must have defined *iface*. For transmission, packets should be send to *brd* address of *iface*.

**Result**

0 for success, -1 for an error.

---

**Function**

int *sk\_set\_ttl* (sock \* *s*, int *tll*) – set transmit TTL for given socket

**Arguments**

sock \* *s*  
    socket

int *tll*  
    TTL value

**Description**

Set TTL for already opened connections when TTL was not set before. Useful for accepted connections when different ones should have different TTL.

**Result**

0 for success, -1 for an error.

---

**Function**

int *sk\_set\_min\_ttl* (sock \* *s*, int *tll*) – set minimal accepted TTL for given socket

**Arguments**

sock \* *s*  
    socket

int *tll*  
    TTL value

**Description**

Set minimal accepted TTL for given socket. Can be used for TTL security. implementations.

**Result**

0 for success, -1 for an error.

**Function**

int *sk\_set\_md5\_auth* (sock \* *s*, ip\_addr *local*, ip\_addr *remote*, struct iface \* *ifa*, char \* *passwd*, int *setkey*) – add / remove MD5 security association for given socket

**Arguments**

sock \* *s*  
     socket

ip\_addr *local*  
     IP address of local side

ip\_addr *remote*  
     IP address of remote side

struct iface \* *ifa*  
     Interface for link-local IP address

char \* *passwd*  
     Password used for MD5 authentication

int *setkey*  
     Update also system SA/SP database

**Description**

In TCP MD5 handling code in kernel, there is a set of security associations used for choosing password and other authentication parameters according to the local and remote address. This function is useful for listening socket, for active sockets it may be enough to set *s->password* field.

When called with *passwd* != NULL, the new pair is added, When called with *passwd* == NULL, the existing pair is removed.

Note that while in Linux, the MD5 SAs are specific to socket, in BSD they are stored in global SA/SP database (but the behavior also must be enabled on per-socket basis). In case of multiple sockets to the same neighbor, the socket-specific state must be configured for each socket while global state just once per src-dst pair. The *setkey* argument controls whether the global state (SA/SP database) is also updated.

**Result**

0 for success, -1 for an error.

**Function**

int *sk\_set\_ipv6\_checksum* (sock \* *s*, int *offset*) – specify IPv6 checksum offset for given socket

**Arguments**

sock \* *s*  
     socket

int *offset*  
     offset

**Description**

Specify IPv6 checksum field offset for given raw IPv6 socket. After that, the kernel will automatically fill it for outgoing packets and check it for incoming packets. Should not be used on ICMPv6 sockets, where the position is known to the kernel.

**Result**

0 for success, -1 for an error.



---

**Function**

`sock * sock_new (pool * p)` – create a socket

**Arguments**

`pool * p`  
pool

**Description**

This function creates a new socket resource. If you want to use it, you need to fill in all the required fields of the structure and call `sk_open()` to do the actual opening of the socket.

The real function name is `sock_new()`, `sk_new()` is a macro wrapper to avoid collision with OpenSSL.

---

**Function**

`int sk_open (sock * s)` – open a socket

**Arguments**

`sock * s`  
socket

**Description**

This function takes a socket resource created by `sk_new()` and initialized by the user and binds a corresponding network connection to it.

**Result**

0 for success, -1 for an error.

---

**Function**

`int sk_send (sock * s, unsigned len)` – send data to a socket

**Arguments**

`sock * s`  
socket  
  
`unsigned len`  
number of bytes to send

**Description**

This function sends `len` bytes of data prepared in the transmit buffer of the socket `s` to the network connection. If the packet can be sent immediately, it does so and returns 1, else it queues the packet for later processing, returns 0 and calls the `tx_hook` of the socket when the transmission takes place.

---

**Function**

`int sk_send_to (sock * s, unsigned len, ip_addr addr, unsigned port)` – send data to a specific destination

**Arguments**

`sock * s`  
socket  
  
`unsigned len`  
number of bytes to send  
  
`ip_addr addr`  
IP address to send the packet to  
  
`unsigned port`  
port to send the packet to

**Description**

This is a `sk_send()` replacement for connection-less packet sockets which allows destination of the packet to be chosen dynamically. Raw IP sockets should use 0 for `port`.

---

**Function**

`void io_log_event (void * hook, void * data)` – mark approaching event into event log

**Arguments**

`void * hook`  
event hook address

`void * data`  
event data address

**Description**

Store info (hook, data, timestamp) about the following internal event into a circular event log (*event\_log*). When latency tracking is enabled, the log entry is kept open (in *event\_open*) so the duration can be filled later.